

M³AV: A Multimodal, Multigenre, and Multipurpose Audio-Visual Academic Lecture Dataset

Zhe Chen¹, Heyang Liu¹, Wenyi Yu², Guangzhi Sun³, Hongcheng Liu¹
Ji Wu², Chao Zhang²✉, Yu Wang^{1,4}✉, Yanfeng Wang^{1,4}

¹Cooperative Medianet Innovation Center, Shanghai JiaoTong University

²Department of Electronic Engineering, Tsinghua University

³University of Cambridge Department of Engineering ⁴Shanghai AI Laboratory

{chenzhe2018, liuheyang, hongcheng_liu, yuwangsJTU, wangyanfeng622}@sjtu.edu.cn
wuy22@mails.tsinghua.edu.cn, gs534@cam.ac.uk, {wuji_ee, cz277}@tsinghua.edu.cn

Abstract

Publishing open-source academic video recordings is an emergent and prevalent approach to sharing knowledge online. Such videos carry rich multimodal information including speech, the facial and body movements of the speakers, as well as the texts and pictures in the slides and possibly even the papers. Although multiple academic video datasets have been constructed and released, few of them support both multimodal content recognition and understanding tasks, which is partially due to the lack of high-quality human annotations. In this paper, we propose a novel multimodal, multi-genre, and multipurpose audio-visual academic lecture dataset (M³AV), which has almost 367 hours of videos from five sources covering computer science, mathematics, and medical and biology topics. With high-quality human annotations of the slide text and spoken words, in particular high-valued name entities, the dataset can be used for multiple audio-visual recognition and understanding tasks. Evaluations performed on contextual speech recognition, speech synthesis, and slide and script generation tasks demonstrate that the diversity of M³AV makes it a challenging dataset¹.

1 Introduction

The rapid progress of technology has brought numerous academic presentations and talks available on the web with open access (Lev et al., 2019; Atri et al., 2021; Lee et al., 2022, 2023). These resources are particularly helpful to researchers as they contain rich specialised knowledge in auditory and visual modalities. With the development of neural-based AI systems, it is desired that AI could comprehend and process these multimodal information resources to assist scientists in accelerating their research. Abilities, including transcribing and

generating speech for a presentation with the aid of slides, as well as presentation slide or script generation, are particularly useful for researchers to conduct investigations and prepare presentations.

As a natural data source containing multimodal content and knowledge, academic lectures have been explored to construct a series of datasets. Some datasets focus on the evaluation of the model’s ability to recognize multimodal content (Dutta et al., 2018; Wang et al., 2023), others focus on promoting models to understand multimodal knowledge in academic videos (Lev et al., 2019; Li et al., 2019; Atri et al., 2021; Lee et al., 2022, 2023). However, none of them has yet simultaneously paid attention to the model’s ability to recognize multimodal content and understand rich academic knowledge, which is partly due to their respective lack of adequate manual labelling. The homogeneity of holding two characteristics is critical to the implementation of end-to-end academic recognition and understanding systems.

To fill the gap, we release the Multimodal, Multi-genre, and Multipurpose Audio-Visual Academic Lecture Dataset (M³AV) as shown in Figure 1. The dataset contains selected academic lectures and presentation videos from multiple fields, including computer science, biomedical science, and mathematics. Each video is endowed with highly qualified speech transcriptions and OCR labels for both printed and handwritten characters, including complex mathematical formulae. In addition, academic papers for some videos are also provided to supplement knowledge more extensively.

Together with the dataset, three benchmark tasks are also proposed that reflect the perception and understanding of the multimodal information in those videos. This includes automatic speech recognition (ASR) with a particular focus on contextual ASR (CASR), spontaneous text-to-speech (TTS) synthesis, and slide and script generation (SSG). Representative baseline models are used to bench-

✉: Corresponding author.

¹Project website: <https://jack-zc8.github.io/M3AV-dataset-page>

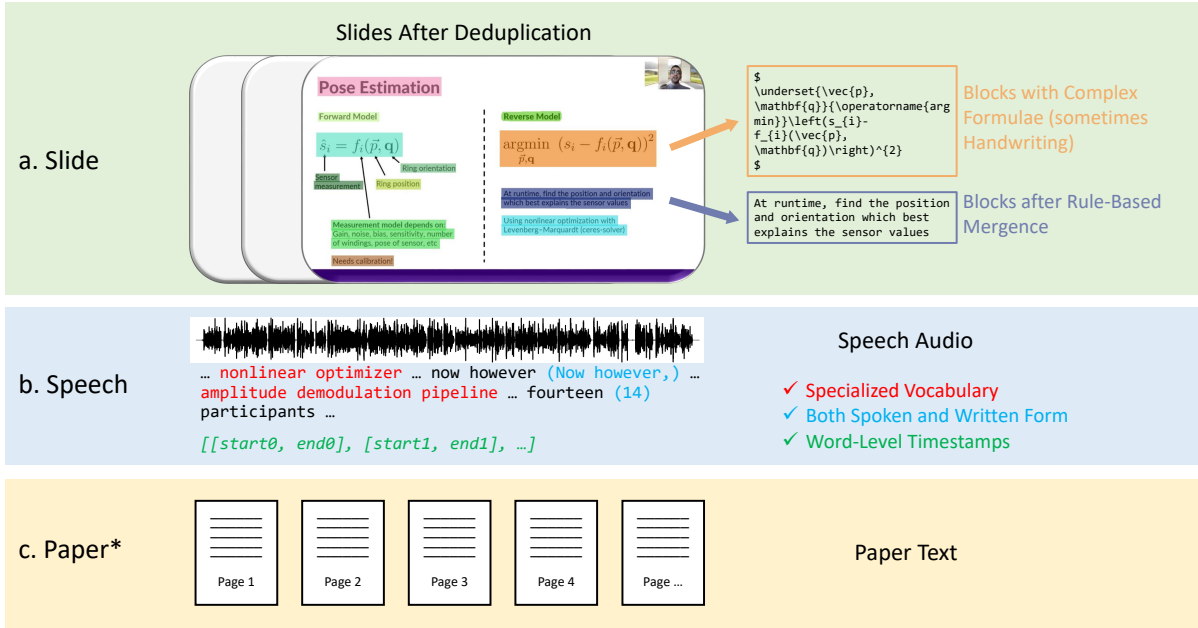


Figure 1: The overview of our M^3AV dataset. The first component is slides annotated with simple and complex blocks. They will be merged following some rules. The second component is speech containing special vocabulary, spoken and written forms, and word-level timestamps. The third component is the paper corresponding to the video. The asterisk (*) denotes that only computer science videos have corresponding papers.

mark the performance for each task. Through our experiments, it is observed that the existing models showcase limited performance in perceiving and understanding the multimodal content in this dataset, and the rich academic knowledge is not utilized effectively.

2 Related Work

The datasets based on academic lectures mainly fall into two categories: measuring the model’s ability to recognize multimodal content and evaluating its capacity to capture academic knowledge. In Table 1, we list the data types and dataset tasks of the relevant academic lecture-based datasets.

Multimodal Content Recognition. LectureVideoDB (Dutta et al., 2018) exploits frames from lecture videos to test the model’s ability to perform handwritten and scene text recognition. SlideSpeech (Wang et al., 2023) enriches the audiovisual corpus with synchronized slides to provide additional textual information. It is worth noting that M^3AV dataset provides both OCR data that is fully manually labelled and speech data that combines multiple ASR system outputs and manual labelling, which is completely qualified to cover tasks in the aforementioned datasets.

Academic Knowledge Understanding. LectureBank (Li et al., 2019) manually annotates the prerequisite relationships between course concepts to carry out prerequisite chain learning. GoogleI/O (Chen et al., 2014) and LaRochelle (Nguyen et al., 2014) both study video-level retrieval using presentation or lecture videos. AVIATE (Atri et al., 2021) adopts multimodal information from academic presentation videos to conduct abstract summarization. Similarly, TalkSumm (Lev et al., 2019) uses a scalable annotation method to construct large-scale paper summaries from corresponding video transcription. LPM (Lee et al., 2022, 2023) introduces crossmodal retrieval and generation tasks around its aligned slides and spoken language. Our M^3AV dataset contains annotated video data for both slides and speeches, which implies little knowledge loss. The addition of paper documentation information also greatly complements the knowledge information. Furthermore, a challenging Slide and Script Generation task is developed to perform the inter-transformation of two kinds of academic materials.

As illustrated in Table 1, our dataset contains the most complete and human-annotated resources of slide, speech, and paper, thus supporting not only the recognition of multimodal content but also the comprehension of high-level academic knowl-

Table 1: Comparison with other academic lecture-based datasets in terms of data types and designed tasks. ‘‘A’’ denotes fully automated processing and ‘‘M’’ denotes fully or partially manual labelling.

	Slide Data					Speech Data	Paper Data	Dataset Tasks	
	Segments	Figures	Text	Handwriting	Formulae	Transcription	Text	Recognition	Understanding
LectureVideoDB (2018)	M	-	M	M	-	-	-	✓	✗
SlideSpeech (2023)	-	-	A	-	-	M	-	✓	✗
GoogleI/O (2014)	A	-	A	-	-	A	-	✗	✓
LaRochelle (2014)	A	-	A	-	-	A	-	✗	✓
LectureBank (2019)	M	-	A	-	-	-	-	✗	✓
TalkSumm (2019)	-	-	-	-	-	A	A	✗	✓
AVIATE (2021)	-	-	A	-	-	A	A	✗	✓
LPM (2022; 2023)	M	M	A	-	-	A	-	✗	✓
M ³ AV (Ours)	M	-	M	M	M	M	A	✓	✓

Table 2: Comparison with other academic lecture-based datasets in terms of data size and availability.

	Size				Available
	# Videos	# Hours	# Slides	# Papers	
LectureVideoDB (2018)	24	-	5474	-	✓
SlideSpeech (2023)	1705	1080	-	-	✓
GoogleI/O (2014)	209	-	-	-	✓
LaRochelle (2014)	47	65	3250	-	✗
LectureBank (2019)	1352	-	51939	-	✓
TalkSumm (2019)	1716	-	-	1716	✓
AVIATE (2021)	8201	~2300	-	8201	✓
LPM (2022; 2023)	334	187	9031	-	✓
M ³ AV (Ours)	1113	367	24956	767	✓

Table 3: Duration of videos from different sources in each data partition.

Sets	CHI	Ubi	NIH	IPP	MLS	Total
Train	44.1	7.9	188.2	34.0	21.1	295.4
Dev	4.4	1.0	23.6	3.6	2.5	35.1
Test	4.7	1.5	23.5	4.1	2.6	36.4

edge. At the same time, the size of our dataset is also relatively rich while accessible as shown in Table 2.

