

MEDIAN-BASED GENERATION OF SYNTHETIC SPEECH DURATIONS USING A NON-PARAMETRIC APPROACH

Srikanth Ronanki, Oliver Watts, Simon King, and Gustav Eje Henter

The Centre for Speech Technology Research (CSTR), The University of Edinburgh, UK

srikanth.ronanki@ed.ac.uk

ABSTRACT

This paper proposes a new approach to duration modelling for statistical parametric speech synthesis in which a recurrent statistical model is trained to output a phone transition probability at each timestep (acoustic frame). Unlike conventional approaches to duration modelling – which assume that duration distributions have a particular form (e.g., a Gaussian) and use the mean of that distribution for synthesis – our approach can in principle model any distribution supported on the non-negative integers. Generation from this model can be performed in many ways; here we consider output generation based on the median predicted duration. The median is more typical (more probable) than the conventional mean duration, is robust to training-data irregularities, and enables incremental generation. Furthermore, a frame-level approach to duration prediction is consistent with a longer-term goal of modelling durations and acoustic features together. Results indicate that the proposed method is competitive with baseline approaches in approximating the median duration of held-out natural speech.

Index Terms—text-to-speech, speech synthesis, duration modelling, non-parametric models, LSTMs.

1. INTRODUCTION

This paper presents a new approach to the modelling and generation of speech-segment durations, a characteristic of speech which contributes to the perception of its rhythm. Generating appropriate rhythm and melody is a challenging but vital step in producing natural-sounding synthetic speech, particularly in expressive or conversational scenarios. Steady improvements have been made in statistical parametric speech synthesis (SPSS), in particular through the adoption of deep and recurrent machine-learning techniques in recent years [1, 2]. As we expect that high-level, long-range dependencies are of importance for the prosodic structure of speech, recurrent models such as LSTMs are well-suited to prosodic sequence modelling problems [3].

In spite of the progress, however, the prosodic characteristics of synthetic speech remain one of its major shortcomings. One factor which we suppose contributes to this shortcoming is that current systems effectively model duration with a Gaussian distribution and generate predictions from the mean of that distribution. This is problematic, as the distribution of speech-segment durations is well known to be skewed, and so using the mean of a normal distribution fit to such observations will not produce predictions that are most typical of the process (that is, have high probability). Another possible weakness of current approaches towards duration generation is that they typically operate as an initial stage, separate from the generation of acoustic features [3]. We consider it desirable

to have a single model whose parameters are learned to simultaneously generate both segment durations and the frames of acoustic features within those segments. A major motivation for this is that such a joint model would allow the simultaneous adaptation [4] and control [5] of rhythmic, melodic and phonetic characteristics in a stable and consistent way. However, one obvious difficulty in implementing such a model is that predictions of segment durations and of acoustic observations are conventionally made on two separate time-scales. Most recently, wavenet [6] has made waveform-level modelling possible in TTS, but still requires an external duration model. Our approach can in principle be integrated with WaveNet, for sample-level duration modelling.

We here present an approach to duration modelling which in essence consists of predicting at each acoustic frame a probability that the model will subsequently advance to the next phone in the phonetic sequence. We show that – assuming a recurrent model is used – this approach may describe any duration distribution on the positive integers. Although it is feasible to model duration by explicitly choosing other distributions which may be more appropriate than a Gaussian (e.g., log-normal or gamma distribution), these still might not be optimal for a given conditional duration, and so the possibility of not having to commit to a predetermined distribution before model training is a powerful advantage of our approach. Furthermore, as our model moves from modelling duration at the level of the phonetic segment to the level of the acoustic frame, our approach is a necessary ingredient of our on-going work in unifying models of durations and acoustics, so that acoustic parameters a phone transition probabilities are predicted jointly for each frame.

2. BACKGROUND

In this section, we give a brief overview of how speech-sound durations have been generated in synthetic speech, with a particular emphasis on the use of statistical methods for this task. This provides context for the proposed method in Section 3.

2.1. A Brief Review of Modelling and Generating Durations

In early, formant-based synthesis systems, phone durations were generated by rule [7]. Rules were commonly hand-crafted, rather than learned from data, meaning that no statistical modelling was used. The concatenative synthesis approaches that emerged next did not require modelling or generating durations, since the units themselves (typically diphones) possess intrinsic durations. That said, some approaches allowed predicted durations to be incorporated into the target cost [8].

