

Generative Spoken Dialogue Language Modeling

Tu Anh Nguyen^{◊,†}, Eugene Kharitonov^{◊*}, Jade Copet[◊], Yossi Adi[§],
Wei-Ning Hsu[¶], Ali Elkahky[¶], Paden Tomasello[¶], Robin Algayres^{◊*},
Benoît Sagot[†], Abdelrahman Mohamed^{◊*}, Emmanuel Dupoux^{◊,†}

[◊]Meta AI Research, France [†]Inria, Paris, France [§]Meta AI Research, Israël

[¶]Meta AI Research, United States [‡]EHESS, ENS-PSL, CNRS, Paris, France

{ntuanh, abdo, dpx}@meta.com

Abstract

We introduce dGSLM, the first “textless” model able to generate audio samples of naturalistic spoken dialogues. It uses recent work on unsupervised spoken unit discovery coupled with a dual-tower transformer architecture with cross-attention trained on 2000 hours of two-channel raw conversational audio (Fisher dataset) without any text or labels. We show that our model is able to generate speech, laughter, and other paralinguistic signals in the two channels simultaneously and reproduces more naturalistic and fluid turn taking compared to a text-based cascaded model.^{1,2}

1 Introduction

In natural conversations, speakers spontaneously coordinate who is currently speaking and when the other person will speak next. As a result, conversations end up being a fluent succession of *turns* without much overlapping speech or long stretches of silence. Of course, silences and overlaps also occur naturally and they carry significant information which is interpreted within the conversation setting. For instance, when overlapping speech occurs it often contains content-neutral verbal information (e.g., “hmm”, “yeah”) or non-verbal vocalization (e.g., laughter), used to convey a listening attitude (back-channelling) (Yngve, 1970; Schegloff, 1982). Short silences between turns do occur and show both cross-cultural variations and universal dependence on dialogue related variables, for instance, straight and positive answers to questions are typically faster than non-responses or negative responses (Stivers et al., 2009).

*Work done while at Meta.

¹Generation samples can be found at <https://speechbot.github.io/dgslm>.

²Code and pre-trained models will be made available at https://github.com/facebookresearch/fairseq/tree/main/examples/textless_nlp/dgslm.

All of this *turn-taking* coordination is natural to humans, and starts to be learned at an early age by infants (Nguyen et al., 2022). In contrast, it remains a challenging area of research in human/machine interactions (Skantze, 2021). One of the reason is that much of the research into natural dialogue modeling is taking place with text-based interfaces. Here, the coordination problem is primarily focused on semantic coherence and appropriateness of the artificial agent in interaction with a human (see Ni et al., 2021, for a review). The turn-taking problem itself is being taken care of by an artificially imposed walkie talkie arrangement; each agent is writing in turn and signalling the end of its turn by pressing carriage return.

Within speech-based systems, it is very similar, as current spoken assistants like Siri or Alexa are triggered by a predetermined wake word, and wait for the end of an utterance followed by sufficient silence to segment the turns of the human interlocutor. This may give rise to slow and unnatural conversations. In fact, in human-human conversation, *pauses* within speaker turns tend to be on average longer than *gaps* between speaker turns (Brady, 1968; Ten Bosch et al., 2005; Heldner and Edlund, 2010), indicating that silence may not be the main cue for humans to switch turns. Because most speech-based systems are based on Automatic Speech Recognition (ASR), and many significant aspects of speech like prosody and nonverbals are typically not annotated in naturalistic speech dialogues, current dialogue systems have been struggling with generating naturalistic dialogue.

Here we capitalize on recent progress in self-supervised learning and textless speech processing (Borgholt et al., 2022; Borsos et al., 2022; Lakhota et al., 2021) to investigate the possibility to directly train a spoken dialogue model from raw audio, bypassing the need for text or

ASR. Briefly, we build on self-supervised discrete speech representations models, which we train on spontaneous conversations with each speaker having his or her own audio channel. After training, the speech units come to represent not only verbal but also nonverbal materials. We can now encode a conversation between two interlocutors as two parallel streams of discrete tokens. We then introduce a novel dual-tower transformer architecture, where each channel is processed by one “tower” of the model that learns via an autoregressive loss, but the two towers also communicate via cross-attention in their hidden units. This cross-attention is critical for the correct synchronization of the two channels and result in a naturalistic distribution of turns, overlap and pauses. While this system is not trained on enough data to capture deep syntactic and semantic aspects of dialogue, and indeed scores below a text-based cascaded ASR+LM+TTS model on semantic content, it does capture better surface characteristics of chitchat in mimicking accurately turn-taking and backchanneling. This can be seen as a proof of principle that previously difficult to capture aspects of spontaneous conversations can be captured with minimally modified language modeling techniques. Finally, our model opens up new possibilities to create more natural naturalistic human-machine dialogue systems in the future.

2 Related Work

Unsupervised Spoken Language Modeling.

Recently, great advances have been achieved in the area of representation learning from raw audio. Models trained with either autoencoder objectives (Ondel et al., 2016; van den Oord et al., 2017) or masked objectives (CPC: van den Oord et al., 2018; APC: Chung and Glass, 2020; wav2vec 2.0: Baevski et al., 2020; HuBERT: Hsu et al., 2021a; MockingJay: Liu et al., 2020) from raw speech can learn audio representation that can be used for a variety of downstream tasks (Yang et al., 2021), see Borgholt et al. (2022) for a review.

Most of these models build a codebook of discrete units, either as latent representation or as targets. The discrete representation can in turn be fed to a standard autoregressive language model, which can then be sampled to generate new speech sequences (Lakhotia et al., 2021; Dieleman et al., 2021). An interesting aspect of this procedure




Figure 1: General Schema for dGSLM: A discrete encoder (HuBERT+kmeans) turns each channel of a dialogue into a string of discrete units (c_1, \dots, c_N). A Dialogue Language Model (DLM) is trained to autoregressively produce units that are turned into waveforms using a decoder (HifiGAN).

is that it can capture aspects of speech that are typically not available in written transcriptions and can therefore model prosody and intonation (Kharitonov et al., 2021), or non verbal vocalizations typical of emotional speech (Kreuk et al., 2021). Up to now, however, no such model has been applied to multi-party conversational speech.

Dialogue Generation. Since the early work on end-to-end neural dialogue generation (Vinyals and Le, 2015; Li et al., 2015; Serban et al., 2016), empowered by scalable methods for language representation (Radford et al., 2018; Lewis et al., 2019), there has been enormous progress in the area of dialogue generation (Roller et al., 2020;

Zhang et al., 2019; Adiwardana et al., 2020). More recent research focused on utilizing retrieval augmented generation methods (Lewis et al., 2020) for long-context, multi-session conversations (Xu et al., 2021), and grounding responses on fresh information from the internet (Komeili et al., 2021; Shuster et al., 2022). However, all the progress in this research work centered around text dialogues leaving out non-lexical information (Schuller et al., 2013; Ang et al., 2002) in human-human dialogues, for example, emotion, pauses, laughter, hesitation, and interruption. Our work builds on end-to-end techniques while taking a speech-first approach to address this shortcoming, where prompts and generated sequences are represented as self-supervised discrete speech representations (Lakhotia et al., 2021). As a result, the capacity of our models is constrained by the amount of publicly available speech dialogues; for example, the LDC English Fisher dialogues corpus (Cieri et al., 2004) contains roughly 12M words compared to tens of billions of words in the case of text-based dialogue systems. There have been recent calls for large-scale end-to-end benchmarks and datasets with spoken input to fill this gap (Faruqui and Hakkani-Tür, 2021).

