



OverFlow: Putting flows on top of neural transducers for better TTS

Shivam Mehta, Ambika Kirkland, Harm Lameris, Jonas Beskow, Éva Székely, Gustav Eje Henter

Division of Speech, Music and Hearing, KTH Royal Institute of Technology, Stockholm, Sweden

{smehta, kirkland, lameris, beskow, szekely, ghe}@kth.se

Abstract

Neural HMMs are a type of neural transducer recently proposed for sequence-to-sequence modelling in text-to-speech. They combine the best features of classic statistical speech synthesis and modern neural TTS, requiring less data and fewer training updates, and are less prone to gibberish output caused by neural attention failures. In this paper, we combine neural HMM TTS with normalising flows for describing the highly non-Gaussian distribution of speech acoustics. The result is a powerful, fully probabilistic model of durations and acoustics that can be trained using exact maximum likelihood. Experiments show that a system based on our proposal needs fewer updates than comparable methods to produce accurate pronunciations and a subjective speech quality close to natural speech. **Index Terms:** Probabilistic TTS, acoustic modelling, hidden Markov models, Glow, invertible post-net

1. Introduction

It was recently demonstrated [1] that sequence-to-sequence neural text-to-speech (TTS) can be improved by replacing conventional neural attention with a so-called neural hidden Markov model, or *neural HMM* [2], which is a kind of *neural transducer* [3]. The resulting model class combines the best features of classic, HMM-based TTS and modern neural sequence-to-sequence models [4, 1]: models are fully probabilistic and can be trained to maximise the exact sequence likelihood, and the use of left-to-right no-skip HMMs ensures each phone in the input is spoken in the correct order, a property known as *monotonicity*. This directly addresses issues due to non-monotonic attention [5], which make many attention-based neural TTS models prone to speaking random gibberish, and generally require a lot of data and updates to learn to speak properly. The sequential nature of autoregressive synthesis is advantageous on devices with limited parallelism, and also lends itself well to incremental and streaming applications of TTS, as shown in [6].

However, TTS based on the neural HMM framework has yet to live up to its full potential. In particular, the systems demonstrated in [7, 6, 1] include restrictive assumptions that state-conditional emission distributions are Gaussian (equivalent to an L2 loss) or Laplace (equivalent to an L1 loss). This makes for weak models of human speech, since natural speech signals follow a highly complex probability distribution.

We present OverFlow, which adds normalising flows on top of neural HMM TTS to better describe the non-Gaussian distribution of speech parameter trajectories. This gives a fully probabilistic model of acoustics and durations that, unlike most flow-based acoustic models [8, 9, 10, 11], uses autoregression to enable long-range memory. Experiments show that the model quickly learns to produce accurate pronunciation and good syn-

Approach:	[12, 13, 14]	[9]	[11]	[15]	[6]	[1, 7]	Prop.
Autoregression	✓			✓	✓	✓	✓
Post-net	✓	N/A	N/A				✓
Monotonic		✓	✓		✓	✓	✓
Attention-free		✓			✓	✓	✓
Fully prob.				✓		✓	✓
Flows		✓	✓	✓			✓

Table 1: Overview of prior acoustic models and the proposed method. “Fully probabilistic” is defined in Sec. 2.2. “Attention-free” refers to the method for aligning input symbols to output frames, and does not consider, e.g., Transformer self-attention.

thesis quality, outperforming comparable systems based on related approaches like Tacotron 2 [12] and Glow-TTS [9]. For audio and code see <https://shivammehta25.github.io/OverFlow/>.

2. Prior work

2.1. TTS with transducers and neural HMMs

Neural HMMs [2] are a kind of autoregressive HMM [16, 17] where the state-conditional emission distribution and transition probability are both defined by a neural net. This makes them vastly more powerful than classical HMMs. The simplest versions of these models represent left-to-right no-skip HMMs, and are therefore monotonic. As these models still satisfy classic hidden Markov assumptions, we can use the forward or Viterbi algorithms [16] to efficiently compute (or lower-bound) the log-likelihood, and then use stochastic gradient ascent on top of automatic differentiation to optimise it. This makes these methods a compelling choice for autoregressive TTS [1].

Mathematically, neural HMMs are a kind of neural transducer [3]. Neural transducers are known for their strong automatic speech-recognition (ASR) performance in recent years (see, e.g., [18]), and were first used for TTS in SSNT-TTS [7]. More recently, they were shown to handle both ASR and streaming TTS within a single model [6]. Our approach differs from [6] in that it is probabilistic (unlike the modified training algorithm in [6]) and from [7] in that integrates quantile-based duration generation [19, 20] for more consistent speech-sound durations. Most importantly, and unlike [7, 6, 1], we integrate normalising flows into this category of models in order to obtain a powerful probabilistic framework capable of describing the behaviour of speech acoustics.

2.2. Normalising flows in TTS acoustic modelling

Normalising flows (a.k.a. *flows*) [21, 22] use deep learning to define highly flexible parametric families of probability distributions. The idea is to create a complex probability distribution

from a simple Gaussian by subjecting it to a series of nonlinear, invertible transformations based on neural networks. Invertibility means that the change-of-variables formula can be used to compute and maximise the likelihood of the resulting model.

Flows have seen many applications across speech technology, but the first TTS acoustic models to integrate flows were Flowtron [15] and Flow-TTS [8]. These were followed by Glow-TTS [9], RAD-TTS [10], and EfficientTTS [11]. Flowtron is an extension of Tacotron 2 with multiple speakers and global style tokens [23] that puts a normalising flow inside the autoregressive decoder that generates the next output frame. Like Tacotron 2, it uses neural attention to learn to speak and align. Flow-TTS is a non-autoregressive architecture reliant on external alignments for training. Glow-TTS improved on Flow-TTS by learning to speak and align at the same time, eliminating the need for external alignment tools. It also introduced a dynamic-programming procedure based on the Viterbi algorithm from HMMs to enforce monotonic alignments during training. RAD-TTS [10] integrated a separate normalising flow for duration modelling along with a more elaborate alignment mechanism, but did not attain the same speech quality as Glow-TTS. EfficientTTS [11], finally, described two efficient mechanisms to encourage or enforce monotonic alignments with conventional dot-product attention, resulting in improved subjective scores. Of these, only Flowtron is *fully probabilistic*, as others lack a discrete-valued probabilistic model of sequence durations, and thus do not define a valid probability distribution over the set of all possible discrete-time output sequences.

Except Flowtron [15], all the above flow-based TTS acoustic models are non-autoregressive (i.e., parallel). This is advantageous on GPU servers. However, autoregressive architectures have often had an edge over non-autoregressive ones in terms of probabilistic-modelling accuracy, e.g., in bits-per-pixel metrics on image benchmarks [24]. A combination of flows and non-autoregressive HMMs recently advanced the state of the art in phone recognition [25, 26]. These findings motivate our innovation to combine flows with autoregression through neural HMMs, since both allow exact maximum-likelihood training and have demonstrated strong probabilistic-modelling results.

Flowtron [15] uses conventional neural attention that does not enforce monotonicity, making it prone to gibberish and difficult to train. Unfortunately, the model cannot learn to speak by training on LJ Speech alone [15], which would be needed for a fair comparison. Glow-TTS [9] enforces monotonic alignments with an algorithm derived from HMMs, like our model. We adopt their decoder architecture in our experiments and also compare to a Glow-TTS system. Their method differs from our proposal in the lack of autoregression, in the use of Viterbi recursion (single path) over the complete forward algorithm (all paths), and in the use of a non-probabilistic duration model.

Flows have also been used in end-to-end TTS, most prominently VITS [27], which adds a flow-based duration model and a neural vocoder to Glow-TTS and trains on a compound loss. This is powerful but, as our results show, much slower to train.

2.3. Flows as an invertible post-net

Many autoregressive TTS systems with regression losses, such as [13, 12, 14], use a greedy, sequential output generation procedure during synthesis. To improve output quality, a so-called *post-net* with non-causal CNNs is then applied to “go back” over the sequence generated by the autoregressive model to enhance it. Unfortunately, that standard post-net architecture is incompatible with maximum-likelihood training of neural

HMMs, and the absence of such a post-net negatively impacts speech quality [1]. By instead passing model output through an *invertible* post-net (thus turning the model into a normalising flow), we are able to incorporate a post-net whilst still training the entire model to maximise exact sequence likelihood. Our experiments use the decoder from Glow-TTS, which leverages an architecture based on non-causal CNNs, evoking similarities to the post-nets in attention-based neural TTS [13, 12, 14].

As an alternative to a post-net, one may instead fine-tune the vocoder to produce more natural waveforms from unenhanced acoustic features. However, this requires an additional training stage, and neural vocoder training is generally slow.

3. Method

Our proposed approach, OverFlow, is shown in Fig. 1. In a nutshell, we apply an invertible neural network to the neural HMM output, greatly increasing the set of probability distributions that can be represented by the model. Experiments in Sec. 4 confirm this significantly improves the output speech. The rest of this section describes the two methods – neural HMM TTS and normalising flows – that make up the proposed approach.

Neural HMMs are a framework of probabilistic encoder-decoder sequence-to-sequence models that provide an alternative to conventional neural attention for aligning input values with output observations. When used for TTS acoustic modelling as in [7, 1], the neural HMM *encoder* converts the input vectors (e.g., phone embeddings) into a sequence of vectors $\mathbf{h}_{1:N}$ that define N states in a left-to-right no-skip HMM. Each input symbol is translated to a fixed number of vectors, for example two states per symbol. The neural HMM *decoder* network takes two inputs and returns two outputs, which we write as $(\theta_t, \tau_t) = \text{Dec}(\mathbf{h}_n, \mathbf{x}_{1:t-1})$. The inputs are a state-defining vector \mathbf{h}_{s_t} from the encoder, corresponding to the HMM state $s_t \in \{1, \dots, N\}$ at timestep t , as well as the previously generated output acoustic frames $\mathbf{x}_{1:t-1}$. The latter input makes the model autoregressive. In practice, it is common to only provide the previous frame \mathbf{x}_{t-1} as explicit input, and then process this input using a so-called *pre-net*, along with an LSTM that ensures outputs from before $t-1$ also can influence the next frame.

The two decoder outputs are θ_t (the parameters of the emission distribution) and a *transition probability* $\tau_t \in [0, 1]$. The parameters θ_t define a probability distribution for the next output frame \mathbf{x}_t . Most neural HMMs assume that \mathbf{x}_t follows a multivariate Gaussian distribution with diagonal covariance matrix. In that case, the parameters $\theta_t = (\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t)$ are two vectors corresponding to the element-wise means and standard deviations of \mathbf{x}_t . Meanwhile, τ_t defines the probability that the HMM transitions to the next state, so that $s_{t+1} = s_t + 1$; else $s_{t+1} = s_t$. Synthesis begins with $s_1 = 1$ and terminates when $s_t = N + 1$. To satisfy the Markov assumption over the hidden states $s_{1:T}$ – and thus be a valid neural HMM – the decoder output may not depend on the input vectors $\mathbf{h}_{1:t-1}$ at previous timesteps, which means that models can be trained to maximise the log-likelihood of training data, computed using the standard forward or Viterbi algorithms from classical HMMs [16]. For a graphical overview of neural HMM TTS, see Fig. 1.

Neural HMMs describe a distribution over discrete-time sequences, typically assuming a Gaussian distribution for each frame. Normalising flows [21, 22], provide a method for turning simple *latent* or *source* distributions \mathbf{Z} (e.g., Gaussians) into much more flexible *target* distributions \mathbf{X} . The idea is to apply an invertible nonlinear transformation \mathbf{f} to the source distribution to obtain the target distribution, as $\mathbf{X} = \mathbf{f}(\mathbf{Z}; \mathbf{W})$,

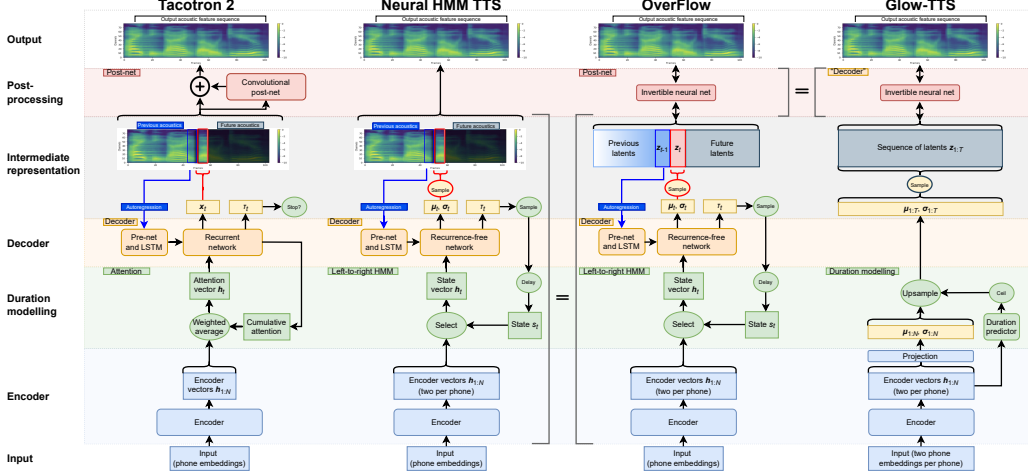


Figure 1: Comparison of four systems from the experiments. Square brackets show how the OverFlow system passes the output of a neural HMM through the invertible neural network from Glow-TTS, similar to the post-net in Tacotron 2. Rectangular boxes are vectors and scalars, rounded boxes are learnt transformations (neural networks), and ovals are fixed mathematical operations.

where \mathbf{f} is parameterised by a neural network with weights \mathbf{W} . Invertibility allows us to compute the (log-)probability of any observed outcome \mathbf{x} using the change-of-variables formula

$$\ln p_{\mathbf{X}}(\mathbf{x}) = \ln p_{\mathbf{Z}}(\mathbf{f}^{-1}(\mathbf{x}; \mathbf{W})) + \ln |\det \mathbf{J}_{\mathbf{f}^{-1}}(\mathbf{x}; \mathbf{W})|,$$

where $\mathbf{J}_{\mathbf{f}^{-1}}(\mathbf{x}; \mathbf{W})$ is the Jacobian matrix of \mathbf{f}^{-1} evaluated at \mathbf{x} . Even if \mathbf{f} is a relatively weak nonlinear transformation, we can obtain highly flexible distributions by chaining together several component transformations \mathbf{f}_i . The combined transformation is often known as an *invertible neural network*. Crucially, both neural HMMs and normalising flows allow exact likelihood maximisation, making them a great fit for strong probabilistic models of speech acoustics.

Glow [28] is one popular normalising-flow architecture. It interleaves learnt global affine transformations with so-called *coupling layers*, that nonlinearly transform half of the elements of the input vector $\mathbf{z}_l \in \mathbb{R}^D$ with respect to the value of the remaining elements (which are left unaltered). Calling the output of the l th coupling layer \mathbf{z}'_l , the layer performs an element-wise invertible affine transformation that can be written

$$\begin{aligned} \mathbf{z}'_{l, 1:D/2} &= \mathbf{z}_{l, 1:D/2} \\ \mathbf{z}'_{l, D/2+1:D} &= \boldsymbol{\alpha}_l(\mathbf{z}_{l, 1:D/2}) \odot \mathbf{z}_{l, D/2+1:D} + \boldsymbol{\beta}_l(\mathbf{z}_{l, 1:D/2}), \end{aligned}$$

where $\boldsymbol{\alpha}_l > \mathbf{0}$ and $\boldsymbol{\beta}_l$ are the outputs of a neural network with weights \mathbf{W} . It is easy to verify that this transformation can be inverted, since $\boldsymbol{\alpha}_l$ and $\boldsymbol{\beta}_l$ can be computed by feeding $\mathbf{z}'_{l, 1:D/2}$ into the neural net. The matrix multiplication (a.k.a. “ 1×1 convolution”) in the affine transformation can be seen as a generalisation of a permutation operation [28], ensuring that all input elements receive multiple nonlinear transformations by the flow.

Most normalising flows in TTS are based on Glow, with Glow-TTS introducing a more parameter-efficient variant of the global affine transformations [9]. Since both Glow and Glow-TTS use CNNs in the coupling layers, the resulting invertible neural network has a finite receptive field. As such, it cannot capture global dependencies between acoustics, which might explain the idiosyncratic utterance-level intonation and emphasis patterns generated by many trained Glow-TTS systems [29]. In our proposal the source distribution \mathbf{Z}_t depends on all previous outputs $\mathbf{z}_{1:t-1}$ through autoregression in a neural HMM (with long-range memory provided by an LSTM), potentially producing more globally consistent synthetic speech.

4. Experiments

To validate our proposal we compared three prior acoustic models to a comparable version of OverFlow in several experiments. We also compared to an end-to-end and a non-parallel TTS baseline. See shivammehta25.github.io/OverFlow/ for audio.

The baseline acoustic models were Tacotron 2 [12] (label **T2**), Glow-TTS [9] (**GTTS**), and the neural HMM TTS (**NHMM**) system in [1]. For Tacotron 2, we used the Nvidia implementation¹ while Glow-TTS used the official GitHub implementation² and sampling temperature (the variance of \mathbf{Z}_t) 0.667 during synthesis. Both used default hyperparameters, for which T2 and GTTS have a very similar model size. The NHMM baseline used the same configuration as system “NH2” in [1], which inherits the encoder and decoder architectures from the T2 system (minus one of the decoder LSTMs). NHMM used greedy deterministic output generation $\mathbf{x}_t = \boldsymbol{\mu}_t$, which is equivalent to random sampling with temperature zero.

For the OverFlow system in the experiments, we added an invertible network adapted from the official Glow-TTS code on top of the NHMM baseline architecture. By reducing the number of nodes in each hidden layer in that added network from 192 to 150, we obtained a total model size very similar to the T2 and GTTS baselines; see Table 2. Our experiments opted for component networks with simple architectures that maximise comparability with prior system implementations, but it is worth noting that OverFlow is a framework within which many different network architectures can be used. It is, e.g., entirely possible to employ Transformers [30] inside any of the blocks in Fig. 1, or, for incremental TTS, one might swap in a causal encoder architecture with optional finite-horizon look-ahead.

All discrete-state acoustic models used two encoder vectors per phone, since this has been found to improve synthesis quality with both neural HMMs and Glow-TTS. Tacotron 2, in contrast, does not need this, since it can represent arbitrary intermediate states using its continuous-valued attention instead.

As additional baselines we included **VITS** [27], an end-to-end model with normalising flows, and **FastPitch v1.1 (FP)** [31, 32], a widely used non-autoregressive neural TTS model. Both of these used the implementation and hyperparameters from the

¹<https://github.com/NVIDIA/tacotron2/>

²<https://github.com/jaywalnut310/glow-tts/>

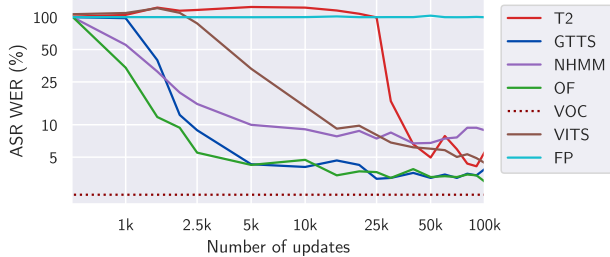


Figure 2: Log-log plot of validation-set TTS WER over time.

Coqui TTS framework³, with `espeak` for text processing.

We trained all systems on LJ Speech⁴ for 100k updates with batch size 32 on a single GPU. Systems were provided the same normalised, phonetised text input and the same mel-spectrogram output features with the same vocoder for waveform generation, namely the pre-trained (V1) universal HiFi-GAN [33] model⁵ followed by a denoising filter as introduced in [34] at a strength of 0.004. Phone durations were generated deterministically and not sampled for any acoustic model; NHMM and OverFlow made use of quantile-based duration generation [19, 20]. OverFlow is seen to improve modelling accuracy very substantially, since it reaches a validation-set per-sequence log-likelihood of 35k after training, compared to 6.5k for NHMM.

For the trained OverFlow system, we generated output under three *conditions* that varied in sampling temperature and in the use of synthesis-time pre-net dropout. The default condition (**OF**) uses a dropout of 0.5 (same as training) and a sampling temperature of 0.667; the no-dropout condition (**OFND**) set synthesis-time pre-net dropout to 0; whilst the zero-temperature condition (**OFZT**) instead set sampling temperature to 0, to allow comparing random sampling to greedy output generation.

Fig. 2 graphs how the ASR word error rate (WER) of synthesising the 100 validation utterances evolves during training, as computed using Whisper [35]. WERs from modern ASR are known to correlate well with speech intelligibility to human listeners [36]. As expected, T2 requires many more updates than other systems to learn to speak since it does not enforce monotonic alignments. We also measured the wall-clock time until reaching 5% WER on the validation utterances, and report the numbers in Table 2. Results show that Glow-TTS and OverFlow reached this level at least seven times faster than other systems. Surprisingly, FastPitch did not learn to speak intelligibly at 100k, perhaps due to v1.1 not using an external aligner.

To evaluate intelligibility on phonetically balanced material, we synthesised the 720 Harvard sentences [37], giving the WERs listed in Table 2. The OverFlow conditions performed substantially better than all other baselines on these sentences.

To evaluate subjective TTS quality, we carried out a web-based listening test on 40 sentences from the LJ Speech test set, enabling us to also compare to vocoded speech (condition **VOC**). VITS and FP were excluded since VITS is not an acoustic model (and uses a different vocoder) whilst FP did not speak intelligibly at 100k updates. We recruited 60 self-reported native speakers of English via Prolific, who all reported wearing headphones during the test. The test used a MUSHRA-like [38] design where listeners were presented with all seven versions of each utterance side by side on a single screen in random order and, similar to a MOS test [39], rated the natural-

³<https://github.com/coqui-ai/TTS/>

⁴<https://keithito.com/LJ-Speech-Dataset/>

⁵<https://github.com/jik876/hifi-gan/>

Me- thod	Con- dition	Temp- erature	Pre-net dropout	Model size	WER @100k	Time to 5%	MOS
[12]	T2	N/A	✓	28.2M	6.36%	54 h	3.25
[9]	GTTS	0.667	N/A	28.6M	3.97%	2.5 h	2.64
[1]	NHMM	0	✓	15.3M	5.96%	125 h	2.97
[27]	VITS	0.667	N/A	83.1M	7.03%	23 h	-
[31]	FP	N/A	N/A	37.5M	100%	21 h	-
Prop.	OF	0.667	✓	28.5M	2.91%	3 h	3.43
	OFND	0.667		"	2.92%	"	3.25
	OFZT	0	✓	"	2.30%	"	3.01

Table 2: Summary of experiment results. 95% confidence intervals for MOS values are ± 0.07 . VOC MOS was 4.18 ± 0.06 .

ness of the audio samples on an integer scale from 1 (“Bad”) to 5 (“Excellent”). Each participant rated 16 utterances, randomly selected from the pool of 40. No test stimuli had any attention issues. Stimuli were loudness-normalised to -20 dB LUFS based on EBU R128 [40]. We additionally included two screens with attention checks, identical to the others except that one audio stimulus instructed the listener to rate that stimulus as 1. Ratings from these screens were excluded from the analysis. Listeners who failed the attention checks or rated $\text{VOC} \leq 2$ more than twice were disqualified and replaced with new listeners.

Mean opinion scores (MOS) from the listening test are found in Table 2. We see that our proposed addition of flows substantially improved on NHMM, and that OverFlow also performs much better than GTTS. Pairwise Wilcoxon signed-rank tests find all pairs of conditions in Table 2 except for (OFND, T2) to be significantly different ($p < 0.05$) after Holm-Bonferroni [41] correction. OF thus performs better than T2, GTTS, and NHMM in the subjective evaluation, and in addition produces intelligible speech much sooner than T2, VITS, or FP, which is important during system development and tuning.

OverFlow is easy to fine-tune on diverse accents; see our webpage for examples. It also achieves good audio quality with both zero and nonzero sampling temperature (OFZT vs. OF/OFND in our listening test). This contrasts against Glow-TTS, where quality degrades substantially at temperature zero (see examples on our webpage), and against NHMM, which instead does much better if the temperature is zero (cf. [1]). Random sampling (OF) is also favoured over not sampling (OFZT), which indicates a highly accurate probabilistic model [42].

5. Conclusions and future work

We have described how the framework of neural HMMs in TTS can be combined with normalising flows (through an invertible post-net) to define strong probabilistic models of speech durations and acoustics. The result is a neural transducer with normalising flows. A simple example system based on our method rapidly learns to speak with fewer mispronunciations than comparable methods, and achieves strong subjective naturalness ratings. Future work includes stronger models using Transformers [30, 14], multi-speaker synthesis, and applications to diverse and challenging data such as spontaneous speech [43, 44].

6. Acknowledgements

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, the Swedish Research Council projects VR-2019-05003 and VR-2020-02396 and Riksbankens Jubileumsfond project P20-0298.

7. References

- [1] S. Mehta, É. Székely, J. Beskow, and G. E. Henter, “Neural HMMs are all you need (for high-quality attention-free TTS),” in *Proc. ICASSP*, 2022, pp. 7457–7461.
- [2] K. M. Tran, Y. Bisk, A. Vaswani, D. Marcu, and K. Knight, “Unsupervised neural hidden Markov models,” in *Proc. Workshop on Structured Prediction for NLP*, 2016, pp. 63–71.
- [3] L. Yu, J. Buys, and P. Blunsom, “Online segment to segment neural transduction,” in *Proc. EMNLP*, 2016, pp. 1307–1316.
- [4] O. Watts, G. E. Henter, J. Fong, and C. Valentini-Botinhao, “Where do the improvements come from in sequence-to-sequence neural TTS?” in *Proc. SSW*, 2019, pp. 217–222.
- [5] M. He, Y. Deng, and L. He, “Robust sequence-to-sequence acoustic modeling with stepwise monotonic attention for neural TTS,” *Proc. Interspeech*, pp. 1293–1297, 2019.
- [6] J. Chen, X. Tan, Y. Leng, J. Xu, G. Wen, T. Qin *et al.*, “Speech-T: Transducer for text to speech and beyond,” *Proc. NeurIPS*, vol. 34, pp. 6621–6633, 2021.
- [7] Y. Yasuda, X. Wang, and J. Yamagishi, “Initial investigation of encoder-decoder end-to-end TTS using marginalization of monotonic hard alignments,” in *Proc. SSW*, 2019, pp. 211–216.
- [8] C. Miao, S. Liang, M. Chen, J. Ma, S. Wang, and J. Xiao, “Flow-TTS: A non-autoregressive network for text to speech based on flow,” in *Proc. ICASSP*, 2020, pp. 7209–7213.
- [9] J. Kim, S. Kim, J. Kong, and S. Yoon, “Glow-TTS: A generative flow for text-to-speech via monotonic alignment search,” in *Proc. NeurIPS*, 2020, pp. 8067–8077.
- [10] K. J. Shih, R. Valle, R. Badlani, A. Łancucki, W. Ping, and B. Catanzaro, “RAD-TTS: Parallel flow-based TTS with robust alignment learning and diverse synthesis,” in *Proc. INNF+*, 2021.
- [11] C. Miao, L. Shuang, Z. Liu, C. Minchuan, J. Ma, S. Wang, and J. Xiao, “EfficientTTS: An efficient and high-quality text-to-speech architecture,” in *Proc. ICML*, 2021, pp. 7700–7709.
- [12] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang *et al.*, “Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions,” in *Proc. ICASSP*, 2018, pp. 4799–4783.
- [13] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly *et al.*, “Tacotron: Towards end-to-end speech synthesis,” in *Proc. Interspeech*, 2017, pp. 4006–4010.
- [14] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, “Neural speech synthesis with transformer network,” in *Proc. AAAI*, 2019, pp. 6706–6713.
- [15] R. Valle, K. J. Shih, R. Prenger, and B. Catanzaro, “Flowtron: an autoregressive flow-based generative network for text-to-speech synthesis,” in *Proc. ICLR*, 2021.
- [16] L. R. Rabiner, “A tutorial on hidden Markov models and selected applications in speech recognition,” *Proc. IEEE*, vol. 77, no. 2, pp. 1989.
- [17] M. Shannon, H. Zen, and W. Byrne, “Autoregressive models for statistical parametric speech synthesis,” *IEEE T. Audio Speech*, vol. 21, no. 3, pp. 587–597, 2013.
- [18] Y. Zhang, J. Qin, D. S. Park, W. Han, C.-C. Chiu, R. Pang *et al.*, “Pushing the limits of semi-supervised learning for automatic speech recognition,” in *Proc. NeurIPS SAS Workshop*, 2020.
- [19] S. Ronanki, O. Watts, S. King, and G. E. Henter, “Median-based generation of synthetic speech durations using a non-parametric approach,” in *Proc. SLT*, 2016, pp. 686–692.
- [20] G. E. Henter, S. Ronanki, O. Watts, and S. King, “Non-parametric duration modelling for speech synthesis with a joint model of acoustics and duration,” in *IEICE Tech. Rep.*, no. 414, 2017, pp. 11–16.
- [21] G. Papamakarios, E. Nalisnick, D. Jimenez Rezende, S. Mohamed, and B. Lakshminarayanan, “Normalizing flows for probabilistic modeling and inference,” *J. Mach. Learn. Res.*, vol. 22, pp. 1–64, 2021.
- [22] I. Kobyzev, S. Prince, and M. Brubaker, “Normalizing flows: An introduction and review of current methods,” *IEEE T. Pattern Anal.*, vol. 43, no. 11, pp. 3964–3979, 2021.
- [23] Y. Wang, D. Stanton, Y. Zhang, R. Skerry-Ryan, E. Battenberg, J. Shor *et al.*, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *Proc. ICML*, 2018, pp. 5180–5189.
- [24] D. Kingma, T. Salimans, B. Poole, and J. Ho, “Variational diffusion models,” in *Proc. NeurIPS*, 2021, pp. 21 696–21 707.
- [25] A. Ghosh, A. Honoré, D. Liu, G. E. Henter, and S. Chatterjee, “Robust classification using hidden Markov models and mixtures of normalizing flows,” in *Proc. MLSP*, 2020.
- [26] —, “Normalizing flow based hidden Markov models for classification of speech phones with explainability,” *arXiv preprint arXiv:2107.00730*, 2021.
- [27] J. Kim, J. Kong, and J. Son, “VITS: Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *Proc. ICML*, 2021, pp. 5530–5540.
- [28] D. P. Kingma and P. Dhariwal, “Glow: Generative flow with invertible 1x1 convolutions,” in *Proc. NeurIPS*, 2018, pp. 10 236–10 245.
- [29] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov, “Grad-TTS: A diffusion probabilistic model for text-to-speech,” in *Proc. ICML*, 2021, pp. 8599–8608.
- [30] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser *et al.*, “Attention is all you need,” in *Proc. NIPS*, 2017, pp. 5998–6008.
- [31] A. Łańcucki, “FastPitch: Parallel text-to-speech with pitch prediction,” in *Proc. ICASSP*, 2021, pp. 6588–6592.
- [32] R. Badlani, A. Łancucki, K. J. Shih, R. Valle, W. Ping, and B. Catanzaro, “One TTS alignment to rule them all,” in *Proc. ICASSP*, 2022, pp. 6092–6096.
- [33] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis,” in *Proc. NeurIPS*, 2020, pp. 17 022–17 033.
- [34] R. Prenger, R. Valle, and B. Catanzaro, “WaveGlow: A flow-based generative network for speech synthesis,” in *Proc. ICASSP*, 2019.
- [35] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” *arXiv preprint arXiv:2212.04356*, 2022.
- [36] J. Taylor and K. Richmond, “Confidence intervals for ASR-based TTS evaluation,” in *Proc. Interspeech*, 2021, pp. 2791–2795.
- [37] E. H. Rothaus *et al.*, “IEEE recommended practice for speech quality measurements,” *IEEE T. Acoust. Speech*, vol. 17, no. 3, pp. 225–246, 1969.
- [38] ITU-R, “Method for the subjective assessment of intermediate quality levels of coding systems,” ITU Recommendation ITU-R BS.1534-3, 2015.
- [39] ITU-T, “Methods for subjective determination of transmission quality,” ITU Recommendation ITU-T P.800, 1996.
- [40] EBU, “Loudness normalisation and permitted maximum level of audio signals,” EBU Recommendation EBU R 128v4, 2020.
- [41] S. Holm, “A simple sequentially rejective multiple test procedure,” *Scand. J. Stat.*, vol. 6, no. 2, pp. 65–70, 1979.
- [42] G. E. Henter, T. Merritt, M. Shannon, C. Mayo, and S. King, “Measuring the perceptual effects of modelling assumptions in speech synthesis using stimuli constructed from repeated natural speech,” in *Proc. Interspeech*, 2014, pp. 1504–1508.
- [43] É. Székely, G. E. Henter, J. Beskow, and J. Gustafson, “Spontaneous conversational speech synthesis from found data,” in *Proc. Interspeech*, 2019, pp. 4435–4439.
- [44] H. Lameris, S. Mehta, G. E. Henter, J. Gustafson, and É. Székely, “Prosody-controllable spontaneous TTS with neural HMMs,” in *Proc. ICASSP*, 2023.