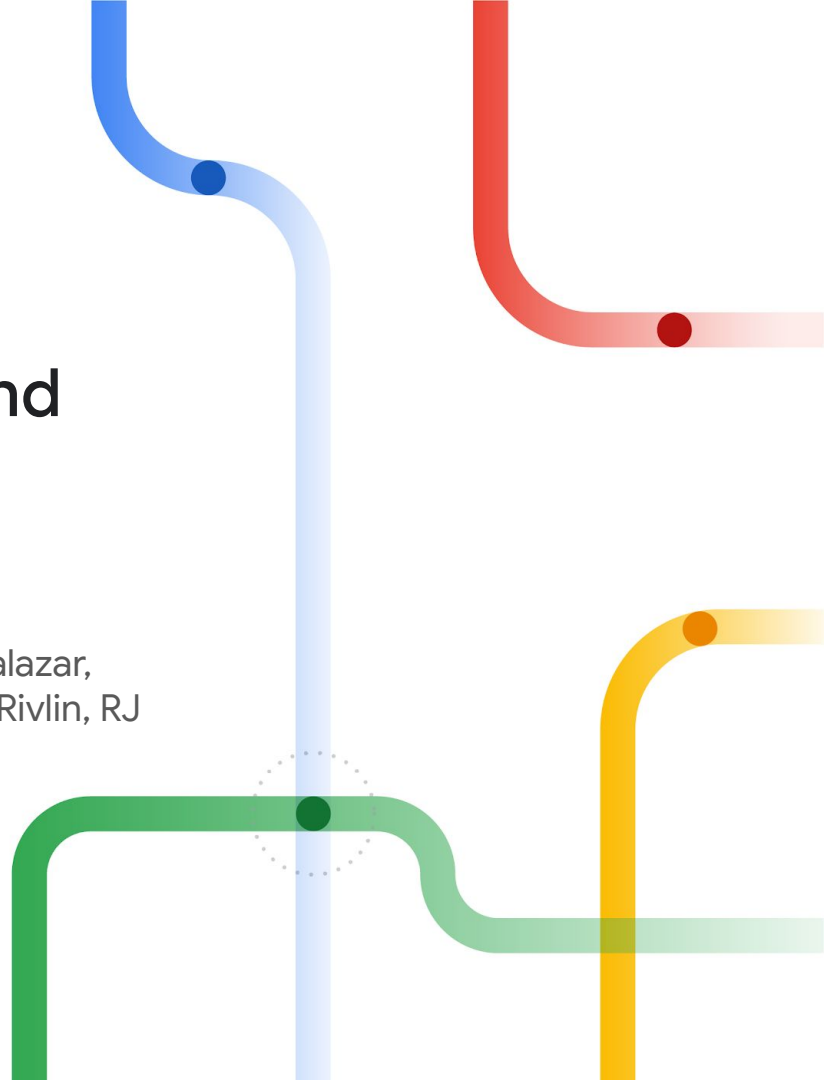


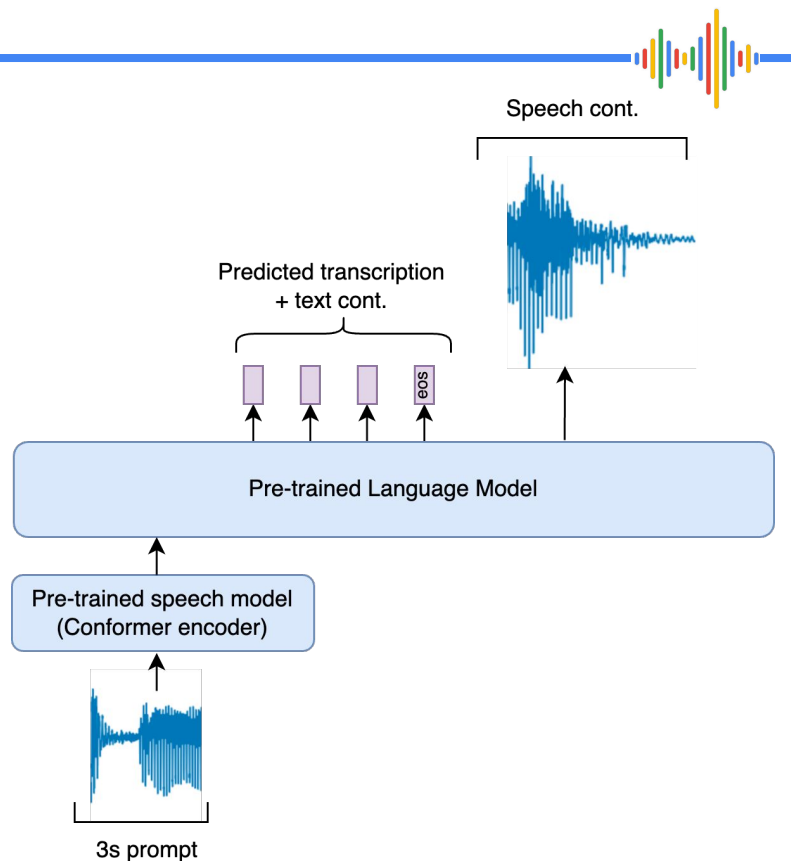
Spoken Question Answering and Speech Continuation Using Spectrogram-Powered LLM

