



# Prosodic Phrase Alignment for Machine Dubbing

Alp Öktem<sup>1,2,\*</sup>, Mireia Farrús<sup>2</sup>, Antonio Bonafonte<sup>3,^</sup>

<sup>1</sup>Col·lectivaT SCCL, Barcelona, Spain

<sup>2</sup>Universitat Pompeu Fabra, Barcelona, Spain

<sup>3</sup>Universitat Politècnica de Catalunya, Barcelona, Spain

{alp.oktem, mireia.farrus}@upf.edu  
antonio.bonafonte@upc.edu

## Abstract

Dubbing is a type of audiovisual translation where dialogues are translated and enacted so that they give the impression that the media is in the target language. It requires a careful alignment of dubbed recordings with the lip movements of performers in order to achieve visual coherence. In this paper, we deal with the specific problem of prosodic phrase synchronization within the framework of machine dubbing. Our methodology exploits the attention mechanism output in neural machine translation to find plausible phrasing for the translated dialogue lines and then uses them to condition their synthesis. Our initial work in this field records comparable speech rate ratio to professional dubbing translation, and improvement in terms of lip-syncing of long dialogue lines.

**Index Terms:** audiovisual translation, dubbing, spoken machine translation, prosody

## 1. Introduction

Introduction of machine mediated methods to audiovisual translation domain has made it possible to obtain transcriptions and translations for multimedia without the huge manual labor that they used to demand. This is especially useful for online video streamers who publish often and lack access to professional translation services. By making their videos transcribed, one could augment visibility to their content; and by making them translated, they can further make them accessible to audiences of different languages. While subtitles remain as the most accessible method of audiovisual translation, it does not suffice as a medium for certain viewers with visual impairments or reading difficulties. Also, due to various cultural and historic reasons, subtitles are not the preferred form of audiovisual translation in certain linguistic communities [1].

The alternative to subtitling is *dubbing*, where the dialogues in an audiovisual are translated in spoken form to replace its original audio track. It is a method widely used for the translation of e.g. movies, commercials and video games. However, unlike subtitling, there are relatively few automated solutions for it. This can be explained mainly through the further complexities that the spoken language output introduces. In dubbing, the translated content does not only need to convey the linguistic and semantic information but also needs to deliver them within a similar prosodic context that matches the unchanged visual content. A key factor in dubbing, referred to as *synchronization* or *lip-sync*, deals with the reproduction of timing, phrasing and phonetic content of the original speech segments in the target language to match with the lip movements of the

original performers [2]. This factor is especially relevant in the cases where the performer is on-screen and their articulations can be visually traced. In a majority of the cases, dubbing translation is made in such a way that it (1) contains a comparable number of syllables, (2) follows a similar phrasing and pausing structure, (3) matches mouth articulation movements (like opening, closing etc.) [3]. These requirements are especially important in movie-domain dubbing scenarios where the final content needs to sound and feel realistic in the target language.

In this work, we focus on the modelling of the cross-lingual prosodic phrasing structure in machine dubbing. In order to carry out experimentation, a prototype pipeline for obtaining dubbed versions of speaker segments of a TV show is presented. Our aim is not only to translate the spoken segments and synthesize them, but also to make the translated and synthesized segments to match the original durations and phrasings. The problem that requires a soft-alignment between translations is tackled by a simple mechanism that exploits the attention mechanism output during neural machine translation. Later, the durations of the input phrases are used to cue the text-to-speech (TTS) system so that the translated output is synthesized in a way that it respects the prosodic phrasing structure of the original take.

## 2. Automating Dubbing

Automatic dubbing has been previously proposed in various works without any translation back-end [4, 5, 6, 7]. General motivation behind these works is to generate an additional track for a given audiovisual for aiding viewers with auditory or visual impairment. This approach uses the segmented transcriptions in already available subtitles of the audiovisual to generate a new audio track with synthesized speech. This additional track containing clearer speech is laid over the original track of the media, which is either muted or turned down.

