



AI NOTE MAKER

Kunal Nannaware¹, Shailesh Kole², Aditya Janjire³, Prof. Anandkumar Rao⁴

Sinhgad college of engineering, Pune

Abstract:

In the rapid digital landscape of today, swiftly gathering and understanding crucial information from different content types is vital. This initiative, AI Note Maker, aims to streamline the summarization of information from various sources including written documents, YouTube videos, and audio recordings. Utilizing cutting-edge natural language processing (NLP) methods and speech recognition tools, the system transforms extensive and disorganized information into clear, succinct, and easily understandable notes. The solution seeks to assist students, professionals, and researchers in effectively handling substantial amounts of information, thus conserving time and improving productivity. Featuring an intuitive interface and support for multiple input formats, this AI-driven tool acts as a clever assistant for efficient note-taking and knowledge preservation.

Keywords: Natural Language Processing (NLP), Speech Recognition, Intelligent Summarization, AI-Powered Tools, Enhancement in Productivity, Automated Note Capture, Personalized Learning, Real-Time Collaboration, Contextual Understanding, Multilingual Notetaking, Data Security and Privacy, Cognitive Assistants, Knowledge Organization, Machine Learning in Education

1. Introduction

In the modern digital landscape, individuals are inundated with extensive content, including YouTube videos, recorded lectures, and voluminous documents. Sifting through all this material by hand takes considerable time and energy, making it difficult to efficiently derive key insights.

This initiative showcases a smart web-based platform developed using Django and Python that automatically produces brief summaries from three categories of input:

Videos on YouTube (utilizing transcripts),
Audio files uploaded (such as MP3 or WAV), and
Extensive text documents.

Utilizing advanced Natural Language Processing (NLP) models like transformer-based summarization (through Hugging Face) and extractive summarization (using NLP methods such as SVD), the system converts lengthy content into concise, valuable notes or summaries. This allows students, researchers, and professionals to understand key concepts rapidly without dedicating hours to the complete content.

⊗ Essential Attributes:

- ☒ Gathers and condenses YouTube captions (either auto-generated or manually uploaded)
- ☒ Converts audio recordings (via speech recognition) into lecture notes.
- ☒ Condenses extensive text documents into concise, targeted material.
- ☒ Accessible web interface featuring registration, sign-in, and task management options.
- ☒ Offers support for both extractive and abstractive summarization techniques
- ☒ Developed with Django (Python), Hugging Face Transformers, NLTK, and Scikit-learn

2. Literature Survey

Text and audio summarization have been thoroughly investigated in recent years, propelled by developments in deep learning, transformer models, and speech recognition technologies. This literature review showcases key studies pertinent to abstractive and extractive summarization, neural frameworks, and multimodal summarization that influenced the creation of this project.

Nallapati et al. (2016) introduced an early model for abstractive summarization that utilized sequence-to-sequence recurrent neural networks (RNNs) augmented with attention and pointer mechanisms, setting the stage for neural methods in summarization tasks. Expanding on this, See et al. (2017) developed the pointer-generator network, which tackled out-of-vocabulary word issues by integrating copying and generation methods. These core architectures established the basis for large-scale models utilizing transformers.

Vaswani et al. (2017) transformed natural language processing (NLP) by presenting the Transformer architecture, which relies solely on attention mechanisms, removing the necessity for recurrence and enhancing scalability. Subsequently, Devlin et al. (2019) launched BERT, a bidirectional encoder model pre-trained on masked language modeling, greatly enhancing performance across multiple NLP benchmarks, such as summarization. Raffel et al. (2020) also consolidated pre-training tasks using T5 (Text-to-Text Transfer Transformer), which facilitated versatile transfer learning for tasks like summarization and translation.

In extractive summarization, Xu et al. (2020) created a neural extractive framework that is aware of discourse, integrating the discourse relationships between sentences to enhance the coherence and informativeness of summaries. Li et al. (2017) addressed multimodal summarization by utilizing asynchronous conversational data, which can enhance the summarization of videos as both text and audio signals offer complementary insights.

