



Dr.VOT : Measuring Positive and Negative Voice Onset Time in the Wild

Yosi Shrem¹, Matthew Goldrick², Joseph Keshet¹

¹Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel

²Department of Linguistics, Northwestern University, Evanston, IL, USA

jkeshet@cs.biu.ac.il

Abstract

Voice Onset Time (VOT), a key measurement of speech for basic research and applied medical studies, is the time between the onset of a stop burst and the onset of voicing. When the voicing onset precedes burst onset the VOT is negative; if voicing onset follows the burst, it is positive. In this work, we present a deep-learning model for accurate and reliable measurement of VOT in naturalistic speech. The proposed system addresses two critical issues: it can measure positive and negative VOT equally well, and it is trained to be robust to variation across annotations. Our approach is based on the structured prediction framework, where the feature functions are defined to be RNNs. These learn to capture segmental variation in the signal. Results suggest that our method substantially improves over the current state-of-the-art. In contrast to previous work, our Deep and Robust VOT annotator, Dr.VOT, can successfully estimate negative VOTs while maintaining state-of-the-art performance on positive VOTs. This high level of performance generalizes to new corpora without further retraining.

Index Terms: structured prediction, multi-task learning, adversarial training, recurrent neural networks, sequence segmentation.

1. Introduction

In acoustic studies of speech, one of the most commonly-examined properties of sounds is the voicing of stop consonants (e.g., the contrast between English /b d g/ and /p t k/). In initial positions this is typically indexed by voice onset time (VOT), the time between the release burst of the obstruent and the onset of voicing (negative when voicing begins before the release, and positive when voicing follows release) [1]. The VOT task makes major contributions to both clinical [2] and theoretical [3] studies, for example: to characterize how a communication disorder affects speech [2], or how languages differ in phonetic cues to stop contrasts [1, 4].

Being a key feature in distinguishing voiced and voiceless consonants across languages, VOT is increasingly used as a feature for ASR tasks, such as stop consonant classification [5, 6, 7]. A core challenge to the use of VOT in acoustic studies is the need for highly accurate measurements. For example, speakers use small variations in VOT to signal contrasts with a similar word (e.g., signaling that the speaker said pill instead of bill [8]). Because such effects are modulated across different contexts [9], such accurate measurements must also be scalable. These issues have been partially addressed in our previous work using discriminative large-margin algorithms to measure positive VOT, trained on data from human expert coders [10, 11]. This approach has a significant shortcoming in failing to address the measurement of negative VOTs, severely limiting the application of this tool. While such VOTs are relatively rare in English, they are the prototypical realization of voiced stops in

many other languages (e.g., many Romance and Semitic languages). The amplitude of voicing occurring before release of closure (as in negative VOTs) can be quite small and highly variable, reflecting the challenges of maintaining voicing during closure [12]. Variation of such weak signals can lead to substantial variation in the application of annotations. Extending algorithms [13, 14] to cover such instances has been challenging, in part, due to variation in the annotation. While there is some variation in the annotation of positive VOT [10], the annotations of negative VOT is extensively diversified. Hence, applying the same measurement is inadequate. A related challenge is the varying structure of different datasets. As noted above, the distribution of VOTs varies across languages, and, within a single language community, there is substantial interspeaker variation in VOT [15], possibly reflecting systematic differences as a function age, gender, and other sociolinguistic variables [16].

Inspired by recent developments in Multi-Task Learning (MTL) [17, 18] and adversarial training techniques [19, 20, 21], we explored these methods separately and combined in a structured prediction framework. MTL [17, 18] is a technique to enhance overall performance and robustness by learning multiple domain-related tasks. Relying on the assumption that both tasks share a familiar set of underlying features, leveraging information from an auxiliary task can improve the main task accuracy. MTL can focus the model's attention to relevant features that single-task training may overlook. For example, in speech recognition, an auxiliary task can be phonetic context identification or gender classification [22]. Our training data is composed of several corpora. Each was recorded in a different environment, and annotated by different annotators using slightly different criteria. Naturally, characteristics of the speaker such as age, gender, and so on, also vary and are known to influence VOT [23]. One technique for addressing this issue is adversarial training, which increases model robustness by making it invariant to a specified criterion. To that end, a set of models are trained together but pursue competing goals. Specifically, a good acoustic representation should be corpus-invariant. It should encode only the relevant features regardless of the recording environment, such that a dataset classifier will not be able to distinguish between examples from different corpora. Since the acoustic features vary across corpora, current approaches [13, 14] are unable to transfer good performance on one dataset to different environment.

