

# MathSpeech: Leveraging Small LMs for Accurate Conversion in Mathematical Speech-to-Formula

Sieun Hyeon<sup>1\*</sup>, Kyudan Jung<sup>2\*</sup>, Jaehee Won<sup>3</sup>, Nam-Joon Kim<sup>1†</sup>,  
Hyun Gon Ryu<sup>4</sup>, Hyuk-Jae Lee<sup>1,5</sup>, Jaeyoung Do<sup>1,5†</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, Seoul National University

<sup>2</sup>Department of Mathematics, Chung-Ang University

<sup>3</sup>College of Liberal Studies, Seoul National University

<sup>4</sup>NVIDIA

<sup>5</sup>Interdisciplinary Program in Artificial Intelligence, Seoul National University

zxc2692@snu.ac.kr , wjdrbeks1021@cau.ac.kr , wonjaehe\_e@snu.ac.kr , knj01@snu.ac.kr ,  
hryu@nvidia.com , hjlee@capp.snu.ac.kr , jaeyoung.do@snu.ac.kr

## Abstract

In various academic and professional settings, such as mathematics lectures or research presentations, it is often necessary to convey mathematical expressions orally. However, reading mathematical expressions aloud without accompanying visuals can significantly hinder comprehension, especially for those who are hearing-impaired or rely on subtitles due to language barriers. For instance, when a presenter reads Euler’s Formula, current Automatic Speech Recognition (ASR) models often produce a verbose and error-prone textual description (e.g.,  $e$  to the power of  $i$   $x$  equals cosine of  $x$  plus  $i$  *side* of  $x$ ), instead of the concise  $\LaTeX$  format (i.e.,  $e^{ix} = \cos(x) + i \sin(x)$ ), which hampers clear understanding and communication. To address this issue, we introduce MathSpeech, a novel pipeline that integrates ASR models with small Language Models (sLMs) to correct errors in mathematical expressions and accurately convert spoken expressions into structured  $\LaTeX$  representations. Evaluated on a new dataset derived from lecture recordings, MathSpeech demonstrates  $\LaTeX$  generation capabilities comparable to leading commercial Large Language Models (LLMs), while leveraging fine-tuned small language models of only 120M parameters. Specifically, in terms of CER, BLEU, and ROUGE scores for  $\LaTeX$  translation, MathSpeech demonstrated significantly superior capabilities compared to GPT-4o. We observed a decrease in CER from 0.390 to 0.298, and higher ROUGE/BLEU scores compared to GPT-4o.

**Code** — <https://github.com/hyeonsieun/MathSpeech>

## Introduction

Taking lectures that cover mathematical content through videos and watching academic presentations via video recordings or online streaming is no longer something out of the ordinary. Online video platforms have revolutionized mathematical education by making high-quality lectures and

resources accessible to a global audience, breaking down geographical and financial barriers. However, these platforms face a critical challenge: accurately generating subtitles for mathematics lectures. Although platforms like YouTube offer automatic subtitle services, their performance deteriorates markedly when handling mathematical content, particularly equations and formulas. As a result, some providers (such as MIT OpenCourseWare<sup>1</sup>) are compelled to manually create subtitles, but this requires a tremendous amount of effort and human labeling resources.

The crux of this issue lies in the limitations of current Automated Speech Recognition (ASR) models when confronted with mathematical expressions. Despite advancements in ASR technology (Radford et al. 2023; NVIDIA 2024a,b; Gandhi, von Platen, and Rush 2023), these models significantly underperform in recognizing and transcribing mathematical speech. This shortcoming became apparent during our initial investigations, highlighting the absence of a benchmark dataset specifically for evaluating ASR models’ proficiency in mathematical contexts.

To address this gap, we first developed a novel benchmark dataset comprising 1,101 audio samples from real mathematics lectures available on YouTube. This dataset<sup>2</sup> serves as a crucial tool for assessing the capabilities of various ASR models in mathematical speech recognition. Our evaluations using this dataset revealed not only poor performance in equation transcription but also a critical lack of  $\LaTeX$  generation ability, which is the standard for typesetting mathematical equations, in existing ASR models. This limitation significantly impedes learners’ comprehension, especially when dealing with complex mathematical expressions. An example of this case is presented in Table 1.

To address these challenges, we introduce MathSpeech, an innovative ASR pipeline specifically designed to transcribe mathematical speech directly into  $\LaTeX$  code instead of plain text. This approach enhances the learning experience by enabling the accurate rendering of mathematical ex-

\*These authors contributed equally.

†Corresponding authors

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup><https://www.youtube.com/@mitocw>

<sup>2</sup><https://huggingface.co/datasets/AAAI2025/MathSpeech>

<b>Formula</b>	$e^{ix} = \cos(x) + i \sin(x)$
<b>Spoken English(SE)</b>	e to the power of i x equals cosine of x plus i sine of x
<b>ASR result with error</b>	e to the power of i x equals cosine of x plus i side of x

Table 1: Comparison between the actual equation and the ASR result

pressions in subtitles, thereby supporting learners in their math education. Rather than incurring the significant costs associated with fine-tuning ASR models on domain-specific mathematical speech data—especially given the lack of publicly available datasets—we developed a novel method that converts mathematical speech into  $\LaTeX$  using small Language Models (LMs). Our methodology corrects ASR outputs containing mathematical expressions (even if they contain errors) and transforms the corrected output into  $\LaTeX$  using fine-tuned small LMs. By employing effective fine-tuning techniques that account for ASR errors and the nuances of spoken English in mathematical contexts, our pipeline has demonstrated superior  $\LaTeX$  generation capabilities compared to commercial large language models like GPT-4o (OpenAI 2024a) and Gemini-Pro (Google Deepmind 2024), despite using a relatively small language model with only 120M parameters.