Turn-taking Modeling. Decades-long research on conversation analysis (Duncan, 1972; Sacks et al., 1974; Schegloff, 2000; Gravano and Hirschberg, 2011; Levinson and Torreira, 2015; Ward, 2019) has shown that human turn-taking relies on a variety of complex signals, or cues, including prosodic cues, linguistic cues, and even non-verbal cues such as gaze or gestures, making turn-taking modeling a challenging problem. Simple turn-taking models using finite-state machines have been proposed to predict the distribution and durations of turn-taking events (Cassell et al., 2001; Thórisson, 2002; Raux and Eskenazi, 2009). More recently, more sophisticated machine learning-based models of turn-taking have been introduced (Meena et al., 2014; Skantze, 2017; Roddy et al., 2018; Masumura et al., 2018). These models used multi-modal features including simple linguistic features and prosodic features extracted from the speech to predict turn shifts. Most recently, Ekstedt and Skantze (2020) has shown the possibility of turn-taking prediction in spoken dialogue using only linguistics features (text input). We use these definitions of turn-taking events to analyse the output of our models.

3 Approach

Our approach is based on the availability of a dataset constructed along the Fisher Telephone conversation collection protocol (Cieri et al., 2004) where each conversation involves two speakers, and each speaker is recorded in a separate audio channel while having a very casual conversation. We follow the textless generative spoken language modeling pipeline of Lakhotia et al. (2021), which decomposes the problem of speech generation into three components: a Speech-to-Units encoder, a Units-to-Units language model, and a Units-to-Speech decoder. For the encoder we adopt HuBERT, (Hsu et al., 2021a) followed by k -means clustering; for the decoder network we use a modified HiFi-GAN neural vocoder (Kong et al., 2020), similarly to Polyak et al. (2021). These models are trained on single channel data from the Fisher dataset and applied to each channel separately, which do not model cross-channel interactions. For the language model, we introduce our new Dialogue Transformer Language Model, or DLM. Figure 1 presents an overview of our system. The following sections (Sections 3.1–3.3) will present at a high level each component of our model and review the turn-taking terminology in this study (Section 3.4).

3.1 Discrete Phonetic Representation

Conversational speech contains casual expressions (filler words like ‘hmm’) and a variety of non verbal sounds (e.g., laughter) that do not appear in formal or read speech. We therefore train a HuBERT model (Hsu et al., 2021a) directly on our conversation dataset in order to obtain domain-appropriate phonetic representation (see Appendix Section A for an analysis). Specifically, it is trained on the collection of voice segments extracted of all speakers in the dataset. The discrete units are then obtained by clustering the representation of the HuBERT model using the k -means algorithm. At inference time, the two-channel speech waveform is encoded channel-wise into two time-aligned streams of discrete units.

3.2 Waveform Generation

For the waveform generation, we used the discrete unit-based HiFi-GAN vocoder from Polyak et al. (2021) trained on a small subset of high-quality single-channel voice segments of our conversation

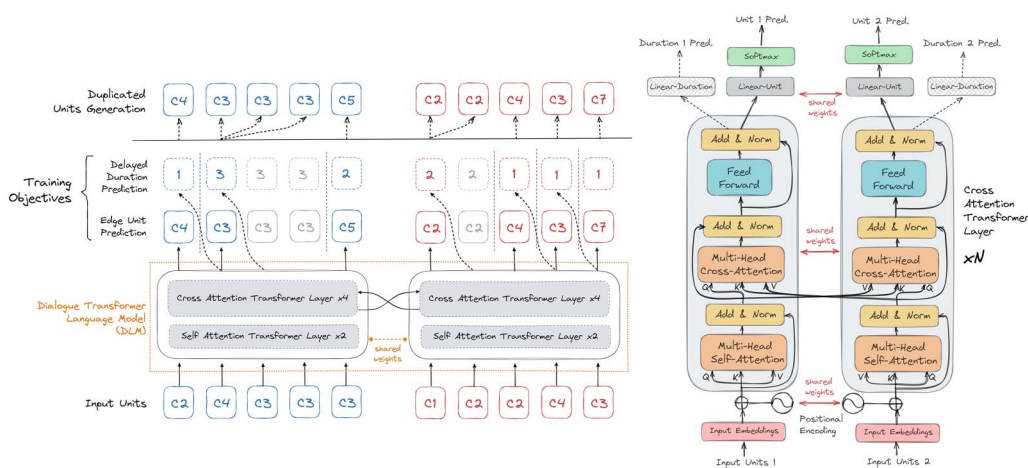


Figure 2: Illustration of the Dialogue Transformer Language Model (DLM). Left: DLM Training Objectives. During training, the loss is applied only to edge units and their durations. During generation, the model duplicates the units with the corresponding predicted durations. Right: The Cross-Attention Transformer Layer Architecture.

dataset, using discrete units obtained from the HuBERT model and 1-hot speaker information from the dataset. During generation, we generate each channel of discrete units with one different speaker, and combine the audio generated from the two channels. Voices for the waveform generation are chosen from the speakers in the HifiGAN training set.

3.3 Dialogue Transformer Language Model

We introduce our Dialogue Transformer Language Model (DLM), which is a two-tower transformer network with *Cross-Attention* and shared weights trained with *Edge Unit Prediction* and *Delayed Duration Prediction* objectives. The model is illustrated in Figure 2 and its components will be detailed below, and we will perform ablations to test for the effects of each of these components.

We will also compare the two-tower model with a simpler single-tower model with dual inputs. This last model is inspired by previous work in multi-stream language model (Kharitonov et al., 2021). It consists of a single transformer, with two embedding heads in the input and two softmax heads in the output. This model combines very early the two speaker channels at the embedding layer and models them jointly, only to separate them again in the last layer. We call this model MS-TLM (Multi-Stream Transformer Language Model).

Cross-Attention Transformer Layer. When modeling separate channels of dialogue, we would like the LM to not only get information from the

history of each channel itself, but also have information from other channels as well. As a result, we add an additional Multi-Head Cross-Attention block after the Multi-Head Self-Attention block to share information between different channels (cf. Figure 2 right). We train a single Transformer model that we clone into the two towers with shared weights, which allows the model to be speaker-independent without having to do permutation invariant training.

