

---

# TrialDura: Hierarchical Attention Transformer for Interpretable Clinical Trial Duration Prediction

---

Ling Yue<sup>1</sup>, Sixue Xing<sup>1</sup>, Jonathan Li<sup>1</sup>, Md Zabirul Islam<sup>1</sup>, Bolun Xia<sup>1</sup>  
Jintai Chen<sup>2</sup>, Tianfan Fu<sup>1\*</sup>

<sup>1</sup>Department of Computer Science, Rensselaer Polytechnic Institute

<sup>2</sup>Hong Kong University of Science and Technology (Guangzhou)

## Abstract

The clinical trial process, a critical phase in drug development, is essential for developing new treatments. The primary goal of interventional clinical trials is to evaluate the safety and efficacy of drug-based treatments on human bodies for specific diseases. However, these trials are often lengthy, labor-intensive, and expensive. The duration of a clinical trial significantly impacts overall costs, making efficient timeline management crucial for controlling budgets and ensuring the economic feasibility of research. To address this issue, we propose TrialDura, a multimodal neural network architecture that estimates the duration of clinical trials using multimodal data (mostly text feature), including disease, drug molecules, trial phases, and eligibility criteria. These are encoded into Bio-BERT embeddings specifically tuned for biomedical contexts to provide a deeper and more relevant semantic understanding of clinical trial data. Finally, the model’s hierarchical attention mechanism integrates all of the text embeddings to capture their interactions and predict clinical trial duration. Our proposed model demonstrated superior performance with a mean absolute error (MAE) of 1.04 years and a root mean square error (RMSE) of 1.39 years compared to other models, indicating more accurate clinical trial duration prediction. The publicly available code can be found at: <https://github.com/LeoYML/TrialDura>.

## 1 Introduction

The clinical trial process, an essential step in drug development, constitutes a critical phase in bringing new medical treatments to fruition. These trials serve as the pivotal gateway for assessing the safety and efficacy of drug-based interventions in addressing various diseases [Vijayanathan and Nawawi, 2008]. However, despite their paramount significance, clinical trials are characterized by many challenges, ranging from their time-consuming nature to their huge costs and relatively low approval rates. Spanning a duration of 7 to 11 years on average, the clinical trial timeline emerges as a prominent obstacle in the drug development landscape [McDowell, 2024]. This protracted duration not only prolongs the time-to-market for potential treatments but also escalates the associated expenses significantly. With an average cost exceeding 2 billion dollars per trial, the financial burden imposed by these endeavors is significant. Furthermore, the low approval rates, around only 15%, underscore the complexity and uncertainty inherent in the clinical trial process [Huang et al., 2022].

One of the primary determinants influencing the overall cost of clinical trials is their duration. Extended trial periods necessitate more resources, more staffing, more medical supplies, and more facility usage, thus increasing operational expenses. Moreover, the prolonged data collection and analysis phases entail more financial investment in technology and personnel. Consequently, effi-

---

\*Corresponding email: fut2@rpi.edu

ciently managing the timeline of clinical trials is paramount in controlling the budget and optimizing the economic feasibility of research endeavors in the pharmaceutical domain. Against this backdrop, integrating machine learning methodologies emerges as a promising avenue for enhancing the efficiency and accuracy of clinical trial duration estimation. By leveraging computational algorithms to analyze intricate datasets, machine learning facilitates the prediction of trial duration with improved precision and reliability.

In this paper, we investigate the application of various machine learning techniques [Yue et al., 2024a], including traditional methodologies such as Linear Regression (LR) [Weisberg, 2005], Gradient Boosting Decision Tree (GBDT) [Ke et al., 2017], Adaptive Boosting (AdaBoost) [Ying et al., 2013], Random Forest (RF) [Breiman, 2001], and a cutting-edge hierarchical deep learning framework that we designed specifically for the clinical trial eligibility criteria to predict the duration of clinical trials.

Our research also produces a comprehensive dataset explicitly tailored for clinical trial duration prediction. This dataset encapsulates diverse eligibility criteria of various clinical trials, including patient demographics, treatment protocols, and regulatory parameters, thereby offering a holistic perspective for predictive modeling. Through extensive experimentation and analysis, we evaluate the efficacy of seven distinct machine learning methodologies in predicting clinical trial duration. The overarching objective of this study is to demonstrate the potential of machine learning in optimizing the clinical trial process, thereby offering insights into the expected duration of clinical trials. By showing the efficacy of various predictive methods, we aim to provide insights that can inform decision-making processes within the pharmaceutical industry, ultimately fostering more efficient and cost-effective drug development practices. Through our empirical findings, we strive to pave the way for adopting data-driven approaches in the landscape of clinical trial management and accelerate the translation of scientific innovations into tangible pharmaceutical solutions.

The main contributions of this paper are summarized as follows:

1. **Problem:** We are the first to identify and propose the problem of clinical trial duration prediction, an essential issue in clinical trial planning, and formulate it as an AI-solvable problem.
2. **Method:** We propose multiple traditional machine learning models to address this problem and design a novel hierarchical attention mechanism, a specialized end-to-end deep learning framework tailored for predicting clinical trial duration. This framework captures multi-modal data elements within clinical trials and leverages large-scale biomedical language models (Bio-BERT).
3. **Results:** The proposed method (TrialDura) demonstrated superior performance, achieving an MAE of 1.044 years and an RMSE of 1.390 years. In comparison to the best baseline method, TrialDura achieves relative reductions of 9% and 7% in MAE and RMSE, respectively.
4. **Interpretability:** Our method exhibits desirable interpretability thanks to the hierarchical attention mechanism, which can assist clinicians in making informed decisions by providing insights into how predictions were reached.

## 2 Methodology

### 2.1 Formulation of Clinical Trial Duration Prediction

A *clinical trial* is an organized endeavor to evaluate the safety and efficacy of a *treatment set* aimed at combating a *target disease set*, according to the guidelines laid out in the *trial eligibility criteria*, for a select group of patients.

**Clinical Trial Phases.** The phases of clinical trials signify distinct stages in the trial process, each with a specific focus and set of objectives. These phases are:

- Phase 1: Initial evaluation of the new drug’s safety and profiling of its pharmacological profile.
- Phase 2: Assessment of the drug’s efficacy and side effects in a larger patient group.
- Phase 3: Confirmation of the drug’s effectiveness, monitoring of side effects, comparison to commonly used treatments, and collection of information that will allow the drug to be used safely.
- Phase 4: Post-marketing studies delineate additional information, including the drug’s risks, benefits, and optimal use.

The phase information is denoted as  $\mathcal{P}$ , which is a one-hot vector.

