

Beyond Negation Detection: Comprehensive Assertion Detection Models for Clinical NLP

Veysel Kocaman¹, Yigit Gul¹, M. Aytug Kaya¹, Hasham Ul Haq¹, Mehmet Butgul¹, Cabir Celik¹ and David Talby¹

¹John Snow Labs inc. 16192 Coastal Highway, Lewes, DE 19958, USA

Abstract

Assertion status detection is a critical yet often overlooked component of clinical NLP, essential for accurately attributing extracted medical facts. Past studies narrowly focused on negation detection, resulting in underperforming commercial solutions such as AWS Medical Comprehend, Azure AI Text Analytics, and GPT-4o due to their limited domain adaptation. To address this gap, we developed state-of-the-art assertion detection models, including fine-tuned LLMs, transformer-based classifiers, few-shot classifiers, and deep learning (DL) approaches and evaluated our models against cloud-based commercial API solutions and legacy rule-based NegEx approach as well as GPT-4o. Our fine-tuned LLM achieves the highest overall accuracy (0.962), outperforming GPT-4o (0.901) and commercial APIs by a notable margin, particularly excelling in *Present* (+4.2%), *Absent* (+8.4%), and *Hypothetical* (+23.4%) assertions. Our DL-based models surpass commercial solutions in *Conditional* (+5.3%) and *Associated with Someone Else* (+10.1%), while few-shot classifier offers a lightweight yet highly competitive alternative (0.929), making it ideal for resource-constrained environments. Integrated within Spark NLP, our models consistently outperform black-box commercial solutions while enabling scalable inference and seamless integration with medical NER, Relation Extraction, and Terminology Resolution. These results reinforce the importance of domain-adapted, transparent, and customizable clinical NLP solutions over general-purpose LLMs and proprietary APIs.

1. Introduction

The widespread adoption of Electronic Health Records (EHRs) has transformed healthcare, with 96% of non-federal acute care hospitals and 78% of office-based physicians in the United States using certified EHR systems by 2021. This digitization has created vast patient data repositories, opening new avenues for clinical applications and research. [1]. To harness this valuable information and discover patterns in EHRs, various Natural Language Processing (NLP) tasks have been performed. Among these, the classification of assertions stands out as a critical but understudied task. Accurate assertion classification allows for the determination of whether a medical concept is *present*, *absent*, *possible*, *hypothetical*, *conditional*, or *associated with someone* other than the patient, crucial for extracting actionable insights from EHRs, driving clinical decision-making, and facilitating healthcare analytics [2]. In other words, the status of an assertion explains how a named entity (e.g. clinical finding, procedure, lab result) pertains to the patient by assigning a label such as *present* ("patient is diabetic"), *absent* ("patient denies nausea"), *conditional* ("dyspnea while climbing stairs"), or *associated with someone else* ("family history of depression"). Table 1 illustrates different assertion classes with their label distribution and sizes.

Although early studies often equated assertion detection with negation detection in time, sophisticated machine learning and deep learning methodologies evolved rudimentary rule-based approaches. Early techniques such as NegEx [3], ConText [4], NegFinder [5] and NegExpander [6] relied on hand-crafted rules and regular expressions, achieving high precision but suffering

In: R. Campos, A. Jorge, A. Jatowt, S. Bhatia, M. Litvak (eds.): *Proceedings of the Text2Story'25 Workshop, Lucca (Italy), 10-April-2025*



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

from low recall due to rigid patterns [7]. In order to learn more about rule-based approaches towards assertion detection, a reader is advised to check out a comprehensive study [8] evaluating these approaches in detail. Deep learning methods, particularly transformer-based models and attention mechanisms, emerged as powerful alternatives, offering more nuanced understanding of clinical text. However, these approaches consistently faced challenges such as requiring large annotated datasets and struggling with minority classes, especially in detecting *possible* medical assertions. Recent developments have focused on addressing these limitations through innovative approaches like multi-task learning, pre-training techniques, and Large Language Models (LLMs).

Bhatia et al. demonstrated the effectiveness of a multitask learning approach for jointly modeling named entity recognition and negation assertion in clinical texts. By utilizing shared parameters, their model achieved improved contextual representation and overcame challenges associated with neural networks in negation detection, outperforming rule-based systems in conjunction with the proposed conditional softmax decoder [9].

Chen et al. explored applying attention-based bi-LSTM architectures for negation and assertion detection in clinical notes, leveraging the ability to selectively focus on relevant information and automatically capture semantic details without relying on external knowledge inputs [10].

Aken et al. proposed a comprehensive study of clinical assertion detection models by manually annotating 5,000 assertions in the MIMIC-III dataset, evaluating medical language models' performance and transferability across different medical domains, and releasing their annotated dataset to address label sparsity and diversity challenges in existing research [11]. Similarly, Wang et al. proposed a novel prompt-based learning approach for assertion classification that addresses existing limitations by leveraging few-shot learning and advanced reasoning techniques [12].

Yuan et al. proposed a deep learning approach for automatic Electronic Medical Record (EMR) sectioning using MIMIC-III data, developing hand-crafted rules to create gold-standard labels and generating multiple note versions with varied section heading formats to train models that achieve robust adaptability and high accuracy in EMR segmentation [13].

Ji et al. proposed a novel method leveraging Large Language Models (LLMs) with advanced reasoning techniques like Tree of Thought (ToT), Chain of Thought (CoT), and Self-Consistency (SC), combined with Low-Rank Adaptation (LoRA) fine-tuning, to transform assertion detection into a generative task that enables more nuanced, contextually aware, and data-efficient medical text understanding across multiple assertion categories [14].