To that end, we propose **Dr.VOT**, a **Deep Robust method for VOT measurement**, incorporating adversarial training and MTL within a structured prediction framework, so as to enhance general performance and directly address the lack of generalization to new environments and domains. In this paper, the classifying of VOT type as positive or negative as an auxiliary task for MTL, while the adversary task is to find the corpus an utterance originated from.

2. Problem Settings

VOT measurement belongs to the family of segmentation problems, as we are provided with a series of time-related samples in an arbitrary length, and required to predict the boundaries for every segment. Let $\bar{\mathbf{x}} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ denote the input speech utterance, represented as a sequence of acoustic features. Every frame \mathbf{x}_t for $t \in [1, T]$ is a D dimensional vector, i.e. $\mathbf{x}_t \in \mathbb{R}^D$. Since the input can be of any length, the number of frames denoted as T is not fixed.

The corresponding segmentation for every input is denoted by $\bar{\mathbf{y}} = (y_1, \dots, y_k)$, where k can vary across different inputs. Each element y_i in the label vector $\bar{\mathbf{y}}$, indicates the start time of a new segment, hence $y_i \in \mathcal{T}$ where $\mathcal{T} = \{1, \dots, T\}$, and the number of segments is denoted by $|\bar{\mathbf{y}}|$.

The VOT is defined as the time difference between the burst onset and the voicing onset. When the voice onset precedes the burst, it is called prevoiced, and the VOT is negative. Otherwise, when the voice comes after the burst, the VOT is positive. In order to measure VOT, each utterance is paired with three indices corresponding to the prevoicing onset, $t_{pv} \in \mathcal{T} \cup \{-1\}$ (in case of positive VOT, no voice precedes the burst), burst onset, $t_b \in \mathcal{T}$, and the vowel onset $t_v \in \mathcal{T}$, s.t. $t_{pv} < t_b < t_v$. In the positive case, the VOT is $t_v - t_b$, and the segmentation is $\bar{\mathbf{y}} = (t_b, t_v)$. Similarly, if negative, the VOT is $t_b - t_{pv}$ and the segmentation is $\bar{\mathbf{y}} = (t_{pv}, t_b)$.

3. MODEL ARCHITECTURE

In this section, we explain our model and the joint optimization process thoroughly. First, we provide basic knowledge about structured prediction. We then describe a basic structured model using an RNN as a feature function, and then combine this with VOT classification and an adversarial branch.

3.1. Structured Prediction

For every pair $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ where $\bar{\mathbf{x}}$ is the input sequence, $\bar{\mathbf{y}}$ is the desired output. Our goal is to find $\bar{\mathbf{y}}'_w$ which is a good approximation to the true output $\bar{\mathbf{y}}$:

$$\bar{\mathbf{y}}'_w(\bar{\mathbf{x}}) = \operatorname{argmax}_{\bar{\mathbf{y}} \in \mathcal{Y}^k} \mathbf{w}^\top \Phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}), \quad (1)$$

where $\mathbf{w} \in \mathbb{R}^D$ is the set of model parameters, which are estimated on a training set. The feature function $\Phi: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^D$, which is described in the following section, maps a real-world object to a real values vector. In order to quantify a prediction's quality, we define a *task loss* denoted by ℓ , as the absolute difference between our prediction and the true segmentation:

$$\ell(\bar{\mathbf{y}}, \bar{\mathbf{y}}') = [|y_1 - y'_1| - \tau]_+ + [|y_2 - y'_2| - \tau]_+, \quad (2)$$

where $[\pi]_+ = \max\{0, \pi\}$, and τ is a user defined tolerance parameter. In order to find the model parameters \mathbf{w} , ideally we would like to minimize the expected task loss. Let ρ be the unknown data distribution, then:

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} \mathbb{E}_{(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \sim \rho} [\ell(\bar{\mathbf{y}}, \bar{\mathbf{y}}'_w(\bar{\mathbf{x}}))]. \quad (3)$$

