



# ASR Error Detection via Audio-Transcript entailment

Nimshi Venkat Meripo, Sandeep Konam

Abridge AI Inc.

venkatm@abridge.com, san@abridge.com

## Abstract

Despite improved performances of the latest Automatic Speech Recognition (ASR) systems, transcription errors are still unavoidable. These errors can have a considerable impact in critical domains such as healthcare, when used to help with clinical documentation. Therefore, detecting ASR errors is a critical first step in preventing further error propagation to downstream applications. To this end, we propose a novel end-to-end approach for ASR error detection using audio-transcript entailment. To the best of our knowledge, we are the first to frame this problem as an end-to-end entailment task between the audio segment and its corresponding transcript segment. Our intuition is that there should be a bidirectional entailment between audio and transcript when there is no recognition error and vice versa. The proposed model utilizes an acoustic encoder and a linguistic encoder to model the speech and transcript respectively. The encoded representations of both modalities are fused to predict the entailment. Since doctor-patient conversations are used in our experiments, a particular emphasis is placed on medical terms. Our proposed model achieves classification error rates (CER) of 26.2% on all transcription errors and 23% on medical errors specifically, leading to improvements upon a strong baseline by 12% and 15.4%, respectively.

**Index Terms:** error-detection, speech recognition, medical conversations, bert, wav2vec2.0, hubert

## 1. Introduction

Electronic Health Record (EHR) documentation results in a high administrative burden for physicians and is a significant cause of physician burnout [1]. Automatic Speech Recognition (ASR) is a promising technology to help physicians with documentation. Various studies have shown that speech recognition for medical dictations increases speed and allows for capturing more comprehensive documentation [2, 3]. Physicians have expressed increased satisfaction using speech recognition for documenting dictations [4]. Beyond dictation solutions, there has also been an uptick in adopting manual scribing solutions to reduce the documentation burden so that physicians can focus more on patient care. More recently, there has been a growing interest in leveraging the ASR transcripts and automating parts of the scribing workflows using spoken conversation understanding technologies such as utterance classification [5], information extraction [6, 7, 8], and conversation summarization [9]. Since speech recognition is upstream of all these dictation and scribing-centered documentation workflows, words mistranscribed by the ASR system [10] can be problematic, leading to users performing time-consuming and laborious corrections [11]. Leaving a mistranscribed word uncorrected may result in unclear documentation, embarrassing errors, and risk to patient safety [11]. Therefore, it is becoming increasingly relevant to detect the transcription errors beforehand to prevent them from going into the downstream applications and workflows.

Table 1: *Mistranscriptions of medical terms*

Source	Transcript
Reference	Keep her on the symbicort for a while.
Hypothesis	Keep her on the civil court for a while.
Reference	Last time I checked did an echocardiogram.
Hypothesis	Last time I checked identical car drive.
Reference	It's September now it's flu season.
Hypothesis	It's September now is who season.

Table 2: *Mistranscriptions of non-medical terms*

Source	Transcript
Reference	Good I'd walk in the doctor's office.
Hypothesis	Could I walk in a box office.
Reference	Probably you won't give a timetable.
Hypothesis	Probably you want to give a timetable.
Reference	It will heal you, go back there a year from now.
Hypothesis	It will heal you, go back the or a year from now.

In recent years, cloud-based speech-to-text services from technology companies such as Google and Amazon have become popular, offering flexibility and convenience. However, while these services exhibit impressive performance for general use-cases, they are not optimal for domain-specific vocabulary [12, 13, 14]. Moreover, these services are provided as black boxes without access to the inner workings of their ASR models, making it impossible to finetune on domain-specific tasks and datasets. Apart from the mismatch in domains, several factors may affect the performance of ASR systems. In particular, the spontaneity of speech and the presence of multiple speakers in medical conversations can be challenging. Furthermore, in an uncontrolled setting such as a hospital, the background noise can also cause recognition errors. In tables 1 and 2, we show domain-specific and generic ASR errors.