**Treatment Set.** The treatment set encompasses a variety of drug candidates, denoted as  $\mathcal{M} = \{m_1, \dots, m_{K_m}\}$ , where  $m_1, \dots, m_{K_m}$  are  $K_m$  drug molecules engaged in this trial. This study concentrates on trials to identify new applications for these drug candidates, while trials focusing on non-drug interventions like surgery or device applications are considered outside the scope of this research.

**Target Disease Set.** For a trial addressing  $K_\delta$  diseases, the Target Disease Set is represented by  $\mathcal{D} = \{\delta_1, \dots, \delta_{K_\delta}\}$ , with each  $\delta_i$  symbolizing the diagnostic code<sup>2</sup> for the  $i$ -th disease.

**Trial Eligibility Criteria.** The trial eligibility criteria encompass both inclusion (+) and exclusion (-) criteria, which respectively outline the desired and undesirable attributes of potential participants. These criteria provide details on various key parameters such as age, gender, medical background, the status of the target disease, and the present health condition.

$$\mathcal{E} = [\pi_1^+, \dots, \pi_Q^+, \pi_1^-, \dots, \pi_R^-], \quad \pi_k^{+/-} \text{ is a criterion,} \quad (1)$$

where  $Q$  ( $R$ ) is the number of inclusion (exclusion) criteria in the trial. The term  $\pi_k^+$  ( $\pi_k^-$ ) designates the  $k$ -th inclusion (exclusion) criterion within the eligibility criteria. Each criterion  $\pi$  is a sentence in unstructured natural language.

**Clinical Trial Duration.** The duration of a clinical trial refers to the number of years the trial lasts. It is represented as a continuous value  $y > 0$ .

**Problem (Clinical Trial Duration Prediction).** The estimation of  $y$ , represented as  $\hat{y}$ , can be formulated through the machine learning model  $f_\Theta$ , such that

$$\hat{y} = f_\Theta(\mathcal{P}, \mathcal{M}, \mathcal{D}, \mathcal{E}),$$

where  $\hat{y} > 0$  denotes the estimated duration of a trial;  $\Theta$  denotes the learnable parameter. In this context,  $\mathcal{P}$ ,  $\mathcal{M}$ ,  $\mathcal{D}$ , and  $\mathcal{E}$  refer to the phase information, treatment set, the target disease set, and the trial eligibility criteria, respectively.

## 2.2 Development of Network

Our novel approach consists of three primary phases that form an end-to-end deep learning framework. These consist of embedding generation, hierarchical attention mechanism, and neural network training. We construct a hierarchical interaction graph  $G$  to connect all input data sources affecting clinical trial duration  $\hat{y}$ . The interaction graph  $G$  is constructed to reflect the real-world trial development process and consists of four inputs, including drugs, diseases, phase, and eligibility criteria with node features of input embedding  $I_m, I_d, I_p, I_e \in \mathbb{R}^d$ . The phase of each clinical trial, represented categorically from 1 to 4, is converted into a one-hot encoded vector.  $I_p$  represents the one-hot encoded phase of a clinical trial:

$$I_p = \begin{cases} [1, 0, 0, 0]^\top, & \text{if Phase} = 1, \\ [0, 1, 0, 0]^\top, & \text{if Phase} = 2, \\ [0, 0, 1, 0]^\top, & \text{if Phase} = 3, \\ [0, 0, 0, 1]^\top, & \text{if Phase} = 4. \end{cases} \quad (2)$$

Formally, we represent drug molecules graphs with node features of input embedding as  $\mathcal{M} = \{m_1, \dots, m_{K_m}\}$  (Equation 3), disease codes as  $\mathcal{D} = \{\delta_1, \dots, \delta_{K_\delta}\}$  (Equation 4), and eligibility criteria as  $\mathcal{E} = [\pi_1^+, \dots, \pi_Q^+, \pi_1^-, \dots, \pi_R^-]$ , where  $\pi_k^{+/-}$  is a criterion (Equation 1) as follows:

$$\text{Molecule Embedding, } I_m = \frac{1}{K_m} \sum_{j=1}^{K_m} f_m(m_j), \quad I_m \in \mathbb{R}^d \quad (3)$$

We need to use a pretrained Bio-language model( [Luo et al., 2022, Yue et al., 2024b]) to embed the drug names, where  $f_m$  is the molecule embedding function. In this paper, we utilized the

<sup>2</sup>In this paper, we use ICD10 codes (International Classification of Diseases) [Anker et al., 2016]

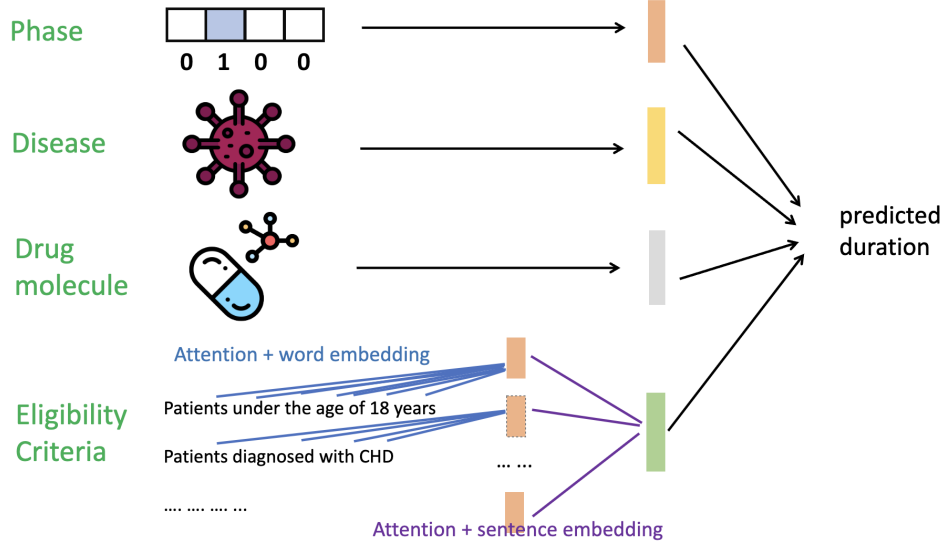


Figure 1: Overview of our model. Our model takes phase information, disease, drug molecule, and eligibility criteria as input and predicts the trial duration. We design a hierarchical attention mechanism to model semantic features from eligibility criteria. The lower-level attention mechanism aggregates word embeddings while higher-level attention gathers sentence embeddings.

widely recognized Bio-BERT model [Lee et al., 2020]. The embedding for each token generated by Bio-BERT is averaged to produce a single 768-dimensional vector representing the drug.

$$\text{Disease Embedding, } I_d = \frac{1}{K_\delta} \sum_{i=1}^{K_\delta} f_\delta(\delta_i), \quad I_d \in \mathbb{R}^d, \quad (4)$$

where  $f_\delta$  represents an embedding of the disease using Bio-BERT, similar to the drug name embedding. The mean of these embeddings provides a 768-dimensional vector representing the disease.

$$\text{Word Embedding, } I_s = \text{BioBERT}_{\text{CLS}}(\pi), \quad I_s \in \mathbb{R}^d \quad (5)$$

where  $I_s$  denotes the sentence embeddings derived from the CLS token outputs of Bio-BERT. The symbol  $\pi$  represents individual sentences within the eligibility criteria.

$$\text{Sentence Embedding, } I_g = f_{\text{transformer}}(I_s), \quad I_g \in \mathbb{R}^d \quad (6)$$

The function  $f_{\text{transformer}}$  encapsulates the transformer model’s mechanism, which processes these sentence embeddings  $I_s$ . This function enhances inter-sentence relationships, culminating in the generation of the paragraph embedding  $I_g$ , which integrates the comprehensive context of the eligibility criteria.