The rise of statistical parametric speech synthesis (SPSS) [9, 10] has introduced a new methodology for duration generation, in which a statistical model (probability distribution) is created to describe

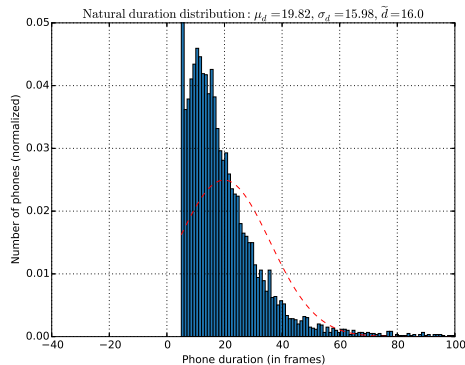


Fig. 1. Distribution of all natural speech-sound durations in our held-out dataset, and a maximum likelihood-fitted Gaussian density. The median duration is 16. The minimum duration is five since the HMM-based forced aligner used five states per phone and a left-to-right topology with no skip transitions.

speech-sound durations. This distribution is supplemented by two things: a machine learning method to predict the properties of the statistical model from the input text (i.e. the linguistic features), and a principle for generating durations from the model distribution, such as random sampling or taking the mean of the distribution.

The first SPSS systems were based on simple hidden Markov models (HMMs), which implies that state durations (commonly five states per phone) are assumed follow a memoryless geometric distribution. Decision trees (DTs) were used to predict the distribution parameter based on training data, with the mean of the predicted distribution used for generation.

In reality, a geometric distribution is quite far removed from how natural speech durations are distributed, as can be seen in Figure 1. Zen et al. [11] introduced the idea of using hidden semi-Markov models (HSMMs) to describe durations in the context of speech synthesis. These models track the amount of time (acoustic frames) spent in the current HMM state, and allow the frame counter to influence the probability of transitioning, making it possible to model a much wider class of duration distributions. Typically, HSMM durations would be assumed to follow a parametric distribution of some sort; the widely-used *HMM-based speech synthesis system* (HTS) [12] defaults to Gaussian duration distributions, for instance, although the skewness and non-negativity of speech durations mean that other choices of distribution (log-normal, gamma) might be more suitable [13, 14]. As with HMMs, decision trees would be used to predict distribution parameters per state, and the mean of the predicted distribution was used for generation.

Recently, SPSS has seen stronger machine-learning techniques such as deep and/or recurrent neural networks replace decision trees for predicting the properties of duration distributions from the text, e.g., [3, 15]. Most often, these systems train DNNs or RNNs to minimise the mean squared prediction error, which is theoretically equivalent to using a Gaussian duration model (with a globally tied standard deviation) in conjunction with mean-based duration generation. The evolution of approaches to duration generation for TTS is summarised in Table 1.

Looking at Table 1, we see that all canonical methods predict only a few numbers per phone; at most 5 means and 5 standard deviations for Gaussian state-level predictions, of which only the means affect output generation. This renders the methods incapable of pre-

TTS type	Distribution	Level	Predictor	Generation
Formant	-	Phone	-	Rule
Concat.	-	Phone	-	Exemplar
HMM	Geometric	State	DT	Mean
HSMM	Parametric	State	DT	Mean
NN	Gaussian	State	DNN/RNN	Mean
Proposed	Non-parametric	Frame	DNN/RNN	Median

Table 1. Methods for generating speech-sound durations (last column) in different types of TTS. Statistical methods incorporate a distribution, a method of predicting its properties, in addition to a method of generating output from the predicted distribution.

dicting general duration distributions with more degrees of freedom. In contrast to these prior methods, we describe a fundamentally non-parametric approach in which a model is trained to predict the phone transition probability at each timestep, i.e., acoustic frame, as outlined on the final line of Table 1. This set-up can in theory describe any probability mass function on the positive integers. (In practice, performance may be limited by biases in the learning algorithm used.) It thus becomes possible to meaningfully represent any and all important properties of the duration distribution, for instance its skewness, which most models hitherto have ignored.

2.2. Median-Based Generation

All methods in Table 1 that predict a distribution of durations automatically support a wide variety of methods for generating output durations from this distribution. While in principle, natural speech is a random sample drawn from the true duration distribution, the output naturalness of sampling methods in speech synthesis has been found to perform poorly unless highly accurate models are used [16]. As a consequence, most SPSS systems use deterministic generation methods, which in practice is synonymous with generating the mean duration.

We here consider a scheme where durations generated at synthesis time are based on the median – rather than the conventional mean – of the predicted duration distribution. Since we are no longer assuming that the duration distribution is symmetric, the median will typically differ from the distribution mean. To our knowledge, no other paper in the statistical speech synthesis literature has considered predicting a median duration different from the mean.