The only durational constraint dealt in previous works is the synchronisation of the synthesized segments with their respective subtitle timestamps. This is obtained by placing the synthesized segment at the beginning of each subtitle entry and then adjusting its speech rate only if it exceeds the duration of the subtitle entry. Matoušek et al. performs an online and offline TTS time compression technique in order to achieve this [5]. Although this methodology can be useful in news or documentary domain, it would lead to limitations for obtaining realistically synthesized segments needed in e.g. movie domain. For instance, subtitle segments that are considered as independent would lead to prosodically incoherent syntheses of long dialogue lines. Also, subtitle segmentation hinders prosodic phrasing only to some extent. Precise prosodic phrasing and phonetic alignment are not given attention in subtitle format translations.

\*Alp Öktem is currently at Translators without Borders.

^Antonio Bonafonte is currently at Amazon Research, UK  
(bonafont@amazon.com)

This paper, to the best of our knowledge, is the first study that approaches the automation of realistic cross-lingual dubbing. We outline a methodology for a possible application that could prove useful for media content creators and focus on the aspect of prosodic phrase synchronization, which is one of the core elements in dubbing.

### 3. Synchronization in Movie-domain Dubbing

As synchronization in dubbing is most relevant in movies, we base our analyses and experimentation on this domain.

#### 3.1. Audio Data

Both for analysis and experiments, we have used the *Heroes Corpus* [8], which is a collection of original and dubbed dialogue segments from a North American TV series. The series is originally in American English and is dubbed into many languages including Spanish. Heroes Corpus contains 7000 parallel English and Spanish single speaker audio segments, together with their transcripts and prosodic annotations. Each audio segment contains a dialogue line of an actor varying between 0.1 to 18 seconds (2.43s in average) uttering a minimum of 1 to a maximum of 45 words (5.43 in average). The prosodic annotations contain word alignment information with word-level silent pause duration information specified.

#### 3.2. Prosodic phrase alignment

We have defined the prosodic phrases that are studied in this paper as voiced segments terminated with a silent pause. Although this is a partial definition of a prosodic phrase, it suffices for modelling of voiced segment alignment of dubbing phrases.

A corpus-level analysis has been performed to see to what extent pausing is reflected between the original and dubbed segments of the dataset. This was done by measuring if a paused interval overlaps in both original and dubbed segment timewise. It was seen that for any pause longer than 0.1 seconds, it was 70% likely that it will be reflected in the dubbing translation. For pauses of 0.2 seconds in length this was 81% and for pauses of 0.4 seconds it was 91%, showing the attention given for the reflection of long pauses in dubbing movie dialogues.

Another key issue in dubbing is making sure that the translated content fits the original take length-wise. This is done by making sure that the translation contains more or less a similar number of syllables with a careful word selection [3]. In our dataset, we have measured that the average ratio of the number of syllables in parallel segments is 1.31 with a standard deviation of 0.39, showing a higher speech rate for Spanish.

## 4. Machine Dubbing Synchronization

The prototype machine dubbing setup that forms the basis of our synchronization methodologies is illustrated in Figure 1. The pipeline performs the translation of segmented and transcribed dialogue lines in the original version of the movie, like in the example below:

MATT: What are you doing in here? Where are we?
MOLLY: This is where he keeps me.
MATT: What?

Having timestamp and word alignment information on these segments makes it possible to infer the prosodic phrasing

structure in the dialogue line where silent pauses that fall between the uttered words are considered as prosodic phrase (PP) cut-off points. In order to differentiate between articulatory effects and linguistically motivated pauses, we have determined a pause duration threshold of 250 *ms* following the conventions of [9, 10, 11].

Dialogue lines are translated using a neural machine translation (NMT) system. After that, translated phrases are synthesized using a text-to-speech (TTS) system with durational cues that reflect the source utterance. These durational and segmentation cues are determined by the two methodologies explained in the following two subsections.

#### 4.1. Cross-lingual phrase alignment using attention

The first stage of phrasing synchronization deals with the determination of the PP boundaries in the target translation. The output of the machine translation is a sequence of tokens (words and punctuation) which needs to be split up into phrases that reflect the original input sentence structure. This necessitates a modelling of alignment between the input and output sentence.

