

Analysis of Head Gesture and Prosody Patterns for Prosody-Driven Head-Gesture Animation

Mehmet Emre Sargin, *Student Member, IEEE*, Yücel Yemez, *Member, IEEE*,
Engin Erzin, *Senior Member, IEEE*, and A. Murat Tekalp, *Fellow, IEEE*

Abstract—We propose a new two-stage framework for joint analysis of head gesture and speech prosody patterns of a speaker toward automatic realistic synthesis of head gestures from speech prosody. In the first stage analysis, we perform Hidden Markov Model (HMM)-based unsupervised temporal segmentation of head gesture and speech prosody features separately to determine elementary head gesture and speech prosody patterns, respectively, for a particular speaker. In the second stage, joint analysis of correlations between these elementary head gesture and prosody patterns is performed using Multistream HMMs to determine an audio-visual mapping model. The resulting audio-visual mapping model is then employed to synthesize natural head gestures from arbitrary input test speech given a head model for the speaker. In the synthesis stage, the audio-visual mapping model is used to predict a sequence of gesture patterns from the prosody pattern sequence computed for the input test speech. The Euler angles associated with each gesture pattern are then applied to animate the speaker head model. Objective and subjective evaluations indicate that the proposed synthesis by analysis scheme provides natural looking head gestures for the speaker with any input test speech, as well as in “prosody transplant” and “gesture transplant” scenarios.

Index Terms—Multimedia information systems, speech analysis, face and gesture recognition, pattern analysis and recognition, animation.

1 INTRODUCTION

STATE-OF-THE-ART visual speaker animation methods are capable of generating synchronized lip movements automatically from speech content; however, they lack automatic synthesis of speaker gestures from speech. Head and face gestures are usually added manually by artists, which is costly and often look unrealistic. Hence, learning the correlation between gesture and speech patterns of a speaker toward automatic realistic synthesis of speaker gestures from speech remains as a challenging research problem.

There exists significant literature on speaker lip animation, that is, rendering lip movements synchronized with the speech signal [1]. Since lip movement is physiologically tightly coupled with acoustic speech, it is relatively an easy task to find a mapping between the phonemes of speech and the visemes of lip movement. Many schemes exist to find such audio-to-visual mappings among which the Hidden Markov Model (HMM)-based techniques are the most common as they yield smooth animations exploiting temporal dynamics of speech [2], [3], [4], [5], [6], [7], [8], [9]. Some of these works also incorporate synthesis of facial expressions along with the lip movements to make animated faces look more natural [3], [6], [8], [9]. The common strategy

in these techniques is to train a joint HMM structure with extracted visual and audio feature vectors and then to use the trained HMM structure to generate speech-driven facial expressions and lip movements.

Despite exhibiting variations from person to person and in time, head and body gestures are also correlated with speech. For example, it has been observed that manual gestures are correlated with prosody [10], [11] and verbal content of the speech [12], whereas head gestures are mostly correlated with the prosody [11], [13], [14]. Although correlations between speech and head/body gestures have been investigated in several works, there are only a limited number of publications addressing speech-driven head and body gesture synthesis. In [15], audio streams from training videos are first segmented using pitch contour information. The same boundaries are also applied to the corresponding video streams for segmenting head motions. The co-occurring audio and head motion segments are stored as pairs in a database. Later, a new test audio stream is segmented, and an optimal head gesture sequence is determined from the database using dynamic programming to create synthetic head motions. A similar methodology is followed in [16], where audio/head motion feature pairs extracted from training videos are stored into a database indexed by audio features. Later, audio features extracted from a new test input speech are used to search for K-nearest neighbors. The optimum nearest neighbor combination, found by dynamic programming, is used to synthesize corresponding head motions. In [17], we presented a preliminary demonstration of natural looking head and arm gesture synthesis from speech using a manually determined audio-visual mapping from speech to head and arm motions.

The aim of this paper is to present a framework for joint analysis of head gesture and speech prosody patterns

• M.E. Sargin is with the Department of Electrical and Computer Engineering, University of California-Santa Barbara, Santa Barbara, CA 93106-9560. E-mail: msargin@ece.ucsb.edu.

• Y. Yemez, E. Erzin, and A.M. Tekalp are with the College of Engineering, Koc University, Sarıyer, İstanbul 34450, Turkey.
E-mail: {yyemez, eerzin, mtekalp}@ku.edu.tr.