Conveniently, the median can be identified from the left tail of the distribution, whereas computing the mean requires evaluating the probability mass function over its entire, infinite support. This property enables median-duration based sequential output generation with no look-ahead or other overhead, which is attractive for incremental synthesis approaches. Unlike the mean, the median is nearly always an integer.

The advantages of generating median durations are more than computational: For skewed distributions, including the distribution of natural speech-sound durations as graphed in Figure 1, the median is frequently closer to the peak of the distribution than the mean is; cf. [17, 18]. This implies that, in the spirit of most likely output parameter generation [19], median-based duration prediction is likely closer to the peak density – the “most typical” outcome – than the mean is. (The mean duration is often atypically long.) Perhaps most importantly, the median is a *statistically robust* quantity, and is not affected by the tails of the real duration distribution. Statistical robustness is compelling for speech synthesis [20], particularly for big and found datasets, as it reduces the sensitivity to errors and unexpected behaviour in the training corpus. Among other things, this

could be of value with the highly expressive and variable training data used for the experiments in Section 4.

2.3. Frame-Level Duration Prediction

While frame-level duration predictions are uncommon, this paper is not the first to suggest it. Watts et al. [21] described a joint model of duration and speech parameters, where a deep neural network was trained to simultaneously output acoustic parameters and a 5-dimensional (phone) state-duration vector for each frame. Synthesis with this type of network results in a chicken-and-egg problem, where state durations are both an input and an output of the sequence prediction network. This was solved by iteratively refining state-duration predictions over multiple passes, which is slow. In our proposed method, a single pass suffices for duration generation. Furthermore, despite the frame-level granularity of the approach in [21], it is not straightforward to identify a probabilistic interpretation of their scheme.

A general advantage of duration predictions made at the frame level is that they can be unified with the (traditionally distinct) prediction of acoustic features, such as pitch and vocal tract filter MGCs, that takes place for each frame. Such joint modelling of durations with acoustic properties of speech was a major factor in motivating the set-up in [21]. It is suspected that generating durations and fundamental frequency contours that are jointly appropriate may be of importance for synthetic speech prosody; cf. [22, 23].

3. THEORY

3.1. Preliminaries

Let $p \in \{1, \dots, P\}$ be a phone index, and let $t \in \{1, \dots, T\}$ be an index into frames. Let further D_p – a random variable – be the duration of phone p , and let $d_p \in \mathbb{Z} > 0$ be an outcome of D_p .¹

Natural speech phones have different duration distributions that depend on the input text. In TTS, the properties of the distribution D_p are predicted from contextual linguistic features \mathbf{l}_p extracted from the text by the synthesiser front-end. Specifically, *duration modelling* is the task of mapping the sequence of linguistic features $(\mathbf{l}_1, \dots, \mathbf{l}_P)$ to a sequence of predicted duration distributions (D_1, \dots, D_P) . *Duration generation*, meanwhile, is the task of mapping $(\mathbf{l}_1, \dots, \mathbf{l}_P)$ to a sequence of generated durations $(\hat{d}_1, \dots, \hat{d}_P)$. In SPSS, one or the other of these mappings is learned from a corpus of parallel text and speech data using a machine learning method; in this paper, deep and recurrent neural networks will be used. We will write \mathcal{D} to denote a dataset of parallel input features \mathbf{l} and random variable outcomes d used to train this predictor.

3.2. Conventional Duration Modelling

In statistical synthesisers that generate durations on the state or phone level, the conventional approach is to assume that the durations D_p follow some parametric family f_D with a parameter $\theta \in \mathbb{R}^N$ (N being the number of degrees of freedom), i.e.,

$$\mathbb{P}(D_p = d) = f_D(d; \theta_p). \quad (1)$$

Predicting the distribution D_p then reduces to the stochastic regression problem of predicting the distribution parameter θ_p from the

¹In many TTS systems, D_p is a vector of per-phone state durations, but we restrict ourselves to phone-level predictions in this paper; a state-based formulation is a straightforward extension.

phone-level parallel input-output dataset

$$\mathcal{D}_p = ((\mathbf{l}_1, \dots, \mathbf{l}_P), (d_1, \dots, d_P)). \quad (2)$$

We will write \mathbf{L}_p to denote the linguistic information available to the predictor at p , which is \mathbf{l}_p for feedforward approaches and $(\mathbf{l}_1, \dots, \mathbf{l}_p)$ for unidirectional RNNs.

