Poster id: 11

# Summarization Assistant for Academic Research with XAI


Karan Kendre, Prof. Tehmina Amjad, MS in Computer Science

Northeastern University
Khoury College of Computer Sciences


## INTRODUCTION

The exponential growth of AI research has created an unprecedented challenge: navigating vast volumes of academic literature efficiently. As researchers struggle with **information overload**, the limitations of current Large Language Models become apparent particularly their tendency to produce unreliable outputs when processing extensive documents. This critical gap in academic workflow demands an intelligent solution that can distill complex research papers into clear, **trustworthy** summaries while maintaining **transparency** about source material.

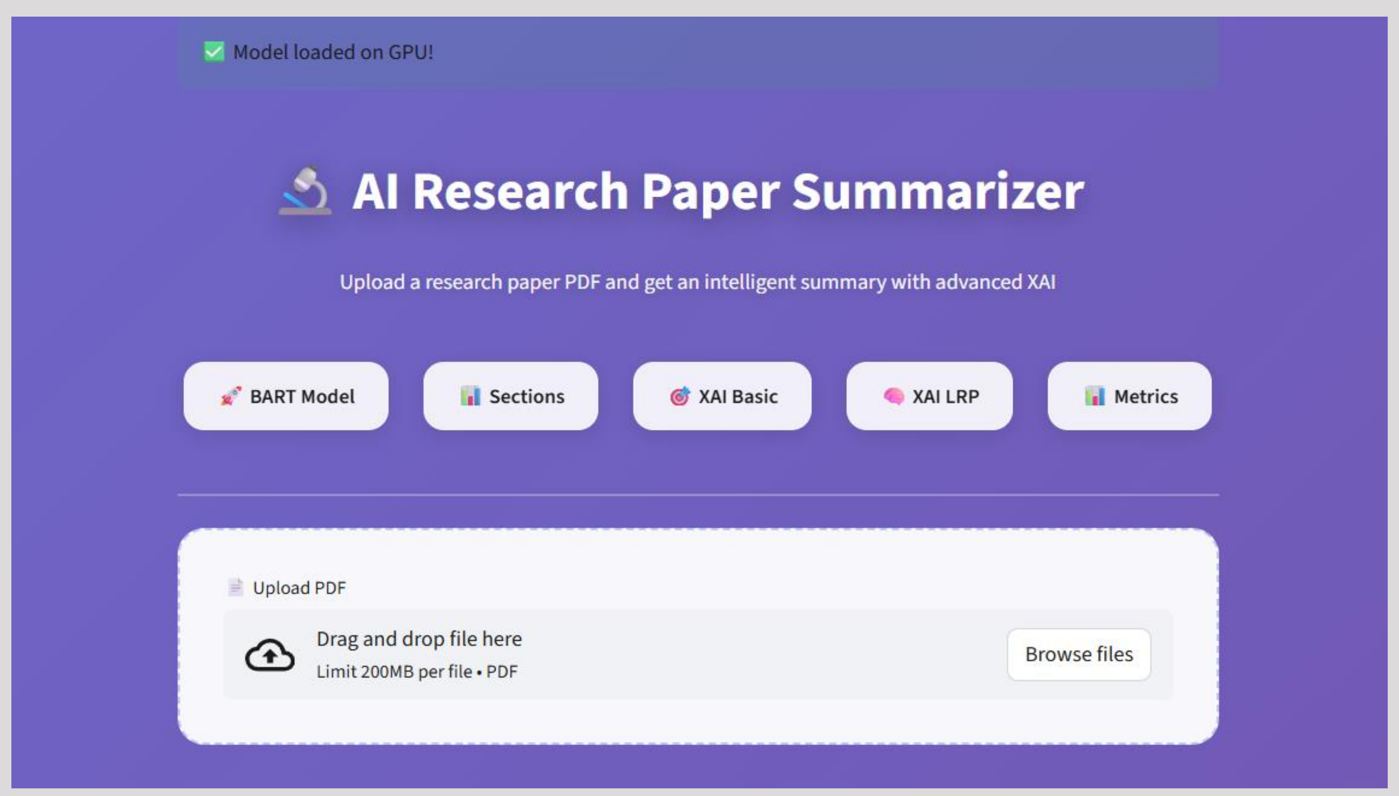


*Figure 1: Streamlit-based interface with animated gradient background and glassmorphism design elements*

## PROBLEM STATEMENT

The absence of a reliable system that combines depth with transparency creates friction in the research process, slowing down knowledge discovery and synthesis. This project addresses this gap by developing a summarization platform that produces contextually **rich outputs** while maintaining **complete traceability** to original sources, enabling researchers to work with both efficiency and confidence.