Manuscript received 29 Aug. 2006; revised 19 June 2007; accepted 27 Sept. 2007; published online 8 Oct. 2007.

Recommended for acceptance by Y. Sato.

For information on obtaining reprints of this article, please send e-mail to: xxxx@computer.org, and reference IEEECS Log Number TPAMI-0634-0806. Digital Object Identifier no. 10.1109/TPAMI.2007.70797.

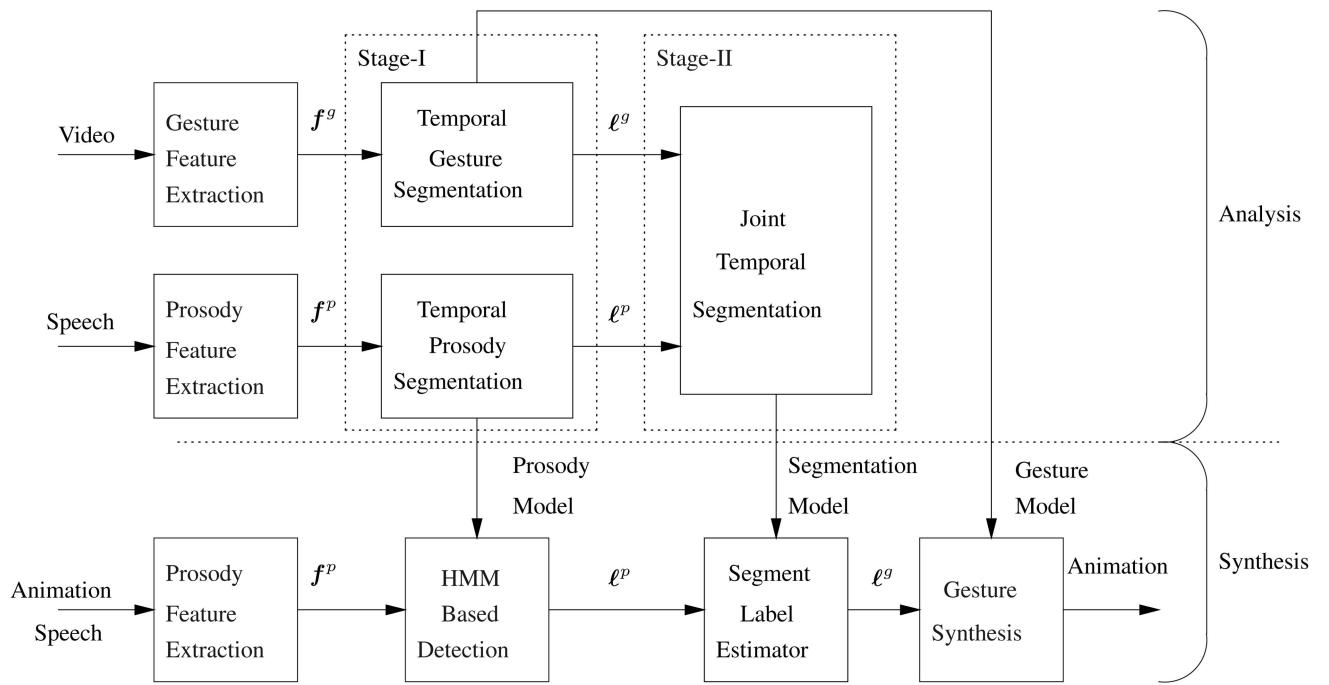


Fig. 1. Overview of the proposed synthesis-by-analysis system.

toward automatic generation of the audio-visual mapping from speech prosody to head gestures. Although the same framework can also be applied to the analysis of co-occurring arm gesture and speech patterns, this is beyond the scope of the current paper. There are some open challenges involved in the joint analysis of head gestures and prosody toward prosody-driven head gesture synthesis: First, unlike phonemes and visemes in speech articulation, there does not exist a well-established set of elementary prosody and gesture patterns for gesture synthesis. Second, synchronicity of gesture and prosody patterns may exhibit variations. For instance, a speaker can move her/his head before the corresponding prosodic utterance with a variable time lag. Moreover, gestural patterns may span time intervals of different duration with respect to their prosodic counterparts. Third, prosody and gesture patterns are speaker dependent and may exhibit variations in time even for the same speaker. Previously reported works [15], [16], [17] do not address any of these challenges, for instance, the asynchrony problem is either ignored or handled by manual alignment. In this work, we address these challenges by first processing the head gesture and prosody features separately by a parallel HMM structure to learn and model the gestural and prosodic elements (elementary patterns), respectively, over training data for a particular speaker. We then employ a multistream parallel HMM structure to find the jointly recurring gesture-prosody patterns and the corresponding audio-to-visual mapping.