In practice, most contemporary DNN-based synthesisers do not perform full distribution modelling, but map directly from \mathbf{L}_p to the property of the distribution D_p that they wish to generate, such as its mean $\mathbb{E}(D_p)$. The dominant principle – and the only one to be considered for the baselines in this paper – is to tune the weights \mathbf{W} of a DNN or RNN $d(\mathbf{L}; \mathbf{W})$ to minimise the mean squared error (MSE) on the training dataset,

$$\widehat{\mathbf{W}}(D_p) = \operatorname{argmin}_{\mathbf{W}} \sum_{(\mathbf{L}_p, d_p) \in \mathcal{D}} (d_p - d(\mathbf{L}_p; \mathbf{W}))^2; \quad (3)$$

synthesis-time durations \hat{d}_p are then generated by

$$\hat{d}(\mathbf{L}_p) = d(\mathbf{L}_p; \widehat{\mathbf{W}}(D_p)). \quad (4)$$

The theoretically optimal predictor \hat{d}^* that minimises the MSE is the conditional mean,

$$\hat{d}_p^*(\mathbf{L}_p) = \operatorname{argmin}_{\hat{d}} \mathbb{E} \left((D_p - \hat{d})^2 \mid \mathbf{L}_p \right) \quad (5)$$

$$= \mathbb{E}(D_p \mid \mathbf{L}_p), \quad (6)$$

so the end result is very similar to fitting a Gaussian duration model f_D and using the mean of that fitted Gaussian to generate durations.

3.3. Non-Parametric Duration Modelling

We will now describe a scheme that, unlike conventional, phone-level approaches with parametric families $f_D(d; \theta)$, is able to model and predict arbitrary duration distributions for D (restricted only by the biases of the machine learning method used). The key idea is to make predictions at the frame level about when phone transitions occur.

Assume phone durations are known up until the current frame t , and let $p(t')$ for $1 \leq t' \leq t$ be a function that maps from a given frame $t' \leq t$ to the phone it has been assigned to. We can then define a frame-level sequence \mathbf{L}_t of linguistic features

$$\mathbf{L}_t = (\mathbf{l}_1, \dots, \mathbf{l}_t) \quad (7)$$

$$= (\mathbf{l}_{p(1)}, \dots, \mathbf{l}_{p(t)}) \quad (8)$$

up until t , which is constant for all frames within each phone. For brevity, we shall write p for the current phone $p(t)$. Finally, we let t_0 be the final frame of the previous phone, so that we can define $n_t = t - t_0 \geq 1$, the duration of the current phone so far.

We now define the *transition probability* π_t for the phone p at time t , given the linguistic features up until this point – that is,

$$\pi_t = \mathbb{P}(D_p = n_t \mid D_p \geq n_t, \mathbf{L}_t). \quad (9)$$

As long as the transition probabilities satisfy $\pi_t \in [0, 1]$ and

$$\prod_{t'=\ell_0+1}^{\infty} (1 - \pi_{t'}) = 0 \quad (10)$$

they induce a unique duration distribution

$$\mathbb{P}(D_p = n_t | \mathbf{L}_t) = \pi_t \prod_{t'=t_0+1}^{t_0+n_t-1} (1 - \pi_{t'}) \quad (11)$$

on the positive integers.

We propose to build a predictor, based on training data, that estimates π_t from the linguistic input features. Specifically, we will train this predictor on the frame-level dataset

$$\mathcal{D}_t = (\mathbf{L}_T, (x_1, \dots, x_T)), \quad (12)$$

where x_t is an indicator variable that equals one if and only if t is the final frame of the current phone, i.e.,

$$x_t = \begin{cases} 1 & \text{if } t = t_0 + d_p \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

In this paper, we consider a deep and unidirectional recurrent neural network $x(\mathbf{L}; \mathbf{W})$, with weights $\widehat{\mathbf{W}}$ trained to minimise the MSE in recursively predicting the indicator variable x_t – that is,

$$\widehat{\mathbf{W}}(\mathcal{D}_t) = \underset{\mathbf{W}}{\operatorname{argmin}} \sum_t (x_t - x(\mathbf{L}_t; \mathbf{W}))^2. \quad (14)$$

The hypothetical predictor \widehat{x}^* that minimises this MSE is the conditional mean,

$$\widehat{x}_t^*(\mathbf{L}_t) = \underset{\widehat{x}}{\operatorname{argmin}} \mathbb{E} \left((X_t - \widehat{x})^2 \mid \mathbf{L}_t \right) \quad (15)$$

$$= \mathbb{E}(X_t | \mathbf{L}_t) \quad (16)$$

$$= \sum_{x \in \{0, 1\}} x \cdot \mathbb{P}(X_t = x | \mathbf{L}_t) \quad (17)$$

$$= \mathbb{P}(X_t = 1 | \mathbf{L}_t), \quad (18)$$

which is mathematically equivalent to the transition probability π_t . This means that, as long as the predictor of X_t is theoretically capable of generating arbitrary outputs for every frame, our approach can describe virtually any transition distribution – and thus any duration distribution – on the positive integers. LSTMs and other RNNs satisfy this requirement, due to their internal state (memory), here denoted \mathbf{c}_t , which evolves from frame to frame.