Since the distribution ρ is unknown, we minimize the empirical loss over a training set of examples that are drawn i.i.d. from ρ . Moreover, the non-convexity of the loss function makes direct optimization a hard task, hence, we minimize a closely related and slightly different function, called *surrogate loss*, denoted $\bar{\ell}_{\text{struct}}(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}})$ that upper bounds the loss:

$$\bar{\ell}_{\text{struct}}(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}}) = \max_{\bar{\mathbf{y}}' \in \mathcal{Y}} [\ell(\bar{\mathbf{y}}, \bar{\mathbf{y}}') - \mathbf{w}^\top \Phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}) + \mathbf{w}^\top \Phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}')]. \quad (4)$$

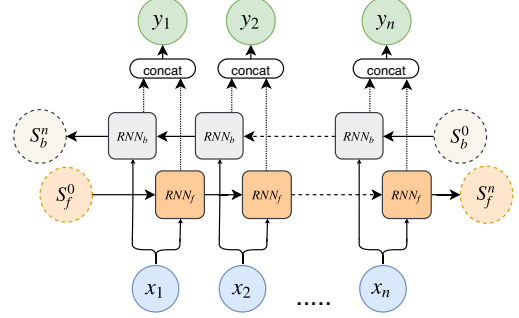


Figure 1: BiRNN scheme where at every index, the outputs of both forward and backward RNNs are concatenated.

3.2. RNNs as Feature Functions

A recurrent neural network (RNN) is a neural architecture that excels in capturing relations between elements within a temporal sequence. Using an inner state which acts as a memory, an RNN leverages insights from previous frames while predicting the next one. Such architectures have excelled in multiple temporal domains such as NLP [24] and speech [25].

Bidirectional RNN (BiRNN) [26] is a model composed of two recurrent networks: the first one reads the input regularly while the later reads it backward. In our model, the outputs of both forward and backward RNNs are concatenated as depicted in Fig. 1. Such a model can predict every frame based on both past and future frames. A common approach to address a segmentation task is to locate the frames in which a new segment begins. To make this search tractable, we decompose $\bar{\mathbf{y}}$ into parts and score every part individually using a feature function. The global score is then the summation of these part scores, $\Phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}) = \sum_{i=1}^k \phi(\bar{\mathbf{x}}, y_i)$. We use BiRNNs to learn the appropriate feature functions, capitalizing on their ability to focus in on the relevant features of the dynamic temporal sequence. We can now re-formulate our prediction rule in Eq. (1) as:

$$\begin{aligned} \bar{\mathbf{y}}'_w(\bar{\mathbf{x}}) &= \operatorname{argmax}_{\bar{\mathbf{y}} \in \mathcal{Y}^k} \mathbf{w}^\top \Phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \\ &= \operatorname{argmax}_{\bar{\mathbf{y}} \in \mathcal{Y}^k} \mathbf{w}^\top \sum_{i=1}^k \phi(\bar{\mathbf{x}}, y_i) \\ &= \operatorname{argmax}_{\bar{\mathbf{y}} \in \mathcal{Y}^k} \mathbf{w}^\top \sum_{i=1}^k \text{BiRNN}(\bar{\mathbf{x}}, y_i) \\ &= \operatorname{argmax}_{\bar{\mathbf{y}} \in \mathcal{Y}^k} \mathbf{w}^\top \sum_{i=1}^k \text{RNN}_f(\bar{\mathbf{x}}, y_i) \oplus^{\text{concat}} \text{RNN}_b(\bar{\mathbf{x}}, y_i) \end{aligned} \quad (5)$$

The notation $\text{RNN}(\bar{\mathbf{x}}, y_i)$, is the RNN's output at index y_i for a given sequence $\bar{\mathbf{x}}$. We use gradient-based methods to minimize the objective Eq. (3), since both the RNN and the structural hinge loss are differentiable with respect to the input and the model parameters.

3.3. VOT Classifier

As noted above, VOT can be either negative or positive. Consider our prediction rule in Eq. (1), the same weight vector \mathbf{w} is used regardless of the VOT type. Having a VOT classification model, which tags whether the utterance has positive or negative VOT may improve our prediction. Specifically, use two different weight vectors: \mathbf{w}_{neg} when dealing with a negative VOT, otherwise, use \mathbf{w}_{pos} . Rather than adapting two separate

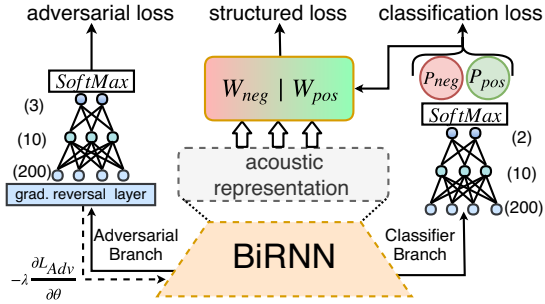


Figure 2: Our structured prediction model combining both classification and adversary branches.