The eligibility criteria text was initially processed using n-word embeddings from the CLS token of Bio-BERT to generate sentence embeddings. To create paragraph embeddings, a transformer with an attention mechanism was used to enhance sentence-to-sentence relations, which is essential for capturing the comprehensive context of the eligibility criteria that often include multiple independent rules. The procedure involved padding sentences to 32 tokens, concatenating the inclusion and exclusion criteria embeddings to form a  $64 \times 768$  matrix, and prefixing with a CLS token embedding to obtain a  $65 \times 768$  matrix. After processing the inputs, the embeddings representing the drug, disease, eligibility criteria (inclusion and exclusion), and trial phases were concatenated to form a unified feature vector. The dimensions of the concatenated vector were  $4 \times 768$ , with the addition of four. The concatenation of input embeddings as an intermediate step:

$$I_{\text{concat}} = \text{CONCAT}(I_p, I_m, I_d, I_g), \quad I_{\text{concat}} \in \mathbb{R}^d. \quad (7)$$

This comprehensive feature vector was then fed into a multiple-layer perceptron (MLP) for regression analysis. The estimated trial duration prediction using the concatenated embedding as input to the MLP:

$$\text{Trial Duration, } \hat{y} = \text{MLP}(I_{\text{concat}}). \quad (8)$$

We employed the Mean Squared Error (MSE) loss function to quantify the discrepancy between predicted and actual trial duration. The objective function is defined by the following equation:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2, \quad (9)$$

where  $y_i$  and  $\hat{y}_i$  represent the true and predicted duration for the  $i^{\text{th}}$  clinical trial, respectively, and  $n$  is the total number of trials in the training dataset. The Adam optimizer [Kingma and Ba, 2015] was utilized to minimize the objective function and update the network weights iteratively based on the training data. The architecture of the whole model is illustrated in Figure 1.

### 3 Experiment

#### 3.1 Dataset Setting

We collected a dataset from <https://clinicaltrials.gov/>, a comprehensive global database of clinical trials, to determine the duration of these trials. We only consider trials that last less than 10 years, which account for 98% of all the trials, because some extremely long values significantly affect model performance. Each trial was represented as an XML file, from which we extracted several pieces of information, including the NCT ID, disease name, drug molecule, trial phase, time duration, and eligibility criteria. After extracting the data, we engaged in data processing to prepare the data for machine learning applications. The statistics for trial duration are summarized in Table 1.

Table 1: Statistics for trial duration (Unit: Year).

Phase	Average	Minimum	1st Quartile	Median	3rd Quartile
All Phases	2.14	0.003	0.83	1.73	3.00
Phase 1	1.52	0.003	0.25	0.87	2.26
Phase 2	2.47	0.003	1.09	2.00	3.34
Phase 3	2.42	0.005	1.08	1.92	3.20
Phase 4	2.06	0.003	0.92	1.67	2.83

The data shows that predicting duration is challenging, with 75% of the training data spanning up to 3 years and a maximum cap of 10 years.

Our final dataset comprised two distinct sets: a training set with 77,818 clinical trial records and a testing set with 36,786 records. We partitioned the data by using January 1, 2019 as boundary, using the data from trials that started before that date for training and validation, and the data from trials that started after that date for testing. This approach ensures that the completion dates of the trials in the training and validation sets precede the start dates of the trials in the testing set.

We studied the general trial duration problem across different trial phases for many diseases, as shown in Table 2.

Table 2: Distribution of clinical trial phases.

Phases	1	2	3	4
# Trials	31,641	44,622	28,250	25,060

#### 3.2 Baseline Methods

We choose multiple widely-recognized models as baselines: Linear Regression (LR), Gradient Boosting Decision Tree (GBDT), Adaptive Boosting (AdaBoost), Random Forest (RF), and Multi-Layer Perceptron (MLP). GBDT, AdaBoost, RF, and MLP have shown initial effectiveness in clinical trial outcome prediction [Fu et al., 2022, 2023, Lu et al., 2024].

We use *MEAN* as a basic reference, which employs the average duration as a direct estimation for all trials and offers a basic point of reference. Linear Regression (LR) was selected due to its simplicity

and effectiveness in revealing linear relationships within the dataset, serving as a fundamental point of reference for performance evaluation. This provides an essential baseline that allows us to appreciate the incremental value added by more complex models. The Gradient Boosting Decision Tree (GBDT) was chosen for its proficiency in managing diverse and complicated datasets through the sequential correction of prediction errors. Its inclusion as a baseline underscores our commitment to assess our model against sophisticated state-of-the-art algorithms. Similarly, the Random Forest (RF) model was incorporated into our baseline comparison due to its robustness and effectiveness in handling high-dimensional data. The risk of overfitting is reduced by aggregating the outcomes of numerous decision trees, offering a reliable baseline for performance comparison.

By benchmarking TrialDura against these distinct and well-established models, we aim to demonstrate the specific advancements our model brings to the domain of clinical trial duration prediction. This comparison not only validates TrialDura’s effectiveness but also situates its performance within the broader landscape of machine learning applications in clinical research, highlighting its contribution to enhancing predictive accuracy and decision-making in clinical trial management.

### 3.3 Experimental Results

Table 3: Summary of model performance (units in years). Results averaged over five runs; standard deviations are shown. Asterisk (\*) denotes statistical significance (p-value < 0.05). NA: not applicable.