Unlike the maximum likelihood objective function, the MSE used here is not disproportionately sensitive to large errors in probability estimates of rare events (outliers). This makes the estimated $\widehat{\mathbf{W}}$ more statistically robust to errors in data and model assumptions.

3.4. From Transitions to Median Durations

Having defined a phone duration distribution in (11) and trained a model to predict this distribution from data, we must consider how predicted durations \widehat{d} are to be generated from this distribution. As discussed in Section 2.2, sampling typically yields poor naturalness, while mean-based generation is sensitive to the tails of the distribution and unsuitable for sequential generation. On the other hand, it is straightforward to derive the right tail probability of the phone duration distribution induced by the values of π_t seen thus far, as

$$\mathbb{P}(D_p > n_t | \mathbf{L}_t) = \prod_{t'=t_0+1}^{t_0+n_t} (1 - \pi_{t'}). \quad (19)$$

This relation straightforwardly enables synthesis based on the predicted *median* duration: by stepping from $n_t = 1$ and upwards, the (estimated) median duration \widehat{d}_p of phone p is reached when $\mathbb{P}(D > d)$ first dips down to 0.5 or below,

$$\widehat{d}_p = \min_{n_t \in \mathbb{Z}} n_t \quad (20)$$

such that $\mathbb{P}(D_p > n_t) \leq 1/2$. (21)

Expression (19) thus allows us to generate frames sequentially, advancing $p(t+1)$ to the next phone $p+1$ when the predicted median duration is reached, with no additional overhead at generation time.

Just as the mean squared error is minimised by the mean, the median is the theoretical minimiser of another error measure, namely the *mean absolute error* (MAE),

$$\operatorname{MAE}(\widehat{d}) = \sum_{p \in \mathcal{D}_p} |d_p - \widehat{d}_p|. \quad (22)$$

If a method improves the MAE, we would expect its prediction to be closer to the (conditional) median of the data. Interestingly, summing absolute rather than squared errors is a common component of the mel-cepstral distortion (MCD) [24] often used to evaluate acoustic models in speech synthesis, but the same idea is less commonly seen for evaluating generated durations.

3.5. Adding External Memory

In the proposal described so far, progress through the current phone is tracked solely via the internal state of the predictor used (\mathbf{c}_t for an LSTM). This is in contrast to the approach used by the hidden semi-Markov models in HTS, which achieved more general duration distributions than regular HMMs by maintaining an external counter of the number of frames spent in each state, and using it to compute the transition probability. However, nothing prevents us from adding our variable n_t , which counts the frame number within the current phone, as an input to the neural network $x(\cdot; \mathbf{W})$ that predicts frame-level transition probabilities, in addition to the regular, linguistic features \mathbf{L}_t . We call these augmented features \mathbf{l}'_t and

$$\mathbf{l}'_t = ([l'_1, n_1]^T, \dots, [l'_t, n_t]^T), \quad (23)$$

so that the augmented predictor is $x(\mathbf{L}'; \mathbf{W})$.

One may surmise that having n_t as an input feature could increase performance, as it provides highly relevant information to the model at all times, instead of hoping the predictor will discover the importance of duration tracking on its own during training. To test this hypothesis, our experiments in Section 4 compare two systems that differ only in whether or not they incorporate an external counter n_t as an input to the central neural network.

As a side note, appending n_t to \mathbf{l}_t makes the inputs to the frame-level neural network $x(\cdot; \mathbf{W})$ different with each frame. This makes it possible for the predicted transition distributions to differ from frame to frame as well, even without using a stateful predictor such as an RNN. This should, for example, enable the prediction of arbitrary transition distributions also when $x(\cdot; \mathbf{W})$ is a non-recurrent, feedforward neural network. We have, however, not explored this possibility in the current paper.

4. EXPERIMENTAL SET-UP

4.1. Data

As a first inquiry into the viability of our frame-level duration-prediction approach, we conducted some experiments on a speech

database created for the 2016 Blizzard Challenge[25].² The database – provided to the Challenge by Usborne Publishing Ltd. – consists of the speech and text of 50 children’s audiobooks spoken by a British female speaker. We made use of a segmentation of the audiobooks carried out by another Challenge participant³ and kindly made available to other participants. The total duration of the audio was approximately 4.33 hours after segmentation. For the purposes of the work presented here, 4% of the data was set aside as a test set. The test set consists of three whole short stories: *Goldilocks and the Three Bears*, *The Boy Who Cried Wolf* and *The Enormous Turnip*, having a total combined duration of approximately 10 minutes.

