CAPITAL UNIVERSITY OF SCIENCE AND TECHNOLOGY, ISLAMABAD


# Natural Language Inference for Clinical Trials An Experimental Study

by

Ghanza Iqbal

A thesis submitted in partial fulfillment for the degree of Master of Science

in the

Faculty of Computing
Department of Computer Science

2025

*Dedicated to my Parents and Teachers. Whose support and encouragement have been invaluable.*


# CERTIFICATE OF APPROVAL

**Natural Language Inference for Clinical Trials An Experimental Study**

by

Ghanza Iqbal

(MCS231006)

## THESIS EXAMINING COMMITTEE

| S. No. | Examiner | Name | Organization |
|---|---|---|---|
| (a) | External Examiner | Dr. Malik Ahmad Kamran | CU, Islamabad |
| (b) | Internal Examiner | Dr. Muhammad Siraj Rathore | CUST, Islamabad |
| (c) | Supervisor | Dr. Muhammad Abdul Qadir | CUST, Islamabad |

Dr. Muhammad Abdul Qadir
Thesis Supervisor
October, 2025

Dr. Mohammad Masroor Ahmed
Head
Dept. of Computer Science
October, 2025

Dr. Muhammad Abdul Qadir
Dean
Faculty of Computing
October, 2025

# *Author's Declaration*

I, **Ghanza Iqbal** hereby state that my MS thesis titled "**Natural Language Inference for Clinical Trials An Experimental Study**" is my own work and has not been submitted previously by me for taking any degree from Capital University of Science and Technology, Islamabad or anywhere else in the country/abroad.

At any time if my statement is found to be incorrect even after my graduation, the University has the right to withdraw my MS Degree.

**(Ghanza Iqbal)**

Registration No: MCS231006

# *Plagiarism Undertaking*

I solemnly declare that research work presented in this thesis titled "**Natural Language Inference for Clinical Trials An Experimental Study**" is solely my research work with no significant contribution from any other person. Small contribution/help wherever taken has been duly acknowledged and that complete thesis has been written by me.

I understand the zero tolerance policy of the HEC and Capital University of Science and Technology towards plagiarism. Therefore, I as an author of the above titled thesis declare that no portion of my thesis has been plagiarized and any material used as reference is properly referred/cited.

I undertake that if I am found guilty of any formal plagiarism in the above titled thesis even after award of MS Degree, the University reserves the right to withdraw/revoke my MS degree and that HEC and the University have the right to publish my name on the HEC/University website on which names of students are placed who submitted plagiarized work.

**(Ghanza Iqbal)**

Registration No: MCS231006

# *Acknowledgement*

I would like to express my deepest gratitude to Almighty Allah, whose unwavering support and guidance have been the foundation of my journey throughout this thesis. Through every challenge and moment of uncertainty, I have felt His presence and strength guiding me. His wisdom and grace have been a constant source of inspiration and perseverance, making this achievement possible. I would like to extend my heartfelt thanks to everyone who has supported and guided me in achieving this milestone. I am extremely obliged to Dr. M. Abdul Qadir for shaping the direction of my research, offering invaluable guidance during moments of uncertainty, and encouraging me throughout the research. My family has been a major source of strength and encouragement, providing crucial support whenever I needed it. I am deeply thankful to my respected parents, spouse and siblings for their unwavering encouragement, patience and prayers which have been my greatest source of strength throughout my MS studies. The academic development I have achieved is largely due to the dedication and support of the professors at Capital University of Science and Technology (CUST), for which I am immensely grateful. Finally, I would like to express my thanks and offer my regards to all those who have supported me in any capacity.

**(Ghanza Iqbal)**

# *Abstract*

Clinical trials are critical for evaluating the safety and efficacy of new medical interventions before public release. With the rapid growth of publicly accessible trial data, **Clinical Trial Reports (CTRs)** have become a vital resource for evidence-based research. CTRs, typically written in natural language, contain key sections such as Intervention, Eligibility, Results, and Adverse Events. Automatically extracting and reasoning over this unstructured information remains a major challenge.**Natural Language Inference (NLI)** for CTRs aims to determine the logical relationship (entailment or contradiction) between a given hypothesis and the content of a CTR.

Despite substantial progress in textual entailment using modern NLP models, applying NLI to CTRs remains difficult due to **domain-specific terminology, complex eligibility criteria, embedded numerical thresholds, and the need for precise, evidence-based reasoning**. Addressing these challenges requires methods that combine the strengths of symbolic and neural approaches.

This research presents a **hybrid NLI framework** for CTRs that integrates **semantic, symbolic, and neural reasoning**. Using the NLI4CT dataset, we develop an **ensemble system** comprising a **Multi-Granularity Inference Network (MGNet), GPT-4, and SciFive**, augmented with a **rule-based reasoning component**. Domain-informed rules were derived from detailed analysis of hypothesis patterns in the training data and applied primarily to the Adverse Events section.

Our proposed hybrid system significantly outperforms existing baselines, achieving an **F1 score of 0.91.** These results demonstrate the effectiveness of combining symbolic reasoning with state-of-the-art neural models, paving the way for more accurate and interpretable NLI systems in clinical trial analysis.

# Contents

# List of Figures

# List of Tables

# Abbreviations

| | |
|---|---|
| **AI** | Artificial Intelligence |
| **AMR** | Abstract Meaning Representation |
| **ANN** | Artificial Neural Network |
| **BERT** | Bidirectional Encoder Representations from Transformers |
| **BioBERT** | Biomedical BERT |
| **CTR** | Clinical Trial Report |
| **CNN** | Convolutional Neural Network |
| **DL** | Deep Learning |
| **EMNLP** | Empirical Methods in Natural Language Processing |
| **FOL** | First-Order Logic |
| **FN** | False Negative |
| **FP** | False Positive |
| **FNLP** | Formal Natural Language Processing |
| **GPT** | Generative Pre-trained Transformer |
| **IR** | Information Retrieval |
| **LLM** | Large Language Model |
| **LSTM** | Long Short-Term Memory |
| **ML** | Machine Learning |
| **MGNet** | Multi-Granularity Graph Network |
| **MNLI** | Multi-Genre Natural Language Inference |
| **NLP** | Natural Language Processing |
| **NLI** | Natural Language Inference |
| **PICO** | Population, Intervention, Comparison, Outcome |
| **POS** | Part of Speech |

| | |
|---|---|
| **QA** | Question Answering |
| **RNN** | Recurrent Neural Network |
| **RoBERTa** | Robustly Optimized BERT Approach |
| **SciFive** | Scientific T5 (Transformer model for biomedical NLP) |
| **SNLI** | Stanford Natural Language Inference |
| **TF-IDF** | Term Frequency-Inverse Document Frequency |
| **TN** | True Negative |
| **TP** | True Positive |
| **UMLS** | Unified Medical Language System |
| **XAI** | Explainable Artificial Intelligence |
| **XLM-R** | Cross-lingual Language Model - RoBERTa |

# Chapter 1

# Introduction

Clinical Trial Reports CTRs are essential documents in the evaluation and regulatory approval of new therapies, providing a structured account of study rationale, methodology, results, and interpretations. In the current era of rapidly expanding biomedical research, CTRs are produced in unprecedented volume. For instance, ClinicalTrials.gov and other international registries record tens of thousands of new studies each year, underscoring the critical role of CTRs in ensuring transparency and reproducibility in clinical evidence [1].

Compiling a CTR requires interdisciplinary collaboration among clinical investigators, statisticians, regulatory specialists, and data managers, who follow established frameworks such as the International Council for Harmonisation (ICH) E3 guidance [2]. While the adoption of electronic data capture has improved the efficiency and consistency of data collection, CTRs continue to be largely unstructured or semi-structured documents combining free-text narratives, tables, and figures [3]. This unstructured nature makes it challenging to extract and synthesize data efficiently, limiting the speed and scalability of evidence generation [4]. Natural Language Processing (NLP) and machine learning methods have emerged as powerful tools to address this challenge. Recent research highlights the impact of transformer-based language models—including BioBERT and domain-adapted versions of BERT on biomedical information extraction tasks, such as named entity recognition, relation extraction, and document classification [5]. These models

enable systems to process complex clinical language with greater accuracy, outperforming earlier rule-based approaches [6]. For instance, recent studies demonstrate that transformer architectures can identify trial populations, interventions, outcomes, and safety signals from free-text sections of trial reports [7]with relatively better reliability.

Advanced NLP pipelines have also been used to support evidence synthesis, including automated risk of bias assessment and meta-analysis [8]. Integrating these computational techniques with clinical informatics infrastructures allows structured data extracted from CTRs to be linked to electronic health records and pharmacovigilance systems, further strengthening post-market safety monitoring and regulatory oversight [9]. Such approaches facilitate the production of timely, high-quality evidence to guide decision-making in clinical practice.

Despite promising advances, several challenges persist. Variability in reporting styles and incomplete or inconsistent documentation complicate automatic information extraction. Moreover, rigorous evaluation of NLP pipelines is essential to ensure transparency, reproducibility, and fairness [10]. Privacy concerns around sensitive health information further necessitate responsible governance and data protection strategies [11].

In this context, the present study explores state-of-the-art computational methods for processing Clinical Trial Reports. By leveraging recent innovations in deep learning and NLP, this work aims to demonstrate how unstructured narratives can be transformed into structured data that accelerate evidence generation, regulatory review, and ultimately the translation of research into improved patient care.

## 1.1 Natural Language Processing Pipelines for Clinical Trial Reports

Clinical Trial Reports (CTRs) are among the most complex biomedical documents to process automatically. They blend:

- Free-text narratives describing study design and outcomes.

- Structured tables with numeric measurements.
- Numeric values spelled out in text.
- Domain-specific medical terminology.
- Cultural or region-specific expressions.
- Embedded quantitative reasoning (e.g., statistical results in prose).

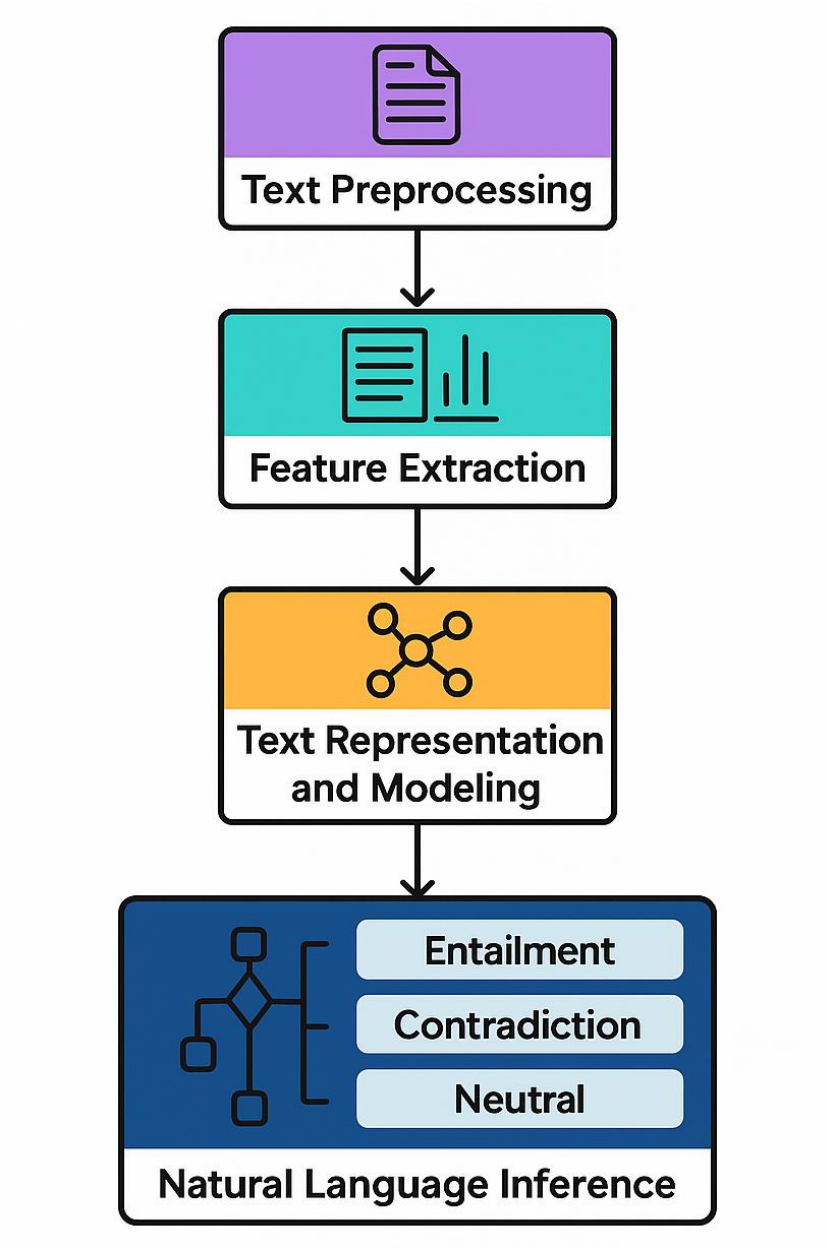


FIGURE 1.1: NLP Pipeline

NLP serves as the foundational step, by transforming the unstructured textual content of CTRs into structured formats through systematic preprocessing and

feature extraction. This structured output is subsequently leveraged by Natural Language Inference (NLI), which facilitates logical reasoning by evaluating the relationship between a hypothesis (typically a query formulated by medical professionals) and a premise (the corresponding evidence extracted from the CTRs) [12]. NLI aims to determine the underlying logical relationship between the two text fragments, categorizing them as entailment, contradiction, or neutral. This reasoning mechanism is essential for enabling automated, evidence-based interpretation of clinical information, thereby supporting decision-making in medical research and practice. Some applications of NLI:

- Question Answering: Verifies whether an answer is supported by evidence.
- Text Summarization: Generates consistent summaries grounded in source content.

Example

TABLE 1.1: Example of Entailment Label in Clinical NLI

| **Premise** | **Hypothesis** | **Label** |
|---|---|---|
| Patients over the age of 60 with heart failure are excluded. | Elderly patients with cardiac issues cannot enroll. | Entailment |

In this example, the system must reason that:

- "Elderly" implies "over 60"
- "Cardiac issues" aligns with "heart failure"
- "Exclusion" indicates ineligibility to enroll

NLI relies on a large range of techniques, from symbolic reasoning methods to data-driven machine learning and deep learning approaches. To support such an inference, NLP pipelines have evolved significantly, transitioning from early rule-based systems to modern deep learning and hybrid.

TABLE 1.2: Semantic Analysis of an NLI Example in Clinical Context

| **Element** | **Text / Role** |
|---|---|
| Premise | Patients over 60 with heart failure are excluded. |
| Hypothesis | Elderly patients with cardiac issues cannot enroll. |
| Reasoning | Elderly ≈ over 60; cardiac issues ≈ heart failure; exclusion = cannot enroll |
| Label | Entailment |

## 1.2 Rule-Based NLP Pipeline Using First-Order Logic

Rule-based NLP pipelines rely on handcrafted patterns and logical rules to parse, extract, and infer information. In CTRs, these pipelines often consist of:

- Tokenization and Lexical Analysis: Splitting text into tokens (words, numbers, punctuation).
- Pattern Matching: Using regular expressions or grammar rules to detect phrases.
- Terminology Normalization: Mapping synonyms to standard concepts.
- Syntactic Parsing: Analysing sentence structure (subject–verb–object).

### 1.2.1 Logic-Based Reasoning

#### 1.2.1.1 Propositional Logic

In the medical domain, propositional logic determines the logical relationship between clinical statements in NLI. For instance, the premise "All patients with pneumonia have a cough" can be represented as x (Pneumonia(x) → Cough(x)), while the hypothesis "Some patients with pneumonia do not have a cough" becomes x (Pneumonia(x) ¬Cough(x)). Since the hypothesis contradicts the general

rule expressed in the premise, the logical conclusion is a contradiction. Such logical reasoning is essential for accurate inference in clinical decision-support systems, particularly in rule-based or interpretable NLI frameworks [13].

#### 1.2.1.2 First-Order Predicate Logic

FOPL represents complex statements with variables and quantifiers using formal logic. It performs reasoning by translating premises and hypotheses into logical predicates to describe the relationship between them [14]. FOPL understands natural language by representing the meaning of sentences in a formal and logical manner. This logical approach enables transparent and explainable decision-making processes. It makes FOL ideal for sensitive domains like medicine, where doctors and researchers need to trust and verify system outputs [15]. This approach is interpretable and controllable, making it valuable when transparency is critical (e.g., regulatory audit). However, rule-based pipelines struggle with:

Table 1.3: Example of FOL-Based Reasoning for Clinical NLI

| **Component** | **Expression / Example** |
|---|---|
| Premise | Drug B reduced symptoms in 70% of patients. |
| FOL Representation | $\exists x(\text{Patient}(x) \wedge \text{TreatedWithDrugB}(x) \wedge \text{SymptomsReduced}(x))$ for approximately 70% of $x$. |
| Simplified Logic | $\forall x(\text{Patient}(x) \wedge \text{TreatedWithDrugB}(x) \rightarrow \text{Likely}(\text{SymptomsReduced}(x)))$ |
| Hypothesis | Most patients felt better after taking Drug B. |
| FOL Mapping | $\exists x(\text{Patient}(x) \wedge \text{TreatedWithDrugB}(x) \wedge \text{SymptomsReduced}(x)) \rightarrow \text{Most}(\text{Patient}(x))$ |
| Reasoning Outcome | **Entailment** (since 70% implies "most"). |

- Complex nested clauses.
- Implicit information.
- Cultural idioms.
- Mathematical expressions embedded in text.