Model	MAE (years) (↓)	RMSE (years) (↓)	$R^2$ (↑)	Pearson Corr. (↑)
MEAN	1.240	1.551	-0.090	NA
Linear Regression	1.162 ± 0.005	1.511 ± 0.003	0.043 ± 0.002	0.275 ± 0.001
GBDT	1.148 ± 0.001	1.494 ± 0.002	0.080 ± 0.001	0.288 ± 0.001
RF	1.167 ± 0.001	1.502 ± 0.001	0.053 ± 0.001	0.232 ± 0.003
XGBoost	1.152 ± 0.004	1.497 ± 0.004	0.069 ± 0.001	0.279 ± 0.001
AdaBoost	1.200 ± 0.003	1.514 ± 0.002	-0.258 ± 0.001	0.169 ± 0.002
MLP	1.144 ± 0.008	1.502 ± 0.011	-0.749 ± 0.049	0.197 ± 0.000
TrialDura	<b>1.044 ± 0.009*</b>	<b>1.390 ± 0.005*</b>	<b>0.191 ± 0.001*</b>	<b>0.463 ± 0.001*</b>

The performance of various predictive models for clinical trial duration is detailed in Table 3. Metrics used for evaluation include Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), the coefficient of determination ( $R^2$ ), and the Pearson correlation coefficient. The MEAN model serves as a basic reference, highlighting the enhanced accuracy of more advanced models. For example, Linear Regression, although simplistic, achieves an MAE of 1.16 years and an RMSE of 1.511 years, demonstrating a basic yet effective capacity to model trial duration.

More sophisticated models such as Gradient Boosted Decision Trees (GBDT), and Random Forest (RF) show further reductions in prediction errors, illustrating their capability in managing complex data structures. Notably, GBDT achieves an MAE of 1.14 years and an RMSE of 1.494 years. Models like XGBoost and Multi-Layer Perceptrons (MLP), while displaying similar MAE values, exhibit variations in other metrics, which may indicate differences in model sensitivity and the distribution of errors.

Our model, TrialDura, significantly outperforms all benchmarks, achieving the lowest MAE and RMSE, which are statistically significant improvements of 9% and 7% respectively over the best baseline model. This superior performance is corroborated by its  $R^2$  and Pearson correlation scores, confirming its robust predictive capability.

We also show TrialDura’s prediction performance for each of the four phases in Table 4. Results show the model’s lowest MAE in Phases 1 and 4, highlighting its strengths and areas for improvement.

Figure 2 visually summarizes these findings, illustrating the superior performance of TrialDura in both MAE and RMSE compared to the other models. The graph provides a clear visual representation of the comparative predictive accuracies, reinforcing the textual data presented in Table 3.

In addition to quantitative metrics, qualitative analysis of these models suggests that the integration of domain-specific features into TrialDura, such as trial phase and therapeutic area, contributes significantly to its enhanced performance. Future work will focus on refining these features and

Table 4: Summary of model performance across different phases.

Phase	MAE (years) ( $\downarrow$ )	RMSE (years) ( $\downarrow$ )	$R^2$ ( $\uparrow$ )	Pearson Corr. ( $\uparrow$ )
All	$1.044 \pm 0.009$	$1.390 \pm 0.005$	$0.191 \pm 0.001$	$0.463 \pm 0.001$
1	$0.836 \pm 0.001$	$1.181 \pm 0.000$	$0.342 \pm 0.000$	$0.593 \pm 0.000$
2	$1.130 \pm 0.001$	$1.459 \pm 0.002$	$0.061 \pm 0.002$	$0.334 \pm 0.001$
3	$1.194 \pm 0.003$	$1.593 \pm 0.002$	$0.159 \pm 0.002$	$0.415 \pm 0.002$
4	$1.099 \pm 0.001$	$1.402 \pm 0.001$	$-0.003 \pm 0.002$	$0.228 \pm 0.002$




Figure 2: Model performance comparison on prediction accuracy.

exploring the incorporation of real-time data updates to further improve the model’s predictive accuracy and operational utility in clinical trial management.

### 3.4 Ablation Study

This ablation study assesses the empirical impact of different model components on predicting clinical trial duration. We evaluate models separately trained for Phases 1 through 4 and compare them against a unified model, TrialDura, which integrates data across all phases.

The results, detailed in Table 5, indicate that the performance varies across phases. For example, Phase 1, tested on 9,816 data points, yielded a Mean Absolute Error (MAE) of approximately 1.12 years and a Root Mean Squared Error (RMSE) of about 1.42 years. The negative  $R^2$  values and negligible Pearson correlation coefficients across all phases suggest challenges in phase-specific predictions. Conversely, the unified TrialDura model demonstrates superior performance with a significant positive  $R^2$  and Pearson correlation, indicating a robust model that effectively captures the underlying patterns across different trial phases.

Table 5: Ablation study: model performance across phases compared to the unified TrialDura model.

Phase	MAE (years) ( $\downarrow$ )	RMSE (years) ( $\downarrow$ )	$R^2$ ( $\uparrow$ )	Pearson Corr. ( $\uparrow$ )
1	$1.117 \pm 0.030$	$1.424 \pm 0.011$	$-0.146 \pm 0.017$	$-0.002 \pm 0.002$
2	$1.227 \pm 0.038$	$1.513 \pm 0.024$	$-0.186 \pm 0.045$	$0.005 \pm 0.006$
3	$1.327 \pm 0.010$	$1.732 \pm 0.008$	$-0.123 \pm 0.017$	$0.005 \pm 0.004$
4	$1.125 \pm 0.010$	$1.386 \pm 0.005$	$-0.148 \pm 0.036$	$-0.001 \pm 0.006$
TrialDura	<b><math>1.044 \pm 0.009</math></b>	<b><math>1.390 \pm 0.005</math></b>	<b><math>0.191 \pm 0.001</math></b>	<b><math>0.463 \pm 0.001</math></b>

Additionally, we explored various methods for aggregating sentence embeddings into inclusion and exclusion criteria embeddings. Table 6 compares the efficacy of maximum pooling, mean pooling, and the CLS token method. The results are closely comparable, with mean pooling slightly outperforming the others by a narrow margin in  $R^2$  value, suggesting that each method effectively captures essential features for the prediction task.

This systematic analysis underscores the importance of each component’s contribution to the model’s overall predictive performance and provides insights for future enhancements.

Table 6: Ablation study: comparison of aggregation methods for sentence representation.

Method	MAE (years) ( $\downarrow$ )	RMSE (years) ( $\downarrow$ )	$R^2$ ( $\uparrow$ )	Pearson Corr. ( $\uparrow$ )
Max Pooling	$1.048 \pm 0.001$	$1.391 \pm 0.002$	$0.189 \pm 0.005$	$0.463 \pm 0.002$
Mean Pooling	$1.048 \pm 0.001$	$1.389 \pm 0.002$	$0.191 \pm 0.001$	$0.463 \pm 0.001$
CLS Token	$1.048 \pm 0.001$	$1.390 \pm 0.002$	$0.191 \pm 0.002$	$0.463 \pm 0.001$

### 3.5 Case Study

Table 7: Comparison of predicted and actual duration of clinical trials.

