

Speech Synthesis Based on Gaussian Conditional Random Fields

Soheil Khorram^(✉), Fahimeh Bahmaninezhad, and Hossein Sameti

Department of Computer Engineering,
Sharif University of Technology, Tehran, Iran
{khorram, bahmaninezhad}@ce.sharif.edu,
sameti@sharif.edu

Abstract. Hidden Markov Model (HMM)-based synthesis (HTS) has recently been confirmed to be the most effective method in generating natural speech. However, it lacks adequate context generalization when the training data is limited. As a solution, current study provides a new context-dependent speech modeling framework based on the Gaussian Conditional Random Field (GCRF) theory. By applying this model, an innovative speech synthesis system has been developed which can be viewed as an extension of Context-Dependent Hidden Semi Markov Model (CD-HSMM). A novel Viterbi decoder along with a stochastic gradient ascent algorithm was applied to train model parameters. Also, a fast and efficient parameter generation algorithm was derived for the synthesis part. Experimental results using objective and subjective criteria have shown that the proposed system outperforms HSMM substantially in limited speech databases. Moreover, Mel-cepstral distance of the spectral parameters has been reduced considerably for any size of training database.

Keywords: Gaussian conditional random field · Statistical parametric speech synthesis · HSMM extension

1 Introduction

Statistical Parametric Speech Synthesis (SPSS) has reportedly been a dominant research area due to its peculiarities since the last decade [1, 2]. Modeling in the domain of SPSS is of prime importance and it is naïve to assume unnecessary simplifying assumptions in modeling as it may reduce the quality of synthetic speech. This work extends Hidden Semi Markov Model (HSMM) synthesis [3] by eliminating some of its simplifying assumptions. In the next subsection we will briefly discuss related works.

1.1 Related Work

Many research activities have already been performed to improve the quality of basic HTS. The progresses such as Hidden Semi Markov Model (HSMM) [3], Trajectory HMM [4] and Multi-Space Distribution HMM [5] have made HTS the most powerful statistical approach. However, these systems do not lead to an acceptable quality with

limited databases (less than 30 min). This deficiency is a direct result of applying decision-tree-based context clustering which cannot exploit contextual information efficiently, because each training sample is associated in modeling only one context cluster. This study is an attempt to improve SPSS quality even for limited training data.

The rest of the paper is organized as follows. In Sect. 2, GCRF is introduced. Sections 3 and 4 propose a context-dependent model for speech using GCRF and its application in speech synthesis. Experimental results are presented in Sect. 5 and final remarks are given in Sect. 6.

2 Gaussian Conditional Random Field

To define GCRF, first a brief description of Markov Random Field (MRF) and Conditional Random Field (CRF) is given.

Definition 1. Let $G = (V, E)$ be an undirected graph, $X = (X_v)_{v \in V}$ be a set of random variables indexed by nodes of G , X is modeled by MRF iff $\forall A, B \subseteq V, P(X_A | X_B) = P(X_A | X_S)$, where S is a border subset of A such that every path from a node in A to a node in B passes through S [6].

Definition 2. (X, C) is a CRF iff for any given set of random variables C , X forms an MRF [6].

In the speech synthesis framework, given an utterance contextual information C , sufficient statistics of speech (acoustic features) can be considered as an MRF.

Hammersley-Clifford's Theorem. Suppose (x, c) is an arbitrary realization of a CRF (X, C) defined based on a graph G with positive probability, then $P(x|c)$ can be factorized by the following Gibbs distribution [7].

$$P(x|c) = \frac{1}{Z(c)} \prod_{\mathcal{A}} \Psi_a(x, c), \quad (1)$$

where \mathcal{A} denotes a set of all maximal cliques of G . $Z(c)$ is called partition function which ensures that the distribution sums to one. In other words,

$$Z(c) = \iint_x \prod_{\mathcal{A}} \Psi_a(x, c). \quad (2)$$

The theorem also states that for any choice of positive local functions $\{\Psi_a(x)\}$ (potential functions) a valid CRF is generated. One of the simplest choices of a potential function is Gaussian function. CRF with Gaussian potential function is named GCRF which is introduced in the next section.

3 Context-Dependent Speech Modeling Using GCRF

For modeling speech, the proposed system primarily splits each segment into a fixed number of states. Then, acoustic and binary contextual features (sufficient statistics) are extracted for each state. The goal is to model and generate acoustic features provided that contextual features are present. The following notations are taken into account henceforth.