While existing clinical NLP approaches have made significant strides in assertion detection, predominantly focusing on negation, they have consistently fallen short of providing a comprehensive, multi-category framework capable of robustly addressing the full spectrum of medical concept assertions. In this paper, we present a comprehensive implementation of assertion detection within Healthcare NLP library [15] (based on Spark NLP [16] [17] [18] ecosystem), utilizing state-of-the-art models and annotators to achieve high accuracy and efficiency in clinical NLP tasks. Our approach transcends traditional negation detection methods by offering a comprehensive, fully integrable end-to-end solution that addresses the entire spectrum of assertion types, including *present*, *absent*, *possible*, *hypothetical*, *conditional*, and *assertions associated with someone other than the patient*. This holistic method leverages advanced deep learning architectures, few-shot learning techniques, and flexible rule-based systems to overcome common challenges in clinical texts, such as class imbalance and ambiguous concept expressions. Specifically, we explore the following architectures/ modules that we developed during this study to detect assertion status from clinical notes:

- **Assertion Detection with LLMs:** To overcome the limitations coming from data collec-

Table 1

Examples from the i2b2 dataset illustrating different assertion classes with their label distribution.

Text	Label	Description	Size
Overnight, the patient became hypoxic , dropping to the 80 's.	<i>present</i>	Confirms the presence of a medical condition.	8622
He gets short of breath with one flight of stairs.	<i>conditional</i>	Represents conditions that might occur under specific circumstances or conditions.	148
Small stroke, nearly recovered, likely embolic from carotid artery .	<i>possible</i>	Suggests uncertainty or potential presence of a condition.	652
There was no evidence of diarrhea during medical Lawrence Memorial Hospital stay.	<i>absent</i>	Indicates the negation or nonexistence of a medical condition.	2594
Mother suffer MI in her 50 's, died at age 59.	<i>associated with someone else (awse)</i>	Refers to medical conditions related to individuals other than the patient, such as family members.	131
Hydrocodone 5 mg with Tylenol , one to two tablets every four hours p.r.n. pain .	<i>hypothetical</i>	Denotes speculative or conjectural conditions that are not currently present.	445

tion and annotations to design ML/ DL based assertion detection models, we experiment with leveraging LLMs pretrained on extensive medical datasets to enhance assertion detection accuracy and comprehensiveness in zero shot settings.

- **Assertion Detection with a DL Model:** A deep learning-based annotator built on a Bi-LSTM architecture, inspired by [19]. This model processes medical concepts and their surrounding tokens using word embeddings within a defined scope window.
- **Assertion Detection with a Bert For Sequence Classification (BFSC):** This approach leverages a transformer-based model, BERT, to classify assertion status in medical texts. By encoding the contextual relationships within sequences, BERT enables accurate detection of negations, affirmations, and other assertion types.
- **Few Shot Assertion Detection with Transformers:** A few-shot learning-based classifier that combines sentence embeddings with lightweight classification models to achieve high accuracy with minimal training data.
- **Rule-based Assertion Detection with Contextual Awareness:** A rule-based annotator designed to enhance assertion detection accuracy in complex clinical contexts. By leveraging customizable keyword sets, regex patterns, and scope windows, this model adapts to diverse clinical scenarios.

The subsequent sections will systematically evaluate these assertion detection architectures on a well-known benchmark dataset and compare them with GPT4o, a rule-based algorithm (NegEx), and cloud-based healthcare-specific APIs offered by commercial providers (AWS Medical Comprehend and Azure AI Text Analytics for Health). Our analysis will showcase a novel combined pipeline that integrates these models, demonstrating their complementary strengths in enhancing assertion detection performance across various computational paradigms and clinical scenarios.

2. Methodology

In this section, we explain the details of various model architectures supported and shipped as a pretrained model into Healthcare NLP library by John Snow Labs (JSL).

2.1. Assertion Detection with LLMs

Traditional approaches to assertion detection in medical text, such as rule-based NLP systems and machine learning or deep learning models, often require significant manual effort to design patterns and frequently fail to capture less common assertion types, resulting in incomplete contextual understanding. To overcome these limitations, we explored finetuning an LLM with assertion detection datasets to enhance assertion detection accuracy and comprehensiveness.

We explored training LLama-3.1-8B[20] on the i2b2 assertion dataset. We fine tuned LLama-3.1-8B[20] model with the i2b2 assertion training dataset using LoRA fine-tuning [21] approach without quantization. LoRA offers parameter efficiency by updating only a small subset of parameters, reducing memory and computational overhead. It minimizes overfitting risk by keeping pre-trained weights fixed, which makes it ideal for small training datasets and it preserves pre-trained knowledge, maintaining generalization capabilities while allowing task-specific tuning. Our final configuration used a LoRA rank of 16, LoRA alpha of 32, and 5 training epochs.

For fine-tuning, a simple and efficient prompt structure is of paramount importance. We explicitly included a detailed description of each assertion status to evidently improve performance. Additionally, we replaced the term *Present* with *Confirmed* which yielded better results, likely due to improved clarity and alignment with the task’s semantics. Including descriptions of assertion statuses in the input prompt also allowed for minor adjustments during inference, enhancing flexibility and adaptability. Our experimentation proved counterintuitive when we studied the context: inputting the whole document created complexity and confusion, impairing performance. We replaced this approach with a context windowing strategy, extracting two sentences before and after the target text. This strategy substantially reduced training time and increased the model’s ability to focus on relevant information.

2.2. Assertion Detection via DL Model

Assertion Detection via DL Model (AssertionDL) is a classification model based on a Bi-LSTM framework, representing a modified version of the architecture proposed by [19]. In this implementation, entities (also referred to as chunks) are processed alongside a context string. The context string and entities are tokenized and embedded before being passed to the Bi-LSTM model. It is important to balance the length of the context string, as excessively long sequences can result in vanishing gradients, which may hinder the model’s performance.

An analysis of the i2b2 dataset revealed that 95% of the relevant scope tokens (neighboring words) are located within a window spanning 9 tokens to the left and 15 tokens to the right of the target tokens. Based on this observation, we adopted the same window size for our model.

The model has been implemented in Healthcare NLP library as an annotator called *AssertionDLModel*, enabling seamless integration into the Spark NLP library for clinical and biomedical text processing.

2.3. Assertion Detection via Bert For Sequence Classification (BFSC)

While the LLM approach generates new tokens as part of its output, we also explored a more direct approach by framing the problem as a classification task. In this setup, the input consists of the entity chunk and its surrounding context, while the output is the predicted assertion status class. Specifically, we implemented a classification layer on top of a transformer model, such as BERT[22], to perform assertion status prediction, a technique known as BERT for Sequence Classification.

