

A Neural Vocoder with Hierarchical Generation of Amplitude and Phase Spectra for Statistical Parametric Speech Synthesis

Yang Ai, *Student Member, IEEE*, Zhen-Hua Ling, *Senior Member, IEEE*

Abstract—This paper presents a neural vocoder named HiNet which reconstructs speech waveforms from acoustic features by predicting amplitude and phase spectra hierarchically. Different from existing neural vocoders such as WaveNet, SampleRNN and WaveRNN which directly generate waveform samples using single neural networks, the HiNet vocoder is composed of an amplitude spectrum predictor (ASP) and a phase spectrum predictor (PSP). The ASP is a simple DNN model which predicts log amplitude spectra (LAS) from acoustic features. The predicted LAS are sent into the PSP for phase recovery. Considering the issue of phase warping and the difficulty of phase modeling, the PSP is constructed by concatenating a neural source-filter waveform generator with a phase extractor. Finally, the outputs of ASP and PSP are combined to reconstruct speech waveforms by short-time Fourier synthesis. Since there are no autoregressive structures in both predictors, the HiNet vocoder can generate speech waveforms with high efficiency. Objective and subjective experimental results show that our proposed HiNet vocoder achieves better naturalness of reconstructed speech than the conventional STRAIGHT vocoder and a 16-bit WaveNet vocoder using open source implementation, and obtains similar performance with a 16-bit WaveRNN vocoder, no matter using natural or predicted acoustic features as input. We also find that the performance of HiNet is insensitive to the complexity of the neural waveform generator in PSP to some extent. After simplifying its model structure, the time consumed for generating 1s waveforms of 16kHz speech can be further reduced from 0.34s to 0.19s without significant quality degradation.

Index Terms—vocoder, neural network, amplitude spectrum, phase spectrum, statistical parametric speech synthesis

I. INTRODUCTION

SPEECH synthesis, a technology that converts texts into speech waveforms, plays a more and more important role in people’s daily life. A speech synthesis system with high intelligibility, naturalness and expressiveness is a goal pursued by speech synthesis researchers. Recently, statistical parametric speech synthesis (SPSS) has become a widely used speech synthesis framework due to its flexibility achieved by acoustic modeling and vocoder-based waveform generation. Hidden Markov models (HMMs) [1], deep neural networks (DNNs) [2], recurrent neural networks (RNNs) [3] and other deep learning models [4] have been applied to build the acoustic models for SPSS. Vcoders [5] which reconstruct speech waveforms from acoustic features (e.g., mel-cepstra and F0) also play an important role in SPSS. Their performance affects the quality of synthetic speech significantly.

Some conventional vocoders, such as STRAIGHT [6] and WORLD [7] which are designed based on the source-filter model of speech production [8], have been popularly applied in current SPSS systems. However, these vocoders still have some deficiencies, such as the loss of spectral details and phase information.

Recently, some neural generative models for raw audio signals [9]–[11] have been proposed and demonstrated good performance. For example, WaveNet [9] and SampleRNN [10] predicted the distribution of each waveform sample conditioned on previous samples and additional conditions using convolutional neural networks (CNNs) and RNNs respectively. These models represented waveform samples as discrete symbols. Although the μ -law quantization strategy [12] has been applied, the neural waveform generators with low quantization bits (e.g., 8-bit or 10-bit) always suffered from perceptible quantization errors. In order to achieve 16-bit quantization of speech waveforms, the WaveRNN model [11] was proposed, which generated 16-bit waveforms by splitting the RNN state into two parts and predicting the 8 coarse bits and the 8 fine bits respectively. However, due to the autoregressive generation manner, these models were very inefficient at generation stage. Therefore, some variants such as knowledge-distilling-based models (e.g., parallel WaveNet [13] and ClariNet [14]) and flow-based models (e.g., WaveGlow [15]) were then proposed to improve the efficiency of generation.

Neural vocoders based on these waveform generation models [16]–[21] have been developed to reconstruct speech waveforms from various acoustic features for SPSS and some other tasks, such as voice conversion [22], [23], bandwidth extension [24] and speech coding [25]. Experimental results confirmed that these neural vocoders performed significantly better than conventional ones. Some improved neural vocoders, such as glottal neural vocoder [26]–[28], LP-WaveNet [29], LPCNet [30], FFTNet [31] and neural source-filter (NSF) vocoder [32], have been further proposed by combining speech production mechanisms with neural networks and have also demonstrated impressive performance.

There are still some limitations with current neural vocoders and the most significant one is that they have much higher computation complexity than conventional STRAIGHT and WORLD vocoders. The autoregressive neural vocoders (e.g., WaveNet, SampleRNN and WaveRNN) are very inefficient at synthesis time due to their point-by-point generation process. The knowledge-distilling-based vocoders (e.g., parallel

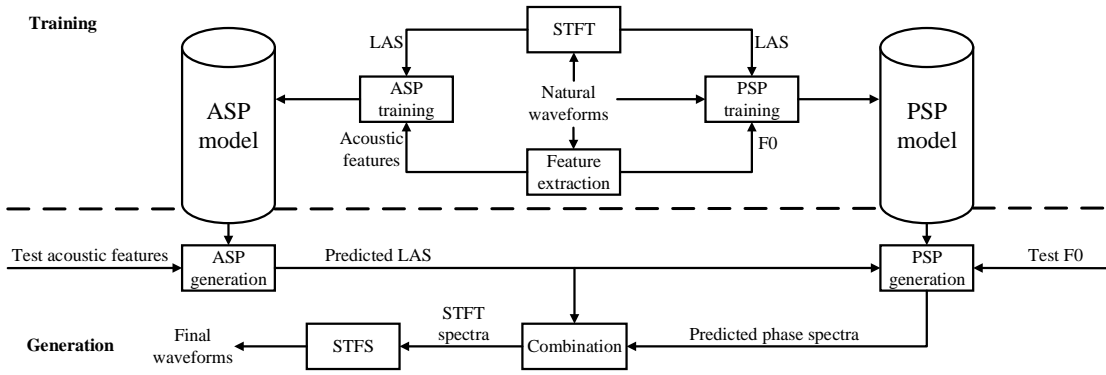


Fig. 1. The flowchart of the training and generation processes of our proposed HiNet vocoder. Here, *ASP*, *PSP* and *LAS* stand for amplitude spectrum predictor, phase spectrum predictor and log amplitude spectra respectively.

WaveNet and ClariNet) and the flow-based vocoders (e.g., WaveGlow) accelerate the generation process by removing autoregressive connections. However, they suffer from the complexity of model structures and the difficulty of model training.

This paper explores the approaches to improve the run-time efficiency of neural vocoders by combining neural waveform generation models with the frequency-domain representation of speech waveforms. Inspired by the knowledge that speech waveforms can be perfectly reconstructed from their short-time Fourier transform (STFT) results which consist of frame-level amplitude spectra and phase spectra, this paper proposes a neural vocoder which recovers speech waveforms by predicting amplitude and phase spectra hierarchically from input acoustic features. We name this vocoder *HiNet* because it is expected to generate waveforms with *high* quality and *high* efficiency by *hierarchical* prediction. Different from existing neural vocoders which directly generate waveform samples using single neural networks, the HiNet vocoder is composed of an amplitude spectrum predictor (ASP) and a phase spectrum predictor (PSP). The ASP is a simple DNN which predicts frame-level log amplitude spectra (LAS) from acoustic features. Then, the predicted LAS are sent into the PSP for phase recovery. Considering the issue of phase warping and the difficulty of phase modeling, the PSP is constructed by concatenating a neural waveform generator with a phase extractor. Since the task of the neural waveform generator in PSP is not to generate the final waveforms but to supplement the amplitude spectra with phase information, some light-weight models can be adopted even if their overall prediction accuracy is not perfect. In our implementation, the neural waveform generator is built by adapting the non-autoregressive NSF vocoder [32] from three aspects. First, LAS are used as the input of PSP rather than spectral features (e.g., mel-cepstra). Second, the initial phase of the sine-based excitation signal is pre-calculated for each voiced segment at the training stage of the PSP to benefit phase modeling. Third, a waveform loss and a correlation loss are introduced into the complete loss function in order to enhance its ability of measuring phase distortion. Finally, the outputs of ASP and PSP are combined to recover speech waveforms by short-

time Fourier synthesis (STFS). Experimental results show that the proposed HiNet vocoder achieves better naturalness of reconstructed speech than the conventional STRAIGHT vocoder and a 16-bit WaveNet vocoder implemented by public source codes, and obtains similar performance with a 16-bit WaveRNN vocoder, no matter using natural or predicted acoustic features as input.