- L, I : Total number of acoustic and linguistic features.
- \mathcal{J} : Total number of states for the current utterance.
- V : All acoustic parameters. (Extracted from frame samples)
- x_{lj} : l -th acoustic feature of state j . (Extracted from V)
- x_l : l -th acoustic feature vector, $x_l \stackrel{\text{def}}{=} [x_{l1}, \dots, x_{l\mathcal{J}}]^T$.
- X : All acoustic features, $X \stackrel{\text{def}}{=} [x_1, \dots, x_L]$.
- c_{ji} : i -th binary linguistic feature of state j .
- c_j : Linguistic feature vector of state j , $c_j \stackrel{\text{def}}{=} [c_{j1}, \dots, c_{jI}]^T$.
- C : All linguistic features, $C \stackrel{\text{def}}{=} [c_1, \dots, c_{\mathcal{J}}]$.

3.1 GCRF Graphical Structure

Factor graph [8] of the proposed GCRF (with order one) is depicted in Fig. 1. As it is obvious in the figure, GCRF is a set of L linear chain CRF [8] (with order one) which are independent when C is given. Each rectangular node Ψ_{lj} represents a potential function describing the effect of a maximal clique $(x_{lj}, x_{l(j-1)}, c_j)$ in the random field distribution. This figure can be extended to higher order linear chain CRFs. As a result, if GCRF extends with order o , Ψ_{lj} becomes a function of $(x_{lj}, \dots, x_{l(j-o)}, c_j)$.

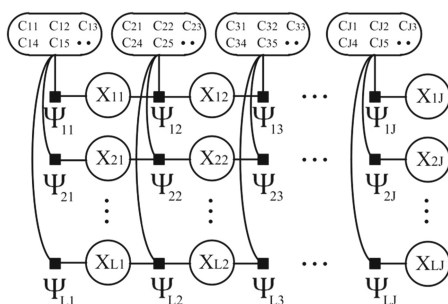


Fig. 1. Factor graph of the first order GCRF.

3.2 GCRF Distribution

Having described the graphical model, this subsection investigates the probability distribution provided by GCRF. Markov property of MRFs implies the following equality.

$$P(X|C; \theta) = \prod_{l=1}^L P(x_l|C; \theta), \tag{3}$$

where θ is the set of all model parameters. This paper assumes that the partition function, Ψ_{lj} , is formulated by Eq. 4 which is a Gaussian function with parameters H_{lji} and u_{lji} .

$$\Psi_{lj} \stackrel{\text{def}}{=} \exp \left\{ -\frac{1}{2} \sum_{i=1}^I \left[(x_l^T H_{lji} x_l + u_{lji}^T x_l) c_{ji} \right] \right\}. \tag{4}$$

In this equation, H_{lji} has to be a symmetric and positive definite matrix. If H_{lji} is not restricted to a positive definite matrix, the distribution may be realized by a number greater than one. Thus, considering positive definite condition seems to be necessary. Moreover, in GCRF with order o , H_{lji} and u_{lji} contain only $(o + 1) \times (o + 1)$ and $(o + 1)$ nonzero elements respectively. The overall form of model parameters is shown as follows.

$$H^{lij} = \begin{bmatrix} 0 & 0 & \dots & 0 & 0 & \dots \\ 0 & h_{(j-o)(j-o)}^{lij} & \dots & h_{(j-o)j}^{lij} & 0 & \dots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots \\ 0 & h_{j(j-o)}^{lij} & \dots & h_{jj}^{lij} & 0 & \dots \\ 0 & 0 & \dots & 0 & 0 & \dots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots \end{bmatrix}, u^{lij} = \begin{bmatrix} 0 \\ u_{j-o}^{lij} \\ \vdots \\ u_j^{lij} \\ 0 \\ \vdots \end{bmatrix}. \tag{5}$$

By considering defined potential function and according to the fundamental theorem of Hammersley and Clifford the final expression for $P(x_l|C; \theta_l)$ is given by

$$P(x_l|C; \theta_l) = \frac{1}{Z_l(C; \theta_l)} \exp \left\{ -\frac{1}{2} (x_l^T H_l x_l + u_l^T x_l) \right\}, \tag{6}$$

where $H_l = \sum_{j=1}^J \sum_{i=1}^I c_{ji} H_{lji}$ and $u_l = \sum_{j=1}^J \sum_{i=1}^I c_{ji} u_{lji}$.