In this paper, we propose a novel approach to ASR error detection by utilizing both audio and text modalities. We fuse an acoustic encoder and linguistic encoder into a single model to predict the entailment between the audio and transcript. In the fused model, we freeze the weights of the linguistic encoder, so the acoustic encoder has to learn discriminative features from speech during the training stage to detect ASR errors. Instead of relying on features from the ASR system, our approach encodes the raw speech audio and the hypothesized transcript for error detection. Since our approach does not depend on the inner workings of the ASR system, it can be applied to the outputs of any ASR model. Furthermore, since our approach uses both acoustic and linguistic features, the model can exploit syntactic,

semantic, and acoustic cues. In summary, the primary contribution of this paper is a novel end-to-end approach to ASR error detection based on entailment between audio and transcript segments. To help understand the impact on downstream clinical use-cases, we pay special attention to the medical term transcription errors. To understand the robustness of the proposed model to linguistic variations, we also evaluate our model by simulating ASR errors different from those in the original hypothesis.

## 2. Related work

ASR error detection is a well-established problem and has been studied extensively. The majority of past research can be classified into two major categories. The first set of approaches only uses features generated from the ASR decoder, such as the posterior probability of a word given an acoustic signal. [15, 16, 17]. The second set of approaches focuses on linguistic features from the hypothesized transcription in addition to ASR decoder features [18, 19].

Early error detection research focused on the use of confidence measures to understand the reliability of ASR transcriptions [16]. In [20], the authors propose using a logistic regression model with features extracted from confusion networks as a post-processing step for ASR error detection. Later [15] uses data mining approaches leveraging features from ASR decoder's confidence scores and linguistic parsing. [18] ventures beyond the features from the ASR decoder and hypothesized transcript by using additional features from statistical machine translation, posteriors from an entity detector, and a word boundary detector. In [21], the authors propose an ASR system-independent error detection and classification model based on a variant RNN-based model. More recently, [22] propose an approach to combine ASR confidence scores and transcription hypotheses using Transformer-based models. Most prior work focuses on extensive feature engineering or relies mainly on the hypothesized transcript. In contrast, our approach doesn't require any feature engineering since we encode the raw audio signal and the hypothesized transcript directly. Also, our approach is the first to combine both audio and transcript features in an end-to-end fashion for ASR error detection.

In the medical domain, there has been an increase in the adoption of ASR transcripts of medical conversations to automate clinical documentation [5, 23]. While there has been significant research on ASR error detection in the past, ASR error detection on medical conversations is not explored in depth. In this work, we build on [12] by integrating acoustic features as the source of truth. However, we focus on error detection rather than error correction. Similar to [12] our current approach does not depend on ASR model features other than the hypothesized transcriptions. Therefore our error detection system can be used regardless of the transcription source. In this work, we develop an end-to-end ASR error detection system using both acoustic and linguistic representations to evaluate the performance in the framework of medical conversations. To the best of our knowledge, we are the first to frame this problem as an end-to-end entailment task between the audio segment and its corresponding transcript segment.

## 3. Error Detection Model

The aim of ASR error detection is to mitigate the impact of mis-transcription by flagging the recognition errors automatically. We model the ASR error detection problem as binary entail-

ment prediction task between hypothesized text and speech signal where positive entailment indicates that the speech signal and corresponding hypothesis text imply each other and a negative entailment indicates they contradict each other. Thus the positive entailment implies no error, and the negative entailment implies an error. In our approach, we use an acoustic encoder and a linguistic encoder to encode the speech and corresponding transcript respectively. Further, we fuse the representations of these two modalities and pass them to an entailment classifier to predict the entailment. Using this model, we aim to automatically detect any systematic speech recognition errors, including medical domain-specific terms.

### 3.1. Acoustic Encoders (AE)

We experiment with Wav2vec2.0 [24] and HuBERT [25] separately as acoustic encoders in our model since they have been well studied and verified. They also have the advantage of operating directly on speech signals without requiring any feature engineering.

#### 3.1.1. Wav2vec2.0

Wav2vec2.0 is a framework for self-supervised learning of speech representations. Taking inspiration from masked language modeling, wav2vec2.0 encodes audio speech using convolutional layers into a latent space representation  $Z$  and masks the spans in latent space. The latent representations  $Z$  are then forwarded to a transformer network to obtain a contextual representation  $C$ . The model is then pretrained via solving a contrastive task. Wav2vec2.0 outperforms previous semi-supervised methods simply by fine-tuning on transcribed speech with connectionist temporal classification criterion [24]. In this work, wav2vec2.0 is used to convert audio speech signals to an acoustic representation. We particularly use the contextualized representations  $C$  as our acoustic representation.