There are two main characteristics of the HiNet vocoder. First, there are no autoregressive structures in both predictors. Thus, the HiNet vocoder is able to generate speech waveforms with high efficiency by parallel computation. Second, the neural waveform generator only contributes to the prediction of phase spectra. Further experimental results reveal that the performance of HiNet is insensitive to the complexity of the neural waveform generator in PSP to some extent. After simplifying its model structure, the time consumed for generating 1s waveforms of 16kHz speech can be further reduced from 0.34s to 0.19s without significant quality degradation.

This paper is organized as follows. In Section II, we present our proposed HiNet vocoder in detail. Section III reports our experimental results and conclusions are given in Section IV.

II. PROPOSED METHODS

The proposed HiNet vocoder consists of an amplitude spectrum predictor (ASP) and a phase spectrum predictor (PSP). The flowchart of its training and generation processes is illustrated in Fig. 1.

A. Amplitude Spectrum Predictor

The ASP predicts LAS from input acoustic features. For better generation efficiency, a simple DNN without any autoregressive structures is adopted to build the ASP as shown in Fig. 2.

Let $\mathbf{a}_n = [a_{n,1}, \dots, a_{n,C}]^T$ and $\mathbf{l}_n = [l_{n,1}, \dots, l_{n,K}]^T$ denote the acoustic features and the LAS at the n -th frame respectively, where n , c and k represent the frame index, the dimension index and the frequency bin index, C and K denote the total numbers of acoustic feature dimensions and frequency bins respectively. For utilizing history information, the model input is a concatenation of current frame and n_p previous

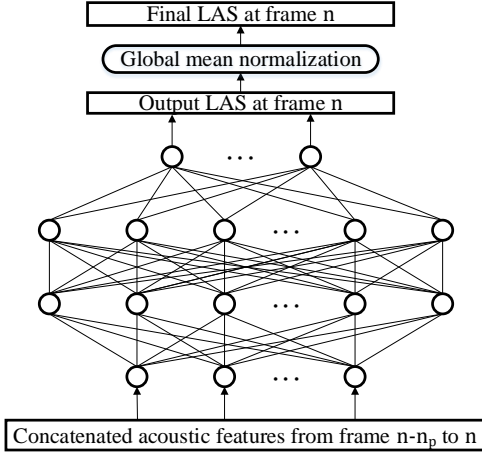


Fig. 2. Model structure of the amplitude spectrum predictor (ASP).

frames (i.e., $[\mathbf{a}_{n-n_p}^\top, \dots, \mathbf{a}_{n-1}^\top, \mathbf{a}_n^\top]^\top$). The model output is the LAS of current frame as shown in Fig. 2.

At the training stage, parallel acoustic features and LAS are extracted from natural waveforms. The training criterion is to minimize the mean square error (MSE) between the predicted and the real LAS as

$$\mathcal{L}_{ASP} = \frac{1}{NK} \sum_n \sum_k (l_{n,k} - \hat{l}_{n,k})^2 + L_2\text{-reg}, \quad (1)$$

where $\hat{l}_n = [\hat{l}_{n,1}, \dots, \hat{l}_{n,K}]$ is the predicted LAS at the n -th frame, and N is the total number of frames. $L_2\text{-reg}$ is an L2 regularization term of all weights in the model for avoiding overfitting.

Then, a global mean normalization (GMN) operation is conducted to compensate the global distortion between the amplitude spectra predicted by the DNN and the natural ones. For the k -th frequency bin, a compensation factor q_k is estimated given the trained DNN as

$$q_k = \frac{\sum_n \exp(l_{n,k})}{\sum_n \exp(\hat{l}_{n,k})}, \quad (2)$$

where n denotes the frame index of the training set. The vector $\mathbf{q} = [q_1, \dots, q_K]^\top$ further passes through a median filter along the frequency axis to get a smoothed curve $\mathbf{q}^{mf} = [q_1^{mf}, \dots, q_K^{mf}]^\top$. At the generation stage, the final LAS at each frame is obtained by

$$\hat{l}_n^{FNL} = \log \left(\exp(\hat{l}_n) \odot \mathbf{q}^{mf} \right), \quad (3)$$

where \odot represents element-wise product. In our informal listening tests, we found that this GMN operation can help to improve the subjective performance of ASP.

We can see that the whole ASP model operates at frame-level without any autoregressive calculations. Therefore, its training and generation processes are very efficient.

B. Phase Spectrum Predictor

The aim of PSP is to recover phase spectra given input amplitude spectra. However, modeling and predicting phase spectra directly are difficult due to the issue of phase warping.

Since temporal waveforms contain the information of both amplitude and phase spectra, this paper proposes to predict phase spectra utilizing a neural waveform generator. For predicting phase spectra efficiently, the waveform generator is designed based on the non-autoregressive NSF vocoder [32]. Several modifications are made in order to focus on recovering phase spectra from input amplitude spectra. Finally, the phase spectra extracted from the waveforms generated by the neural waveform generator are used as the outputs of PSP.

The structure of PSP is illustrated in Fig. 3, which first converts input LAS and F0 sequences into waveforms $\hat{\mathbf{x}} = [\hat{x}_1, \dots, \hat{x}_T]^\top$ using a neural waveform generator and then extracts phase spectra $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_N]$ from $\hat{\mathbf{x}}$ by STFT analysis. At the training stage, LAS and F0 sequences are calculated from natural waveform $\mathbf{x} = [x_1, \dots, x_T]^\top$ and the loss functions are defined between \mathbf{x} and $\hat{\mathbf{x}}$. At the generation stage, the PSP adopts the test F0 sequence and the LAS predicted by ASP as inputs. Similar to the NSF vocoder [32], the neural waveform generator in PSP consists of a source module and a filter module. The details of these two modules and the loss functions will be introduced in the following subsections.

1) *Source Module*: The upsampled F0 sequence $\mathbf{f} = [f_1, \dots, f_T]^\top$ is obtained by repeating the F0 values within each frame, and is used as the input of the source module. The output of the source module is an excitation signal $\mathbf{e} = [e_1, \dots, e_T]^\top$, which is a sine-based signal for voiced segments and a DNN-transformed Gaussian white noise for unvoiced segments. Mathematically, for time step t , the excitation signal e_t is defined as

$$e_t = \begin{cases} \alpha \sin\left(\sum_{h=1}^t 2\pi f_h \frac{1}{N_s} + \phi_j\right) + n_t, & f_t > 0, t \in V_j \\ g\left(\frac{1}{3\sigma} n_t\right), & f_t = 0 \end{cases}, \quad (4)$$

where $f_t = 0$ denotes that the t -th sampling point belongs to an unvoiced frame, $g(\cdot)$ represents a DNN-based transformation, $n_t \sim \mathcal{N}(0, \sigma^2)$ is a Gaussian white noise at time t , N_s is the sampling rate of waveforms, V_j is the j -th voiced segment that the t -th sampling point belongs to, $\phi_j \in (-\pi, \pi]$ is the initial phase of the j -th voiced segment, α and σ are hyperparameters. At the training stage, we estimate the initial phase ϕ_j of each voiced segment for better phase modeling. First, the j -th voiced segment of the natural waveform \mathbf{x} passes through a low-pass filter whose cut-off frequency is the maximal F0 of this segment in order to obtain a reference waveform without formant influence. Then, ϕ_j is determined by maximizing the correlation coefficient between the sine wave in Eq. (4) and the reference waveform for each voiced segment. At the generation stage, ϕ_j is set as a random initial phase. Only the DNN for noise transformation has model parameters in the source module.