3 Summary of M³AV Dataset

3.1 Metadata

We treat each video as a unit and rename each video in the form of ‘‘{source}-{id}’’. For the raw data, we provide the YouTube URL of each video for users to download.

Speech Annotation. The speech annotations are saved in JSON files. Each unit contains the start and end timestamps, the transcribed text in the spoken and written form, and the corresponding word-level timestamps.

OCR Annotation. In OCR annotations, the coordinates and manually corrected text (formulae are in LaTeX format) of each block are provided in JSON files. Results of merging (detailed in Appendix A.1) are also provided.

Dataset Split. The division of the sets is determined by the speech data, as the duration can measure the amount of information consistently. The video durations are illustrated in Table 3.

3.2 Characteristics of Speech Data

Variety of Academic Fields. As shown in Figure 2(a), M³AV contains academic presentations and talk videos from three different fields including computer science, biomedical science, and mathematics. The Conference on Human Factors in Computing Systems (CHI) 2021 Presentations and ACM International Conference on Ubiquitous Computing (Ubi) 2020 Presentations focus on research in computer science. The National Institutes of Health Director’s Wednesday Afternoon Lectures (NIH) presents weekly science talks given by top biomedical scientists. The Introduction to the Principles and Practice of Clinical Research (IPP) teaches how to conduct clinical research effectively and safely. Oxford Mathematics (MLS) contains mathematics lectures. The total number of videos is 1113, with 366.9 hours of speech.

Abundance of Rare Words. The size of our spoken form word table is 47865, while the word frequency within 1k reaches 47483 words (99.20%) as shown in Figure 2(e). It effectively represents the typical scenario met in academic presentations where new terminologies constantly appear and are crucial to the overall understanding.

Quality of Transcription. M³AV is labelled by combining multiple high-performing ASR systems together with manual labels. As human often fails to label rare terminologies, the assistance of a high-performing ASR is indispensable. The multimodal labelling pipeline is described in Section 4.2.

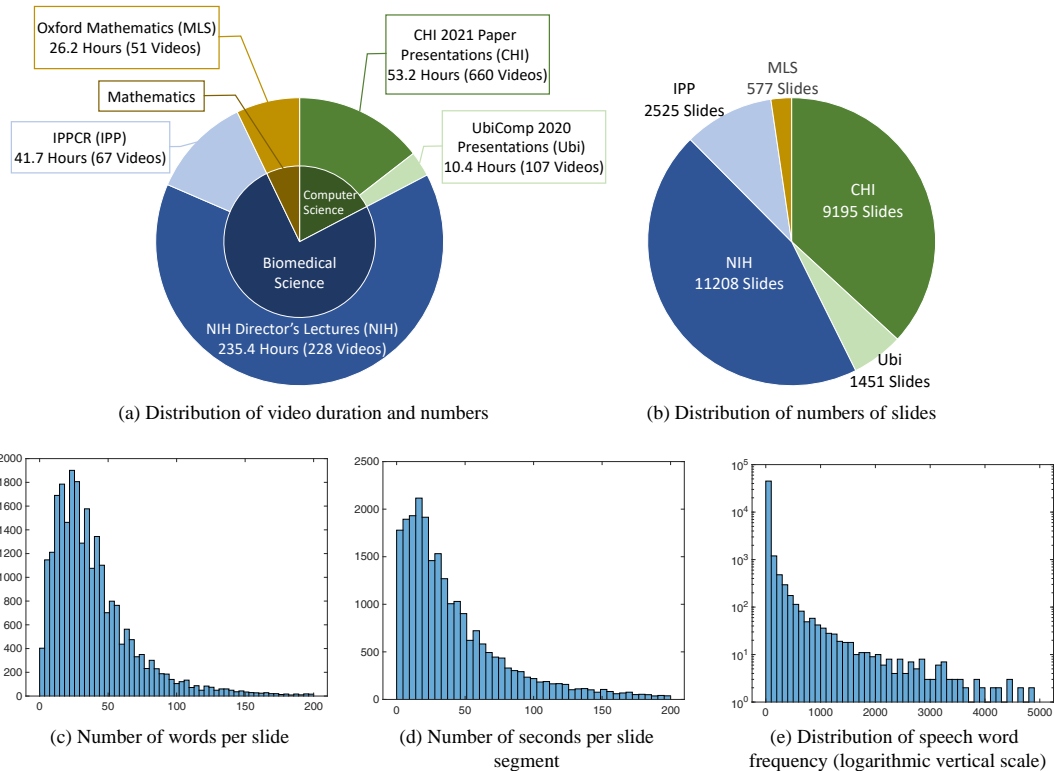


Figure 2: Statistics of our dataset. (a) shows video duration and numbers. (b) shows the number of slides. (c) shows the number of words per slide. (d) shows the duration per slide segment. (e) shows the speech word frequency.

3.3 Characteristics of Visual Data

Sufficient OCR Text Labelling. The total number of slides has reached 24956 as illustrated in Figure 2(b). The average number of words per slide is 40.96 as shown in Figure 2(c). We can observe that there is a sufficient number of slides and the density of the words is high, which indicates the richness of the textual information in our slides. The duration of every slide segment is 50.67 as indicated in Figure 2(d).

Complex Handwriting and Formulae. Significant numbers of pages and blocks in Table 4 show the enrichment of our complex texts. Furthermore, handwriting and formula texts which both contain lots of presentation content and academic knowledge have already been processed into standard text or LaTeX format and thus can be directly utilised by large language models.

4 Data Creation Pipeline

4.1 Data Collection

Videos of open-source conferences and lectures are collected from YouTube, together with descriptions and subtitles. Papers for computer science videos are downloaded and subsequently parsed

Table 4: Statistics of handwriting and formulae. Blocks that appear on the same page as the complex block are also counted into "# Blocks" due to the same processing.

Sources	# Pages	# Blocks*
CHI	58	656
Ubi	63	906
NIH	53	695
IPP	21	224
MLS	577	5736
Total	772	8217

using science-parse².

4.2 Speech Transcription

The diagram of creating speech transcription is shown in Figure 3, comprising four steps.

4.2.1 Step 1: Candidate Generation

The downloaded subtitles are not high-quality as expected by our observation. Therefore, we introduce expert ASR systems following Zhang et al. (2022). In particular, we select Microsoft STT³,

²<https://github.com/allenai/science-parse>

³<https://azure.microsoft.com/en-us/products/ai-services/speech-to-text>

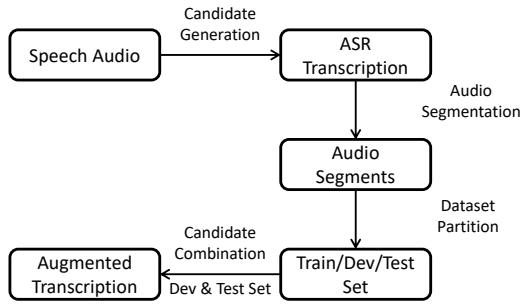


Figure 3: Diagram illustration of the process of creating speech transcription.

and Whisper-large-v2 (Radford et al., 2023)⁴. They both achieve a WER of less than 10% or even 5% on multiple test sets indicated in the SpeechColab ASR Benchmark (EN)⁵ in Oct. 2022.

The Microsoft transcriptions in written and spoken form⁶ are denoted as M_w and M_s respectively, and the Whisper transcriptions result in written form are denoted as W_w (only the written form exists in the Whisper result).

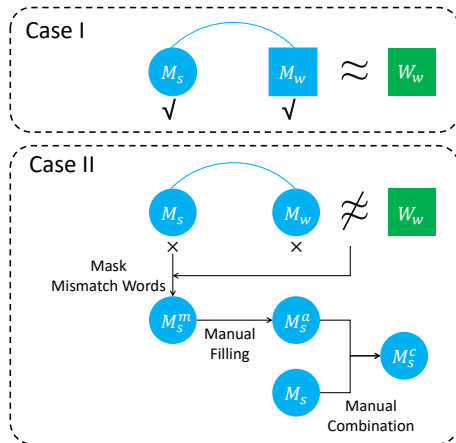


Figure 4: Diagram illustration of the process of candidate combination.

4.2.2 Step 2: Audio Segmentation

Most ASR models expect that the input audio comes in the form of a relatively short duration, usually up to a few seconds in length. Therefore, we conduct re-segmentation using the word-level timestamps in M_w . The split is allowed to be in silence and punctuation. And the duration of each segment is kept to 10 seconds or less. Detailed rules can be found in Appendix A.1.

⁴<https://github.com/openai/whisper>

⁵<https://github.com/SpeechColab/Leaderboard>

⁶“Spoken form” refers to the representation of language as it is actually spoken or pronounced, while “written form” refers to the representation of language in its written or textual format.

4.2.3 Step 3: Dataset Partition

We divide all the videos into training, development, and test sets, with a ratio of roughly 8:1:1 for total duration. We also avoid a speaker appearing in different sets, which is identified by the introductory information and description text.

4.2.4 Step 4: Candidate Combination

For the training set, we directly adopt the Microsoft transcription M_w and M_s , because it shows high quality verified by checking against manual labelling on a subset and has both spoken and written form. For the development and test set, we design the candidate combination process shown in Figure 4.

In Case I where M_w is “approximately equal”⁷ to W_w , M_s and M_w will be directly adopted.

In Case II where M_w is not “approximately equal” to W_w , we design the following steps:

1. We use the misalignment to mask the words in M_s as [???].
2. Annotators are assigned to fill in the masks. The result, M_s^a can be regarded as a type of data augmentation.
3. Annotators are instructed to combine M_s and M_s^a to provide the final corrected results M_s^c .

The postprocessing of M_s^c is detailed in Appendix A.3. The manual labelling process is detailed in Appendix A.4. The comparison with manual labelling on a subset for the training set and the manual processing for the dev/test set ensure the accuracy and consistency of the speech annotation.

4.3 Slide Annotation

The diagram of creating slide annotation is shown in Figure 5, comprising four steps.

4.3.1 Step 1: Extensive Pre-Annotation

We conduct frame sampling for the video at 1 frame per second and use PaddleOCR⁸ for large-scale text detection and recognition. PaddleOCR is selected for its completeness and high accuracy of the recognition results with little sticking between words.

4.3.2 Step 2: Slide Deduplication

Since the frames are repetitive, we first perform automatic deduplication. Specifically, we remove uninformative frames and blocks and filter out frames with too few or too many blocks. Next, we compare

⁷It forgives misalignments due to pauses and common words, as they are mainly short-pronounced, and M_w and M_s have high confidence in them.