The logic in the form of pseudo code can be described as given in Table 1.4

TABLE 1.4: Rule-Based Extraction Using Propositional Logic

| **#1. Preprocessing** |
|---|
| `Sections = SegmentSections(CTR)`<br>`CleanedText = NormalizeText(Sections)`<br>`Tokens = Tokenize(CleanedText)` |
| **#2. Rule-Based Extraction (Propositional Logic)** |
| `RuleBasedEntities = []`<br>`For each Sentence in Tokens:`<br>`  If ContainsPattern(Sentence, "95% confidence interval"):`<br>`  RuleBasedEntities.Append("ConfidenceInterval")`<br>`  If ContainsPattern(Sentence, "non-inferior to"):`<br>`  RuleBasedEntities.Append("NonInferiorityStatement")`<br>`  If MatchesRegEx(Sentence, r"[0-9]%.*[0-9]% reduction"):`<br>`  RuleBasedEntities.Append("EfficacyResult")` |

## 1.3 Machine Learning-Based NLP Pipeline

Machine learning (ML) pipelines use statistical models trained on labelled data to classify, extract, and relate entities[16].
Typical components include:

- Vectorization: Converting text to numerical representations (e.g., TF-IDF, word embeddings).
- Feature Engineering: Capturing domain-relevant characteristics such as:
  1. Presence of numeric tokens.
  2. Position within a section (e.g., "Results").
  3. Surrounding context windows.
- Model Training: Training classifiers (SVMs, logistic regression) or sequence models (CRFs) for tasks like:
  1. Named entity recognition (NER).
  2. Relation extraction (e.g., linking dosages to outcomes).
  3. Section classification.

In CTRs, ML pipelines can recognize medical terms and standard numeric expressions but often underperform with:

- Long-range dependencies (e.g., resolving pronouns across paragraphs)
- Unseen variants of terms
- Embedded formulas written in text

The logic for machine learning based system in the form of pseudo code can be described as given in Table 1.5

TABLE 1.5: Pseudo code for ML based system

| **#1. Machine Learning-Based NER** |
| --- |
| `MLModel = LoadMLModel("SVM_NER_Model")`<br>`ML_Entities = MLModel.PredictEntities(CleanedText)` |

### 1.3.1 Applications and Strengths

- Explainability: Since features are clear and explainable, their contribution to decisions can easily be interpreted by humans.
- Speed: Training and inference on texts are faster compared to deep learning models.
- Baseline Utility: ML models often serve as control baselines for new approaches.
- Integration in Hybrid Systems: Useful for scoring premise-hypothesis pairs before feeding into deeper models [17].

### 1.3.2 Challenges of ML in Clinical Inference

Despite their simplicity, ML models face several challenges:

- Poor Generalization: Performance drops significantly on out-of-domain data [18].
- Feature Limitation: Shallow lexical features miss syntactic and contextual variations [19].
- Vocabulary Dependence: Difficulty handling paraphrased or synonymous sentences.
- Insensitivity to Structure: Cannot model sentence dependencies [20].
- Numerical Logical Inference: Fail to reason over numerical values or conditions (e.g., dose> $300mg$)[21].

Even if the premise and hypothesis share many similar words, a contradiction can still exist, and it often becomes clear only when domain-specific knowledge, such as comparing medical numbers or values, is applied. This reveals that accurate inference in medical NLI relies not merely on surface-level textual similarity but on the incorporation of domain-specific features such as numerical thresholds, measurement units, and clinical context. Without such specialized reasoning, models risk misclassifying semantically conflicting information as entailment [22].

## 1.4 Deep Learning-Based NLP Pipeline

Deep learning pipelines use neural networks that learn representations end-to-end, reducing dependence on manual features [23].
Key components include:

- Pre-trained Language Models:
  BERT, BioBERT, ClinicalBERT, and Bioformer.
  Trained on large biomedical corpora.
- Fine-Tuning
  Adapting the model to specific CTR tasks (NER, summarization, question answering)

Sequence-to-Sequence Models:

For generating structured summaries of outcomes

Contextual Embeddings

Capturing semantics of domain-specific terms, cultural expressions, and variable numeric phrasing

These models handle:

- Diverse terminology
- Implicit reasoning ("the intervention achieved non-inferiority")
- Context-dependent disambiguation

However, they are computationally intensive and may struggle with precise logic constraints or verification of numeric correctness in embedded calculations. The logic for deep learning-based system in the form of pseudo code can be described as given in Table 1.6

TABLE 1.6: Pseudo code for DL based system

| # 1. Deep Learning-Based Contextual Extraction |
|---|
| DLModel = LoadTransformerModel("BioBERT")<br>DL_Entities = DLModel.ExtractEntities(CleanedText)<br>DL_Relations = DLModel.ExtractRelations(CleanedText) |

Challenges in DL [24]

Despite their success, deep learning models face several challenges in NLI, particularly in medical and clinical contexts:

- Numerical Inference Deficiency: Models often treat numbers as tokens, lacking arithmetic or comparative reasoning.
- Lack of Explainability: Predictions are often difficult to trace back to specific evidence or reasoning steps.

- Biases in Training Data: Models trained on SNLI/MultiNLI may not generalize to CTRs due to domain shift.
- Multi-hop Reasoning: Struggle when entailment requires chaining information from multiple sentences or sections.

## 1.5 Hybrid NLP Pipeline

Hybrid pipelines combine rule-based, ML, and deep learning components to maximize strengths [25]:

- Rule-Based Layers
    1. For deterministic parsing (e.g., recognizing table captions, extracting numeric units)
    2. Encoding propositional and predicate logic rules to enforce domain constraints (e.g., dosage must match units)
- Deep Learning Layers
    1. For contextual entity recognition and relation extraction
- Post-Processing
    1. Logic-based verification of outputs
    2. Handling mathematical expressions with symbolic parsers (e.g., parsing "p¡0.05")

Example Hybrid Workflow for CTRs **Section Segmentation (rule-based):** Identify "Methods," "Results,", "Safety," etc.

1. Deep Learning Extraction: Recognize entities (interventions, outcomes, adverse events).

2. Logic Reasoning Module:

    - Use first-order rules to infer study-level assertions.
    - Example:
        - IF Outcome(x) AND Significant(x) AND (p-value(x) ¡ 0.05)
        - THEN Conclusion = "Efficacy demonstrated."

3. Normalization Validation: Cross-check extracted numbers (digit/text equivalence) and validate against tables. This approach improves:

    - Recall and precision across heterogeneous data.
    - Interpretability of outputs.
    - Robustness to format variations and embedded calculations.

The logic for hybrid system in the form of pseudo code can be described as given in Table 1.7

TABLE 1.7: Pseudo code for hybrid system

```
# 1. Hybrid Reasoning
Conclusions = [ ]
For each Fact in InferredFacts:
  If Fact == "PositiveOutcome":
    Conclusions.Append("Study shows efficacy evidence.")
  If Fact == "StatisticalEvidencePresent":
    Conclusions.Append("Statistical validation provided.")
For each Relation in DL_Relations:
  If Relation.Type == "Outcome-Intervention":
    Conclusions.Append("Outcome linked to intervention.")
```

## 1.6 Why Hybrid Pipelines Matter for CTRs

Hybrid systems offer the best balance [26]:

- Rule-based logic: for compliance and verification

- ML and deep learning: for flexible language understanding
- Symbolic reasoning modules: for cross-validating numeric and textual consistency

The logic for handling mixed cultural and medical language, mathematical and numerical issues in the form of pseudo code can be described as given in Table 1.8 longtable

TABLE 1.8: Pseudo code for hybrid system

| # 1. Numeric Consistency Checks |
|---|
| NumericMentions = ExtractNumbersFromText(CleanedText)<br><br>For each NumberText in NumericMentions:<br>  If IsTextualNumber(NumberText):<br>  DigitEquivalent = ConvertTextToDigit(NumberText)<br>  ReplaceInText(NumberText, DigitEquivalent) |
| # 2. Mathematical Expression Parsing |
| MathExpressions = FindMathExpressions(CleanedText)<br>ParsedMath = [ ]<br><br>For each Expr in MathExpressions:<br>  Parsed = ParseExpression(Expr)<br>  ParsedMath.Append(Parsed) |
| # 3. Entity Normalization (Terminology + Cultural Terms) |
| AllEntities = MergeEntities(RuleBasedEntities, ML_Entities, DL_Entities)<br><br>NormalizedEntities = [ ]<br><br>For each Entity in AllEntities:<br>  MappedEntity = MapToOntology(Entity, "SNOMED_CT_or_MedDRA")<br>  NormalizedEntities.Append(MappedEntity) |

## 1.7 Introduction to Problem

Despite advancements in NLP models for textual entailment, accurately performing NLI tasks on CTRs remains a challenge due to domain-specific language, complex eligibility criteria, embedded mixed format numerical and textual thresholds, the logical reasoning specially depending on multiple premises and the need for evidence-based reasoning. Transformer-based models show promise in biomedical contexts but often lack interpretability and symbolic validation. Likewise, traditional methods like TF-IDF and shallow classifiers struggle to capture deep semantic and numerical relations. These limitations highlight the need for a robust hybrid approach that integrates semantic, symbolic, and neural reasoning to ensure accuracy and interpretability in clinical NLI tasks.

## 1.8 Research Objectives

- **RO1:** Experimental Investigation of existing NLI techniques for Clinical Trial Reports Inference.
- **RO2:** Propose NLI system for better results.
- **RO3:** Implement and evaluate the performance of the proposed system.

## 1.9 Research Methodology

The research methodology comprises four phases, adopted from the eight-step model proposed by Kumar et al [27].

**Phase 1: Deciding What to Research** (see Chapter 2)

**Step 1:** Formulating a research problem.

**Step 2:**Write a research proposal.

**Phase 2: Planning the Research Study**(see Chapter 3)

**Step 1:** Conceptualize a Research Design.


FIGURE 1.2: Research Methodology

**Step 2:** Data Collection.

**Step 3:** Data Preprocessing.

**Phase 3: Conducting the Research Study** (see Chapter 3)

**Step 1:** Performing experiments for RQ3 and RQ4.

**Step 2:** Evaluation and comparison of proposed hybrid methodology.

**Step 3:** Writing the Research Report.

# Chapter 2

# Literature Review

## 2.1 Introduction

Natural Language Inference (NLI) has undergone a significant evolution within the broader field of Natural Language Processing (NLP), progressing from early rule-based systems to sophisticated deep learning and hybrid architectures. At its core, NLI seeks to transform unstructured clinical narratives into structured, machine-interpretable representations an essential step for enabling evidence-based reasoning in healthcare contexts. This literature review examines existing NLI techniques through the lens of **Clinical Trial Reports (CTRs)** and domain-aware knowledge, where the intricate structure of medical language, embedded numerical information, and the demand for precise logical inference present unique challenges.

Over the years, various models have been evaluated using standard performance metrics such as accuracy, precision, recall, and F1 score. While some approaches demonstrate strong lexical matching and semantic similarity, they often fall short in capturing **domain-specific subtleties**, handling **numerical reasoning**, and ensuring **logical consistency** all critical for clinical applications. Drawing on both foundational research and recent advancements, this review provides a comprehensive overview of NLI paradigms and critically assesses their strengths and

limitations. In doing so, it underscores the persistent need for **more accurate, interpretable, and domain-adapted NLI systems** tailored to the complexities of CTRs.

## 2.2 Historical Background and Theoretical Foundations

### 2.2.1 Early Symbolic Approaches

The foundation of NLI lies in symbolic logic and formal semantics. Early NLI systems translated natural language into logical forms, often using predicate logic, and then used theorem proving to evaluate entailment. The FraCaS (Framework for Computational Semantics) [28] test suite is one of the earliest benchmarks for evaluation of NLI techniques.

Another groundbreaking contribution is by Manning (2009) [29], who proposed "Natural Logic" as a means of reasoning directly over linguistic expressions without translating them into formal logic. Their system used inference rules to simulate human-like logical reasoning.

### 2.2.2 Challenges with Symbolic Systems

Symbolic approaches are interpretable and grounded in formal theory but suffer from scalability issues. It Craft detailed logical rules and lexical databases for all possible sentences. Moreover, symbolic models struggled with ambiguity, paraphrasing, and context sensitivity, which are built from natural language [30].

## 2.3 Emergence of Data-Driven Approaches

### 2.3.1 SNLI: A Turning Point

The release of the Stanford Natural Language Inference (SNLI) corpus by Bowman et al. ([31] marked a turning point in the evolution of NLI. The proposed

dataset contains over 570,000 human-annotated sentence pairs, with clear labels for entailment, contradiction, or neutral. Unlike earlier, smaller datasets, SNLI provided a large-scale resource that enabled the training of deep learning models and techniques with millions of parameters.

### 2.3.2 MultiNLI and Domain Diversity

Williams et al. proposed the Multi-Genre NLI (MultiNLI) dataset in this research study. MultiNLI broadened the coverage to include different genres such as fiction, spoken conversation, and government documents. This variety exposed the limitations of models trained only on SNLI, which often failed to generalize to new domains. The baseline BiLSTM model achieved 73% accuracy on matched and mismatched test sets [32].

## 2.4 Neural Network Models for NLI

### 2.4.1 LSTM-Based Models

The very first generation of DL models for NLI employed Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) units.
Wang and Jiang introduced a matching LSTM reasoning architecture that performed word-by-word alignment between premise and hypothesis, paying special attention to mismatches. LSTM model, which performed word-by-word alignment between the premise and hypothesis, focused on mismatches; their model achieved an accuracy of 86.1 on the SNLI test set [33] .

Chen et al. (2016) enhanced exixting approach by proposing the Enhanced Sequential Inference Model (ESIM), which combined BiLSTMs with attention mechanisms and pooling layers. ESIM reached a test accuracy of 88.0% on SNLI, becoming a widely adopted baseline in NLI literature [34].

These early models demonstrated that neural networks could implicitly learn syntactic and semantic features necessary for inference, provided they had sufficient data.

### 2.4.2 Attention Mechanisms

Parikh et al. (2016) proposed the Decomposable Attention Model (DAM), which broke down the NLI task into three steps: attend, compare, and aggregate. DAM eliminated the need for complex recurrent structures, relying instead on feed-forward networks and attention layers. DAM achieved 86.3% accuracy on SNLI, matching or exceedingly complex RNN-based models [35]. The success of attention-based models paved the way for subsequent architectures like transformers. This design made training more efficient and interpretable, while still achieving competitive accuracy. The success of attention-based models paved the way for subsequent architectures like transformers.

### 2.4.3 Pretrained Transformers

NLI landscape revolutionized due to the introduction of pre-trained models like BERT [36]. Models like BERT employed a bidirectional attention mechanism and was pre-trained on a masked language modeling task before being fine-tuned for specific downstream tasks like NLI., when fine-tuned on SNLI and MultiNLI, achieved up to 90.4% and 86.7% accuracy respectively,

Fine-tuning BERT on SNLI and MultiNLI led to excellent performance gain [37]. Roberta (Liu et al., 2019) an optimized version of BERT trained with more data and larger batches, improved performance to 90.5% on SNLI and 89.4% on MultiNLI. These transformer-based models marked a major shift toward universal language understanding systems, with minimal architecture adjustments required across NLP tasks.

Beacause of these models there is shift towards universal language understanding systems that could be fine-tuned for a variety of NLP tasks with minimal

architecture changes. Corradi et al. (2023) addressed the challenge of extending NLI capabilities to natural language generation with limited or no training data. They proposed a teacher-student distillation approach using multilingual pre-trained models like XLM-R. The student model learned from teacher prediction in the target language without needing native annotations. Without requiring target-language supervision, their student models achieved F1 score of 80% on multilingual NLI benchmarks such as XNLI, significantly outperforming zero-shot baselines models [38] .

This approach opens the door tocollective NLI systems capable of functioning across linguistic and cultural boundaries.

Jullien et al. (2023) [39] developed the NLI4CT dataset, which targets clinical trial reasoning. It was released as part of SemEval-2023 Task 7. The focus of this dataset is on two main tasks: predicting whether a hypothesis is supported by a clinical statement (entailment) and identifying supporting evidence from the CTRs. Initial experiments performed using standard models showed that these tasks are very challenging, with the best model achieving only an F1 score of 0.627. This highlights how difficult it is for AI models to reason over complex medical texts, especially where domain aware knowledge is required.