Our methodology is based on exploiting the attention weights that is used in neural network based machine translation (NMT). Attention was first introduced for encoder-decoder based NMT [12] as a mechanism to help focus on different parts of the input at each step of decoding [13]. It has also been shown that it represents a structure relating with the alignment between the input and output phrases [14]. Designed to learn the relevance between parts of input and target sequence in training, it can be accessed during inference for obtaining alignment information between the source and predicted sentence. Figure 2 shows the visualization of the attention weights output in translation. The alignment of the phrases can be traced on the distribution of attention weights.

We define the cross-lingual phrase alignment process as follows: Given a source sentence with tokens  $W_e : \langle w_e^1, w_e^2, \dots, w_e^M \rangle$  and their respective PP labels  $L_e : \langle l_e^1, l_e^2, \dots, l_e^M \rangle$ , find the best PP label sequence  $L_f : \langle l_f^1, l_f^2, \dots, l_f^N \rangle$  of the translated sequence of tokens  $W_f : \langle w_f^1, w_f^2, \dots, w_f^N \rangle$ .

The procedure for automatic cross-lingual phrase alignment consists of two stages: (1) Population of the set of possible target PP label sequences  $S : \{L_f^1, \dots, L_f^X\}$  and (2) Scoring of the sequences in  $S$  in terms of which leads to a more appropriate segmentation. Finally, the best scoring sequence is selected as the target PP sequence.

Possible target prosodic phrasings need to follow certain requirements in order to obtain meaningful phrase alignments. Firstly, it is necessary that  $L_f$  contains the same number of unique PP labels as  $L_e$  and in the same order. This is obtained by listing all possible sequences starting from  $l_e^1$  and then eliminating the ones that don't follow a monotonic progression until last label  $l_e^M$  is reached. Secondly, it is made sure that tokens that need to fall inside the same PP are not assigned different labels. This is especially relevant in the cases where tokens represent sub-linguistic units.

Second stage consists of ranking each candidate PP label sequence  $s = \langle l_f^1, l_f^2, \dots, l_f^N \rangle \in S$ . The alignment information residing in attention matrix holds is used to define a score as follows:

$$score(s) = \prod_n \sum_m W_{masked}^l(n; m) \quad (1)$$

$W_{masked}^l$  is a masked version of the attention matrix  $W$  with respect to label  $l$ . This ensures obtaining of a higher score if

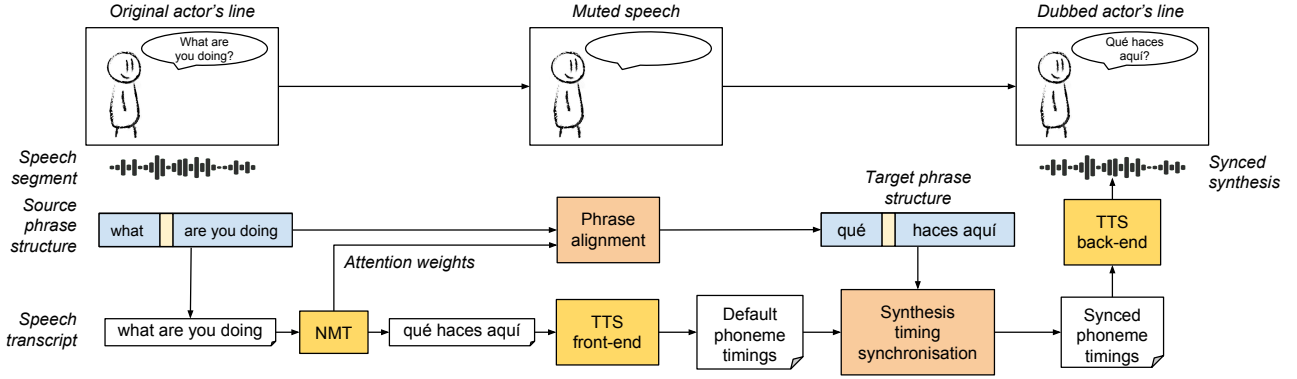


Figure 1: Prototype machine dubbing pipeline with prosodic phrase synchronisation.