Rather than using the standard BERT model, we utilized the pre-trained BERT models from [23], which have been fine-tuned on biomedical text. Among these, we selected the model trained on BioBert [24], as it demonstrated the best performance for our task. This approach has previously shown promising results for assertion detection [11] and helps the model focus on the target entity, even in contexts that contain multiple entities.

The input text was prepared by a novel approach as explained in [11]. In addition, we experimented with varying context lengths by incorporating additional sentences around the target chunk. However, this approach yielded minimal performance improvements while evidently increasing training and processing time.

2.4. Few Shot Assertion Detection via Transformers

Few Shot Assertion Detection via Transformers (FewShotAssertion) in this study is built on a modified version of SetFit (Sentence Transformer Fine-Tuning) framework [25], which leverages sentence-transformer embeddings and a lightweight classifier for few-shot learning. SetFit enables efficient fine-tuning by coupling a pre-trained sentence-transformer model with a classifier trained on task-specific data using contrastive learning.

The model takes as input the assertion context and the target entity, embedding them using a pre-trained transformer encoder. These embeddings are then fine-tuned using contrastive learning to align positive examples while separating negative ones in the embedding space. A lightweight linear classifier is subsequently trained on the refined embeddings to predict the assertion status. This approach is particularly well-suited for assertion detection in the i2b2 dataset, as it effectively handles limited labeled data while maintaining robust performance.

2.5. Rule-based Assertion Detection with Contextual Awareness (ContextualAssertion)

The model based on this architecture enables assertion detection by labeling entities (chunks) based on user-defined rules and contextual patterns, building upon principles similar to ConText [26] and the widely used NegEx framework [27]. Unlike NegEx, which focuses on negation detection using fixed lexical patterns, the Contextual Assertion module provides advanced configurability through prefix and suffix keywords, regex patterns, exception handling, and customizable scope windows. These enhancements enable the establishment of complex linguistic rules, allowing the annotator to function as a robust and flexible guardrail for NLP pipelines. Following are some of its features:

2.6. Using Assertion Detection Models within Healthcare NLP Pipeline

While the i2b2 dataset provides pre-annotated named entities (including their indices), practical applications require extracting these entities directly from unstructured text. To address this, we propose an end-to-end, flexible pipeline with component sharing, as illustrated in Figure A1.

In this pipeline, named entities are identified using Healthcare NLP’s NER models and subsequently passed to assertion models for assertion status detection. The assertion model utilizes the same embeddings as the NER model, enabling embedding sharing for improved memory management and reduced latency.

The pipeline also supports a stacking approach, allowing multiple assertion models to coexist within a single framework. To enhance performance, we developed a merging mechanism that combines predictions from three assertion models and prioritizes them to produce a unified label for each entity based on the performance of each assertion models on certain

entities. The key components of this pipeline includes **AssertionDL**, **FewShotAssertion** and **ContextualAssertion**.

To resolve conflicts in predictions across models, a **majority voting mechanism** is applied. This approach ensures the final label reflects the consensus among models, mitigating the impact of outlier predictions.

2.7. Pretrained Models Offered in Healthcare NLP

The Healthcare NLP library by JSL offers a range of domain-specific pretrained assertion models (e.g., oncology, radiology) that have been fine-tuned or trained using the architectures explored in this study. These models are fully optimized for integration within a Healthcare NLP pipeline, enabling scalable and efficient deployment. For a detailed list of pretrained clinical assertion models and their corresponding benchmarks, refer to the Table A4 that showcases the best performance scores achieved by these models across multiple assertion categories (12 categories, more than what is covered in this study) including *Present*, *Past*, *Possible*, *Absent*, *Hypothetical*, *Family*, *Someone Else*, *Planned*, *Conditional*, *Confirmed*, *Negative*, and *Suspected*.

3. Experiments and Results

3.1. Experimental Setup

In this study, we benchmarked the performance of our assertion classification approaches — AssertionDL, FewShotAssertion, ContextualAssertion, BFSC and a combined pipeline against available counterparts – NegEx, AWS Comprehend Medical, Azure AI Text Analytics, and GPT-4o.

NegEx is a rule-based algorithm designed to identify negation in clinical text, particularly to determine whether a medical concept is *absent* or not. Introduced by [3], NegEx uses regular expressions and predefined linguistic patterns to detect negation cues (e.g., “no,” “denies”) and their scope within a sentence.

GPT-4o was employed to benchmark assertion detection for medical conditions in the i2b2 dataset. As the disclosure statement of i2b2 dataset prohibits sharing the data via cloud based APIs, we obfuscated the i2b2 dataset both for PHI and medical terms using Healthcare NLP tools provided by John Snow Labs, and then run the evaluation. A carefully crafted prompt (see figure A2) guided the model to classify assertion statuses for specified medical entities.

AWS Comprehend Medical is an NLP service offered by Amazon Web Services, designed to automate the extraction of medical information from unstructured text. **Azure AI Text Analytics** is a natural language processing (NLP) service provided by Microsoft, designed to analyze and extract insights from unstructured text.

Both AWS and Azure services extract entities at first and then annotate them with assertion labels (e.g., *present*, *absent*, *hypothetical*). We aligned these annotations with i2b2 dataset taxonomies via label mapping to ensure consistency in evaluation. Since these services assign assertion labels only to the entities extracted by them at first, the evaluation is run over the partially or fully overlapped common entities from i2b2 dataset. The overlapping rates can be seen at Table A1 in Appendix.

To maintain consistency, labels from Azure AI and AWS Comprehend were mapped to i2b2 equivalents. Matches were categorized into Full Match, Partial Match, and No Match, focusing the evaluation on full and partial matches. Statistics for matching outcomes are summarized in Table A1 with label mapping details available in Table A2 in Appendix for a reference.

3.2. Dataset Description

The evaluation and benchmarking in this study are conducted exclusively on the **official 2010 i2b2 dataset (test split)**[28], which represents a comprehensive resource for assessing assertion detection frameworks in real-world clinical scenarios. The results focus on both individual models and combined pipelines, showcasing their relative strengths and collective impact on performance.