HMM-based segmentation techniques are commonly employed in modeling multistream correlations; for example, for speech-driven lip animation in [7], [8], and [9] and for audio-visual event detection in [18]. We can classify HMM-based modeling techniques as supervised and unsupervised. Speech and lip motion correlation modeling can be thought of as a supervised analysis/segmentation problem, since phonemes and visemes

constitute well-established elementary units for these modalities. Hence, speech-driven lip animation task is often equivalent to find a mapping between the phonemes of speech and the visemes of lip movement. On the other hand, we shall consider the audio-visual gesture modeling/mapping as an unsupervised segmentation problem, where the recurrent joint events are not well defined and to be extracted from the joint feature streams.

The organization of this paper is given as follows: In Section 2, we first provide an overview of the proposed HMM-based analysis-synthesis framework and then describe the computation of head gesture and speech prosody features. Robust and accurate tracking of the speaker head motions is an integral part of the overall system; hence, it is described in detail. Section 3 presents the proposed two-stage unsupervised analysis procedure to identify and model jointly recurring head gesture and prosody patterns. Section 4 explains the HMM-based synthesis of head gesture parameters from input test speech. In Section 5, we describe the experiments conducted and present objective and subjective evaluation of the prosody-driven head gesture synthesis results. Finally, Section 6 provides conclusions.

2 OVERVIEW OF THE PROPOSED SYSTEM AND FEATURE EXTRACTION

A block diagram of the proposed system for prosody-driven head gesture animation, which consists of analysis and synthesis parts, is depicted in Fig. 1. The analysis part includes two feature extraction modules and two-stages of analysis. Feature extraction modules compute the head gesture features f^g and speech prosody features f^p , respectively, from training stereo video sequences of a speaker. At the first stage analysis, individual feature streams are used to train separate parallel HMM structures,

which provide probabilistic models for temporal recurrent patterns in the corresponding modalities, respectively. The segments corresponding to these patterns are detected and labeled over the training video streams, where pattern labels for prosody and gesture are denoted by l^p and l^g , respectively. At the second stage, the labels of temporally segmented gesture and prosody streams are used together to train a discrete multistream parallel HMM to identify jointly recurring patterns. The resulting joint HMM structure models the correlation between speech prosody and head gestures. The synthesis part makes use of the joint HMM to predict the gesture labels from the prosody labels computed for a test input speech using the prosody HMM obtained by the first stage analysis. The corresponding gesture features, that is, head motion parameters, are synthesized using the gesture HMM obtained at the first stage analysis and finally animated on a 3D head model. The details of the two stages of the analysis, shown by Stage-I and Stage-II blocks in Fig. 1, are presented in Section 3, whereas the gesture synthesis part is described in detail in Section 4. In the remainder of this section, we describe our methodology for extraction of head gesture and speech prosody features.

2.1 Extraction of Head Gesture Features

We define the head gesture feature vector, f_k^g , for frame k to include the Euler angles associated with the 3D head rotation and their first differences:

$$f_k^g = [\theta_k, \phi_k, \psi_k, \Delta\theta_k, \Delta\phi_k, \Delta\psi_k]^T, \quad (1)$$

where θ_k , ϕ_k , and ψ_k are the Euler angles of rotation, with respect to a reference frame k_r , around the x -, y -, and z -axes, respectively, and $\Delta\theta_k$, $\Delta\phi_k$, and $\Delta\psi_k$ denote their respective first differences. The reference frame k_r can be selected as the first frame in which the subject's head is assumed to be at neutral position.

2.1.1 3D Point Tracking

For video recording, we use a rectified stereo camera system with two identical cameras, and assume that the intrinsic camera parameters are known a priori. For each frame k , we initially detect a rectangular head region from one of the stereo views (for example, the right or the left but not both) using a boosted Haar-based cascade classifier structure, which was initially proposed by Viola and Jones [19] and later improved by Lienhart and Maydt [20]. The detected rectangular head region is used to initialize the search window within which facial pixels are segmented based on a Gaussian skin color distribution model computed over a training set of sampled skin colors. An ellipse \mathcal{E}_k is then fitted to the facial skin region.