Z_l is the partition function and is computed by Eq. 2. Fortunately, for Gaussian distribution of Eq. 4 there is a closed formula for the partition function as:

$$Z_l(C; \theta_l) = (2\pi)^{\frac{J}{2}} (\det(H_l^{-1}))^{\frac{1}{2}} \exp \left(\frac{1}{8} u_l^T H_l^{-1} u_l \right). \tag{7}$$

A marvelous point is that conventional CD-HSMM can be considered as a type of GCRF with order zero and mutually exclusive contextual features.

4 Speech Synthesis Based on GCRF

Figure 2 shows an overview of the proposed GCRF-based speech synthesis system. All blocks in the figure are identical to classical SPSS [1], except the three further blocks added with a different color. In the training part, acoustic sufficient statistics or features (X) are extracted according to both speech parameters (V) and state boundaries (\mathcal{T}). State boundaries are latent and the added Viterbi block is employed to train them in an unsupervised manner. It should be noted that only sufficient statistics are modeled in the training phase; therefore synthesis phase has to generate them first. After generating features, speech parameters and speech signal are successively synthesized.

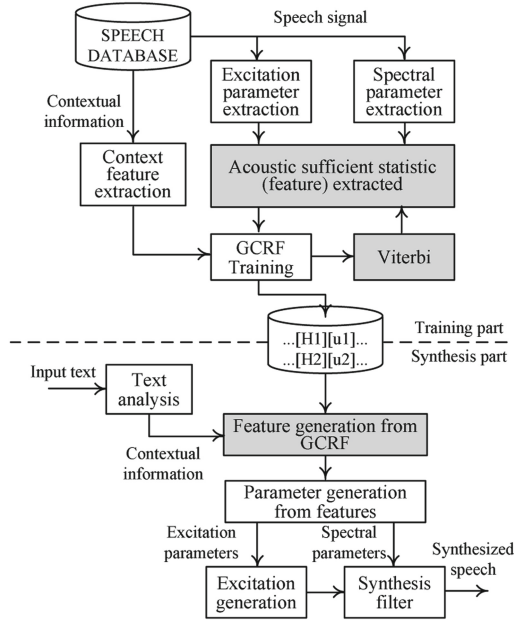


Fig. 2. An overview of the proposed architecture.

4.1 Estimation of Model Parameters

In this section, we discuss how to train model parameters θ . We are given a set of T iid training data $\{X^t, C^t\}_{t=1}^T$, the goal is to find the best set of parameters, $\hat{\theta}$, which maximizes the conditional log likelihood:

$$\hat{\theta} = \operatorname{argmax}_{\theta} L(\theta), \quad (8)$$

$$L(\theta) \stackrel{\text{def}}{=} \frac{1}{T} \sum_{t=1}^T \log P(X^t | C^t; \theta). \quad (9)$$

The problem is that, acoustic feature Matrix X^t , wholly depends on the state boundaries which are latent. Hence, it is impossible to compute $L(\theta)$. A correct solution for this problem that converges to the Maximum Likelihood (ML)-estimate is given by the Expectation Maximization (EM) algorithm; however, EM is computationally expensive. Another commonly used method which is computationally efficient and works well in practice is to compute first X^t and then $L(\theta)$ on the Viterbi path. Applying this approach and substituting $P(X^t|C^t; \theta)$ with Eq. 6 gives

$$L(\theta) = -\frac{1}{2T} \sum_{t=1}^T \sum_{l=1}^L \{L_l^t(\theta_l)\}, \quad (10)$$

$$L_l^t(\theta_l) \stackrel{\text{def}}{=} x_l^{tT} H_l^t x_l^t + u_l^{tT} x_l^t + \mathcal{J} \log 2\pi - \log \det H_l^t + \frac{1}{4} u_l^{tT} H_l^{t-1} u_l^t. \quad (11)$$

In general, this function cannot be maximized in closed form, therefore numerical optimization is used. The partial derivatives of $L(\theta)$ are calculated as follows.