The dataset utilized in this study **covers all six assertion categories**: *Absent*, *Associated with someone else*, *Conditional*, *Hypothetical*, *Possible*, and *Present*. However, the fine-tuned LLM excludes the *Conditional* label due to its ambiguity with the *Hypothetical* label, which could complicate fine-tuning. This exclusion simplifies training and sharpens the model’s focus on the remaining categories. In contrast, other models, including LLMs, retain all six categories to ensure a thorough evaluation of performance across the full range of assertion types.

3.3. Comparative Results

Table 2 presents the experimental results, highlighting the performance of each model across relevant categories.

Table 2

Comparison of assertion models across various categories. Best performing model for each category is represented with bold characters. The models in the first section of this table are developed by JSL. In LLM and GPT-4o experiments, *hypothetical* and *conditional* labels are merged/treated as a single label.

Model	present	absent	possible	hypothetical	conditional	awse*	weighted avg
<i>Combined Pipeline**</i>	0.963	0.951	0.755	0.875	0.511	0.922	0.941
<i>AssertionDL</i>	0.941	0.898	0.672	0.761	0.599	0.886	0.907
<i>FewShotAssertion</i>	0.955	0.942	0.748	0.872	0.293	0.809	0.929
<i>ContextualAssertion</i>	-	0.929	0.708	-	-	0.835	0.883
Fine Tuned LLM	0.976	0.975	0.759	0.911	-	0.943	0.962
BFSC (BioBert)	0.975	0.972	0.787	0.918	0.590	0.913	0.957
GPT-4o	0.937	0.891	0.692	0.677	-	0.805	0.901
Azure Ai Text Analytics	-	0.761	0.583	0.763	0.569	0.800	0.727
AWS Med Comprehend	0.882	0.788	0.659	0.617	-	0.737	0.839
NegEx	-	0.897	-	-	-	-	0.897
BFSC latest best [11]	0.979	0.972	0.786	-	-	-	0.952
Prompt-based Bert [29]	0.971	0.968	0.763	0.921	0.485	0.875	0.951

*awse: associated with someone else. **Combined pipeline elements denoted in italics. ***BFSC latest best is benchmarked only on 3 labels by its authors [11]; hence excluded from comparison.

Our fine-tuned LLM, based on the LLaMA 3.1-8B model and trained using LoRA on the i2b2 dataset, demonstrates superior performance in most categories compared to other models. This approach aligns with recent research in domain adaptation for clinical NLP tasks [30]. The results emphasize the efficacy of smaller, domain-specific models, which, when coupled with carefully engineered prompts, can often outperform much larger, general-purpose models. Our experimental findings indicate near-perfect performance across most categories, with only minor underperformance in the *possible* and *hypothetical* labels. Notably, our model excels not only in covering a broader range of categories but also in evidently outperforming commercial solutions such as GPT-4o, Azure AI Text Analytics, and AWS Comprehend.

The combined pipeline, which integrates rule-based methods with machine learning techniques, closely mirrors the performance of the fine-tuned LLM across most categories. This hybrid approach, which captures the strengths of both deep learning and rule-based systems, outperforms comparable solutions offered by Azure and AWS in every category except the *conditional* label. Unlike Azure AI Text Analytics for Health and AWS Medical Comprehend, which are API-based black-box solutions, our pipeline offers customization and fine-tuning

options, allowing for potential performance improvements across all categories, including the *conditional* label. This flexibility represents a significant advantage in adapting the system to meet specific healthcare needs and optimizing performance across various clinical NLP tasks.

For use cases where deploying the full combined pipeline is not feasible, users can still achieve exceptional results by leveraging its individual components. AssertionDL, in particular, stands out as a versatile solution, effectively handling all assertion categories with its advanced deep learning architecture. It performs particularly well in the *conditional* and *associated with someone else* categories, demonstrating superior results in the *conditional* label. Notably, AssertionDL outperforms GPT-4o in most categories, making it a robust standalone option for clinical assertion tasks.

The FewShotAssertion model can be used both standalone and as part of the pipeline, offering an ideal solution for rapid training and inference in resource-constrained clinical NLP environments where efficiency is crucial. It performs comparably to the fine-tuned LLM across most categories, with the exception of the “conditional” category. However, when integrated into the Healthcare NLP pipeline, its contribution of *absent* and *hypothetical* labels helps mitigate this limitation.

The BFSC model highlights the power of domain adaptation in clinical NLP tasks. By leveraging the domain-specific BioBERT language model and employing a sequence classifier, this approach demonstrates superior performance due to its fine-tuning on meticulously curated training data. While the BFSC model slightly underperforms compared to AssertionDL in the *conditional* label category, its performance is close to the benchmark, and it holds potential for further improvement through strategic augmentation of the training dataset.

Despite its superior performance, LLM-based solutions come with substantial computational costs, requiring GPUs to run efficiently while still being slower. In our benchmarks, what takes around 3 seconds using our deep-learning-based approach on a CPU requires around 300 seconds on a GPU-powered LLM, which is 100× slower. **Given that GPU instances cost more than CPU instances—often 10–50× higher per hour—the operational cost of running LLM-based assertion detection can be thousands of times more expensive for only a 1-2% accuracy gain.** This highlights the trade-off between accuracy and feasibility, where our lightweight, domain-adapted models provide a far more scalable and cost-effective alternative for real-world clinical NLP applications (see Table A3).

4. Limitations

While this study demonstrates notable advancements in clinical assertion detection, several limitations should be acknowledged. The models were benchmarked exclusively on the i2b2 dataset, which may limit generalizability to diverse clinical contexts. However, beyond this study, there are numerous models that have been trained for use in various domains such as oncology and radiology. Their F1 scores are available in Section A.5, Pre-trained Assertion Models in Healthcare NLP, and Table A4. These models can be accessed via the JohnSnowLabs Model Hub page[31]. Performance on underrepresented assertion types (e.g., conditional, associated with someone else) could vary in real-world settings with different label distributions.