⁸<https://github.com/PaddlePaddle/PaddleOCR>

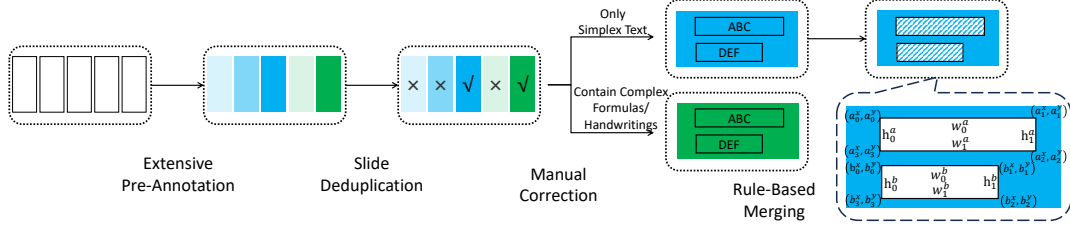


Figure 5: Diagram illustration of the process of slide annotation. Shades of the same colour represent the amount of slide content in the same segment. For example, the dark green page adds some content to the light green page. The right sign (✓) represents reservation, while the wrong sign (×) represents discarding.

the OCR contents of neighbouring frames and only keep the latter one when the content of the former is entirely included. Detailed criteria are shown in Appendix A.2. After completing the automatic deduplication, manual deduplication is performed to retain the last frame in each segment.

4.3.3 Step 3: Manual Correction

For frames with simple text, corrections are directly executed by experienced researchers for the location and content of the OCR results. Note that blocks of presentation-irrelevant text such as URLs and email addresses are omitted. For frames with complex handwriting or formulae, we pre-label them using the Mathpix OCR API⁹. Then, the annotator performs careful corrections for the OCR result due to their complexity. The formulae are labelled in the format of in-line LaTeX. The details of manual labelling are shown in Appendix A.4. To control the quality of labelling, we sampled the results of 9 annotators individually and checked whether 120 of the annotated images had a high number of missing or incorrect items, with no more than 10 wrong items.

4.3.4 Step 4: Rule-Based Merging

To provide comprehensive semantic information, we devise three rules to perform merging¹⁰:

- Similar Height:

$$\max(h_0^a, h_0^b) - \min(h_0^a, h_0^b) \leq a \times \max(h_0^a, h_0^b) \quad (1)$$

- High Overlap in the Horizontal Direction:

$$\min(a_2^x, b_1^x) - \max(a_3^x, b_0^x) \geq b \times \min(b_1^x - b_0^x, a_2^x - a_3^x) \quad (2)$$

⁹<https://mathpix.com/ocr>

¹⁰As demonstrated in the top part of Figure 1, the recognition results of the general OCR model are in rows. It is required to merge these lines into paragraphs to provide complete semantic information.

- Proximity of the Upper and Lower:

$$b_0^y - a_3^y \leq c \times \min(h_0^a, h_0^b) \quad (3)$$

where the notations are shown as Figure 5. We set $a = 0.8$, $b = 0.8$, and $c = 0.6$.

5 Benchmarks and Experiments

5.1 ASR and CASR Task

5.1.1 Task Description

E2E approaches have gained substantial attention for ASR, which enable direct mapping from raw speech audio to its transcription. However, they often struggle with the accurate recognition of rare words that carry important information for understanding. CASR aims to improve the model’s recognition performance for these rare words by incorporating them into the biasing list.

5.1.2 Benchmark Systems

We adopt two widely-used benchmark systems for ASR and CASR, namely Attention-based Encoder-Decoder (AED) (Chan et al., 2016) and RNN-Transducer (RNN-T) (Graves, 2012). Both systems employ a Conformer (Gulati et al., 2020) encoder. The AED system is jointly trained and decoded with Connectionist Temporal Classification (CTC) (Graves et al., 2006; Watanabe et al., 2017).

For ASR, we also introduce a variety of state-of-the-art pretrained ASR models. Note that the results of Whisper-v3 are for reference only, as the Whisper-series model is involved in dataset construction. The performance of the other open-source ASR models can be fairly evaluated, given that our videos are properly transcribed for the first time, avoiding the issue of training data contamination. These pretrained models are selected from the SpeechCollab ASR Benchmark (EN)¹¹ and Hugging Face

¹¹<https://github.com/SpeechCollab/Leaderboard>

Table 5: Evaluation results on ASR and CASR tasks.

Model	Biasing List	Dev			Test		
		CER (↓)	WER (↓)	BWER (↓)	CER (↓)	WER (↓)	BWER (↓)
ASR							
AED	-	4.6	10.0	35.5	4.5	10.0	35.6
RNN-T	-	4.7	10.3	36.8	4.7	10.6	38.0
<i>Whisper-v3</i> _{1.5B}	-	3.8	7.1	19.4	3.8	7.2	18.9
Wav2Vec2+CTC _{0.3B}	-	10.6	28.3	60.1	10.6	28.5	62.4
Conformer+Trans _{0.6B}	-	4.6	10.7	37.8	4.6	10.5	38.2
Conformer+CTC _{0.6B}	-	4.4	11.1	40.6	4.3	11.1	42.2
SeamlessM4T-v2 _{2.3B}	-	12.4	17.1	36.7	12.4	18.6	40.7
CASR							
RNN-T + TCPGen	Random Distractors	4.7	9.7	26.7	4.7	9.7	27.1
RNN-T + TCPGen	OCR + Same Topic Rare Words	4.8	9.8	30.5	4.7	10.0	31.7
RNN-T + TCPGen	OCR	4.5	9.4	22.9	4.6	9.6	25.0

Community¹² which include Whisper-v3 (Radford et al., 2023), Wav2Vec2+CTC (Grosman, 2021), Conformer+Transducer¹³, Conformer+CTC¹⁴ and SeamlessM4T-v2 (Barrault et al., 2023)¹⁵.

For CASR, the Tree-Constrained Pointer Generator (TCPGen) (Sun et al., 2023b) with GNN tree encodings (Sun et al., 2023a) is used which employs a combination of symbolic prefix-tree search and a neural pointer generator for contextual biasing. More details are illustrated in Appendix B.1 and Appendix B.2.

5.1.3 Experimental Results

In addition to the character error rate (CER) and word error rate (WER) for ASR, we use biasing WER (BWER) for CASR. BWER is defined as the total rare word errors divided by the total number of rare words, including insertions of rare words. Table 5 shows the experimental results. The AED and RNN-T systems achieve around 10.0% WER on both the Dev and Test sets. However, such end-to-end models suffer from rare word recognition as reflected by the BWER where a more than two times increase in the error rate is observed comparing BWER to the WER.

When 1000 distractors¹⁶ are randomly selected

¹²https://huggingface.co/models?pipeline_tag=automatic-speech-recognition&sort=trending

¹³https://catalog.ngc.nvidia.com/orgs/nvidia/teams/nemo/models/stt_en_conformer_transducer_xlarge

¹⁴https://catalog.ngc.nvidia.com/orgs/nvidia/teams/nemo/models/stt_en_conformer_ctc_xlarge

¹⁵<https://huggingface.co/facebook/seamless-m4t-v2-large>

¹⁶We select a certain number of significant words, which are the words in the bias list we aim to recognize accurately. When decoding, a particular utterance only includes a few of these words, while the rest will appear as distractors. The use

of distractors follows Le et al.; Sun et al. from all rare words and are added to the biasing list, significant improvements in rare word recognition can be observed. When using OCR information together with rare words from the same topic to form the biasing list of each utterance during decoding, the recognition performance slightly deteriorates. This is because strongly related words are included in the biasing list, which may cause more confusion since words with related meanings often have similar pronunciations and word forms. When only OCR words are used, we achieve the best performance of the CASR task. In conclusion, by using TCPGen utilizing the OCR information, we achieve a relative BWER decrease of 37.8% and 34.2% on dev and test sets respectively.

5.2 Spontaneous TTS Task

5.2.1 Task Description

Spontaneous Text-to-Speech (TTS) (Guo et al., 2021; Yan et al., 2021) focuses on generating speech with a natural and conversational session. Most TTS systems use clean speech. This form makes it easier to extract pronunciation features, while it places higher requirements on the quality of the training audio and often leads to a discernible gap between the synthesized output and authentic speech. We are devoted to using real speech data in M³AV for the TTS system to produce speech that aligns more closely with natural conversational patterns, as it is spontaneous and has regulated pronunciations.

5.2.2 Benchmark Systems

Multi-codebook vector quantized TTS (MQTTS) (Chen et al., 2023) is introduced as our spontaneous

of distractors follows Le et al.; Sun et al.

Table 6: Evaluation results on Spontaneous TTS task. “GT” denotes the ground truth.

Model	Objective		Subjective
	FFE (\downarrow)	MCD (\downarrow)	MOS (\uparrow)
GT	-	-	3.86 ± 1.02
Bark	0.4269	10.57	3.31 ± 1.02
SpeechT5	0.4153	9.43	3.43 ± 1.09
MQTTS	0.3518	9.40	3.67 ± 0.94

TTS baseline. It employs a two-stage training architecture. In the first stage, a quantizer is trained to discretize the raw waveform into a sequence of codes, minimizing quantization and reconstruction losses. A discriminator guides reconstruction via adversarial learning. In the second stage, the fixed quantizer is combined with a Transformer that performs autoregressive generation on the code sequence, conditioned on a speaker embedding for multi-speaker capability. More details can be found in Appendix B.3.

5.2.3 Experimental Results

For the objective metrics, we adopt the F0 Frame Error (FFE) (Chu and Alwan, 2009) and Mel-Cepstral Distortion (MCD) (Kubichek, 1993). For the subjective metrics, the Mean Opinion Score (MOS) test on a scale of 1 to 5 is performed, which is detailed in Appendix A.4. We randomly choose 20 samples and generate speech segments with the MQTTS model and well-known pre-trained models, including Bark¹⁷ and SpeechT5 (Ao et al., 2022)¹⁸. Table 6 shows the experimental results, in which the MQTTS model shows the best performance within all the evaluation metrics. It indicates that the real speech in our dataset can drive AI systems to simulate more natural speech.

5.3 Slide and Script Generation Task

5.3.1 Task Description

We define the task formally as follows. Let I , T_i and T_s be the slide picture, slide text and the speech transcription of every segment, respectively. For the “Slide→Script” task, we need to generate T_s based on I . Note that T_i can be viewed as the intermediate result generated from I by the OCR model. For the “Script→Slide” task, we need to generate T_i based on T_s . The Slide and Script Generation (SSG) Task is designed to promote AI models to understand and reconstruct advanced academic knowledge, thus assisting researchers in

¹⁷<https://huggingface.co/suno/bark>

¹⁸https://huggingface.co/microsoft/speecht5_tts

creating slides and presentations. The selected paper sentences are provided for retrieval-augmented generation (RAG), supplementing rich knowledge. Such models with academic comprehension can facilitate researchers to handle frequently updated academic materials to acquire knowledge and innovation efficiently.