2) *Filter Module*: The excitation signal \mathbf{e} generated by the source module and the upsampled LAS sequence $\mathbf{L} = [l_1, \dots, l_T]$ are input into the filter module. Before upsampling, the frame-level LAS features pass through GRU-based

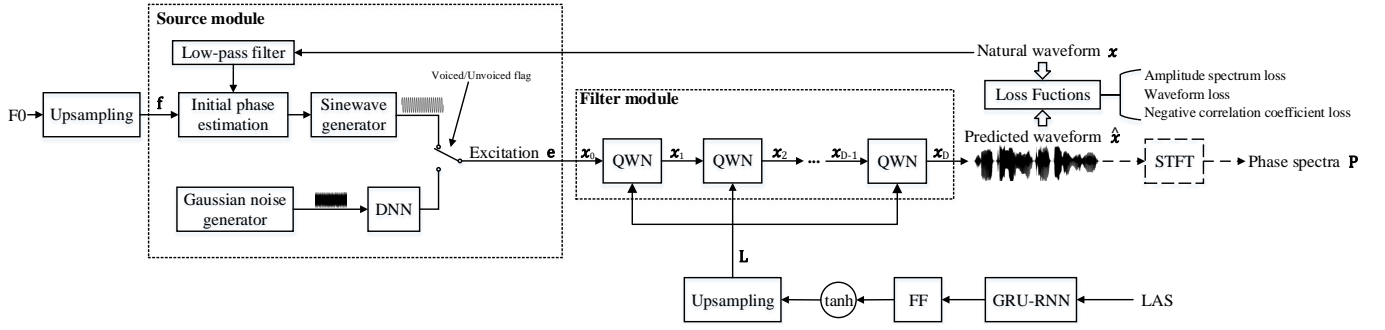


Fig. 3. Model structure of the phase spectrum predictor (PSP). Here, FF and $GRU-RNN$ represent feed-forward and unidirectional GRU-based recurrent layers respectively, DNN represents a deep neural network with multiple FF layers, \tanh denotes the hyperbolic tangent functions, QWN is a quasi WaveNet structure as shown in Fig. 4. The dotted lines indicates the source module and filter module. The dashed lines represent the operations only used at the generation stage.

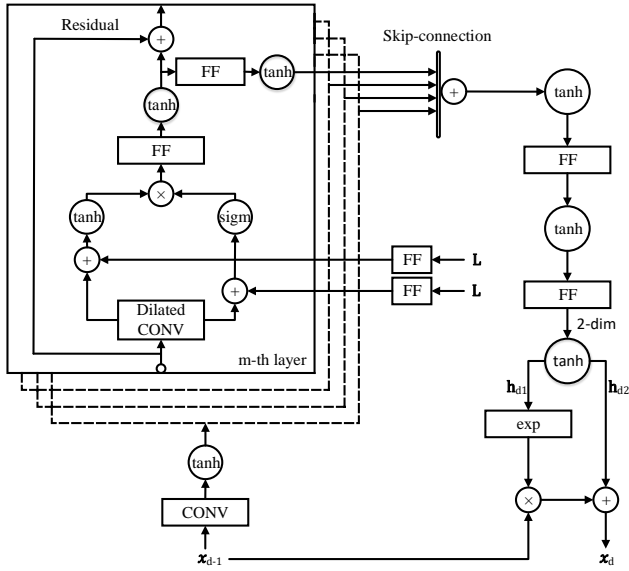


Fig. 4. The structure of the d -th QWN block in the PSP. Here, FF , $CONV$, $Dilated CONV$ represent the feed-forward, convolutional and dilated convolutional layers respectively, \tanh and sigm represent the hyperbolic tangent and sigmoid functions respectively.

recurrent layers and feed-forward layers for pre-processing. The output of the filter module is the predicted waveform \hat{x} .

As shown in Fig. 3, the filter module is a concatenation of D quasi WaveNet (denoted by QWN) blocks. Assume $x_0 = e$ and $x_D = \hat{x}$. The d -th QWN uses sequence x_{d-1} and L as input and predicts sequence x_d . The detailed structure of the d -th QWN is illustrated in Fig. 4. A QWN block is similar to a WaveNet model [9] whose key elements include dilated convolutions, gated activation units, residual connections and skip connections. The difference is that QWNs are non-autoregressive with non-causal convolution because the whole sequence x_{d-1} is already known for the d -th block. The LAS features are connected to the gated activation units after passing through two FF layers. The hyperbolic tangent activation function is used in QWNs because the range of waveform samples is from -1 to 1. The two FF layers after the skip connections are employed to reduce the dimensionality of the skip output and to generate sequences h_{d1} and h_{d2} both

with length T . Finally, the output sequence x_d is calculated as

$$x_d = x_{d-1} \odot \exp(h_{d1}) + h_{d2}, \quad (5)$$

where \odot represents element-wise product. The output of the last QWN $x_D = \hat{x}$ is used to define the loss function at the training stage and to extract phase spectra at the generation stage.

3) *Model Training*: Three loss functions are defined between the predicted waveform \hat{x} and the natural reference x , including amplitude spectrum loss, waveform loss and negative correlation coefficient loss. Comparing with the original NSF vocoder [32], the last two losses are added for indirectly evaluating the phase accuracy of the predicted waveforms.

The amplitude spectrum loss is the MSE between the natural amplitude spectra and the predicted ones which are derived from x and \hat{x} using STFT respectively. Similar to the NSF vocoder [32], multiple sets of frame length (FL), frame shift (FS), and FFT point number (FN) are adopted for STFT in our implementation. For the i -th set of (FL^i, FS^i, FN^i) , the amplitude spectrum loss is calculated as

$$\mathcal{L}_{ASi} = \frac{1}{N^i K^i} \sum_{n=1}^{N^i} \sum_{k=1}^{K^i} (A_{n,k}^i - \hat{A}_{n,k}^i)^2, \quad (6)$$

where $A_{n,k}^i$ and $\hat{A}_{n,k}^i$ are the spectral amplitude at frame n and frequency bin k of x and \hat{x} respectively, N^i denotes the total number of frames and $K^i = \frac{FN^i}{2} + 1$.

The waveform loss is defined as the MSE between the natural waveform samples and the predicted ones, i.e.,

$$\mathcal{L}_W = \frac{1}{T} \sum_{t=1}^T (x_t - \hat{x}_t)^2. \quad (7)$$

The negative correlation coefficient loss is calculated as the negative correlation coefficient between the natural waveform and the predicted waveform, i.e.,

$$\mathcal{L}_C = -\frac{E[(x - E(x))(\hat{x} - E(\hat{x}))]}{\sqrt{Var(x)Var(\hat{x})}}, \quad (8)$$

where the functions $E(\cdot)$ and $Var(\cdot)$ calculate mean and variance respectively.

Finally, the training criterion of the waveform generator in PSP is to minimize the combined loss function as

$$\mathcal{L}_{PSP} = \sum_i \mathcal{L}_{ASi} + \mathcal{L}_W + \mathcal{L}_C. \quad (9)$$

III. EXPERIMENTS

A. Experimental Setup

The recordings of the female speaker *slt* and the male speaker *bdl* in CMU-ARCTIC databases [33] which contained English speech with 16 kHz sampling rate and 16bits resolution were adopted in our experiments. For each speaker, we chose 1000 and 66 utterances to construct the training set and the validation set respectively, and the remaining 66 utterances were used as the test set. The acoustic features at each frame were 43-dimensional including 40-dimensional mel-cepstra, an energy, an F0 and a voiced/unvoiced (V/UV) flag. The natural acoustic features were extracted by STRAIGHT and the window size was 400 samples (i.e., 25ms) and the window shift was 80 samples (i.e., 5ms). This paper focuses on neural vocoders, thus a simple acoustic model for SPSS was used in our experiments. A bidirectional LSTM-RNN acoustic model [3] having 2 hidden layers with 1024 units per layer (512 forward units and 512 backward units) was trained to predict acoustic features from linguistic features. The input linguistic features were the same as the ones used in Merlin toolkit [34] for CMU-ARCTIC databases which were 425-dimensional. The output of the acoustic model contained the acoustic features together with their delta and acceleration counterparts, which were totally 127 dimensions (the V/UV flag had no dynamic components). Then, the predicted acoustic features were generated from the output by maximum likelihood parameter generation (MLPG) algorithm [35] considering global variance (GV) [36].