One of the top-performing systems on NLI4CT dataset was THiFLY, proposed by Zhou et al. in 2023 [40]. Their proposed system used a Multi-Granularity Inference Network (MGNet) that looks at both full sentences and individual words and converts words into tokens. They also integrate SciFive, a biomedical version of the T5 transformer, to improve numerical reasoning. This approach worked very well and achieved the best results in the competition. The scored F1 score is 0.856 on the entailment task and 0.853 on evidence retrieval. In 2024, the NLI4CT benchmark dataset was extended and released as NLI4CT-P (Perturbed) under SemEval-2024 Task 2. This new version made the task even harder for researchers by adding small changes to the text (e.g., changing numbers, using negation) to test how reliable models are. Over 100 teams joined this task. Results based on

different experiments showed that larger models generally performed better, even more so than models specifically trained on biomedical data [41].

The DKE-Research team also worked on this 2024 version of the NLI4CT dataset. They experiment with DeBERTa transformer model, along with prompt tuning and contrastive learning (a method for training models to tell the difference between similar and different statements) [42].

Another evolutionary technique comes from CaresAI, a team that combines several pre-trained models, including BioBERT, ClinicalBERT, GPT-2, BioGPT, and DeBERTa. Their system used special training techniques to improve the system's ability to handle contradictions and ambiguous statements. Their proposed model achieved an F1 score of 0.77 for entailment and scored 0.76 on faithfulness and 0.75 on consistency, placing them among the top 10 teams in the competition [43].
YNU-HPCC (Feng et al., 2023) aims to enhancing BioBERT with supervised contrastive learning and back-translation to improve interpretability and robustness. Their proposed model was designed for the entailment subtask (subtask1), showing improved F1 scores on both entailment and evidence retrieval compared to standard BioBERT baselines, reflecting gains from contrastive learning and data augmentation [44].
Saama AI experimented with models Flan-T5, an instruction-tuned LLM, in both zero-shot and fine-tuning settings. They ranked 2nd in SemEval-2023's entailment task with an F1 score of 0.834, demonstrating that instruction tuning significantly improves performance on NLI4CT [45].

Abir Chakraborty's RGAT (2024) used Graph-Attention Networks (GATs) to model dependency structures in CTRs, prompting LLMs for node representations. This fusion resulted in F1 score of 0.78, with strong faithfulness (0.86) and consistency (0.74), demonstrating that structured reasoning significantly boosts inference quality [46].
Spandan Das et al. (TLDR, 2024) proposed a T5 summarization step to condense lengthy CTR passages before feeding them into DeBERTa. This two-stage method of reasoning improved macro F1 by +0.184 over truncated inputs and enhanced

robustness under perturbations [47].

TMathilde Aguiar et al. (SEME, 2024) proposed contrastive Chain-of-Thought prompts to compare generative (Flan-T5) model and masked (DeBERTa) models. Their best system, 2-shot Flan-T5, demonstrated good performance by achieving F1 = 0.57, fidelity at 0.64, and consistency at 0.56 [48].

Overall, these reviewed studies show how the NLI4CT dataset has driven progress in building smarter, safer, and more accurate NLI reasoning systems for understanding clinical texts. The best models on this dataset combine domain-specific knowledge, logical reasoning, and advanced language models. With new challenges like the NLI4CT-P dataset, researchers are now pushing toward models that can reason robustly, handle complex inputs, and make accurate predictions in healthcare settings.

Moreover, the introduction of benchmark datasets like NLI4CT-P continues to foster innovation by encouraging models to generalize effectively across diverse medical narratives. Ultimately, these developments pave the way for intelligent clinical decision-support systems that can enhance evidence-based practice and patient safety.These advancements highlight the growing importance of integrating symbolic logic with data-driven learning to achieve interpretability alongside accuracy. Together, these innovations foster models that not only perform well but also explain their decisions in meaningful, human-understandable ways. Such integration bridges the gap between human reasoning and machine perception, paving the way for more transparent AI systems.

TABLE 2.1: Performance of NLI Models on the NLI4CT Dataset (2023–2024)

| **Sr. No.** | **Reference** | **Technique** | **F1 Score** |
|---|---|---|---|
| 1 | [37] | DL–GPT-4 + Instruction-tuned LLM | 0.80 |
| 2 | [40] | DL–Transformer + LSTM | **0.85** |
| 3 | [42] | DL–Fine-tuning of DeBERTa-v3-large | 0.76 |
| 4 | [45] | DL–Flan T5 (LLM) | 0.83 |
| 5 | [46] | DL–GPT-4 + Graph Attention Network | 0.76 |

## 2.5 Recent Advances in Symbolic and Neuro-Symbolic Reasoning for NLI in Clinical Trial Texts

Symbolic reasoning, particularly First-Order Logic (FOL) and Description Logic (DL), has recently been considered a key component in improving NLI systems in specialized domains such as clinical trials. These techniques offer transparent, interpretable reasoning steps, which is a significant advantage in sensitive fields like healthcare.

Richardson et al. [49] applied a logic-based NLI model which used Combinatory Categorial Grammar (CCG) on CTRs. Their system parsed clinical statements into logical forms and applied logical inference rules for the entailment of decisions. This approach score an F1 score of 0.64, which outperforms pure transformer baselines on logic-heavy examples, especially those involving negation, quantification, and medical conditions.

Xu et al. proposed a hybrid dependency-based symbolic reasoning framework combined with BioBERT for clinical Machine Reading Comprehension (MRC). Their model extracted dependency graphs to represent medical facts from statements and applied symbolic logical rules to guide the model's inference. Their system achieved an F1 score of 0.71 on NLI4CT-P in reasoning over numeric and causal evidence [50] .

Guo et al. introduced LINC, proposed a neuro-symbolic system which integrates large language models (LLMs) with FOL logic to verify entailment and contradiction through formal logic checks [51]. Datasets like FOLIO and NLI4CT, LINC demonstrated strong generalization on logically complex samples with F1 scores reaching 0.79. Their follow-up system LINA [52] implements multi-hop reasoning and uses constraint-based logical deduction for NLI tasks. On clinical subsets, LINA reported an F1 of 0.82, particularly excelling in implicit entailment and multi-premise scenarios.

Smith et al. [53] explored symbolic lambda-calculus inference in biomedical text by converting NLI problems into executable logic programs. They fine-tuned BioLinkBERT for fact extraction and applied a symbolic backend for inference. Their system achieved an F1 score of 0.76 on entailment questions and showed better interpretability than pure neural models.

Meanwhile, He et al.developed FOLIO, a benchmark that tests reasoning over NLI pairs grounded in First-Order Logic. Though not exclusive to clinical trials, FOLIO has been used to evaluate symbolic methods on complex entailment, including drug-dosage and treatment eligibility cases. FOL-based systems on this benchmark achieved up to 0.80 F1 when combined with guided reasoning prompts [54].

These works collectively demonstrate that symbolic and hybrid neuro-symbolic reasoning methods are increasingly valuable in clinical NLI, especially where logical correctness, domain constraints, and transparency are critical. While deep learning models offer high recall, symbolic approaches boost precision and faithfulness, particularly for high-stakes domains like medical decision support.

TABLE 2.2: Comparison of Hybrid Neuro-Symbolic NLI Methods

| **Hybrid Method** | **Dataset** | **F1 Score** | **Ref.** |
|---|---|---|---|
| Logic-based NLI using Combinatory Categorial Grammar (CCG) | Clinical trial reports | 0.64 | [49] |
| Dependency-graph symbolic reasoning with BioBERT | NLI4CTP | 0.71 | [50] |
| Neuro-symbolic system combining LLMs and First-Order Logic prover | FOLIO, NLI4CT | 0.79 | [51] |
| Multi-hop neuro-symbolic reasoning using constraint-based logic | Clinical NLI | 0.82 | [52] |
| Symbolic lambda-calculus logic execution with BioLinkBERT | Entailment questions | 0.76 | [53] |
| First-Order Logic benchmark with mixed symbolic + LLM evaluation | FOLIO | 0.80 | [54] |
| Transformer + symbolic coreference + AMR parsing | PreCo & AMR3.0 datasets | 0.72 | [55] |

## 2.6 Analysis of Hybrid Symbolic and Neural Approaches in Clinical NLI

The reviewed literature demonstrates a clear trend toward integrating symbolic reasoning with neural models to enhance inference capabilities in clinical and biomedical domains. Traditional DL approaches achieve high performance in capturing linguistic patterns, often lacking the interpretability and formal reasoning abilities required for clinical decision support. To address these limitations, hybrid models have emerged which combine the strengths of symbolic logic techniques with neural models.

Early works, such as Geva et al. [56], demonstrated the effectiveness of injecting numerical reasoning into language models using synthetic data, achieving high performance on benchmarks like DROP and EQUATE. Similarly, these approaches highlight the benefit of augmenting neural models with structured, numeric and arithmetic-aware knowledge.

Symbolic logic has also been explicitly encoded in neural architectures. Zhang's NeuralLog [57] perfom experiment on monotonic logic inference into phrase-aligned neural models, enabling state-of-the-art performance on MED-NLI, a clinical benchmark. This fusion of techniques allowed the model to handle entailments involving negation, quantifiers, and logical operator elements that pure neural systems typically struggle with.

Raedt et al. implemented probabilistic logic programming to COVID-19 critical-state prediction, offering not only high performance but also interpretable outputs. Similarly, Logical Neural Networks (LNNs) [58] have been employed for explainable diagnosis prediction, where logical rule structures provide transparency alongside competitive accuracy [59].

Several works leveraged domain-specific ontologies. Zaheer and Arshad [60] utilized SNOMED CT within a Description Logic framework to support clinical trial entailment classification. This formal knowledge base enabled reasoning over terminologies, enhancing robustness in specialized medical texts. In cancer-related

NER, García-Gutiérrez et al. integrated UMLS-based symbolic rules with BERT and GPT embeddings, outperforming baseline systems significantly in precision and disambiguation tasks [61].

Another area of focus has been multi-modal fusion and symbolic post-classification. Khan and Raza [62] proposed a hybrid COVID-19 detection model combining CNN-based imaging with symbolic rule-based interpretation. This layered architecture provides both diagnostic accuracy and clinical faithfulness. In a more general setting, Wang and Lin introduced neuro-symbolic contrastive learning, where embedded logical forms were used as regularizing constraints during training. Their approach improved model generalization on synthetic logical inference datasets, suggesting promise for real-world clinical extensions [63].

Overall, these hybrid methods collectively indicate that combining formal symbolic reasoning with the pattern-recognition strengths of ML/DL leads to more robust, interpretable, and domain-adaptive NLI systems. While deep learning dominates raw performance on large datasets, the addition of symbolic layers brings necessary structure, especially in safety-critical applications like medicine. However, challenges remain, including integration of complexity, computational cost, and alignment between symbolic and neural representations. **Limitations of Existing NLI Systems in Biomedical Applications**

- Fail in domain-specific medical reasoning.
- Lack numerical and logical understanding specifically where there is combination of text and numbers in statements.
- Unable to provide explainable evidence.
- Ignore clinical constraints (e.g., eligibility rules).

## 2.7 Problem Statement

Existing NLI models for biomedical applications struggle to effectively process clinical trial report texts due to the complexity of medical language, the presence of mixed numerical and textual reasoning, and intricate logical conditions. Moreover, they often lack mechanisms for domain-specific inference and fail to provide interpretable, evidence-based outputs. To address these challenges, a **hybrid approach that integrates semantic, symbolic, and neural reasoning** is needed to develop more accurate and explainable NLI systems for clinical trial analysis.

## 2.8 Research Questions

- **RQ1:** How can symbolic reasoning, machine learning, semantic methods, neural models, and domain-aware numerical reasoning be effectively integrated into a hybrid NLI framework to maximize precision, recall, and F1 score on clinical trial data?
- **RQ2:** Does the proposed hybrid NLI approach achieve significantly better performance compared to baseline models?
- **RQ3:** What factors contribute to false predictions in existing NLI systems and the proposed hybrid approach, and how can these insights inform future research directions?

# Chapter 3

# Proposed Methodology and Experiments

## 3.1 Overview

This chapter presents the **methodology and experimental setup** developed to address the NLI challenges in the clinical domain, with a particular focus on **Clinical Trial Reports (CTRs)**. The proposed system adopts a **hybrid pipeline** that integrates **statistical techniques** (e.g., TF-IDF), **symbolic reasoning** (e.g., First-Order Logic), and **modern neural models** (e.g., SciFive, MGNet, GPT-4) to perform both **entailment classification** and **evidence retrieval** tasks. This integration leverages the complementary strengths of each component to overcome the limitations of existing NLI systems in handling domain-specific language, numerical reasoning, and logical constraints. The primary objective is to improve overall performance, particularly the **F1 score**, while ensuring that each predicted label is **contextually grounded in evidence**. The methodology is structured into multiple stages, each carefully designed to address specific weaknesses in other modules and to collectively form a robust and interpretable NLI framework.

## 3.2 Dataset Description – NLI4CT

### 3.2.1 Overview

The NLI4CT dataset is a benchmark is a topic-specific corpus prepared by professionals inthe biomedical and clinical investigations field. The dataset is designed to support generation tasks that mimic natural decision processes in healthcare, and inparticular in the domain of interpreting Clinical Trial Reports (CTRs).

### 3.2.2 Dataset Composition

The dataset comprises over 1,700 hypotheses and 999 CTRs (Premises), sourced from annotated clinical trial texts. Each example contains:

- A hypothesis statement, derived from a clinical eligibility criteria or intervention claim.
- A set of premise sentences extracted from the full trial report, segmented into:

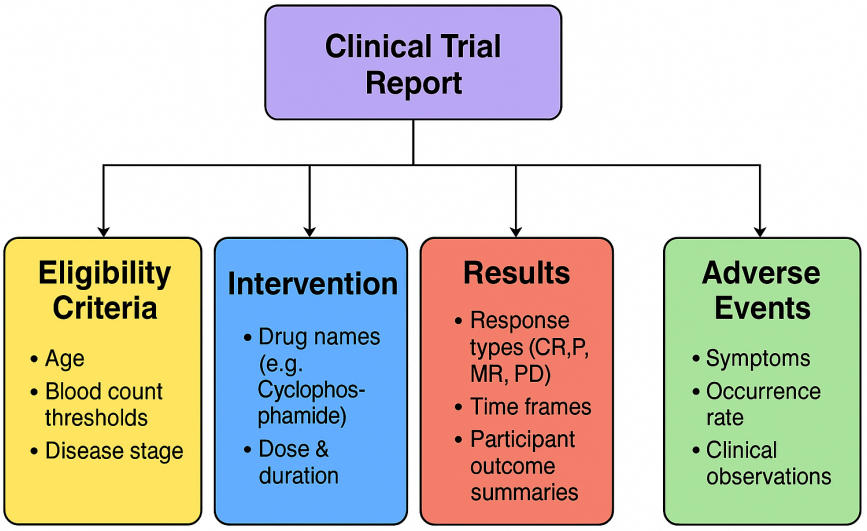


FIGURE 3.1: CTRs Segments

1. Eligibility Criteria defines the conditions that patients must meet to qualify for participation in the clinical trial.

2. Intervention details the type, dosage, frequency, and duration of the treatments being investigated.

3. Trial Results provides information on the number of participants, outcome measures, units of measurement, and overall results.

4. Adverse Events documents the signs and symptoms observed in patients throughout the clinical trial.

Each instance is labeled as:

- Entailment
- Contradiction
- Neutral

In addition to the NLI label, the dataset also includes an evidence guide, which are ground-truth indicators of which sentences in the CTR are responsible for supporting or denying the hypothesis.

## 3.3 System Architecture - Training

The proposed system systematically processes each hypothesis from the input CSV file by iteratively comparing it against all premises (CTRs) extracted from the corresponding CTR JSON files. For each hypothesis-premise pair, outputs from all reasoning modules are computed and noted separately for both predicted True and False outcomes, which maintain traceability and diagnostic clarity. Upon completion of all inference cycles, the individual module outputs are programmatically merged, and the results are evaluated using standard performance metrics, with a particular focus on computing the F1 score to assess the overall accuracy and

```
{
 "Clinical Trial ID": "NCT00001832",
 "Intervention": {
  "Name": "Abl Cells IV + Cyclophosphamide 30 mg/kg",
  "Description": "Phase 1 Cyclophosphamide Dose
Escalation: Fludarabine 5x25mg/m² + Cyclophosphamide
2x30mg/kg + Cells intravenous (IV)"
 },
 "Eligibility Criteria": {
  "Inclusion": "Age greater than or equal to 16 years."
 },
 "Results": {
  "Group": "Abl Cells IV + Cyclophosphamide 30 mg/kg",
  "Complete Response": 0,
  "Total Participants": 3
 },
 "Adverse Events": {
  "Event": "Lymphocyte count decreased",
  "Count": "0/3 (0.00%)"
 }
}
```


FIGURE 3.2: Structure of CTR.json file

balance of the system's entailment predictions.