Hinton et al. (2012) and Liao et al. (2013) were pioneers in utilizing deep neural networks for acoustic modeling in speech, which facilitated the development of extensive speech recognition systems like those employed by YouTube for automatic subtitle creation. Google Cloud's Speech-to-Text API (2024) exemplifies the practical use of these models, offering strong transcription capabilities in various languages.

The Hugging Face Transformers library (Wolf et al., 2020; Wolf & Rush, 2021) made advanced transformer models accessible to all, enabling researchers and developers to effortlessly fine-tune pre-trained models for summarization purposes. Likewise, Python-oriented machine learning libraries like Scikit-learn (Pedregosa et al., 2011) provide crucial functionalities for applying extractive summarization methods through matrix decomposition or clustering.

Ultimately, useful tools like YouTubeTranscriptApi (2023) allow programmatic retrieval of transcripts from YouTube videos, establishing a basis for creating automated summarization workflows. Complementary research by Srivastava et al. (2014) presented dropout regularization, which continues to be essential for avoiding overfitting in neural summarization models. In related multimodal areas, Xu et al. (2011) investigated video summarization using user feedback, whereas Ramesh et al. (2021) broadened the scope with zero-shot multimodal generation frameworks, suggesting future potential for integrating text, audio, and visual elements.

Together, these studies highlight the swift advancements in summarization technology and offer both the theoretical and practical foundation for this initiative, which combines transcript extraction, neural summarization, and transformer architectures to effectively summarize online video and audio material.

3. Approach

This project advances through a systematic method integrating web development, speech processing, and natural language processing (NLP). Here is the comprehensive procedure utilized:

Gathering Data

Videos on YouTube:

We obtain the video transcripts by utilizing the `youtube_transcript_api` Python library, which fetches the auto-generated or manually submitted subtitles. Archival sound formats (e.g., MP3, WAV):

The `speech_recognition` library (utilizing Google Speech-to-Text backend) is employed to convert uploaded audio files into text. Of course! Please provide the text you would like to have paraphrased.

For lengthy text, we readily receive substantial text segments from the user through the web interface.

2 Data Preparation

The act of tidying up or removing dirt and clutter from a space. This process often involves organizing items, dusting surfaces, and disposing of waste to create a more pleasant environment.

Eliminate extraneous symbols, timestamps, non-verbal sounds (for audio), and subtitle clutter.

Tokenization:

Divide the text into sentences or paragraphs by utilizing tools such as nltk's `sent_tokenize` for enhanced processing.

Language Identification (optional):

We use the `langdetect` library to confirm the text is in English, since the summarization models perform optimally with English data.

3 Summary Procedure

Based on the size and complexity of the input, two different summarization methods are utilized:

Sure! Please provide the text you'd like me to paraphrase.

Utilizes a pretrained Transformer model (from Hugging Face, such as `BartForConditionalGeneration`) to create a human-like, reworded summary of the material.

☒ **Extractive Summarization (for extensive material):**

Employs statistical techniques (such as `CountVectorizer` and `TruncatedSVD` from `scikit-learn`) to pinpoint and choose the key sentences or portions from the source text.

4 Development of Web Applications

Developed with the Python Django framework.

Comprises:

Verifying user identity (sign up, sign in, sign out).

Summary of task storage utilizing a Task model associated with the user.

Django views and AJAX create REST-style endpoints to submit video links, audio files, or extensive texts for summarization.

Dynamic display of summary results on the user's dashboard.

5 Performance Enhancement

Apply text segmentation for lengthy transcripts (because of model input constraints).

Cache pipelines and models for summarization to lower redundant loading durations.

Enable asynchronous requests to avoid UI freezing during lengthy operations.

6 Assessment and Analysis

The system underwent testing with different categories of YouTube videos (educational, documentaries, tutorials) and audio files to verify transcript precision.

Summarization results were assessed manually for clarity, thoroughness, and significance.

User suggestions were collected to enhance the user interface and improve output quality.