$$\frac{\partial L(\theta)}{\partial u_{lij}} = -\frac{1}{2T} \sum_{t=1}^T \frac{\partial L_l^t(\theta_l)}{\partial u_{lij}}, \quad (12)$$

$$\frac{\partial L_l^t(\theta)}{\partial u_{lij}} = \left[\left(x_l^t + \frac{1}{2} H_l^{t-1} u_l^t \right) c_{ji}^t \right] \star \mathbb{b}(\mathcal{J}, j, o), \quad (13)$$

$$\frac{\partial L(\theta)}{\partial H_{lij}} = -\frac{1}{2T} \sum_{t=1}^T \frac{\partial L_l^t(\theta_l)}{\partial H_{lij}}, \quad (14)$$

$$\frac{\partial L_l^t(\theta_l)}{\partial H_{lij}} = \left[\left(x_l^t x_l^{tT} - H_l^{t-1} - \frac{1}{4} H_l^{t-1} u_l^t u_l^{tT} H_l^{t-1} \right) c_{ji}^t \right] \star \mathbb{B}(\mathcal{J}, j, o). \quad (15)$$

where o denotes the order of model, \star denotes element-by-element product operator and \mathbb{B} (\mathbb{b}) is a \mathcal{J} -by- \mathcal{J} (\mathcal{J}) Boolean matrix (vector) defined by an indicator function I as:

$$\mathbb{b}(\mathcal{J}, j, o) \stackrel{\text{def}}{=} [\mathbb{b}_m(\mathcal{J}, j, o)]_{\mathcal{J} \times 1}, \quad (16)$$

$$\mathbb{b}_m(\mathcal{J}, j, o) \stackrel{\text{def}}{=} I(j - o \leq m \leq j),$$

$$\mathbb{B}(\mathcal{J}, j, o) \stackrel{\text{def}}{=} [\mathbb{B}_{mn}(\mathcal{J}, j, o)]_{\mathcal{J} \times 1}, \quad (17)$$

$$\mathbb{B}_{mn}(\mathcal{J}, j, o) \stackrel{\text{def}}{=} I((j - o \leq m \leq j) \& (j - o \leq n \leq j)).$$

A common solution of this optimization problem is to take entire training samples into account and update model parameters using an optimization algorithm such as *BFGS*. Unfortunately, this in turn leads to large computational complexity. This paper proposes the application of *stochastic gradient ascent* [9] method which is faster than

above-mentioned algorithm by orders of magnitude. This method has proven to be effective [9]. Following equations express its updating rule:

$$\mathbf{u}_{ij}^t = \mathbf{u}_{ij}^{t-1} - \alpha^t \frac{\partial \mathcal{L}_l^k(\theta_l)}{\partial \mathbf{u}_{ij}} \Big|_{\mathbf{u}_{ij}^t, \mathbf{H}_{ij}^t}, \quad (18)$$

$$\mathbf{H}_{ij}^t = \mathbf{H}_{ij}^{t-1} - \alpha^t \frac{\partial \mathcal{L}_l^l(\theta_l)}{\partial \mathbf{H}_{ij}} \Big|_{\mathbf{u}_{ij}^t, \mathbf{H}_{ij}^t}. \quad (19)$$

A variable step size algorithm described by [10] is utilized in our experiments.

4.2 Viterbi Algorithm for GCRF

Given a sequence of acoustic parameters (V), sentence contextual features (C) and a trained GCRF parameters (θ), this section presents an algorithm to find the most likely state boundaries ($\hat{\mathcal{T}}$). Thus the aim is to estimate $\hat{\mathcal{T}}$ such that

$$\hat{\mathcal{T}} \stackrel{\text{def}}{=} \operatorname{argmax}_{\mathcal{T}} P(\mathcal{T}|V, C; \theta) = \operatorname{argmax}_{\mathcal{T}} P(X(\mathcal{T}, V)|V, C; \theta). \quad (20)$$

From Eq. 6 we have

$$\hat{\mathcal{T}} = \operatorname{argmin}_{\mathcal{T}} \sum_{j=1}^{\mathcal{J}} \phi_j(\mathcal{T}, V, C, \theta), \quad (21)$$

where $\phi_j(\mathcal{T}, V, C, \theta) \stackrel{\text{def}}{=} \sum_{i=1}^L \sum_{i=1}^I \left(\mathbf{x}_{ij}^T \mathbf{H}_{ij} \mathbf{x}_{ij} + \mathbf{b}_{ij}^T \mathbf{x}_{ij} \right) c_{ji}$.