In summary, our contributions are as follows:

- We constructed and released the first benchmark dataset for evaluating ASR models’ ability to transcribe mathematical equations.
- We identified and demonstrated the poor performance of existing ASR models in reading mathematical equations.
- We proposed a pipeline that corrects ASR errors and converts the output into  $\LaTeX$
- We confirmed that our pipeline, despite being significantly smaller (120M parameters) than commercial LLMs, outperformed GPT-4o (OpenAI 2024a) and Gemini-Pro (Google Deepmind 2024).

## Related Works

### ASR correction with LM

As ASR (Radford et al. 2023; NVIDIA 2024a,b; Gandhi, von Platen, and Rush 2023) systems have advanced, the need to correct ASR errors has become increasingly important. To enhance the quality of ASR output, there has been a significant amount of prior research focused on post-processing using language models. Since ASR outputs are in text form, many studies have employed sequence-to-sequence techniques.

In the past, statistical machine translation (Cucu et al. 2013; D’Haro and Banchs 2016) was used for this purpose. With the development of neural network-based language models, autoregressive sequence-to-sequence models are used for error correction (Tanaka et al. 2018; Liao et al. 2023), like neural machine translation. Moreover, with the

advancement of attention mechanisms (Chan et al. 2016; Inaguma and Kawahara 2023), research utilizing the Transformer architecture (Vaswani et al. 2017) for error correction has demonstrated strong performance (Mani et al. 2020; Leng et al. 2021b, 2023). Additionally, research on ASR error correction has been conducted using various language models, such as BERT (Yang, Li, and Peng 2022), BART (Zhao et al. 2021), ELECTRA (Futami et al. 2021; Yeen, Kim, and Koo 2023), and T5 (Ma et al. 2023a; Yeen, Kim, and Koo 2023).

Moreover, with the emergence of various Large Language Models (LLMs), which have shown remarkable performance across diverse domains, research on post-processing for ASR correction has also been actively conducted (Hu et al. 2024b; Sachdev, Wang, and Yang 2024; Ma et al. 2023c; Hu et al. 2024a). Several studies have shown that LLMs can be effectively used for ASR correction. However, LLMs have certain drawbacks, such as their large size and slow inference speed.

When using LLMs for ASR error correction, some research has demonstrated that utilizing multiple candidates generated during beam search can result in a voting effect (Leng et al. 2021a), leading to improved performance. This method, known as N-best, allows the N candidates obtained from the ASR output to provide clues to the language model regarding potential errors. Many studies (Imamura and Sumita 2017; Zhu et al. 2021; Ganesan et al. 2021; Ma et al. 2023b,a; Leng et al. 2023; Weng et al. 2020) on ASR correction have adopted this N-best approach.

### $\LaTeX$ translation and generation

Research related to  $\LaTeX$  has mainly focused on converting formula images to  $\LaTeX$  (Blecher et al. 2023; Blecher, L 2024) using Optical Character Recognition(OCR). Recently, studies have also been conducted on generating  $\LaTeX$  from spoken English that describes formulas (Jung et al. 2024a). However, this research has not explored correcting ASR results obtained from actual speech before translating them into  $\LaTeX$ . The focus of this study has been on translating clean, error-free spoken English into  $\LaTeX$ , without considering ASR errors

Additionally, research has been conducted on using Large Language Models (LLMs) to enhance efficiency and quality in popular academic writing tools like Overleaf (Wen et al. 2024), which includes generating  $\LaTeX$  with LLMs. In fact, well-known state-of-the-art commercial LLMs such as GPT series (OpenAI 2024b,a), Gemini-Pro (Google Deepmind 2024) have demonstrated remarkable abilities in generating  $\LaTeX$ . However, the previous study (Jung et al. 2024a,b; Hyeon et al. 2025) indicates that comparing the  $\LaTeX$  generation and translation capabilities of different LLMs presents a significant challenge, as there is no suitable metric to measure  $\LaTeX$  generation performance.

## Motivation

Subtitle services are often used when individuals watch academic videos or lectures. For the general public, subtitles serve as an auxiliary tool to help them understand video

content. However, for individuals with hearing impairments or students who speak a different language than the lecturer, subtitles are essential. Inaccurate subtitle services can severely hinder content comprehension, leading to a significant decrease in learning effectiveness. With recent advancements in Automatic Speech Recognition (ASR) models, the ability to convert speech into text has become highly accurate, greatly benefiting these users. However, the accuracy of ASR models remains significantly lower for academic videos in fields such as mathematics and physics than for other subjects.

	Models	Params	WER(%) (Leaderboard)	WER(%) (Formula)
OpenAI	Whisper-base	74M	10.3	34.7
	Whisper-small	244M	8.59	29.5
	Whisper-largeV2	1550M	7.83	31.0
	Whisper-largeV3	1550M	7.44	33.3
NVIDIA	Canary-1B	1B	6.5	35.2

Table 2: The WER for Leaderboard was from the HuggingFace Open ASR Leaderboard, while the WER for Formula was measured using our MathSpeech Benchmark.

Table 2 presents the results of measuring the WER (Word Error Rate) of ASR models using mathematical speech collected from actual lecture videos.

When comparing the WER results of Whisper (Radford et al. 2023) and Canary-1b (NVIDIA 2024a) on the HuggingFace Open ASR Leaderboard<sup>34</sup> with the MathSpeech benchmark dataset WER results, we observed that the WER for formula speech was significantly higher. The reasons for the elevated error rates are as follows.

**(1) Severe Noise:** Our benchmark dataset consists of audio from lecture videos recorded 10 to 20 years ago, resulting in relatively poor audio quality and higher levels of noise. Additionally, since the audio is taken from real lectures, it includes sounds such as chalk writing on the blackboard and students’ chatter mixed into the speech.