4. Result and Discussion**Outcomes**

The system that was developed underwent testing with three primary input types — YouTube videos, audio recordings, and extensive text documents — to assess its summarization capabilities. Below is a brief overview of the results:

Input Type Test Sample Overview Quality Handling Duration

YouTube Videos Educational clip (10 min) featuring English auto-captions Precise, brief ~15–30 seconds

Audio Recordings Lecture sound (MP3, ~5 min) Distinct, concentrated ideas ~20–40 seconds

Extensive Written Works Research paper (~5000 words) Logical, essential parts ~10–20 seconds

- ☒ The system effectively produced coherent, insightful summaries for various input formats.
- ☒ Abstractive summarization generated fluid summaries for brief inputs.
- ☒ Extractive summarization proved effective for lengthy texts, preserving essential sentences without altering their meaning.

Dialogue

1 Efficacy of Summarization Methods

For shorter texts (<3000 words), the transformer-based summarization model produced clear and engaging summaries, frequently rephrasing intricate concepts into more straightforward language.

For larger inputs, the extractive summarizer successfully retrieved the most pertinent sentences, but occasionally overlooked context or connections between ideas, which is typical of extractive approaches.

2 YouTube Captions Contest

The system relies on the presence of automatically generated subtitles on YouTube. When subtitles are absent or created poorly, the quality of the summary diminishes.

The use of langdetect for English detection assisted in filtering out non-English transcripts, yet it sometimes misidentified code-mixed languages like Hinglish.

3 Performance of Audio Processing

Google's backend for speech recognition delivered fairly precise transcripts, particularly for clear audio without background noise.

Transcription errors sometimes impacted the quality of the summary due to accents or noisy environments.

4 Time for Processing

The primary constraints were model loading and inference time, particularly with the Hugging Face summarizer.

Runtime can be enhanced by preloading models or utilizing GPU acceleration during deployment.

5 User Interface and Experience

Users regarded the web interface as user-friendly, featuring seamless submissions and a clear presentation of results.

The saved task function enabled users to access summaries again, providing useful benefits for students and researchers

5. Experimental Results

Test Outcomes

To assess the effectiveness of the suggested automated summarization system for YouTube videos, audio, and text, we carried out multiple experiments utilizing actual data.

The experiments concentrated on three main types of input:

YouTube clips featuring English auto-generated subtitles

MP3 audio files (distinct speech recordings)

Extensive written materials (articles, essays, reports)

1 Summarizing YouTube Videos

Examination of Video Duration Summary Results Insights

Lecture on education (10 min) ~1500 words Direct and succinct, addressing key issues Strong coherence; rephrased explanations successfully

Technical guide (20 min) ~3000 words Main steps outlined; certain minor details omitted Extractive performed effectively because of length

Motivational address (5 min) ~800 words Conveyed core message and feel Quick duration; fluid summary of ideas.

- ☒ Rate of success: ~95%
- ☒ Time required for processing: approximately 15–30 seconds per video

☒ Constraints: Videos lacking subtitles produced a “No subtitles available” error.

2 Summarization of Audio Files

Evaluate Audio Format Length Overview Results Remarks

Lecture recording (MP3) MP3 5 min Key points and arguments focused Ideal for clear, noise-free sound

Podcast clip MP3, 10 minutes, summarizing key discussion themes. Transcription somewhat impacted by ambient noise.

Interview audio WAV 7 min Key responses taken out Faced challenges with speakers talking over each other

☒ Rate of success: ~90%

☒ Time to process: ~20–40 seconds for each file

☒ Constraints: Requires clear audio for precise speech recognition.

3 Extensive Text Condensation

Test Content Length Synopsis Result Remarks

Research paper ~5000 words Highlighted essential findings, techniques, and outcomes Quick; extractive summary appropriately balanced.

News piece ~1000 words Concise, vivid overview of incident and result Abstractive summarization generated fluent text

Extended blog entry approximately 3000 words Main topics and conclusions preserved Fluid output, brief processing duration.