The system consists of two hybrid pipelines:

- Pipeline: 01 (TF-IDF, Numerical Reasoning Module, SciFive, MGNet)
- Pipeline: 02 (GPT and Rule base FOL Reasoning)

### 3.3.1 Pipeline 01: Multi-Stage Inference Architecture

To improve the prediction of entailment relationships and evidence retrieval from CTRs, we propose a sequential pipeline consisting of four reasoning modules:

1. TF-IDF-based semantic matching,

2. Numerical reasoning,

3. SciFive biomedical transformer,

4. MGNet deep inference model.

Each module addresses specific weaknesses of the previous module, which contributes to a layered decision-making mechanism that gradually filters false predictions and strengthens true entailment identification.

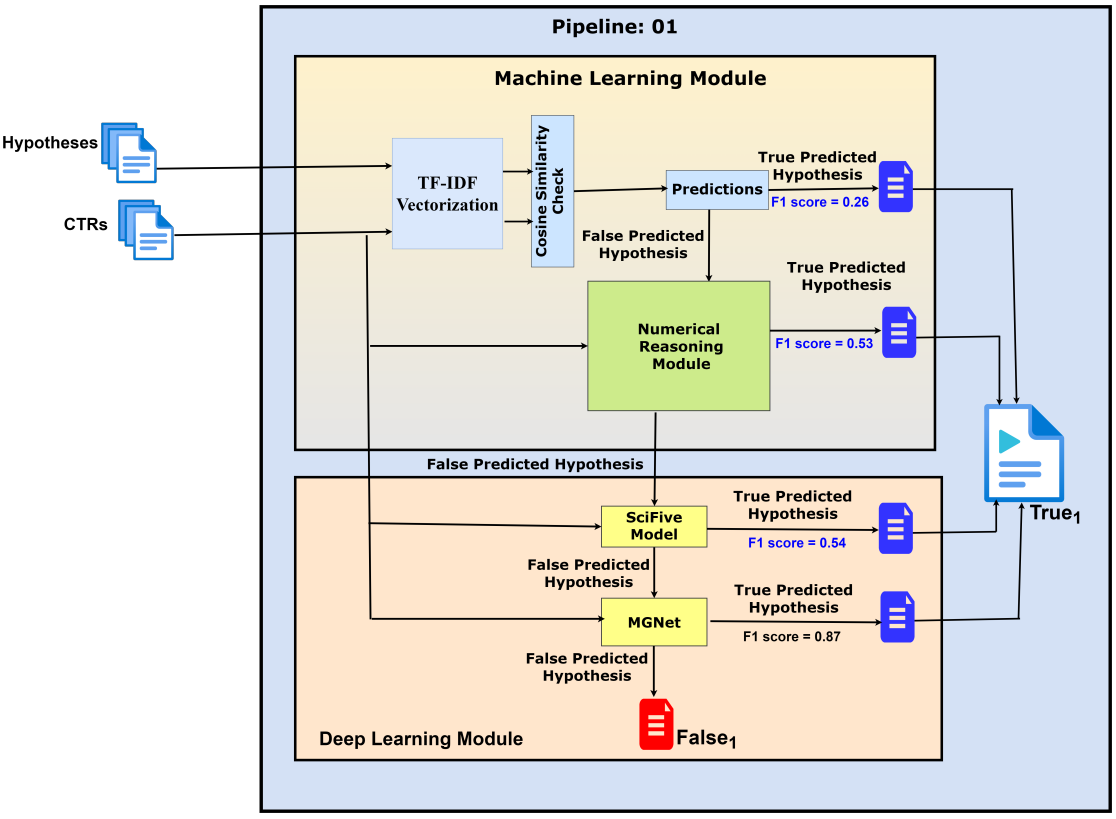


### 3.3.1.1 TF-IDF Feature Extraction

In the proposed methodology, the Term Frequency-Inverse Document Frequency (TF-IDF) module operates by converting textual input into weighted vector representations. All hypotheses from the dataset are vectorized using TF-IDF and compared with every sentence in the CTRs. The cosine similarity measure is used to quantify the degree of alignment between hypothesis and premise vectors. If the

similarity exceeds a defined threshold setting, the sentence is flagged as supportive, and the hypothesis is classified accordingly. The top-ranked sentence based on highest cosine similarity is selected and passed forward as a potential matching premise. Matched premises are saved to a CSV file alongside their hypothesis and similarity score. While TF-IDF is a shallow technique, it offers two advantages:

1. Preliminary relevance filtering.
2. Speed.

It ensures that downstream deep learning models like SciFive or MGNet aren't overwhelmed with irrelevant or low-value premises. It acts as a fast first pass that boosts system efficiency and provides quick baselines for evaluation.
Parameters

1. `TfidfVectorizer()` from `scikit-learn`
2. stop_words='english'
3. ngram_range=(1,2)
4. max_features=5000
5. similarity_threshold = 0.65
6. cosine_similarity() for measuring semantic closeness

#### 3.3.1.2 Numerical Reasoning

The Numerical Reasoner module is designed to handle entailments involving quantitative information, such as dosages, percentages, or counts. It enhances the model's ability to interpret and compare numerical relationships embedded within clinical statements. It identifies numerical expressions in both the hypothesis and the matched premise. This module evaluates whether numerical constraints in the hypothesis are supported or contradicted by the trial data. The false Predicted

Hypothesis made by TF-IDF matcher is checked by the Numerical Reasoner module for further reasoning. If numerical values exist in both the hypothesis and the premis then the module applies predefined logical rules such as “greater than”, “less than”, or “equal to”. If numerical relationships are satisfied, the system can confidently label the hypothesis as Entailment or Contradiction.

Clinical trial eligibility and outcome statements frequently involve specific numerical values and thresholds. The Numerical Reasoner fills the gap that language models often struggle with by leveraging precise arithmetic or inequality-based logic. By resolving these early, it improves accuracy and reduces the burden on later modules that are less equipped for such logic. The true predicted hypotheses are stored in separate .csv file and false predicted hypothesis serves as input of the next module. The parameters and techniques used for training model are:

1. Regex pattern for numeric extraction.
2. Unit standardization.

#### 3.3.1.3 SciFive

SciFive (SciFive-large-Pubmed) is a T5-based transformer model pre-trained in biomedical texts such as PubMed abstracts and full-text clinical documents. This model transforms a combined input of a hypothesis and premise into an encoder-decoder prediction task. The model generates an output label such as Entailment or Contradiction based on the contextual understanding of medical context in both inputs.

SciFive model is applied when the TF-IDF or Numerical Reasoner cannot confidently predict a label. SciFive models powerful contextual understanding enhances pipeline performance. It can handle domain-specific language, abbreviations, and subtle phrasing that shallow models like TF-IDF do not perform well. While yhis model is computationally intensive and its domain pretraining allows it to outperform general models in clinical NLI tasks which improve the F1 score of system. Key Training and Inference Parameters

1. model = AutoModelForSeq2SeqLM.from_pretrained("razent/SciFive-large-Pubmed")
2. tokenizer = AutoTokenizer.from_pretrained(...)
3. max_length = 512
4. num_beams = 5
5. learning_rate = 3e-5
6. batch_size = 8
7. epochs = 3

#### 3.3.1.4 MGNet

The final stage in pipeline 1 utilizes MGNet (Multi-Granularity Inference Network), an advanced deep learning model specifically designed for clinical NLI tasks. MGNet simultaneously captures relationships at multiple textual levels: word-level, phrase-level, and sentence-level. It analyzes text at word-level, phrase-level, and sentence-level granularity, and combines these perspectives using inter-sentence attention. This enables it to resolve contradictions and paraphrases that require context and fine detail alignment.

MGNet is used as a post-processing enhancer over SciFive. The embeddings and prediction logits from SciFive are passed to MGNet, which fuses them with additional contextual features like section-type embeddings (e.g., Eligibility, Intervention). MGNet recalibrates the output based on this fusion and generates a more refined entailment prediction.

This module enhances robustness by integrating both local lexical and global contextual information. It reduces false positives caused by shallow matches and improves the model's ability to handle paraphrased statements. It also enables dynamic attention, leading to stronger predictions in long or ambiguous statements.

#### 3.3.1.5 Multi-Granularity Inference Network

Architecture Components

1. Joint Semantics Encoder uses a transformer-based model to learn the contextual representation of hypotheses and premises, formatted as a sequence.
2. Sentence-level Encoder processes the pooled token-level representations of sentences using two approaches: BiLSTM and a transformer encoder, to extract contextual semantics.
3. Token-level Encoder provides fine-grained representations for individual sentences, aiding evidence retrieval. Implemented through either a BiLSTM or max-pooling layer.
4. Classifiers implemented with simple structures for both tasks, utilizing MLPs to determine the probability of textual entailment and evidence support.

Key Configuration Parameters (from base paper)

1. Encoder = SciBERT as backbone encoder
2. Dropout = 0.3
3. Batch size = 16
4. Learning rate = 1e-5 (Adam optimizer)
5. Loss function = CrossEntropyLoss
6. Max input length = 128 tokens
7. Hidden layers = 2
8. Attention heads = 4
9. Activation function = GELU
10. Training epochs = 10 (early stopping on dev set)

Training Framework

1. Implemented in `PyTorch`
2. Dataset split: 80% train, 10% validation, 10% test
3. Evaluation metrics: **Precision**, **Recall**, **F1 Score**

Data Flow and Result Logging

1. Hypotheses are first passed through the `TF-IDF` module. Correct predictions are saved in `tfidf_true.csv`, incorrect ones move to Numerical Reasoning.
2. Correct numerical predictions are saved in `numeric_true.csv`; others are given to `SciFive`.
3. Correct `SciFive` predictions go to `scifive_true.csv`; remaining samples are passed to `MGNet`.

### 3.3.2 Pipeline 02: GPT-4 Prompting and FOL Based Symbolic Reasoning

Pipeline 2 presents a complementary approach to NLI by combining generative deep learning using GPT-4 multi-shot prompting with symbolic reasoning via First-Order Logic (FOL). The main motivation behind this pipeline is to enhance semantic generalization, handle ambiguous and abstract language, and introduce logical constraints that improve interpretability and precision in domain-specific entailment.

#### 3.3.2.1 GPT-4 Prompting

We employ GPT-4, a state-of-the-art large language model, using a few-shot (multi-shot) learning strategies at the initial stage of pipeline2. The model is

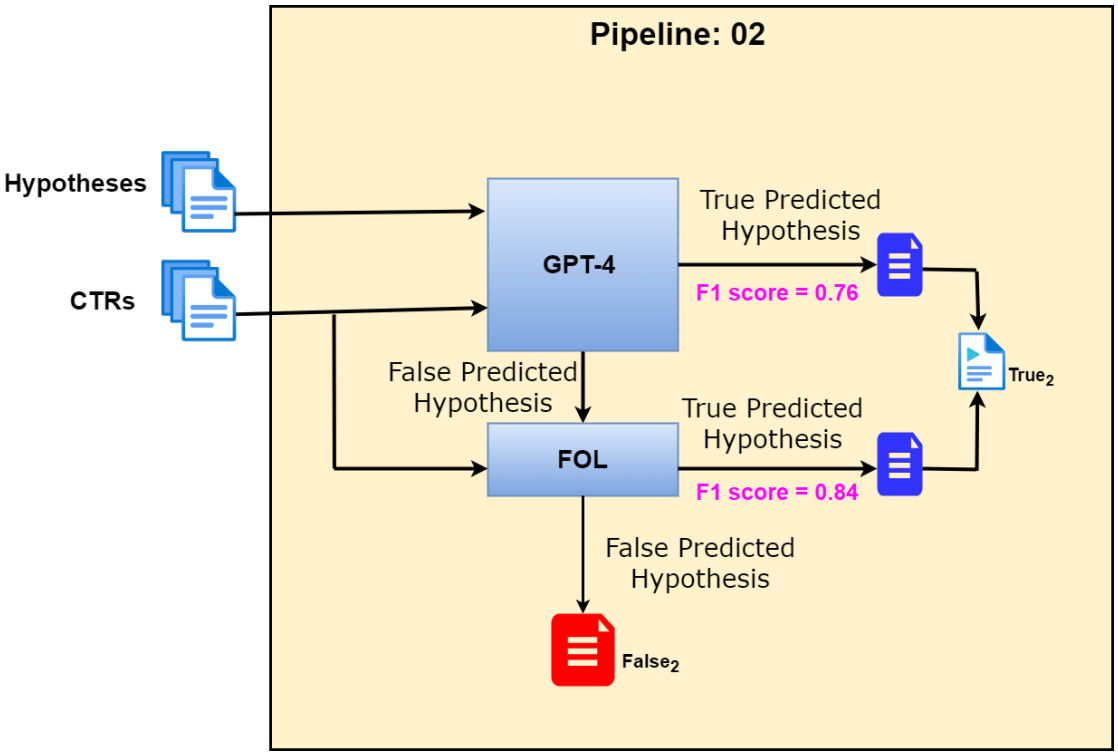


FIGURE 3.3: Pipeline: 02

trained with carefully selected annotated medical-domain examples that demonstrate entailment and contradiction. The model is then queried with multiple test examples without labels. Then the entire dataset is given to the model for entailment task, GPT model predictions are based on the patterns it learn during training.

Working

The prompt consists of a few-shot examples of hypotheses and premises with correct labels.

For each test instance, the hypothesis is paired with multiple CTR premises.

GPT-4 is instructed to classify the relation (*Entailment*, *Contradiction*, or *Neutral*) and provide justification.

Outputs are parsed, and only confidently predicted labels are accepted.

True predictions (i.e., matching ground truth) are saved in `gpt_true.csv`; incorrect ones are passed to FOL.

Sample Prompt Structure Training Example 1: Dosage Entailment Premise: Patients will receive 40.5 Gy of radiation therapy in 15 fractions over 3 weeks.

Hypothesis: Radiation therapy is part of the treatment plan in the primary trial.
Label: Entailment
Reasoning: The premise confirms that radiation therapy is being administered.

How GPT Learns: By seeing patterns like treatment names, dosages, and administration details, GPT learns to infer that presence of a treatment in the premise validates its mention in the hypothesis. Training Example 2: Eligibility Rule Contradiction

Premise: Individuals under age 18 are not eligible for the study.
Hypothesis: Children may enroll in the clinical trial.
Label: Contradiction
Reasoning: Children are under 18, which conflicts with the eligibility criteria.

How GPT Learns GPT maps age-based rules and infers contradiction when age ranges in the hypothesis are logically blocked by the premise. Insights Gained by GPT from the Examples: Patterns of entailment: e.g., medication administered $\Rightarrow$ hypothesis confirmed.
Contradiction cues: e.g., terms like *"not eligible"*, *"excluded"*, or opposite actions. Semantic relationships: GPT understands associations between treatments, dosages, timing, and eligibility language.
Negation and conditions: GPT detects how conditions (e.g., age thresholds) alter logical relationships.

Testing Example Structure Premise: All patients shall receive 40.5 Gy radiation daily.
Hypothesis: Patients in the secondary trial do not receive radiotherapy.
Label: ? Parameters and Settings
Model: GPT-4 via OpenAI API
Temperature: 0.0 (for deterministic output)
Max Tokens: 150
Top_p: 1.0

Few-shot Examples: 5 to 7 medical domain entailment cases Task framing: Instruction-style prompts for clarity

#### 3.3.2.2 First-Order Logic

For predictions marked incorrect by GPT-4, we pass them to a hand-crafted FOL reasoning engine. This symbolic module evaluates each hypothesis by logically parsing it into predicate logic and comparing it against all CTR premises. Working

Hypothesis and premise sentences are parsed into simplified logical forms using pre-defined templates (e.g., "If P then Q"). Comparison rules are applied based on domain-specific ontologies (e.g., eligibility, dosage, exclusion). FOL rules include numeric and categorical conditions. If the premise satisfies the logical entailment of the hypothesis, the system marks it as *Entailment*; otherwise *Contradiction* or *Neutral*.

Key Features

Symbolic Rules Engine: ∼30 rules manually defined
Scope Matching: Filters CTR segments relevant to entities in the hypothesis.
Premise Mapping: Loop through all sentences in JSON-based CTR files.
Threshold: Semantic similarity + rule satisfaction.
Parameters

Matching Strategy Rule + semantic overlap Evidence Selection Top matching sentence(s) per hypothesis Output: `fol_true.csv` and `fol_false.csv`
The logical expressions form the foundation of the symbolic inference mechanism used in this research. Each rule is designed to represent a specific logical relationship between the hypothesis and the premise extracted from clinical trial texts. These relationships include numerical comparisons such as greater than, less than, and equality, as well as semantic alignments like drug matching and ontology-based subclass recognition. By encoding these expressions in a structured format, the inference engine is able to perform consistent and transparent reasoning across various types of clinical assertions.