#### 3.1.2. HuBERT

HuBERT is a self-supervised speech representation learning approach and follows an architecture similar to Wav2vec2.0. It has state of the art performance either matching or improving upon Wav2vec2.0. The speech representations are learned by pre-training with a masked prediction task. HuBERT uses an offline k-means clustering step and learns the speech representations by alternating between clustering and prediction processes by predicting the proper cluster for the masked audio signal. It relies mainly on the consistency of the unsupervised clustering step rather than the intrinsic quality of the allocated cluster labels [25].

### 3.2. Linguistic Encoder (LE)

We choose BERT [26] as the linguistic encoder in our model. BERT is the language model trained using masked language modeling objective. During pre-training BERT utilizes both the left and right context and tries to predict the masked words. BERT is trained on large amounts of text data and pre-trained BERT has been shown to achieve impressive performance even in cases of limited data for downstream tasks. Previous probing analyses have also shown that pre-trained BERT captures syntactic dependency structures and coreferences with remarkable accuracy [27, 28]. Thus, the priors on BERT can provide a robust semantic and syntactic context in encoding the ASR hypothesis.

### 3.3. Entailment Classifier

Since the speech signal and transcript can be of varying lengths, the representations obtained from the acoustic encoder and linguistic encoder are averaged to a fixed dimensional representation. We also experimented with the self-attention mechanism to obtain fixed-size representations, but found these did not yield better results than simple averaging.

$$\begin{aligned} a_i &= W_a.AE(s_i) + b_a \\ l_i &= W_l.LE(t_i) + b_l \\ c_i &= \text{concat}([a_i, l_i]) \\ e_i &= \text{sigmoid}(W_e c_i + b_e) \quad \forall i = 1 \dots M \end{aligned}$$

Because the representations from acoustic and linguistic encoders are from different modalities, we use a projection layer to bring these modalities to the same representation space. Let  $s_i$  be the speech segment normalized with zero-mean unit-variance and  $t_i$  be the hypothesized transcript segment. We obtain the representations  $a_i$  and  $l_i$  from the projection layer as shown in the above equations. We then concatenate these projected representations to obtain the context vector  $c_i$ . Finally, the context vector is applied to a fully connected layer followed by a sigmoid layer to predict the entailment.

## 4. Experiments

### 4.1. Dataset

Using reference human-written and ASR transcripts of the same speech signal, we automatically label the recognition errors and construct a dataset of speech and ASR transcript segments. We use a fully-consented and deidentified dataset containing 3807 doctor-patient conversations with an average conversation length of 9 minutes. The ASR hypothesis is generated using the Google Speech-to-Text API service<sup>1</sup>. For comparison, we show the error rates computed on the same dataset of hypothesized transcripts in Table 3 from our previous work [12]. We use the sentence tokenizer from Spacy [29] to chunk both the reference and hypothesis transcripts into segments. Google Speech-to-Text API provided word-level timestamps can be used to align reference and hypothesis segments. However, the timestamps in the human-annotated reference transcripts are often off by a few seconds.

**Alignment** Since the timestamps in the reference transcript are not reliable, specific word-level alignment techniques are required to match the output of the ASR system. The alignment is achieved using the Smith-Waterman algorithm, which obtains optimal local alignments between reference and hypothesis transcripts. A lower gap penalty and higher match score is used to favor local alignments when possible. To obtain good alignments, a series of text cleaning steps, casing, stemming, and punctuation stripping are performed before applying the alignment algorithm.

After the alignment process, we perform simple text matching to label ASR errors. Since we focus on downstream clinical documentation applications, we also label ASR errors particularly in medical terms. To annotate medical term-related errors, we identify medical terms in aligned reference and hypothesis transcript segments using the biomedical ontology UMLS [30] and perform string matching.

<sup>1</sup><https://cloud.google.com/speech-to-text>

Table 3: Word error rate and medical word error rates on hypothesized transcripts

WER	BLEU	Medical WER		
		P	R	F1
41.0	52.1	90.1	58.4	67.3

**UMLS** UMLS is a biomedical ontology with comprehensive coverage of medical terms often used in clinical information extraction systems. The medical terms are manually categorized and organized into semantic groups such as body system, clinical drug, etc. For evaluation purposes, we consider five major semantic groups: Chemicals & Drugs, Disorders, Procedures, Anatomy, and Physiology. To identify medical terms in each transcript segment, we extract n-grams from the transcript and apply a dictionary matching algorithm to match the medical terms in the UMLS database. With the diversity of clinical specialties in the dataset, we obtain segments with various medical terms.