Let \mathcal{t}_j be the j -th state boundary (j -th element of \mathcal{T}), then for a GCRF with order o , ϕ_j becomes a function of $\mathcal{t}_{j-o-1}, \dots, \mathcal{t}_j$ instead of entire elements of \mathcal{T} . This fact gives us an ability to exploit dynamic programming for performing a complete search on \mathcal{T} . Inspired by the other Viterbi algorithms, we need to define an auxiliary variable δ_j .

$$\delta_j(\mathcal{t}_{j-o}, \dots, \mathcal{t}_j) \stackrel{\text{def}}{=} \min_{\mathcal{t}_1, \dots, \mathcal{t}_{j-o-1}} \sum_{j=1}^j \Phi_j(\mathcal{t}_{j-o-1}, \dots, \mathcal{t}_j). \quad (22)$$

δ_j can be calculated from δ_{j-1} by following recursion.

$$\delta_{j+1}(\mathcal{t}_{j-o+1}, \dots, \mathcal{t}_{j+1}) = \min_{\mathcal{t}_{j-o}} [\delta_j(\mathcal{t}_{j-o}, \dots, \mathcal{t}_j) + \Phi_{j+1}(\mathcal{t}_{j-o}, \dots, \mathcal{t}_{j+1})]. \quad (23)$$

Using this recursion, it is straightforward to obtain Viterbi algorithm.

4.3 Parameter Generation Algorithm

This section, for a given GCRF, derives an algorithm to estimate the best synthesized speech parameters (\hat{V}) by maximizing the likelihood criteria, i.e.

$$\hat{V} \stackrel{\text{def}}{=} \operatorname{argmax}_V P(V|\theta) = \operatorname{argmax}_V \sum_{\mathcal{T}} P(X(V, \mathcal{T})|\theta). \quad (24)$$

The synthesis part needs to respond quickly, however, solving this problem directly is challenging. Hence, the algorithm derived from Eq. 24 is not practical.

A two-step algorithm is proposed here which approximates \hat{V} fast.

Step 1. For a given θ , compute the ML-estimate of X :

$$\hat{X} \stackrel{\text{def}}{=} \operatorname{argmax}_X P(X|\theta). \quad (25)$$

Step 2. For a given X , compute the ML-estimate of V :

$$\hat{V} \stackrel{\text{def}}{=} \operatorname{argmax}_V P(V|X). \quad (26)$$

The first step is simply obtained by considering the distribution discussed in Sect. 3. Since different acoustic features are statistically independent (given in Eq. 3), the algorithm can generate features independently, i.e.

$$\hat{x}_l = \operatorname{argmax}_{x_l} P(x_l|C; \theta_l). \quad (27)$$

Optimizing the Gaussian distribution $P(x_l|C; \theta_l)$, expressed by Eq. 6, results in the set of linear equations below:

$$H^l \hat{x}_l = -\frac{1}{2} b^l. \quad (28)$$

H^l is symmetric and positive definite, so Eq. 28 can be efficiently solved using the Cholesky decomposition.

Second step depends heavily on the selected acoustic features. For the set of acoustic features extracted in our system, Tokoda et al. [11] algorithm was used in this step.

5 Experiments

5.1 Experimental Conditions

To evaluate the proposed system, a Persian speech database [12] consisting of 1000 utterances with an average length of 8 s was employed. Experiments were conducted on a fixed test set of 200 utterances and 5 different training sets with remaining 50, 100, 200, 400 and 800 utterances. It should be noted that the average length of each utterance is about 8 s. Speech parameters including mel-cepstral coefficients, bandpass aperiodicity and fundamental frequency were extracted by STRAIGHT [13]. Sample mean and variance of each static and dynamic parameter, in addition to the voicing probability and duration are computed as the acoustic state features. For contextual

state features a set of 150 well designed binary questions are employed. Following subsections evaluate the proposed method in contrast to the HSMM-based technique.

5.2 Objective Evaluation

As Fig. 3 shows, three objective measures were calculated to evaluate the proposed and HSMM-based systems, namely the average mel-cepstral distortion (expressed in dB) [14], the Root-Mean-Square (RMS) error of fundamental frequency logarithm (expressed in cent) and the RMS error of phoneme durations (expressed in terms of number of frames). Computing the first and second measures needs an assumption about state boundaries that was estimated here using the Viterbi algorithm. Since F0 value is not observed in unvoiced regions, only voiced frames of speech were taken into account for the second measure.

From Fig. 3, it is noticeable that GCRF always outperforms HSMM in generating mel-cepstral and duration parameters, but HSMM is superior in synthesizing fundamental frequency when the number of training data is larger than 200 utterances. This drawback is a result of weak estimation of F0 parameters during the training process. Table 1 compares the accuracy of voiced/unvoiced detection in proposed system with its counterpart in HSMM-based synthesis.




Fig. 3. Objective evaluation of HSMM-based and proposed speech synthesis systems. (Left) Mel-cepstral distance [dB]; (Middle) RMSE of log F0 [cent]; (Right) RMSE of phoneme duration [frame].