models (VOT predictor and VOT classifier), we integrate the tagging task with the main prediction. We do this by using an MTL technique where VOT tagging as positive or negative is the auxiliary task. Practically, the last hidden state of the RNN is forwarded to a small fully connected network, ending with a softmax layer of size two as depicted in Fig. 2. The output of the softmax is a 2-valued vector which represents the probabilities for positive and negative, $[P(pos|\bar{x}), P(neg|\bar{x})]$. The weight vector in Eq. (5) is then set to be the one associated to the higher probability:

$$\mathbf{w} := \begin{cases} \mathbf{w}_{pos} & P(pos|\bar{x}) > P(neg|\bar{x}) \\ \mathbf{w}_{neg} & \text{else} \end{cases} \quad (6)$$

The classifier parameters are optimized by minimizing the negative log-likelihood loss. As shown in Fig. 2, we allow the gradients to backpropagate from the classifier and update the BiRNN parameters as well; thus, the tasks are jointly trained.

3.4. Adversarial Multi-Task Branch

During training, a model will learn to map the features of seen examples to the desired outcome. However, if the seen examples do not include all the possible dynamics and behavior of the features, performance may be degraded on a different dataset with a different distribution over the features. To make our model robust in the face of this issue – to make our model dataset-invariant – we must develop a method by which the model can apprehend the relevant features while ignoring the rest. Conceptually, the RNN should “filter” the acoustic representation, removing those features that are not relevant to the task.

To implement this, we feed the last state of the RNN to another small fully connected network which acts as a dataset classifier. This is similar to the VOT tagger in the previous section but has a different impact on the RNN. To encourage the main model to extract only those features which are task-related – to avoid extraction of features that alter across datasets – we add a gradient reversal layer [27] before feeding the adversarial net. This layer allows the normal flow of information in the forward step, but in the backward step, it changes the sign of the gradients. Hence, while the adversarial net tries to minimize its loss in dataset classification, the feature extractor is tuned to maximize this loss. At inference, the adversarial branch is irrelevant and can be ignored.

Overall our model’s loss, $\bar{\ell}_{\text{model}}$ is given by:

$$\bar{\ell}_{\text{model}} = \bar{\ell}_{\text{struct}} + \bar{\ell}_{\text{tagger}} + \lambda \bar{\ell}_{\text{adversarial}}, \quad (7)$$

where $\bar{\ell}_{\text{struct}}$ is the structured max-margin loss defined in Eq. 4, $\bar{\ell}_{\text{tagger}}$ is the NLL loss of the tagger, $\bar{\ell}_{\text{adversarial}}$ is the NLL loss

of the adversarial branch, and λ is a hyper-parameter to control the weight of the adversarial loss in the training objective.

4. Experiments

First, we describe the datasets and then proceed by comparing the performance of each model on both seen and unseen datasets to assess generalization ability.

4.1. Datasets

In order to estimate a model’s cross-datasets performance, we used four corpora of isolated productions of words with initial voiced and voiceless stops, characterized by different speakers and recording environments. The first dataset named PGWORDS [3] consists of English picture naming and reading aloud by L1 English speakers and L1 Portuguese/L2 English bilinguals from the Northwestern University community. The other three ([28], [29], [30, 31]) consist of recordings of single words read aloud by L1 English participants. The corpus described in [28] is composed of recordings from the Northwestern University community where [29] and [30, 31] are recordings of speakers from the Purdue University community. Each corpus was divided into train, validation, and test; each speaker was uniquely assigned to one of these sets, allowing us to ensure that performance generalizes across talker differences. The input features are described in detail in [10] and have been utilized in previous research [13, 14]. Broadly speaking, these features measure small time scale (1 and 5 msec windows) spectral, duration, and amplitude measures utilized by human annotators. Overall we have 63 features per frame. Among the entire training corpus consisting of 7,483 utterances, the ratio of positive to negative VOT is 80:20 respectively.

