



HAL
open science

BERT, can HE predict contrastive focus? Predicting and controlling prominence in neural TTS using a language model

Brooke Stephenson, Laurent Besacier, Laurent Girin, Thomas Hueber

► To cite this version:

Brooke Stephenson, Laurent Besacier, Laurent Girin, Thomas Hueber. BERT, can HE predict contrastive focus? Predicting and controlling prominence in neural TTS using a language model. Interspeech 2022 - 23rd Annual Conference of the International Speech Communication Association, Sep 2022, Incheon, South Korea. pp.3383-3387, 10.21437/Interspeech.2022-10116 . hal-03791472

HAL Id: hal-03791472

<https://hal.science/hal-03791472>

Submitted on 29 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

BERT, can HE predict contrastive focus? Predicting and controlling prominence in neural TTS using a language model

Brooke Stephenson^{1,2}, Laurent Besacier^{2,3}, Laurent Girin¹, Thomas Hueber¹

¹Université Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, 38000 Grenoble, France

²LIG, UGA, G-INP, CNRS, INRIA, Grenoble, France

³NAVER LABS Europe, Meylan, France

brooke.stephenson@grenoble-inp.fr, thomas.hueber@grenoble-inp.fr,
laurent.girin@grenoble-inp.fr, laurent.besacier@univ-grenoble-alpes.fr

Abstract

Several recent studies have tested the use of transformer language model representations to infer prosodic features for text-to-speech synthesis (TTS). While these studies have explored prosody in general, in this work, we look specifically at the prediction of contrastive focus on personal pronouns. This is a particularly challenging task as it often requires semantic, discursive and/or pragmatic knowledge to predict correctly. We collect a corpus of utterances containing contrastive focus and we evaluate the accuracy of a BERT model, finetuned to predict quantized acoustic prominence features, on these samples. We also investigate how past utterances can provide relevant information for this prediction. Furthermore, we evaluate the controllability of pronoun prominence in a TTS model conditioned on acoustic prominence features.

Index Terms: text-to-speech, language model, BERT, prosody, contrastive focus, control.

1. Introduction

“HE loves her”, “he LOVES her”, and “he loves HER” all have the same textual content, but three distinct communicative goals. Indeed, such *contrastive focus* is used by speakers to evoke alternative sets in the discourse [1]. This can be utilized to make explicit intended discourse relations between clauses/paragraphs/sections, to highlight a fact that the listener may find surprising, or to express a specific semantic or pragmatic meaning. The prediction of contrastive focus placement therefore often requires high-level linguistic understanding. Current vanilla neural text-to-speech (TTS) synthesis systems lack this understanding and will always pronounce the above sentences in the same way, irrespective of the context. In this work, we investigate methods to predict the placement of contrastive focus and to control it in a TTS system. Figure 1 illustrates our overall approach which addresses both *predicting* and *controlling* prominence.

One way to insert sophisticated linguistic information may be through the use of contextualized word embeddings. While other works have explored the use of transformer language models to predict prosodic and stylistic features [2, 3, 4, 5, 6, 7, 8], it has not been fully explored how much the encoded word representations actually imbue high-level knowledge. In other words, do they provide information about the content and context of the message or do they only provide/reinforce low-level linguistic features such as the likelihood of lexical prominence, parts of speech and position in the sentence? For prominence/pitch accent prediction, a fairly high baseline can be achieved using word majority/accent-ratio alone (i.e., if a

lexical item is usually prominent in the training set, it is likely to be prominent in the test set) [9]. Moreover, [10] found that in a binary pitch accent prediction task, using broad word category distinctions (open/content or closed/function) could achieve 68% accuracy; more fine-grained division of the closed class category brought that number up to 77%. In this work, we probe a language model, in the present case BERT [11], by choosing a testing ground that cannot rely on simple heuristics to achieve good results: the prediction of contrastive focus on personal pronouns. Personal pronouns, at least in the corpora typically used to train TTS systems, are majority non-prominent.¹

Contrastive focus in English is mainly realized prosodically through increases in pitch, energy and/or duration. Suni et al. [12] have proposed extracting prosodic information automatically through a combined representation of these signals. They use continuous wavelet transforms, which analyze the speech signal at different timescales, to identify acoustically prominent words. In [4], they tested the prediction capabilities of a language model on these features. Here, we extend this investigation by looking at a challenging set of contrastively focused pronouns and by looking at the use of an extended context (current sentence, +previous sentence, +2 previous sentences). Moreover, because the *prediction* of contrastive focus is complicated by the one-to-many issue that hinders TTS evaluation in general, we have collected a corpus of audiobooks where we have three separate renditions of each book, each read by a separate speaker. We use these three realizations to study consensus among speakers in the placement of contrastive focus.