4.2. Feature extraction

We obtained a state-level forced alignment of the sentence-segmented data described above using context independent HMMs. Festvox’s `ehmm` [26] was used to insert pauses into the annotated phone sequences based on the acoustics, to improve the accuracy of forced-aligned durations. Each phone was then characterised by a vector of 481 text-derived binary and numerical features: these features are a subset of the features used in decision-tree clustering questions from the HTS public demo [12]; numerical features queried by those questions were used directly where possible. All inputs were normalised to the range [0.01, 0.99]. The phone durations used to train conventional baselines was obtained by counting the frames in a phone. Contrary to our previous work [20], sub-phone states were not used in either DNN training or prediction. This is consistent with our longer-term goal of freeing ourselves from depending on often arbitrary HMM sub-phone alignments.

4.3. System training

Four systems were trained: two phone-level predictors (which we identify as *Phone-DNN* and *Phone-LSTM*) to provide benchmarks and assess the impact of recurrent modelling on duration prediction, and two experimental systems implementing the new idea (which we identify as *Frame-LSTM-I* and *Frame-LSTM-E*).

For training, all networks were initialised using small random weights, with no pre-training. Each duration prediction system was trained with a fixed learning rate, manually tuned to yield close-to-optimal results on the development set in 25 epochs or less. Early stopping was used to avoid overfitting, by aborting training once the objective function on the development set had failed to improve for five epochs.

4.3.1. *Phone-DNN* and *Phone-LSTM*

The two baselines use phone-level linguistic features as input and are optimised to predict the (mean and variance normalised) duration of the phones in the training data. *Phone-DNN* is a feedforward DNN with six layers of 1024 nodes each. The hidden nodes used tanh activation and the output was linear. *Phone-LSTM* was configured with five feed-forward layers of 1024 nodes each and a final unidirectional SLSTM [27] hidden layer consisting of 512 nodes.

4.3.2. *Frame-LSTM-I* and *Frame-LSTM-E*

Both proposed systems used the same architecture as that of *Phone-LSTM* but – unlike the baseline systems – were trained with 1 data-point per frame instead of per phone. The targets presented in train-

ing were 0.0 for non-phone-final frames, and 1.0 for phone-final frames. *Frame-LSTM-I* took only the vector representing the current phone as input; at run-time it was therefore required to rely on its *internal* memory (hidden state and LSTM cell state) to determine the phone transition probability at any given frame. The inputs to *Frame-LSTM-I* are identical for all frames in any given phone. *Frame-LSTM-E* implemented the idea described in Section 3.5: inputs encoding the current phone context are augmented with a frame counter indicating the number of frames that have passed since the start of the phone. *Frame-LSTM-E* can therefore also rely on this *external* memory of how much time has elapsed within the current phone when making phone transition predictions, in addition to internal memory.

4.4. Synthesis

At synthesis time, `ehmm` phone sequences derived from the test data were used as input to each duration prediction model. This corresponds to using an oracle pausing strategy, but providing no other acoustically-derived information to the predictors. Generation from *Phone-DNN* and *Phone-LSTM* is straightforward: an (effectively mean) output is generated for each phone.

Median-based duration predictions were obtained from *Frame-LSTM-I* and *Frame-LSTM-E* by running the models recurrently over frame-level inputs and using the technique explained in Section 3.4 and summarised in pseudocode in Algorithm 1, to decide when the median had been reached and thus when to move to the next phone. By keeping track of the time spent in each phone, we obtain median-based predictions of phone duration either with external memory of time elapsed in current phone (*Frame-LSTM-E*) or without it (*Frame-LSTM-I*).

Algorithm 1 Switching criterion for *Frame-LSTM-I*

```

1: procedure SWITCHPROB(linguisticInput)
2:   nof  $\leftarrow$  length of linguisticInput
3:   for each phoneNum in nof do
4:     remMass  $\leftarrow$  1
5:     for each frame t in order do
6:       ph  $\leftarrow$  linguisticInput[phoneNum]
7:       probSwitchAtT  $\leftarrow$  FrameLSTM(ph)
8:       remMass  $\leftarrow$  remMass * (1 - probSwitchAtT)
9:       if remMass  $\leq$  0.5 then
10:        predictedDur[phoneNum]  $\leftarrow$  t + 1
11:        break;
12:       end if
13:     end for
14:   end for
15:   return predictedDur
16: end procedure

```

5. RESULTS

Table 2 gives RMSE, mean absolute error (MAE) and Pearson correlation computed by comparing the predicted durations against the held-out reference. Silence segments were excluded from these calculations.