5.3.2 Benchmark Systems

Having witnessed the excellent performance of large language models (LLMs) and large multi-modal models (LMMs) nowadays, we explore their capabilities there. For the “Slide→Script” task, we adopt multimodal cascading pipelines and integrated systems. The cascading pipeline consists of an OCR model to extract the text (we directly adopt the OCR annotations) and a language model to generate the answer. Specifically, the cascading pipelines include fine-tuned LLaMA-2 (Touvron et al., 2023) and GPT-4 (OpenAI, 2023a) using slide OCR text and the integrated systems include fine-tuned InstructBLIP (Dai et al., 2023) and GPT-4V (OpenAI, 2023b,c). For the “Script→Slide” task, we adopt fine-tuned LLaMA-2 and GPT-4. For the selection of related paper sentences, We calculate the similarity between the script text and each paper sentence and filter from them. More details can be found in Appendix B.4.

5.3.3 Experimental Results

ROUGE (Lin, 2004), BERTScore (Zhang et al., 2019), and BARTScore (Yuan et al., 2021) results are reported in Table 7 and Table 8. We also perform manual scoring, where each sample is scored on a scale from 1 (worst) to 5 (best) by 3 annotators on multiple dimensions, including coherence, consistency, fluency and relevance (Fabbri et al., 2021). We show their average values, while the detailed values are in Appendix B.5. Scoring is performed on a subset of size 100, detailed in Appendix A.4.

We get significant observations: 1) The open-source models (LLaMA-2, InstructBLIP) show a limited performance improvement when raised from 7B to 13B. Their performances are far from the closed-source models (GPT-4 and GPT-4V). We believe that high-quality pre-training data, e.g., informative corpus and visual QA data which encapsulate multimodal information, is required to enhance their SSG performances beyond just boosting the model size. 2) The latest LMM (GPT-4V) has already exceeded the cascaded pipeline composed of unimodal expert models. It sug-

Table 7: Evaluation results on SSG tasks. The upper part of "Slide→Script" shows cascading pipelines, while the lower part shows integrated systems.

Model	ROUGE-1 (↑)	ROUGE-2 (↑)	ROUGE-L (↑)	BERTScore (↑)	BARTScore (↑)	Human (↑)
<i>Slide → Script</i>						
LLaMA-2 _{7B} +OCR	24.6	5.6	21.6	0.073	-4.970	3.6
LLaMA-2 _{13B} +OCR	26.9	6.4	23.6	0.102	-4.823	3.9
GPT-4+OCR	26.2	5.8	23.7	0.085	-4.621	4.6
InstructBLIP _{7B}	24.3	2.2	17.9	0.041	-5.504	2.4
InstructBLIP _{13B}	24.7	2.3	18.2	0.046	-5.514	2.5
GPT-4V	26.4	6.2	23.7	0.092	-4.630	4.7
<i>Script → Slide</i>						
LLaMA-2 _{7B}	20.3	6.2	17.3	0.064	-6.689	3.3
LLaMA-2 _{13B}	23.3	7.5	19.8	0.092	-6.484	3.8
GPT-4	26.8	7.7	23.0	0.077	-6.136	4.8

Table 8: Performance improvements of LLaMA-2_{7B} brought by retrieving paper information. "Subset" denotes that only Computer Science videos are contained in all sets for they are the only ones with downloadable papers.

Range	ROUGE-1 (↑)	ROUGE-2 (↑)	ROUGE-L (↑)	BERTScore (↑)	BARTScore (↑)	Human (↑)
<i>Slide → Script</i>						
Subset	24.9	6.4	21.8	0.114	-5.110	3.7
+ Paper	29.2	8.7	25.8	0.161	-4.806	4.1
<i>Script → Slide</i>						
Subset	20.9	6.8	18.1	0.061	-6.737	3.3
+ Paper	22.3	7.9	19.3	0.061	-6.651	3.6

gests that the LMM not only maintains the ability to process textual information but also possesses multi-sensory capabilities, such as the perception and recognition of the slides (Yang et al., 2023). 3) RAG substantially enhances the generation, as shown in Table 8. Qualitative analyses in Appendix C provide a more fine-grained view.

As for future development, we are urged to strengthen the open-source model’s ability to understand high-level knowledge and perceive multimodal information. They can be achieved by increasing the model size and supplying high-quality data. The issue that the open-source multimodal models struggle with recognising text in images can be mitigated by increasing the OCR input. Furthermore, how to effectively bring in external knowledge is an essential topic.

6 Conclusion

We release the Multimodal, Multigenre, and Multipurpose Audio-Visual Dataset with Academic Lectures (M³AV) covering a range of academic fields. This dataset contains manually annotated speech transcriptions, slide text, and additional extracted papers, providing a basis for evaluating AI models for recognizing multimodal content and understanding academic knowledge. We detail the cre-

ation pipeline and conduct various analyses of the dataset. Furthermore, we build benchmarks and conduct experiments around the dataset. We find there is still large room for existing models to improve perceptions and understanding of academic lecture videos.

Limitations

There are some limitations in our work. First, there may exist biases in the dataset, such as the types of presentations and the demographics of the speakers. It is hard to completely avoid biases, although we have incorporated a diverse range of data sources. The users of the dataset should be aware of the impact brought to the AI models by these biases.

Secondly, the data we collected has limitations in terms of academic domains and multilingualism. We will consider adding videos from more research directions, such as humanities disciplines (e.g. economics, law, and sociology) to an extended dataset in future work. And we also plan to collect more data in other languages to prompt multilingual academic research. Nevertheless, we believe our current collection would already be useful for the community to develop some initial approaches.

Thirdly, we have some deficiencies in mining and labelling for visual information of the dataset.

The illustrations in the video frames are not extracted separately like Lee et al. (2022, 2023). And, the talking heads of the speakers are not picked up, which can be utilised to research the generation of virtual presenters. In summary, further exploitation of the visual elements in the slides would be a future extension of the dataset.

Finally, the benchmark systems for the Slide Generation task do not involve the generation of visual pictures, instead, we only get the generated slides in the form of text. The generation of slide images requires complex multimodal models, and the corresponding evaluation metrics require more sophisticated designs.

Ethical Consideration

This work presents a dataset based on academic speeches with manual annotation in terms of slides and speech. The presentations and talks included in the dataset, are open accessed on the YouTube website. The copyright remains with the original owners of the videos. Our dataset is provided under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. The annual annotation is done by the team’s researchers and workers on the Amazon Mechanical Turk. All annotators received clear annotation rules before the annotation and fair payments upon completion. No personal information is requested from any of the annotators.

Acknowledgements

This work is supported by National Key R&D Program of China (No. 2022ZD0162101), National Natural Science Foundation of China (No. 62106140) and STCSM (No. 21511101100, No. 22DZ2229005).

References

Petar Aleksic, Mohammadreza Ghodsi, Assaf Michaely, Cyril Allauzen, Keith Hall, Brian Roark, David Rybach, and Pedro Moreno. 2015. Bringing contextual information to Google speech recognition. In *Proc. Interspeech 2015*, pages 468–472.

Junyi Ao, Rui Wang, Long Zhou, Chengyi Wang, Shuo Ren, Yu Wu, Shujie Liu, Tom Ko, Qing Li, Yu Zhang, et al. 2022. SpeechT5: Unified-modal encoder-decoder pre-training for spoken language processing. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5723–5738.

Yash Kumar Atri, Shraman Pramanick, Vikram Goyal, and Tanmoy Chakraborty. 2021. See, hear, read: Leveraging multimodality with guided attention for abstractive text summarization. *Knowledge-Based Systems*, 227:107152.

Loïc Barrault, Yu-An Chung, Mariano Coria Meglioli, David Dale, Ning Dong, Mark Duppenthaler, Paul-Ambroise Duquenne, Brian Ellis, Hady Elsahar, Justin Haaheim, et al. 2023. Seamless: Multilingual expressive and streaming speech translation. *arXiv preprint arXiv:2312.05187*.

William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. 2016. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4960–4964.

Guoguo Chen, Shuzhou Chai, Guan-Bo Wang, Jiayu Du, Wei-Qiang Zhang, Chao Weng, Dan Su, Daniel Povey, Jan Trmal, Junbo Zhang, Mingjie Jin, Sanjeev Khudanpur, Shinji Watanabe, Shuaijiang Zhao, Wei Zou, Xiangang Li, Xuchen Yao, Yongqing Wang, Zhao You, and Zhiyong Yan. 2021. GigaSpeech: An evolving, multi-domain ASR corpus with 10,000 hours of transcribed audio. In *Proc. Interspeech 2021*, pages 3670–3674.

Huizhong Chen, Matthew Cooper, Dhiraj Joshi, and Bernd Girod. 2014. Multi-modal language models for lecture video retrieval. In *Proceedings of the 22nd ACM International Conference on Multimedia*, page 1081–1084.

Li-Wei Chen, Shinji Watanabe, and Alexander Rudnicky. 2023. A vector quantized approach for text to speech synthesis on real-world spontaneous speech. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11):12644–12652.

Ha-Yeong Choi, Sang-Hoon Lee, and Seong-Wan Lee. 2023. DDDM-VC: decoupled denoising diffusion models with disentangled representation and prior mixup for verified robust voice conversion. *CoRR*, abs/2305.15816.

Wei Chu and Abeer Alwan. 2009. Reducing f0 frame error of f0 tracking algorithms under noisy conditions with an unvoiced/voiced classification frontend. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3969–3972.

Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. InstructBLIP: Towards general-purpose vision-language models with instruction tuning.

Kartik Dutta, Minesh Mathew, Praveen Krishnan, and C.V Jawahar. 2018. Localizing and recognizing text in lecture videos. In *2018 16th International Conference on Frontiers in Handwriting Recognition (ICFHR)*, pages 235–240.

- Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. Summeval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Alex Graves. 2012. Sequence transduction with recurrent neural networks. In *International Conference of Machine Learning 2012 Workshop on Representation Learning*.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376.
- Jonatas Grosman. 2021. Fine-tuned XLSR-53 large model for speech recognition in English. <https://huggingface.co/jonatasgrosman/wav2vec2-large-xlsr-53-english>.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. Conformer: Convolution-augmented transformer for speech recognition. In *Proc. Interspeech 2020*, pages 5036–5040.
- Haohan Guo, Shaofei Zhang, Frank K. Soong, Lei He, and Lei Xie. 2021. Conversational end-to-end TTS for voice agents. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 403–409.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis. In *Advances in Neural Information Processing Systems*, volume 33, pages 17022–17033.
- R. Kubichek. 1993. Mel-cepstral distance measure for objective speech quality assessment. In *Proceedings of IEEE Pacific Rim Conference on Communications Computers and Signal Processing*, volume 1, pages 125–128 vol.1.
- Taku Kudo and John Richardson. 2018. SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71.
- Duc Le, Mahaveer Jain, Gil Keren, Suyoun Kim, Yangyang Shi, Jay Mahadeokar, Julian Chan, Yuan Shangguan, Christian Fuegen, Ozlem Kalinli, Yatharth Saraf, and Michael L. Seltzer. 2021a. Contextualized streaming end-to-end speech recognition with trie-based deep biasing and shallow fusion. In *Proc. Interspeech 2021*, pages 1772–1776.
- Duc Le, Gil Keren, Julian Chan, Jay Mahadeokar, Christian Fuegen, and Michael L. Seltzer. 2021b. Deep shallow fusion for RNN-T personalization. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 251–257.
- Dong Won Lee, Chaitanya Ahuja, Paul Pu Liang, Sanika Natu, and Louis-Philippe Morency. 2022. Multimodal lecture presentations dataset: Understanding multimodality in educational slides. *arXiv preprint arXiv:2208.08080*.
- Dong Won Lee, Chaitanya Ahuja, Paul Pu Liang, Sanika Natu, and Louis-Philippe Morency. 2023. Lecture presentations multimodal dataset: Towards understanding multimodality in educational videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20087–20098.
- Guy Lev, Michal Shmueli-Scheuer, Jonathan Herzig, Achiya Jerbi, and David Konopnicki. 2019. TalkSum: A dataset and scalable annotation method for scientific paper summarization based on conference talks. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2125–2131. Association for Computational Linguistics.
- Irene Li, Alexander R Fabbri, Robert R Tung, and Dragomir R Radev. 2019. What should I learn first: Introducing LectureBank for NLP education and prerequisite chain learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6674–6681.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*.
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. Montreal Forced Aligner: Trainable text-speech alignment using Kaldi. In *Proc. Interspeech 2017*, pages 498–502.

- Nhu-Van Nguyen, Mickaël Coustaty, and Jean-Marc Ogier. 2014. Multi-modal and cross-modal for lecture videos retrieval. In *22nd International Conference on Pattern Recognition, ICPR 2014, Stockholm, Sweden, August 24-28, 2014*, pages 2667–2672. IEEE Computer Society.
- OpenAI. 2023a. [GPT-4 technical report](#).
- OpenAI. 2023b. [GPT-4V\(ision\) system card](#).
- OpenAI. 2023c. [GPT-4V\(ision\) technical work and authors](#).
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: An ASR corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le. 2019. SpecAugment: A simple data augmentation method for automatic speech recognition. In *Proc. Interspeech 2019*, pages 2613–2617.
- Golan Pundak, Tara N. Sainath, Rohit Prabhavalkar, Anjali Kannan, and Ding Zhao. 2018. Deep context: End-to-end contextual speech recognition. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 418–425.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. [Robust speech recognition via large-scale weak supervision](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 28492–28518. PMLR.
- Nils Reimers and Iryna Gurevych. 2019. SentenceBERT: Sentence embeddings using siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *CoRR*, abs/1910.01108.
- Guangzhi Sun, Chao Zhang, and Phil Woodland. 2022. Tree-constrained pointer generator with graph neural network encodings for contextual speech recognition. In *Proc. Interspeech 2022*, pages 2043–2047.
- Guangzhi Sun, Chao Zhang, and Philip C. Woodland. 2023a. Graph neural networks for contextual ASR with the tree-constrained pointer generator. *arXiv preprint arXiv:2305.18824*.
- Guangzhi Sun, Chao Zhang, and Philip C. Woodland. 2023b. Minimising biasing word errors for contextual ASR with the tree-constrained pointer generator. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:345–354.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Haoxu Wang, Fan Yu, Xian Shi, Yuezhong Wang, and Shiliang Zhang. 2023. SlideSpeech: A large-scale slide-enriched audio-visual corpus. *arXiv preprint arXiv:2309.05396*.
- Shinji Watanabe, Takaaki Hori, Shigeaki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai. 2018. ESPnet: End-to-end speech processing toolkit. In *Proc. Interspeech 2018*, pages 2207–2211.
- Shinji Watanabe, Takaaki Hori, Suyoun Kim, John R. Hershey, and Tomoki Hayashi. 2017. Hybrid ctc/attention architecture for end-to-end speech recognition. *IEEE Journal of Selected Topics in Signal Processing*, 11(8):1240–1253.
- Yuzi Yan, Xu Tan, Bohan Li, Guangyan Zhang, Tao Qin, Sheng Zhao, Yuan Shen, Wei-Qiang Zhang, and Tie-Yan Liu. 2021. Adaptive text to speech for spontaneous style. In *Proc. Interspeech 2021*, pages 4668–4672.
- Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. 2023. [The dawn of LMMs: Preliminary explorations with GPT-4V\(ision\)](#). *CoRR*, abs/2309.17421.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. BARTScore: Evaluating generated text as text generation. *Advances in Neural Information Processing Systems*, 34:27263–27277.
- Binbin Zhang, Hang Lv, Pengcheng Guo, Qijie Shao, Chao Yang, Lei Xie, Xin Xu, Hui Bu, Xiaoyu Chen, Chenchen Zeng, Di Wu, and Zhendong Peng. 2022. WENETSPEECH: A 10000+ hours multi-domain mandarin corpus for speech recognition. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6182–6186.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. BERTScore: Evaluating text generation with bert. In *International Conference on Learning Representations*.
- Ding Zhao, Tara N. Sainath, David Rybach, Pat Rondon, Deepti Bhatia, Bo Li, and Ruoming Pang. 2019. Shallow-fusion end-to-end contextual biasing. In *Proc. Interspeech 2019*, pages 1418–1422.

A Creation Pipeline Details

A.1 Audio Segmentation Rules

We establish the following rules of segmentation:

1. Split is allowed when the silence exceeds 0.2 seconds or there is sentence break punctuation (".", "!", "?").
2. If the accumulated "small segments" obtained are greater than or equal to 8 seconds or the accumulation is greater than or equal to 10 seconds after the next "small segment" is added, then the accumulated segments will form a new segment.
3. If the silence exceeds 5 seconds, then this silence is discarded and a new segment is formed from the accumulated "small segments".

The first rule aims to preserve boundary words as well as the semantics of sentences inspired by LibriSpeech (Panayotov et al., 2015) and GigaSpeech (Chen et al., 2021). The second rule controls the duration of each segment, while the third rule avoids too long silence within a segment.

A.2 Auto Slide Deduplication Details

1. For the frames that contain some OCR block representing the video cover B_{cover} , we delete these images directly. And we remove fixed appearing logo blocks B_{excluded} , such as "NIH".
2. Then, we set a threshold \min_{alpha}^B to filter blocks with too few letters. For the filtering process for images, we set the thresholds \min_B^I and \max_B^I to filter images with too few blocks and too many blocks, respectively.
3. After the above filtering, we merge all the blocks of each image into long text by their coordinates.
4. Initially, we assign the first image to segment 0. Next, we will iterate through each image chronologically: if the error rates (ER) between the current image (segment i) content and all past image contents are all greater than the threshold \min_{ER} , then it is considered as a new image and assigned to segment $i + 1$, otherwise, it is not and kept as segment i .
5. Eventually, the last image in each segment will be the result of the deduplication since they are not repeating each other and contain the most text of the images in their segment.

We would like to emphasize that the thresholds mentioned above depend on the characteristics of each source of videos, such as the clarity of the picture, the density of the text, and the presence of handwriting and formulae. Moreover, we adopt a modified Word Error Rate (WER) or Char Error Rate (CER) in the calculation of ER mainly considering the incremental content of slides. They are formulated as follows:

$$\text{WER}^{\text{mod}} = \frac{S + D + 0 \times I}{H + S + D} \quad (4)$$

$$\text{CER}^{\text{mod}} = \frac{S + D + 0.1 \times I}{H + S + D} \quad (5)$$

where S , D , I , and H represent the number of substitutions, deletions, insertions, and hits, respectively. The choice of WER or CER also depends on the aforementioned characteristics of the videos.

A.3 Postprocessing for Manual Combination Results

To preserve the consistency of the data after obtaining M_s^c shown in Figure 4, we perform inverse text normalization¹⁹, add punctuations, and capitalize words using tools²⁰. Finally, we obtain the word-level timestamps of M_s^c using Montreal Forced Aligner (McAuliffe et al., 2017)²¹.

A.4 Manual Annotation Details

The master-granted annotators are assigned to complete manual filling and combination in creating speech transcription. In these two tasks, each worker receives 0.06 USD for each finished sample, respectively. The annotation interfaces shown in Figure 12 and Figure 13 describe the rules as well as the rejection cases of manual filling and combination in detail, respectively. The failed samples are passed to other workers for annotation.

The experienced researchers in our team are assigned to annotate the OCR data, and the annotation of each sample includes slide deduplication and OCR content correction. Each worker received the appropriate remuneration. The interface for OCR content correction is shown in Figure 14,

¹⁹https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/main/nlp/text_normalization/wfst/wfst_text_normalization.html#inverse-text-normalization

²⁰https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/main/nlp/punctuation_and_capitalization.html

²¹<https://github.com/MontrealCorpusTools/Montreal-Forced-Aligner>

where PPOCRLabel²² is adopted. To control the quality of labelling, we sampled the results of 9 annotators individually and checked whether 120 of the annotated images had a high number of missing or incorrect items, with no more than 10 wrong items.

The master-granted annotators are assigned to complete the MOS scoring in the TTS task evaluation. A batch of 6 TTS results from the same audio clip (with the addition of a trap sample used to detect attentiveness). Each worker receives 0.2 USD for each completed batch respectively. Similarly, the annotation interface shown in Figure 15 details the rules as well as the rejection cases, which is inspired by Choi et al. (2023). The failed samples will be passed to other workers for annotation.

As for the manual scoring of SSG tasks, we posted the annotation task on Amazon Mechanical Turk initially but found that the quality was unsatisfactory due to the complexity of the scripts and speech texts. Therefore, three experienced researchers were assigned to carry out the annotation task. Specifically, annotators need to rate the generated slides/scripts on multiple dimensions, including coherence, consistency, fluency and relevance, as referenced in SummEval (Fabbri et al., 2021). Each sample is scored on a scale from 1 (worst) to 5 (best) by 3 annotators. Note that we sampled 100 items due to cost and multiple models. The annotation interfaces are shown in Figure 16 and Figure 17.