NCTID	Phase	Disease	Drug	$y$ (Actual Years)	$\hat{y}$ (Predicted Years)
NCT03553810	II	Hypertensive heart disease	Entresto	5.00	4.64
NCT03681093	III	Asthma	Fevipirant	4.00	3.95
NCT02198209	IV	Type 2 diabetes	Sitagliptin	2.37	2.11
NCT03636373	IV	Gout attack	Etanercept	1.77	1.80
NCT04126317	II	Neovascular macular degeneration	Aflibercept	2.51	1.92
NCT04249882	III	Depression	Naltrexone	3.22	2.49
NCT04050488	II	Retinopathy of prematurity	Bevacizumab	3.00	2.87
NCT04388852	I	Prostate carcinoma	Bendamustine	4.03	4.04
NCT04483778	II	Neuroblastoma	Fludarabine	4.99	5.00
NCT04023669	I	Brain tumor	Prexasertib	3.71	3.73

Clinical trials are crucial for advancing medical science, providing essential data on the efficacy and safety of new treatments. Structured into four phases—Phase I (safety), Phase II (efficacy and side effects), Phase III (comparative effectiveness), and Phase IV (post-market surveillance)—each phase plays a vital role in the comprehensive assessment of new treatments. Accurate predictions of trial duration are essential for effective planning, resource allocation, and regulatory compliance, making this a key area for the application of machine learning.

**Examples of successful trials:** Our model, TrialDura, was applied to a series of recent clinical trials, demonstrating its utility in accurately estimating trial duration. The trials, detailed in Table 7, include various drugs and conditions, showing the model’s broad applicability. For instance, the trial for Entresto, a heart failure drug developed by Novartis, involved 4,822 patients and was predicted to last 4.9 years, closely matching the actual duration of five years. Despite the trial not meeting its primary endpoints, the accurate prediction underscores the financial implications and potential for improved planning and resource allocation. Similarly, the phase III trial for Fevipirant, aimed at treating asthma, lasted four years, with our model’s prediction being just under 3.95 years. Although the trial did not meet its anticipated goals, the precision of our predictions highlights the model’s capability to aid the financial forecasting and strategic planning of clinical trials.




Figure 3: Visualization of text segments in the TrialDura model’s output, illustrating Shapley values derived from Clinical Trials (NCT03553810). Shapley values correspond to attention weights, with darker colors indicating higher weights. The number and color at the beginning of each sentence represent the attention weight for the entire sentence.

### 3.6 Interpretability Analysis

The visualization of word and sentence importance in the TrialDura model’s eligibility criteria can be clearly demonstrated using Shapley values, which are depicted in Figure 3. The Shapley



value [Winter, 2002], from game theory, fairly distributes total gains or costs among players based on their contributions. In machine learning, Shapley values measure each feature's impact on a model's prediction. These values show the incremental contribution of each word and sentence in the text as it transitions from a fully masked to an unmasked state, highlighting the impact of each component on the model's final output. This process is quantified starting from a base logit value obtained when all the input text is masked, moving to specific word and sentence contributions as each is revealed.

In our model, attention is focused on words and sentences with higher Shapley values, as shown by darker colors in the visual representation. This highlights the impact of specific terms in the eligibility criteria on the model's predictions. By visualizing these values, we identify the most critical words or phrases influencing trial duration, offering insights into the model's decision-making process and guiding further refinement or feature engineering. This method enhances interpretability, ensuring the model's focus aligns with clinical expectations and relevant input data.

## 4 Discussion

This study advances the use of machine learning to predict clinical trial durations by integrating multimodal data through TrialDura's architecture. The model processes complex datasets, including drug characteristics, disease categories, and eligibility criteria, offering valuable insights into factors influencing trial timelines.

**Technical Implications** Technically, TrialDura introduces a novel use of hierarchical attention mechanisms, specifically designed to address the multifaceted nature of clinical trial data. This approach enhances interpretability, a key feature in clinical applications where understanding the rationale behind model predictions is as critical as the predictions themselves. Additionally, the incorporation of Bio-BERT for processing biomedical text significantly improves the model's ability to interpret complex medical documentation. This specificity could be beneficial in other healthcare applications where accurate interpretation of technical language is essential.

**Clinical Implications** From a clinical perspective, the predictive accuracy of TrialDura can significantly improve the management of clinical trials. By providing reliable estimates of trial durations, the model aids in optimizing resource distribution, refining recruitment strategies, and reducing overall costs and development times for new medical treatments. These capabilities are crucial for enhancing the efficiency of clinical trials, especially in an environment where delays are both common and costly.

**Limitations** Despite its advancements, TrialDura faces challenges in achieving precise predictions, evidenced by a mean absolute error (MAE) of over one year. This inaccuracy stems from the inherent complexity of clinical trial data, the limited availability of robust training datasets, and the model's current inability to adapt to ongoing changes in regulatory and legal frameworks. Future improvements should focus on refining data processing techniques, enhancing the model's sensitivity and specificity, and increasing its adaptability to continuously evolving conditions.

In conclusion, TrialDura represents a key advancement in applying machine learning to clinical trials, improving accuracy and offering deeper insights, thus promoting a data-driven approach to more efficient trial management.