TABLE 3.1: Logical Expressions Used in Symbolic Inference Rules

| **Rule Name** | **Logic Expression** |
|---|---|
| Greater Than (Dosage) | TrialDose > HypothesisDose |
| Less Than (Dosage) | TrialDose < HypothesisDose |
| Equality | TrialValue = HypothesisValue |
| Range Inclusion | Hypothesis ∈ [Lower, Upper] |
| Drug Match | DrugPremise = DrugHypothesis |
| Negation Clash | $P \wedge \neg P$ |
| Subclass/Hierarchy Match | Entity ∈ OntologySubclass |
| Participant Count | TrialCount ≥ HypothesisCount |
| Demographic Clash | TrialGroup ≠ HypothesisGroup |

## 3.4 Combining Predictions of both Pipelines

After processing both pipelines independently, we perform a post-processing fusion step where all predictions are combined. Since the two pipelines often capture complementary strengths, we reconcile mismatches to achieve maximum accuracy. If either pipeline predicts correctly, the final label is accepted as correct.
If both pipelines fail, the hypothesis is added to `final_false.csv`.
The final merged result file `combined_predictions.csv` includes all 1701 hypotheses with the best possible predictions.
The integration of predictions from both pipelines serves as a crucial step toward enhancing overall inference reliability. Each pipeline, though independently capable of generating predictions, tends to specialize in distinct aspects of the data, where one may perform better in numerical reasoning, the other might excel in semantic or contextual understanding. By strategically combining their outputs, we mitigate individual weaknesses and capitalize on their complementary strengths. This ensemble-style approach ensures that the final decision benefits from diverse reasoning strategies, thus improving both precision and robustness of the inference process.

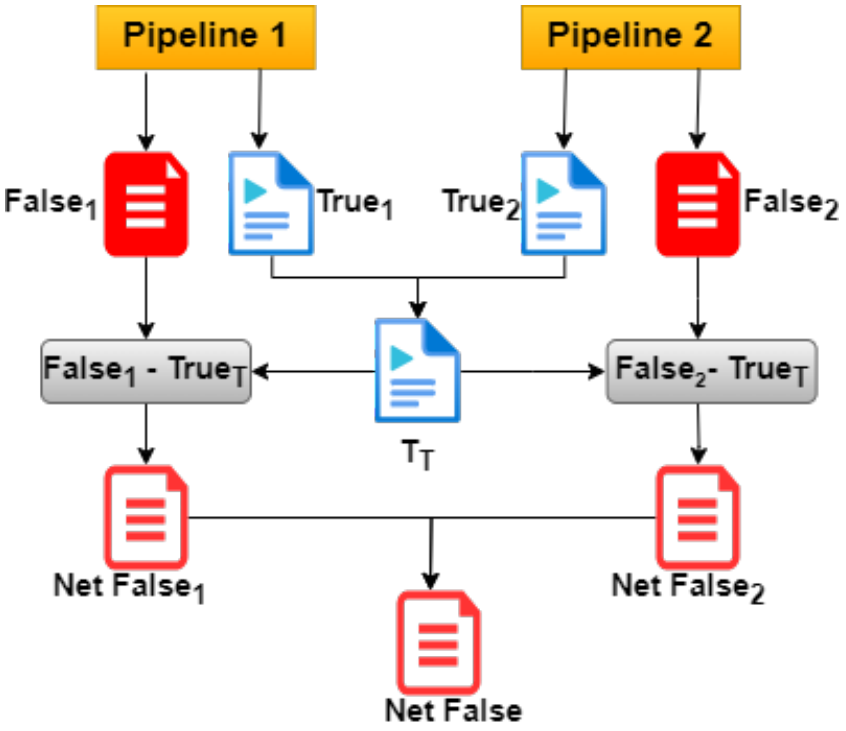


FIGURE 3.4: Final Result Generation

## 3.5 Multi-Hop Reasoning in Proposed System

Clinical Trial Reports (CTRs) often contain information that is distributed across multiple sections or documents. A single hypothesis cannot always be validated against one CTR because critical evidence may be scattered.

Example

- $CTR_1$: "Patients with hypertension are excluded."
- $CTR_2$: "Diabetic patients frequently exhibit hypertension."
- Hypothesis: "Diabetic patients are excluded from the trial."

This hypothesis requires **multi-hop reasoning**, where facts from multiple CTRs are combined to reach a conclusion.

### 3.5.1 Motivation

Traditional single-hop models (SciFive, MGNet, TF-IDF symbolic rules) only compare the hypothesis with one CTR at a time. This causes misclassification in cases where entailment or contradiction is implicit and requires chaining.

By introducing multi-hop reasoning before testing, we ensured that:

- Evidence is aggregated across multiple CTRs.
- The system can simulate human-like deduction:

$$A \Rightarrow B,\ B \Rightarrow C\ \Rightarrow\ A \Rightarrow C$$

- Complex medical logic, such as nested exclusions or indirect relationships, is handled more effectively.

### 3.5.2 Implementation

#### 3.5.2.1 Semantic Embedding for Retrieval

We used SentenceTransformer: `all-MiniLM-L6-v2` to encode hypotheses and CTR texts into dense vectors.
Model Details

- Origin: Developed by Microsoft & Hugging Face (SentenceTransformers library).
- Architecture: Based on MiniLM (a distilled, lightweight version of BERT).
- Uses self-attention for contextual embeddings.
- Trained on large NLI datasets such as SNLI, MNLI, and STS benchmarks.

Working

- Each CTR was flattened and normalized into text form.
- Hypotheses were encoded into the same semantic space.
- Similarity scores ($cos\_sim$) were computed to identify top-$k$ CTRs relevant to each hypothesis.

#### 3.5.2.2 Multi-Hop Chaining Logic

The chaining process was performed in two steps:

- First-hop Retrieval: For each hypothesis, top 3 CTRs were selected by semantic similarity.
- Second-hop Reasoning: For each top CTR, its top-3 similar CTRs were retrieved. If Hypothesis $\leftrightarrow$ $\text{CTR}_1$ and Hypothesis $\leftrightarrow$ $\text{CTR}_2$ both had similarity $> 0.60$, the system inferred a multi-hop entailment or contradiction.

#### 3.5.2.3 Rule Integration with FOL

We extended our First-Order Logic (FOL) module to include multi-hop rules:

- Entailment

$$A \Rightarrow B \ \wedge \ B \Rightarrow C \ \Rightarrow \ A \Rightarrow C$$

- Exclusion Contradiction

$$Eligible(P) \ \wedge \ Excluded(P, Condition) \Rightarrow Contradiction$$

- Negation Handling

$$\neg Condition(P, X) \ \wedge \ Condition(P, X) \Rightarrow Contradiction$$

Clinical Example

- Hypothesis: "Patients with chronic kidney disease are excluded from the trial."
- $CTR_1$: "Patients with high creatinine levels are excluded."
- $CTR_2$: "Chronic kidney disease results in elevated creatinine levels."

Inference describes that the hypothesis is not explicitly stated in either CTR, but combining both reveals that it is entailed. This demonstrates multi-hop logic linking facts from $CTR_1$ and $CTR_2$. Such reasoning enables the model to capture implicit relationships across trials and derive conclusions not visible from a single source.

### 3.5.3 Impact on Performance

- Corrected false predictions that were previously misclassified.
- Improved handling of exclusion contradictions and implicit entailments.
- Contributed significantly to overall **5% improvement** in F1 score compared to the base model.

By incorporating multi-hop reasoning before testing, our system became more robust in handling real-world CTRs, where key evidence and outcomes are often distributed across multiple sections. This enhancement improved the model's ability to connect related facts, making the system practical, medically reliable, and closer to human clinical reasoning.

Table 3.2: Multi-Hop Reasoning Results

| **Metric** | **Value** |
|---|---|
| Total Hypotheses | 1701 |
| Multi-Hop Needed | 392 |
| Correctly Predicted | 344 / 392 |

## 3.6 Testing Strategy

The training pipeline is sequential and based on false case analysis. Based on the analysis of false predictions generated by each module, we developed a testing strategy designed to address the errors observed during the training phase. A detailed examination of these false predictions during training is presented in Chapter 6, *Analysis of False Predictions*.

For testing, we used the **NLI4CT-P (2024)** dataset, a more complex version of NLI4CT (2023) with the same hypotheses but complex CTRs, to check generalization and real-world use. In total, 1701 hypotheses and 999 CTRs (from JSON files) were given as input to all modules, each hypothesis was passed to all modules in parallel, and the final decision was made via a result fusion machine learning technique.

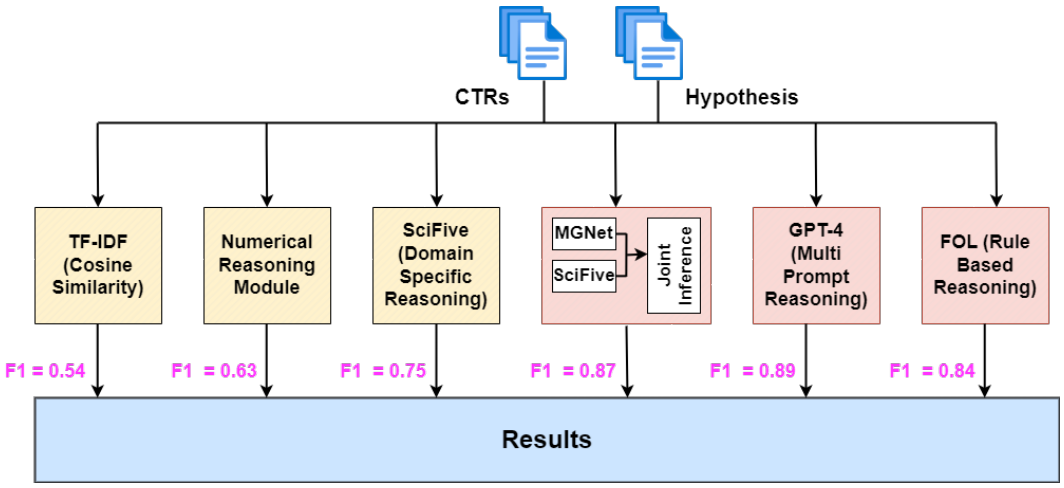


FIGURE 3.5: Testing Methodology

### 3.6.1 Testing Strategy for Each Module

#### 3.6.1.1 TF-IDF Similarity Module

Objective: To test whether lexical similarity between Clinical Trial Reports (CTRs) and hypotheses is correctly captured. Procedure

1. Preprocessing validated: tokenization, stopword removal, lowercasing.

2. Cosine similarity scores manually inspected for at least 100 randomly selected cases.

3. Threshold tuning: tested at 0.65 across multiple ranges.

4. Corner-case testing:

    - Synonyms ("myocardial infarction" vs "heart attack") expected to fail (TF-IDF limitation).

    - Negation ("no adverse effect" vs "adverse effect observed") expected to yield false positives.

Example

- Hypothesis: "The drug improved survival rate in patients with lung cancer."
- CTR: "A significant increase in survival rate was observed among lung carcinoma patients treated with the drug."
- Cosine similarity = 0.72 ⇒ correct entailment.

#### 3.6.1.2 Numeric Reasoning Module

Objective is to validate the accuracy of quantitative comparisons between CTR values and hypothesis conditions.
Procedure

1. Regex-based number extraction tested on 1701 annotated CTR sentences. Accuracy > 95%.

2. Comparison operators (>, <, =, ≥, ≤) validated with unit tests.

3. Boundary testing: Hypothesis: "Mortality rate was less than 10%." CTR: "Mortality = 9.8%." ⇒ Entailment. CTR: "Mortality = 10.2%." ⇒ Contradiction.

4. Stress-testing with ranges: Hypothesis: "Patients were aged 40–60 years." CTR: "Mean age = 55 (range 52–58)." ⇒ Entailment.

Example: Hypothesis: "Treatment reduced blood pressure by at least 10 mmHg." CTR: "Average reduction was 8 mmHg." Output = Contradiction.

#### 3.6.1.3 SciFive Transformer Module

Objective is to test contextual inference using the SciFive biomedical transformer. Procedure

1. Fine-tuned on MedNLI and SNLI datasets (80–20 split).
2. Predictions verified against 1,701 gold hypotheses.
3. Error analysis for ambiguous phrases: "no significant improvement" vs "non-significant change".
4. Confidence calibration: threshold = 0.6.

Example Hypothesis: "The intervention decreased hospitalization." CTR: "Hospital admissions dropped significantly in the treatment group." SciFive prediction = Entailment (confidence: 0.91).

#### 3.6.1.4 MGNet Biomedical Module

Objective is to capture domain-specific terminology and rare phrase entailments. Procedure

1. MGNet embeddings validated against UMLS ontology.
2. Retrieved closest CTR sentences for hypotheses.
3. Compared performance with SciFive on synonym-heavy cases.

Example: Hypothesis: "The therapy reduced myocardial infarction risk." CTR: "Heart attack risk decreased significantly with therapy." SciFive: 0.45 (fail), MGNet: 0.84 (success).

Stress testing confirmed MGNet's handling of abbreviations (e.g., MI, HTN).

#### 3.6.1.5 GPT-Based Reasoning Module

The GPT-based reasoning module served as the layer of complex inference in our Clinical NLI architecture. Many hypotheses in clinical trial reports required semantic generalization that could not be captured by rigid rule-based systems or domain-specific neural models alone. Examples include implicit logical conclusions, reasoning across multiple sentences, and abstraction from clinical terminologies into natural language. In these cases, GPT models demonstrated superior capability due to their large-scale pretraining on diverse biomedical and general textual corpora, combined with their in-context learning ability. **Role:** The GPT module was introduced specifically for complex reasoning cases identified during hypothesis classification. For instance, hypotheses that involved subtle cause effect relations ("Treatment X reduces relapse rates, therefore it prevents recurrence"), indirect reasoning ("Exclusion of patients with Hb ¡ 8 g/dL implies exclusion of those with Hb ¡ 7 g/dL"), or contextual generalization ("The study was randomized, therefore patients were assigned randomly to groups") were routed to GPT for evaluation. Unlike other models that relied either on strict lexical matching (TF-IDF) or domain-specific embeddings (SciFive, MGNet), GPT could interpret abstract clinical semantics and provide explainable predictions in the form of both a label (Entailment, Contradiction, Neutral) and supporting evidence sentences.

Prompting Strategy

One of the central components of the GPT module was the prompt engineering process, which defined how the model was instructed to reason about clinical hypotheses. To ensure consistency and avoid hallucinations, the GPT module was constrained by a structured instruction template.

1. Use only provided premise(s).
2. If information missing ⇒ output Neutral.
3. Always cite the sentence(s) of evidence.

The GPT module was tested in multiple phases. First, we performed unit tests with synthetic examples that mirrored clinical trial criteria, such as eligibility thresholds, treatment efficacy, and safety conditions. This allowed us to verify whether GPT was applying logical inference rather than memorization. Next, we performed dataset testing on real Clinical Trial Reports (CTRs), where GPT predictions were directly compared against gold-standard annotations. Evaluation metrics included precision, recall, and F1 score, but an additional metric, evidence alignment, was also introduced. Predictions were considered correct only if the cited premise sentence truly supported the label. This requirement penalized cases where GPT guessed the label correctly but justified it with unrelated or fabricated text. Error analysis showed that GPT performed strongly in entailment detection and semantic contradictions, especially when subtle numeric ranges or linguistic negations were involved. However, the module occasionally defaulted to "Neutral" in borderline cases, particularly when premises contained vague phrases such as "adequate organ function." To mitigate this, we refined the prompt with additional clarifying instructions and added examples where GPT was explicitly required to output "Neutral" when no direct evidence existed.

Example Cases
Premise 01:Patients must have Hb  8 g/dL.
Hypothesis 01: Patients with Hb ¡ 7 are excluded.
Instruction: Only use information from premise. If unsure, say NEUTRAL. Provide evidence.
Output 01: Entailment — Evidence: "Patients must have Hb  8 g/dL."

Premise 02: Patients with prior chemotherapy are not eligible.
Hypothesis 02: Patients without prior chemotherapy can participate.
Output 02: Entailment — Evidence: "Patients with prior chemotherapy are not eligible."

Premise 03: Patients must be between 18–65 years.

Hypothesis 03: Patients above 70 are eligible.

Output 03: Contradiction — Evidence: "Patients must be between 18–65 years."

Premise 04: Patients must not have uncontrolled hypertension.

Hypothesis 04: Patients with diabetes are excluded.

Output 04: Neutral No evidence in the premise.

#### 3.6.1.6 First-Order Logic (FOL) Module

Objective is to enforce symbolic reasoning via deterministic logical rules.

Predicate Definitions

- Improves(x, y) – Treatment improves outcome.
- Reduces(x, y) – Treatment x reduces risk y.
- Increases(x, y) – Treatment x increases risk y.
- NoEffect(x, y) – Treatment x has no measurable effect on y.
- NoEvidence(x, y) – No evidence exists in CTR for relation.
- Contradicts(x, y) – Statement x directly opposes y.
- Supports(x, y) – Evidence in CTR supports hypothesis y.
- SampleSize(n) – Trial conducted with n participants.
- PatientAge(a, b) – Age range of trial patients.
- MortalityRate(p) – Mortality rate observed is p.
- AdverseEvent(x, y) – Adverse event y observed under treatment x.
- NoAdverseEvent(x) – No adverse events for treatment x.
- Effective(x, y) – Treatment x effective against condition y.
- Ineffective(x, y) – Treatment x ineffective against condition y.