B Experiment Details

B.1 ASR Benchmark Details

We construct the Attention-based Encoder-Decoder (AED) with the help of the well-known ESP-net (Watanabe et al., 2018) toolkit. We use Conformer (Gulati et al., 2020) as the encoder, which consists of 12 blocks, with hidden linear units $d^{\text{ff}} = 2048$ and attention output size $d^{\text{att}} = 512$. Attention head H is set to 8 and the front CNN kernel size is 31. As for the decoder, we use the Transformer with 6 blocks ($d^{\text{ff}} = 2048, H = 8$). We use BPE as our recognition unit, which consists of 1k tokens obtained by SentencePiece (Kudo and Richardson, 2018). In addition to the attention training objectives, We use a certain degree of CTC loss function. CTC weight λ is 0.3 during training, while 0.2 for inferring. SpecAugment (Park et al.,

2019) is applied with time mask width $T = 40$ and frequency mask with $F = 30$. Our model is trained on four 24GB 3090 RTX GPUs for 60 epochs. The top 10 checkpoints after one epoch training are preserved for model averaging.

For the RNN-Transducer (RNN-T), we also adopt Conformer as the encoder, with an intermediate joint network of 320 dimensions. The Conformer has 15 blocks, 1024 hidden linear units, 4 attention heads, and 256 attention output dimensions. After that, a single layer of LSTM is performed.

B.2 CASR Benchmark Details

Our CASR task is performed under two paradigms. One paradigm is to form the biasing list for an utterance using target biasing words together with distractors that are randomly selected from all rare words, as a simulation of the real-world scenario (Aleksic et al., 2015; Zhao et al., 2019; Pundak et al., 2018; Le et al., 2021b). Our experiment uses 1000 distractors to form the biasing list for each utterance, and rare words are defined by removing the top 5k most frequent words from the vocabulary. Another paradigm is audio-visual CASR (Sun et al., 2022) which extracts the biasing list from visual information, such as OCR annotation.

As for TCPGen, we use a 6-layer GCN (Kipf and Welling, 2017) for future information representation. Following the method proposed in (Le et al., 2021a), the biasing lists are extracted by identifying words in the speech transcription of each utterance that also appears in the total rare words list, followed by the addition of plenty of distractors to simulate a real-life setting. The length of the biasing list is set to 500 during training and 1000 during inferencing. Model training lasts for 70 epochs, and the best 10 of the checkpoints are saved for model averaging. Our models are trained on 24GB 3090 RTX GPUs.

Before model training, the rare word selection is first performed. Our speech corpus includes approximately 43k different words in total. We conduct a word frequency analysis on the training set, identifying words with a frequency ranking exceeding 5k as members of the total rare word set. The number of our rare word list is about 38k.

B.3 Spontaneous TTS Benchmark Details

We followed the settings recommended by the architecture authors, and their experiments were trained on a subset of GigaSpeech (Chen et al., 2021), which is a speech corpus containing tran-

²²<https://github.com/PaddlePaddle/PaddleOCR/tree/release/2.7/PPOCRLabel>

scribed audio from audiobooks, Podcasts, and YouTube. We use the training set of the M³AV for MQTTS model training and arbitrarily extract a sufficient number of samples from the Dev and Test Set for speech synthesis to verify the effectiveness of our model. We adopt HiFi-GAN (Kong et al., 2020) as the backbone architecture for the quantizer and discriminator, which we trained for 600k steps. After that, a Transformer is trained to autoregressively generate the code representation based on the past code sequence. The last checkpoints from both steps are preserved. It is worth noting that MQTTS training may take a long time. Our models are trained on 24GB 3090 RTX GPUs.

B.4 Slide and Script Generation Benchmark Details

For the selection of related paper sentences, we adopt the DistilBERT (Sanh et al., 2019) model in the repository of sentence-transformers (Reimers and Gurevych, 2019)²³ to calculate the similarity between the script text and each paper sentence. Then we select the top 3 candidates with similarity greater than 0.5 as the selection results. In the training phase of benchmark systems, we use LoRA (Hu et al., 2022) to fine-tune LLaMA-2 (Touvron et al., 2023). The optimizer we use is AdamW (Loshchilov and Hutter, 2018). The learning rate curve is 1 epoch of warm-up and then drops with a peak learning rate of 2×10^{-4} . The total number of training epochs is 3. Our models are trained on 24GB 3090 GPUs and 48GB A6000 GPUs. Fine-tuned InstructBLIP (Dai et al., 2023) is another baseline. Specifically, we fine-tune the officially released checkpoint for 20 epochs with the learning rate increasing linearly from 1×10^{-7} to 5×10^{-5} during the first 1000 steps and dropping to 0 as a cosine function. The model is trained on 80GB A100 GPUs with AdamW (Loshchilov and Hutter, 2018) as the optimizer. The best checkpoints are chosen according to the performance of the development set. As for the generation settings, we use nucleus sampling (Holtzman et al., 2020) with a top_p of 0.9. The maximum number of tokens generated is 300. These generation settings are consistent across LLaMA-2, InstructBLIP, GPT-4 and GPT-4V.

²³<https://huggingface.co/sentence-transformers/msmarco-distilbert-base-v4>

Table 9: Detailed manual scoring results on SSG tasks.

Model	Coh. (↑)	Cons. (↑)	Flu. (↑)	Rel. (↑)	Avg. (↑)
<i>Slide → Script</i>					
LLaMA-2 _{7B} +OCR	3.2	3.0	4.2	4.1	3.6
LLaMA-2 _{13B} +OCR	3.7	3.4	4.1	4.3	3.9
GPT-4+OCR	4.0	4.4	5.0	5.0	4.6
InstructBLIP _{7B}	1.6	2.0	2.6	3.2	2.4
InstructBLIP _{13B}	2.0	2.1	2.6	3.3	2.5
GPT-4V	4.2	4.5	5.0	5.0	4.7
<i>Script → Slide</i>					
LLaMA-2 _{7B}	3.1	3.4	3.4	3.3	3.3
LLaMA-2 _{13B}	3.5	3.6	3.9	4.0	3.8
GPT-4	4.6	4.4	5.0	5.0	4.8

Table 10: Detailed manual scoring results on SSG tasks when paper information is introduced.

Range	Coh. (↑)	Cons. (↑)	Flu. (↑)	Rel. (↑)	Avg. (↑)
<i>Slide → Script</i>					
Subset	3.4	3.2	4.2	4.1	3.7
+ Paper	4.1	4.2	4.1	4.1	4.1
<i>Script → Slide</i>					
Subset	3.2	3.4	3.3	3.3	3.3
+ Paper	3.6	3.7	3.5	3.4	3.6

B.5 Detailed Manual Scoring Results on SSG Tasks

We show the detailed scoring results in Table 9 and Table 10, which include the coherence (Coh.), consistency (Cons.), fluency (Flu.), relevance (Rel.) and average value (Avg.) between the hypothesis and its reference text.

C Qualitative Analyses of SSG Benchmark Systems

The comparisons of the benchmark systems in “Slide→Script” are shown in Figure 6 and Figure 7. LLaMA-2 and InstructBLIP below refer to models of size 7B. We can find that the scripts generated by GPT-4+OCR and GPT-4V are highly qualified and roughly match the references. The results of LLaMA-2+OCR are acceptable, although slightly weaker than the previous two. All three can develop detailed descriptions around topics in slides, such as the “NIH Outreach Toolkit” in Figure 6 and “Aging and Migration” in Figure 7. InstructBLIP, on the other hand, fails to find the right theme, as shown in Figure 6 and Figure 7, suggesting that today’s open-source LMM models struggle in handling text in slides.

The comparisons of the benchmark systems in “Script→Slide” are shown in Figure 9 and Figure 10. The slides generated by GPT-4 are closer to references than those generated by LLaMA-2, and the contents are more accurate and condensed. For example, GPT-4 presents “Drosophila ToLL” in Figure 9 and “Pipeline” in Figure 10, while

LLaMA-2 does not. We can conclude that GPT-4 has a better understanding of complex academic knowledge than LLaMA-2.

Figure 8 and Figure 11 show the improvement in the quality of the generated scripts and slides brought by external knowledge. The addition of external knowledge enables AI models to produce necessary details, such as “WiFi CSI” in Figure 8 and “13 long-term users” in Figure 11. It shows that the introduction of external knowledge leads to more comprehensive information captured by AI models.

Slide

Prompt There is a single slide and the speaker's speech within a video clip. The clip is a part of the whole speech video. Please act like a speaker and generate the corresponding speech text based on the text (picture) in the given single slide.

Reference For more information on inclusion, the NIH Office of Research on Women's Health has created the Outreach Inclusion Toolkit to help principal investigators and their teams fulfill their responsibilities related to inclusion of women in clinical research. Supported by NIH, we provide information on best practices, federal laws, regulations, and NIH policies, as well as a variety of case studies featuring researchers' actual experiences with including women. In their studies, they share lessons learned.

LLaMA-2+ OCR doctor. It's an amazing job. The NIH puts out a really nice set of documents that they call the NIH outreach toolkit. And you can find those and they provide a lot of, and I mean a lot of help in terms of helping you with how to engage, recruit and retain women in clinical research. And it includes several things that we just talked about in terms of

GPT-4 + OCR Hello everyone, Today we will be delving into the importance of engaging, recruiting, and retaining women in clinical research and the role of the NIH Outreach Toolkit in this process. The NIH, or National Institutes of Health, has developed an Outreach Toolkit which serves as a guide for principal investigators and their research teams. Its primary purpose is to aid in fulfilling their responsibilities to ensure the inclusion of women in their clinical studies. The inclusion of women in clinical research isn't just about being fair or balanced. It's about...

InstructBLIP research so the nih outreach toolkit is a resource that we developed to help engage retire women and retain racial ethnic minority woman in clinical trials, particularly for studies looking at gender specific diseases like breast cancer it includes 5 key steps

GPT-4V Welcome, everyone. Today, we're going to talk about an essential resource provided by the NIH, the NIH Outreach Toolkit, which is specifically designed to assist principal investigators and their research teams in engaging, recruiting, and retaining women in clinical research. As we delve into this toolkit, it's important to understand that it's not just a collection of suggestions; it's a comprehensive guide built on a strong foundation of best practices and research. The toolkit offers a review of research on recruitment and outreach best practices...

Figure 6: “Slide→Script” Case I where scripts generated by benchmark systems are compared.

Slide

Prompt There is a single slide and the speaker's speech within a video clip. The clip is a part of the whole speech video. Please act like a speaker and generate the corresponding speech text based on the text (picture) in the given single slide.

Reference So this suggesting that suggested then that quiescent neural stem cells may become that are normally very stuck lining the ventricle. They may become abnormal migratory with age and that might we don't know, but we it's very speculative, but we speculate that this might actually be problematic for their normal function. Maybe you know those cells are considered the mother stem cells for glioblastoma, so the fact that they are more migratory could also like accentuate the bad features of the glioblastoma. Whereas in contrast the neuro stem cells, the activated...