Assuming we are able to predict good candidates for contrastive focus, it remains to be verified if we can *control* the contrastive prominence on pronouns in a neural TTS system. Other works have proposed utterance-level control systems [13, 14] but for this task we need to target the word-level. This was achieved previously within an HMM framework [15]. [3] evaluated the expression of prosodic features from ToBI-label conditioned TTS, but only for ground-truth input labels. [16] report successful control over other word categories from systems conditioned on human or automatically annotated data (although this was not evaluated with a perceptual test). More recently, [17] proposed to use control tags for prominence in an end-to-end TTS system. However, their system relied on the availability of a specific (and limited) pre-annotated corpus whereas we demonstrate controllability from prominence labels obtained automatically.

¹We conflate the notions of prominence and contrastive focus for personal pronouns; when they are prominent, they typically possess a contrastive meaning.

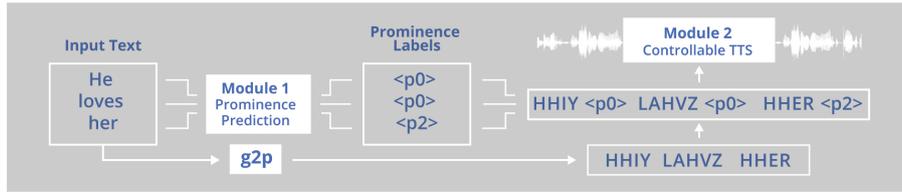


Figure 1: TTS overview: The system is split into two modules. The first uses a language model to predict prominence labels. The second controls the prominence in the synthetic speech in accordance with the predicted labels.

2. Prepared Datasets

The selected audiobooks for our corpus are literary texts sourced from Librivox² and from the Blizzard Challenge 2013 dataset.³ Criteria for book selection included open-source status, the availability of multiple recordings with different speakers (minimum 3), audio quality and the subject matter (we favoured books dealing with interpersonal relationships as they are more likely to contain contrastively focused pronouns). Five novels (x 3 speakers) were selected for the training set (41,593 utterances, approx. 66 hours of audio/speaker⁴ and one novel (x 3 speakers) was selected for testing (6838 utterances, approx. 11 hours of audio per speaker). The corresponding transcriptions were obtained from Project Gutenberg.⁵ Transcripts were split into chapters and then sentences using [18] and [19]. The audio files were segmented into utterances using the Aeneas library⁶ and phoneme alignment was obtained using the Montreal Forced Aligner [20].

Prosodic feature extraction. The audio files were analyzed with the continuous wavelet transform (CWT) method of [12] implemented in the Wavelet Prosody Toolkit.⁷ This method assigns a prominence score to each word in the corpus. This is done by combining f0, energy and duration into a composite signal, performing the CWT, establishing lines of maximum amplitude connecting the various timescales and then calculating a weighted sum of the points in this line (see [12] for details). We then quantize this score into three categories: <p2> (strong prominence), <p1> (intermediate prominence) and <p0> (no prominence).⁸

Contrastive personal pronoun subcorpus. With the processed data from the previous section, we searched for utterances containing <p2> labelled personal pronouns (strong prominence) in the test set. With manual verification, we collected positive examples of contrastive focus on pronouns. We randomly selected an equal number of negative samples where the pronoun was tagged as <p0> for all three speakers; these samples were also manually checked. We then enlisted three native English speakers to validate 200 pronouns (x 3 speakers) from the collected samples (100 positive samples and 100 negative samples, just over 20% of the full pronoun corpus). Validators were presented with the audio clips (for all three speakers) and a transcription of the text with the pronoun of interest highlighted in red. They were asked to assign a value of 1 if they deemed the speaker had used prosody to convey contrastive focus and