Four vocoders were compared in our experiments¹. The descriptions of these vocoders are as follows.

- **STRAIGHT** The conventional STRAIGHT vocoder. At synthesis time, the spectral envelope at each frame was first reconstructed from input mel-cepstra and frame energy, and was then used to generate speech waveforms together with input source parameters (i.e., F0 and V/UV flag) [6].
- **WaveNet** A 16-bit WaveNet-based neural vocoder, which is the teacher model used in parallel WaveNet [13]. Two speaker-dependent vocoders were trained using an open source implementation². 3 upsampling layers with upsampling rates {5,4,4} were adopted. Other configurations remained the same as that of the open source implementation. The built model was a mixture density network, outputting the parameters for a mixture of 10 logistic distributions at each timestep and had 24 dilated casual convolutional layers which were divided into 4 convolutional blocks. Each block contained 6 layers and their dilation coefficients were $\{2^0, 2^1, \dots, 2^5\}$. The filter width was 3. The number of gate channels in gated

activation units was 512. For the residual architectures, the number of residual channels was 512 and the number of skip channels was 256. An *Adam* optimizer [37] was used to update the parameters by minimizing the negative log likelihood. Models were trained and evaluated on a single Nvidia 1080Ti GPU using PyTorch framework [38].

- **WaveRNN** A 16-bit WaveRNN-based neural vocoder implemented by ourselves. The structure was the same as the one used in our previous work [20] which did not adopt the efficiency optimization strategies introduced in [11]. The built model had one hidden layer of 1024 nodes where 512 nodes for coarse outputs and another 512 nodes for fine outputs. The waveform samples were quantized to discrete values by 16-bit linear quantization. Truncated back propagation through time (TBPTT) algorithm was employed to improve the efficiency of model training and the truncated length was set to 480. An *Adam* optimizer [37] was used to update the parameters by minimizing the cross-entropy. Models were trained and evaluated on a single Nvidia 1080Ti GPU using TensorFlow framework [39].
- **HiNet** Our proposed HiNet neural vocoder. When extracting LAS, the frame length and frame shift of STFT were 640 samples (i.e., 40ms) and 80 samples (i.e., 5ms) respectively and the FFT point number was 1024. For ASP, the acoustic features at current frame along with 6 previous frames (i.e., $n_p = 6$) were concatenated to form the complete input which was 301-dimensional. There were two hidden layers with 2048 nodes per layer, and a 513-dimensional linear output layer which predicted the LAS at current frame. The activation function was rectified linear units (ReLU) for hidden layers. For PSP, an unidirectional GRU layer with 1024 nodes and an FF layer with 128 nodes were used to pre-process LAS. When extracting phase spectra from the predicted waveform at the generation stage, the STFT parameter settings were consistent with the ones used for extracting LAS. In the source module, the DNN for transforming Gaussian noise had two FF layers with 512 nodes per layer and hyperbolic tangent activation function together with a 1-dimensional linear output layer. Hyperparameters α and σ were set as 0.1 and 0.003 respectively. Referring to the configuration of original NSF model [32], the filter module consisted of 5 QWN blocks (i.e., $D = 5$). Each QWN had a non-causal convolutional layer for processing the input sequence and 10 dilated non-causal convolution layers and their dilation coefficients were $\{2^0, 2^1, \dots, 2^9\}$. The filter width was 5. The number of gate channels in gated activation units was 128. The additional inputs L were connected to the gated activation units after passing through two FF layers both having 128 nodes. For the residual architectures, the number of residual channels was 128 and the number of skip channels was 256. After the skip connections, an FF layer with 16 nodes and an FF layer with 2 nodes were used to reduce the dimensionality of the skip output. For the loss function of PSP, two sets of

¹Examples of generated speech can be found at http://home.ustc.edu.cn/~ay8067/IEEETran_2019/demo.html.

² https://github.com/r9y9/wavenet_vocoder.

TABLE I

OBJECTIVE EVALUATION RESULTS OF FOUR VOCODERS ON THE TEST SETS OF TWO SPEAKERS. HERE, ‘‘R’’ STANDS FOR USING NATURAL ACOUSTIC FEATURES AS INPUT AND ‘‘P’’ STANDS FOR USING PREDICTED ACOUSTIC FEATURES AS INPUT.

		R				P			
		<i>STRAIGHT</i>	<i>WaveNet</i>	<i>WaveRNN</i>	<i>HiNet</i>	<i>STRAIGHT</i>	<i>WaveNet</i>	<i>WaveRNN</i>	<i>HiNet</i>
<i>slt</i>	SNR(dB)	0.5357	3.5228	6.0568	6.2937	–	–	–	–
	LAS-RMSE(dB)	5.5800	6.0681	6.2489	5.5937	–	–	–	–
	MCD(dB)	1.3315	1.5950	1.6042	1.5036	1.5793	1.6335	1.5210	1.5910
	F0-RMSE(cent)	14.8430	71.9886	12.1309	8.0286	12.7086	93.5050	15.4046	7.0443
	V/UV error(%)	3.3994	4.6260	3.3756	2.1971	3.0498	4.5810	3.1356	2.2715
<i>bdl</i>	SNR(dB)	1.0987	2.7105	3.8993	4.5905	–	–	–	–
	LAS-RMSE(dB)	5.6434	6.0581	6.1812	5.7486	–	–	–	–
	MCD(dB)	1.3097	1.4093	1.5150	1.5528	1.6037	1.4575	1.5187	1.4960
	F0-RMSE(cent)	25.7898	98.3218	21.0020	10.5880	20.1130	122.6858	21.5037	8.9084
	V/UV error(%)	4.5588	8.7091	5.5817	2.7663	4.3126	7.7423	5.4913	3.0644

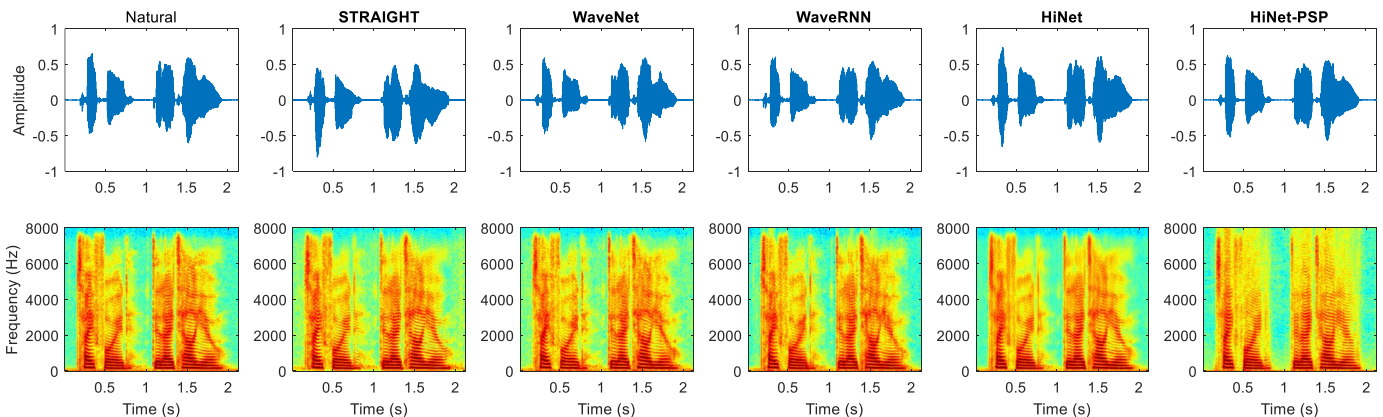


Fig. 5. The waveforms and spectrograms of natural speech and the speech generated by different vocoders when using natural acoustic features as input for an example sentence (arctic_b0536) in the test set of speaker *slt*. Here, *HiNet-PSP* denotes the waveforms generated by the PSP in HiNet.