In this study, we train our error detection models at the segment level rather than word level for two reasons. First, interpolating word-level timestamps for ASR errors may not be reliable, and second, transcript segments can provide valuable contextual cues regarding recognition error. To allow for a fair comparison between the error detection task and the domain-specific error detection task, we balance the datasets for each type of error and finally obtain nearly 100 hours of speech data with 65k speech segments and corresponding hypothesized transcript segments and labels. The minimum and maximum length of audio segments are chosen to be between 1 and 30 seconds. We set aside 80% of the data for training purposes, and the remaining 20% data is split equally for validation and testing. To evaluate the robustness of the model to linguistic variations, we construct a simulated error test set by replacing mistranscribed words in the test set using the ASR error simulation approach proposed in [31].

We use base models pre-trained on the librispeech 960-hour dataset for acoustic encoders. For the linguistic encoder, we use the uncased base BERT model. While training, we freeze the weights of the linguistic encoder and latent representation extractor network in the acoustic encoder. The contextual Transformer network in the acoustic encoder is fine-tuned while training the entailment classifier. The intuition is that the BERT model provides strong syntactic and semantic context. At the same time, the acoustic encoder has to learn representations specific to the error detection task to entail the linguistic representation.

We construct a baseline model by replacing the acoustic encoder in our proposed model with ASR confidence scores as features (CS + BERT).

We use a learning rate of 1e-3 to train all models.

## 5. Results and Discussion

For evaluation, we choose classification error rate (CER) along with precision, recall, and F1 metrics. We evaluate our model both qualitatively and quantitatively. To understand the impact on downstream clinical documentation applications, we also focus on major semantic groups of medical terms. In addition, we also evaluate the model for robustness by simulating ASR errors.

Table 4: Evaluation of error detection entailment task using confidence scores (CS) and acoustic features

Model	Precision	Recall	F1	CER
Medical Errors				
CS + BERT	72.14	80.54	76.11	27.2
wav2vec2.0 + BERT	<b>75.57</b>	84.18	79.63	23.3
HuBERT + BERT	74.00	<b>88.24</b>	<b>80.46</b>	<b>23.0</b>
All Errors				
CS + BERT	70.08	70.65	70.36	29.8
wav2vec2.0 + BERT	<b>75.51</b>	70.44	72.89	<b>26.2</b>
HuBERT + BERT	74.28	<b>72.29</b>	<b>73.27</b>	26.4

Table 5: Evaluation of best performing HuBERT + BERT model on simulated ASR errors

Error Type	Precision	Recall	F1	CER
Medical Errors	71.99	91.86	80.72	23.5
All Errors	77.86	81.01	79.41	21.0

**Error type detection** We present our results on both medical and all error detection tasks using different acoustic encoders and the baseline confidence score features (CS + BERT) in Table 4. When compared with the baseline model using confidence scores, we see an improvement across all metrics showing the difficulty of the task and effectiveness of using linguistic and acoustic representations together. In both the domain-specific and general error detection tasks, the models using the HuBERT acoustic encoder have improved recall performance, while those using the Wav2vec2.0 acoustic encoder have slightly better precision. Overall, we observe that the domain-specific medical term classification error rate is lower compared to the generic error detection task. We hypothesize that medical terms have specific linguistic and acoustic patterns and hence are easier to recognize than general errors which are more nuanced and can happen in an unpredictable manner.

**Semantic Groups and Medical Terms** We measure the medical error detection metrics specifically for the five major semantic groups. To compare errors based on semantic groups, we evaluate our chosen metrics on segments containing medical terms of each corresponding semantic group. In Table 6, we notice that the ASR error detection performs best on the Chemicals & Drugs semantic group. This semantic group is the most frequent error type and mainly captures medications. Apart from having slightly better performance on Chemicals & Drugs semantic group, the metrics are very consistent across the remaining semantic groups. In Table 7, we specifically focus on the top five most frequent and top five least frequent medical terms. For the least frequent medical terms, we randomly sample five medical terms at a cut off frequency of ten. We notice a relatively higher performance on frequently occurring medical terms, although the performance can be varied in least frequent terms as the number of samples is low.

**Robustness** In Table 5, we evaluate the best performing HuBERT + BERT model on the simulated error dataset. On the medical error detection task, the model performs comparably on both original and simulated error datasets. On the other hand,

there’s a significant improvement of performance on all errors with the simulated dataset suggesting that simulated errors are much easier to detect than the actual ASR hypothesized errors. Overall the metrics suggest that the model can also work well in case of different linguistic variations in ASR errors.