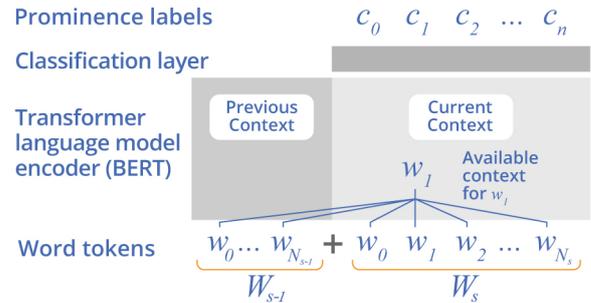


Figure 2: Module 1: Predicting prominence

0 if they did not. Cohen-kappa scores [21] were used to evaluate inter-annotator agreement between the three raters. These scores range from 0.85–0.90; this shows strong to almost perfect agreement. For evaluation purposes, we sorted the positive samples into two groups: 1) those where the majority of speakers (at least 2 out of 3) used contrastive focus (**Pronoun maj.**). This group contains 393 pronouns; 310 of which are prosodically contrasted by all three speakers; 2) those where only 1 out of 3 speakers used contrastive focus (**Pronoun min.**). This group contains 100 pronouns. All contrastive pronouns come from 406 utterances as several utterances contain multiple examples.

In this study, our intention is to find challenging examples for prominence prediction (i.e., words that are not frequently prominent). Subjective pronouns (e.g., I, we), objective pronouns (e.g., me, us) and possessive determiners (e.g., my, our) all fit this requirement, but possessive pronouns (e.g., mine, ours) are more often prominent than not. We decided to include these “easy” words in the corpus because they may be of interest for future work on contrastive focus. However, they do not present a particular challenge to prominence prediction.

3. Predicting Prominence

BERT. For our prominence prediction task, we used English BERT [11], available on HuggingFace.⁹ BERT is a transformer encoder that is trained on a masked language modeling task (i.e., it learns to predict masked words using information from the other words in the sentence). Bert makes use of self-attention layers and positional embeddings to form contextualized representations of words.

Prediction task. For the sequence of words $\{w_0, w_1, \dots, w_N\}$ our objective is to predict a sequence of prominence labels $\{c_0, c_1, \dots, c_N\}$, where c_n is either <p0>, <p1> or <p2>. Since knowledge about the previous context is sometimes es-

²<https://librivox.org>

³https://www.synsig.org/index.php/Blizzard_Challenge_2013

⁴Two of the five book sets are single speaker. The third set is made up of multiple speakers.

⁵<https://www.gutenberg.org>

⁶<https://github.com/readbeyond/aeneas>

⁷https://github.com/asuni/wavelet_prosody_toolkit

⁸Dataset available at <https://doi.org/10.5281/zenodo.6646827>.

⁹<https://huggingface.co/bert-base-cased>

Table 1: Results for the prominence prediction task for the <p2> (high prominence) category. Recall (R), Precision (P) and F1 are reported for the full test set (all POS categories combined).

Model	Current sentence			+previous sentence			+2 previous sentences			# p2 tags	# words
	F1	R	P	F1	R	P	F1	R	P		
Word majority	0.458	0.394	0.546							79,793	365,330
Randomly initialized BERT	0.572	0.539	0.609	0.565	0.546	0.586	0.561	0.558	0.565	79,793	365,330
Fine-tuned BERT	0.588	0.552	0.629	0.580	0.536	0.632	0.579	0.535	0.629	79,793	365,330

Table 2: Results on pronouns only for the prominence prediction task. Recall is reported for the two manually verified subsets of contrastively focused personal pronouns and the non-contrastive pronoun subset (Neg: 493 samples). Maj: group of 393 samples where majority of speakers used contrastive focus. Min: group of 100 samples where only 1 out of 3 speakers used contrastive focus. Context: current (Curr), previous (Prev)