STFT configurations (FL, FS, FN), i.e., (320, 80, 512) and (80, 40, 128), were used for the amplitude spectrum loss. An *Adam* optimizer [37] was used to update the parameters by minimizing \mathcal{L}_{ASP} and \mathcal{L}_{PSP} for ASP and PSP respectively. Truncated waveform sequences with 16000 samples were used for training PSP to avoid the overflow of GPU memory. Models were trained and evaluated on a single Nvidia 1080Ti GPU using TensorFlow framework [39].

B. Objective Evaluation

In this section, we compared the performance of the four vocoders mentioned in Section III-A by objective evaluation.

When using natural acoustic features as input, we compared the distortions between natural speech and the speech reproduced by these four vocoders. Five metrics used in [16] were adopted here, including signal-to-noise ratio (SNR) which reflected the distortion of waveforms, root MSE (RMSE) of LAS (denoted by LAS-RMSE) which reflected the distortion in frequency domain, mel-cepstrum distortion (MCD) which described the distortion of mel-cepstra, MSE of F0 which reflected the distortion of F0 (denoted by F0-RMSE), and V/UV error which was the ratio between the number of frames with mismatched V/UV flags and the total number of frames. Among these metrics, SNR can be considered as an overall measurement on the distortions of both amplitude and

phase spectra, while LAS-RMSE and MCD mainly present the distortion of amplitude spectra. STRAIGHT was used to extract acoustic features from both original and reproduced speech waveforms for calculating all these metrics. When using the acoustic features predicted by the acoustic model as vocoder input, only the metrics of MCD, F0-RMSE and V/UV error were adopted because the calculation of SNR and LAS-RMSE relied on natural speech waveforms.

The results on the test sets of the two speakers are listed in Table I. It is obvious that the STRAIGHT vocoder achieved the lowest SNR for both speakers due to the neglect of natural phase information. Our proposed HiNet vocoder outperformed the WaveNet vocoder and the WaveRNN vocoder on the SNR metric for both speakers. This indicated that the HiNet vocoder restored the shape of waveforms more accurately than other vocoders. Besides, our proposed HiNet vocoder achieved the lowest LAS-RMSE among the three neural vocoders which implied the advantage of using a separate ASP in our proposed method. Regarding with MCD, the results on these two speakers were inconsistent which needed further investigation. Our proposed HiNet vocoder achieved the lowest F0-RMSE and V/UV error among all four vocoders and their differences were significant, no matter natural or predicted acoustic features were used as inputs. This advantage can be attributed to the explicit excitation signal determined by F0s and U/V flags in the PSP of HiNet.

Fig. 5 shows the waveforms and spectrograms of natural

TABLE II
REAL TIME FACTORS (RTFs) OF THREE NEURAL VOCODERS.

Vocoder	WaveNet	WaveRNN	HiNet
RTF	222.3656	100.9148	0.3420




Fig. 6. Average MUSHRA scores with 95% confidence interval of the four vocoders for speaker *slt*. “R” stands for using natural acoustic features as input and “P” stands for using predicted acoustic features as input.

speech and the speech generated by different vocoders when using natural acoustic features as input for an example sentence in the test set of speaker *slt*. We can see that there was observable difference between the overall contours of the waveforms generated by STRAIGHT and the natural waveforms due to the neglect of natural phase information in STRAIGHT. In contrast, the neural vocoders (i.e., WaveNet, WaveRNN and HiNet) restored the overall waveform contours much better. Besides, our proposed HiNet vocoder was better at reconstructing the high-frequency harmonic structures of some voiced segments (e.g., 1.4~1.6s and 4000~6000Hz in Fig. 5) as shown in the spectrograms.

In order to evaluate the run-time efficiency of different neural vocoders, real time factor (RTF) which is defined as the ratio between the time consumed to generate speech waveforms and the duration of the generated speech was utilized as the measurement. In our implementation, the RTF value was calculated as the ratio between the time consumed to generate all test sentences using a single Nvidia 1080Ti GPU and the total duration of the test set. The results are listed in Table II. It can be observed that our proposed HiNet vocoder achieved the highest generation efficiency with RTF= 0.3420. A further analysis showed that 94.4% of the time used by HiNet was spent on the PSP. This inspires us to simplify the waveform generator in PSP for further improving the efficiency of HiNet in Section III-E. For the WaveNet and WaveRNN vocoders, they were very inefficient due to the point-by-point autoregressive generation.

C. Subjective Evaluation

Four MUSHRA (MUltiple Stimuli with Hidden Reference and Anchor) tests [40] were conducted to compare the naturalness of these four vocoders with natural recordings as references for both speakers and using both natural and predicted acoustic features as input. In each test, twenty test sentences synthesized by the four vocoders were evaluated by at least 30 English native listeners on the crowdsourcing




Fig. 7. Average MUSHRA scores with 95% confidence interval of the four vocoders for speaker *bdl*. “R” stands for using natural acoustic features as input and “P” stands for using predicted acoustic features as input.

platform of Amazon Mechanical Turk³ with anti-cheating considerations [41]. Listeners were asked to give a naturalness score between 0 and 100 to each sample and the reference natural recording had the maximum score of 100.

The average naturalness scores and their 95% confidence intervals of these four vocoders are shown in Fig. 6 and Fig. 7 for speaker *slt* and *bdl* respectively. The results of paired *t*-test showed that the HiNet vocoder outperformed STRAIGHT and the WaveNet vocoder significantly at significance level of 0.01 and the differences between the HiNet and the WaveRNN vocoders were not significant for both speakers, no matter using natural or predicted acoustic features as input. Besides, the differences between STRAIGHT and the WaveNet vocoder for speaker *bdl* when using predicted acoustic features as input were also not significant. This may be attributed to the severe F0 distortion (F0-RMSE=122.6858 cent in Table I) of the WaveNet vocoder for speaker *bdl* when using predicted acoustic features as input. Although our proposed HiNet vocoder achieved similar performance with that of the WaveRNN vocoder, its run-time efficiency was about 300 times higher as shown in Table II.

D. Performance of ASP and PSP in HiNet

In this section, we compared the performance of ASP and PSP in our proposed HiNet vocoder with other neural vocoders and phase spectrum reconstruction algorithms such as Griffin-Lim [42] by some combination experiments. Four variants of the HiNet vocoder were adopted for comparison as shown in Table III and IV and their descriptions are as follows.

- **HiNet-PSP** The waveforms generated by PSP of the HiNet vocoder (i.e., the predicted waveform \hat{x} in Fig. 3).
- **HiNet-ASP+WaveNet** The waveforms reconstructed by combining the amplitude spectra generated by ASP in the HiNet vocoder and the phase spectra extracted from the output of the WaveNet vocoder.
- **HiNet-ASP+WaveRNN** The waveforms reconstructed by combining the amplitude spectra generated by ASP in the HiNet vocoder and the phase spectra extracted from the output of the WaveRNN vocoder.

³<https://www.mturk.com>.

TABLE III

OBJECTIVE EVALUATION RESULTS OF THE HiNET VOCODER AND ITS FOUR VARIANTS USING NATURAL ACOUSTIC FEATURES AS INPUT ON THE TEST SETS OF TWO SPEAKERS.

		<i>HiNet</i>	<i>HiNet-PSP</i>	<i>HiNet-ASP+WaveNet</i>	<i>HiNet-ASP+WaveRNN</i>	<i>HiNet-ASP+GL</i>
<i>slt</i>	SNR(dB)	6.2937	6.1603	4.5333	6.2025	2.3445
	LAS-RMSE(dB)	5.5937	11.2823	5.5058	5.3876	6.0350
	MCD(dB)	1.5036	3.9385	1.4696	1.2614	1.3232
	F0-RMSE(cent)	8.0286	8.8295	63.9517	11.3556	15.1105
	V/UV error(%)	2.1971	2.6730	3.2803	2.5572	2.9047
<i>bdl</i>	SNR(dB)	4.5905	4.0259	3.3204	4.3568	1.4059
	LAS-RMSE(dB)	5.7486	10.9305	5.7996	5.6270	6.1129
	MCD(dB)	1.5528	3.7918	1.4753	1.3131	1.3471
	F0-RMSE(cent)	10.5880	14.3587	81.2990	18.2293	21.9268
	V/UV error(%)	2.7663	3.4705	7.8202	4.1802	4.4503

TABLE IV

OBJECTIVE EVALUATION RESULTS OF THE HiNET VOCODER AND ITS FOUR VARIANTS USING PREDICTED ACOUSTIC FEATURES AS INPUT ON THE TEST SETS OF TWO SPEAKERS.