Firstly, results for the benchmark systems show that *Phone-LSTM* outperforms *Phone-DNN* in every respect, confirming the importance of the LSTM’s recurrence for duration modelling. Turning to the experimental systems, it can be seen that while the mean-based

²http://www.synsig.org/index.php/Blizzard_Challenge_2016

³Innoetics: <https://www.innoetics.com>

Model	RMSE	MAE	Corr.
Phone-DNN	8.037	4.759	0.750
Phone-LSTM	7.789	4.556	0.765
Frame-LSTM-I	8.254	4.610	0.761
Frame-LSTM-E	8.294	4.574	0.754

Table 2. RMSE and MAE (both in frames per phone), along with Pearson correlation of the predicted durations, all measured w.r.t. force-aligned durations.

Phonetic class	Phone-LSTM			Frame-LSTM-E		
	RMSE	MAE	Corr.	RMSE	MAE	Corr.
Vowels	8.516	4.848	0.809	9.027	4.891	0.799
Consonants	7.313	4.378	0.709	7.815	4.382	0.694
Plosives	5.206	3.608	0.732	5.610	3.612	0.720
Fricatives	6.489	4.380	0.769	6.859	4.246	0.764
Nasals	6.833	4.459	0.568	7.376	4.398	0.550
Affricates	5.658	4.220	0.797	5.432	3.746	0.821
Glides + liquids	8.013	5.260	0.569	8.235	5.075	0.599

Table 3. Objective measures from Table 2 broken down by phonetic class.

predictions of the LSTM and the DNN baselines clearly outperform the proposed methods in terms of RMSE, the gap is much smaller when it comes to MAE. Median-based generation methods are expected to score better on MAE than RMSE, because – as discussed in Section 3.4 – the median is the theoretical minimiser of MAE, while the mean minimises the (R)MSE.

The breakdown of results by phonetic class given in Table 3 shows an interesting phenomenon: while the proposed system *Frame-LSTM-E*’s MAE is worse than the LSTM benchmark for vowels and slightly worse for consonants overall, for all classes of consonant except plosives the proposed method performs better.

While the best-performing experimental system does not overall outperform the LSTM baseline even on MAE, the gap in performance is effectively closed (4.556 vs. 4.574). This means that we have devised a system which gives similar performance in the MAE sense as our existing one, but provides greater compatibility with our acoustic models, as it too operates at the frame level. The next logical step, therefore, is to apply this idea to joint modelling of duration and acoustic features.

6. CONCLUSIONS AND FUTURE WORK

We have described a new duration modeling paradigm with LSTMs which operates at frame-level for duration prediction in speech synthesis. Experiments conducted on an audiobook data were found to perform competitively when compared with a baseline phone-level LSTM system. As a next step, we will explore the generation based on other quantiles such as mean so as to allow the change in rate of speaking. Also, the frame-level joint acoustic/duration model should allow iterative refinement of the data alignments. Future work includes joint modeling of duration and acoustic features along with subjective evaluation of synthesised speech.

7. ACKNOWLEDGEMENTS

This research was supported by EPSRC Programme Grant EP/I031022/1, Natural Speech Technology (NST). The NST research data collection may be accessed at <http://hdl.handle.net/10283/786>.