## References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, page 2623–2631, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362016. doi: 10.1145/3292500.3330701. URL <https://doi.org/10.1145/3292500.3330701>.
- Alex Aliper, Roman Kudrin, Daniil Polykovskiy, Petrina Kamyra, Elena Tutubalina, Shan Chen, Feng Ren, and Alex Zhavoronkov. Prediction of clinical trials outcomes based on target choice and clinical trial design with multi-modal artificial intelligence. *Clinical Pharmacology & Therapeutics*, 114(5):972–980, 2023. doi: <https://doi.org/10.1002/cpt.3008>. URL <https://ascpt.onlinelibrary.wiley.com/doi/abs/10.1002/cpt.3008>.
- Abdullah Alsultan, Wael A. Alghamdi, Jahad Alghamdi, Abeer F. Alharbi, Abdullah Aljutayli, Ahmed Albassam, Omar Almazroo, and Saeed Alqahtani. Clinical pharmacology applications in clinical drug development and clinical care: A focus on Saudi Arabia. *Saudi Pharmaceutical Journal*, 28(10):1217–1227, 2020. ISSN 1319-0164. doi: <https://doi.org/10.1016/j.jsps.2020.08.012>. URL <https://www.sciencedirect.com/science/article/pii/S1319016420301833>.
- Stefan D Anker, John E Morley, and Stephan von Haehling. Welcome to the ICD-10 code for sarcopenia. *Journal of Cachexia, Sarcopenia and Muscle*, 7(5):512–514, 2016.
- Scott Askin, Denis Burkhalter, Gilda Calado, and Samar Dakrouni. Artificial intelligence applied to clinical trials: opportunities and challenges. *Health and Technology*, 13, 02 2023. doi: 10.1007/s12553-023-00738-2.
- Kara Baskin. Using machine learning to better predict clinical trial outcomes, 2019. <https://mitsloan.mit.edu/ideas-made-to-matter/using-machine-learning-to-better-predict-clinical-trial-outcomes> [Accessed: (9 March 2024)].
- Paola Berchiolla, Corrado Lanera, Veronica Sciannameo, Dario Gregori, and Ileana Baldi. Prediction of treatment outcome in clinical trials under a personalized medicine perspective. *Scientific Reports*, 12(1):4115, 2022. ISSN 2045-2322. doi: <https://doi.org/10.1038/s41598-022-07801-4>. URL <https://www.nature.com/articles/s41598-022-07801-4>.
- Leo Breiman. Random forests. *Machine Learning*, 45:5–32, 2001.
- Phillip B. Chappell Charles A. Knirsch. Risk assessment and mitigation, 2012. <https://www.appliedclinicaltrialsonline.com/view/risk-assessment-and-mitigation> [Accessed: (9 March 2024)].
- Tianfan Fu, Kexin Huang, Cao Xiao, Lucas M Glass, and Jimeng Sun. HINT: Hierarchical interaction network for clinical-trial-outcome predictions. *Patterns*, 3(4):100445, 2022.
- Tianfan Fu, Kexin Huang, and Jimeng Sun. Automated prediction of clinical trial outcome, February 2 2023. US Patent App. 17/749,065.
- Zhen Yu Hong, Jooyong Shim, Woo Chan Son, and Changha Hwang. Predicting successes and failures of clinical trials with an ensemble ls-svr. *medRxiv*, 2020.
- Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Artificial intelligence foundation for therapeutic science. *Nature Chemical Biology*, pages 1–4, 2022.
- Ronald W Hansen Joseph A DiMasi, Henry G Grabowski. Innovation in the pharmaceutical industry: New estimates of R&D costs. *Journal of Health Economics*, 47:20–33, 2016. doi: <https://doi.org/10.1016/j.jhealeco.2016.01.012>. URL <https://pubmed.ncbi.nlm.nih.gov/26928437/>.
- Ece Kavalcı and Anthony Hartshorn. Improving clinical trial design using interpretable machine learning based prediction of early trial termination. *Scientific Reports*, 13(1):121, 2023. ISSN 2045-2322. doi: <https://doi.org/10.1038/s41598-023-27416-7>. URL <https://www.nature.com/articles/s41598-023-27416-7>.

- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30, 2017.
- Henry Kerali. Forecasting clinical trials: The essential checklist, 2018. <https://www.clinicaltrialsarena.com/comment/how-to-forecast-a-clinical-trial-5783321-2/> [Accessed: (9 March 2024)].
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations, ICLR*, 2015.
- Cecilia S. Lee and Aaron Y. Lee. How Artificial Intelligence Can Transform Randomized Controlled Trials. *Translational Vision Science & Technology*, 9(2):9–9, 02 2020. ISSN 2164-2591. doi: 10.1167/tvst.9.2.9. URL <https://doi.org/10.1167/tvst.9.2.9>.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020.
- Andrew W. Lo, Kien Wei Siah, and Chi Heem Wong. Machine Learning With Statistical Imputation for Predicting Drug Approvals. *Harvard Data Science Review*, 1(1), jul 1 2019. <https://hdsr.mitpress.mit.edu/pub/ct67j043>.
- Yingzhou Lu, Tianyi Chen, Nan Hao, Capucine Van Rechem, Jintai Chen, and Tianfan Fu. Uncertainty quantification and interpretability for clinical trial approval prediction. *Health Data Science*, 2024.
- Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. Biogpt: generative pre-trained transformer for biomedical text generation and mining. *Briefings in bioinformatics*, 23(6):bbac409, 2022.
- Amanda McDowell. How long do clinical trial phases take?, 2024. <https://www.antidote.me/blog/how-long-do-clinical-trial-phases-take> [Accessed: (9 March 2024)].
- Rajiv Prasad. Navigating the complexities of clinical trial budget forecasting and payments, 2024.
- Youran Qi and Qi Tang. Predicting phase 3 clinical trial results by modeling phase 2 clinical trial subject level data using deep learning. volume 106, page 288 – 303, 2019. URL <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85097921119&partnerID=40&md5=ebecbca42397d5ed01d32725146f49ae>. Cited by: 8.
- Pranav Rajpurkar, Jingbo Yang, Nathan Dass, Vinjai Vale, Arielle S Keller, Jeremy Irvin, Zachary Taylor, Sanjay Basu, Andrew Ng, and Leanne M Williams. Evaluation of a machine learning model based on pretreatment symptoms and electroencephalographic features to predict outcomes of antidepressant treatment in adults with depression: a prespecified secondary analysis of a randomized clinical trial. *JAMA network open*, 3(6):e206653–e206653, 2020.
- Neetu Sangari and Yanzhen Qu. A comparative study on machine learning algorithms for predicting breast cancer prognosis in improving clinical trials. In *2020 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 813–818, 2020. doi: 10.1109/CSCI51800.2020.00152.
- Kien Wei Siah, Nicholas W. Kelley, Steffen Ballerstedt, Björn Holzhauer, Tianmeng Lyu, David Mettler, Sophie Sun, Simon Wandel, Yang Zhong, Bin Zhou, Shifeng Pan, Yingyao Zhou, and Andrew W. Lo. Predicting drug approvals: The novartis data science and artificial intelligence challenge. *Patterns*, 2(8):100312, 2021. ISSN 2666-3899. doi: <https://doi.org/10.1016/j.patter.2021.100312>. URL <https://www.sciencedirect.com/science/article/pii/S2666389921001550>.
- Anushya Vijayanathan and Ouzreiah Nawawi. The importance of good clinical practice guidelines and its role in clinical trials. *Biomedical imaging and intervention journal*, 4(1), 2008.
- Yue Wang, Yingzhou Lu, Yinlong Xu, Zihan Ma, Hongxia Xu, Bang Du, Honghao Gao, and Jian Wu. Twin-gpt: Digital twins for clinical trials via large language model. *arXiv preprint arXiv:2404.01273*, 2024.