Figure 2: Visualization of attention output during translation of a sample from the dataset. Tokens belonging to different prosodic phrases are denoted with the colors. Attention weights are higher in the areas where source and target phrases align.

the target tokens are labelled with the same label as the source tokens that they are aligned with.

$$W_{masked}^l(i; j) = \begin{cases} W(i; j) & \text{if } l = l_e^j \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

Finally, the best scoring PP label sequence is chosen as optimal PP label sequence.

$$L_f = \arg \max_{s \in S} score(s) \quad (3)$$

The ranking methodology is a way to average alignment scores in order to get a parallel phrasing structure. By defining allowed sequences first and then ranking them, complicated token alignments do not result in broken phrasings. For example, a target token that aligns with an outlier token (or doesn't align at all) gets evened out among its neighbouring tokens and thus end up in the same phrase with them.

#### 4.2. Synthesis timing synchronization

Synchronized audio segments are obtained by conditioning the synthesis with respect to a predefined phrasing and durational structure, as illustrated in Figure 3. First step in achieving this

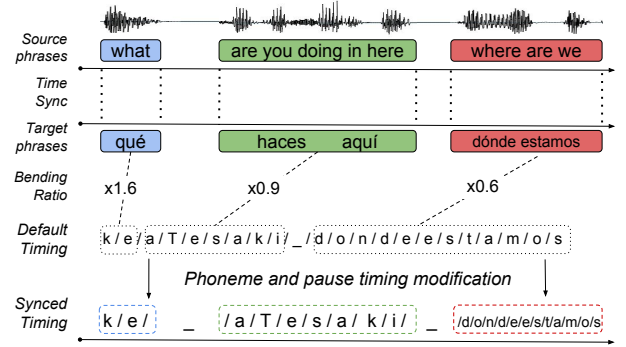


Figure 3: Durational conditioning on the target prosodic phrases with respect to source prosodic phrase and pause durations used for phoneme timing synchronization.

is to define the desired durational structure of the translation by directly mapping the durational parameters of the source PPs to the target PPs that align with them. This is to ensure that voice activations overlap as much as possible so that lip movements align with the synthesized tracks. The desired durational structure is then used to condition the phoneme and pause durations that are assigned by the TTS back-end. A *bending ratio* is calculated for each phrase by dividing the desired phrase duration by the duration of the same segment that TTS predicts. This ratio later determines how much the phonemes belonging to that PP need to be speed-up or slow-down. Finally, the modified phoneme durations are sent to the TTS front-end to obtain the synchronized audio segment.

## 5. Evaluation

### 5.1. Evaluation setup

**Machine translation** models were trained initially on the WITTED corpus [15] and then fine-tuned to our domain. Half of the Heroes Corpus was allocated for this purpose. As in the examples shown in previous section, we have only focused on the direction English→Spanish. OpenNMT-py toolkit [16] was used for obtaining the models. The MT model consists of a 3-layer attentional encoder-decoder network with 512-unit bidirectional LSTMs [12, 17, 13] and employs beam search, coverage attention [18] and input feeding [19]. A shared *subword*

Table 1: Datasets used for training and in-domain adaptation of MT models, their training/testing set sizes and BLEU scores obtained in test sets.

Dataset	# Tokens (Train/test set)	BLEU (%)
WIT-TED	8.4M/366K	34.04
Heroes	59K/62K	24.14




Figure 4: Speech rate comparison with respect to syllable count (blue) and synthesis time modification ratio (red) distribution between source and automatically aligned segments. Mean values are shown with dashed lines.

vocabulary was trained for both languages using *SentencePiece* byte-pair encoding (BPE) [20, 21]. The sizes of the training corpora and the BLEU scores obtained with them on separate test sets is listed in Table 1.

**Text-to-speech** system was chosen among free alternatives that would allow us to modify phoneme durations before the synthesis takes place. For this, we have opted for the *mbrola* [22] system as a TTS back-end. Mbrola is a research-oriented TTS system that performs synthesis from pre-calculated linguistic and prosodic features without any text processing. As a front-end we used the *espeak*<sup>1</sup> tool to perform morpheme-to-phoneme conversion, phoneme time duration and *f0* contour prediction.