LLaMA-2+ OCR its own genetic defects in order to provide guidance to these very late dividing neuroblasts? And then, you know, now I've shown you some bad things happening in neuroblastoma, but there are lots of bad things happening in normal aging. Aging has a deleterious effect on cell migration and therefore on tissue organization. And I really believe that that that that the aging of the neural stem cell niche is partly because of deleterious changes in stem cell migration. So we are now collaborating with Dr. Murry Panoragam, who has shown that aging and injury...

GPT-4 + OCR In this part of our discussion, we are going to explore a fascinating area, which is the correlation between aging and migration within the neurogenic niche. Before we proceed, it's essential to understand the significance of the neurogenic niche, particularly the subventricular zone, located along the lateral ventricles in our brains. This area is where neural stem cells reside, contributing to the continuous generation of new neurons and glial cells throughout our lives. In instances of stroke or traumatic brain injury, the functioning and behavior of these neural...

InstructBLIP cells and it's thought that there are a bunch of other things we know from human data for example in the aging brain you can see all sorts or regions where ra is localized so one would assume as i just mentioned this means they may not be accessible to

GPT-4V Welcome everyone, and thank you for joining today's session. Our focus in this segment is on the intriguing topic of 'Aging and migration in neurogenic niches.' As we delve into this subject, we'll be exploring the relationship between aging and the behavior of neural stem cells within the brain. Let's take a look at our slide here, where we have a simplified representation of the brain. You can see the lateral ventricle and the subventricular zone highlighted. This area, known as the neural stem cell niche, plays a critical role in the production and migration...

Figure 7: “Slide→Script” Case II where scripts generated by benchmark systems are compared.

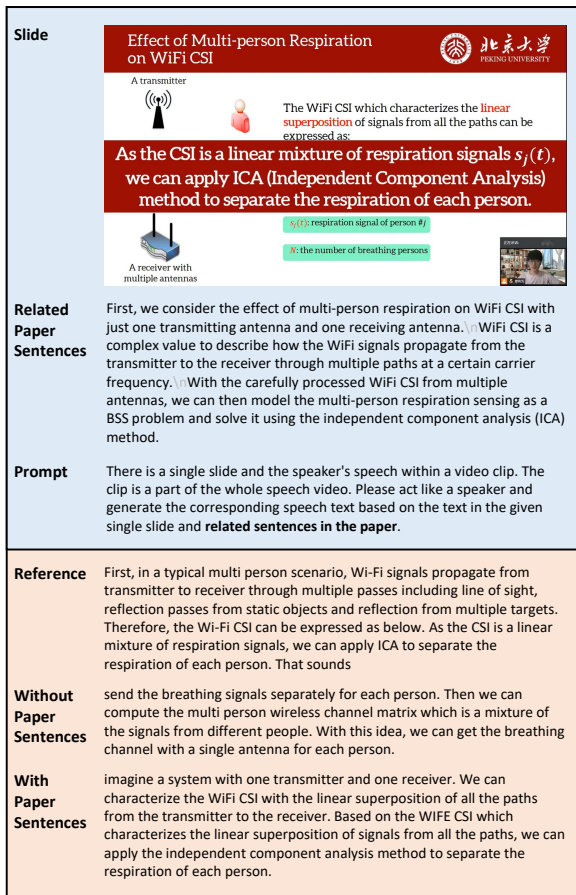


Figure 8: "Slide to Script" Case III where paper sentences are provided as external knowledge.

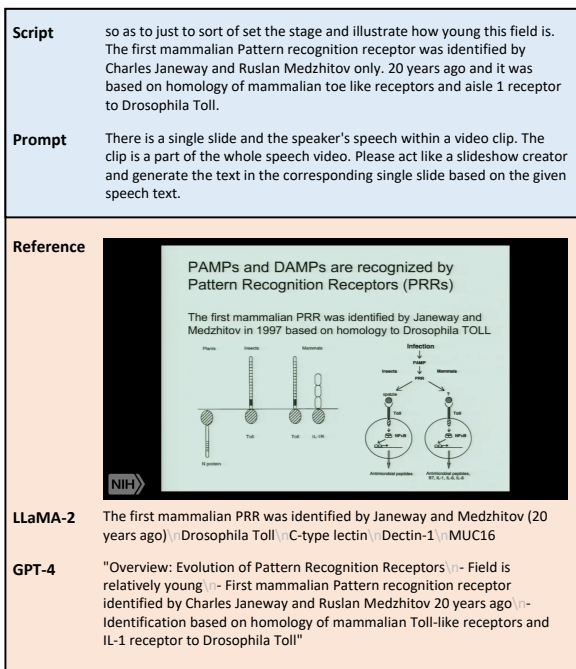


Figure 9: "Script to Slide" Case I where slides generated by benchmark systems are compared.

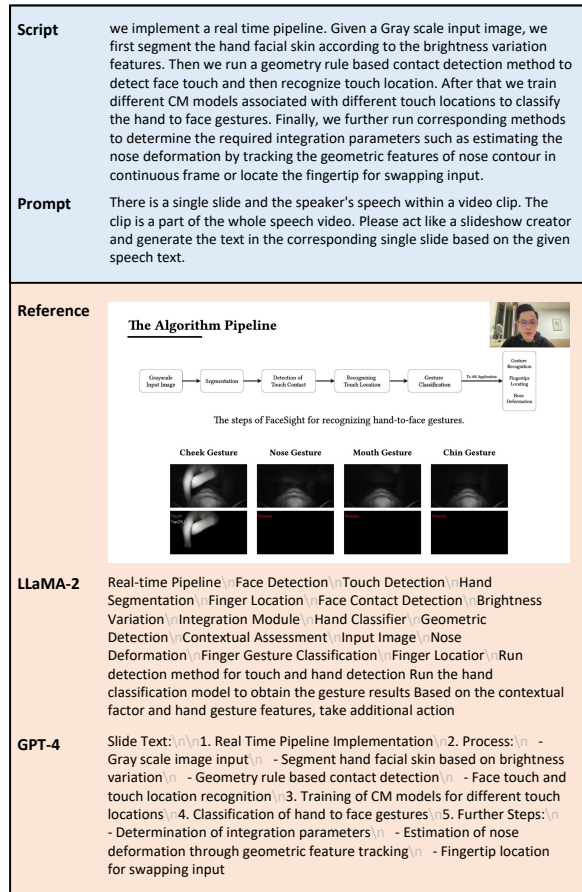


Figure 10: "Script to Slide" Case II where slides generated by benchmark systems are compared.

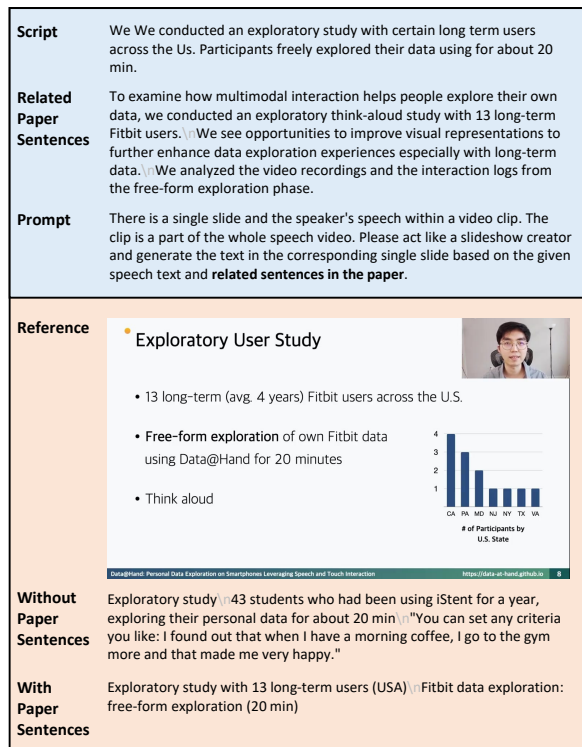


Figure 11: "Script to Slide" Case III where paper sentences are provided as external knowledge.

Transcription rules

1. Complete Missing Content: In this annotation task, you will be provided with a video and a **Reference** text with the missing content marked as [???]. Your task is to **COPY** the original sentence and then fill in the missing content represented by [???]. The missing content [???] may consist of a single word, multiple words, or even be empty.
2. Pay Special Attention to Transcribing **RARE WORDS** in the Missing Content: When transcribing the audio, please ensure that you watch the video and not just rely on listening to the speech. The video may contain specific names and professional vocabulary that may be difficult to recognize solely by auditory means. So pay special attention to **RARE WORDS** as their accuracy is crucial for the acceptance of the annotation.
3. Spoken Form Annotation: All annotation results should be in **SPOKEN FORM** and NOT in written form. For example, write 'two' instead of '2', 'one hundred' instead of '100' and 'dollar' instead of '\$'.
4. Avoid Punctuation and Capitalization: Do not include any punctuation marks or capitalize any letters when completing the missing content.
5. Rejection Criteria:
 - Divergence from Estimated Results: If your annotation significantly differs from our estimated results, or if [???] has content, but the annotator directly deletes [???] without providing the missing content, it will be rejected (If there is no content for a certain [???] indeed, please directly delete the [???] at that location).
 - Partial Response: Please provide the entire sentence, and not provide the missing part in [???] without copying the rest of the sentence.

Reference:

we first investigate [???] sourcing needs and practices when sourcing [???] and discuss two ugc sourcing approaches deep reporting and [???] reporting

Input the whole sentence instead of only the missing content...

Submit

Takeaways

We investigate journalists' sourcing needs and practices as they source UGC content

- Two UGC sourcing approaches
 - *deep reporting* and *wide reporting*
- What journalists look for in UGCs:
 - Personal experiences and expertise
 - Community responses and trends
 - Questions
 - Different viewpoints and opinions

0:07 / 0:09

Figure 12: Annotation interface for manual filling in constructing speech transcription.

Transcription rules

1. Combine two different transcriptions: In this annotation task, you will be provided with a video and two transcriptions: **Reference1** and **Reference2**. Your task is to combine the two reference texts to provide the best transcription results.
2. Pay Special Attention to Transcribing **RARE WORDS**: When transcribing the audio, please ensure that you watch the video and not just rely on listening to the speech. The video may contain specific names and professional vocabulary that may be difficult to recognize solely by auditory means. So pay special attention to **RARE WORDS** as their accuracy is crucial for the acceptance of the annotation.
3. Spoken Form Annotation: All annotation results should be in **SPOKEN FORM** and NOT in written form. For example, write 'two' instead of '2', 'one hundred' instead of '100' and 'dollar' instead of '\$'.
4. Avoid Punctuation and Capitalization: Do not include any punctuation marks or capitalize any letters in the results.
5. Rejection Criteria:
 - Avoid Direct Copying of References: If an individual's submissions are the **SAME** (punctuation excluded) with one of the references on multiple occasions, his(her) submissions will be **REJECTED**.