Model	Data	Cat.	R		
			Curr	+1 Prev	+2 Prev
Word Majority	Maj	<p2>	0.079		
	Min	<p2>	0.000		
	Neg	<p0>	1.000		
Randomly Initialized BERT	Maj	<p2>	0.178	0.160	0.135
	Min	<p2>	0.060	0.060	0.060
	Neg	<p0>	0.980	0.980	0.982
Fine-tuned BERT	Maj	<p2>	0.239	0.239	0.216
	Min	<p2>	0.060	0.070	0.070
	Neg	<p0>	0.990	0.986	0.988

essential for determining whether a word should be contrastive or not, we experiment with models conditioned on different degrees of past context. We note $W_s = \{w_1, \dots, w_{N_s}\}$ the sequence of N_s words in the current sentence s which we want to synthesize. The BERT model is given either the current sentence only, i.e., W_s , or both the previous and the current sentences, i.e., $\{W_{s-1}, W_s\}$, or both the two previous sentences and the current one, i.e., $\{W_{s-2}, W_{s-1}, W_s\}$. See Figure 2 for an illustration of the model’s components.

Models and linguistic knowledge. We evaluate three methods for prominence prediction with increasing access to linguistic knowledge:

- Our baseline is a simple word majority method: word statistics from the training corpus are computed; we count how often a lexical item belongs to each of the three prominence categories and the majority category is used for all predictions in the test set.
- The second method involves the use of the BERT architecture, but instead of using weights pretrained on a masked language modelling task, we randomly initialize the model and train it to predict the prominence labels for each word in the input sequence. This model can presumably learn the same word statistics available to the word majority method and additionally the self-attention layers and positional embeddings provide the model with information about the surrounding lexical items and the positional context of each word. We expect this model will be able to learn canonical patterns of English prosody (i.e. that prominent, nuclear accents are typically found at the end of an intonational phrase) even if the semantic knowledge about the content of the sentences will be non-existent or at best, very naive.
- The third method involves finetuning a pretrained BERT

model on the prominence features (learning rate = $5e-5$). From the beginning of training, this model has access to the syntactic and semantic representations learned from training on huge amounts of textual data; during the finetuning process it must find the optimal way to use this information for the prosody/prominence prediction task.

Results. The results for the full dataset are shown in Table 1 and the results for the pronoun subsets are shown in Table 2. Analyzing these results, we notice that performance on the contrastive pronoun sets is significantly lower than the full dataset (Recall with finetuned BERT is 0.239 for the pronoun majority group and 0.552 for all POS). The tokens in the pronoun minority group are very rarely predicted to be prominent (highest recall score = 0.07). Furthermore, we see that BERT does provide some improved prediction accuracy over the two baselines, but the improvement is fairly small. Word majority method correctly predicts possessives (e.g., mine, yours), and the randomly initialized BERT learns structurally/positionally prominent positions; it correctly predicts prominence at the ends of prosodic phrases (immediately preceding punctuation marks) and following the word ‘for’ (e.g., “For YOU are...”). We performed McNemar’s Test [22] to compare models and we find significant differences (p-values < 0.05) between each of the three models. With regards to the use of previous context, we do not see any improvements in prediction performance (p-values > 0.05).

We can imagine several possible causes for low prediction scores. It may be that we have an insufficient number of samples of contrastively focused pronouns to train the model to recognize focus patterns; there is an average of 7893.6 <p2> labelled personal pronouns/speaker in the training set, and given the complexity of the task, this may not be enough. Alternatively, learnt representations may not be sophisticated enough to encode the higher level linguistic information required for this task. The recent research trend in language modeling is to scale models bigger and bigger and this increase in size results in better quality on tasks such as text generation. In future work, we will explore using larger LMs. Finally, as can be expected with any automatic annotation method, there is some noise in the data: we did find examples for which prominence was questionable, predominately at phrase boundaries (tagged <p2> because of a sharp rise in f0). Hence, human intervention may still be necessary for better fine-grained annotation/control of prosodic data.