## 5.2. Analysis

For evaluation, we have allocated 3490 segments of the Heroes corpus for evaluation and performed various manual and experimental analyses on them. We first wanted to see how much the aligned prosodic phrases match in terms of their length and how much of a speed-up or slow-down rate has been applied to make their durations fit during synthesis synchronization. Figure 4 shows the density graph of the speech rate ratios and bending ratios calculated for each matching prosodic phrase in the evaluation set. Speech rate ratio was calculated with respect to syllables. It can be seen that average speech rate ratio (1.27) corresponds to the original ratio measured in Heroes Corpus (1.31). It is also seen that the phrase durations predicted by the TTS front-end were more likely to be sped-up than slowed-down.

A perception test has also been prepared to compare our methodology to a subtitle reading approach similar to [4, 5, 6]. We have selected 10 samples from our evaluation set and cut out the video portions corresponding to them. Two dubbed audio tracks were prepared for each sample: One with the machine translated and synced synthesis using our methodology and another one with synthesized Spanish subtitles with their starting times dictated by their original timestamps<sup>2</sup>. 18 Spanish speak-

<sup>1</sup><http://espeak.sourceforge.net/>

<sup>2</sup>Samples and source code can be accessed from the project repository: <https://github.com/alpoktem/MachineDub>

Table 2: Mean opinion scores (MOS) for translation and lip-syncing quality between synthesis of translated subtitle segments (subtitle) and our methodology (synced)

System	MOS		Preference
	Translation	Lip-sync	
subtitle	4.08	3.44	%68
synced	2.96	3.58	%32

ing participants were asked to compare two machine dubbed versions of each sample in terms of (1) translation quality, and (2) lip-syncing precision. Also, they were asked which version they preferred for each sample. Mean opinion scores (MOS) results are shown in Table 2.

## 5.3. Discussion

On manual inspection, our methodology works generally well in determining meaningful prosodic phrasing for the translations only by using the attention output. A similar speech rate ratio was obtained in average for automatically aligned phrases as the original dubbed segments. However, manual inspection on the synthesized samples show that bending ratios that are higher or lower than a certain limit lead to unnatural sounding syntheses. This calls for further modalities that could be integrated in translation to achieve similar length segments both in terms of syllable count and articulation duration. Also, the system as it is only manages to align syntheses to spoken intervals. Further lip-syncing guidelines such as alignment of open, closed phones should also be given attention to achieve further realistic synchronicity.

Another bottleneck that we observed both through manual analyses and perception test results is the effect of poor translation quality. Test participants clearly opt for professional translations even in average their lip-syncing performs more poorly.

## 6. Conclusions

In this paper, we have introduced a methodology for the synchronization of prosodic phrases in a machine dubbing scenario. Our main contribution is a simple but effective exploitation of the by-product of neural machine translation to achieve a mapping between prosodic phrases in the source sentence with tokens in the target sentence. We have further demonstrated its use in a complete machine dubbing pipeline that fuses this transferred phrasing information with the durational structure of the source sentence in order to obtain a synchronized audio translation. In average, the length-wise comparison of automatically aligned phrases outputted by our methodology was able to get close to the speech rate ratio of professionally dubbed segments. Also, qualitative evaluations suggest that it is possible to obtain acceptable or better lip-syncing quality with automatically translated and synchronized dialogue lines. Future work will involve a scoring mechanism for phoneme level alignment and also exploiting of the N-best translation results that could help obtain better speech rate ratios using alternative translations.

## 7. Acknowledgements

The second author is funded by the Spanish Ministry of Economy, Industry and Competitiveness through the *Ramón y Cajal* program.