Reference1:

the first contribution is the vince data set which has a total of four eight oh oh annotated ui images across two design stages

Reference2:

the first contribution is the vince data set which has a total of four thousand eight hundred annotated ui images across two design stages

Input the best whole sentence...

Submit

Data Collection

• VINS dataset* contains **4,800** images of annotated UI designs screens:

- Abstract wireframes: 257 images
- High-fidelity screens: 4,543 images

0:08 / 0:09

Figure 13: Annotation interface for manual combination in constructing speech transcription.

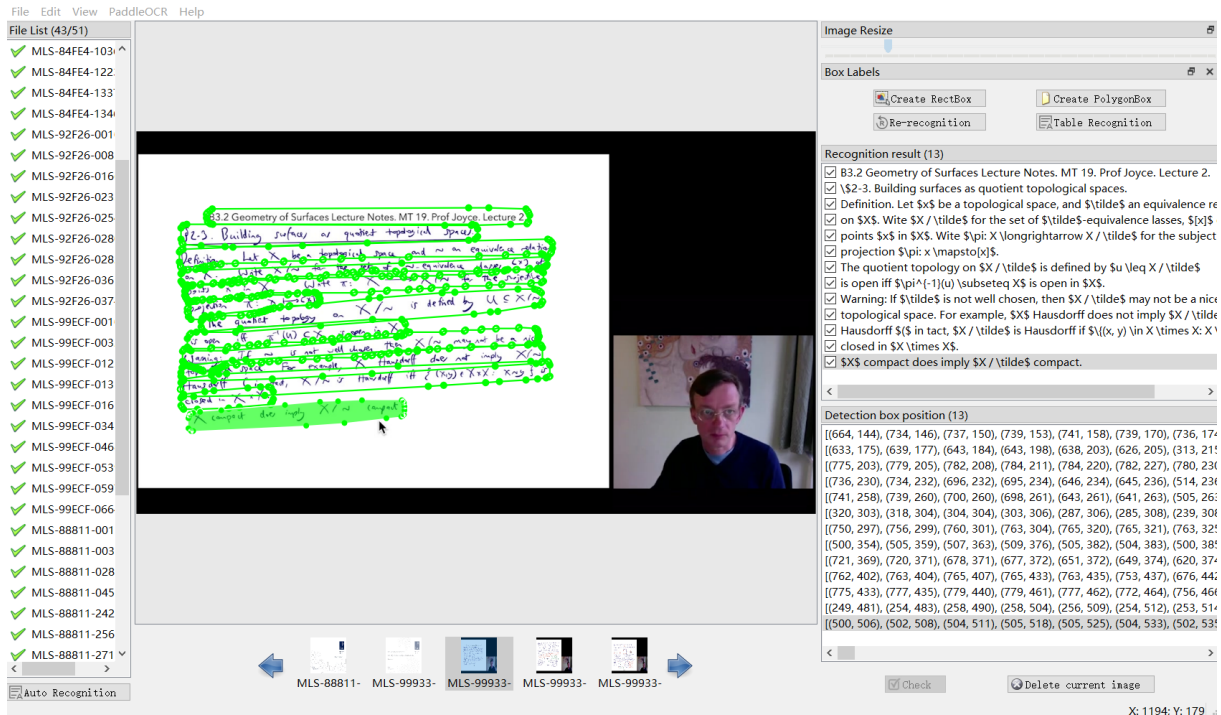


Figure 14: Annotation interface for correcting OCR results.

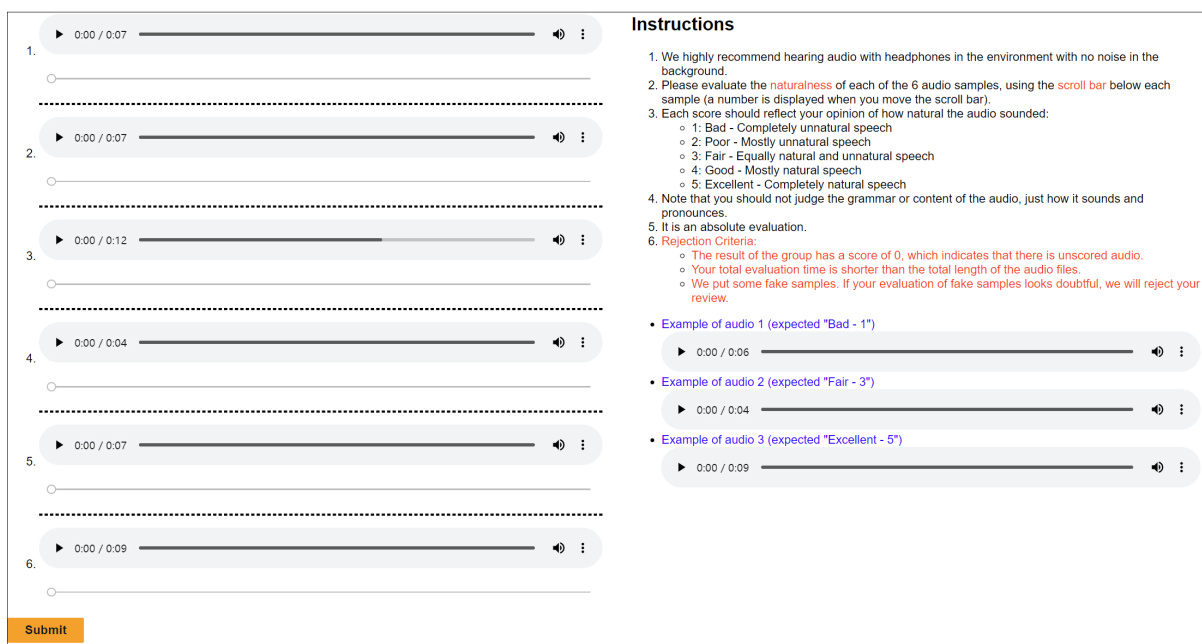


Figure 15: Annotation interface for MOS scoring in evaluating TTS.

Instructions

1. In this task, you will evaluate the quality of the generated scripts. We provide the **true script** (in text form) as a standard answer, please rate the **generated script** (in text form).
2. Each item is rated by multiple annotators, and **assignments that differ significantly from the overall rating will be REJECTED**.
3. Rate the generated script on a scale from **1 (worst)** to **5 (best)** by its **relevance, consistency, fluency, and coherence**. Please rate by dragging the corresponding sliders.
 - **Relevance:** The rating measures how well the generated script captures the key points of the true script. Consider whether all and only the important aspects are contained in the generated script.
 - **Consistency:** The rating measures whether the facts in the generated script are consistent with the facts in the true script. Consider whether the generated script reproduces all facts accurately and does not make up untrue information.
 - **Fluency:** This rating measures the quality of the generated script and, whether it is well-written and grammatically correct. Consider the quality of individual sentences.
 - **Coherence:** The rating measures the quality of all sentences in the generated script collectively, to fit together and sound naturally. Consider the quality of the generated script as a whole.

True Script

and the stromal cells and how they work. And so I want to just finish off by thanking Rosa, whom I've mentioned Mariana, who's a super talented computational postdoc and this was really a dream team that that ripped through this project in in supersonic time in my lab, and I've mentioned Muzz Haniffa Lab, with whom we worked very closely on this up in Newcastle, close to human developmental biology resource. And also, I want to say a big thank you to Ashley Moffett, who was really an intellectual partner in this work from the very beginning. She's an emeritus professor in the pathology department and has what's a. Has worked on NK cells in the Doors for decades, and also Mike Stubbington, Angel Gonzalez and the

Generated Script

Welcome to this segment of our presentation on the Human Cell Atlas. As we delve further into the complexities and the intricacies of our project, I want to pause and pay due credit to those who have been instrumental in its development and implementation. First and foremost, I'd like to extend a profound thank you to the Teichmann Lab for their exemplary work. Your team's motto - "Be Bold. Be Brilliant. Be Kind." - has truly echoed throughout this endeavor. Our partnership with the Pathology Department at Cambridge University has been invaluable. Your expertise and guidance have greatly contributed to the overall success of our project. Additionally, the European Molecular Biology Organisation or EMBO, we are deeply grateful for your involvement. I want to express my appreciation to our talented researchers Roser Vento, Miriana Efreanova, Margherita Y. Turco, Lucy Gardner, and Ashley Moffett. Your dedication to science is truly inspiring. Moreover, the Human Frontier Science Program, or HFSR your support has truly been instrumental. To Angela Zou, Kerstin Meyer, Johan Henriksson, and Raghd Rostom from Newcastle University, your contribution is greatly appreciated. Our collaboration with the Wellcome Sanger Institute has been immensely helpful in advancing our research. The MuzzHaniffa Lab and HDBR, your insights have been indispensable. Let me also take a moment to acknowledge Mike Stubbington, Angela Goncalves, Gavin Wright, Sam Behjati, and

1. Relevance



2. Consistency



3. Fluency



4. Coherence



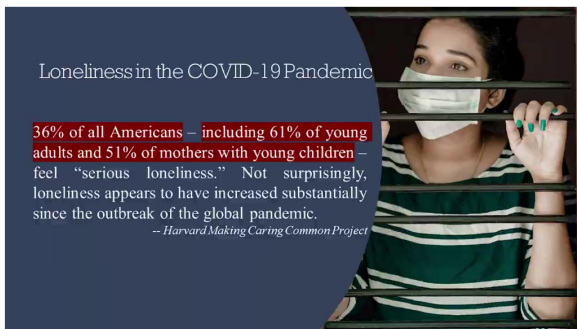
Submit

Figure 16: Annotation interface for manual scoring in Slide → Script task.

Instructions

1. In this task, you will evaluate the quality of the generated slides. We provide the **true slide** (in picture form) as a standard answer, please rate the **generated slide** (in text form).
2. Each item is rated by multiple annotators, and **assignments that differ significantly from the overall rating will be REJECTED**.
3. Rate the generated slide on a scale from **1 (worst)** to **5 (best)** by its **relevance, consistency, fluency, and coherence**. Please rate by dragging the corresponding sliders.
 - **Relevance:** The rating measures how well the generated slide captures the key points of the true slide. Consider whether all and only the important aspects are contained in the generated slide.
 - **Consistency:** The rating measures whether the facts in the generated slide are consistent with the facts in the true slide. Consider whether the generated slide reproduces all facts accurately and does not make up untrue information.
 - **Fluency:** This rating measures the quality of the generated slide and, whether it is well-written and grammatically correct. Consider the quality of individual sentences.
 - **Coherence:** The rating measures the quality of all sentences in the generated slide collectively, to fit together and sound naturally. Consider the quality of the generated slide as a whole.

True Slide



Generated Slide

"Loneliness: A Severe Side Effect of COVID-19 Social Distancing"

1. Relevance



2. Consistency



3. Fluency



4. Coherence



Submit

Figure 17: Annotation interface for manual scoring in Script → Slide task.