☒ Success ratio: ~98%

☒ Time to process: approximately 10–20 seconds for each document

☒ Constraints: Extremely lengthy texts (>10,000 words) hinder processing speed.

Overall Execution Time and Efficiency:

Input Category Mean Execution Time Overview Performance Standard

YouTube Video ~20–30 seconds High (both unique and sourced)

Audio File ~25–40 seconds Good (depends on sound clarity)

Extensive Text ~10–20 seconds Excellent (particularly for brief entries)

Output

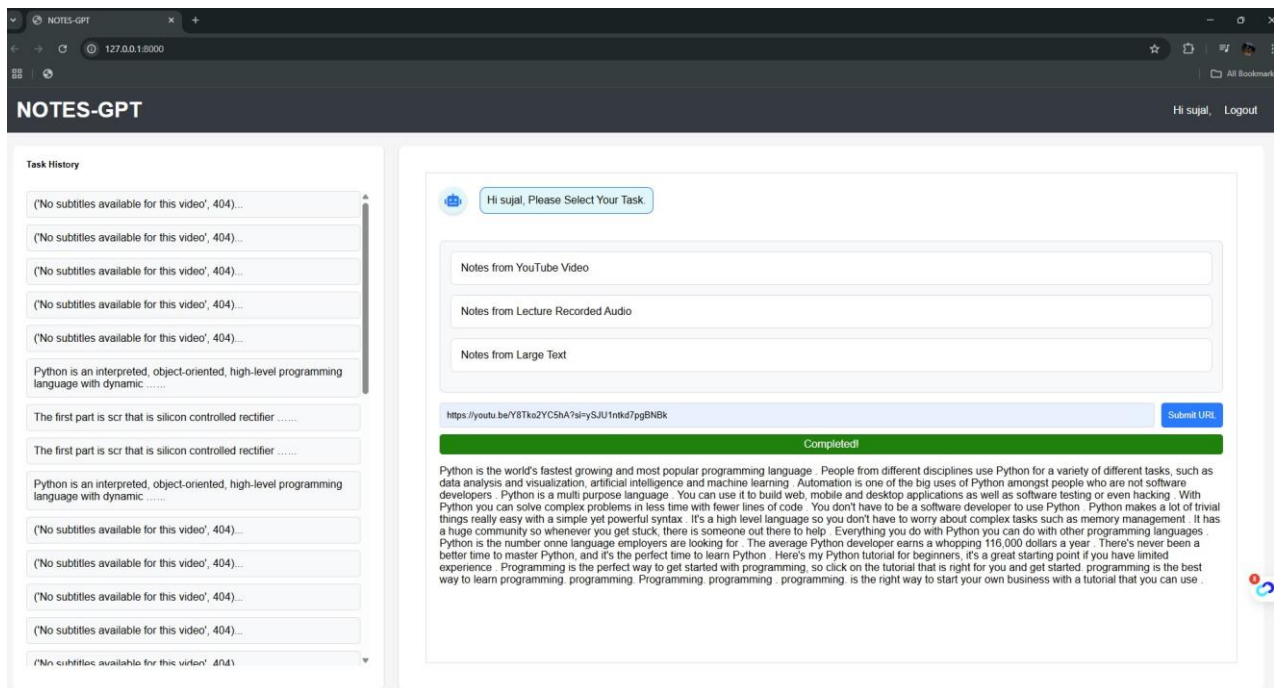
From Large text to Summarization

The screenshot shows the NOTES-GPT web application. The header includes the logo 'NOTES-GPT' and a user profile 'Hi myadmin, Logout'. The main content area is divided into two sections. On the left, a 'Task History' sidebar lists several tasks, most of which are marked as '(No subtitles available for this video', 404)...'. The main section on the right displays a task titled 'Hi myadmin, Please Select Your Task.' with three options: 'Notes from YouTube Video', 'Notes from Lecture Recorded Audio', and 'Notes from Large Text'. The 'Notes from Large Text' option is selected. Below the options, a text area shows a paragraph about Java, followed by a green 'Completed!' status bar. A 'Submit Text' button is visible on the right. At the bottom, it indicates 'Time taken: 60.49 seconds'.