Although the fine-tuned LLM achieved state-of-the-art accuracy, its GPU dependency and 100× slower inference speed compared to CPU-based DL models raise practical scalability concerns, potentially hindering deployment in resource-constrained healthcare environments. Despite addressing label skew (e.g., absent and present dominate the dataset), minority classes like hypothetical (3.5% prevalence) and conditional (1.2%) still showed lower F1 scores, suggesting residual bias in model predictions.

Commercial APIs (AWS, Azure) were evaluated only on overlapping entities detected by their

proprietary NER systems, introducing selection bias, and partial matches (28–35% of cases) may have skewed performance metrics for these systems.

5. Ethical Considerations

The development of clinical assertion detection models necessitates ethical scrutiny due to potential biases, privacy risks, and implications for patient care. The i2b2 dataset may contain demographic biases, risking inequitable model performance across populations. Future work should incorporate fairness audits and demographic stratification to mitigate these risks. Privacy concerns arise from processing sensitive patient data, particularly with cloud-based APIs (e.g., GPT-4o, AWS, Azure), necessitating transparent data governance frameworks for compliance with HIPAA and GDPR. The lack of interpretability in black-box models threatens clinical trust, underscoring the need for explainability tools to audit model decisions. Over-reliance on automation may lead to uncritical adoption in healthcare workflows, necessitating human-in-the-loop validation mechanisms. Additionally, the high computational cost of LLM training raises sustainability concerns, warranting efficiency-focused approaches. Future research should prioritize bias mitigation, open fairness benchmarks, ethical model documentation, and federated learning to enhance privacy and equity.

6. Conclusion

In this study, we present a comprehensive evaluation of JSL’s state-of-the-art assertion detection models, covering architectures from lightweight deep learning (DL) models to advanced fine-tuned LLMs. Overall, our fine-tuned LLM achieves the highest overall accuracy (0.962), outperforming GPT-4o (0.901) and commercial APIs by a notable margin, particularly in *Present*, *Absent*, and *Hypothetical* assertions. However, this comes at a high computational cost; Our DL-based models run 100× faster on a CPU than the LLM on a GPU, while the LLM is thousands of times more expensive for just 1-2% better accuracy. This highlights the impracticality of LLM-based assertion detection for real-time, scalable clinical NLP.

Our AssertionDL and FewShotAssertion models provide strong, efficient alternatives, excelling in categories like *Conditional* and *Associated with someone else* assertions, while BFSC achieves near-parity with our fine-tuned LLM. The Combined Pipeline outperforms all commercial solutions and offers a balance of accuracy and efficiency. As part of a scalable, production-ready Healthcare NLP library, these models seamlessly integrate with other clinical NLP components, enabling robust, high-performance assertion detection at scale. **Our results highlight that smaller, domain-specific models outperform commercial black-box solutions like GPT-4o, Azure AI, and AWS Medical Comprehend in both accuracy and scalability.** Integrated within Spark NLP, our pretrained assertion models and model architectures provide production-ready, cost-effective alternatives for clinical text analysis, filling a critical gap in extracting accurate medical insights.

References

- [1] Office of the National Coordinator for Health Information Technology, Adoption of electronic health records by hospital service type 2019-2021, health it quick stat 60, 2022. Available at: <https://www.healthit.gov/data/quickstats/adoption-electronic-health-records-hospital-service-type-2019-2021>.
- [2] A. Wen, L. Wang, H. He, S. Fu, S. Liu, D. A. Hanauer, D. R. Harris, R. Kavuluru, R. Zhang, K. Natarajan, et al., A case demonstration of the open health natural language processing toolkit from the national covid-19 cohort collaborative and the researching covid to enhance recovery programs for a natural language processing system for covid-19 or postacute sequelae of sars cov-2 infection: Algorithm development and validation, *JMIR medical informatics* 12 (2024) e49997.
- [3] W. W. Chapman, W. Bridewell, P. Hanbury, G. F. Cooper, B. G. Buchanan, A simple algorithm for identifying negated findings and diseases in discharge summaries, *Journal of biomedical informatics* 34 (2001) 301–310.
- [4] W. Chapman, J. Dowling, D. Chu, Context: An algorithm for identifying contextual features from clinical text, in: *Biological, translational, and clinical language processing*, 2007, pp. 81–88.
- [5] P. G. Mutalik, A. Deshpande, P. M. Nadkarni, Use of general-purpose negation detection to augment concept indexing of medical documents: a quantitative study using the umls, *Journal of the American Medical Informatics Association* 8 (2001) 598–609.
- [6] D. B. Aronow, F. Fangfang, W. B. Croft, Ad hoc classification of radiology reports, *Journal of the American Medical Informatics Association* 6 (1999) 393–411.
- [7] N. Perez, M. Cuadros, G. Rigau, Negation and speculation processing: A study on cue-scope labelling and assertion classification in spanish clinical text, *Artificial Intelligence in Medicine* 145 (2023) 102682. doi:10.1016/j.artmed.2023.102682.
- [8] Ö. Uzuner, X. Zhang, T. Sibanda, Machine learning and rule-based approaches to assertion classification, *Journal of the American Medical Informatics Association* 16 (2009) 109–115.
- [9] P. Bhatia, Joint entity extraction and assertion detection for clinical text, in: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, 2019, pp. 954–959.
- [10] L. Chen, Attention-based deep learning system for negation and assertion detection in clinical notes, *International Journal of Artificial Intelligence and Applications (IJAIA)* 10 (2019). URL: <https://ssrn.com/abstract=3342402>.
- [11] B. van Aken, I. Trajanovska, A. Siu, M. Mayrdorfer, K. Budde, A. Loeser, Assertion detection in clinical notes: Medical language models to the rescue?, in: C. Shivade, R. Gangadharaiah, S. Gella, S. Konam, S. Yuan, Y. Zhang, P. Bhatia, B. Wallace (Eds.), *Proceedings of the Second Workshop on Natural Language Processing for Medical Conversations*, Association for Computational Linguistics, Online, 2021, pp. 35–40. URL: <https://aclanthology.org/2021.nlpmc-1.5/>. doi:10.18653/v1/2021.nlpmc-1.5.
- [12] S. Wang, L. Tang, A. Majety, J. F. Rousseau, G. Shih, Y. Ding, Y. Peng, Trustworthy assertion classification through prompting, *Journal of Biomedical Informatics* 132 (2022) 104139. URL: <https://www.sciencedirect.com/science/article/pii/S1532046422001538>. doi:<https://doi.org/10.1016/j.jbi.2022.104139>.
- [13] H. Yuan, S. Yu, Black-box segmentation of electronic medical records, 2024. URL: <https://arxiv.org/abs/2409.19796>. arXiv:2409.19796.
- [14] Y. Ji, Z. Yu, Y. Wang, Assertion detection large language model in-context learning lora fine-tuning, 2024. URL: <https://arxiv.org/abs/2401.17602>. arXiv:2401.17602.
- [15] V. Kocaman, D. Talby, Accurate clinical and biomedical named entity recognition at scale,