		<i>HiNet</i>	<i>HiNet-PSP</i>	<i>HiNet-ASP+WaveNet</i>	<i>HiNet-ASP+WaveRNN</i>	<i>HiNet-ASP+GL</i>
<i>slt</i>	MCD(dB)	1.5910	3.8496	1.6043	1.3353	1.3197
	F0-RMSE(cent)	7.0443	7.3981	80.7298	14.8798	19.5369
	V/UV error(%)	2.2715	2.3762	4.4068	3.0550	4.1179
<i>bdl</i>	MCD(dB)	1.4960	3.7426	1.6080	1.3598	1.3656
	F0-RMSE(cent)	8.9084	9.9549	103.1390	18.9066	22.7011
	V/UV error(%)	3.0644	3.1576	9.0204	4.3158	4.3134

TABLE V

RTFs OF THE ORIGINAL HiNET VOCODER AND THE ONES AFTER MODEL SIMPLIFICATION.

Vocoder	<i>HiNet</i>	<i>HiNet-4QWN</i>	<i>HiNet-3QWN</i>	<i>HiNet-2QWN</i>	<i>HiNet-1QWN</i>	<i>HiNet-HC</i>	<i>HiNet-1QWN-HC</i>
RTF	0.3420	0.2996	0.2682	0.2559	0.2124	0.3057	0.1929

TABLE VI

OBJECTIVE EVALUATION RESULTS OF THE ORIGINAL HiNET VOCODER AND THE ONES AFTER MODEL SIMPLIFICATION ON THE TEST SETS OF SPEAKER *slt*. “R” STANDS FOR USING NATURAL ACOUSTIC FEATURES AS INPUT AND “P” STANDS FOR USING PREDICTED ACOUSTIC FEATURES AS INPUT.

		<i>HiNet</i>	<i>HiNet-4QWN</i>	<i>HiNet-3QWN</i>	<i>HiNet-2QWN</i>	<i>HiNet-1QWN</i>	<i>HiNet-HC</i>	<i>HiNet-1QWN-HC</i>
R	SNR(dB)	6.2937	6.2904	6.2539	6.1673	6.1559	6.1690	6.1952
	LAS-RMSE(dB)	5.5937	5.6040	5.6396	5.6235	5.6077	5.5770	5.6291
	MCD(dB)	1.5036	1.4996	1.5261	1.5093	1.4914	1.5243	1.5030
	F0-RMSE(cent)	8.0286	8.1837	8.0784	7.3329	7.3888	7.7502	7.4500
	V/UV error(%)	2.1971	2.2240	2.0337	2.0305	2.0891	2.1264	2.1019
P	MCD(dB)	1.5910	1.5711	1.6038	1.5854	1.5694	1.6036	1.5703
	F0-RMSE(cent)	7.0443	7.6475	6.9059	6.9587	6.5577	6.8257	6.6374
	V/UV error(%)	2.2715	2.3531	2.3668	2.1937	2.2319	2.3927	2.0910

- ***HiNet-ASP+GL*** The waveforms reconstructed by sending the amplitude spectra generated by ASP into the Griffin-Lim [42] algorithm with random initialization and 100 iterations.

We first analyzed the performance of ASP in our proposed HiNet vocoder. By comparing *HiNet* with *HiNet-PSP* in Table III and IV, it is obvious that the amplitude spectra generated by ASP outperformed the amplitude spectra of the waveforms generated by PSP on all metrics. As shown in Fig. 5, the spectrogram generated by *HiNet* for an example sentence in the test set of speaker *slt* was much more natural than that of *HiNet-PSP*. By comparing *WaveNet* and *WaveRNN* in Table I with *HiNet-ASP+WaveNet* and *HiNet-ASP+WaveRNN* in Table III and IV, we found that replacing the original amplitude spectra of the WaveNet vocoder and the WaveRNN vocoder with the amplitude spectra generated by the ASP of HiNet improved their performances

on most metrics. This result demonstrated the feasibility and effectiveness of predicting amplitude spectra using a simple frame-level DNN in our proposed HiNet vocoder.

We then analyzed the performance of PSP in our proposed HiNet vocoder. Although *HiNet-PSP* suffered from large LAS-RMSE and MCD, it achieved quite high waveform SNR as shown in Table III. This implied the effectiveness of recovering phase spectra using PSP. Besides, *HiNet-PSP* also achieved similar performance on F0-RMSE and V/UV error with *HiNet*. Furthermore, we compared *HiNet* with *HiNet-ASP+WaveNet*, *HiNet-ASP+WaveRNN* and *HiNet-ASP+GL*, since they shared the same amplitude spectra and employed different phase spectra. From Table III, we found that *HiNet* achieved the best waveform SNR performance among them. This result indicated that the phase spectra generated by the PSP of the HiNet vocoder were more precise than the ones generated by either the WaveNet or the WaveRNN




Fig. 8. The spectrograms of the speech generated by *HiNet* and *HiNet-1QWN-HC* when using natural acoustic features as input for an example sentence (arctic_b0536) in the test set of speaker *slt*.

TABLE VII

AVERAGE PREFERENCE SCORES (%) ON SPEECH QUALITY BETWEEN *HiNet* AND *HiNet-1QWN-HC* OF SPEAKER *slt*, WHERE N/P STANDS FOR “NO PREFERENCE” AND p DENOTES THE p -VALUE OF A t -TEST BETWEEN TWO SYSTEMS. “R” STANDS FOR USING NATURAL ACOUSTIC FEATURES AS INPUT AND “P” STANDS FOR USING PREDICTED ACOUSTIC FEATURES AS INPUT.

	<i>HiNet</i>	<i>HiNet-1QWN-HC</i>	N/P	p
R	29.67	30.66	39.67	0.7528
P	27.42	24.19	48.39	0.2639

vocoder and the ones recovered by the Griffin-Lim algorithm. Besides, it can be observed from Table III and IV that *HiNet* obtained much lower F0-RMSE and V/UV error than *HiNet-ASP+WaveNet*, *HiNet-ASP+WaveRNN* and *HiNet-ASP+GL*.

E. Model Simplification of PSP

As mentioned in Section III-B, the PSP model consumed most of the computation at the generation stage of the *HiNet* vocoder. Therefore, we explored whether a reduced-scale neural waveform generator is enough for predicting phase spectra in order to further decrease the computation complexity of the *HiNet* vocoder. Here, only speaker *slt* was used for experiments. Six vocoders with simplified structures were built for comparison and their descriptions are as follows.

- *HiNet-iQWN* A *HiNet* vocoder built by reducing the number of QWN blocks in PSP from 5 to i ($i = 4, 3, 2, 1$).
- *HiNet-HC* A *HiNet* vocoder built by halving the numbers of the gate channels, residual channels and skip channels mentioned in Section III-A.
- *HiNet-1QWN-HC* A *HiNet* vocoder built by reducing the number of QWN blocks in PSP from 5 to 1, and halving the numbers of the gate channels, residual channels and skip channels.

Table V listed the RTFs of the original *HiNet* vocoder and the *HiNet* vocoders after model simplification. By gradually simplifying the PSP model, the RTF also decreased from 0.34 of *HiNet* to 0.19 of *HiNet-1QWN-HC*. The objective performance evaluation results of all these vocoders are listed in Table VI. We found that there were no significant degradations on all metrics after simplifying the structure of the neural waveform generator in PSP. Fig. 8 shows the spectrograms of the speech generated by *HiNet* and *HiNet-1QWN-HC* for an example sentence. We can see that they were very similar and had no obvious differences.