- Sanford Weisberg. *Applied linear regression*, volume 528. John Wiley & Sons, 2005.
- Eyal Winter. The shapley value. *Handbook of game theory with economic applications*, 3:2025–2054, 2002.
- Yonghui Wu, Mia Levy, Christine Micheel, Paul Yeh, Buzhou Tang, Michael Cantrell, Stacy Coorman, and Wang Qi. Identifying the status of genetic lesion in cancer clinical trial documents using machine learning. *BMC Genomics*, 13:S21, 12 2012. doi: 10.1186/1471-2164-13-S8-S21.
- Cao Ying, Miao Qi-Guang, Liu Jia-Chen, and Gao Lin. Advance and prospects of adaboost algorithm. *Acta Automatica Sinica*, 39(6):745–758, 2013.
- Mengjia Yu, Sheng Zhong, Yunzhao Xing, and Li Wang. Enrollment forecast for clinical trials at the planning phase with study-level historical data. *Therapeutic Innovation & Regulatory Science*, 58: 42–52, 2024. doi: <https://doi.org/10.1007/s43441-023-00564-8>. URL <https://doi.org/10.1007/s43441-023-00564-8>.
- Ling Yue and Tianfan Fu. Ct-agent: Clinical trial multi-agent with large language model-based reasoning. *arXiv preprint arXiv:2404.14777*, 2024.
- Ling Yue, Sixue Xing, Jintai Chen, and Tianfan Fu. Trialenroll: Predicting clinical trial enrollment success with deep & cross network and large language models. *arXiv preprint arXiv:2407.13115*, 2024a.
- Ling Yue, Sixue Xing, Yingzhou Lu, and Tianfan Fu. Biomamba: A pre-trained biomedical language representation model leveraging mamba. *arXiv preprint arXiv:2408.02600*, 2024b.
- Alex Zhavoronkov, Quentin Vanhaelen, and Tudor I. Oprea. Will artificial intelligence for drug discovery impact clinical pharmacology? *Clinical Pharmacology & Therapeutics*, 107(4):780–785, 2020. doi: <https://doi.org/10.1002/cpt.1795>. URL <https://ascpt.onlinelibrary.wiley.com/doi/abs/10.1002/cpt.1795>.

## A Appendix

### A.1 Evaluation Metrics

To evaluate the prediction performance of our TrialDura model, we employed the Mean Absolute Error (MAE) and Root Mean Square Error (RMSE), defined by the following equations:

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|, \quad (10)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}, \quad (11)$$

where  $y_i$  and  $\hat{y}_i$  represent the true and predicted duration for the  $i^{\text{th}}$  clinical trial, respectively, and  $n$  is the total number of trials in the test dataset.

**R-squared ( $R^2$ ) score** is defined as the proportion of the variation in the dependent variable that is predictable from the independent variable(s). It is also known as the coefficient of determination in statistics. Suppose we have  $N$  continuous ground truth  $y_1, \dots, y_N$  and  $N$  corresponding predictions  $\hat{y}_1, \dots, \hat{y}_N$ . The difference  $y_i - \hat{y}_i$  is called residual. Then, we defined the residual sum of squares as

$$\text{SS}_{\text{res}} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2, \quad (12)$$

and define the total sum of squares as

$$\text{SS}_{\text{total}} = \frac{1}{N} \sum_{i=1}^N (y_i - \bar{y})^2, \quad (13)$$

where  $\bar{y} = \frac{1}{N} \sum_{i=1}^N y_i$  is the mean of the ground truth. Then  $R^2$  score is defined as

$$R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{total}}}. \quad (14)$$

Higher  $R^2$  scores indicate better performance. In a perfect prediction model, the prediction exactly matches the ground truth, then  $SS_{\text{res}} = 0$  and  $R^2 = 1$ . A weak predictor that always predicts  $\bar{y}$  would cause  $R^2 = 0$ . Some predictors would even lead to a negative  $R^2$  score.

**Pearson Correlation (PC)** is defined as the covariance of the prediction and the ground truth divided by the product of their standard deviations. For two random variables  $x$  and  $y$ , Pearson Correlation is formally defined as

$$\text{PC} = \frac{\mathbb{E}[(x - \mu_x)(y - \mu_y)]}{\sigma_x \sigma_y}, \quad (15)$$

In the regression task, suppose there are  $N$  data points in the test set,  $y_i$  is the ground truth of the  $i$ -th data sample,  $\hat{y}_i$  is the prediction for  $i$ -th data, Pearson Correlation becomes

$$\text{PC} = \frac{\sum_{i=1}^N ((y_i - \mu_y)(\hat{y}_i - \mu_{\hat{y}}))}{\sigma_y \sigma_{\hat{y}}}, \quad (16)$$

where  $\mu_y = \frac{1}{N} \sum_{j=1}^N y_j$  and  $\mu_{\hat{y}} = \frac{1}{N} \sum_{j=1}^N \hat{y}_j$  are mean of ground truth and prediction, respectively.  $\sigma_y = \sqrt{\sum_{i=1}^N (y_i - \mu_y)^2}$  and  $\sigma_{\hat{y}} = \sqrt{\sum_{i=1}^N (\hat{y}_i - \mu_{\hat{y}})^2}$  are the standard deviations of ground truth and prediction, respectively. The value ranges from -1 to 1. A higher Pearson correlation value indicates better performance.

Also, we conduct hypothesis testing by evaluating the p-value to showcase the statistical significance of our method over the best baseline results. If the p-value is smaller than 0.05, we reject the null hypothesis and claim our method significantly outperforms the best baseline method.

## A.2 Data Example

For ease of exposition, Table 8 shows an example of a real clinical trial and all the related features.

## A.3 Implementation Details and Hyperparameter Tuning

### Hardware and Software Configuration

- **CPU:** Intel(R) Xeon(R) Gold 6248
- **RAM:** 128GB
- **GPU:** NVIDIA RTX A5000
- **Operating System:** Ubuntu 20.04
- **Python Version:** 3.8
- **PyTorch Version:** 2.3

We employed Bayesian Optimization with TPE (Tree-structured Parzen Estimators) for hyperparameter tuning using the Optuna [Akiba et al., 2019] library. In our configuration, we set the number of epochs to 1500 and allowed the number of attention heads to vary among [2, 3, 4]. The number of hidden units was set to (2\*768) for the first linear layer, 768 units for the second linear layer, and 768 units for the transformer feed-forward layer. We maintained the dropout ratio at 0.1 and set the learning rate at 0.001 to optimize our model’s performance efficiently. The model was trained for 200 epochs at a learning rate of 0.001 to optimize the network weights and biases to effectively reduce the loss function.

## A.4 Broad Impact

Accurately forecasting the duration of clinical trials offers significant benefits for trial management. By predicting trial duration, resource allocation such as staffing, budget, and facilities can be optimized, ensuring resources are available when needed to prevent inefficiencies and bottlenecks [Kerali,

Table 8: A real example of a clinical trial record.