## APPROACH

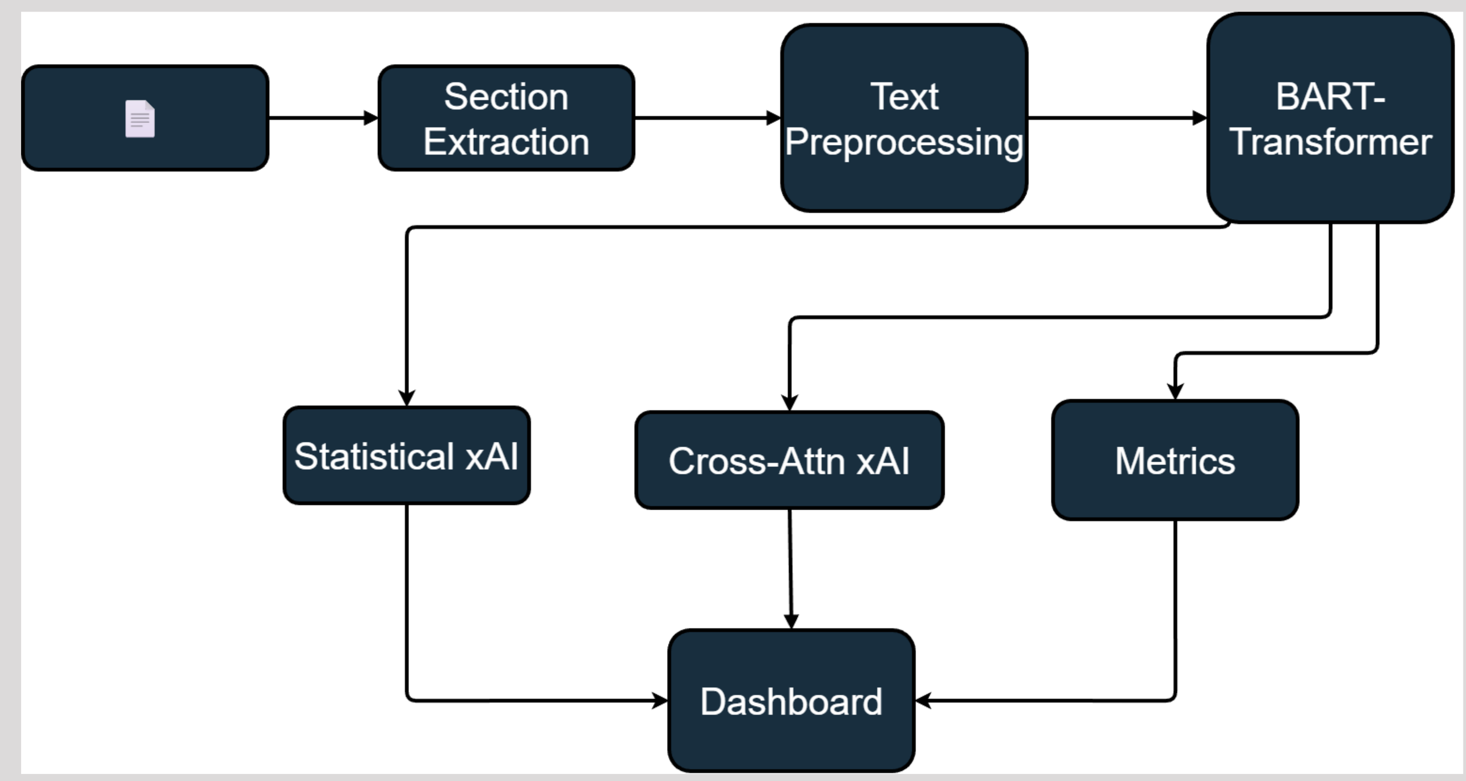


*Figure 2: End-to-end processing pipeline*

1. **Basic Features**
   - Model: facebook/bart-large-cnn, pretrained on CNN/DailyMail dataset [1]
   - Smart chunking for long sections (> 900 tokens), recursive summarization for coherence
   - Streamlit UI, with section extraction and custom summary length
2. **Basic xAI (TF – IDF Analysis)**
   - Calculates cosine similarity between vectors of source sentences and summary
   - Performance: Fast (~1-2 seconds), no GPU required
3. **Advance xAI (Cross-attention LRP)**
   - Extracts decoder's cross attention weights showing actual model focus when it's writing summary [2]
   - Performance: Compute takes time, GPU required
4. **Evaluation Metrics**
   - ROUGE Scores: Unigram, Bi-gram, LCS overlap score
   - BERTScore: Semantic similarity using BERT embeddings for contextual meaning preservation of the summary