- Software Impacts 13 (2022) 100373.
- [16] V. Kocaman, D. Talby, Spark nlp: natural language understanding at scale, *Software Impacts* 8 (2021) 100058.
 - [17] H. U. Haq, V. Kocaman, D. Talby, Deeper clinical document understanding using relation extraction, 2021. URL: <https://arxiv.org/abs/2112.13259>. arXiv:2112.13259.
 - [18] H. U. Haq, V. Kocaman, D. Talby, Mining adverse drug reactions from unstructured mediums at scale, 2022. URL: <https://arxiv.org/abs/2201.01405>. arXiv:2201.01405.
 - [19] F. Fancellu, A. Lopez, B. Webber, Neural networks for negation scope detection, in: K. Erk, N. A. Smith (Eds.), *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Berlin, Germany, 2016, pp. 495–504. URL: <https://aclanthology.org/P16-1047/>. doi:10.18653/v1/P16-1047.
 - [20] A. Grattafiori, A. Dubey, A. J. et al, The llama 3 herd of models, 2024. URL: <https://arxiv.org/abs/2407.21783>. arXiv:2407.21783.
 - [21] X. Wang, L. Aitchison, M. Rudolph, Lora ensembles for large language model fine-tuning, arXiv preprint arXiv:2310.00035 (2023).
 - [22] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, in: *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 2019, pp. 4171–4186.
 - [23] E. Alsentzer, J. R. Murphy, W. Boag, W.-H. Weng, D. Jin, T. Naumann, M. B. A. McDermott, Publicly available clinical bert embeddings, 2019. URL: <https://arxiv.org/abs/1904.03323>. arXiv:1904.03323.
 - [24] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, J. Kang, Biobert: a pre-trained biomedical language representation model for biomedical text mining, *Bioinformatics* 36 (2020) 1234–1240.
 - [25] L. Tunstall, N. Reimers, U. E. S. Jo, L. Bates, D. Korat, M. Wasserblat, O. Pereg, Efficient few-shot learning without prompts, 2022. URL: <https://arxiv.org/abs/2209.11055>. arXiv:2209.11055.
 - [26] H. Harkema, J. N. Dowling, T. Thornblade, W. W. Chapman, Context: an algorithm for determining negation, experience, and temporal status from clinical reports, *Journal of biomedical informatics* 42 (2009) 839–851.
 - [27] W. W. Chapman, W. Bridewell, P. Hanbury, G. F. Cooper, B. G. Buchanan, A simple algorithm for identifying negated findings and diseases in discharge summaries, *Journal of biomedical informatics* 34 5 (2001) 301–10. URL: <https://api.semanticscholar.org/CorpusID:6315215>.
 - [28] O. Uzuner, B. R. South, S. Shen, S. L. DuVall, 2010 i2b2/va challenge on concepts, assertions, and relations in clinical text, *Journal of the American Medical Informatics Association* 18 (2011) 552–556. URL: <https://doi.org/10.1136/amiajnl-2011-000203>. doi:10.1136/amiajnl-2011-000203.
 - [29] S. Wang, L. Tang, A. Majety, J. F. Rousseau, G. Shih, Y. Ding, Y. Peng, Trustworthy assertion classification through prompting, *Journal of biomedical informatics* 132 (2022) 104139.
 - [30] J. Zhao, et al., Lora land: 310 fine-tuned llms that rival gpt-4, a technical report, arXiv preprint arXiv:2405.00732 (2024).
 - [31] JohnSnowLabs, Johnsnowlabs model hub, <https://nlp.johnsnowlabs.com/models>, 2025. Accessed: March 5, 2025.

Appendix

A.1 A Spark NLP pipeline

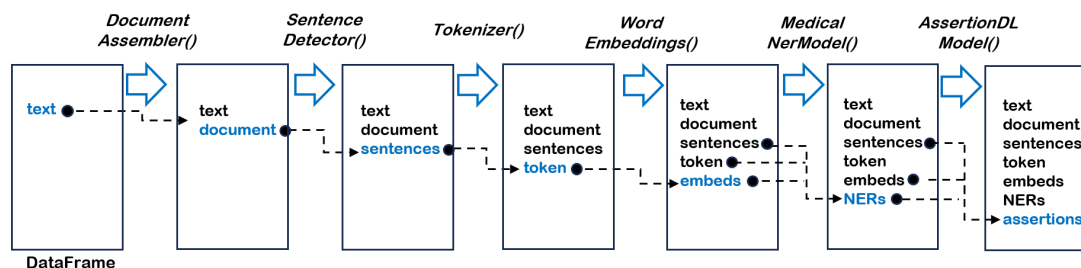


Figure A1: The flow diagram of a Spark NLP pipeline. When we fit() on the pipeline with a Spark data frame, its text column is fed into the DocumentAssembler() transformer and a new column document is created as an initial entry point to Spark NLP for any Spark data frame. Then, its document column is fed into the SentenceDetector(), Tokenizer() and WordEmbeddings(). Now data is ready to be fed into NER models and then to the assertion model.

A.2 Entity Overlapping Rates

Table A1

Entity overlapping rates across the methods evaluated in this study. Percentages are calculated relative to Total Rows.