4.2. Evaluation and discussion

We started by evaluating our proposed model *Dr.VOT* compared to previous known models, namely *AutoVOT* [10], *DeepVOT* [32], and *DeepSegmentor* [14]). In all the experiments we used a 2 layers BiLSTM [33] with a hidden size of 100 for each network, as the feature function, and optimized parameters with *Adagrad* [34] and a learning rate of 0.01; λ was set to 0.1. Table 1 summarizes the performance achieved by each model as reported in previous papers [32, 14], when trained and tested on the PGWORDS dataset. The performances were measured by the proportion of differences between automatic and manual measures falling at or below a given tolerance value. As noted in the introduction, the resolution needed for this task is quite fine-grained; we therefore focus primarily on the 2, 5, and 10 msec tolerances in order to reproduce phonological and clinical research. For example, for *DeepSegmentor*, the difference between automatic and manual measurements in the test set was 2 msec or less in 78.2% of examples. As one can see, our model, which is denoted as *Dr.VOT*, performed almost identically to the state-of-the-art model *DeepSegmentor*.

Table 1: Proportion of differences between automatic and manual measures falling at or below a given tolerance value (results for the PGWORDS dataset).

Model	$\tau \leq 2$ msec	$\tau \leq 5$ msec	$\tau \leq 10$ msec	$\tau \leq 15$ msec
<i>AutoVOT</i>	49.1	81.3	93.9	96.0
<i>DeepVOT</i>	53.8	91.6	97.6	98.7
<i>DeepSeg.</i>	78.2	94.1	97.0	98.6
<i>Dr.VOT</i>	78.5	94.6	98.0	98.8

Table 2: Proportion of differences between automatic and manual measures falling at or below a given tolerance value. Performance at each tolerance level is separated by within-corpus (first column) vs. unseen corpora (second column).

Model	$\tau \leq 2$ msec		$\tau \leq 5$ msec		$\tau \leq 10$ msec		$\tau \leq 15$ msec	
<i>DeepSeg.</i>	64.8	50.8	86.2	78.8	93.6	89.5	96.3	93.2
<i>TGM</i>	70.9	53.8	90.7	81.3	96.5	91.5	98.0	94.6
<i>ADM</i>	69.1	55.4	89.4	81.1	95.8	90.9	97.8	94.2
<i>Dr.VOT</i>	71.1	56.5	90.3	82.5	96.2	91.5	98.0	94.7

Using the previous models “in the wild,” i.e., on an entirely new dataset with different talkers from which the model was trained on, and a new acoustic environment, they face a severe performance degradation. For example, when testing *AutoVOT* and *DeepSegmentor* on the unseen corpus [28], the results in the $\tau \leq 2$ msec section; the percentage of files in which the predicted VOT was off by up to 2 msec, dived from 49.1 to 31.9 and from 78.2 to 46.8, respectively.

To better understand the effect of each component of our proposed model, we compared the following four models:

- i *DeepSegmentor* [14];
- ii Tagging model, denoted with *TGM*, an MTL extension of *DeepSegmentor* with the VOT classifier;
- iii Adversarial model, denoted with *ADM*, is an adversarial side-branch added to *DeepSegmentor*;
- iv Both tagging and adversarial models, denoted as *Dr.VOT*.

To check the performance on unseen datasets, we used a 4-fold cross validation method on the datasets. In each fold one of the 4 datasets was left out for evaluation while the rest were used for training. Each model was evaluated on unseen examples from the same dataset it was trained on as well as on examples from the evaluation dataset. We balanced the data proportionately giving a higher probability to examples of negative VOT. Those probabilities were calculated at the beginning of each run to ensure the training set was equally distributed. We report in Table 2 the average performance achieved by each model using the task loss in Eq. (2).

The results suggest that the proposed model, *Dr.VOT*, handled untrained items from both seen and unseen corpora with greater accuracy than *DeepSegmentor* at all tolerance values. At the lowest tolerance values on the trained dataset, both of our additions improve performance relative to the benchmark (64.8%). Adding the tagging task via MTL had a slightly better impact (70.9%) than adversarial training (with dataset classification as the adversarial task: 69.1%). Combining both methods barely disrupts the model with respect to the tagger alone. When looking at the reported results on an unseen corpus (the second of two columns at each tolerance value), the benchmark consistently underperforms the other models. Adding an adversarial branch contributes more to the performance at the lowest range of tolerances (from 50.8% to 55.4%) compared to the classification branch (from 50.8% to 53.8%). Across levels of tolerance, the *Dr.VOT* exhibits the highest level of generalization.