**(2) Non-native Accent:** When the speaker is not a native English speaker, the model often misinterprets some words as other words due to the speaker’s accent. Our benchmark dataset includes speakers with distinctive accents.

**(3) Label Ambiguity Problem:** Mathematical expressions read aloud can be ASR-transcribed in multiple ways, which increases the WER. In other words, while the meaning is correct, the text output differs, leading it to be counted as an error. For example, in the case of speech reading number **1**, we transcribed it as **1**, but the ASR model outputs it as **one**. Although the ASR output for mathematical speech differed from the labels we assigned, it was semantically equivalent, which contributed to the higher WER measurements.

To accurately convert the spoken English of mathematical expressions into  $\LaTeX$ , it is necessary to carefully consider

<sup>3</sup>[https://huggingface.co/spaces/hf-audio/open\\_asr\\_leaderboard](https://huggingface.co/spaces/hf-audio/open_asr_leaderboard)

<sup>4</sup>This value is based on results as of 2024-08-16.

these factors and adjust the ASR errors accordingly. Therefore, we propose a method for training small language models that considers both the  $\LaTeX$  output and ASR transcription results. This will be discussed in the following section.

## Methodology

### Datasets

Our goal is to implement a lightweight, fast, and high-performance pipeline that outputs  $\LaTeX$ . We determined that the most effective approach is fine-tuning small LMs, which involves collecting appropriate datasets. To achieve this, we collected two types of data.

#### (1) (Spoken English, $\LaTeX$ ) pairs

As mentioned earlier, ASR models convert spoken mathematical expressions into plain English text. Since our goal is to translate such Spoken English(SE) into  $\LaTeX$ , we decided to use a dataset of (Spoken English,  $\LaTeX$ ) pairs. We were able to obtain a publicly available dataset (Jung et al. 2024a) on HuggingFace and used it in our work. This data was collected by web crawling and OCR, extracting only the  $\LaTeX$  portions from arxiv papers and textbooks. This is a large dataset containing 23 million (SE,  $\LaTeX$ ) data pairs. We fine-tuned the T5-small (Raffel et al. 2023) with this dataset to convert Spoken English(SE) to  $\LaTeX$ .

#### (2) ASR Error Correction Dataset

In our initial experiments, we attempted to fine-tune T5 (Raffel et al. 2023) only using the (Spoken English(SE),  $\LaTeX$ ) pair dataset and connect it to the ASR as a post-LM. However, the performance was not satisfactory because the ASR model itself made significant errors when converting spoken mathematical expressions into plain text (Table 3).

Spoken English(SE)	ASR error result
x plus 5y plus 10z equals 0	x plus 5y plus 10z école 0
cosine of psi sub i, psi sub j	Posing of psi sub i, psi sub j

Table 3: The examples of ASR error results

Since such erroneous texts did not exist in the dataset (Jung et al. 2024a), the fine-tuned T5 also produced incorrect  $\LaTeX$  outputs. Therefore, we determined that adding a process to correct errors that occur in the ASR system would significantly improve performance. Thus, we built another dataset for fine-tuning the ASR error correction model. As shown in Figure 1, we converted Spoken English (SE) into speech using TTS and then fed the speech into the ASR model to obtain the erroneous Spoken English ASR outputs.

In other words, let  $y_{se}$  represent Spoken English and  $y_{latex}$  represent the  $\LaTeX$  expression for  $y_{se}$ . Then, we can denote the data pair (Jung et al. 2024a) as  $(y_{se}, y_{latex})$ .

Using TTS, we generate audio for  $y_{se}$ . If we denote this audio as  $x_{tts}$ , then we can denote the text obtained by inputting  $x_{tts}$  into an ASR model as  $y_{asr^i}$ , where  $i$  is used to distinguish between different models. Therefore, using the method shown in Figure 1, we transform  $(y_{se}, y_{latex})$  into  $(y_{asr^i}, y_{se}, y_{latex})$ .



Figure 1: Method for Collecting ASR Error Results.




Figure 2: This figure compares 2-beam search and our method. The left shows top-2 beam search by a single ASR model, while the right shows top-1 beam search by two ASR models.

At this stage, we used the VITS (Kim, Kong, and Son 2021) model for TTS and collected the voices using 2-speaker. For the ASR model, we used Whisper-base, small, largeV2 (Radford et al. 2023) and Canary-1b (NVIDIA 2024a) to collect the Error ASR results. We collected 6M ASR error results using Whisper-base and small, and 1M ASR error results using Whisper-largeV2 and Canary-1b.

## Models

Our MathSpeech pipeline can be seen in Figure 3. We used two models configured in two stages.

### (1) Error Corrector

The purpose of the Error Corrector is to fix the errors that occur in the ASR model. In other words, the goal of the Error Corrector is to make the ASR error results similar to the input data of the  $\LaTeX$  Translator, which is called Spoken English. In previous studies, T5 (Ma et al. 2023a,b; Yeen, Kim, and Koo 2023) has demonstrated good performance in ASR error correction. To minimize the model size, we used T5-small for the Error Corrector.

Additionally, recent research has shown that inputting multiple candidates (Imamura and Sumita 2017; Zhu et al. 2021; Ganesan et al. 2021; Ma et al. 2023b,a; Leng et al. 2023; Weng et al. 2020) generated by ASR beam search can yield good performance in various ASR Corrector models due to the voting effect (Leng et al. 2021a). We have further developed this idea. Instead of using multiple candidates generated by ASR beam search, we use multiple candidates from the top-1 results of different models (Figure 2).

An advantage of this method is its versatility. We believe that training the Error Corrector with error information from various models will make it universally applicable to different ASR models. For example, if the corrector is trained with two ASR beam search results from Whisper-small, it may perform well on Whisper-small, but not on Whisper-base. However, if the corrector is trained with the ASR results from both Whisper-small and Whisper-base, we believe that it will effectively correct the errors for both models. Experimental results are presented in the following section.