Table 6: Evaluation of medical error detection based on 5 major semantic groups

Semantic Group	Precision	Recall	F1	CER
Chemicals & Drugs	<b>92.21</b>	90.11	<b>91.15</b>	<b>16.0</b>
Disorders	86.75	91.52	89.07	19.0
Anatomy	86.20	90.61	88.35	20.4
Procedures	85.42	91.23	88.23	20.4
Physiology	86.19	<b>91.59</b>	88.80	19.4

Table 7: Evaluation based on most and least frequent medical term mistranscriptions

Medical term	Precision	Recall	F1	CER
Most frequent				
heart	85.42	93.51	89.28	18.8
blood pressure	83.09	91.47	87.08	21.4
coumadin	<b>92.59</b>	93.47	93.02	13.0
cholesterol	86.95	97.90	92.10	14.3
atrial fibrillation	90.52	<b>98.85</b>	<b>94.50</b>	<b>10.4</b>
Least frequent				
chest pain	75.00	80.00	77.41	35.0
stress	90.90	66.66	76.92	35.3
angiogram	<b>93.33</b>	73.68	82.35	30.0
A1c	88.00	<b>95.65</b>	<b>91.66</b>	<b>15.3</b>
vitamin	91.30	87.50	89.36	19.2

## 6. Conclusion

In this work, we introduced a novel approach for automatic ASR error detection in an end-to-end manner. Our approach jointly encodes speech features from audio and linguistic features from transcription hypotheses into a contextualized representation. The entailment classifier detects the ASR errors by modeling features at semantic, syntactic, and acoustic levels. Since our method doesn’t rely on any ASR model, it is viable for many domain-specific applications and is easy to implement as a post-processing step. We also evaluated our error detection approach on medical terms and a simulated ASR error dataset for robustness. With good recall on major semantic groups, including disorders and procedures, the proposed model can be a potential candidate for detecting errors. In future work, we would like to investigate pre-training the acoustic feature extractor with domain-specific data. Further, we plan to extend our work for ASR error correction by jointly training a language decoder and evaluate under different acoustic conditions.

## 7. Acknowledgements

We would like to acknowledge Shruti Palaskar for helpful discussions during the preparation of this work and, Elisa Ferracane and Sai Prabhakar Pandi Selvaraj for valuable feedback on the manuscript.