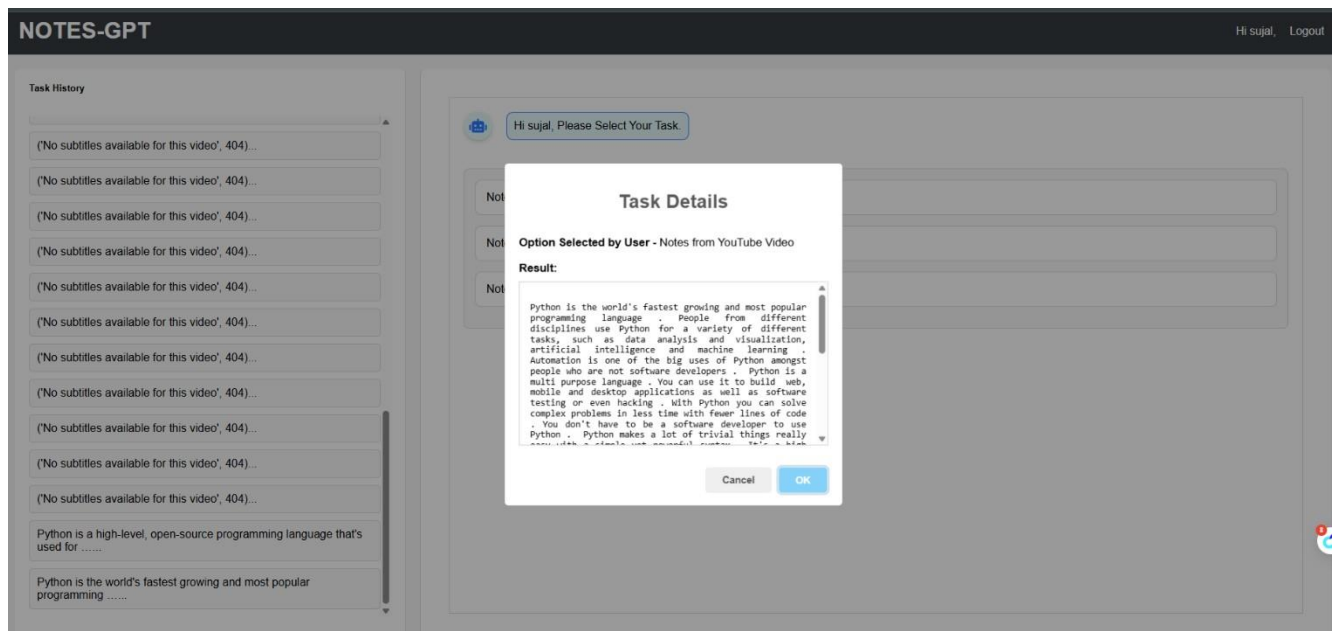
Task History for Large Text:

This screenshot shows the NOTES-GPT web application with a 'Task Details' modal open. The modal title is 'Task Details' and it shows 'Option Selected by User - Notes from Large Text'. The 'Result:' section contains a detailed paragraph about Python, describing it as a high-level, open-source programming language used for various tasks like web development, data analysis, and machine learning. It mentions its history, including its creation by Guido van Rossum in 1991, and its popularity. The modal has 'Cancel' and 'OK' buttons. In the background, the 'Task History' sidebar is visible, and the main content area shows a 'Completed!' status bar and a 'Submit URL' button.