### (2) $\LaTeX$ Translator

According to a previous study (Jung et al. 2024a), the T5 model achieved the best performance for the  $\LaTeX$  translator. To minimize the model size, we used the smallest version, T5-small. We fine-tuned T5-small using the (SE,  $\LaTeX$ ) pair dataset so that when Spoken English is input,  $\LaTeX$  is output.

## Training

We implemented a pipeline connecting two T5-small models. Instead of simply chaining the fine-tuned Error Corrector and  $\LaTeX$  Translator, we performed end-to-end training (Figure 4). Considering the characteristics of ASR error results and  $\LaTeX$ , we constructed the loss function as follows.

For a given input audio  $x_{tts}$ , let the inference results of two ASR models be  $y_{asr^1}$  and  $y_{asr^2}$ . Here,  $y_{asr^1}$  and  $y_{asr^2}$  are the outputs from the ASR models, which may contain errors. We then provide these two ASR outputs to the Error Corrector  $F$ . The resulting text can be denoted as  $F(y_{asr^1}, y_{asr^2})$ . Next, we take  $\hat{y}_{se} = F(y_{asr^1}, y_{asr^2})$  and feed it into the  $\LaTeX$  Translator  $G$  to produce the corresponding  $\LaTeX$ , denoted as  $\hat{y}_{latex} = G(F(y_{asr^1}, y_{asr^2}))$ . The loss function is then defined as follows:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{se} + \lambda_2 \mathcal{L}_{latex} \quad (1)$$

And  $\mathcal{L}_{se}$  and  $\mathcal{L}_{latex}$  are calculated as cross-entropy losses of the tokenized outputs  $\hat{y}_{se}$  and  $y_{se}$ , and  $\hat{y}_{latex}$  and  $y_{latex}$ , respectively.

$$\mathcal{L}_{se} = - \sum_{t=1}^T \sum_{i=1}^V y_{se_t^i} \log(\hat{y}_{se_t^i}) \quad (2)$$

$$\mathcal{L}_{latex} = - \sum_{t=1}^T \sum_{i=1}^V y_{latex_t^i} \log(\hat{y}_{latex_t^i}) \quad (3)$$

The cross-entropy loss is calculated as a negative sum over all time steps  $t$  from 1 to  $T$  (the length of the sequence) and over all possible words  $i$  in the vocabulary  $V$ . At each

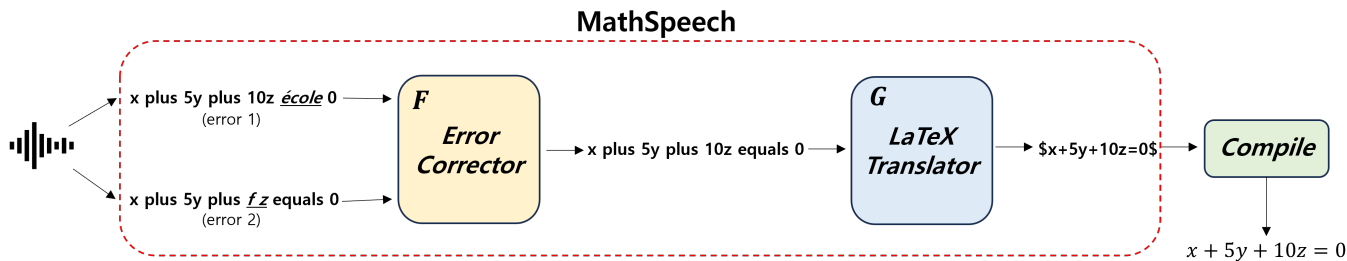


Figure 3: Our pipeline that converts the lecturer’s voice into  $\LaTeX$ .

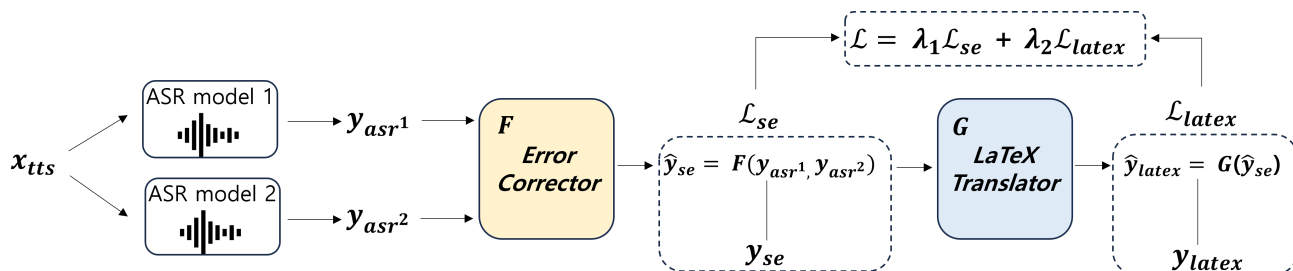


Figure 4: The method of training MathSpeech in an end-to-end manner

time step  $t$ ,  $y_{se_t}$  is a one-hot encoded vector over the vocabulary, where  $y_{se_t^i}$  is its  $i$ -th element. Specifically,  $y_{se_t^i} = 1$  if the correct ground-truth word at time step  $t$  is the  $i$ -th word in the vocabulary, and 0 otherwise. Correspondingly,  $\hat{y}_{se_t}$  is the predicted probability distribution vector over the vocabulary at time  $t$ , and  $\hat{y}_{se_t^i}$  is the predicted probability that the correct word at time  $t$  is the  $i$ -th word in the vocabulary.

We calculated the final loss  $\mathcal{L}$  by assigning different weights to the two cross entropy loss functions. The weight assigned to the  $\hat{y}_{latex}$  was set higher than the weight for the loss on  $\hat{y}_{se}$  because there can be different SE for the same  $\LaTeX$  result (Table 4).