## 8. References

- [1] E. Gesner, P. Gazarian, and P. Dykes, "The burden and burnout in documenting patient care: an integrative literature review," *MED-INFO 2019: Health and Wellbeing e-Networks for All*, pp. 1194–1198, 2019.
- [2] M. P. Koivikko, T. Kauppinen, and J. Ahovuo, "Improvement of report workflow and productivity using speech recognition—a follow-up study," *Journal of Digital Imaging*, vol. 21, no. 4, pp. 378–382, 2008.
- [3] M. Vogel, W. Kaisers, R. Wassmuth, E. Mayatepek *et al.*, "Analysis of documentation speed using web-based medical speech recognition technology: randomized controlled trial," *Journal of Medical Internet Research*, vol. 17, no. 11, p. e5072, 2015.
- [4] F. R. Goss, S. V. Blackley, C. A. Ortega, L. T. Kowalski, A. B. Landman, C.-T. Lin, M. Meteer, S. Bakes, S. C. Gradwohl, D. W. Bates *et al.*, "A clinician survey of using speech recognition for clinical documentation in the electronic health record," *International Journal of Medical Informatics*, vol. 130, p. 103938, 2019.
- [5] B. Schloss and S. Konam, "Towards an automated SOAP note: classifying utterances from medical conversations," in *Machine Learning for Healthcare Conference*. PMLR, 2020, pp. 610–631.
- [6] S. P. Selvaraj and S. Konam, "Medication regimen extraction from clinical conversations," *CoRR*, vol. abs/1912.04961, 2019. [Online]. Available: <http://arxiv.org/abs/1912.04961>
- [7] D. Patel, S. Konam, and S. P. Selvaraj, "Weakly supervised medication regimen extraction from medical conversations," *arXiv preprint arXiv:2010.05317*, 2020.
- [8] N. V. Meripo and S. Konam, "Extracting appointment spans from medical conversations," in *Proceedings of the Second Workshop on Natural Language Processing for Medical Conversations*. Online: Association for Computational Linguistics, Jun. 2021, pp. 41–46. [Online]. Available: <https://aclanthology.org/2021.nlpmc-1.6>
- [9] K. Krishna, S. Khosla, J. P. Bigham, and Z. C. Lipton, "Generating SOAP notes from doctor-patient conversations," *CoRR*, vol. abs/2005.01795, 2020. [Online]. Available: <https://arxiv.org/abs/2005.01795>
- [10] F. R. Goss, L. Zhou, and S. G. Weiner, "Incidence of speech recognition errors in the emergency department," *International journal of medical informatics*, vol. 93, pp. 70–73, 2016.
- [11] Y. A. Kumah-Crystal, C. J. Pirtle, H. M. Whyte, E. S. Goode, S. H. Anders, and C. U. Lehmann, "Electronic health record interactions through voice: a review," *Applied Clinical Informatics*, vol. 9, no. 03, pp. 541–552, 2018.
- [12] A. Mani, S. Palaskar, N. V. Meripo, S. Konam, and F. Metzger, "ASR error correction and domain adaptation using machine translation," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6344–6348.
- [13] A. Mani, S. Palaskar, and S. Konam, "Towards understanding ASR error correction for medical conversations," in *Proceedings of the First Workshop on Natural Language Processing for Medical Conversations*. Online: Association for Computational Linguistics, Jul. 2020. [Online]. Available: <https://aclanthology.org/2020.nlpmc-1.2>
- [14] K. Khandelwal, P. Jyothi, A. Awasthi, and S. Sarawagi, "Black-Box Adaptation of ASR for Accented Speech," in *Proc. Inter-speech 2020*, 2020, pp. 1281–1285.
- [15] R. Errattahi, S. Deena, A. El Hannani, H. Ouahmane, and T. Hain, "Improving ASR error detection with rnnlm adaptation," in *2018 IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 190–196.
- [16] H. Jiang, "Confidence measures for speech recognition: A survey," *Speech Communication*, vol. 45, no. 4, pp. 455–470, 2005.
- [17] R. Errattahi, A. El Hannani, and H. Ouahmane, "Automatic speech recognition errors detection and correction: A review," *Procedia Computer Science*, vol. 128, pp. 32–37, 2018.
- [18] W. Chen, S. Ananthakrishnan, R. Kumar, R. Prasad, and P. Natarajan, "ASR error detection in a conversational spoken language translation system," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 7418–7422.
- [19] T. Pellegrini and I. Trancoso, "Improving ASR error detection with non-decoder based features," in *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- [20] A. Allauzen, "Error detection in confusion network," in *Eighth Annual Conference of the International Speech Communication Association*, 2007.
- [21] M. Topaz, A. Schaffer, K. Lai, T. Korach, J. Einbinder, and L. Zhou, "Medical malpractice trends: Errors in automated speech recognition," *Journal of Medical Systems*, vol. 42, 08 2018.
- [22] Z. Gekhman, D. Zverinski, J. Mallinson, and G. Beryozkin, "Red-ace: Robust error detection for ASR using confidence embeddings," *arXiv preprint arXiv:2203.07172*, 2022.
- [23] K. Krishna, S. Khosla, J. P. Bigham, and Z. C. Lipton, "Generating SOAP notes from doctor-patient conversations using modular summarization techniques," *arXiv preprint arXiv:2005.01795*, 2020.
- [24] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [25] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [26] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [27] G. Penha and C. Hauff, "What does BERT know about books, movies and music? probing BERT for conversational recommendation," in *Fourteenth ACM Conference on Recommender Systems*, 2020, pp. 388–397.
- [28] Z. Wu, Y. Chen, B. Kao, and Q. Liu, "Perturbed masking: Parameter-free probing for analyzing and interpreting BERT," *arXiv preprint arXiv:2004.14786*, 2020.
- [29] M. Honnibal and I. Montani, "spaCy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing," *To appear*, vol. 7, no. 1, pp. 411–420, 2017.
- [30] P. L. Schuyler, W. T. Hole, M. S. Tuttle, and D. D. Sherertz, "The UMLS metathesaurus: representing different views of biomedical concepts," *Bulletin of the Medical Library Association*, vol. 81, no. 2, p. 217, 1993.
- [31] R. Voleti, J. M. Liss, and V. Berisha, "Investigating the effects of word substitution errors on sentence embeddings," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7315–7319.