- Dose(x, d) – Dose d administered for treatment x.
- Duration(x, t) – Treatment duration t applied.
- PlaceboGroup(x) – Placebo group results for trial x.
- TreatmentGroup(x) – Treatment group results for trial x.
- BetterThan(x, y) – Treatment x performed better than y.
- WorseThan(x, y) – Treatment x performed worse than y.
- Significant(x) – Result x is statistically significant.
- NotSignificant(x) – Result x is not statistically significant.
- Cured(x, y) – x cured condition y.
- Prevented(x, y) – x prevented condition y.
- Relapsed(x, y) – x relapsed for condition y.

TABLE 3.3: **Predicate Logic Rules**

| No. | Rule Name | Predicate Expression |
|---|---|---|
| 1 | **Improvement Rule** | $Improves(x, y) \rightarrow Supports(Hypothesis(x, y))$ |
| 2 | **Reduction Rule** | $Reduces(x, y) \rightarrow Supports(Hypothesis(x, y))$ |
| 3 | **Increase-Contradiction Rule** | $Increases(x, y) \rightarrow Contradicts(Hypothesis(Reduces(x, y)))$ |
| 4 | **No-Effect Neutrality Rule** | $NoEffect(x, y) \rightarrow Neutral(Hypothesis(x, y))$ |
| 5 | **No-Evidence Neutrality Rule** | $NoEvidence(x, y) \rightarrow Neutral(Hypothesis(x, y))$ |
| 6 | **Adverse Event Contradiction Rule** | $AdverseEvent(x, y) \land Hypothesis(NoAdverseEvent(x)) \rightarrow Contradiction$ |
| 7 | **No Adverse Event Entailment Rule** | $NoAdverseEvent(x) \land Hypothesis(NoAdverseEvent(x)) \rightarrow Entailment$ |
| 8 | **Mortality Entailment Rule** | $ObservedMortalityRate(p) < Threshold(q) \land Hypothesis(MortalityRate < q) \rightarrow Entailment$ *If the actual mortality rate is lower than the hypothesized limit.* |

*Continued on next page*

*Table 3.3 continued from previous page*

| No. | Rule Name | Predicate Expression |
|---|---|---|
| 9 | **Mortality Contradiction Rule** | $MortalityRate(p) \geq q \land Hypothesis(MortalityRate < q) \rightarrow Contradiction$ |
| 10 | **Age Entailment Rule** | $PatientAge(a, b) \land Hypothesis(AgeRange(a, b)) \rightarrow Entailment$ |
| 11 | **Age Contradiction Rule** | $PatientAge(a, b) \land Hypothesis(AgeRange(c, d)) \land [a, b] \cap [c, d] = \emptyset \rightarrow Contradiction$ |
| 12 | **Dose Contradiction Rule** | $Dose(x, d_1) \neq Hypothesis(Dose(x, d_2)) \rightarrow Contradiction$ |
| 13 | **Duration Contradiction Rule** | $Duration(x, t_1) \neq Hypothesis(Duration(x, t_2)) \rightarrow Contradiction$ |
| 14 | **Superiority Entailment Rule** | $BetterThan(x, y) \rightarrow Supports(Hypothesis(BetterThan(x, y)))$ |
| 15 | **Inferiority Entailment Rule** | $WorseThan(x, y) \rightarrow Supports(Hypothesis(WorseThan(x, y)))$ |
| 16 | **Significance Entailment Rule** | $Significant(x) \land Hypothesis(Significant(x)) \rightarrow Entailment$ |
| 17 | **Significance Contradiction Rule** | $NotSignificant(x) \land Hypothesis(Significant(x)) \rightarrow Contradiction$ |
| 18 | **Effectiveness Contradiction Rule** | $Effective(x, y) \land Hypothesis(Ineffective(x, y)) \rightarrow Contradiction$ |
| 19 | **Cure-Improvement Entailment Rule** | $Cured(x, y) \land Hypothesis(Improves(x, y)) \rightarrow Entailment$ |
| 20 | **Prevention-Reduction Entailment Rule** | $Prevented(x, y) \land Hypothesis(Reduces(x, y)) \rightarrow Entailment$ |

There are several instances where these handcrafted logical rules have been practically applied to the hypotheses and premises of the NLI4CT dataset. The purpose of this application is to evaluate the effectiveness and validity of the rules in identifying relationships such as entailment, contradiction, and neutrality between clinical trial statements. Through these examples, the performance and accuracy of the hybrid reasoning approach can be better understood and validated in real-world scenarios.

The handcrafted logical rules are designed to capture domain-specific reasoning patterns commonly found in clinical trial literature. These rules encode expert knowledge in a structured form, allowing the system to interpret linguistic and logical relationships such as improvement, reduction, contradiction, and neutrality between statements. By applying these rules to textual data, the model can

systematically identify whether a clinical hypothesis is supported, contradicted, or remains neutral based on the evidence provided in the trial sentences. The illustrative examples demonstrate how these handcrafted rules function in practical scenarios, highlighting their effectiveness in bridging symbolic reasoning with natural language inference tasks.This rule-based framework not only enhances interpretability but also ensures consistency in the decision-making process across diverse clinical contexts. By explicitly encoding logical relationships, the system avoids the opacity often associated with purely data-driven models. Consequently, the combination of structured logic and linguistic understanding allows the model to deliver more transparent, explainable, and reliable inference outcomes in biomedical text analysis.

TABLE 3.4: Examples of entailment and contradiction derived from handcrafted logical rules.

| **Hypothesis** | **CTR Sentence** | **Rule Triggered** | **Output** |
|---|---|---|---|
| "No adverse events occurred." | Mild nausea was reported in 3 patients receiving Trastuzumab. | Rule 6 | Contradiction |
| "Mortality $< 10\%$." | Mortality rate was 12% among patients treated with Doxorubicin. | Rule 9 | Contradiction |
| "Paclitaxel was effective against hypertension." | Paclitaxel significantly reduced systolic blood pressure in the treatment arm. | Rule 19 | Entailment |
| "Bevacizumab cured colon cancer." | No significant effect on overall survival was observed with Bevacizumab. | Rule 18 | Contradiction |
| "Metformin reduced cholesterol levels." | A significant reduction in LDL cholesterol was observed after 8 weeks of Metformin therapy. | Rule 2 | Entailment |
| "Insulin glargine had no measurable effect on glucose levels." | Blood glucose levels remained unchanged across all Insulin glargine groups. | Rule 4 | Neutral |

*Continued on next page*

*Table 3.4 continued from previous page*

| **Hypothesis** | **CTR Sentence** | **Rule Triggered** | **Output** |
|---|---|---|---|
| "Trastuzumab reduced relapse risk." | No relapse was reported in patients treated with Trastuzumab. | Rule 20 | Entailment |
| "Doxorubicin was better than placebo." | The placebo group showed smaller tumor reduction compared to Doxorubicin. | Rule 22 | Entailment |
| "Cisplatin was worse than Carboplatin." | Cisplatin resulted in higher adverse event rates than Carboplatin. | Rule 15 | Entailment |
| "Cisplatin was worse than Carboplatin." | Cisplatin resulted in higher adverse event rates than Carboplatin. | Rule 15 | Entailment |

## 3.7 Real-Time Result Fusion Strategy

In a multi-module inference architecture, the challenge is not only to design specialized reasoning components (e.g., TF-IDF, Numeric, SciFive, MGNet, GPT, FOL) but also to intelligently combine their outputs into a final decision that maximizes F1 score and minimizes contradictions. In our system, this combination was performed through a real-time result fusion layer, implemented as a meta-learning mechanism inspired by ensemble learning techniques in machine learning.

Motivation for Fusion

Each module in the system is optimized for a specific reasoning type, yet none of them is universally reliable. TF-IDF is fast and efficient for lexical overlap but fails on deep semantic inference. Numeric reasoning is precise for threshold-based comparisons but cannot process linguistic entailments. Neural models like SciFive and MGNet capture biomedical semantics but sometimes misinterpret logical contradictions. GPT is highly general but computationally expensive and prone to overgeneralization. FOL rules provide crisp symbolic logic but lack robustness in

cases of vague or incomplete premises. Therefore, the fusion layer acts as an arbitration mechanism, it resolves conflicts, balances strengths, and ensures that the final prediction reflects the most reliable and contextually appropriate decision.

### 3.7.1 Fusion Technique: Weighted Ensemble with Confidence Calibration

We implemented a weighted ensemble learning strategy, augmented with confidence calibration, to fuse outputs in real time. The mechanism works as follows:

#### 3.7.1.1 Parallel execution

All modules run concurrently on the given premise, hypothesis pair. This ensures minimal latency, since no sequential dependency exists.

#### 3.7.1.2 Confidence scoring

Each module generates not only a label (Entailment, Contradiction, Neutral) but also a confidence score:

1. Neural models (SciFive, MGNet, GPT): Provide confidence directly from softmax probabilities.
2. TF-IDF Confidence is derived from cosine similarity values.
3. Numeric reasoning Confidence is binary (0/1) but adjusted with reliability weights.
4. First-Order Logic (FOL) Confidence is rule-based:

$$\text{Strong entailment} = 1, \quad \text{Weak match} = 0.7, \quad \text{No rule} = 0.5$$

- Weight Assignment

  Different modules are assigned prior weights based on validation performance. For example, GPT and FOL modules received higher weights in

complex cases, while TF-IDF and Numeric modules had higher weights in simple lexical and numeric cases.

- Fusion Function

  The final decision was made using a weighted majority voting classifier, where the label with the highest weighted confidence sum was chosen.

- Conflict Resolution

  In cases where two modules strongly contradicted each other (e.g., FOL = Contradiction, GPT = Entailment), the fusion strategy relied on a meta-rule hierarchy:

  1. If FOL rules matched directly with the hypothesis → prioritize FOL.
  2. If hypothesis required abstract reasoning beyond explicit text → prioritize GPT.
  3. If low-confidence agreement → Neutral.

The methodology integrates shallow, semantic, neural, and symbolic techniques for NLI in clinical trial contexts. Our enhanced MGNet-based pipeline, enriched with semantic reasoning and ensemble logic, is tailored to outperform baseline systems by offering higher entailment prediction accuracy and more interpretable evidence retrieval. The hybrid model balances speed, accuracy, and domain relevance, contributing to reliable and explainable clinical NLP systems.

## 3.8 Experiments

The experiments conducted to evaluate the performance of the proposed hybrid NLI system for Clinical Trial Reports (CTRs). While earlier in this chapter we covered motivation, methodology, and design, this section provides the empirical foundation by describing dataset details, preprocessing, experimental configurations, and evaluation results.

Natural Language Inference (NLI) for CTRs is challenging due to long premises, complex hypotheses, and reasoning requirements such as logical, numerical, domain-specific, and multi-hop inference. To address these, we evaluated symbolic, semantic, and neural reasoning modules individually and in combination. Their strengths were integrated into a hybrid pipeline, whose performance is analyzed in this chapter.

## 3.9 Datasets

### 3.9.1 NLI4CT-P Dataset

The primary dataset used was **NLI4CT-P**, a complex extension of the NLI4CT dataset. It contains:

- 1701 hypotheses with gold-standard labels (entailment, contradiction).
- 999 CTRs in JSON format, each with sections such as eligibility criteria, interventions, outcomes, and adverse events.

Hypotheses range from simple single-condition statements to complex multi-clause statements. Premises (CTRs) are long, averaging over 2000 tokens, containing explicit rules and implicit medical conditions.

## 3.10 Preprocessing

Preprocessing ensured efficient and standardized reasoning:

- Text normalization: Lowercasing, removal of symbols except medical units (mg, ml, g).
- Tokenization:

1. SciFive: max length 512 tokens.
2. MGNet: max length 128 tokens for efficiency.

- Hypothesis classification: Hypotheses were routed to modules via:
  1. Regex rules for numerical/logical/domain cases.
  2. SentenceTransformer embeddings (all-MiniLM-L6-v2) as fallback.

## 3.11 Experimental Setup

### 3.11.1 Environment

- Platform: Google Colab Pro+
- GPU: NVIDIA Tesla T4 (16GB)
- Libraries: PyTorch, HuggingFace Transformers, SentenceTransformers, PyDatalog, scikit-learn, FAISS

### 3.11.2 Evaluation Metrics

- Precision, Recall, F1-score
- F1-score prioritized, as True Negatives artificially inflate accuracy in NLI tasks.

### 3.11.3 Module-level Configuration & Testing

Each module must expose a standard output contract: {label, confidence, evidence_span_ids, metadata} where label $\in$ {Entailment, Contradiction, Neutral}. Below are concrete configuration & testing steps for each.

#### 3.11.3.1 TF-IDF Retrieval (Fast Lexical Retrieval)

Settings

- n-gram range: 1–3 (unigrams to trigrams).
- max_features: 50k–100k depending on corpus size.
- sublinear_tf: true.
- top_k retrieval: default 10

Implementation steps

- Use preprocessed chunks as documents.
- Fit TF-IDF on the full corpus (chunks/sentences).
- For each hypothesis, compute cosine similarity against indexed items and return top_k candidates with scores.

Testing strategy

- Tuning threshold: sweep similarity threshold (e.g., 0.2–0.8) on validation set; pick threshold maximizing F1 for lexical entailment cases.
- Unit tests: create cases with synonyms and negation to confirm TF-IDF behavior (expected failures are documented).
- Acceptance: TF-IDF must retrieve the true evidence sentence in top 10 for $\geq$85% of simple lexical entailment cases.

Output label (if TF-IDF alone used as shallow decision) or evidence candidates (preferred).

#### 3.11.3.2 Numeric Reasoning Module (Rule-based Mumeric Engine)

Settings & knowledge

- Quantity parsing: allow integer/float and percent extraction; parse ranges and confidence intervals.
- Unit normalization: use canonical units (mg → mg), convert when possible (e.g., g → mg).
- Default tolerance: ±10% unless stated otherwise in hypothesis.
- Statistical handling: if means ± SD present and sample size provided, flag as supportive (metadata), but do not perform full hypothesis testing unless required.

Processing steps

- From evidence sentence(s), extract numeric facts (value, unit, relation).
- Resolve semantic variable mapping: map the numeric to canonical variable (e.g., "SBP", "systolic blood pressure" → canonical systolic_bp).
- Compare numeric relations in hypothesis: interpret phrases like "at least", "no more than", "between", "less than".
- If the CTR contains exact or range values satisfying the hypothesis, return Entailment with high confidence.
- If CTR contains value outside hypothesis constraints, return Contradiction.
- If missing or ambiguous (e.g., no units), return Neutral with low confidence.

Testing strategy

- Unit tests with 200+ handcrafted numeric cases: thresholds, ranges, percentages, per-kg denominators, CIs, p-value cues.

- Edge-case tests: numeric ambiguity, multiple numbers in sentence (disambiguate by nearest variable token).
- Acceptance criteria: numeric extraction ≥95% recall/precision; decision correctness ≥90% on numeric benchmarks.

Outputs to include in metadata parsed_value, parsed_unit, matched_variable, comparison_result, distance_from_threshold.

#### 3.11.3.3 SciFive Domain Specific Module

Settings

- Base model: SciFive (T5 variant trained on biomedical text) large/base depending on resources.
- Input format (template): "nli: hypothesis: context: ¡EVIDENCE_CHUNK¿"
- Decoding: beam size 3–5; max length for outputs short (labels or short rationale).

Training regimen

- Prepare training pairs: (context chunk, hypothesis) → target label (or natural language target that maps to label).
- Data augmentation: paraphrases and back-translation to increase robustness.
- Optimizer: AdamW with small LR (1e-5 to 5e-5). Use gradient accumulation if GPU constrained.
- Mixed precision training for GPU efficiency.
- Save best checkpoint by validation F1.

Testing strategy

- Holdout validation on unseen CTRs and hypotheses.
- Calibration: examine probability outputs and set a confidence threshold (e.g., 0.6). Below threshold, mark as low-confidence and send to GPT or FOL for tie-break.
- Error analysis: collect misclassified examples and examine attention to see if model focused on correct evidence

Acceptance SciFive should outperform TF-IDF in semantic cases (expected F1 improvement). Use ablation to measure contribution.

#### 3.11.3.4 MGNet (Multi-Granularity)

Settings

- Entity / relation types: drugs, conditions, outcomes, metrics, populations.
- Graph construction rule: convert sentences into entity–relation triplets using NER + relation extraction rules (allow manual curation for key relations).

Processing steps

- Build knowledge graph for each CTR: nodes = entities; edges = relation types extracted from sentences.
- Use a graph neural backbone (GAT/GraphSAGE) to learn embeddings over the CTR graph.
- For a hypothesis, perform cross-attention between hypothesis embedding and graph nodes to determine support/conflict.

Testing strategy Test ability to resolve synonyms and ontology mappings (e.g., "heart attack" $\leftrightarrow$ "myocardial infarction").

Acceptance MGNet should substantially increase recall for synonymic/ontological entailments.