4. Controlling Prominence

Controllable TTS Model. To synthesize speech with controllable prominence, we follow the method proposed in [16] where the TTS is conditioned on prominence labels. Our implementation differs in that we use FastSpeech 2 [23] (as implemented by [24]) instead of DC-TTS and we refrain from using boundary tags as we are primarily interested in prominence control. FastSpeech 2 is a non-autoregressive, transformer encoder-decoder model that makes explicit duration, f0 and energy predictions

between the encoder and decoder. The input to our model is a sequence of phonemes and prominence labels. Each word w_n in the utterance is converted into a sequence of phonemes $\{p_{n,0}, \dots, p_{n,m}\}$ and this phoneme sequence is followed by the prominence label c_n for w_n (phonemes and prominence labels are converted into embeddings in the first layer of the model). The output of the modified FastSpeech 2 model is a Mel-spectrogram, which is converted into a waveform using a Parallel WaveGAN vocoder [25].¹⁰

To train and test the TTS model, we use the data described in Section 2, but only for a single speaker (Blizzard Challenge 2013; this speaker has read all 6 books). While the training corpus is read by a single speaker, this speaker portrays several different characters with different accents and pitch ranges. This tends to introduce fuzziness into the synthetic speech. To help the model learn these characteristics and improve quality, we added a speaker embedding to the TTS input. To obtain a speaker embedding, we 1) encoded each utterance in the training set with a pretrained speaker identification model (ECAPA-TDNN [26] available at [27]), 2) used k-means clustering on these embeddings to obtain 30 different ‘speakers’ and 3) used these speaker labels as an additional input to FastSpeech 2; the speaker embeddings are concatenated with the phoneme and prominence label embeddings.

Listening Test. To test the controllability of our TTS model in terms of prominence, we conducted an ordinal ranking listening test using the Web Audio Evaluation Tool [28], following this procedure: 1) we randomly sampled 100 utterances containing pronouns from our full test set (not solely from the contrastive subset); 2) from this selection, we took the first 10 utterances containing a subjective pronoun, the first 10 with an objective pronoun and the first 10 with a possessive determiner (for a total of 30 utterances); 3) using our pretrained TTS system, we synthesized three versions of each utterance, changing only the prominence label for the relevant pronoun ($c \in \{\langle p0 \rangle, \langle p1 \rangle, \langle p2 \rangle\}$). The prominence labels for all other words in the sentence were kept constant with the ground truth values (extracted from the original audio with the CWT method); 4) 30 native English evaluators, recruited on Prolific,¹¹ were presented with the three versions of the synthesized utterances (in random presentation order) and a transcript of the audio with the pronoun of interest in uppercase letters. Participants were asked to rank the prominence of the pronoun by dragging and dropping the movable audio clips so that they were arranged from most prominent to least; 5) clips ranked as most prominent are assigned a score of 1; clips ranked as second most prominent are assigned a score of 0.5 and clips ranked as least prominent are given a score of 0. Sample audio files are available in the supplementary multimedia materials.

Results. The results of the listening test are shown in Figure 3. We see the median values align with the prosody labels used (median: $\langle p0 \rangle = 0$, $\langle p1 \rangle = 0.5$, $\langle p2 \rangle = 1.0$). We do however see a wide distribution in the responses. This, and the examination of the ratings for individual utterances, indicates that this method works, but not consistently (i.e., there are some utterances for which the evaluators could not detect a difference). The only pronoun category for which we see a fairly clear distinction between $\langle p0 \rangle$, $\langle p1 \rangle$, and $\langle p2 \rangle$ in perceived prominence is for possessive determiners. This may be because there is more natural variation within the training corpus for this category. Or, it may be due to the labelling errors at phrase bound-

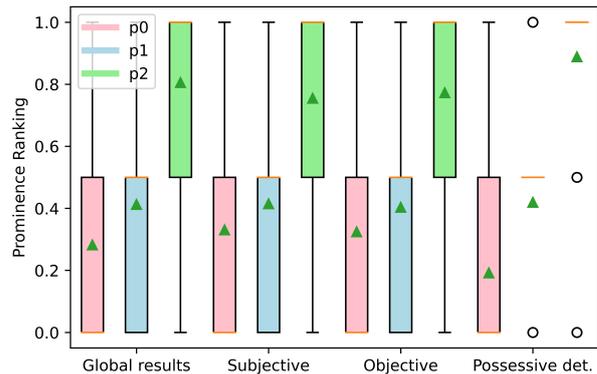


Figure 3: Results from the prominence ordinal ranking listening test. Audio ranked as most prominent was assigned a score of 1, second most prominent 0.5 and least prominent 0. Orange lines show the median and green triangles show the mean.

aries discussed in the previous section: the sampled subjective and objective pronouns were found more often at the beginning and end of phrases than the possessive determiners. More work is thus needed to disentangle the global prosodic representations from that of individual words, but this separation is difficult because it may result in less natural utterances.