Eliya Nachmani, Alon Levkovitch, Roy Hirsch, Julian Salazar,
Chulayuth Asawaroengchai, Soroosh Mariooryad, Ehud Rivlin, RJ
Skerry-Ryan, Michelle Tadmor Ramanovich

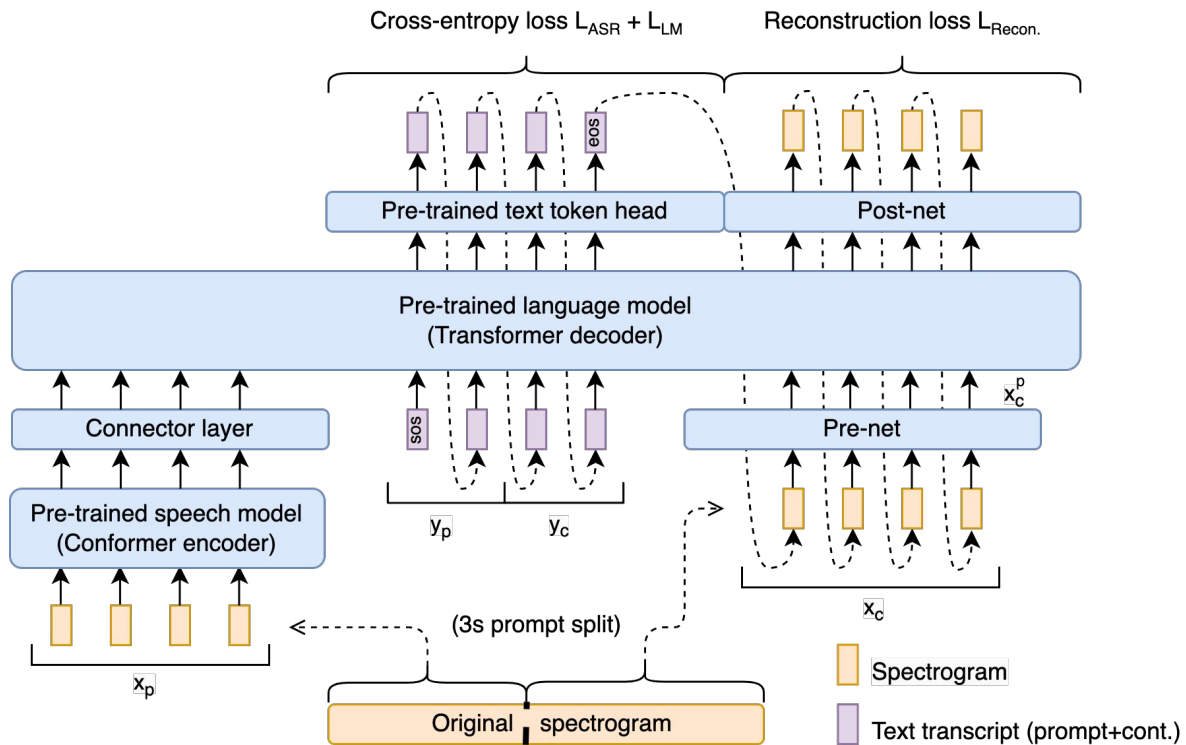


SPECTRON

- ❑ Directly process spectrograms - input and output.
- ❑ Leverage the capabilities of a pre-trained speech encoder.
- ❑ Retain generative natural language abilities from a pre-trained LLM.
- ❑ End-to-end training.
- ❑ End-to-end spoken question answering.