TABLE VIII

SNRS AMONG THE ORIGINAL *HiNet* VOCODER, THE *HiNet* VOCODER WITHOUT PRE-CALCULATED INITIAL PHASE AND THE *HiNet* VOCODERS TRAINED WITH DIFFERENT LOSS FUNCTIONS ON THE TEST SETS OF SPEAKER *slt* WHEN USING NATURAL ACOUSTIC FEATURES AS INPUT.

	<i>HiNet</i>	<i>HiNet-woPCIP</i>	<i>HiNet-L1</i>	<i>HiNet-L2</i>	<i>HiNet-L3</i>
SNR(dB)	6.2937	1.8579	6.1380	6.2899	6.1377

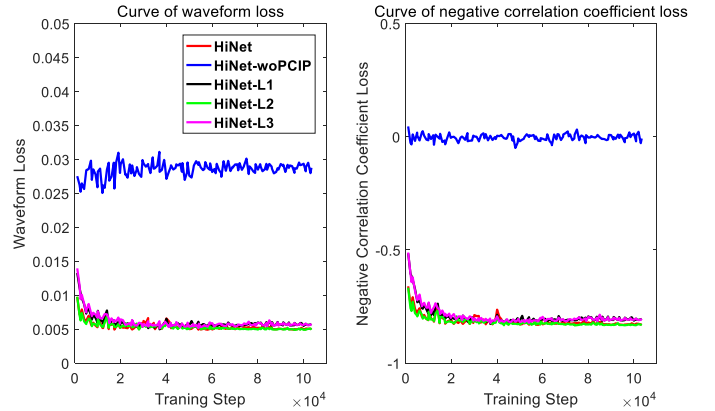


Fig. 9. The waveform loss and the negative correlation coefficient loss of PSPs in different vocoders on the validation set of speaker *slt*, where the x-axis shows training steps.

In order to examine whether there were significant subjective differences between the waveforms generated by *HiNet* and *HiNet-1QWN-HC*, two groups of ABX preference tests were conducted using natural and predicted acoustic features as input respectively. In each subjective test, twenty sentences randomly selected from the test set were synthesised by two comparative vocoders. Each pair of generated speech were evaluated by at least 30 English native listeners on the crowdsourcing platform of Amazon Mechanical Turk in random order. The listeners were asked to judge which utterance in each pair had better speech quality or there was no preference. In addition to calculating the average preference scores, the p -value of a t -test was used to measure the significance of the difference between two vocoders. The results are listed in Table VII. We can see that there was no significant difference ($p > 0.05$) between the subjective quality of *HiNet* and *HiNet-1QWN-HC*, no matter using natural or predicted acoustic features as input. These results indicated that the performance of *HiNet* was insensitive to the complexity of the NSF-based waveform generator in PSP to some extent. A neural waveform generator with much smaller scale than the ones for direct waveform generation may be enough for phase recovery.

F. Discussions

1) *Effects of Pre-Calculated Initial Phase*: As introduced in Section II-B1, we pre-calculated the initial phase ϕ_j for the sine-based excitation signal of each voiced segment at the training stage of PSP, expecting to benefit the recovery of phase spectra. To confirm the effectiveness of the pre-calculated initial phase, the *HiNet-woPCIP* vocoder was built for comparison. This vocoder adopted random initial phase

ϕ_j for the sine-based excitation signal of all voiced segments at the training stage of PSP. The speaker *slt* was used for experiments.

Here we focused on the SNR metric which reflected the performance of phase prediction and the results are listed in Table VIII. It is obvious that **HiNet-woPCIP** achieved much lower waveform SNR than **HiNet**. Fig. 9 draws the curves of the waveform loss and the negative correlation coefficient loss of PSPs on the validation set as a function of training steps. In our implementation, a training step generated a truncated sequence with 16000 samples. An epoch contained 2462 training steps for speaker *slt*. A validation was performed every 1000 training steps and at the end of an epoch during the training process. We can see from Fig. 9 that the waveform loss and the negative correlation coefficient loss of **HiNet** gradually decreased and both converged eventually, which implied that the PSP in the original HiNet vocoder gradually learnt the phase information during model training. However, the waveform loss of **HiNet-woPCIP** was almost unchanged and its negative correlation coefficient loss remained close to zero (i.e., no correlation), which indicated that discarding the pre-calculated initial phase prevented the PSP from learning the phase information through the waveform loss and the negative correlation coefficient loss. Therefore, the initial phase pre-calculation was crucial in our proposed method.

2) *Effects of Loss Functions*: As introduced in Section II-B3, a combination of amplitude spectrum loss, waveform loss and negative correlation coefficient loss was used to train the waveform generator in PSP. In this subsection, we explored the effects of the components in the combined loss function by ablation tests. Three vocoders with different loss functions for PSP were built and their descriptions were as follows.

- **HiNet-L1** The HiNet vocoder removing the negative correlation coefficient loss from the combined loss function for PSP (i.e., $\mathcal{L}_{PSP} = \sum_i \mathcal{L}_{ASi} + \mathcal{L}_W$).
- **HiNet-L2** The HiNet vocoder removing the waveform loss from the combined loss function for PSP (i.e., $\mathcal{L}_{PSP} = \sum_i \mathcal{L}_{ASi} + \mathcal{L}_C$).
- **HiNet-L3** The HiNet vocoder removing the waveform loss and the negative correlation coefficient loss from the combined loss function for PSP (i.e., $\mathcal{L}_{PSP} = \sum_i \mathcal{L}_{ASi}$).

Similar with Section III-F1, only the SNR results on speaker *slt* are listed in Table VIII. Comparing **HiNet** with **HiNet-L1**, it can be observed that removing the negative correlation coefficient loss led to the degradation of waveform SNR. In contrast, removing the waveform loss did not cause a significant degradation on waveform SNR. The curves of the waveform loss and the negative correlation coefficient loss on the validation set are also drawn in Fig. 9. We can see that there were no significant differences between **HiNet** and **HiNet-L2**. However, the converged losses of the HiNet vocoders trained without the negative correlation coefficient loss (i.e., **HiNet-L1** and **HiNet-L3**) were slightly higher than that of the other two HiNet vocoders (i.e., **HiNet** and **HiNet-L2**). In summary, the negative correlation coefficient loss played a more important role than the waveform loss for training PSPs in our experiments.

IV. CONCLUSION

In this paper, we have proposed a novel neural vocoder named HiNet which adopts hierarchical generation of amplitude and phase spectra for statistical parametric speech synthesis. The HiNet vocoder consists of an amplitude spectrum predictor (ASP) and a phase spectrum predictor (PSP). The former employs a DNN model to generate the amplitude spectra and the latter utilizes a neural source-filter (NSF) waveform generator to predict the phase spectra given amplitude spectra. The experimental results show that our proposed HiNet vocoder outperformed the conventional STRAIGHT vocoder and a 16-bit WaveNet vocoder using open source implementation, and achieved similar performance with a 16-bit WaveRNN vocoder, no matter using natural or predicted acoustic features as input. Because there are no autoregressive structures in both ASP and PSP, our proposed HiNet vocoder can reconstruct speech waveforms very efficiently. Through model simplification, the proposed HiNet vocoder can generate 1s waveforms of 16kHz speech in about 0.19s. Further improving the performance of ASP and PSP by using generative adversarial networks (GANs) [43] and applying the HiNet vocoder to other tasks such as voice conversion will be our future work.