5. Conclusion and Perspectives

In this paper, we investigated the difficult tasks of prominence prediction and control for contrastively focused personal pronouns. **Prediction:** we compared models with varying degrees of access to linguistic knowledge and past context on a word-level prominence label prediction task. We found that a fine-tuned BERT gave the best prediction performance, but that the improvement over baselines was very small. In the future, we will investigate the use of word representations from larger language models with more sophisticated linguistic understanding. **Control:** with a perceptive test, we evaluated the control of prominence in a TTS model conditioned on prominence labels. The results show the model is able to provide some control but the performance is not consistent over all samples. **Natural variation:** we used multiple spoken versions of the same written text to see the agreement among speakers on the use of contrastive focus. But we must keep in mind that while the textual context remains the same, the interpretation of the text can vary. For example, we infer that one of the speakers interprets some of the characters in the novel to be passive aggressive and they convey this through the frequent use of contrastive focus on ‘I’ (e.g., **I** am going to Gretna Green (intended meaning: and **YOU** are not). Removing the contrastive focus here is not wrong, but gives a very different impression of character/situation/relationship. This illustrates the difficulty of the prediction task and therefore, depending on the intended usage of the TTS system, it might be fruitful to explore other sources of input to the prominence prediction model (e.g., the source speech in a speech-to-speech translation system) in order to be as faithful to the intended meaning as possible.

6. Acknowledgements

Work was funded by the Multidisciplinary Institute in Artificial Intelligence MIAI@Grenoble-Alpes (ANR-19-P3IA-0003).