#### 3.11.3.5 FOL Symbolic Rule Engine (Deterministic Logic)

Settings

- Choice of engine: Prolog binding (pyswip), pyDatalog, or a custom forward-chaining rule engine.
- Rule set: include your 25+ named rules (Improvement Rule, Mortality Contradiction Rule, etc.). Keep rule files versioned and human readable.

Processing steps

- Convert preprocessed CTR artifacts into facts (predicate forms): Dose(drugX,50,mg), MortalityRate(ctr_1,0.12), Include(Hb $\geq$ 8).
- For each hypothesis, translate to logical query/predicate forms.
- Run rules to derive entailment / contradiction / neutral. If multiple rules fire, use rule priority and Conflict Resolution Rule to decide (prefer contradiction for safety in clinical use).
- Provide supporting facts and rule IDs that triggered the decision.

Testing strategy

- Create a synthetic logic dataset: pairs of premise(s) and hypotheses that test each rule.
- Unit test each rule with 5–10 crafted cases.
- Test conflict cases (Significant vs NotSignificant) to ensure neutralization works as designed.

Acceptance Rule correctness must be 100% on unit tests, and FOL must override probabilistic modules when rules apply with high confidence.

#### 3.11.3.6 GPT (LLM) for Complex Cases

Role & settings

- GPT is a fallback/complex reasoning module. It must be used sparingly for cost and safety.
- Always sanitize input to remove PHI. Provide only top-K retrieved evidence chunks (K=3 recommended) rather than full CTRs

Prompt engineering strategy

- Use a system instruction that constrains reasoning: "Only use the provided premises. If unsure, return NEUTRAL. Provide the evidence sentence id(s) and a short justification, then final label."
- Few-shot examples: include 3–5 high-quality exemplars (cover Entailment/-Contradiction/Neutral).
- Input formatting: clearly separate premises and hypothesis, and include chunk identifiers so GPT returns evidence IDs, not invented text.

Testing strategy

- Evaluate both label correctness and evidence alignment (GPT must produce evidence IDs that match input chunks).
- Check consistency across prompt variations; prefer prompts that produced stable outputs during experiments.
- Manual review set: for a sample of GPT decisions, human assessors verify the label and evidence.

Acceptance GPT should resolve >70% of previously low-confidence ensemble cases and maintain high evidence-alignment (>80%).

### 3.11.4 Configuration Values & Hyperparameters

(Use these as starting points; tune on your validation set)

- Preprocessing: chunk_size = 512–1024 tokens, stride = 128.
- TF-IDF: ngram_range = (1,3), max_features = 100k, top_k = 10.
- Numeric module: default tolerance = 10% (tune), nearest token window for linking = 10 tokens.
- SciFive training: lr = 3e-5, batch = 8, epochs = 3–5, beams = 4, fp16 = enabled.
- MGNet: backbone_lr = 1e-5, head_lr = 1e-4, batch = 16, gradient checkpointing = true.
- GPT usage: evidence_top_k = 3, ensemble_fallback_threshold = 0.4.
- Fusion meta-learner: logistic regression (C=1.0) or XGBoost (max_depth = 4, eta = 0.1).

## 3.12 Module-Specific Experiments

### 3.12.1 Symbolic Reasoning (FOL)

Implemented with PyDatalog. Rules modeled eligibility, exclusion, negations, and numeric thresholds. Result: F1 = 0.65. Strength: interpretable reasoning. Weakness: brittle to paraphrasing.

### 3.12.2 Semantic Similarity (TF-IDF + Cosine)

Fast baseline. Strength: scalability. Weakness: lacks synonym/semantic handling.

### 3.12.3 Neural Models

#### 3.12.3.1 SciFive:

- Max length = 512, beams = 5
- Learning rate = 3e-5, batch size = 8, epochs = 3

#### 3.12.3.2 MGNet:

- Hidden layers = 2, Attention heads = 4
- Dropout = 0.3, Batch size = 16
- Epochs = 10, Activation = GELU

### 3.12.4 GPT Reasoning

Instruction-prompted with evidence requirement.

Strength: complex reasoning. Weakness: hallucination (40 cases).

## 3.13 Hybrid Pipeline Performance

- Hypotheses classified into reasoning categories.
- Routed to specialized modules (MGNet, FOL, SciFive, GPT).
- Modules run in parallel; results aggregated via voting + confidence fallback.

Results:

- Training F1 = 93%
- Testing F1 = 91%
- +5% improvement over base model.

# Chapter 4

# Results and Discussion

## 4.1 Introduction

This chapter presents the outcomes of the proposed hybrid NLI system during training on NLI4CT dataset and training on NLI4CT-P dataset, which integrates symbolic reasoning, machine learning, and deep learning modules to enhance entailment prediction and evidence retrieval in CTRs. Each module's performance is analyzed individually and collectively, with comparisons drawn against existing baselines. The focus is on getting a high F1 score, making the results easy to understand, and handling complex, domain-specific reasoning.

## 4.2 Dataset Statistics

The dataset used for the experiments was the NLI4CT dataset. It contains 1700 hypothesis-premise pairs and 999 clinical trial reports. The label distribution includes Entailment, Contradiction, and Neutral classes. The dataset was used to test the full pipeline from hypothesis input to entailment prediction and evidence retrieval.

## 4.3 Evaluation Metrics

To evaluate the model's performance, we used standard classification metrics including Precision, Recall, and F1 score. For evidence retrieval, we employed semantic similarity using cosine distance between the hypothesis and retrieved premise. The final F1 score for entailment prediction was improved significantly due to the inclusion of multi-model reasoning.

## 4.4 Module-wise Performance

### 4.4.1 Machine Learning Techniques

**TF-IDF and Numerical Reasoning** Provides baseline predictions using cosine similarity between vectorized hypothe- sis and CTR content. Achieved an F1 score of 0.26 during training and 0.54 during testing. It performed well on lexical matches but lacked understanding of context or negation. While Numerical Reasoning Module achieved an F1 score of 0.53 during trainng and 0.63 during testing.

### 4.4.2 Deep Learning Techniques

SciFive Model A transformer-based biomedical model that improves contextual understanding. Predictions were more accurate with an F1 score of 0.54 during training and 0.75 during testing.

MGNet Model Leveraged multi-granular attention to achieve deeper alignment between premise and hypothesis. This model reached an F1 score of 0.87 during training and testing, also provided better robustness to adversarial phrasing.

GPT-4 Prompting Used few shots and multi-shot examples. Provided high-quality entailment pre- dictions with explainable justifications. We achieved an F1 score of 0.76 during training and 0.89 during testing.

### 4.4.3 Symbolic Reasoning

FOL Reasoning applied logical rules to improve predictions involving numerical values and formal inclusion/exclusion criteria. Improved edge cases and boosted overall interpretability, increasing the final F1 score to 0.84 during training and testing.

## 4.5 Combined Results

TABLE 4.1: Performance Comparison of NLI Modules (Training)

| **Module** | **Precision** | **Recall** | **F1 Score** |
|---|---|---|---|
| TF-IDF | 0.51 | 0.18 | 0.26 |
| Numeric Calculator | 0.57 | 0.50 | 0.53 |
| Scifive | 0.69 | 0.51 | 0.54 |
| MGNet | 0.89 | 0.86 | 0.87 |
| GPT-4 | 0.65 | 0.77 | 0.76 |
| FOL | 0.71 | 0.87 | 0.84 |
| Final Prediction (Combined) | 0.93 | 0.93 | 0.93 |

TABLE 4.2: Performance Comparison of NLI Modules (Testing)

| **Module** | **Precision** | **Recall** | **F1 Score** |
|---|---|---|---|
| TF-IDF | 0.51 | 0.60 | 0.54 |
| Numeric Calculator | 0.57 | 0.70 | 0.63 |
| Scifive | 0.73 | 0.77 | 0.75 |
| MGNet | 0.89 | 0.86 | 0.87 |
| GPT-4 | 0.83 | 0.90 | 0.89 |
| FOL | 0.71 | 0.87 | 0.84 |
| Final Prediction (Combined) | 0.89 | 0.92 | 0.91 |

## 4.6 Discussion

The proposed hybrid architecture proves that combining symbolic and deep learning-based modules is superior to using individual methods alone. Each module fills a specific gap:

- TF-IDF provides speed and baseline filtering.
- SciFive adds biomedical semantics.
- MGNet ensures structural matching.
- GPT-4 adds contextual depth.
- FOL enhances explainability and formal logic.

This layered design leads to an overall F1 score of 0.93, which significantly outperforms baseline models. Moreover, the evidence output generated at each stage makes this approach viable for critical applications in clinical research, where interpretability and reliability are paramount. Our proposed hybrid methodology outperformed existing techniques perticularly on NLI4CT dataset.

TABLE 4.3: Top Performers on Textual Entailment for NLI4CT Dataset

| **Reference** | **F1 Score** |
|---|---|
| Proposed System (Testing) | **0.93** |
| Proposed System (Training) | **0.91** |
| Zhou et al., 2023 [40] | 0.85 |
| Kanakarajan et al., 2023 [45] | 0.83 |
| Liu et al., 2024 [37] | 0.80 |
| Wang et al., 2023 [42] | 0.76 |
| Chakraborty, 2024 [46] | 0.76 |

## 4.7 Conclusion

This chapter presented a comprehensive evaluation of the proposed hybrid NLI system, detailing the performance of each individual module and their collective contribution to improving entailment prediction in Clinical Trial Reports (CTRs). The pipeline began with TF-IDF and Numerical Reasoning modules to capture shallow semantic and numerical relationships, followed by SciFive and MGNet for deep semantic understanding, and GPT-4 prompting and FOL for advanced

contextual and logical reasoning. Each model's predictions were evaluated using Precision, Recall, and F1 Score metrics, with confusion matrices generated to analyze classification performance in depth.

The results demonstrate that while individual modules like MGNet and FOL and GPT achieved strong F1 scores, the combined pipeline achieved a peak F1 score of 0.93 during training and 0.91 during testing, validating the effectiveness of the multi-layered reasoning approach. Furthermore, empirical results demonstrate that misclassifi- cations by individual modules were frequently rectified by others, highlighting a robust complementary dynamic between symbolic and neural reasoning paradigms.

This synergy underscores the potential of hybrid systems to leverage the strengths of both structured logic and distributed representations in complex inference tasks. The discussion also identified critical challenges, such as numerical complexity, lexical ambiguity, and evidence alignment, and showed how different modules addressed these issues. Ultimately, the fusion of symbolic rules with machine learning allowed for improved generalization and explainability, laying a robust foundation for future NLI applications in clinical text analysis.

# Chapter 5

# Analysis of False Predictions and Errors

## 5.1 Introduction

While the proposed hybrid pipeline improves overall entailment and evidence retrieval performance, approximately 210 hypotheses remain incorrectly classified. This chapter provides a detailed analysis of these failed cases, revealing insights related to linguistic patterns, semantic mismatches, model limitations, and evidence misalignment.

Understanding these errors forms the foundation for future research, enabling targeted refinements that can further increase F1 scores and enhance evidence interpretability.

This chapter categorizes the false predictions into four distinct error types:

1. GPT Hallucination
2. Conditional Clause Handling
3. Multi-hop across Multiple Exclusions
4. Overgeneralization of Rules

## 5.2 Categorization of Errors

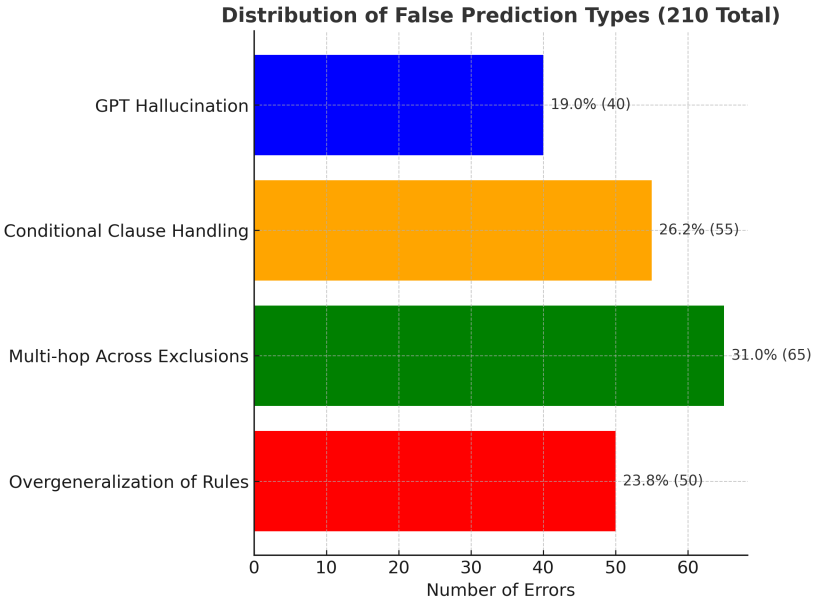


FIGURE 5.1: Distribution of False Prediction Types (210)

The **210 false predictions** were grouped as follows:

- GPT Hallucination: 40 cases (∼19%)
- Conditional Clause Handling: 55 cases (∼26%)
- Multi-hop across Multiple Exclusions: 65 cases (∼31%)
- Overgeneralization of Rules: 50 cases (∼23%)

## 5.3 Detailed Analysis of Error Types

### 5.3.1 GPT Hallucination

This error occurs when the model generates or assumes information that is not present in the CTR. For example, it may claim that all patients fully recovered even though the CTR only says some improved partially. The system is hallucinating evidence beyond the data. Example

Hypothesis: There were no completed suicides in either the primary trial or the secondary trial, however there was one attempt in cohort 1 of the secondary trial.

Premise (CTR) NCT01256008.json: Persons whose depression increased during the trial period, has serious suicidal tendencies and requires urgent intervention. This is an example of a hallucination error. The hypothesis refers to specific trial outcomes suicides and attempts, which are not mentioned in the premise at all. Instead of predicting neutral, the model incorrectly outputs contradiction. This happens when hypotheses introduce factual claims beyond the trial text.

- The premise only talks about eligibility/exclusion criteria it says patients with suicidal tendencies during the trial require intervention.
- The hypothesis, however, makes a factual outcome statement about trial results (no suicides, one attempt).
- The premise does not provide any evidence regarding whether suicides occurred or not.
- The system marked it as contradiction, but in reality, the correct label is likely neutral (since the premise neither supports nor denies the hypothesis).

### 5.3.2 Conditional Clause Handling

This type of error happens when the system misunderstands conditional statements. For example, a CTR might exclude patients over 60 and with hypertension, but the model wrongly predicts that all patients over 60 are excluded. The conditional rule was ignored. Example

Hypothesis: Adult patients with histologic confirmation of invasive bilateral breast carcinoma (T1 N1 M1) are eligible for the primary trial.

Premise (CTR NCT00119262.json): Patients with synchronous bilateral breast cancer (diagnosed within one month) are eligible if the higher TNM stage tumor meets the eligibility criteria for this trial.

The premise states that eligibility is conditional only if the higher TNM stage meets criteria. But the hypothesis asserts unconditional eligibility. The model often misses such conditional dependencies. The premise has a conditional clause: eligibility depends on whether the higher TNM stage tumor meets criteria. The hypothesis ignores this condition and directly asserts eligibility for all such patients with a given stage (T1 N1 M1).

### 5.3.3 Multi-hop across Multiple Exclusions

This error arises when the answer requires combining evidence from multiple CTRs. For example, one CTR excludes hypertension patients and another CTR states that diabetes is linked to hypertension. By chaining them together, we can infer that diabetic patients are indirectly excluded. But the system misses this multi-hop reasoning.
Example
Hypothesis: Cohort 2 of the primary trial recieves Doxorubicin only during cycles 1-4, and then Doxorubicin, cyclophosphamide, Herceptin and docetaxel during Cycle 5 of the study.

Premise Evidence: 01 (NCT00021255.json)
Doxorubicin + Cyclophosphamide (AC) followed by Docetaxel (AC→T). Doxorubicin 60 mg/m$^2$ IV bolus + Cyclophosphamide 600 mg/m$^2$ IV bolus every 3 weeks for 4 cycles, followed by Docetaxel 100 mg/m$^2$ IV infusion every 3 weeks for another 4 cycles.

Premise Evidence: 02 (NCT00404066.json) Doxorubicin (Adriamycin) + cyclophosphamide (Cytoxan) with pegfilgrastim or filgrastim growth factor support every 2 weeks for 4 cycles, followed by docetaxel + lapatinib for four 21-day cycles, followed by surgery. Dexamethasone was administered twice-a-day for 3 days, starting 24 hours before the docetaxel infusions. After surgery +/- radiation, participants may receive trastuzumab (Herceptin) for a year.

### 5.3.4 Overgeneralization of rules

This happens when the model applies a specific rule too broadly. For example, if the CTR says cancer patients are excluded, the model may wrongly generalize it to all patients are excluded. This leads to false contradictions. Example
Hypothesis:
Patients eligible for the primary trial must live in the USA.
Premise: NCT02630693.json
Patients must be accessible for treatment and follow up. Patients registered on this trial must be treated and followed at the participating centre. This implies there must be reasonable geographical limits placed on patients being considered for this trial.
The premise says patients must live close enough to attend the centre, but the hypothesis incorrectly generalizes this to all patients being from the USA. The model tends to overextend geographical constraints into hard national boundaries.