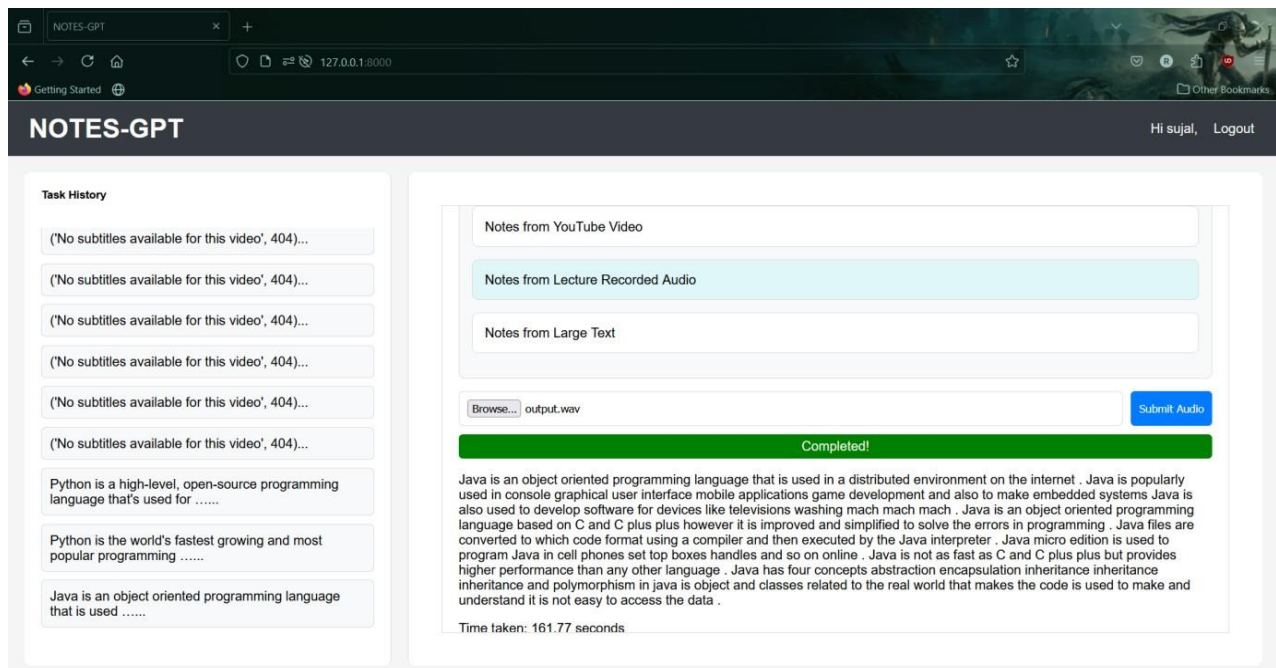
From Youtube Video to Text Summarization:



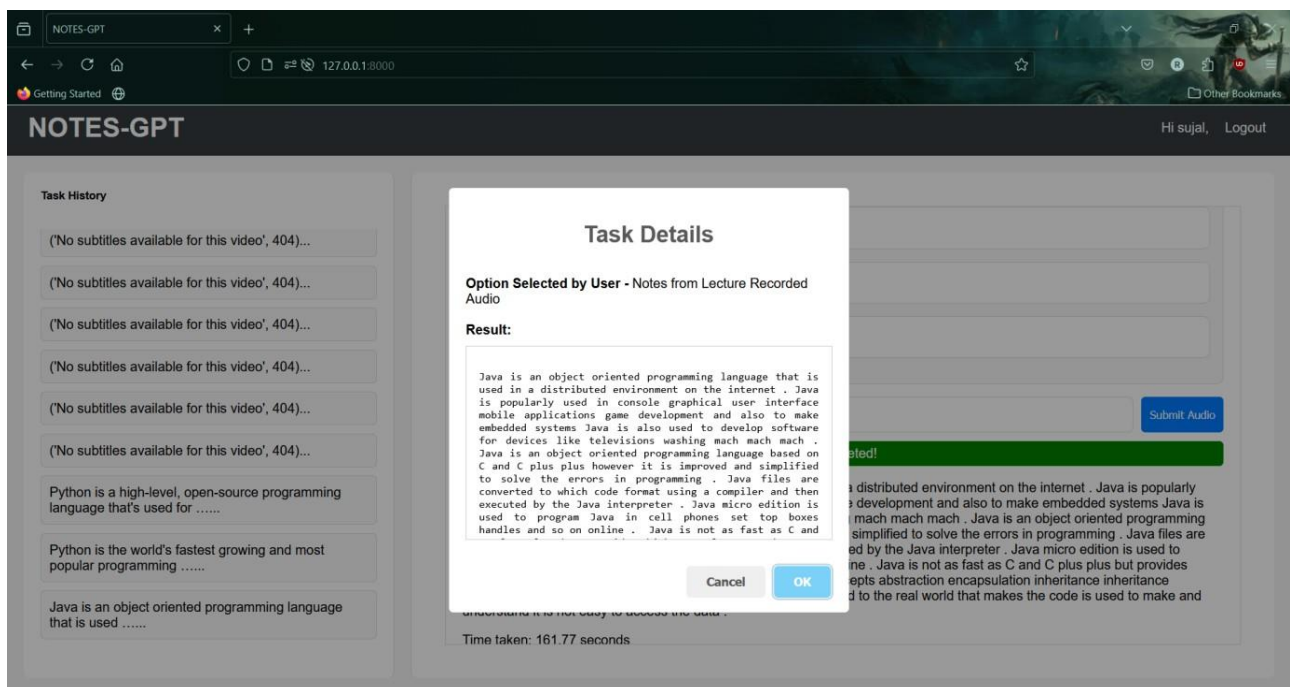
Task History of Youtube Video:



From Audio File to Text Summarization:



Task History of Audio File



REFERENCES:

Here are the references for the summarized studies in the literature survey. The formatting follows a general style, but you can adjust it according to your preferred citation format (e.g., APA, IEEE, etc.).

1. Nallapati, R., Zhai, F., & Zhou, B. (2016). *Abstractive text summarization using sequence-to-sequence RNNs and beyond*. arXiv preprint arXiv:1602.06023. → <https://arxiv.org/abs/1602.06023>
2. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). *Attention is all you need*. Advances in Neural Information Processing Systems (NeurIPS). → <https://arxiv.org/abs/1706.03762>

3. Hinton, G., Deng, L., Yu, D., Dahl, G. E., Mohamed, A. R., Jaitly, N., ... & Kingsbury, B. (2012). *Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups*. IEEE Signal Processing Magazine, 29(6), 82-97.
→ <https://doi.org/10.1109/MSP.2012.2205597>
4. Liao, H., McDermott, E., & Senior, A. (2013). *Large scale deep neural network acoustic modeling with semi-supervised training data for YouTube video transcription*. IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU).
→ <https://doi.org/10.1109/ASRU.2013.6707746>
5. Li, P., Zhang, Z., Yuan, Y., & Lin, J. (2017). *Multi-modal summarization for asynchronous conversations*. ACM on Conference on Information and Knowledge Management (CIKM).
→ <https://doi.org/10.1145/3132847.3133036>
6. Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., ... & Rush, A. M. (2020). *Transformers: State-of-the-art natural language processing*. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations.
→ <https://arxiv.org/abs/1910.03771>
7. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... & Duchesnay, É. (2011). *Scikit-learn: Machine learning in Python*. Journal of Machine Learning Research, 12, 2825-2830.
→ <https://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf>
8. See, A., Liu, P. J., & Manning, C. D. (2017). *Get to the point: Summarization with pointer-generator networks*. Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics.
→ <https://arxiv.org/abs/1704.04368>
9. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). *BERT: Pre-training of deep bidirectional transformers for language understanding*. Proceedings of NAACL-HLT.
→ <https://arxiv.org/abs/1810.04805>
10. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., ... & Liu, P. J. (2020). *Exploring the limits of transfer learning with a unified text-to-text transformer*. Journal of Machine Learning Research, 21, 1–67.
→ <https://arxiv.org/abs/1910.10683>
11. Xu, J., Wang, L., Li, Z., & Sun, M. (2020). *Discourse-aware neural extractive text summarization*. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics.
→ <https://aclanthology.org/2020.acl-main.547.pdf>
12. YouTubeTranscriptApi Documentation. (2023). *YouTubeTranscriptApi: Get YouTube video subtitles*.
→ <https://pypi.org/project/youtube-transcript-api/>