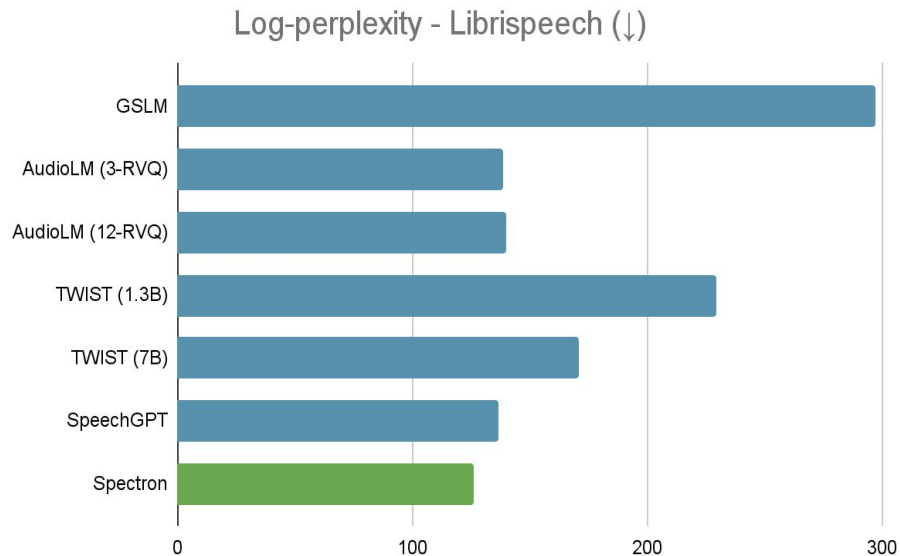
Architecture



Speech Continuation - Semantic Quality





- ❑ Semantic Quality is measured by Log-Perplexity.
- ❑ Log-Perplexity computed using GPT2 on transcribed continuations.



Speech Continuation - Acoustic Quality



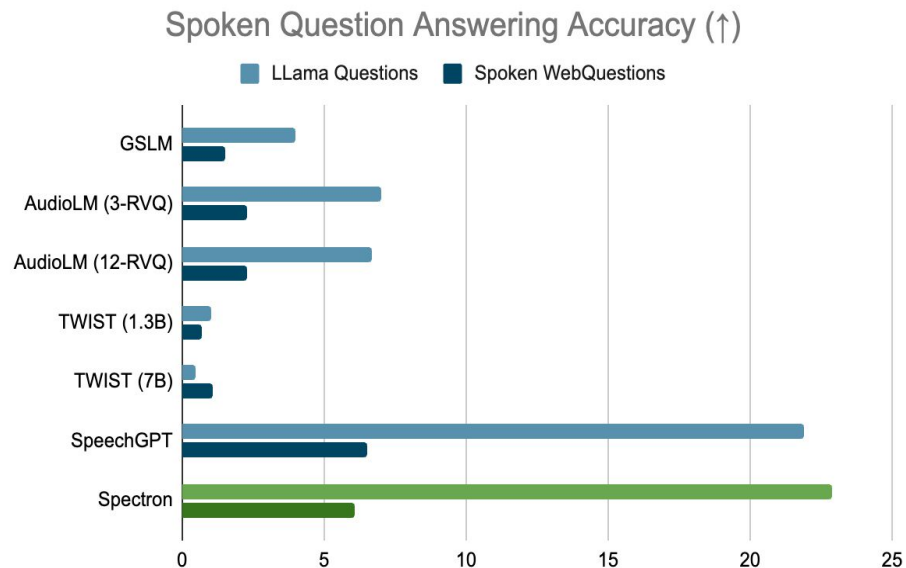
- ❑ Acoustic Naturalness is measured by Mean Opinion Score (MOS).
- ❑ Speaker fidelity is measured by Speaker Similarity.



Spoken Question Answering



- ❑ **Spoken WebQuestions: Spoken questions generated WebQuestions.**
- ❑ **LLama Questions: open-domain QA dataset that we had gathered from LLama-7B model.**



Audio Samples



Prompt	Generated Continuation 