

Transformer-Based Models for Domain-Specific Natural Language Processing

Geetu Soni, Student, Department of Computer Science, Singhania University (Rajasthan)

Abstract

Natural Language Processing (NLP) has experienced significant advancements with the emergence of Transformer-based architectures, which have demonstrated remarkable capabilities in understanding and generating human language. However, general-purpose language models often face limitations when applied to specialized domains such as healthcare, finance, law, education, and cybersecurity, where domain-specific terminology, contextual relationships, and specialized knowledge are essential. This paper investigates the application of Transformer-based models for domain-specific Natural Language Processing and examines their effectiveness in addressing these challenges. Models based on architectures such as BERT, RoBERTa, and domain-adapted Transformer variants can be fine-tuned on specialized datasets to improve performance on tasks including text classification, named entity recognition, information extraction, question answering, and sentiment analysis. The study discusses key approaches for domain adaptation, including continued pre-training, task-specific fine-tuning, transfer learning, and parameter-efficient adaptation techniques. Particular attention is given to the challenges of limited labelled data, domain shift, computational requirements, model interpretability, and potential bias in specialized datasets. Performance can be evaluated using task-specific metrics such as accuracy, precision, recall, F1-score, and perplexity. The paper highlights how domain-adapted Transformer models can provide more accurate and context-aware language understanding compared with conventional NLP approaches. The findings emphasize the potential of Transformer-based architectures to develop efficient, scalable, and domain-aware NLP applications across diverse real-world environments.

Keywords: Transformer, Natural Language Processing, BERT, RoBERTa, Domain Adaptation, Transfer Learning, Fine-Tuning, Named Entity Recognition, Text Classification, Large Language Models.