1st SE	x plus five y plus ten z equals zero
2nd SE	x plus 5 y plus 10 z equals 0
$\LaTeX$	$\$x + 5y + 10z = 10\$$

Table 4: The reason why errors between  $\LaTeX$  are more critical than errors between SE.

In other words, even if the SEs are slightly different, the  $\LaTeX$  can still produce the correct answer because they are semantically equivalent. Therefore, our goal is to obtain the correct  $\LaTeX$ , so we assigned a higher weight to  $\mathcal{L}_{latex}$  when constructing the loss function. Our experimental results showed that setting the weight  $\lambda_1$  for SE to 0.3 and the weight for  $\LaTeX$  to 0.7 yielded the best performance, so we used these values.

### Evaluation Metrics

In previous studies (Jung et al. 2024a), when evaluating  $\LaTeX$ , metrics commonly used in translation tasks, such as ROUGE (Lin 2004) and BLEU (Papineni et al. 2002), were

employed. Based on this idea, we used ROUGE-1, BLEU, and ROUGE-L. Furthermore, we employed CER, a traditional ASR metric. However, WER was not measured in our experiments. This is because the spaces are often ignored in  $\LaTeX$ . For example,  $\$A B\$$  and  $\$AB\$$  result in the same  $\LaTeX$  compilation output. So the evaluation was conducted after removing all spaces.

### Experiments

To evaluate the  $\LaTeX$  translation capabilities of our pipeline, we compared its performance against existing commercial large language models.

#### Setup

We fixed the same hyperparameters for the fine-tuned models. The maximum number of training epochs was set to 20, and the model with the lowest validation loss was selected. The learning rate was set to a maximum of  $1e-4$  and a minimum of  $1e-6$ , adjusted using a linear learning rate scheduler. For the Error Corrector, which requires two ASR outputs as input, the maximum input sequence length was set to 540, with an output length of 275. For the  $\LaTeX$  translator, both input and output sequence lengths were set to 275. T5-small was trained with a batch size of 48 on an NVIDIA A100, and T5-base with a batch size of 84 on an NVIDIA H100.

As a comparison group for our pipeline, we selected GPT-3.5 (OpenAI 2024b), GPT-4o (OpenAI 2024a), and Gemini-Pro (Google Deepmind 2024), using 1-shot prompting with one example for all. To observe  $\LaTeX$  translation results across various ASR models, we used five ASR models.

#### Result

The experimental results are presented in Tables 5 and 6, respectively. Table 5 lists the outcomes of inputting two ASR

ASR	LM	CER ↓	ROUGE-1 ↑	ROUGE-L ↑	BLEU ↑
whisper-base	GPT-3.5	0.443	0.782	0.775	0.483
	GPT-4o	<i>0.410*</i>	<i>0.813*</i>	<i>0.808*</i>	<i>0.487*</i>
	Gemini-Pro	0.424	0.756	0.749	0.418
	<b>MathSpeech</b>	<b>0.336</b>	<b>0.824</b>	<b>0.819</b>	<b>0.662</b>
whisper-small	GPT-3.5	0.391	0.815	0.809	<i>0.519*</i>
	GPT-4o	<i>0.384*</i>	<i>0.840*</i>	<i>0.835*</i>	0.516
	Gemini-Pro	0.390	0.799	0.792	0.455
	<b>MathSpeech</b>	<b>0.309</b>	<b>0.852<sup>++</sup></b>	<b>0.847<sup>++</sup></b>	<b>0.689<sup>++</sup></b>
whisper-largeV2	GPT-3.5	0.401	0.820	0.813	0.519
	GPT-4o	<i>0.380*</i>	<i>0.844*</i>	<i>0.839*</i>	<i>0.520*</i>
	Gemini-Pro	0.393	0.800	0.792	0.458
	<b>MathSpeech</b>	<b>0.298<sup>++</sup></b>	<b>0.848</b>	<b>0.844</b>	<b>0.683</b>
whisper-largeV3	GPT-3.5	0.404	0.797	0.792	0.493
	GPT-4o	<i>0.398*</i>	<b>0.826</b>	<b>0.822</b>	<i>0.495*</i>
	Gemini-Pro	0.409	0.772	0.767	0.431
	<b>MathSpeech</b>	<b>0.317</b>	<i>0.817*</i>	<i>0.812*</i>	<b>0.673</b>
canary-1b	GPT-3.5	0.458	0.795	0.787	0.434
	GPT-4o	<i>0.422*</i>	<i>0.824*</i>	<i>0.819*</i>	<i>0.445*</i>
	Gemini-Pro	0.442	0.781	0.774	0.392
	<b>MathSpeech</b>	<b>0.325</b>	<b>0.832</b>	<b>0.828</b>	<b>0.674</b>

Table 5: Having only one model in the ASR column means that the top-2 beam search ASR outputs of a single model are input into the LLM. The bold text indicates the result with the best score for the same ASR output, while an asterisk (\*) indicates the second-best score for the same ASR output. The double plus sign (++) indicates the case with the best score.

prediction candidates that require correction into the LM. In this regard, MathSpeech achieved the best scores for the CER, ROUGE-1, ROUGE-L, and BLEU. The LLM with the second-best performance was GPT-4o, while GPT-3.5 and Gemini-Pro showed similar results.

The main factor that lowered the performance scores of the commercial large language models was hallucination. When the ASR results were unusual or ambiguous, the baseline LLMs would either output completely different  $\LaTeX$  formulas or produce non- $\LaTeX$  texts (e.g., Sorry, I can't understand). Additionally, whisper-small and whisper-largeV2 outperformed whisper-base, whisper-largeV3, and canary-1b in overall metric scores. This result aligns with the WER measurements of Spoken English for whisper-small and whisper-largeV2, which were relatively better, as observed in Table 2. In other words, better ASR results lead to improved  $\LaTeX$  translation performance.