REFERENCES

- [1] K. Tokuda, Y. Nankaku, T. Toda, H. Zen, J. Yamagishi, and K. Oura, "Speech synthesis based on hidden Markov models," *Proceedings of the IEEE*, vol. 101, no. 5, pp. 1234–1252, 2013.
- [2] H. Zen, A. Senior, and M. Schuster, "Statistical parametric speech synthesis using deep neural networks," in *Proc. ICASSP*, 2013, pp. 7962–7966.
- [3] Y. Fan, Y. Qian, F.-L. Xie, and F. K. Soong, "TTS synthesis with bidirectional LSTM based recurrent neural networks," in *Proc. Interspeech*, 2014, pp. 1964–1968.
- [4] Z.-H. Ling, S.-Y. Kang, H. Zen, A. Senior, M. Schuster, X.-J. Qian, H. M. Meng, and L. Deng, "Deep learning for acoustic modeling in parametric speech generation: A systematic review of existing techniques and future trends," *IEEE Signal Processing Magazine*, vol. 32, no. 3, pp. 35–52, 2015.
- [5] H. Dudley, "The vocoder," *Bell Labs Record*, vol. 18, no. 4, pp. 122–126, 1939.
- [6] H. Kawahara, I. Masuda-Katsuse, and A. De Cheveigne, "Restructuring speech representations using a pitch-adaptive time–frequency smoothing and an instantaneous-frequency-based f0 extraction: Possible role of a repetitive structure in sounds," *Speech communication*, vol. 27, no. 3, pp. 187–207, 1999.
- [7] M. Morise, F. Yokomori, and K. Ozawa, "WORLD: A vocoder-based high-quality speech synthesis system for real-time applications," *IEICE Transactions on Information and Systems*, vol. 99, no. 7, pp. 1877–1884, 2016.
- [8] F. Gunnar, "The acoustic theory of speech production," *SGravenhage, Mouton*, 1960.
- [9] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "WaveNet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [10] S. Mehri, K. Kumar, I. Gulrajani, R. Kumar, S. Jain, J. Sotelo, A. Courville, and Y. Bengio, "SampleRNN: An unconditional end-to-end neural audio generation model," *arXiv preprint arXiv:1612.07837*, 2016.
- [11] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. v. d. Oord, S. Dieleman, and K. Kavukcuoglu, "Efficient neural audio synthesis," *arXiv preprint arXiv:1802.08435*, 2018.
- [12] I. Recommendation, "G. 711: Pulse code modulation (PCM) of voice frequencies," *International Telecommunication Union*, 1988.

- [13] A. v. d. Oord, Y. Li, I. Babuschkin, K. Simonyan, O. Vinyals, K. Kavukcuoglu, G. v. d. Driessche, E. Lockhart, L. C. Cobo, F. Stimberg *et al.*, “Parallel WaveNet: Fast high-fidelity speech synthesis,” *arXiv preprint arXiv:1711.10433*, 2017.
- [14] W. Ping, K. Peng, and J. Chen, “ClariNet: Parallel wave generation in end-to-end text-to-speech,” *arXiv preprint arXiv:1807.07281*, 2018.
- [15] R. Prenger, R. Valle, and B. Catanzaro, “WaveGlow: A flow-based generative network for speech synthesis,” *arXiv preprint arXiv:1811.00002*, 2018.
- [16] A. Tamamori, T. Hayashi, K. Kobayashi, K. Takeda, and T. Toda, “Speaker-dependent WaveNet vocoder,” in *Proc. Interspeech*, 2017, pp. 1118–1122.
- [17] T. Hayashi, A. Tamamori, K. Kobayashi, K. Takeda, and T. Toda, “An investigation of multi-speaker training for WaveNet vocoder,” in *Proc. ASRU*, 2017, pp. 712–718.
- [18] N. Adiga, V. Tsiasaras, and Y. Stylianou, “On the use of WaveNet as a statistical vocoder,” in *Proc. ICASSP*, 2018, pp. 5674–5678.
- [19] Y. Ai, H.-C. Wu, and Z.-H. Ling, “SampleRNN-based neural vocoder for statistical parametric speech synthesis,” in *Proc. ICASSP*, 2018, pp. 5659–5663.
- [20] Y. Ai, J.-X. Zhang, L. Chen, and Z.-H. Ling, “DNN-based spectral enhancement for neural waveform generators with low-bit quantization,” in *Proc. ICASSP*, 2019, pp. 7025–7029.
- [21] J. Lorenzo-Trueba, T. Drugman, J. Latorre, T. Merritt, B. Putrycz, and R. Barra-Chicote, “Robust universal neural vocoding,” *arXiv preprint arXiv:1811.06292*, 2018.
- [22] L.-J. Liu, Z.-H. Ling, Y. Jiang, M. Zhou, and L.-R. Dai, “WaveNet vocoder with limited training data for voice conversion,” in *Proc. Interspeech*, 2018, pp. 1983–1987.
- [23] K. Kobayashi, T. Hayashi, A. Tamamori, and T. Toda, “Statistical voice conversion with WaveNet-based waveform generation,” in *Proc. Interspeech*, 2017, pp. 1138–1142.
- [24] Z.-H. Ling, Y. Ai, Y. Gu, and L.-R. Dai, “Waveform modeling and generation using hierarchical recurrent neural networks for speech bandwidth extension,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 5, pp. 883–894, 2018.
- [25] J. Klejsa, P. Hedelin, C. Zhou, R. Fejgin, and L. Villemoes, “High-quality speech coding with SampleRNN,” *arXiv preprint arXiv:1811.03021*, 2018.
- [26] Y. Cui, X. Wang, L. He, and F. K. Soong, “A new glottal neural vocoder for speech synthesis,” in *Proc. Interspeech*, 2018, pp. 2017–2021.
- [27] L. Juvela, V. Tsiasaras, B. Bollepalli, M. Airaksinen, J. Yamagishi, and P. Alku, “Speaker-independent raw waveform model for glottal excitation,” *arXiv preprint arXiv:1804.09593*, 2018.
- [28] L. Juvela, B. Bollepalli, V. Tsiasaras, and P. Alku, “GlottNet-A raw waveform model for the glottal excitation in statistical parametric speech synthesis,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2019.
- [29] M.-J. Hwang, F. Soong, F. Xie, X. Wang, and H.-G. Kang, “LP-WaveNet: Linear prediction-based WaveNet speech synthesis,” *arXiv preprint arXiv:1811.11913*, 2018.
- [30] J.-M. Valin and J. Skoglund, “LPCNet: Improving neural speech synthesis through linear prediction,” in *Proc. ICASSP*, 2019, pp. 5891–5895.
- [31] Z. Jin, A. Finkelstein, G. J. Mysore, and J. Lu, “FFNet: A real-time speaker-dependent neural vocoder,” in *Proc. ICASSP*, 2018, pp. 2251–2255.
- [32] X. Wang, S. Takaki, and J. Yamagishi, “Neural source-filter-based waveform model for statistical parametric speech synthesis,” in *Proc. ICASSP*, 2019, pp. 5916–5920.
- [33] J. Kominek and A. W. Black, “The CMU Arctic speech databases,” in *Proc. ISCA Workshop*, 2004.
- [34] Z. Wu, O. Watts, and S. King, “Merlin: An open source neural network speech synthesis system,” in *Proc. ISCA Speech Synthesis Workshop*, 2016, pp. 202–207.
- [35] K. Tokuda, T. Yoshimura, T. Masuko, T. Kobayashi, and T. Kitamura, “Speech parameter generation algorithms for HMM-based speech synthesis,” in *Proc. ICASSP*, vol. 3, 2000, pp. 1315–1318.
- [36] T. Toda and K. Tokuda, “A speech parameter generation algorithm considering global variance for HMM-based speech synthesis,” *IEICE Transactions on Information and Systems*, vol. 90, no. 5, pp. 816–824, 2007.
- [37] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [38] A. Paszke, S. Gross, S. Chintala, and G. Chanan, “Pytorch: Tensors and dynamic neural networks in python with strong gpu acceleration,” 2017.
- [39] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin *et al.*, “Tensorflow: Large-scale machine learning on heterogeneous distributed systems,” *arXiv preprint arXiv:1603.04467*, 2016.
- [40] I. Recommendation, “Method for the subjective assessment of intermediate sound quality (MUSHRA),” *ITU, BS*, pp. 1543–1, 2001.
- [41] S. Buchholz and J. Latorre, “Crowdsourcing preference tests, and how to detect cheating,” in *Proc. Interspeech*, 2011, pp. 3053–3056.
- [42] D. Griffin and J. Lim, “Signal estimation from modified short-time fourier transform,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, no. 2, pp. 236–243, 1984.
- [43] 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.