Feature	Descriptions
NCTID	NCT00610792
disease	Ovarian Cancer
phase	II
title	Phase 2 Study of Twice Weekly VELCADE and CAELYX in Patients With Ovarian Cancer Failing Platinum Containing Regimens
summary	This is a Phase 2, multicenter open-label, uncontrolled 2-step design. Patients will be arranged in two groups based on the response to their last platinum-containing therapy. The two groups are: 1) Platinum-Resistant Patients: patients with progressive disease while on platinum-containing therapy or stable disease after at least 4 cycles; patients relapsing following an objective response while still receiving treatment; patients relapsing after an objective response within 6 months from the discontinuation of the last chemotherapy, and 2) Platinum-Sensitive Patients: patients who relapsed following an objective response.
study type	Interventional
Inclusion Criteria	ECOG performance status grade 0 or 1; Age $\geq 18$ and $\leq 75$ yrs; Life expectancy of at least 3 months; LVEF must be within normal limits; ...
Exclusion Criteria	Chemotherapy, hormonal, radiation, or immunotherapy or participation in any investigational drug study within 4 weeks of study entry; Pre-existing peripheral neuropathy $>$ Grade 1; Presence of cirrhosis or active or chronic hepatitis; ...
drug	Bortezomib and pegylated liposomal doxorubicin
start date	July 2006
completed date	September 2009
duration	3.2 years
sponsor	Millennium Pharmaceuticals, Inc.
outcome	Withdrawn

2018]. This capability is crucial for planning financial aspects, enhancing the accuracy of budgets, and ensuring efficient use of capital [Baskin, 2019].

Understanding trial timelines aids in devising effective participant recruitment and retention strategies, which are vital in maintaining engagement and compliance in longer trials. This strategic approach minimizes dropout rates and improves trial outcomes. Accurate forecasts also enable precise budgeting for personnel, equipment, and operational expenses, thus managing costs more effectively and reducing waste [Prasad, 2024].

Further, forecasting facilitates improved communication with stakeholders such as sponsors, regulatory bodies, and participants, setting realistic expectations and fostering trust, which leads to more effective collaborations [Yu et al., 2024]. It also aids in securing funding by providing reliable financial plans to sponsors and grant committees.

Recognizing potential delays and their impact allows for proactive risk management, including the development of contingency plans for recruitment challenges, regulatory delays, or unforeseen events. This proactive approach is supported by retrospective analyses of past trials and potentially enhanced by prospective data collection during ongoing trials [Charles A. Knirsch, 2012].

For pharmaceutical companies, accurate forecasting of trial duration is essential for strategic planning, including market entry and product launch strategies. It determines the timing of drug approval and market availability, crucial for competitive positioning and financial planning [Joseph A DiMasi, 2016]. Forecasting also supports planning for regulatory submissions, leading to smoother interactions with regulatory bodies and a more efficient approval process, ultimately speeding up market access for new treatments [Alsultan et al., 2020].

In summary, the ability to forecast the duration of clinical trials provides strategic advantages in managing the complexity, costs, and risks associated with clinical research. It enhances the efficiency and effectiveness of trial execution and plays a critical role in advancing medical innovation by bringing new treatments to market more reliably and swiftly.

## B Related Works

Earlier works in the area of machine learning for clinical trial predictions involve two major directions: predicting individual patient outcomes and predicting the probability of trial success [Askin et al., 2023], [Yue and Fu, 2024].

In terms of predicting individual patient outcomes, the early approaches use traditional methods in machine learning. For instance, Wu et al. [2012] utilized support vector machines (SVMs) to forecast genetic lesions based on cancer clinical trial documents. Rajpurkar et al. [2020] employed gradient-boosted decision trees (GBDTs) to predict improvements in symptom scores by integrating treatment symptom scores and EEG measures in the context of antidepressant treatment for depressive symptoms. Meanwhile, Hong et al. [2020] focused on projecting clinical drug toxicity based on features related to drug properties and target properties, employing an ensemble classifier consisting of weighted least squares support vector regression. Machine learning can also play a pivotal role in generating simulated data to identify more efficient statistical outcome measures [Sangari and Qu, 2020]. One study proposes that employing an AI algorithm to forecast individual patient outcomes and pinpoint those likely to progress rapidly, leading to earlier trial endpoints, could result in shorter trial duration [Lee and Lee, 2020]. Moreover, Berchiolla et al. [2022] developed a machine learning framework to conduct a heterogeneous treatment analysis on type 2 diabetes using randomized clinical trials that share similar inclusion and exclusion criteria as well as a set of common clinical patient characteristics.

In terms of predicting the probability of trial success, machine learning has been utilized to generate early detections and prognoses of disease, which can increase the probabilities of success in a clinical trial [Sangari and Qu, 2020]. Beyond that, machine learning can be employed to predict molecular features, target sensitivity, bioavailability, and toxicity [Zavoronkov et al., 2020] in order to minimize potential trial failures in later stages. This can contribute to designing phase-II and phase-III trials that have a higher chance of success. Qi and Tang [2019] designed a Residual Semi-Recurrent Neural Network (RS-RNN) to forecast phase III trial outcomes based on phase II results. Lo et al. [2019] explored different imputation techniques and a range of traditional machine learning models, including logistic regression, random forest, and SVM, to forecast drug approval within 15 disease categories. Similarly, Siah et al. [2021] assessed various traditional machine learning models for predicting clinical trial outcomes. In more recent works, Fu et al. [2023, 2022], Lu et al. [2024] predicted outcomes across all clinical trial phases, utilizing comprehensive trial features such as drug molecules and trial eligibility criteria and integrating multimodal data sources. Wang et al. [2024] designed a TWIN-GPT model to synthesize patient electronic records to simulate clinical trials. Kavalci and Hartshorn [2023] introduced an improvement to the predictive model for clinical trial termination by incorporating eligibility criteria search capabilities derived from unstructured text and disease categorization characteristics into the existing study features, showing a statistically significant improvement in predicting early termination of clinical trials. Moreover, Aliper et al. [2023] created a platform that utilizes transformer technology and generative AI, incorporating multimodal data such as omics, text, clinical trial design, and small-molecule properties to forecast the results of phase II clinical trials.

Earlier works in machine learning for clinical trial predictions primarily focus on two aspects: predicting individual patient outcomes and predicting the probability of trial success. However, these approaches, while innovative and beneficial in their respective domains, did not directly address the issue of clinical trial duration. Our approach, however, seeks to fill this gap by specifically targeting the trial duration problem. We cannot simply adapt the methodologies from prior works because our task encompasses a broader scope. Our model integrates Bio-BERT embeddings, specifically tuned for biomedical contexts, providing a deeper and more relevant semantic understanding of clinical trial data. Additionally, the hierarchical attention mechanism in our model enables a nuanced analysis of relationships between various data elements, a feature not commonly addressed in previous studies.