Edge Unit Prediction. Previous work (Kharitonov et al., 2021) disentangles the content modeling problem from the duration modeling problem by training the language model on deduplicated discrete units and the corresponding unit durations with different objectives. However, in our setting, units from different channels are time-aligned and there would be no easy way to keep the alignment if we were to deduplicate each input stream. On the other hand, training a language model on duplicated units is more difficult as content and duration information are entangled and learned simultaneously, resulting in a poor modeling performance. From this point of view, we introduce an edge unit prediction objective, which forces the model to predict the next unit only if it is different from the current one (i.e., edge unit). We use cross-entropy loss for this objective, and the edge unit prediction loss is then defined as:

$$\mathcal{L}_{EU} = \sum_{c=1}^2 \sum_{\substack{t \\ u_t^{(c)} \neq u_{t-1}^{(c)}}} \log p(u_t^{(c)} | u_{1:t-1}^{(1,2)}; \theta),$$

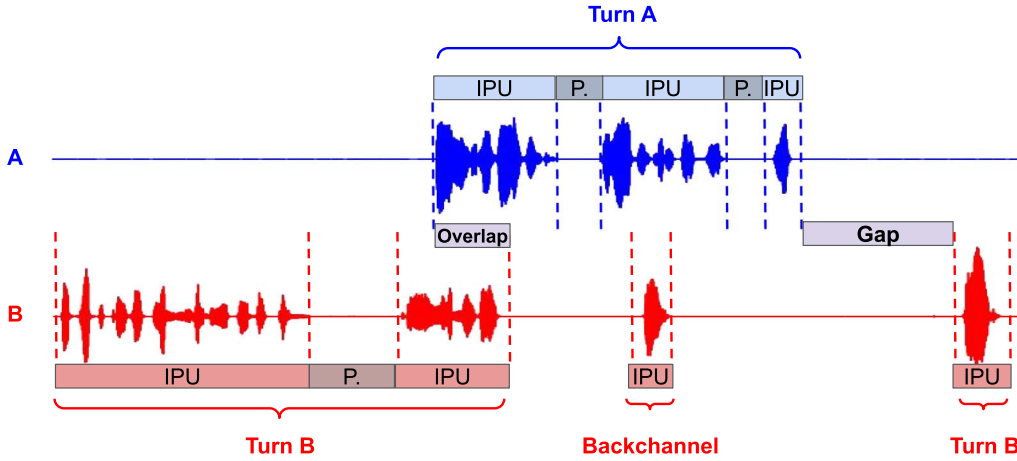


Figure 3: Illustration of turn-taking events: IPU (Interpausal Unit), Turn (for speaker A and Speaker B, resp.), P. (within-speaker Pause), Gap, Overlap and Backchannel.

where $u_t^{(c)}$ represents the discrete unit from channel c at time t and θ denotes the model parameters.

Delayed Duration Prediction. Besides the unit prediction objective, DLM models the duration of the edge units with a duration prediction objective. As unit durations are highly varied, we output a continuous duration prediction and employ an L1 loss. Due to the high correlation between the duration and the unit itself, we follow Kharitonov et al., (2021) and perform a delayed unit duration prediction, which predicts the duration of an edge unit at time t given the first $t - 1 + \Delta$ units, where Δ is a delay factor ($\Delta \geq 0$). The delayed duration prediction loss is then defined as:

$$\mathcal{L}_{ED} = \sum_{c=1}^2 \sum_{\substack{t \\ u_t^{(c)} \neq u_{t-1}^{(c)}}} \left| d_t^{(c)} - \hat{d}_t^{(c)} \left(u_{1:t-1+\Delta}^{(1,2)}; \theta \right) \right|,$$

where $d_t^{(c)}$ represents the target duration (number of repetitions) of the edge unit $u_t^{(c)}$ and $\hat{d}_t^{(c)}$ is the continuous duration prediction of the DLM model.

Training Objective. The training loss of DLM is the sum of the edge unit prediction loss and the delayed duration prediction loss:

$$\mathcal{L}_{DLM} = \mathcal{L}_{EU} + \mathcal{L}_{ED}. \quad (1)$$

Model Inference for Generation. For generation, we autoregressively generate edge units and the corresponding durations in both channels. Even though the loss is applied only at the edge

units, the model may generate spurious and inconsistent data at other non-edge time steps. We give precedence to the predicted duration associated with the first edge unit predicted in each channel and overwrite the network output with this edge units for the corresponding number of steps. It is this overwritten content which is used as input to the network till the next edge unit. For example, if we predict a unit $u_t^{(c)}$ at time t and the corresponding duration $d_t^{(c)}$ at time $t + 1$, we replace the next $d_t^{(c)}$ units of channel c by $u_t^{(c)}$ and only alter the unit at time $t + d_t^{(c)}$. The duration prediction is rounded during generation.

3.4 Definitions of Turn-taking Metrics

Because our model generates two audio channels in parallel, it is possible to use simple Voice Activity Detection (VAD) tools on the output to derive turn-taking metrics. Following Figure 3, we define an *Inter-Pausal Unit* (IPU) as continuous stretch of speech in one speaker’s channel, delimited by a VAD silence of more than 200 ms on both side. We define *silence* as sections of the recording with no voice signals on either channel and *overlap* as sections where there are voice signals on both channels. Silences can be subdivided into *gaps* (when it occurs between two IPUs by distinct speakers) and *pauses* (when they occur for the same speaker). Successive IPUs by the same speaker separated by a pause are regrouped into a *turn*. Overlap could also theoretically be subdivided into *backchannel* (when it is rather short IPU contained within an IPU of the other speaker) and *interruption* (when it starts within

an IPU of the other channel and continues after its end), but the exact definition is dependant on high-level linguistic features, which we will not attempt to extract here. In our analysis, we will therefore tally the distribution of duration of IPUs, gaps, pauses and overlaps in the training corpus and in generated dialogues of our various models.

3.5 Cascaded Dialogue Baseline System

We compare our textless-based dialogue models with a traditional cascaded dialogue system which consists of an ASR model, followed by a text-based language model and a Text-To-Speech (TTS) module. We first transcribe each channel of the dialogue with the ASR model, we then combine the transcribed text into a turn-based conversation³, we ignore any turns that are completely contained inside an other turn. We train a Transformer Language Model on these conversations and we finally employ a TTS module to synthesize the generated text into a turn-based conversation.

4 Experimental Setup

4.1 Training Set

We use in this work the Fisher Dataset (Cieri et al., 2004), a conversation corpus consisting of more than 16,000 English telephone conversations averaging 10 minutes in duration and focusing on various topics. The audio was recorded separately in two channels resulting in 2000 hours of transcribed speech.⁴

For the training of HuBERT and HifiGAN models, we follow the preprocessing steps of Kuchaiev et al. (2019)⁵ to obtain a collection of single-channel voice segments of the Fisher dataset. The segments vary mostly from 10–15 seconds, with a total duration of about 1800 hours. We divide the Fisher dataset into train/valid/test sets with a 98/1/1 split (different speakers in each split).

³example: <A> hi hi how you doing <A> great good good my name is marine.

⁴The transcription was done using the Quick Transcription specification (Cieri et al., 2004), resulting in some inaccuracies and untranscribed portions. Here, we only used the transcriptions to obtain speech segments containing vocal activity to train the HifiGan and HuBERT model. The DLM was trained on the unsegmented raw data.

⁵https://gitlab.nrp-nautilus.io/ar-noc/nemo/-/blob/master/scripts/process.fisher_data.py.

4.2 Model Training

We train a HuBERT Base model (Hsu et al., 2021a) from raw audio. The encoder contains seven 512-channel CNN layers with strides [5,2,2,2,2,2,2] and kernel widths [10,3,3,3,3,2,2], converting the signal rate from 16000 samples/sec down to 50 frames/sec. It is followed by 12 Transformer blocks. The model is trained with a masked objective for 3 iterations following the same recipe as in (Hsu et al., 2021a). The model alternates between feature extraction/quantization and masked-prediction training in each iteration. We used the k -means algorithm with codebook sizes of 100, 500, and 500 to quantize the MFCC features, the 6th transformer layer features, and the 9th transformer layer features for the three HuBERT training iterations. After training, we quantize the final transformer layer features into 500 units for the DLM training. We choose a large codebook size of 500 to model various kinds of vocalizations beyond broad phonetic classes. Following Hsu et al. (2021a), we use 250k training updates in the first iteration and 400k model updates in subsequent training iterations using 32 V100 32GB GPUs. As the transformer does not change the input frame rate, the encoded discrete units have a frame rate of 50 units per second (one every 20 ms). We show in Table A1 that our HuBERT model trained on the Fisher dataset learns better phonetic information suitable for conversations than the publicly available HuBERT model trained on audiobooks (Hsu et al., 2021a).

We train the HifiGAN model on a small subset of the Fisher dataset segments consisting of 120 speakers with 10 minutes each. These speakers were selected to be of high intelligibility using the average perplexity of a phone recognizer trained on the clean Librispeech 100h training subset (Rivière and Dupoux, 2021). The model is trained to generate the audio waveform given HuBERT units of a segment and a speaker embedding vector.

For the DLM models, we use a transformer model consisting of 6 layers, with 8 attention heads per layer, and an embedding size of 512. When cross-attention is used, it is added to the top 4 transformer layers. We show in Table A2 the effect of number of cross-attention layers on language modeling metrics. We train the DLM model on the parallel unit streams encoded from 2000 hours of stereo audio, each sample contains




Figure 4: Distributions of durations of turn-taking events in prompted continuations across models, compared to the prompts’ continuation ground truth segments (see models ids in Table 1). The green line and the red triangle represent the mean and the median of the events respectively.

up to 6144 unit pairs, an equivalent of 123 seconds. The models are trained on a total of 32 V100 32GB GPUs, with a batch size of 370 seconds of audio per GPU for a total number of 250k steps. We used an Adam optimizer (Kingma and Ba, 2015) with a max learning rate of 5×10^{-4} . The implementation of the DLM model is done using the fairseq (Ott et al., 2019) toolkit. It took us 66 hours on average to train 100k steps of DLM models without edge unit prediction, and 95 hours with additional edge unit prediction objective.

We also train a Multi-Stream Transformer Language Model (MS-TLM, Kharitonov et al., 2021), a single transformer model taking two streams of units as input and autoregressively predict the next units in both streams. It is a standard Transformer Language Model, with 6 layers, 8 attention heads per layer and an embedding size of 512, with the difference that the embedding layer concatenates the two embeddings of the two parallel units, and the output layer produces two softmax layers to predict the next units in both streams. We train the MS-TLM model similarly to the DLM models as previously mentioned. Training 100k steps of MS-TLM model took us 40 hours.

For the cascaded system, we use a pre-trained ASR model⁶ to decode the Fisher dataset. We then train a standard 6-layer Transformer Language

⁶We use the robust wav2vec2-large model fine-tuned on Switchboard dataset (Hsu et al., 2021a). For decoding, we use the 4-gram KenLM language model trained on Switchboard dataset.

Model on the turn-based conversations obtained from the ASR. We pre-process the text using a byte pair encoding (BPE, Sennrich et al., 2016) with 20k iterations and limit each sample to have 512 tokens. We trained the language model for 100k steps on 32 V100 32GB GPUs with a batch size of 2048 tokens per GPU. We use the same optimizer as for other models. Finally, we use the Google TTS API to synthesize generated conversations, with two different voices indicating two different speakers.

4.3 Evaluation Metrics

This section presents the evaluation metrics used to assess our dialogue models on two dimensions: Training and Generation.

4.3.1 Training Metrics

These metrics evaluate the dialogue modeling performance in each channel separately using metrics close to the training loss. They are computed by encoding files from the development set and extracting statistics on the predicted outputs at each time steps. They are used to compare the different versions of the DLMs and therefore not applied to the cascaded model.

Edge Unit Prediction. We report the Negative Log-Likelihood (NLL), or Cross Entropy loss when predicting edge units. We also compute the Prediction Accuracy.

Id	CA	EP	DP	Δ	Edge Unit		Duration	
					NLL \downarrow	Acc \uparrow	MAE \downarrow	Acc \uparrow
MS-TLM								
0	—	—	—	—	3.05	34.14	—	—
DLM								
1	×	×	×	—	3.07	34.13	—	—
2	✓	×	×	—	2.95	35.68	—	—
3	✓	✓	×	—	2.49	48.36	—	—
4	✓	✓	✓	0	2.26	54.09	1.47	51.90
5	✓	✓	✓	1	2.25	54.27	1.23	58.18

Table 1: Training Metrics across the DLM models that differ in Cross-Attention Layer (CA), Edge Unit Prediction (EP), Duration Prediction (DP) and Duration Delayed Factor (Δ). The MS-TLM model used a single transformer with two input and output streams.

Edge Duration Prediction. We use the Mean Absolute Error (MAE), or L1 Loss when evaluating edge duration prediction (a MAE of 1 corresponds to 20ms of error). The Duration Accuracy is also reported.

4.3.2 Dialogue Generation Metrics

We evaluate the generation properties of our models using descriptive statistics, automatic metrics and human-based judgements. Unless otherwise written, we perform conditional generation and generate 90-second long continuations using 117 30-second long prompts extracted from the development set and use the default generation temperature of 1.0. We generate the units by sampling among the top 20 possible units.

Turn-taking Event Statistics. We compute the turn-taking events as defined in Section 3.4 using the samples generated by the models with VAD using the pyannotate library⁷ (Bredin et al., 2020). We then analyze the statistics of these turn-taking events (number of events and their durations) across different models.

Turn-taking Event Consistency. We evaluate the model’s capacity to generate consistent conversations in terms of turn-taking events. We measure the Pearson correlation between the total duration of events in each prompt and in the corresponding continuation.

⁷<https://github.com/pyannotate/pyannotate-audio>.

Natural Dialogue Event Statistics. We evaluate the naturalness of the generated speech by focusing on the Speaking Rate (WPM, words per minute), Laughter Frequency (LPM, laughs per minute), Filler Word Rate (FWR, filler words per 100 words), and Floor Transfer Offset (FTO, duration between two consecutive turns of the two speakers, a positive FTO represents a gap while a negative FTO represents an overlap). For this evaluation, we use the same ASR model used to decode the Fisher dataset⁶ to transcribe the generated speech. To detect laughs in the speech, we use an open-source model described in Gillick et al., (2021).⁸ To compute the FWR, we use the following filler words set: {‘uh’, ‘um’, ‘like’, ‘i mean’, ‘you know’}.

Semantic Evaluation. We use two evaluation metrics proposed in Lakhotia et al. (2021), perplexity (PPL) and VERT, to assess the generation quality and diversity of the models. We first transcribe the generated speech using the ASR system. As these metrics are calculated on text sequences, we combine the text from two channels into a single turn-based text sequence,³ ignoring any turns that are completely contained inside an other turn. We employ the open-source DialoGPT model⁹ (Zhang et al., 2019) to compute the perplexity on the turn-based sequences. We simply replace the speaker tokens (<A>,) with the <|endof-text|> token, indicating a turn switch. For the VERT metrics, we also compute the self-BLEU and auto-BLEU on the turn-based text sequences. As the conversation texts contain a lot of repetitions, we report the VERT-4 score instead of VERT-2 score as in Lakhotia et al. (2021).

Since the PPL and VERT scores highly depend on the generation temperature, we perform generation on different temperatures ranging from 0.3–2.0. We then compute the PPL and VERT for each temperature and fit the points corresponding to different temperatures with an exponential line and report the PPL@GT (PPL with respect to the ground truth VERT) score (cf. Figure 7). For the conditional generation case, we compute instead the conditional perplexity (cond. PPL), which is the perplexity of the generated sequence given the concatenation of the prompt sequence

⁸<https://github.com/jrgillick/laughter-detection>.

⁹<https://huggingface.co/microsoft/DialoGPT-medium>.

Id	Model	Number of occurrences / min				Cumulated duration / min			
		IPU	Pause	Gap	Overlap	IPU	Pause	Gap	Overlap
0	MS-TLM	19.4	10.6	5.1	3.3	49.4s	8.9s	2.9s	1.3s
1	DLM-1	17.7	7.9	3.9	5.5	41.4s	13.8s	10.7s	6.1s
2	DLM-2	20.0	10.4	5.5	3.6	48.9s	9.1s	3.6s	1.7s
3	DLM-3	19.0	1.8	4.9	11.7	65.0s	1.1s	1.8s	8.1s
4	DLM-4	18.9	3.2	5.6	9.4	60.7s	2.4s	2.9s	6.1s
5	DLM-5	24.2	5.4	7.2	10.9	59.1s	3.6s	2.9s	5.8s
6	Cascaded	17.5	0.0	14.9	0.0	54.8s	0.0s	5.3s	0.0s
	Ground Truth	21.6	7.0	7.5	6.5	53.5s	5.5s	4.4s	3.6s
	Training Set	25.9	7.2	8.6	10.0	54.5s	5.6s	4.6s	4.7s

Table 2: Number of turn-taking events and cumulated durations per minute across models for prompted continuations, compared to ground truth continuations, and to the same statistics in the training set.




Figure 5: Correlation between the duration of events in the prompts and in the continuations across models, compared to ground truth (GT), where the correlation is computed between the first 30 seconds and the following 90 seconds of the samples.

and generated sequence as input to the DialoGPT model.

Human Opinion Score. We perform a human evaluation on the generated examples. The opinions are based on two dimensions: *N-MOS* (*naturalness Mean Opinion Score*) representing naturalness and turn-taking conversationality, and *M-MOS* (*meaningfulness Mean Opinion Score*) for meaningfulness and content quality. For N-MOS, we asked the participants to concentrate on the fluidity and naturality of the interaction as well as the expressiveness of the speakers regardless of meaning. For M-MOS, they should focus on what is being said and if it is semantically coherent. For these two evaluations, we used a scale of 1-5 (1: worst, 5: best). The CrowdMOS package (Ribeiro et al., 2011) was used for all subjective evaluations using the recommended recipes for

detecting and discarding inaccurate scores. Indeed, we remove all workers whose correlation with the mean scores is lower than 0.25, and then filter out outlier workers whose correlation with the mean scores is lower than 0.6. We enforced at least six raters for each of the generated samples. Participants were recruited using a crowdsourcing platform.

5 Results

5.1 Content and Duration Modeling

Table 1 reports the modeling evaluation metrics on our development subset of the Fisher dataset. In rows Id 1–5, we compare different DLM models, while row Id 0 represents the MS-TLM model, which takes as input multiple unit streams from different channels, and predicts the next-step units only. We note that for models Id 1–3, the next-step

unit prediction objective is also included in the training process, but when the duration prediction objective is employed (models Id 4–5), the next-step unit prediction objective is omitted.

We observe that by using the self cross-attention layers, the edge unit prediction metrics slightly improve (u NLL: 3.07 vs 2.95). On considering models Id 2 & 3, we observe a huge improvement in edge unit NLL & Accuracy when introducing the edge unit prediction objective (u NLL: 2.95 vs 2.49). By introducing the duration prediction objective and removing the next-step unit prediction objective, we see that the model performs even better on the edge unit prediction metrics (u NLL: 2.26), and finally the duration metrics greatly improves when we apply a delayed duration prediction (d MAE: 1.47 vs 1.23).

On comparing with the MS-TLM model, we see that our best DLM model perform much better on content modeling. The reason, we believe, is related to the entangled modeling of content and duration in the MS-TLM model.

5.2 Turn-taking Event Statistics

In this section, we analyze the distribution of the turn-taking events (as described in Section 3.4) in the dialogue continuations generated by our models. The statistics are computed over 3 hours of generated speech per model.

Figure 4 shows the distribution of each of the 4 turn-taking events: IPU, pause, gap, and overlap. In this figure, the Ground truth corresponds to the true continuation of the prompts in the original corpus. Despite having a reasonably good modeling score (cf. Table 1), DLM-1, which has no cross-attention layers between the two transformer towers, has poor performance on turn-takings events, except for the IPU event. The lack of communication between the two channels during generation creates huge gaps and overlaps in the generated samples. The MS-TLM and DLM-2 models have similar distributions of shorter overlaps and longer pauses and gaps. They were trained using the next-step prediction loss on duplicated unit sequences, which could lead to repeated unit generation, causing a slow pace and more extended silences in the generated audio. The opposite effect happens when we introduce the edge unit prediction (DLM-3-5). These models manage to generate more overlaps, with pauses

Id	Model	WPM	LPM	FWR
0	MS-TLM	139.17	1.88	9.36
1	DLM-1	123.60	1.98	9.39
2	DLM-2	141.09	2.06	10.36
3	DLM-3	281.41	7.08	3.40
4	DLM-4	244.13	6.05	3.38
5	DLM-5	211.98	3.62	5.50
6	Cascaded	216.73	0.00	7.08
	Ground Truth	181.46	3.60	7.25

Table 3: Natural Dialogue Event Statistics. Speaking Rate (WPM, words per minute), Laughter Frequency (LPM, laughs per minute) and Filler Word Rate (FWR, filler words per 100 words) of the prompted continuation speech across models, compared to ground truth continuations.

and gaps of shorter duration. These observations are further reinforced in Table 2, which details the number of events and their total durations per minute. It is interesting to note that all models, except DLM-1, manage to capture the empirical fact that *intra-turn* pauses tend to be longer than *between-turn* gaps (Brady, 1968; Ten Bosch et al., 2005; Heldner and Edlund, 2010).

The cascaded model only produces alternating speech turns and therefore has almost no overlap and pause. This also results in low variance in the gap distribution, making the generated speech turn sound like machine conversation.

5.3 Turn-taking Event Consistency

Figure 5 shows the correlation between the total duration of turn-taking events in the prompts and in the generated continuations. For the ground truth, we compute the correlation of the events’ duration between the first 30 seconds and the following 90 seconds in each sample. We observe that in general all models except DLM-1 and cascaded have good correlations, showing their ability to maintain the dialogue consistency. Unsurprisingly, the cascaded model has no correlation with the prompt events, except for the gaps, which are proportional to the number of turn changes.

5.4 Natural Dialogue Event Statistics

Table 3 reports the naturalness statistics on the generated samples of our models. We first notice that, compared to ground truth, models that don’t have edge unit prediction (MS-TLM, DLM-1–2)




Figure 6: Histogram of Floor Transfer Offset (FTO) in the generated speech across models, compared to ground truth continuations and the training set.

tend to produce speech with less information and more hesitations (lower rate, less laughter, more filler words) than those with edge unit prediction (DLM-3–5). Adding duration prediction can effectively help to produce more natural speech, but it still produces more words than ground truth. The cascaded model is unable to produce laughter as the ASR and TTS modules are not able to capture these information, it also generate nearly “non-stop” speech at a faster rate than natural speech. Looking at Figure 6, we see indeed that the cascaded model has no negative FTO (overlap), and the positive FTOs (gaps) fall mostly in the range of one second. In general, other models seem to have good FTO distribution compared to the reference ground truth and training set.

5.5 Semantic Evaluation

For semantic metrics, we perform both conditional and unconditional generations. For conditional generation, we select 50 10-second long prompts in the validation set. For each model and temperature, we generate 50 samples and limit the transcribed turn-based text sequences to 50 words.

We found that for certain models is not possible to obtain PPL@GT as they tend to generate repeated units at low temperatures, creating complete noise in the synthesis. We therefore report the PPL scores for the default temperature 1.0 (@t1).

Id	Model	unconditional		conditional	
		PPL↓		cond. PPL↓	
		@t1	@GT	@t1	@GT
0	MS-TLM	190.59	144.82	741.86	–
1	DLM-1	145.85	–	195.89	–
2	DLM-2	218.30	–	453.73	–
3	DLM-3	155.17	161.58	463.27	329.74
4	DLM-4	290.07	231.00	693.48	314.49
5	DLM-5	179.65	187.16	605.84	365.08
6	Cascaded	32.23	80.80	45.93	117.06
	Ground Truth	100.85	100.85	65.00	65.00

Table 4: Semantic Evaluation. Perplexity of ASR-transcribed generated speech at default temperature (@t1) and at ground truth VERT (@GT) in both unconditional and conditional generation across models compared to ground truth transcriptions. We limit the transcribed turn-based sequences to 50 words.

As shown in Table 4, we see that the dialogue models fail to generate semantically coherent speech, resulting in high perplexity, especially in prompted generation. The cascaded model has a very good perplexity as the language model was trained on word and sub-word levels, it even has a higher PPL@GT than the ground truth in the unconditional case. When it comes to conditional generation, the cascaded model has a good PPL, but is still way below the ground truth.

5.6 Human Evaluations

For this evaluation, we filter the prompts to contain genuine alternations between the two interlocutors and balanced gender. We retained 50 10-second long prompts and generated 10 20-second long continuations for each prompt. Human evaluation results are reported in Table 5. The naturalness and meaningfulness MOS scores correlate well with results in previous sections. The DLM-5 model has the best performance among dialogue models, while the DLM-1 performs significantly worse on both scores. Interestingly, whereas there is a large gap between our best model and ground truth on meaningfulness (1.73 points on the 5-point scale) this gap is much reduced on turn-taking (.53 points). The cascaded model shows a lack of naturalness, while having better scores on meaningfulness than all dialogue models. However, it is still far below the ground truth despite having a very good semantic scores. Overall, our models can generate dialogues mimicking natural turn-taking, while fail maintaining cross-sentence

Id	Model	N-MOS \uparrow	M-MOS \uparrow
0	MS-TLM	3.31 ± 0.43	2.29 ± 0.49
1	DLM-1	2.25 ± 0.60	1.70 ± 0.44
2	DLM-2	2.95 ± 0.37	2.24 ± 0.47
3	DLM-3	3.29 ± 0.43	2.20 ± 0.44
4	DLM-4	3.36 ± 0.44	2.18 ± 0.46
5	DLM-5	3.70 ± 0.46	2.48 ± 0.49
6	Cascaded	2.38 ± 0.63	2.70 ± 0.38
	Ground Truth	4.23 ± 0.26	4.21 ± 0.25

Table 5: Human Evaluations. Conversation naturalness (N-MOS) and conversation meaningfulness (M-MOS) on a 5-point scale (5 is best) with 95% CI.

meaningfulness. We believe the lack of semantic coherence in generated dialogues results from the fine-grained acoustic units used for modeling and the small training corpus size.

6 Conclusion and Future Work

We have presented dGSLM, the first model for spoken dialogue generation trained from raw audio. This model has been shown to reproduce naturalistic intelligible speech, while trained on only 2k hours of audio from telephone conversations. Informal inspection of the generated samples¹ shows that it is able to reproduce non-verbal vocalizations (laughter, backchannels). Detailed analysis of the turn-taking events show that the model is able to reproduce accurate synchronization including distribution and duration of turn-taking events like IPU, gaps, pauses, and overlaps. In particular, it is able to reproduce the rather puzzling observation that inter-turn pauses tend to be on average longer than between-turn gaps, suggesting the pauses alone are not a sufficient signal to indicate a change of turn.

Although the model lacks the ability to produce semantically coherent speech, it paves the way for the construction of more naturalistic human-machine dialogue systems. The logic and timing of turn-taking, which has been up to now very difficult to model artificially emerges naturally from our system, while it is clearly not yet able to process speech at a deep semantic level. This indicates that a model that correctly predicts synchronization between turns can be learned from relatively a small amount of data.




Figure 7: PPL vs VERT scores with unconditioned generation for MS-TLM, DLM-5 and Cascaded models compared to ground truth transcriptions. The sizes of the points correspond to the temperature used for generation (0.3–2.0), squares mean default temperature 1.0. The turn-based sequences are limited to 50 words.

This is surprising given that one major paralinguistic information, intonation, was not explicitly encoded in the input (or the output) of the system. Further work incorporating pitch (Kharitonov et al., 2021) could potentially improve the current results. Results from the cascaded system also suggest that either using larger linguistic units (like BPE) from raw audio (Borsos et al., 2022) or combining our model with text-based models would create systems which could generate more natural and meaningful conversations.

Acknowledgments

In this work, E.D. in his academic role (EHES, ENS-PSL, CNRS) was supported by the Agence Nationale pour la Recherche (ANR-17-EURE-0017 Frontcog, ANR-10-IDEX-0001-02 PSL*, ANR-19-P3IA-0001 PRAIRIE 3IA Institute), a grant from CIFAR (Learning in Machines and Brains). B.S. was also supported by the Agence Nationale pour la Recherche (ANR-19-P3IA-0001 PRAIRIE 3IA Institute).

Appendix

A Phonetic Quality of HuBERT Fisher

In Table A1, we compare the HuBERT Base model (Hsu et al., 2021a) trained on 2000 hours of Fisher dataset versus 1000 hours of Librispeech dataset on the machine-ABX phonetic test. We used Libri-light ABX (Kahn et al., 2020) for the Librispeech test. For the Fisher, we generated a Fisher ABX dataset using the phonetic

	Fisher		LibriSpeech	
	within↓	across↓	within↓	across↓
HuBERT Base	7.77	12.57	3.95	4.69
HuBERT Fisher	5.50	8.35	11.17	14.70

Table A1: Within and Across-Speaker ABX error on Fisher dev and LibriSpeech dev-clean datasets for HuBERT Base and HuBERT Fisher models.

n cross layers	NLL↓	Acc↑
0/6	1.387	71.77
2/6	1.341	72.06
4/6	1.338	72.10
6/6	1.337	72.11

Table A2: Unit prediction loss (NLL) and accuracy metrics of DLM models trained with different number of cross-attention layers. When the number of cross-attention layers is less than 6, they are put on top of self-attention layers. The models are trained with the Next-step Unit Prediction Objective on the parallel unit streams of the Fisher stereo audio dataset.

alignments obtained from Fisher development set. The results clearly show a domain effect, whereby the Fisher dataset is a better training set than the Librispeech dataset for ABX discriminations in Fisher.

B Effects of Cross-Attention Layer

In Table A2, we show the NLL and Accuracy scores of Transformer Language Models as a function of number of Cross-Attention layers. The models are Two-tower Transformer Systems but are trained with only Next-Step Prediction Objective. We find that more layers give better scores, but that 4 layers of cross-attention gives almost the same performance as 6 for less complexity.

References

Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. Towards a human-like open-domain chatbot. *CoRR*, abs/2001.09977.

Jeremy Ang, Rajdip Dhillon, Ashley Krupski, Elizabeth Shriberg, and Andreas Stolcke. 2002. Prosody-based automatic detection of annoyance and frustration in human-computer dialog. In *INTERSPEECH*.

Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *arXiv preprint arXiv:2006.11477*.

Lasse Borgholt, Jakob Drachmann Havtorn, Joakim Edin, Lars Maaløe, and Christian Igel. 2022. A brief overview of unsupervised neural speech representation learning. *arXiv preprint arXiv:2203.01829*.

Zalán Borsos, Raphaël Marinier, Damien Vincent, Eugene Kharitonov, Olivier Pietquin, Matt Sharifi, Olivier Teboul, David Grangier, Marco Tagliasacchi, and Neil Zeghidour. 2022. AudioLM: A language modeling approach to audio generation. <https://arxiv.org/abs/2209.03143>

Paul T. Brady. 1968. A statistical analysis of on-off patterns in 16 conversations. *Bell System Technical Journal*, 47(1):73–91. <https://doi.org/10.1002/j.1538-7305.1968.tb00031.x>

Hervé Bredin, Ruiqing Yin, Juan Manuel Coria, Gregory Gelly, Pavel Korshunov, Marvin Lavechin, Diego Fustes, Hadrien Titeux, Wassim Bouaziz, and Marie-Philippe Gill. 2020. pyannote.audio: Neural building blocks for speaker diarization. In *ICASSP 2020, IEEE International Conference on Acoustics, Speech, and Signal Processing*. Barcelona, Spain. <https://doi.org/10.1109/ICASSP40776.2020.9052974>

Justine Cassell, Tim Bickmore, Lee Campbell, Hannes Vilhjálmsón, and Hao Yan. 2001. *Human Conversation as a System Framework: Designing Embodied Conversational Agents*, pages 29–63. MIT Press, Cambridge, MA, USA.

Yu-An Chung and James Glass. 2020. Generative pre-training for speech with autoregressive predictive coding. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3497–3501. IEEE.

- Christopher Cieri, David Miller, and Kevin Walker. 2004. The Fisher corpus: A resource for the next generations of speech-to-text. In *LREC*.
- Sander Dieleman, Charlie Nash, Jesse Engel, and Karen Simonyan. 2021. Variable-rate discrete representation learning. *arXiv preprint arXiv:2103.06089*.
- Starkey Duncan. 1972. Some signals and rules for taking speaking turns in conversations. *Journal of Personality and Social Psychology*, 23(2):283. <https://doi.org/10.1037/h0033031>
- Erik Ekstedt and Gabriel Skantze. 2020. Turngpt: A transformer-based language model for predicting turn-taking in spoken dialog. *arXiv preprint arXiv:2010.10874*. <https://doi.org/10.18653/v1/2020.findings-emnlp.268>
- Manaal Faruqui and Dilek Hakkani-Tür. 2021. Revisiting the boundary between ASR and NLU in the age of conversational dialog systems. *CoRR*, abs/2112.05842.
- Jon Gillick, Wesley Deng, Kimiko Ryokai, and David Bamman. 2021. Robust laughter detection in noisy environments. In *Proceedings of Interspeech 2021*, pages 2481–2485.
- Agustín Gravano and Julia Hirschberg. 2011. Turn-taking cues in task-oriented dialogue. *Computer Speech & Language*, 25(3):601–634. <https://doi.org/10.1016/j.csl.2010.10.003>
- Mattias Heldner and Jens Edlund. 2010. Pauses, gaps and overlaps in conversations. *Journal of Phonetics*, 38(4):555–568. <https://doi.org/10.1016/j.wocn.2010.08.002>
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhota, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021a. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *CoRR*, abs/2106.07447.
- Wei-Ning Hsu, Anuroop Sriram, Alexei Baevski, Tatiana Likhomanenko, Qiantong Xu, Vineel Pratap, Jacob Kahn, Ann Lee, Ronan Collobert, Gabriel Synnaeve, and Michael Auli. 2021b. Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training. In *Proceedings of Interspeech 2021*, pages 721–725.
- J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P. E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux. 2020. Libri-light: A benchmark for asr with limited or no supervision. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7669–7673. <https://doi.org/10.1109/ICASSP40776.2020.9052942>
- Eugene Kharitonov, Ann Lee, Adam Polyak, Yossi Adi, Jade Copet, Kushal Lakhota, Tu-Anh Nguyen, Morgane Rivière, Abdelrahman Mohamed, Emmanuel Dupoux, and Wei-Ning Hsu. 2021. Text-free prosody-aware generative spoken language modeling. *arXiv preprint arXiv:2109.03264*. <https://doi.org/10.18653/v1/2022.acl-long.593>
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. <http://arxiv.org/abs/1412.6980>
- Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2021. Internet-augmented dialogue generation. *CoRR*, abs/2107.07566.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis. In *Proceedings of NeurIPS*.
- Felix Kreuk, Adam Polyak, Jade Copet, Eugene Kharitonov, Tu-Anh Nguyen, Morgane Rivière, Wei-Ning Hsu, Abdelrahman Mohamed, Emmanuel Dupoux, and Yossi Adi. 2021. Textless speech emotion conversion using decomposed and discrete representations. *arXiv preprint arXiv:2111.07402*.
- Oleksii Kuchaiev, Jason Li, Huyen Nguyen, Oleksii Hrinchuk, Ryan Leary, Boris Ginsburg, Samuel Krizan, Stanislav Beliaev, Vitaly Lavrukhin, Jack Cook, Patrice Castonguay, Mariya Popova, Jocelyn Huang, and Jonathan M. Cohen. 2019. Nemo: A toolkit for building ai applications using neural modules.
- Kushal Lakhota, Evgeny Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, and Emmanuel

- Dupoux. 2021. On Generative Spoken Language Modeling from Raw Audio. *Transactions of the Association for Computational Linguistics*.
- Stephen C. Levinson and Francisco Torreira. 2015. Timing in turn-taking and its implications for processing models of language. *Frontiers in Psychology*, 6:731.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *CoRR*, abs/1910.13461. <https://doi.org/10.18653/v1/2020.acl-main.703>
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. *CoRR*, abs/2005.11401.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *CoRR*, abs/1510.03055.
- Andy T. Liu, Shu-wen Yang, Po-Han Chi, Po-chun Hsu, and Hung-yi Lee. 2020. Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6419–6423. IEEE.
- Ryo Masumura, Tomohiro Tanaka, Atsushi Ando, Ryo Ishii, Ryuichiro Higashinaka, and Yushi Aono. 2018. Neural dialogue context online end-of-turn detection. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 224–228.
- Raveesh Meena, Gabriel Skantze, and Joakim Gustafson. 2014. Data-driven models for timing feedback responses in a map task dialogue system. *Computer Speech & Language*, 28(4):903–922. <https://doi.org/10.1016/j.csl.2014.02.002>
- Vivian Nguyen, Otto Versyp, Christopher Cox, and Riccardo Fusaroli. 2022. A systematic review and Bayesian meta-analysis of the development of turn taking in adult-child vocal interactions.
- Jinjie Ni, Tom Young, Vlad Pandelea, Fuzhao Xue, Vinay Adiga, and Erik Cambria. 2021. Recent advances in deep learning based dialogue systems: A systematic survey. *arXiv preprint arXiv:2105.04387*.
- Lucas Ondel, Lukás Burget, and Jan Cernocký. 2016. Variational inference for acoustic unit discovery. In *SLTU*, volume 81 of *Procedia Computer Science*, pages 80–86. Elsevier.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N19-4009>
- Adam Polyak, Yossi Adi, Jade Copet, Eugene Kharonov, Kushal Lakhotia, Wei-Ning Hsu, Abdelrahman Mohamed, and Emmanuel Dupoux. 2021. Speech resynthesis from discrete disentangled self-supervised representations. In *Proceedings of INTERSPEECH*. <https://doi.org/10.21437/Interspeech.2021-475>
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Antoine Raux and Maxine Eskenazi. 2009. A finite-state turn-taking model for spoken dialog systems. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 629–637. <https://doi.org/10.3115/1620754.1620846>
- Flávio Ribeiro, Dinei Florêncio, Cha Zhang, and Michael Seltzer. 2011. Crowdmos: An approach for crowdsourcing mean opinion score studies. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2416–2419. IEEE.

- <https://doi.org/10.1109/ICASSP.2011.5946971>
- Morgane Rivière and Emmanuel Dupoux. 2021. Towards unsupervised learning of speech features in the wild. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 156–163. IEEE.
- Matthew Roddy, Gabriel Skantze, and Naomi Harte. 2018. Investigating speech features for continuous turn-taking prediction using LSTMs. *arXiv preprint arXiv:1806.11461*. <https://doi.org/10.21437/Interspeech.2018-2124>
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2020. Recipes for building an open-domain chatbot. *CoRR*, abs/2004.13637.
- Harvey Sacks, Emanuel A. Schegloff, and Gail Jefferson. 1974. A simplest systematics for the organization of turn-taking for conversation. *Language*, 50(4):696–735. <https://doi.org/10.1353/lan.1974.0010>
- Emanuel A. Schegloff. 1982. Discourse as an interactional achievement: Some uses of ‘uh huh’ and other things that come between sentences. *Analyzing Discourse: Text and Talk*, 71:71–93.
- Emanuel A. Schegloff. 2000. Overlapping talk and the organization of turn-taking for conversation. *Language in Society*, 29(1):1–63. <https://doi.org/10.1017/S0047404500001019>
- Björn Schuller, Stefan Steidl, Anton Batliner, Felix Burkhardt, Laurence Devillers, Christian Müller, and Shrikanth Narayanan. 2013. Paralinguistics in speech and language—state-of-the-art and the challenge. *Computer Speech & Language*, 27(1):4–39. Special issue on Paralinguistics in Naturalistic Speech and Language. <https://doi.org/10.1016/j.csl.2012.02.005>
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P16-1162>
- Iulian Vlad Serban, Ryan Lowe, Laurent Charlin, and Joelle Pineau. 2016. Generative deep neural networks for dialogue: A short review. *CoRR*, abs/1611.06216.
- Kurt Shuster, Mojtaba Komeili, Leonard Adolphs, Stephen Roller, Arthur Szlam, and Jason Weston. 2022. Language models that seek for knowledge: Modular search & generation for dialogue and prompt completion.
- Gabriel Skantze. 2017. Towards a general, continuous model of turn-taking in spoken dialogue using LSTM recurrent neural networks. In *SIGdial*.
- Gabriel Skantze. 2021. Turn-taking in conversational systems and human-robot interaction: A review. *Computer Speech & Language*, 67:101178. <https://doi.org/10.1016/j.csl.2020.101178>
- Tanya Stivers, Nicholas J. Enfield, Penelope Brown, Christina Englert, Makoto Hayashi, Trine Heinemann, Gertie Hoymann, Federico Rossano, Jan Peter De Ruiter, Kyung-Eun Yoon, et al. 2009. Universals and cultural variation in turn-taking in conversation. *Proceedings of the National Academy of Sciences*, 106(26):10587–10592. <https://doi.org/10.1073/pnas.0903616106>, PubMed: 19553212
- Louis Ten Bosch, Nelleke Oostdijk, and Lou Boves. 2005. On temporal aspects of turn taking in conversational dialogues. *Speech Communication*, 47(1-2):80–86. <https://doi.org/10.1016/j.specom.2005.05.009>
- Kristinn R. Thórisson. 2002. Natural turn-taking needs no manual: Computational theory and model, from perception to action. https://doi.org/10.1007/978-94-017-2367-1_8
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *CoRR*, abs/1807.03748.
- Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. 2017. Neural discrete representation learning. *CoRR*, abs/1711.00937.

- Oriol Vinyals and Quoc V. Le. 2015. A neural conversational model. *CoRR*, abs/1506.05869.
- Nigel G. Ward. 2019. *Prosodic Patterns in English Conversation*. Cambridge University Press.
- Jing Xu, Arthur Szlam, and Jason Weston. 2021. Beyond goldfish memory: Long-term open-domain conversation. *CoRR*, abs/2107.07567.
- Shu-wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhotia, Yist Y. Lin, Andy T. Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, Tzu-Hsien Huang, Wei-Cheng Tseng, Ko-tik Lee, Da-Rong Liu, Zili Huang, Shuyan Dong, Shang-Wen Li, Shinji Watanabe, Abdelrahman Mohamed, and Hung-yi Lee. 2021. Superb: Speech processing universal performance benchmark. *arXiv preprint arXiv:2105.01051*. <https://doi.org/10.21437/Interspeech.2021-1775>
- Victor H. Yngve. 1970. On getting a word in edgewise. In *Chicago Linguistics Society, 6th Meeting, 1970*, pages 567–578.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2019. Dialogpt: Large-scale generative pre-training for conversational response generation. *CoRR*, abs/1911.00536.