Model	Predicted Labels	Full Match (%)	Partial Match (%)	No Match (%)	Predicted Rows	Total Rows
Fine-Tuned LLM	present, absent, possible, hypothetical, associated with someone else	12444 (100.00%)	0 (0.00%)	0 (0.00%)	12444	12444
BFSC (BioBert)	present, absent, possible, hypothetical, associated with someone else	12592 (100.00%)	0 (0.00%)	0 (0.00%)	12592	12592
AssertionDL	present, absent, possible, hypothetical, conditional, associated with someone else	12592 (100.00%)	0 (0.00%)	0 (0.00%)	12592	12592
FewShotAssertion	present, absent, possible, hypothetical, conditional, associated with someone else	12592 (100.00%)	0 (0.00%)	0 (0.00%)	12592	12592
Combined Pipeline	present, absent, possible, hypothetical, conditional, associated with someone else	12592 (100.00%)	0 (0.00%)	0 (0.00%)	12592	12592
ContextualAssertion	absent, possible, associated with someone else	3525 (100.00%)	0 (0.00%)	0 (0.00%)	3377	3377
GPT-4o	present, absent, possible, hypothetical, associated with someone else	12444 (100.00%)	0 (0.00%)	0 (0.00%)	12444	12444
Azure Ai Textanalytics	absent, possible, hypothetical, associated with someone else	2714 (68.37%)	1150 (28.96%)	103 (2.59%)	3867	3970
AWS Comprehend	present, absent, possible, hypothetical, associated with someone else	7644 (61.41%)	4427 (35.57%)	373 (3.00%)	12071	12444
Negex	absent	1443 (55.64%)	535 (20.63%)	616 (23.73%)	1978	2594

*associated with someone else: Refers to medical conditions related to individuals other than the patient, such as family members.

Table A2

Label Correspondences: AWS model does not represent *conditional* labels. Azure AI Text Analytics for Health has 3 distinct *possible* labels - *positive*, *negative* and *neutral*, while *present* is not represented.

i2b2 Label	AWS Comprehend Label	Azure Ai Textanalytics
Present	SIGN/SYMPTOM/DIAGNOSIS	N/A
Absent	NEGATION	Negative
Possible	LOW_CONFIDENCE	Possible
Hypothetical	HYPOTHETICAL	Hypothetical
Conditional	N/A	Conditional
Associated_with_someone_else	PERTAINS_TO_FAMILY	Other

A.3 Label Correspondence

A.4 Latency Comparison

Table A3

Mean latency per 100 rows, measured in seconds for various assertion methods. Experiments were run on Google Colab servers, with CPU tasks performed on a CPU instance (8vCPU @ 2.2 GHz, 50.99 GB RAM) and GPU tasks executed on an NVIDIA A100 GPU (40 GB HBM2).

Methods	Parameter Size	CPU(seconds)	GPU(seconds)
Fine-Tuned LLM	8 Billion	N/A	294
Combined Pipeline	N/A	12	4
AssertionDL	11 Million	3	2
FewShotAssertion	109 Million	5	2
ContextualAssertion	N/A	2	1
BFSC	110 Million	5	4
Negex	N/A	1	1

A.5 Pretrained Assertion Models in Healthcare NLP

Table A4

Comparison of pretrained models across various assertion categories. Each row corresponds to a specific pretrained model—such as Clinical, Oncology, Radiology, Smoking, Menopause, and Social Determinants of Health (SDOH)—highlighting their strengths in different assertion labels. Overall, this comparison serves as a valuable resource for practitioners and researchers seeking to select the most appropriate model for their specific application in medical text analysis.

Categories	Clinical	Oncology	Radiology	Smoking	Menopause	SDOH
Present	0.95	0.90	-	0.92	0.90	0.87
Past	0.91	0.93	-	0.91	0.43	0.71
Possible	0.86	0.77	-	-	0.52	0.74
Absent	0.97	0.86	-	0.97	0.84	0.94
Hypothetical	0.89	0.71	-	-	0.75	0.83
Family	0.92	0.92	-	-	0.57	-
Someone Else	0.94	-	-	-	-	0.81
Planned	0.82	-	-	-	0.73	-
Conditional	0.30	-	-	-	-	-
Confirmed	-	-	0.95	-	-	-
Negative	-	-	0.96	-	-	-
Suspected	-	-	0.86	-	-	-

A.6 GPT Prompt

GPT-4o Prompt

You are a highly experienced medical data expert specializing in patient medical records.

In this context, an assertion refers to the sentiment or condition associated with a specific medical entity within the context of a patient's record. This helps determine whether symptoms or conditions are present, absent, possible, hypothetical, or related to someone else, enhancing the precision of medical documentation and analysis.

Your task is to detect the assertion status of medical conditions mentioned in notes. The possible assertion types are:

- **absent**: condition is explicitly negated
- **associated_with_someone_else**: condition refers to someone other than the patient
- **conditional**: condition is mentioned as contingent on another factor
- **hypothetical**: condition is part of a hypothetical scenario
- **possible**: condition is suggested as a possibility but not confirmed
- **present**: condition is clearly present for the patient

Instructions:

- 1 Analyze the input TEXT and identify the assertion status of the TARGET condition.
- 2 Format your answer in valid JSON, using double quotes for both keys and values.
- 3 If multiple assertions are required, choose the most confident one.

EXAMPLE INPUT

```
{
  "TEXT": "She was then started on Heparin with transition to Coumadin (goal INR of 2-3 secondary to h/o bilateral DVTs).",
  "TARGET": "bilateral DVT"
}
```

INPUT

```
{
  "TEXT": "text",
  "TARGET": "target"
}
```

Your Answer in JSON:

Provide a JSON object where the text and assertion type are the key-value pairs.

Example Output Format:

```
{
  "TARGET": "bilateral dvt",
  "ASSERTION_STATUS": "present"
}
```

Figure A2: Example of GPT-4o prompt for detecting assertion status in medical records

A.7 Fine-tuned LLM Prompt

Fine-tuned LLM Prompt

You are provided with a document and an extracted entity (chunk).