When splitting a dataset into training, validation, and test sets, each set should represent the distribution of the entire dataset. Consequently, results on a test set should be similar to performance in the “real-world.” As mentioned before, a primary goal of this paper is to address a key challenge when measuring VOT: producing accurate measurements when the VOT

Table 3: For negative VOT examples from unseen corpora only, proportion of differences between automatic and manual measures falling at or below a given tolerance value.

Model	$\tau \leq 2$ msec	$\tau \leq 5$ msec	$\tau \leq 10$ msec	$\tau \leq 15$ msec
<i>AutoVOT</i>	3.7	7.9	14.2	20.2
<i>DeepSeg.</i>	23.2	46.1	63.4	71.4
<i>Dr.VOT</i>	32.4	55.3	75.7	84.5

is negative. As we break down the reported results in Table 2 to the performance on negative VOT utterances, we can clearly see the contribution of our proposed algorithm. Since the positive case is far more common, the results on the positive VOT examples are similar to those reported in Table 2. Hence, our discussion focuses on predictions for negative VOT.

The results for examples from unseen corpora are summarized in Table 3. It is clear that it was more difficult to predict negative VOT than positive VOT. This may be due to difficulties in a consistent annotation of the onset of voicing, given that this often has an extremely small amplitude in negative VOT. However, even considering the overall lower performance, it is clear that our model exhibits dramatic improvement over existing approaches. In 32.4% of the utterances which have negative VOT, our model prediction was off by at most 2 msec, a 9.2% absolute improvement and a roughly +40% relative improvement to the 23.2% achieved by the current state-of-the-art. When allowing larger threshold values, our method consistently shows a relative improvement of around +20%. (Note the *AutoVOT* model is limited to the positive case only, leading to very poor performance on the negative-VOT utterances.) This dramatic performance improvement on negative VOTs accounts for most of the overall performance boost of the current approach. For example, considering all VOTs, at a 10 msec threshold *Dr.VOT* outperforms *DeepSegmentor* by 2% (91.5% over 89.5%); focusing only on the less common negative VOTs, there is a 12.3% gain in performance (75.7% over 63.4%).

Finally, in addition to producing better predictions of segmentations, our model also successfully classifies whether the VOT is negative or positive with an accuracy of 94-97% (depending on the dataset) on both seen and unseen corpora.

5. Conclusions

We have presented two different approaches for generalization in a deep learning structured prediction framework. Focusing on enhancing overall performance, we combined MTL within the prediction rule. Focusing on generalization, we used adversarial learning to “filter” unnecessary properties from the feature space. As far as we are aware of, this is the first work to tackle, in an integrated way, VOT classification along with prediction of both positive and negative VOTs – critically, yielding well-balanced results on both. In future work we would like to further explore the adversarial effect by testing it over multilingual corpora, and add handcrafted auxiliary related tasks. Such an approach may allow the first-ever application of these techniques to the analysis of bilingual language production.

After acceptance *Dr.VOT* will be publicly available at <https://github.com/MLSpeech/Dr.VOT>.

6. Acknowledgements

Research supported in part by NIH grant 1R21HD077140. Thanks to Olga Dmitrieva and Alexander Francis for providing access to VOT datasets.