## 8. References

- [1] A. R. Micola, A. A. Fenoll, A. Banal-Estaol, and A. Bris, “TV or not TV? The impact of subtitling on English skills,” *Journal of Economic Behavior and Organization*, vol. 158, pp. 487–499, 2019.
- [2] G.-M. Luyken, *Overcoming language barriers in television : dubbing and subtitling for the European audience*. Düsseldorf: The European Institute for the Media, 1991.
- [3] F. Chaume, *Audiovisual Translation: Dubbing*. St. Jerome Pub., 2012.
- [4] Z. Hanzlíček, J. Matoušek, and D. Tihelka, “Towards automatic audio track generation for Czech TV broadcasting: Initial experiments with subtitles-to-speech synthesis,” in *2008 9th International Conference on Signal Processing*, Oct 2008, pp. 2721–2724.
- [5] J. Matoušek, Z. Hanzlíček, D. Tihelka, and M. Mner, “Automatic dubbing of TV programmes for the hearing impaired,” in *IEEE 10th INTERNATIONAL CONFERENCE ON SIGNAL PROCESSING PROCEEDINGS*, Oct 2010, pp. 589–592.
- [6] S. Derbring, P. Ljunglöf, and M. Olsson, “SubTTS: Lightweight automatic reading of subtitles,” in *Proceedings of the 17th Nordic Conference of Computational Linguistics (NODALIDA 2009)*. Northern European Association for Language Technology (NEALT), 2009, pp. 272–274.
- [7] P. Ljunglöf, S. Derbring, and M. Olsson, “A free and open-source tool that reads movie subtitles aloud,” in *Proceedings of the Third Workshop on Speech and Language Processing for Assistive Technologies*, ser. SLPAT ’12. Stroudsburg, PA, USA: Association for Computational Linguistics, 2012, pp. 1–4.
- [8] A. Öktem, M. Farrús, and A. Bonafonte, “Bilingual prosodic dataset compilation for spoken language translation,” in *Proceedings of Iberspeech*, Barcelona, Spain, 2018, pp. 20–24.
- [9] P. Deputy, H. Nakasone, and O. Tosi, “Analysis of pauses occurring in the speech of children with consistent misarticulations,” *Journal of Communication Disorders*, vol. 15, no. 1, pp. 43–54, 1982.
- [10] F. Goldman-Eisler, “Pauses, clauses, sentences,” *Language and speech*, vol. 15, pp. 103–113, 04 1972.
- [11] M. Oliveira, “The role of pause occurrence and pause duration in the signaling of narrative structure,” in *Advances in Natural Language Processing*, E. Ranchhod and N. J. Mamede, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 43–51.
- [12] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 3104–3112.
- [13] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *CoRR*, vol. abs/1409.0473, 2014.
- [14] H. Ghader and C. Monz, “What does attention in neural machine translation pay attention to?” in *Proceedings of the Eighth International Joint Conference on Natural Language Processing, IJCNLP 2017, Taipei, Taiwan, November 27 - December 1, 2017 - Volume 1: Long Papers*, 2017, pp. 30–39.
- [15] M. Cettolo, C. Girardi, and M. Federico, “Wit<sup>3</sup>: Web inventory of transcribed and translated talks,” in *Proceedings of the 16<sup>th</sup> Conference of the European Association for Machine Translation (EAMT)*, Trento, Italy, May 2012, pp. 261–268.
- [16] G. Klein, Y. Kim, Y. Deng, J. Senellart, and A. M. Rush, “OpenNMT: Open-source toolkit for neural machine translation,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017)*, Vancouver, Canada, 2017.
- [17] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [18] Z. Tu, Z. Lu, Y. Liu, X. Liu, and H. Li, “Modeling coverage for neural machine translation,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany: Association for Computational Linguistics, Aug. 2016, pp. 76–85.
- [19] T. Luong, H. Pham, and C. D. Manning, “Effective approaches to attention-based neural machine translation,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, Sep. 2015, pp. 1412–1421.
- [20] T. Kudo and J. Richardson, “Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Brussels, Belgium: Association for Computational Linguistics, Nov. 2018, pp. 66–71.
- [21] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany: Association for Computational Linguistics, Aug. 2016, pp. 1715–1725.
- [22] T. Dutoit, V. Pagel, N. Pierret, F. Bataille, and O. van der Vrecken, “The MBROLA project: towards a set of high quality speech synthesizers free of use for non commercial purposes,” in *Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP ’96*, vol. 3, Oct 1996, pp. 1393–1396 vol.3.