5.3 Subjective Evaluation

We conducted preference score measure to compare the proposed and HSMM-based systems subjectively. 20 subjects were presented with 10 randomly chosen pairs of synthesized speech from the two models and then asked for their preference.

Figure 4 shows the average preference score. The result confirms that the synthetic speech generated by proposed system has been favorable when training data are limited.




Fig. 4. Subjective evaluation of HSM and proposed systems using preference score.

Table 1. Accuracy of Voiced/Unvoiced Detector.

| # train data | Proposed accuracy | HSM accuracy |
|--------------|-------------------|--------------|
| 50 | 0.9184 | 0.8851 |
| 100 | 0.9241 | 0.8828 |
| 200 | 0.9157 | 0.8903 |
| 400 | 0.9104 | 0.8783 |
| 800 | 0.9037 | 0.8809 |

6 Conclusion

This paper improves HSM-based synthesis in the following ways:

1. The independence assumption of states distribution in HTS is removed.
2. In contrast to HMM, the proposed model does not limit its potential functions to be a probability distribution.
3. CD-HMM uses decision-tree-based context clustering that does not provide efficient generalization in limited training data, because each speech parameter vector is associated in modeling of only one context cluster. In contrast, our method contributes each training vector in many clusters to offer an efficient generalization.

Despite the advantages, which made our system to outperform in small training data, a drawback such as difficult training procedure is noticed in large databases.

References

1. Black, A.W., Zen, H., Tokuda, K.: Statistical Parametric Speech Synthesis. In: ICASSP'2007, Honolulu, Hawai'i, USA, pp. IV-1229–IV-1232 (2007)
2. Zen, H., Tokuda, K., Black, A.W.: Statistical parametric speech synthesis. *Speech Commun.* Elsevier **51**(11), 1039–1064 (2009)
3. Zen, H., Tokuda, K., Masuko, T., Kobayashi, T., Kitamura, T.: Hidden semi-markov model based speech synthesis. In: Interspeech'2004, Jeju Island, Korea, pp. 1393–1396, October 4–8 2004
4. Zen, H., Tokuda, K., Kitamura, T.: An introduction of trajectory model into hmm-based speech synthesis. In: SSW5, pp. 191–196. Carnegie Mellon University, June 2004
5. Tokuda, K., Masuko, T., Miyazaki, N., Kobayashi, T.: Multi-space probability distribution HMM. *IEICE Trans. Inf. Syst.* **E85-D**(3), 455–464 (2002)
6. Lafferty, J., McCallum, A., Pereira, F.: Conditional random fields: probabilistic models for segmenting and labeling sequence data. In: Proceedings of the Eighteenth International Conference on Machine Learning, pp. 282–289 (2001)
7. Grimmett, G.R.: A theorem about random fields. *Bull. Lond. Math. Soc.* **5**, 81–84 (1973)

8. Sutton, C., McCallum, A.: An introduction to conditional random fields for relational learning. In: Getoor, L., Taskar, B. (eds.) *Introduction to statistical Relational Learning*. MIT Press, Cambridge (2006)
9. Gardner, W.A.: Learning characteristics of stochastic-gradient-descent algorithms: a general study, analysis and critique. *Sig. Process.* **6**(2), 113–133 (1984)
10. Vrahatis, M.N., Androulakis, G.S., Lambrinos, J.N., Magoulas, G.D.: A class of gradient unconstrained minimization algorithms with adaptive stepsize. *J. Comput. Appl. Math.* **114**(2), 367–386 (2000)
11. Tokuda, K., Yoshimura, T., Masuko, T., Kobayashi, T., Kitamura, T.: Speech parameter generation algorithms for HMM-based speech synthesis. In: *ICASSP'2000*, vol. 3, Istanbul, pp. 1315–1318, June 2000
12. Bijankhan, M., Sheikhzadegan, J., Roohani, M.R., Samareh, Y., Lucas, C., Tebiani, M.: The speech database of farsi spoken language. In: *Proceedings of 5th Australian International Conference on Speech Science and Technology (SST'94)*, pp. 826–831 (1994)
13. Kawahara, H., Masuda-Katsuse, I., de Cheveigné, A.: 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 Commun.* **27**(3–4), 187–207 (1999)
14. Kubichek, R.F.: Mel-cepstral distance measure for objective speech quality assessment. In: *Proceedings of the IEEE Pacific Rim Conference on Communications, Computers, and Signal Processing*, pp. 125–128 (1993)