Table 6 lists the outcomes when the top-1 results from different ASR models were used as inputs. The highest performance scores were achieved when the ASR results from Whisper-base & small were used as inputs. This can be attributed to the fact that our ASR error results dataset contains relatively more information from Whisper-base & small. The second highest performance was observed when using the ASR results from Whisper-small & largev2, which can be attributed to the relatively lower WER of these two models on our benchmark dataset. Furthermore, the strong performance on Whisper-largeV3, for which we did not collect ASR error results, demonstrates that MathSpeech can perform well even on ASR models that were not used during training. Moreover, our MathSpeech model, being a small-

sized model with 120M parameters, has low latency. When inference latency was measured on an NVIDIA V100 GPU, it took 0.45 seconds to convert the ASR result of 5 seconds of speech into  $\LaTeX$ .

## Ablation Study

To demonstrate the effectiveness of the MathSpeech structure, we conducted ablation studies (Table 7).

To show that correcting ASR outputs is crucial for  $\LaTeX$  translation, we removed the corrector and conducted an experiment in which ASR outputs were directly translated into  $\LaTeX$ . As a result, both T5-small and T5-base showed significant performance degradation.

To demonstrate that implementing the corrector and translator in a 2-stage structure is effective, we trained a single T5 model to perform both correction and translation in a 1-stage process and observed its performance. The experiment, where we trained the model to translate two ASR error outputs into  $\LaTeX$  with the same setup, showed lower performance compared to our pipeline.

To validate the effectiveness of our end-to-end training method, we compared it with a method where the corrector and translator were trained separately and simply concatenated. The results confirm that the proposed training method is more effective.

Since a single T5-base (220M) is larger than two T5-small models (120M), we did not apply our end-to-end training method to the 2-stage T5-base pipeline. However, we can infer that if we were to apply our training method to the simple concatenation of two T5-base models, the performance could potentially improve further.

ASR	LM	CER ↓	ROUGE-1 ↑	ROUGE-L ↑	BLEU ↑
whisper-base & whisper-small	GPT-3.5	0.407	0.801	0.792	0.502
	GPT-4o	0.379*	0.843*	0.839*	0.518*
	Gemini-Pro	0.382	0.802	0.796	0.457
	<b>MathSpeech</b>	<b>0.243<sup>++</sup></b>	<b>0.870<sup>++</sup></b>	<b>0.864<sup>++</sup></b>	<b>0.718<sup>++</sup></b>
whisper-small & whisper-largeV2	GPT-3.5	0.388	0.825	0.818	0.523
	GPT-4o	0.373*	0.852*	0.848*	0.532*
	Gemini-Pro	0.374	0.806	0.800	0.469
	<b>MathSpeech</b>	<b>0.269</b>	<b>0.864</b>	<b>0.859</b>	<b>0.708</b>
canary-1b & whisper-largeV2	GPT-3.5	0.399	0.826	0.820	0.524*
	GPT-4o	0.386*	<b>0.849</b>	<b>0.844</b>	0.516
	Gemini-Pro	0.393	0.800	0.793	0.458
	<b>MathSpeech</b>	<b>0.294</b>	0.848*	0.843*	<b>0.694</b>
canary-1b & whisper-largeV3	GPT-3.5	0.394	0.813	0.806	0.504
	GPT-4o	0.375*	0.849*	0.844*	0.514
	Gemini-Pro	0.399	0.810	0.804	0.445
	<b>MathSpeech</b>	<b>0.292</b>	<b>0.853</b>	<b>0.846</b>	<b>0.698</b>
<i>Total Average</i>	GPT-3.5	0.409	0.808	0.801	0.500
	GPT-4o	0.390*	0.838*	0.833*	0.505*
	Gemini-Pro	0.400	0.792	0.785	0.442
	<b>MathSpeech</b>	<b>0.298</b>	<b>0.845</b>	<b>0.840</b>	<b>0.689</b>

Table 6: Having two models in the ASR column means that the top-1 beam search ASR output from each of the two different models is obtained and both are input into the LLM. *Total Average* refers to the average of all the results from Table 5 and 6.

ASR	LM	CER ↓	ROUGE-1 ↑	ROUGE-L ↑	BLEU ↑
whisper-base	T5-small (1 stage, w/o corrector)	0.693	0.734	0.729	0.483
	T5-base (1 stage, w/o corrector)	0.530*	0.749*	0.744*	0.554*
	<b>MathSpeech</b>	<b>0.336</b>	<b>0.824</b>	<b>0.819</b>	<b>0.662</b>
whisper-small	T5-small (1 stage, w/o corrector)	0.630	0.769	0.765	0.521
	T5-base (1 stage, w/o corrector)	0.423*	0.783*	0.780*	0.618*
	<b>MathSpeech</b>	<b>0.309</b>	<b>0.852</b>	<b>0.847</b>	<b>0.689</b>
whisper-base & whisper-small	T5-small (1 stage, fine-tuned with errors)	0.403	0.824*	0.819*	0.635
	T5-base (1 stage, fine-tuned with errors)	0.358	0.813	0.809	0.650
	T5-small (2 stage, Just connect)	0.357	0.820	0.815	0.656*
	T5-base (2 stage, Just connect)	0.343*	0.817	0.814	0.651
	<b>MathSpeech</b>	<b>0.243</b>	<b>0.870</b>	<b>0.864</b>	<b>0.718</b>

Table 7: Ablation Study Results. 1 stage refers to using a single T5 model, while 2 stage refers to using two T5 models.