## RESULT

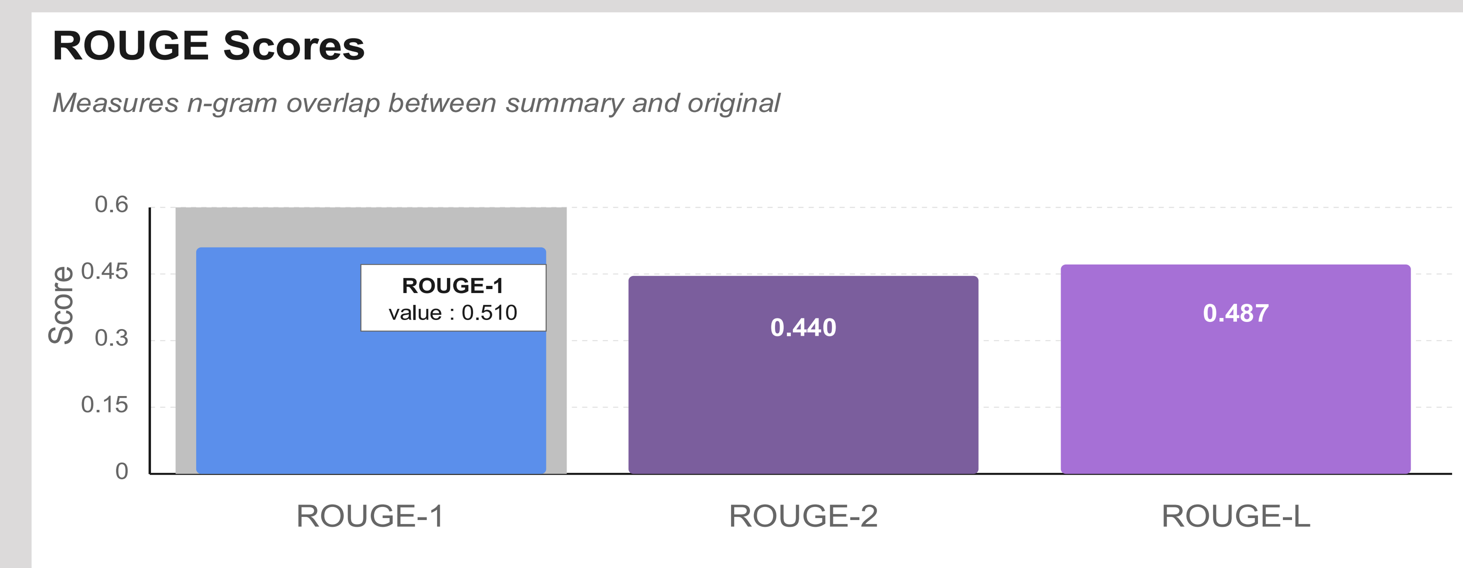


*Figure 3: ROUGE Scores evaluations*

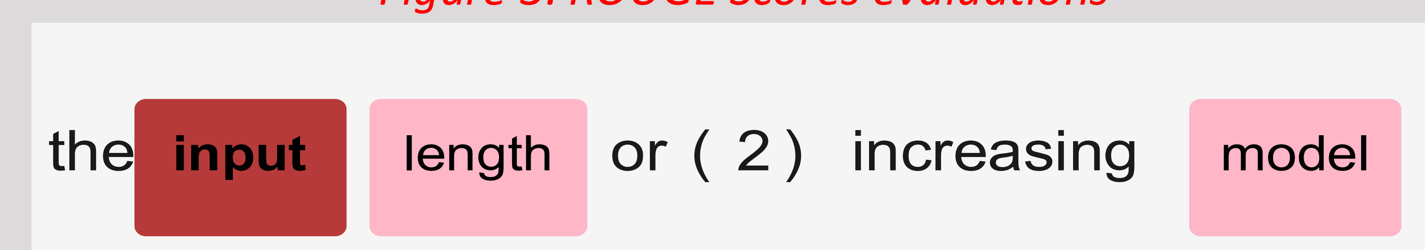


*Figure 4: Layer-wise cross-attention token visualization*

The system provides comprehensive metrics (ROUGE, BERTScore, Readability). Our model focuses on technical terms and methodological phrases, with dark red tokens (attention >0.7) highlighting concepts included in the generated summary.

## CONCLUSION

By combining state-of-the-art neural summarization with transparent explainability, this system presents that AI can now be both powerful as well as trustable.

## REFERENCES


[1] Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., & Zettlemoyer, L. (2020). BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension.

[2] Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K. R., & Samek, W. (2015). On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation.