



Vocal effort modeling in neural TTS for improving the intelligibility of synthetic speech in noise

Tuomo Raitio, Petko Petkov, Jiangchuan Li, Muhammed Shifas, Andrea Davis, Yannis Stylianou

Apple

Abstract

We present a neural text-to-speech (TTS) method that models natural vocal effort variation to improve the intelligibility of synthetic speech in the presence of noise. The method consists of first measuring the spectral tilt of unlabeled conventional speech data, and then conditioning a neural TTS model with normalized spectral tilt among other prosodic factors. Changing the spectral tilt parameter and keeping other prosodic factors unchanged enables effective vocal effort control at synthesis time independent of other prosodic factors. By extrapolation of the spectral tilt values beyond what has been seen in the original data, we can generate speech with high vocal effort levels, thus improving the intelligibility of speech in the presence of masking noise. We evaluate the intelligibility and quality of normal speech and speech with increased vocal effort in the presence of various masking noise conditions, and compare these to well-known speech intelligibility-enhancing algorithms. The evaluations show that the proposed method can improve the intelligibility of synthetic speech with little loss in speech quality.

Index Terms: vocal effort, speech synthesis, speech-in-noise intelligibility, spectral tilt

1. Introduction

The intelligibility of natural and synthetic speech decreases in adverse listening conditions [1, 2, 3]. To improve speech intelligibility in the presence of noise, humans adapt their voice, referred to as the Lombard effect [4, 5], by increasing the intensity of the signal, and therefore the signal-to-noise ratio (SNR), and by modifying voice characteristics as a result of increased vocal effort [6, 7]. Change in the vocal effort can also be triggered by the need to communicate over a distance [8] or a change in emotional expression [9, 10]. Similarly, it is desirable for speech synthesizers to adapt to adverse environments and different contexts for improved intelligibility and more natural delivery of the message. Increasing the signal intensity increases the SNR and intelligibility, but this approach alone is not ideal for either the listener (discomfort) or the output device (power consumption, mechanical limits). Therefore, a preferable approach is to change the speech characteristics to increase the intelligibility, also called near-end listening enhancement [11, 12, 13, 14]. This is especially important for text-to-speech (TTS) where the intelligibility is generally lower than in natural speech [3].

Assuming an equal power constraint, there are two main approaches that have been successfully used to improve the intelligibility of synthetic speech in noise. First, using signal processing techniques, such as spectral shaping and dynamic range compression [15, 14], to shift the signal power to frequencies that are most relevant for intelligibility and to amplify low-energy speech sounds, such as plosives and stops. Second, mimicking human speech production in noise, which can be achieved by training or adapting a TTS synthesizer with specifi-

cally recorded Lombard or high vocal effort speech [16, 17, 18].

Both approaches have been shown to improve the intelligibility of speech in noise. The signal processing based methods have the benefit that no additional speech data are required to enhance the intelligibility, and the intelligibility gains can be substantial [1, 2]. On the other hand, modeling vocal effort in speech synthesis not only increases the intelligibility of speech in noise, but also makes the synthetic voice more natural and more suitable for the listening condition [16, 17]. However, specifically recorded Lombard or high vocal effort speech data is usually required for training the TTS system.

To overcome the limitations of requiring specifically recorded and labeled Lombard or high vocal effort speech, we propose a method that automatically estimates the vocal effort level from normal speech data using spectral tilt [19, 20]. The natural variation of vocal effort is then modeled and controlled at synthesis time. Moreover, the model is able to extrapolate vocal effort level from the observed neutrally spoken speech data to reproduce speech with very low (soft) or high (loud) vocal effort. Also, the vocal effort modification at synthesis time has minimal effect on other prosodic factors, such as duration and pitch. Finally, the proposed method provides additional benefits, such as overall control of the prosodic space independently from other prosodic factors [21, 22], and higher naturalness in comparison to signal processing based methods¹.

2. Background

The ability of speakers to adapt the speech generation process to the environment has been studied thoroughly in the past. We summarize some of the main findings, relevant to this work, in Sec. 2.1. The results of such studies have contributed to developing various techniques for artificially enhancing the intelligibility of natural and synthetic speech signals. To provide the context for evaluating vocal effort adjustment for intelligibility enhancement, we summarize the relevant methods in Sec. 2.2.

2.1. Speech modification for enhanced intelligibility

Changes in vocal effort induced by background noise are referred to as the Lombard effect [4, 23, 5]. Communication over a distance [8], as well as emotional expression [9, 10] also modulate vocal effort as a means of efficient message delivery. The associated effects on naturally-produced speech have been studied, e.g., in [24, 25, 23, 8].

Key speech characteristics affected in the process of modulating vocal effort include pitch, duration, intensity and spectral tilt [4]. Isolating each one of these in turn and probing its impact on intelligibility in background noise has shown that pitch and duration do not have a significant contribution [26, 2]. Such changes are explained in view of physiological constraints in the voice production mechanism. Intensity and spectral tilt, on

¹Speech samples can be found at <https://apple.github.io/vocal-effort-modeling-tts-intelligibility/>.

the other hand, have a direct and significant impact on human ability to understand speech in noise.

In the case of intensity, intelligibility increases as a result of an overall increase in SNR. Reduction of spectral tilt, on the other hand, boosts the mid and high frequencies at the expense of the low frequencies. This provides an increased amount of intelligibility cues to the listener due to: 1) typical low-pass characteristics of environmental noise, and 2) the higher importance of the mid relative to the low frequency range of the speech signal [27]. Fixing the average intensity and exploring time-variant as well as time-invariant frequency-dependent speech modifications has produced a numbers of methods for intelligibility enhancement. Key methods, from the perspective of TTS, are summarized in the following section.

2.2. Methods for intelligibility enhancement of TTS signals

Intelligibility enhancement of TTS can be performed at synthesis time, e.g., by parametrizing and adjusting properties of the speech generation process as a function of environment, which is the approach explored in this work. Alternatively, it can be achieved by post-processing a pre-synthesized signal. While any existing method for intelligibility enhancement can be applied in the latter case, the quality of synthetic speech must be sufficiently high to prevent signal quality degradation [28]. State-of-the-art neural TTS satisfies the requirement for quality and is suitable for combination with post-processing methods.

One of the more successful post-processing methods combines dynamic range compression (DRC) with high-pass filtering [15]. More recently, spectral shaping and dynamic range compression (SSDRC) similarly exploits the combination of a linear time-invariant (LTI) filter and DRC [14]. A multi-band parametric time-varying DRC model recovering the spectral dynamics of the signal is presented in [29]. A multi-band time-invariant DRC approach tuned to mitigate the combined effect of noise and reverberation is considered in [30].

Large-scale evaluations through headphones indicate that DRC-based techniques outperform alternative methodologies in terms of intelligibility gain [31, 2]. On the other hand, LTI and linear time variant (LTV) systems offer the benefit of lower signal processing distortion and robustness to noise in the input signal. Thus, where DRC would decrease SNR in the input signal, LTI and LTV would preserve it. Among others, some recently-proposed techniques are found in [32, 33, 34].

Intelligibility enhancement at synthesis time in the parametric space of a neural TTS system is a novel paradigm and differs from prior art due to its capability to produce both linear (filtering) and non-linear (DRC) signal modifications by non-linear transformations (neural networks). Existing work in this domain relies on reference training data, which is either created with a pre-existing signal-processing approach [35], or taken from a categorical speaking style such as clear or Lombard speech [16, 17, 36, 18]. The effectiveness of such methods is constrained by the quality and amount of the training data and the capability of the neural network model for adaptation. To our knowledge, this work is the first example of neural TTS providing both: 1) the capability for parametric adaptation of TTS synthesis to environmental conditions, and 2) the possibility of adaptation without dedicated and labeled training data.

3. Modeling of vocal effort in neural TTS

In this section, we describe the measurement of vocal effort using spectral tilt, and how that can be used in the modeling of

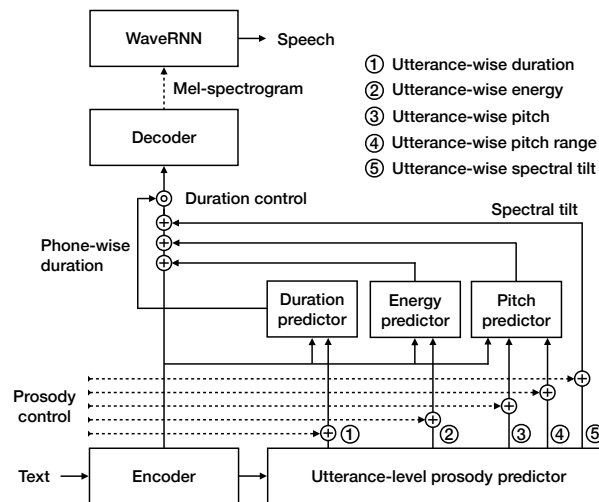


Figure 1: Proposed TTS architecture with vocal effort control.

vocal effort in neural TTS.

3.1. Measuring vocal effort using spectral tilt

To model vocal effort without changing other prosodic aspects of speech, we extract and model the following set of prosodic features, averaged per utterance: pitch, pitch range, phone duration, speech energy, and spectral tilt. These features are easy to calculate and are robust against background noise or other recordings conditions. These features are also disentangled so that they can be varied independently [21, 22].

The extraction of the acoustic features is described in [22], however, we will describe the extraction of spectral tilt here in detail. First, we resample the speech signal to 16 kHz and use a high-pass filter with a cut-off frequency of 70 Hz to remove any low-frequency fluctuations. Then we calculate the frame-wise spectral tilt of voiced speech using the predictor coefficient of a first order all-pole filter. The spectral tilt values are averaged per utterance and then normalized to $[-1, 1]$ by first calculating the median (M) and the standard deviation (σ) and then projecting the data in the range $[M - 3\sigma, M + 3\sigma]$ into $[-1, 1]$. Finally, we clip values $|x| > 1$ so that all data is in the range $[-1, 1]$.

3.2. Neural TTS architecture

A typical neural text-to-speech (TTS) system consists of two components: a neural front-end [37, 38, 39] that maps from phoneme input to Mel-spectrograms, and a neural back-end [40, 41] that maps from the Mel-spectrogram into a sequence of speech samples. These two networks are trained with a large amount of speech data to generate high-quality speech. Since prosody and vocal effort are mainly modeled by the neural front-end, this paper will focus on the front-end architecture.

Previously, we presented hierarchical prosody modeling and control using Tacotron 2 [21] and non-autoregressive parallel TTS [22], both including spectral tilt modeling. In this work we focus specifically on modeling vocal effort using spectral tilt and its effect on speech intelligibility in noise.

The proposed front-end architecture, shown in Fig. 1, is similar to FastSpeech 2 [39], but we extend the model with a new variance adaptor at the utterance level [22]. We predict the utterance-wise prosodic features from the encoder outputs, namely, pitch, pitch range, duration, energy, and spectral tilt. The architecture of the new utterance-level variance adaptor is

identical to the phone-level adaptors except that we predict all five features using the same module. We use teacher-forcing of both utterance-wise and phone-wise prosodic features using ground-truth targets to efficiently train the network. At synthesis time, vocal effort control is achieved by adding bias to the spectral tilt predictions while keeping the other prosodic features unchanged to minimize undesired changes in prosody. Factorizing the prosody and only changing intended prosodic aspects is crucial in order to achieve desired changes with high quality. We observed overall lower synthesis quality with models that were conditioned only with spectral tilt due to the high correlation between speaking style and spectral tilt.

4. Experiments

4.1. Data

We use data from two American English speakers, a high-pitched speaker with 36-hour dataset (Voice 1), and a low-pitched speaker with 23-hour dataset (Voice 2). All speech data were sampled at 24 kHz. For evaluating the intelligibility and quality of synthetic speech, we used the 720 phonetically balanced Harvard sentences [42] as the source text for synthesis.

4.2. Model training

We trained the proposed vocal effort control models for both voices. We use phone-wise duration, pitch, and energy as the fine-grained features, and utterance-wise pitch, pitch range, phone duration, speech energy, and spectral tilt as the higher-level prosodic features. 80-dimensional Mel-spectrograms are computed from pre-emphasized speech using STFT with 25 ms frame length and 10 ms shift. The encoder has 4 feed-forward Transformer layers each with a self-attention layer having 2 attention heads and 256 hidden units, and two 1-D convolution layers each having a kernel size of 9 and 1024 filters. The decoder has 2 dilated convolution blocks with six 1-D convolution layers with dilation rates of 1, 2, 4, 8, 16 and 32, respectively, kernel size of 3, and 256 filters. The feature predictors have two 1-D convolution layers each having a kernel size of 3 and 256 filters. The dropout rate of 0.2 and ϵ for layer normalization of 10^{-6} were used. We trained the models for 300k steps using 16 GPUs and a batch size of 512. We use WaveRNN [41, 43] to generate speech from the Mel-spectrograms, trained separately for each speaker. The $[M-3\sigma, M+3\sigma]$ spectral tilt values for Voice 1 and 2 are $[-0.984, -0.926]$ and $[-0.990, -0.931]$, respectively. These are mapped to $[-1, 1]$, where a positive value represents flatter spectral tilt and thus higher vocal effort.

4.3. Systems

We use the following systems in our evaluation:

1. **Baseline:** Synthetic speech without modification [22].
2. **Proposed:** Synthetic speech with increased vocal effort.
3. **SS:** Baseline processed with spectral shaping (SS) [14].
4. **SSDRC:** Baseline processed with spectral shaping and dynamic range compression (SSDRC) [14].

The baseline represents high-quality neural TTS [22] without any adaptation to listening conditions. The proposed system uses an increased vocal effort of +3, which means that the aimed vocal effort is 9σ higher than the median of the data. To compare the effectiveness of the proposed method to traditional approaches, we integrated two post-processing intelligibility enhancement modules to the baseline system, SS and SSDRC, which have been shown to provide state-of-the-art results [1, 2].




Figure 2: Means and 95% confidence intervals of normalized spectral tilt (vocal effort) measured for all systems, including various target spectral tilt values for the proposed TTS system.

The SS algorithm comprises three-stage filtering of the input speech in the short-time Fourier transform (STFT) domain. Two of the filtering stages perform phoneme-adaptive spectral shaping based on the probability of voicing, while the final filter is a fixed spectral shaper. On the adaptive spectral shaping, the local maxima (reflective of formants) are sharpened by a spectral sharpening filter followed by a high-frequency booster. Both of the filters update their coefficients based on the voicing probability of each frame. A non-adaptive pre-emphasis filter then modifies the spectra by enhancing the frequency components between 1000 Hz and 4000 Hz by a factor of 12 dB, while reducing the energy for the frequencies below 500 Hz by 6 dB/octave. Inverse Fourier transform and overlap-add are applied to get the final spectral-shaped waveform.

The DRC of the SSDRC is a time-domain operation where the objective is to reduce the envelope variation of speech. This is done by modifying the speech waveform in each frame, adaptive to the temporal envelopes. DRC is also a two-step process. In the first stage, the time-domain envelope is dynamically compressed with recursive smoothing. The smoothed envelope is projected onto the input-output envelope characteristic (IOEC) curve which gives the DRC gain. Finally, the spectral-shaped waveform from the SS module is multiplied by the estimated gains in the DRC to generate the SSDRC waveform.

4.4. Objective measures

We used the 720 Harvard utterances [42] to synthesize speech from each system and voice. In addition, to demonstrate the change in the vocal effort in the proposed system, we generated synthetic speech with various target spectral tilt values from -3 to +3, and then measured the output spectral tilt. The results, illustrated in Fig. 2, show that the (normalized) spectral tilt, and therefore the vocal effort, increases as the target vocal effort is increased. However, the SS and SSDRC systems have higher (flatter) spectral tilt in comparison to the proposed TTS system with adjusted spectral tilt of +3. We also calculated the long-term average spectra (LTAS) over all the speech samples for each system and voice. The LTAS, shown in Fig. 3, illustrate that all methods show a shift of energy from low to mid and high frequencies in comparison to the baseline.

4.5. Speech intelligibility in noise

We generated speech samples for each system and each speaker using the 720 the Harvard utterances [42]. We used two types of masking noises in the test: a fluctuating masker using competing speaker (CS) and a stationary masker using speech shaped noise (SSN). The CS masker was recorded read speech with alternate voice (Voice 2 for evaluating Voice 1, and vice versa) with SNRs of -7, -14, and -21 dB. The SSN masker was con-




Figure 3: Long-term average spectra over all speech samples.

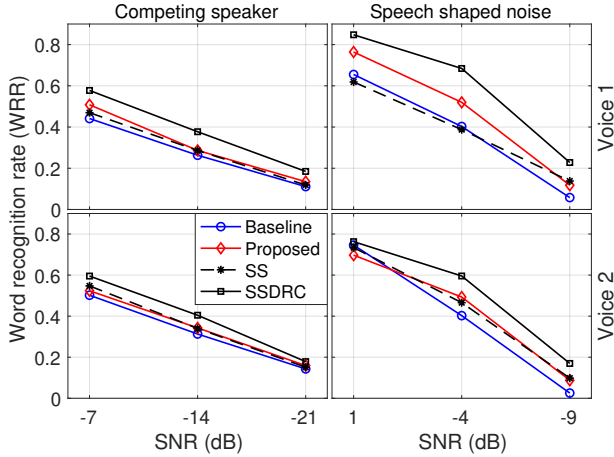


Figure 4: Word recognition rates (WRR) for all conditions.

structured by measuring the LTAS of the alternate voice with linear prediction and then filtering white noise with the resulting spectrum. The SNRs used were 1, -4, and -9 dB. The active speech level of the target samples were normalized according to ITU-T P.56 standard [44] before adding the masking noise according to the SNRs. The target speech content was centrally embedded in the noise with 0.5 seconds of lead and lag time. Overall, the test design was similar to [1].

We organized a large intelligibility test where test subjects used a graphical UI and were asked to listen to a speech sample in noise only once, and then type what they heard. Each listener was restricted to evaluate only a maximum of 720 utterances so that they never heard the same target utterance twice. The total number of samples in the test was 34,560, consisting of 2 voices \times 4 systems \times 2 noise types \times 3 SNRs \times 720 utterances. We added 10 reference speech samples without noise for each listener to evaluate their performance. The listeners had to achieve 80% word recognition rate (WRR) on the reference samples and more than 10% WRR on the test samples in order to qualify.

A total of 29,520 ratings were given by 41 qualified listeners using earbuds (47%), headphones (38%), or loudspeakers (16%). The WRR was measured as the proportion of content words that were recognized correctly and in the right order. The intelligibility test results are shown in Fig. 4.

Linear mixed effects models with *system* as the fixed effect and *voice*, *listener*, *utterance* and *audio device* as random intercepts were fit to the data using the lme4 package [45]. Comparing models with and without *system*, a likelihood ratio test ($\chi^2 = 1105.2$, $p < .0001$) and a drop in Akaike information criterion showed that *system* significantly predicted WRR. Dunnett’s multiple comparison test from the emmeans package [46] shows that the proposed system had significantly better intelligibility than baseline and SS on average ($p < .0001$), while SSDRC had significantly better intelligibility ($p < .0001$) on average than




Figure 5: Means and 95% confidence intervals of MOS.

other systems. As expected, the audio device had a clear effect on intelligibility, with the following average WRRs: earbuds 37.9 ± 0.6 %, headphones 42.5 ± 0.7 %, speakers 32.8 ± 1.0 %.

4.6. Quality without noise

To measure the effect of the intelligibility enhancement on the overall quality of speech without noise, we performed a mean opinion score (MOS) test with each of the 4 systems and 2 speakers using the 720 Harvard sentences. A 5-point MOS test was performed by 68 individual American English native speakers (51 female, 17 male, 18-69 years old) using either headphones (33) or loudspeakers (35). Listeners provided 5 ratings per utterance, resulting in a total of 28,800 responses.

The results are shown in Fig. 5. Linear mixed effects models with *system* and *voice* as fixed effects and *listener* and *utterance* as random intercepts were fit to the data. Likelihood ratio tests showed no significant interaction between *system* and *voice* ($p = .57$), and significant effects of *system* ($p < .0001$). Pairwise comparisons using the Tukey adjustment method showed significant differences between all systems ($p < .0001$). In the absence of noise, synthetic speech without any modification is preferred by listeners, which we would expect, but the proposed method was significantly preferred above other intelligibility enhancement methods (across both voices).

5. Discussion and conclusions

The focus of our experiment was on the average effect of intelligibility, and we do observe a consistent directional improvement for the proposed method across SNRs for both voices. Nonetheless, the data suggest that the strength of the effect may be contextual. In particular, the difference of most systems from baseline seems to generally be higher for Voice 1 than Voice 2. It may be that intelligibility enhancement methods work better with some voices than others. Our experiment was designed with sufficient power to detect main effects only, as that is the focus of this paper, and thus any exploration of the effect of context or voice is more speculative and worthy of exploration in future work. Also, developing a framework for automatically determining optimal vocal effort level with respect to masking noise type and level is left for future work.

We proposed a neural TTS system that models the natural vocal effort variation in the training data without specifically designed or labeled database. We use spectral tilt to measure vocal effort, and condition our models with spectral tilt among other prosodic features to factorize the prosodic space. By altering the input spectral tilt at inference time, we can successfully control the vocal effort of synthetic speech. Our experiments show that the proposed system achieved significantly higher word recognition rate than the baseline, while in the absence of noise, the proposed method achieves higher quality than other intelligibility enhancement methods.

6. References

- [1] M. Cooke, C. Mayo, and C. Valentini-Botinhao, “Intelligibility-enhancing speech modifications: The Hurricane Challenge,” in *Interspeech*, 2013.
- [2] J. Rennie, H. Schepker, C. Valentini-Botinhao, and M. Cooke, “Intelligibility-enhancing speech modifications – The Hurricane Challenge 2.0,” in *Interspeech*, 2020.
- [3] B. Langner and A. W. Black, “Improving the understandability of speech synthesis by modeling speech in noise,” in *ICASSP*, 2005.
- [4] E. Lombard, “Le signe de l’elevation de la voix,” *Ann. Maladies Oreille, Larynx, Nez, Pharynx*, vol. 37, no. 101–119, p. 25, 1911.
- [5] M. Cooke and Y. Lu, “Spectral and temporal changes to speech produced in the presence of energetic and informational maskers,” *J. Acoust. Soc. Am.*, vol. 128, pp. 2059–69, 10 2010.
- [6] B. Lindblom, *Explaining Phonetic Variation: A Sketch of the H&H Theory*. Springer Netherlands, 1990, pp. 403–439.
- [7] —, “Role of articulation in speech perception: Clues from production,” *The Journal of the Acoustical Society of America*, vol. 99, no. 3, pp. 1683–1692, 1996.
- [8] H. Traummüller and A. Eriksson, “Acoustic effects of variation in vocal effort by men, women, and children,” *J. Acoust. Soc. Am.*, vol. 107, no. 6, pp. 3438–3451, 2000.
- [9] C. T. Ishi, H. Ishiguro, and N. Hagita, “Analysis of the roles and the dynamics of breathy and whispery voice qualities in dialogue speech,” *EURASIP J. Audio Speech Music Process.*, vol. 3, pp. 1–12, 2010.
- [10] C. Gobl and A. Ní Chasaide, “The role of voice quality in communicating emotion, mood and attitude,” *Speech Comm.*, vol. 40, no. 1–2, pp. 189–212, 2003.
- [11] B. Sauert and P. Vary, “Near end listening enhancement: Speech intelligibility improvement in noisy environments,” in *ICASSP*, 2006.
- [12] M. Skowronski and J. Harris, “Applied principles of clear and Lombard speech for automated intelligibility enhancement in noisy environments,” *Speech Comm.*, vol. 48, pp. 549–558, 2006.
- [13] S. D. Yoo, J. R. Boston, A. El-Jaroudi, C.-C. Li, J. D. Durrant, K. Kovacyk, and S. Shaiman, “Speech signal modification to increase intelligibility in noisy environments,” *J. Acoust. Soc. Am.*, vol. 122, no. 2, pp. 1138–1149, 2007.
- [14] T.-C. Zorilá, V. Kandia, and Y. Stylianou, “Speech-in-noise intelligibility improvement based on spectral shaping and dynamic range compression,” in *Interspeech*, 2012.
- [15] R. J. Niederjohn and J. H. Grotelueschen, “The enhancement of speech intelligibility in high noise levels by high-pass filtering followed by rapid amplitude compression,” *IEEE Trans. Acous., Speech and Sig. Proc.*, vol. 24, no. 4, pp. 277–282, 1976.
- [16] T. Raitio, A. Suni, M. Vainio, and P. Alku, “Analysis of HMM-based Lombard speech synthesis,” in *Interspeech*, 2011.
- [17] —, “Synthesis and perception of breathy, normal, and Lombard speech in the presence of noise,” *Computer Speech & Language*, vol. 28, no. 2, pp. 648–664, 2014.
- [18] Q. Hu, T. Bleisch, P. Petkov, T. Raitio, E. Marchi, and V. Lakshminarasimhan, “Whispered and lombard neural speech synthesis,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*, 2021.
- [19] N. Campbell and P. Mokhtari, “Voice quality: The 4th prosodic dimension,” in *15th Intl. Congress of Phonetic Sciences*, 2003.
- [20] A. M. C. Sluijter and V. J. van Heuven, “Spectral balance as an acoustic correlate of linguistic stress,” *J. Acoust. Soc. Am.*, vol. 100 4 Pt 1, pp. 2471–2485, 1996.
- [21] T. Raitio, R. Rasipuram, and D. Castellani, “Controllable neural text-to-speech synthesis using intuitive prosodic features,” in *Interspeech*, 2020.
- [22] T. Raitio, J. Li, and S. Seshadri, “Hierarchical prosody modeling and control in non-autoregressive parallel neural TTS,” in *ICASSP*, 2022.
- [23] J.-C. Junqua, “The Lombard reflex and its role on human listeners and automatic speech recognizers,” *J. Acoust. Soc. Am.*, vol. 93, no. 1, pp. 510–524, 1993.
- [24] D. Rostolland, “Acoustic features of shouted voice,” *Acustica*, vol. 50, pp. 118–125, 1982.
- [25] W. V. Summers, D. Pisoni, R. Bernacki, R. Pedlow, and M. Stokes, “Effects of noise on speech production: Acoustic and perceptual analyses,” *J. Acoust. Soc. Am.*, vol. 84, no. 3, pp. 917–928, 1988.
- [26] E. Godoy, M. Koutsogiannaki, and Y. Stylianou, “Approaching speech intelligibility enhancement with inspiration from Lombard and clear speaking styles,” *Computer Speech & Language*, vol. 28, no. 2, pp. 629–647, 2014.
- [27] D. O’Shaughnessy, *Speech communications: Human and Machine, 2nd Edition*. IEEE, 2000.
- [28] C. V. Botinhao, J. Yamagishi, S. King, and R. Maia, “Intelligibility enhancement of HMM-generated speech in additive noise by modifying Mel cepstral coefficients to increase the Glimpse Proportion,” *Computer Speech & Language*, vol. 28, no. 2, pp. 665–686, 2013.
- [29] P. N. Petkov and W. B. Kleijn, “Spectral dynamics recovery for enhanced speech intelligibility in noise,” *IEEE Trans. Audio Speech Lang. Proc.*, vol. 23, no. 2, pp. 327–338, 2015.
- [30] C. Chermaz and S. King, “A sound engineering approach to near end listening enhancement,” in *Interspeech*, 2020.
- [31] M. Cooke, C. Mayo, C. V. Botinhao, Y. Stylianou, B. Sauert, and Y. Tang, “Evaluating the intelligibility benefit of speech modifications in known noise conditions,” *Speech Comm.*, vol. 55, pp. 572–585, 2013.
- [32] P. S. Chanda and S. Park, “Speech intelligibility enhancement using tunable equalization filter,” in *ICASSP*, 2007.
- [33] B. Sauert and P. Vary, “Near end listening enhancement optimized with respect to speech intelligibility index and audio power limitations,” in *Proc. Europ. Sig. Proc. Conf.*, 2010, pp. 1919–1923.
- [34] C. Taal, J. Jensen, and A. Leijon, “On optimal linear filtering of speech for near-end listening enhancement,” *IEEE Sig. Proc. Letters*, vol. 20, no. 3, pp. 225–228, 2013.
- [35] D. Paul, M. P. Shifas, Y. Pantazis, and Y. Stylianou, “Enhancing speech intelligibility in text-to-speech synthesis using speaking style conversion,” in *Interspeech*, 2020.
- [36] G. Li, R. Hu, S. Ke, R. Zhang, X. Wang, and L. Gao, “Speech intelligibility enhancement using non-parallel speaking style conversion with StarGAN and dynamic range compression,” in *ICME*, 2020, pp. 1–6.
- [37] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Agiomyriannakis, and Y. Wu, “Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions,” in *ICASSP*, 2017.
- [38] W. Ping, K. Peng, A. Gibiansky, S. O. Arik, A. Kannan, S. Narang, J. Raiman, and J. Miller, “Deep Voice 3: 2000-speaker neural text-to-speech,” in *ICLR*, 2018.
- [39] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, “FastSpeech 2: Fast and high-quality end-to-end text to speech,” *arXiv:2006.04558*, 2020.
- [40] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” in *SSW9*, 2016.
- [41] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. van den Oord, S. Dieleman, and K. Kavukcuoglu, “Efficient neural audio synthesis,” *arXiv:1802.08435*, 2018.
- [42] IEEE Subcommittee on Subjective Measurements, “IEEE recommended practices for speech quality measurements,” *IEEE Transactions on Audio and Electroacoustics*, vol. 17, pp. 227–46, 1969.
- [43] S. Achanta, A. Antony, L. Golipour, J. Li, T. Raitio, R. Rasipuram, F. Rossi, J. Shi, J. Upadhyay, D. Winarsky, and H. Zhang, “On-device neural speech synthesis,” in *ASRU*, 2021.
- [44] ITU-T, “Objective measurement of active speech level,” Recommendation P.56, 1993.
- [45] D. Bates, M. Mächler, B. Bolker, and S. Walker, “Fitting linear mixed-effects models using lme4,” *Journal of Statistical Software*, vol. 67, no. 1, pp. 1–48, 2015.
- [46] R. V. Lenth, *emmeans: Estimated marginal means, aka least-squares means*, 2022, R package version 1.7.2.