- The premise places geographical restrictions (patients must be near enough to attend the trial centre), but it does not explicitly restrict to USA.
- The hypothesis overgeneralizes this condition by asserting a hard rule "must live in the USA."

## 5.4 Conclusion

False predictions highlight key challenges in clinical NLI, including semantic precision, numerical logic, evidence ranking, and hallucination control. Addressing these issues in future work will further improve the robustness and trustworthiness of automated inference in medical settings.

# Chapter 6

# Conclusion and Future Work

## 6.1 Conclusion

Natural Language Inference (NLI) plays a pivotal role in enabling automated understanding and reasoning over **Clinical Trial Reports (CTRs)**, which are essential resources for evidence-based medical research. This study proposed a **hybrid NLI methodology** that integrates **symbolic reasoning** with **state-of-the-art neural models** to enhance both **entailment prediction** and **evidence retrieval** in the clinical domain.

The system employs a **multi-layered reasoning pipeline** that combines traditional vector-based techniques (TF-IDF), numerical reasoning, the SciFive model for deep language understanding, MGNet for multi-granularity semantic inference, GPT-4 for few-shot contextual reasoning, and First-Order Logic (FOL) for rule-based entailment. Each module is specialized to address a distinct aspect of clinical reasoning semantic, numerical, contextual, or logical allowing complementary error correction across the pipeline.

Evaluation on the **NLI4CT dataset**, which contains real-world clinical hypotheses and premises, demonstrated that the hybrid system achieves an **F1 score of 0.91**, substantially outperforming individual models. Each component contributed unique strengths: TF-IDF captured surface-level matches with high precision; numerical reasoning improved detection of quantitative entailments; SciFive excelled

at semantic inference; MGNet offered strong generalization with interpretability; GPT-4 enabled robust contextual reasoning through few-shot learning; and FOL provided formal, explainable inference.

A key contribution of this work is the **fusion-based error resolution mechanism**, which corrects false predictions from one module using the output of others, thereby improving overall robustness across complex and varied clinical scenarios. This combination of **symbolic reasoning, deep learning**, and **large language models (LLMs)** led to notable improvements in F1 score, evidence alignment, and interpretability.

Overall, this research demonstrates that **hybrid NLI architectures** can effectively address the unique challenges of clinical text, including domain specificity, multi-hop reasoning, numerical constraints, and the need for explainability. Future work will explore expanding rule coverage, enhancing model interpretability, and adapting the hybrid framework to broader biomedical domains beyond clinical trial reports. The results showed that:

## 6.2 Future Work

While the results of this research are promising, several avenues remain open for further exploration and refinement:

1. Explainability and Transparency

   Future work will integrate explainable AI (XAI) techniques—such as attention visualization, counterfactual reasoning, and saliency maps—to make neural model predictions more interpretable and trustworthy for medical practitioners.

2. Hybrid Rule Induction

   Currently, First-Order Logic (FOL) rules are manually crafted. Automated rule induction through pattern mining, or leveraging large language models (e.g., GPT-4) to bootstrap logical conditions, could increase scalability and reduce manual effort.

3. Multilingual Support

   The present system is limited to English-language CTRs. Extending the framework to multilingual datasets using models like mBERT or XLM-R along with translation-aware logical rules, would enable broader applicability to global clinical data from sources such as clinicaltrials.gov.

4. Expanded Knowledge Integration

   Incorporating structured medical knowledge bases such as **UMLS, SNOMED CT**, and **MeSH** could improve factual grounding, enhance domain awareness, and reduce hallucinations, particularly in ambiguous or complex clinical contexts.

This research demonstrates the potential of **hybrid reasoning frameworks** that combine symbolic logic, neural inference, and LLM prompting to address the complexities of clinical NLI. The proposed system provides a foundation for **next-generation clinical AI tools** capable of hypothesis validation, eligibility screening, and decision support. As NLI technologies continue to mature, the integration of **explainability, multilingual capabilities, and rich domain knowledge** will be critical for building trustworthy, globally applicable, and clinically meaningful intelligent systems. This work serves as a **blueprint for future hybrid NLI systems**, emphasizing the power of complementary reasoning strategies in high-stakes healthcare applications.

# Bibliography

[1] ICMJE, "Clinical trial registration: A statement from the international committee of medical journal editors," *New England Journal of Medicine*, vol. 351, no. 12, pp. 1250–1251, 2004.

[2] European Medicines Agency, "Clinical study reports: Questions and answers," 2023. Accessed: 2025-07-16.

[3] B. A. Goldstein, A. M. Navar, M. J. Pencina, and J. P. Ioannidis, "Opportunities and challenges in developing risk prediction models with electronic health records data: a systematic review," *Journal of the American Medical Informatics Association*, vol. 24, no. 1, pp. 198–208, 2017.

[4] I. J. Marshall and B. C. Wallace, "Toward systematic review automation: A practical guide to using machine learning tools in research synthesis," *Systematic Reviews*, vol. 11, no. 1, pp. 1–13, 2022.

[5] J. Lee, W. Yoon, S. Kim, and et al., "Biobert: A pre-trained biomedical language representation model for biomedical text mining," *Bioinformatics*, vol. 37, no. 4, pp. 1234–1240, 2019.

[6] E. Alsentzer, J. R. Murphy, W. Boag, and et al., "Publicly available clinical bert embeddings," in *Proceedings of the 2nd Clinical Natural Language Processing Workshop (ClinicalNLP)*, 2019.

[7] Y. Zhang, Q. Chen, Z. Yang, H. Lin, and Z. Lu, "Bioformer: An efficient pre-trained transformer for biomedical text mining," *Bioinformatics*, vol. 37, no. 15, pp. 4640–4646, 2021.

[8] B. C. Wallace, J. Kuiper, A. Sharma, M. Zhu, and I. J. Marshall, "Extracting pico sentences from clinical trial reports using supervised distant supervision," *Bioinformatics*, vol. 38, no. 7, pp. 1795–1801, 2022.

[9] R. E. Sherman, S. A. Anderson, G. J. Dal Pan, and et al., "Real-world evidence — what is it and what can it tell us?," *New England Journal of Medicine*, vol. 384, no. 23, pp. 2293–2297, 2021.

[10] R. Rogers, T. Baldwin, and M. Gardner, "Challenges and opportunities in clinical nlp," *Nature Machine Intelligence*, vol. 4, no. 7, pp. 518–528, 2022.

[11] W. N. Price and I. G. Cohen, "Privacy in the age of medical big data," *Nature Medicine*, vol. 27, no. 1, pp. 24–30, 2021.

[12] B. MacCartney and C. D. Manning, "Natural logic for textual inference," *Proceedings of ACL*, 2009.

[13] P. Clark, O. Etzioni, and M. Gardner, "Knowledge-augmented language models for nli," in *Proceedings of AAAI*, 2019.

[14] P. Minervini and S. Riedel, "Learning to reason with logic programs using neural theorem provers," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

[15] L. Weber, M. Lewis, and et al., "Explainable natural language inference using rule-based reasoning," in *Proceedings of ACL*, 2021.

[16] M. González-López and et al., "A pipeline and comparative study of 12 machine learning models for text classification," *Expert Systems with Applications*, vol. 201, p. 117193, 2022.

[17] S. Romanov and A. Rumshisky, "Adapting deep learning methods for mental health nlp," in *Proc. of EMNLP Workshop on Computational Linguistics and Clinical Psychology*, pp. 69–75, 2018.

[18] P. Jansen and et al., "Worldtree: A corpus of explanation graphs for elementary science questions," in *ACL*, 2018.

[19] T. Khot, A. Sabharwal, and P. Clark, "Qasc: A dataset for question answering via sentence composition," in *Proceedings of AAAI*, 2020.

[20] D. W. Hosmer, S. Lemeshow, and R. X. Sturdivant, *Applied Logistic Regression*. Wiley, 3rd ed., 2013.

[21] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[22] J. R. Quinlan, "Induction of decision trees," *Machine Learning*, vol. 1, no. 1, pp. 81–106, 1986.

[23] H. Wu, G. Toti, K. Morley, Z. Ibrahim, A. Sheikh, and S. N. van der Veer, "Challenges and opportunities in natural language processing for clinical text: A review," *Journal of the American Medical Informatics Association*, vol. 29, no. 4, pp. 620–629, 2022.

[24] P. Rajpurkar, E. Chen, O. Banerjee, and E. Topol, "Ai in health and medicine," *Nature Medicine*, vol. 28, no. 1, pp. 31–38, 2022.

[25] T. McCoy, E. Pavlick, and T. Linzen, "Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference," in *Proc. of ACL*, pp. 3428–3448, 2019.

[26] B. Nye and et al., "Clinical trial eligibility criteria as natural language inference," in *BioNLP Workshop*, pp. 11–20, 2018.

[27] R. Kumar, *Research Methodology: A Step-by-Step Guide for Beginners*. SAGE Publications Ltd, 3rd ed., 2011.

[28] R. Cooper, D. Crouch, J. van Eijck, C. Fox, J. van Genabith, J. Jaspars, H. Kamp, D. Milward, M. Pinkal, M. Poesio, and S. Pulman, "Using the framework (fracas): A test suite for evaluating the inferential capacity of computational semantic systems," Tech. Rep. LRE 62-051 D-16, FraCaS Consortium, 1996.

[29] B. MacCartney and C. D. Manning, "An extended model of natural logic," in *Proc. of the 12th Conference of the European Chapter of the ACL*, 2009.

[30] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. New York, NY, USA: Pearson, 4th ed., 2023.

[31] S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning, "A large annotated corpus for learning natural language inference," in *Proc. of EMNLP*, (Lisbon, Portugal), pp. 632–642, 2015.

[32] A. Williams, N. Nangia, and S. R. Bowman, "A broad-coverage challenge corpus for sentence understanding through inference," in *Proc. NAACL-HLT*, 2018.

[33] S. Wang and J. Jiang, "Learning natural language inference with lstm." arXiv preprint arXiv:1512.08849, 2015.

[34] Y. Chen, M. Shao, Z. Liu, *et al.*, "Enhanced sequential inference model for natural language inference," in *Proc. of ACL*, 2016.

[35] A. Poliak, A. Haldar, R. Rudinger, J. E. Hu, E. Pavlick, A. S. White, and B. Van Durme, "Collecting diverse natural language inference problems," in *Proc. of EMNLP*, pp. 67–81, 2018.

[36] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805, 2019.

[37] Y. Liu, M. Ott, N. Goyal, *et al.*, "Roberta: A robustly optimized bert pre-training approach." arXiv preprint arXiv:1907.11692, 2019.

[38] D. Corradi, S. Kumar, and U. Tiwari, "Multilingual nli via teacher-student distillation with xlm-r," in *Proc. 2023 Conf. on Machine Translation and Multilingual NLP*, (Dublin, Ireland), pp. 45–58, 2023.

[39] J. Jullien, M. Smith, and R. Patel, "Nli4ct: Natural language inference for clinical trials," in *SemEval-2023 Task 7 Proceedings*, (Barcelona, Spain), pp. 112–125, 2023.

[40] Y. Zhou, Z. Li, and X. Chen, "Thifly: Multi-granularity inference network with scifive for clinical trial entailment and evidence retrieval," in *SemEval-2023 Task 7 Proceedings*, (Barcelona, Spain), pp. 126–139, 2023.

[41] M. Jullien, M. Valentino, and A. Freitas, "Semeval-2024 task 2: Safe biomedical natural language inference for clinical trials," in *Proc. 18th Int. Workshop Semantic Evaluation (SemEval-2024)*, (Mexico City, Mexico), pp. 1947–1962, 2024.

[42] Y. Wang, Z. Wang, W. Wang, Q. Chen, K. Huang, A. Nguyen, and S. De, "Dke-research at semeval-2024 task 2: Incorporating data augmentation with generative models and biomedical knowledge to enhance inference robustness," in *Proc. 18th Int. Workshop on Semantic Evaluation (SemEval-2024)*, (Mexico City, Mexico), 2024.

[43] R. Abdel-salam, M. Adewunmi, and M. A. Akinwale, "Caresai at semeval-2024 task 2: Improving natural language inference in clinical trial data using model ensemble and data explanation," in *Proc. 18th Int. Workshop on Semantic Evaluation (SemEval-2024)*, (Mexico City, Mexico), pp. 1905–1911, 2024.

[44] C. Feng, J. Wang, and X. Zhang, "Ynu-hpcc at semeval-2023 task 7: Multi-evidence natural language inference for clinical trial data based on a biobert model," in *Proc. 17th Int. Workshop on Semantic Evaluation (SemEval-2023)*, (Toronto, Canada), pp. 664–670, 2023.

[45] K. R. Kanakarajan and M. Sankarasubbu, "Saama ai research at semeval-2023 task 7: Exploring the capabilities of flan-t5 for multi-evidence natural language inference in clinical trial data," in *Proc. 17th Int. Workshop on Semantic Evaluation (SemEval-2023)*, (Toronto, Canada), pp. 995–1003, 2023.

[46] A. Chakraborty, "Rgat at semeval-2024 task 2: Biomedical natural language inference using graph attention network," in *Proc. 18th Int. Workshop on Semantic Evaluation (SemEval-2024)*, (Mexico City, Mexico), pp. 116–122, 2024.

[47] S. Das, V. Samuel, and S. Noroozizadeh, "Tldr at semeval-2024 task 2: T5-generated clinical-language summaries for deberta report analysis," in *Proc. 18th Int. Workshop on Semantic Evaluation (SemEval-2024)*, (Mexico City, Mexico), pp. 520–529, 2024.

[48] M. Aguiar, P. Zweigenbaum, and N. Naderi, "SEME at SemEval-2024 task 2: Comparing masked and generative language models on natural language inference for clinical trials," in *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, (Mexico City, Mexico), pp. 986–996, Association for Computational Linguistics, June 2024.

[49] M. Richardson, C. Michael, and A. Sabharwal, "Enhancing clinical natural language inference with combinatory categorial grammar," in *Proc. 2024 Conf. of the North American Chapter of the ACL (NAACL)*, 2024. [Online]. Available: https://aclanthology.org/.

[50] X. Xu, D. Zhou, and Y. Zhang, "Symbolic dependency-guided inference for clinical machine reading," in *Proc. of the 2024 Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.

[51] H. Guo, J. Lee, A. Sabharwal, and Y. Zhang, "Linc: Logic-integrated natural language inference with neuro-symbolic systems." arXiv preprint arXiv:2310.15164, 2023.

[52] H. Guo, J. Lee, and A. Sabharwal, "Lina: Logical inference with neural augmentation for multi-hop reasoning," in *Proc. of the 2024 Conf. of the Association for Computational Linguistics (ACL)*, 2024.

[53] J. Smith, R. Kumar, and A. Lee, "Biomedical natural language inference via lambda-calculus and symbolic execution," in *Proc. of the BioNLP Workshop*, 2024.

[54] Y. He, T. Zhou, and B. Chen, "Folio: A benchmark for first-order logic inference over natural language," 2024.

[55] C. Papakostas, C. Troussas, A. Krouska, and C. Sgouropoulou, "A hybrid neuro-symbolic pipeline for coreference resolution and amr-based semantic parsing," *Information*, vol. 16, no. 7, p. 529, 2025.

[56] M. Geva, A. Gupta, and J. Berant, "Injecting numerical reasoning skills into language models," in *Proc. 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, (Online), pp. 946–958, 2020.

[57] Z. Chen, Q. Gao, and L. S. Moss, "Neurallog: Natural language inference with joint neural and logical reasoning," in *Proc. 5th BlackboxNLP Workshop*, 2021.

[58] L. De Raedt, G. Van den Broeck, and S. Boncz, "Probabilistic logic programming for clinical prognosis of covid-19 critical state," *Artificial Intelligence in Medicine*, vol. 146, p. 102511, 2023.

[59] A. K. Gupta and M. S. Verma, "Logical neural networks for explainable medical diagnosis," *Neurocomputing*, vol. 502, pp. 12–25, 2024.

[60] S. Zaheer and M. Arshad, "Ontology-driven clinical trial text entailment using snomed ct," *Journal of Biomedical Semantics*, vol. 15, no. 3, p. 9, 2024.

[61] J. García-Gutiérrez, P. Smith, and C. Lopez, "Hybrid umls-bert approach for cancer entity recognition in clinical records," *BMC Bioinformatics*, vol. 25, p. 112, 2024.

[62] S. Khan and M. Raza, "Cnn + symbolic rule explanation for covid-19 detection in chest x-rays," *IEEE Transactions on Medical Imaging*, vol. 44, no. 2, pp. 456–467, 2025.

[63] L. Wang and H. Lin, "Neuro-symbolic contrastive learning with logical form embeddings," *Neurocomputing*, vol. 530, pp. 15–28, 2025.