¹⁰<https://github.com/kan-bayashi/ParallelWaveGAN>

¹¹<https://www.prolific.co>

7. References

- [1] M. Rooth, "A theory of focus interpretation," *Natural Language Semantics*, vol. 1, no. 1, pp. 75–116, 1992.
- [2] T. Kenter, M. Sharma, and R. Clark, "Improving the prosody of RNN-based english text-to-speech synthesis by incorporating a BERT model," in *Proc. of Interspeech*, Shanghai, China (virtual conference), 2020.
- [3] Y. Zou, S. Liu, X. Yin, H. Lin, C. Wang, H. Zhang, and Z. Ma, "Fine-grained prosody modeling in neural speech synthesis using ToBI representation," in *Proc. of Interspeech*, Brno, Czech Republic, 2021.
- [4] A. Talman, A. Suni, H. Celikkanat, S. Kakouros, J. Tiedemann, and M. Vainio, "Predicting prosodic prominence from text with pre-trained contextualized word representations," in *Proc. of Nordic Conference on Computational Linguistics (NoDaLiDa)*, Turku, Finland, 2019.
- [5] Y. Xiao, L. He, H. Ming, and F. K. Soong, "Improving prosody with linguistic and BERT derived features in multi-speaker based Mandarin Chinese neural TTS," in *Proc. of IEEE ICASSP*, Barcelona, Spain (virtual conference), 2020.
- [6] Z. Hodari, A. Moinet, S. Karlapati, J. Lorenzo-Trueba, T. Merritt, A. Joly, A. Abbas, P. Karanasou, and T. Drugman, "CAMP: a two-stage approach to modelling prosody in context," in *Proc. of IEEE ICASSP*, Toronto, Canada, 2021.
- [7] D. Stanton, Y. Wang, and R. Skerry-Ryan, "Predicting expressive speaking style from text in end-to-end speech synthesis," in *Proc. of IEEE Spoken Language Tech. Workshop (SLT)*, Athens, Greece, 2018.
- [8] T. Hayashi, S. Watanabe, T. Toda, K. Takeda, S. Toshniwal, and K. Livescu, "Pre-trained text embeddings for enhanced text-to-speech synthesis," in *Proc. of Interspeech*, Graz, Austria, 2019.
- [9] A. Nenkova, J. Brenier, A. Kothari, S. Calhoun, L. Whitton, D. Beaver, and D. Jurafsky, "To memorize or to predict: Prominence labeling in conversational speech," in *Proc. of the Conf. of the North Am. Chapter of the Assoc. for Computational Linguistics*, Rochester, New York, 2007.
- [10] J. Hirschberg, "Pitch accent in context: Predicting intonational prominence from text," *Artificial Intelligence*, vol. 63, no. 1–2, pp. 305–340, 1993.
- [11] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. of ACL*, Florence, Italy, 2019.
- [12] A. Suni, J. Simko, D. Aalto, and M. Vainio, "Hierarchical representation and estimation of prosody using continuous wavelet transform," *Computer Speech and Language*, vol. 45, pp. 123–136, 2017.
- [13] T. Raitio, R. Rasipuram, and D. Castellani, "Controllable neural text-to-speech synthesis using intuitive prosodic features," in *Proc. of Interspeech*, Shanghai, China (virtual conference), 2020.
- [14] Y. Wang, D. Stanton, Y. Zhang, R.-S. Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. Ren, and R. A. Saurous, "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," in *Int. Conf. on Machine Learning*. PMLR, 2018.
- [15] L. Badino, J. S. Andersson, J. Yamagishi, and R. A. Clark, "Identification of contrast and its emphatic realization in HMM based speech synthesis," in *Proc. of Interspeech*, Brighton, UK, 2009.
- [16] A. Suni, S. Kakouros, M. Vainio, and J. Šimko, "Prosodic prominence and boundaries in sequence-to-sequence speech synthesis," in *Proc. of ISCA Int. Conf. on Speech Prosody*, Tokyo, Japan, 2020.
- [17] S. Latif, I. Kim, I. Calapodescu, and L. Besacier, "Controlling prosody in end-to-end TTS: A case study on contrastive focus generation," in *Proc. of the Conf. on Computational Natural Language Learning (CoNLL)*, Punta Cana, Dominican Republic, 2021.
- [18] J. Reeve, "Chapterize," <https://github.com/JonathanReeve/chapterize>, 2016.
- [19] M. Honnibal, I. Montani, S. Van Landeghem, and A. Boyd, "spaCy: Industrial-strength Natural Language Processing in Python," <https://spacy.io/>, 2020.
- [20] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal Forced Aligner: Trainable text-speech alignment using Kaldi," in *Proc. of Interspeech*, Stockholm, Sweden, 2017.
- [21] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and Psychological Measurement*, vol. 20, no. 1, pp. 37–46, 1960.
- [22] Q. McNemar, "Note on the sampling error of the difference between correlated proportions or percentages," *Psychometrika*, vol. 12, no. 2, pp. 153–157, 1947.
- [23] Y. Ren, C. Hu, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "FastSpeech 2: Fast and high-quality end-to-end text-to-speech," *arXiv preprint arXiv:2006.04558*, 2020.
- [24] C. Wang, W.-N. Hsu, Y. Adi, A. Polyak, A. Lee, P.-J. Chen, J. Gu, and J. Pino, "FAIRSEQ S²: A scalable and integrable speech synthesis toolkit," *arXiv preprint arXiv:2109.06912*, 2021.
- [25] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *Proc. of ICASSP*, Barcelona, Spain (virtual conference), 2020.
- [26] B. Desplanques, J. Thienpondt, and K. Demuyne, "ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Proc. of Interspeech*, Shanghai, China (virtual conference), 2020.
- [27] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, "SpeechBrain: A general-purpose speech toolkit," 2021, arXiv:2106.04624.
- [28] N. Jillings, D. Moffat, B. De Man, and J. D. Reiss, "Web Audio Evaluation Tool: A browser-based listening test environment," in *Proc. of the Sound and Music Computing Conf.*, Maynooth, Ireland, 2015.