## Future Works

### (1) Defining a Metric to Solve the Label Ambiguity

Since  $\mathcal{L}_{\text{TeX}}$  can represent the same formula in multiple ways, it is necessary to consider various possible cases when evaluating the performance of  $\mathcal{L}_{\text{TeX}}$  translation.  $\mathcal{L}_{\text{TeX}}$  is closer to a computer language, like SQL, than to natural language, so metrics like BLEU or CER are not perfect for evaluating  $\mathcal{L}_{\text{TeX}}$ . Therefore, it is necessary to implement a metric that is more suitable for evaluating  $\mathcal{L}_{\text{TeX}}$ .

### (2) Formula Detection in Practice and $\mathcal{L}_{\text{TeX}}$ Conversion

This research focuses only on the ability of ASR models and LMs to generate  $\mathcal{L}_{\text{TeX}}$ . However, to apply this to actual subtitle services, it is essential to develop the ability to detect and separate formulaic parts from speech. This can likely be achieved by training the LM not only on  $\mathcal{L}_{\text{TeX}}$  but also on mixed general text. Additionally, in real-world situations,

the speaker may not finish verbally expressing the entire formula. It must be possible to complete such interrupted formulas into full formulas. This could be implemented through the inference capabilities of large language models (LLMs).

## Conclusion

In this paper, we confirmed through a self-constructed benchmark dataset that existing ASR models lack the ability to read mathematical formulas and are unable to generate  $\mathcal{L}_{\text{TeX}}$ . To address this, we propose MathSpeech, a pipeline that connects ASR models with small LMs to generate  $\mathcal{L}_{\text{TeX}}$ . By effectively connecting two T5-small models and training them end-to-end, our approach demonstrated superior  $\mathcal{L}_{\text{TeX}}$  translation capabilities compared to existing commercial large language models. Our research opens up the possibility of more accurate subtitles in the field of math.

## Acknowledgments

This work was supported in part by National Research Foundation of Korea (NRF) grant (RS-2024-00414981), Institute of Information & communications Technology Planning & Evaluation (IITP) grant (RS-2024-00399936, RS-2021-II211343, RS-2024-00454666, IITP-2024-RS-2024-00397085, IITP-2024-RS-2024-00441407), and Research Grant (0418-20240053) from Seoul National University. Samsung Memory Research Center (SMRC) provided research facilities for this work. J. Do is with ASRI, Seoul National University.

## References

- Blecher, L.; Cucurull, G.; Scialom, T.; and Stojnic, R. 2023. Nougat: Neural Optical Understanding for Academic Documents. arXiv:2308.13418.
- Blecher, L. 2024. pix2tex - LaTeX OCR. <https://github.com/lukas-blecher/LaTeX-OCR>. Accessed: 2024-08-15.
- Chan, W.; Jaitly, N.; Le, Q.; and Vinyals, O. 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)*, 4960–4964.
- Cucu, H.; Buzo, A.; Besacier, L.; and Burileanu, C. 2013. Statistical error correction methods for domain-specific ASR systems. In *Statistical Language and Speech Processing: First International Conference, SLSP 2013, Tarragona, Spain, July 29-31, 2013. Proceedings 1*, 83–92. Springer.
- D’Haro, L. F.; and Banchs, R. E. 2016. Automatic Correction of ASR Outputs by Using Machine Translation. In *Proc. Interspeech 2016*, 3469–3473.
- Futami, H.; Inaguma, H.; Mimura, M.; Sakai, S.; and Kawahara, T. 2021. ASR Rescoring and Confidence Estimation with Electra. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 380–387.
- Gandhi, S.; von Platen, P.; and Rush, A. M. 2023. Distilwhisper: Robust knowledge distillation via large-scale pseudo labelling. *arXiv preprint arXiv:2311.00430*.
- Ganesan, K.; Bamdev, P.; B, J.; Venugopal, A.; and Tushar, A. 2021. N-Best ASR Transformer: Enhancing SLU Performance using Multiple ASR Hypotheses. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, 93–98. Online: Association for Computational Linguistics.
- Google Deepmind. 2024. Gemini Pro. Our best model for general performance across a wide range of tasks. <https://deepmind.google/technologies/gemini/pro/>. Accessed: 2024-08-15.
- Hu, Y.; Chen, C.; Qin, C.; Zhu, Q.; Chng, E.; and Li, R. 2024a. Listen Again and Choose the Right Answer: A New Paradigm for Automatic Speech Recognition with Large Language Models. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Findings of the Association for Computational Linguistics ACL 2024*, 666–679. Bangkok, Thailand and virtual meeting: Association for Computational Linguistics.
- Hu, Y.; CHEN, C.; Yang, C.-H. H.; Li, R.; Zhang, C.; Chen, P.-Y.; and Chng, E. 2024b. Large Language Models are Efficient Learners of Noise-Robust Speech Recognition. In *The Twelfth International Conference on Learning Representations*.
- Hyeon, S.; Jung, K.; Kim, N.-J.; Ryu, H. G.; and Do, J. 2025. MathReader : Text-to-Speech for Mathematical Documents. arXiv:2501.07088.
- Imamura, K.; and Sumita, E. 2017. Ensemble and Reranking: Using Multiple Models in the NICT-2 Neural Machine Translation System at WAT2017. In Nakazawa, T.; and Goto, I., eds., *Proceedings of the 4th Workshop on Asian Translation (WAT2017)*, 127–134. Taipei, Taiwan: Asian Federation of Natural Language Processing.
- Inaguma, H.; and Kawahara, T. 2023. Alignment Knowledge Distillation for Online Streaming Attention-Based Speech Recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31: 1371–1385.
- Jung, K.; Hyeon, S.; Youn, K. J.; Kim, N.-J.; Ryu, H. G.; Lee, H.-J.; and Do, J. 2024a. MathBridge: A Large-Scale Dataset for Translating Mathematical Expressions into Formula Images. arXiv:2408.07081.
- Jung, K.; Kim, N.-J.; Ryu, H.; Hyeon, S.; jun Lee, S.; and jae Lee, H. 2024b. TeXBLEU: Automatic Metric for Evaluate LaTeX Format. arXiv:2409.06639.
- Kim, J.; Kong, J.; and Son, J. 2021. Conditional Variational Autoencoder with Adversarial Learning for End-to-End Text-to-Speech. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 5530–5540. PMLR.
- Leng, Y.; Tan, X.; Liu, W.; Song, K.; Wang, R.; Li, X.-Y.; Qin, T.; Lin, E.; and Liu, T.-Y. 2023. Softcorrect: Error correction with soft detection for automatic speech recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 13034–13042.
- Leng, Y.; Tan, X.; Wang, R.; Zhu, L.; Xu, J.; Liu, W.; Liu, L.; Li, X.-Y.; Qin, T.; Lin, E.; and Liu, T.-Y. 2021a. FastCorrect 2: Fast Error Correction on Multiple Candidates for Automatic Speech Recognition. In Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Findings of the Association for Computational Linguistics: EMNLP 2021*, 4328–4337. Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Leng, Y.; Tan, X.; Zhu, L.; Xu, J.; Luo, R.; Liu, L.; Qin, T.; Li, X.; Lin, E.; and Liu, T.-Y. 2021b. FastCorrect: Fast Error Correction with Edit Alignment for Automatic Speech Recognition. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 21708–21719. Curran Associates, Inc.
- Liao, J.; Eskimez, S.; Lu, L.; Shi, Y.; Gong, M.; Shou, L.; Qu, H.; and Zeng, M. 2023. Improving Readability for Automatic Speech Recognition Transcription. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 22(5).
- Lin, C.-Y. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*,