Your job is to analyze the document and the chunk, understand the context, and assign one of the following statuses to the chunk:

- ****present****: If the chunk is mentioned in the context of the person. *Example*: “He has a fractured ankle.”
- ****absent****: If the chunk is explicitly negated by the person. *Example*: “He did not suffer from pain.” (In this case, “pain” is absent/negated.)
- ****hypothetical****: If the chunk is mentioned in a hypothetical scenario or as part of guidelines. *Example*: “Adults above 70 are at greater risk of cancer.” (Here, “cancer” is hypothetical.)
- ****possible****: If the chunk is mentioned in a way that implies possibility. *Example*: “Possible fracture.”
- ****associated_with_someone_else****: If the condition refers to someone other than the patient. *Example*: “Her mother has breast cancer.”

Document:

```
{  
  "DOCUMENT": "doc"  
}
```

Chunk:

```
{  
  "CHUNK": "chunks"  
}
```

Your Answer in JSON:

Provide a JSON object where the chunk and assertion status are the key-value pairs.

Example Output Format:

```
{  
  "CHUNK": "fractured ankle",  
  "ASSERTION_STATUS": "present"  
}
```

Figure A3: Example of Fine-tuned LLM prompt for detecting assertion status in medical records

A.8 Healthcare NLP Pipeline

```
1 from sparknlp.base import DocumentAssembler
2 from sparknlp.annotator import Tokenizer, WordEmbeddingsModel, AssertionDLModel
3 from sparknlp_jsl.annotator import AssertionChunkConverter
4
5 # Document Assembler
6 document_assembler = DocumentAssembler() \
7     .setInputCol("text") \
8     .setOutputCol("sentence")
9
10 # Tokenizer
11 tokenizer = Tokenizer() \
12     .setInputCols(["sentence"]) \
13     .setOutputCol("token")
14
15 # Assertion Chunk Converter
16 converter = AssertionChunkConverter() \
17     .setInputCols("token")\
18     .setChunkTextCol("chunk")\
19     .setChunkBeginCol("begin")\
20     .setChunkEndCol("end")\
21     .setOutputTokenBeginCol("token_begin")\
22     .setOutputTokenEndCol("token_end")\
23     .setOutputCol("ner_chunk")
24
25 # Word Embeddings Model
26 word_embeddings_100 = WordEmbeddingsModel.pretrained("embeddings_healthcare_100d", "en", "clinical/
27     models")\
28     .setInputCols(["sentence", "token"])\
29     .setOutputCol("embeddings")
30
31 # AssertionDL Model
32 clinical_assertion_100 = AssertionDLModel.pretrained("assertion_dl_healthcare", "en", "clinical/
33     models") \
34     .setInputCols(["sentence", "ner_chunk", "embeddings"]) \
35     .setOutputCol("assertionDL")\
36     .setEntityAssertionCaseSensitive(False)
37
38 #FewShotAssertion Model
39 few_shot_assertion_converter = FewShotAssertionSentenceConverter()\
40     .setInputCols(["sentence", "token", "ner_chunk"])\
41     .setOutputCol("assertion_sentence")
42
43 e5_embeddings = E5Embeddings.pretrained("e5_base_v2_embeddings_medical_assertion_i2b2", "en", "
44     clinical/models")\
45     .setInputCols(["assertion_sentence"])\
46     .setOutputCol("assertion_embedding")
47
48 few_shot_assertion_classifier = FewShotAssertionClassifierModel()\
49     .pretrained("fewhot_assertion_i2b2_e5_base_v2_i2b2", "en", "clinical/models")\
50     .setInputCols(["assertion_embedding"])\
51     .setOutputCol("assertion_fewshot")
```

```

1 #Contextual Assertion Models
2 contextual_assertion_possible = ContextualAssertion.pretrained("contextual_assertion_possible", "en"
3     , "clinical/models")\
4     .setInputCols("sentence", "token", "ner_chunk") \
5     .setOutputCol("ca_possible")
6
7 contextual_assertion_conditional = ContextualAssertion.pretrained("contextual_assertion_conditional
8     ", "en", "clinical/models")\
9     .setInputCols("sentence", "token", "ner_chunk") \
10    .setOutputCol("ca_conditional")
11
12 #Merger
13 assertionMerger_fewshot = AssertionMerger()\
14    .setInputCols("assertion_fewshot")\
15    .setOutputCol("assertion_merger_fewshot")\
16    .setWhiteList(["absent", "hypothetical"])
17
18 assertionMerger_dl = AssertionMerger()\
19    .setInputCols("assertionDL")\
20    .setOutputCol("assertion_merger_dl")\
21    .setWhiteList(["associated_with_someone_else", "conditional"])
22
23 assertionMerger_all = AssertionMerger()\
24    .setInputCols("assertionDL", "assertion_fewshot", "ca_possible")\
25    .setOutputCol("assertion_merger_all")\
26    .setMergeOverlapping(True)\
27    .setMajorityVoting(False)\
28    .setOrderingFeatures(["confidence"])\
29    .setWhiteList(["present", "possible"])\
30    .setApplyFilterBeforeMerge(True)
31
32 assertionMerger_final = AssertionMerger()\
33    .setInputCols("assertion_merger_fewshot", "assertion_merger_dl", "assertion_merger_all", "
34        ca_conditional")\
35    .setOutputCol("assertion_merger")\
36    .setMergeOverlapping(True)\
37    .setMajorityVoting(True)\
38    .setOrderingFeatures(["confidence"])\
39
40 #Pipeline
41 pipeline = Pipeline(stages=[
42     document_assembler,
43     tokenizer,
44     converter,
45     few_shot_assertion_converter,
46     e5_embeddings,
47     few_shot_assertion_classifier,
48     word_embeddings_100,
49     clinical_assertion_100,
50     assertionMerger_fewshot,
51     contextual_assertion_conditional,
52     contextual_assertion_possible,
53     assertionMerger_dl,
54     assertionMerger_all,
55     assertionMerger_final
56 ])
57
58 empty_data = spark.createDataFrame([[""]]).toDF("text")
59 dataframe = spark.createDataFrame(data_df)
60 model = pipeline.fit(empty_data)
61 result = model.transform(dataframe)

```