7. References

- [1] L. Lisker and A. S. Abramson, "A cross-language study of voicing in initial stops: Acoustical measurements," *Word*, vol. 20, no. 3, pp. 384–422, 1964.
- [2] P. Auzou, C. Ozsancak, R. J. Morris, M. Jan, F. Eustache, and D. Hannequin, "Voice onset time in aphasia, apraxia of speech and dysarthria: a review," *Clinical Linguistics & Phonetics*, vol. 14, no. 2, pp. 131–150, 2000.
- [3] N. Paterson, "Interactions in bilingual speech processing," Ph.D. dissertation, Northwestern University, 2011.
- [4] T. Cho and P. Ladefoged, "Variation and universals in vot: evidence from 18 languages," *Journal of phonetics*, vol. 27, no. 2, pp. 207–229, 1999.
- [5] P. Niyogi and P. Ramesh, "The voicing feature for stop consonants: Recognition experiments with continuously spoken alphabets," *Speech Communication*, vol. 41, no. 2-3, pp. 349–367, 2003.
- [6] V. Stouten and H. van Hamme, "Automatic voice onset time estimation from reassignment spectra," *Speech Commun.*, vol. 51, pp. 1194–1205, 2009.
- [7] J. H. Hansen, S. S. Gray, and W. Kim, "Automatic voice onset time detection for unvoiced stops (*/p/,/t/,/k/*) with application to accent classification," *Speech Communication*, vol. 52, no. 10, pp. 777–789, 2010.
- [8] M. Baese-Berk and M. Goldrick, "Mechanisms of interaction in speech production," *Language and cognitive processes*, vol. 24(4), pp. 527–554, 2009.
- [9] E. Buz, M. K. Tanenhaus, and T. F. Jaeger, "Dynamically adapted context-specific hyper-articulation: Feedback from interlocutors affects speakers subsequent pronunciations," *Journal of memory and language*, vol. 89, pp. 68–86, 2016.
- [10] M. Sonderegger and J. Keshet, "Automatic measurement of voice onset time using discriminative structured prediction," *The Journal of the Acoustical Society of America*, vol. 132, no. 6, pp. 3965–3979, 2012.
- [11] M. Goldrick, J. Keshet, E. Gustafson, J. Heller, and J. Needle, "Automatic analysis of slips of the tongue: Insights into the cognitive architecture of speech production," *Cognition*, vol. 149, pp. 31–39, 2016.
- [12] A. S. Abramson and D. Whalen, "Voice onset time (vot) at 50: Theoretical and practical issues in measuring voicing distinctions," *Journal of Phonetics*, vol. 63, pp. 75–86, 2017.
- [13] Y. Adi, J. Keshet, and M. Goldrick, "Vowel duration measurement using deep neural networks," in *MLSP*, 2015.
- [14] Y. Adi, J. Keshet, E. Cibelli, and M. Goldrick, "Sequence segmentation using joint rnn and structured prediction models," in *ICASSP*, 2017.
- [15] M. Sonderegger, M. Bane, and P. Graff, "The medium-term dynamics of accents on reality television," *Language*, vol. 93, no. 3, pp. 598–640, 2017.
- [16] P. Torre III and J. A. Barlow, "Age-related changes in acoustic characteristics of adult speech," *Journal of communication disorders*, vol. 42, no. 5, pp. 324–333, 2009.
- [17] R. Caruana, "Multitask learning," *Machine learning*, vol. 28, no. 1, pp. 41–75, 1997.
- [18] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa, "Natural language processing (almost) from scratch," *Journal of Machine Learning Research*, vol. 12, no. Aug, pp. 2493–2537, 2011.
- [19] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [20] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, M. Laviolette, F. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [21] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *arXiv preprint arXiv:1706.06083*, 2017.
- [22] G. Pironkov, S. Dupont, and T. Dutoit, "Multi-task learning for speech recognition: an overview," in *Proceedings of the 24th European Symposium on Artificial Neural Networks (ESANN)*, vol. 192, 2016.
- [23] R. Singh, J. Keshet, D. Gencaga, and B. Raj, "The relationship of voice onset time and voice offset time to physical age," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5390–5394.
- [24] T. Mikolov, M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur, "Recurrent neural network based language model," in *Interspeech*, 2010.
- [25] A. Graves, A. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Acoustics, speech and signal processing (icassp), 2013 IEEE international conference on*. IEEE, 2013, pp. 6645–6649.
- [26] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997.
- [27] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," *arXiv preprint arXiv:1409.7495*, 2014.
- [28] M. Goldrick, C. Vaughn, and A. Murphy, "The effects of lexical neighbors on stop consonant articulation," *The Journal of the Acoustical Society of America*, vol. 134, no. 2, pp. 172–177, 2013.
- [29] O. Dmitrieva and A. N. Tews, "First language phonetic drift in second language instructional environment," *The Journal of the Acoustical Society of America*, vol. 143, no. 3, pp. 1952–1952, 2018.
- [30] A. A. Shultz, "Individual differences in cue weighting of stop consonant voicing in perception and production," Master's thesis, Purdue University, 2011.
- [31] A. A. Shultz, A. L. Francis, and F. Llanos, "Differential cue weighting in perception and production of consonant voicing," *The Journal of the Acoustical Society of America*, vol. 132, no. 2, pp. 95–101, 2012.
- [32] Y. Adi, J. Keshet, O. Dmitrieva, and M. Goldrick, "Automatic measurement of voice onset time and prevoicing using recurrent neural networks," *Interspeech 2016*, pp. 3152–3155, 2016.
- [33] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [34] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Research*, vol. 12, no. Jul, pp. 2121–2159, 2011.