- 74–81. Barcelona, Spain: Association for Computational Linguistics.
- Ma, R.; Gales, M. J. F.; Knill, K. M.; and Qian, M. 2023a. N-best T5: Robust ASR Error Correction using Multiple Input Hypotheses and Constrained Decoding Space. In *Proc. INTERSPEECH 2023*, 3267–3271.
- Ma, R.; Qian, M.; Gales, M. J. F.; and Knill, K. M. 2023b. Adapting an Unadaptable ASR System. In *Proc. INTERSPEECH 2023*, 989–993.
- Ma, R.; Qian, M.; Manakul, P.; Gales, M.; and Knill, K. 2023c. Can Generative Large Language Models Perform ASR Error Correction? [arXiv:2307.04172](https://arxiv.org/abs/2307.04172).
- Mani, A.; Palaskar, S.; Meripo, N. V.; Konam, S.; and Metze, F. 2020. ASR Error Correction and Domain Adaptation Using Machine Translation. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6344–6348.
- NVIDIA. 2024a. New Standard for Speech Recognition and Translation from the NVIDIA NeMo Canary Model. <https://developer.nvidia.com/blog/new-standard-for-speech-recognition-and-translation-from-the-nvidia-nemo-canary-model/>. Accessed: 2024-08-15.
- NVIDIA. 2024b. Pushing the Boundaries of Speech Recognition with NVIDIA NeMo Parakeet ASR Models. <https://developer.nvidia.com/blog/pushing-the-boundaries-of-speech-recognition-with-nemo-parakeet-asr-models/>. Accessed: 2024-08-15.
- OpenAI. 2024a. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>. Accessed: 2024-08-15.
- OpenAI. 2024b. OpenAI’s GPT-3.5 Turbo. <https://platform.openai.com/docs/models/gpt-3-5-turbo>. Accessed: 2024-08-15.
- Papineni, K.; Roukos, S.; Ward, T.; and Jing Zhu, W. 2002. BLEU: a Method for Automatic Evaluation of Machine Translation. 311–318.
- Radford, A.; Kim, J. W.; Xu, T.; Brockman, G.; McLeavey, C.; and Sutskever, I. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, 28492–28518. PMLR.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2023. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. [arXiv:1910.10683](https://arxiv.org/abs/1910.10683).
- Sachdev, R.; Wang, Z.-Q.; and Yang, C.-H. H. 2024. Evolutionary Prompt Design for LLM-Based Post-ASR Error Correction. [arXiv:2407.16370](https://arxiv.org/abs/2407.16370).
- Tanaka, T.; Masumura, R.; Masataki, H.; and Aono, Y. 2018. Neural Error Corrective Language Models for Automatic Speech Recognition. In *Proc. Interspeech 2018*, 401–405.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wen, H.; Wei, Z.; Lin, Y.; Wang, J.; Liang, Y.; and Wan, H. 2024. OverleafCopilot: Empowering Academic Writing in Overleaf with Large Language Models. *arXiv preprint arXiv:2403.09733*.
- Weng, Y.; Miryala, S. S.; Khatri, C.; Wang, R.; Zheng, H.; Molino, P.; Namazifar, M.; Papangelis, A.; Williams, H.; Bell, F.; and Tur, G. 2020. Joint Contextual Modeling for ASR Correction and Language Understanding. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6349–6353.
- Yang, J.; Li, R.; and Peng, W. 2022. ASR Error Correction with Constrained Decoding on Operation Prediction. In *Proc. Interspeech 2022*, 3874–3878.
- Yeen, H.-Y.; Kim, M.-J.; and Koo, M.-W. 2023. I Learned Error, I Can Fix It! : A Detector-Corrector Structure for ASR Error Calibration. In *Proc. INTERSPEECH 2023*, 2693–2697.
- Zhao, Y.; Yang, X.; Wang, J.; Gao, Y.; Yan, C.; and Zhou, Y. 2021. BART Based Semantic Correction for Mandarin Automatic Speech Recognition System. In *Proc. Interspeech 2021*, 2017–2021.
- Zhu, L.; Liu, W.; Liu, L.; and Lin, E. 2021. Improving ASR Error Correction Using N-Best Hypotheses. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 83–89.