8. REFERENCES

- [1] Heiga Zen, Andrew Senior, and Mike Schuster, “Statistical parametric speech synthesis using deep neural networks,” in *Proc. ICASSP*, 2013, pp. 7962–7966.
- [2] Zhizheng Wu, Cassia Valentini-Botinhao, Oliver Watts, and Simon King, “Deep neural networks employing multi-task learning and stacked bottleneck features for speech synthesis,” in *Proc. ICASSP*, 2015, pp. 4460–4464.
- [3] Heiga Zen and Haşim Sak, “Unidirectional long short-term memory recurrent neural network with recurrent output layer for low-latency speech synthesis,” in *Proc. ICASSP*, 2015, pp. 4470–4474.
- [4] Zhizheng Wu, Pawel Swietojanski, Christophe Veaux, Steve Renals, and Simon King, “A study of speaker adaptation for DNN-based speech synthesis,” in *Proc. Interspeech*, 2015.
- [5] Oliver Watts, Zhizheng Wu, and Simon King, “Sentence-level control vectors for deep neural network speech synthesis,” in *Proc. Interspeech*, 2015, pp. 2217–2221.
- [6] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu, “Wavenet: A generative model for raw audio,” *ArXiv*, 2016.
- [7] Dennis H. Klatt, “Review of text-to-speech conversion for English,” *J. Acoust. Soc. Am.*, vol. 82, no. 3, pp. 737–793, 1987.
- [8] Ivan Bulyko and Mari Ostendorf, “Joint prosody prediction and unit selection for concatenative speech synthesis,” in *Proc. ICASSP*, 2001, vol. 2, pp. 781–784.
- [9] Heiga Zen, Keiichi Tokuda, and Alan W. Black, “Statistical parametric speech synthesis,” *Speech Commun.*, vol. 51, no. 11, pp. 1039–1064, 2009.
- [10] Simon King, “An introduction to statistical parametric speech synthesis,” *Sadhana*, vol. 36, no. 5, pp. 837–852, 2011.
- [11] Heiga Zen, Keiichi Tokuda, Takashi Masuko, Takao Kobayashi, and Tadashi Kitamura, “Hidden semi-markov model based speech synthesis,” in *Proc. Interspeech*, 2004, pp. 1393–1396.
- [12] Heiga Zen, Takashi Nose, Junichi Yamagishi, Shinji Sako, Takashi Masuko, Alan Black, and Keiichi Tokuda, “The HMM-based speech synthesis system (HTS) version 2.0,” in *Proc. SSW*, 2007, vol. 6, pp. 294–299.
- [13] W. Nick Campbell, “Syllable-level duration determination,” in *Proc. Eurospeech*, 1989, pp. 2698–2701.
- [14] K. Huber, “A statistical model of duration control for speech synthesis,” in *Proc. EUSIPCO*, 1990, pp. 1127–1130.
- [15] Heiga Zen, Yannis Agiomyriannakis, Niels Egberts, Fergus Henderson, and Przemysław Szczepaniak, “Fast, compact, and high quality LSTM-RNN based statistical parametric speech synthesizers for mobile devices,” in *Proc. Interspeech*, 2016.
- [16] Gustav Eje Henter, Thomas Merritt, Matt Shannon, Catherine Mayo, and Simon King, “Measuring the perceptual effects of modelling assumptions in speech synthesis using stimuli constructed from repeated natural speech,” in *Proc. Interspeech*, 2014, pp. 1504–1508.
- [17] H. L. MacGillivray, “The mean, median, mode inequality and skewness for a class of densities,” *Australian Journal of Statistics*, vol. 23, no. 2, pp. 247–250, 1981.

- [18] Sanjib Basu and Anirban DasGupta, “The mean, median, and mode of unimodal distributions: a characterization,” *Theory Probab. Appl.*, vol. 41, no. 2, pp. 210–223, 1997.
- [19] Keiichi Tokuda, Takayoshi Yoshimura, Takashi Masuko, Takao Kobayashi, and Tadashi Kitamura, “Speech parameter generation algorithms for HMM-based speech synthesis,” in *Proc. ICASSP*, 2000, vol. 3, pp. 1315–1318.
- [20] Gustav Eje Henter, Srikanth Ronanki, Oliver Watts, Mirjam Wester, Zhizheng Wu, and Simon King, “Robust TTS duration modelling using DNNs,” in *Proc. ICASSP*, 2016, vol. 41, pp. 5130–5134.
- [21] Oliver Watts, Srikanth Ronanki, Zhizheng Wu, Tuomo Raitio, and Antti Suni, “The NST–GlottHMM entry to the Blizzard Challenge 2015,” in *Proc. Blizzard Challenge Workshop*, 2015.
- [22] Raul Fernandez, Asaf Rendel, Bhuvana Ramabhadran, and Ron Hoory, “Prosody contour prediction with long short-term memory, bi-directional, deep recurrent neural networks,” in *Proc. Interspeech*, 2014, pp. 2268–2272.
- [23] Srikanth Ronanki, Gustav Eje Henter, Zhizheng Wu, and Simon King, “A template-based approach for speech synthesis intonation generation using LSTMs,” in *Proc. Interspeech*, 2016.
- [24] Robert F. Kubichek, “Mel-cepstral distance measure for objective speech quality assessment,” in *Proc. IEEE Pac. Rim Conf. Commun. Comput. Signal Process.*, 1993, vol. 1, pp. 125–128.
- [25] Simon King and Vasilis Karaiskos, “The Blizzard Challenge 2016,” in *Proceedings of Blizzard Challenge Workshop*, 2016.
- [26] Kishore Prahallad, Alan W. Black, and Ravishankhar Mosur, “Sub-phonetic modeling for capturing pronunciation variations for conversational speech synthesis,” in *Proc. ICASSP*, 2006, pp. I-853–I-856.
- [27] Zhizheng Wu and Simon King, “Investigating gated recurrent networks for speech synthesis,” in *Proc. ICASSP*, 2016, pp. 5140–5144.