Let P_{k_r} denote the set of image points within the ellipse \mathcal{E}_{k_r} of the reference frame k_r so that $P_{k_r} = \{p_{k_r,1}, p_{k_r,2}, \dots, p_{k_r,N}\}$ and $p_{k_r,n} = [x_n, y_n]^T$. For each frame k , we employ the hierarchical Lukas-Kanade technique [21] to find the optical flow vectors, $\{v_{k,1}, v_{k,2}, \dots, v_{k,N}\}$, from frame k_r to frame k . The set P_k of the corresponding image points in frame k is then obtained by $p_{k,n} = p_{k_r,n} + v_{k,n}$, $n = 0, 1, \dots, N$.

In order to find the 3D world coordinates of the image points in each set P_k , we compute the disparity vectors at these points using bandpass images and a cross correlation measure based on the sum of absolute differences [22]. The

disparity vectors are also validated using several criteria [23]. Given the disparity vectors for each frame and the intrinsic parameters of the rectified stereo camera system, the 3D world coordinates of the 2D points from both sets P_{k_r} and P_k are calculated by the well-known triangulation technique. Let W_k denote the $3 \times M$ matrix formed by the 3D world coordinates of the points associated with P_k , so that $W_k = [w_{k,1}, w_{k,2}, \dots, w_{k,M}]$ and $w_{k,m} = [X_m, Y_m, Z_m]^T$. Although forming the matrix W_k , we exclude those points in P_k that fall outside the ellipse \mathcal{E}_k due to possible erroneous optical flow vectors. The excluded points are outliers that may corrupt the 3D motion capture process. Hence, the dimension M of the matrices W_k and W_{k_r} are redetermined at each frame k according to the number of points that fall within the detected ellipse \mathcal{E}_k .

2.1.2 Computation of the Euler Angles

Let R_k and t_k denote the rotation matrix and the translation vector, respectively, of the rigid head motion from frame k_r to k . Then, W_k and W_{k_r} are related by

$$W_k = [R_k \ t_k] \begin{bmatrix} W_{k_r} \\ 1^T \end{bmatrix}. \quad (2)$$

The rotation matrix R_k and translation vector t_k are estimated by a unitary constraint optimization technique as explained in the Appendix. Once estimated, the rotation matrix R_k can be decomposed into three matrices:

$$R_k = [r_{ij}^k] = R_x(\theta_k)R_y(\phi_k)R_z(\psi_k), \quad (3)$$

where $R_x(\theta_k)$, $R_y(\phi_k)$, and $R_z(\psi_k)$ are the matrices that specify rotations around x -, y -, and z -axes, respectively [24], [25]. The Euler angle vector $e_k = [\theta_k, \phi_k, \psi_k]^T$ that maps W_{k_r} to W_k is finally extracted from this decomposition by

$$e_k = [\arctan(-r_{23}^k/r_{33}^k), \arcsin(r_{13}^k), \arctan(-r_{12}^k/r_{11}^k)]^T. \quad (4)$$

In cases where the head rotation between the current frame k and reference frame k_r is larger than a threshold angle (for example, if $|\theta_k| > 25$ degrees, $|\phi_k| > 25$ degrees, or $|\psi_k| > 25$ degrees), the optical flow vectors and, hence, the 3D point correspondences between two frames may become unreliable. In such cases, we switch to incremental motion estimation, where the reference frame for frame k is set to frame $k - 1$. Thus, we recompute optical flow vectors with respect to frame $k - 1$; hence, the new 3D point correspondences and the resulting incremental Euler angle vector δ_{k-1} , which defines the rotation between frames k and $k - 1$ are computed. Then, the Euler angle vector with respect to the reference frame k_r is given by

$$e_k = e_{k-1} + \delta_{k-1}. \quad (5)$$

2.1.3 Smoothing of the Feature Vector by Kalman Filtering

We finally employ a Kalman filter for post smoothing of the computed (estimated) Euler angles, which are input as observations z_k to the Kalman filter. The measurement noise r_k models the estimation errors in the Euler angles. The head gesture feature vector, f_k (the superscript g is omitted for ease of notation), consisting of the Euler angles and their

first differences, is selected as the state vector. The state-space representation of the Kalman filter is given by

$$\begin{aligned} \mathbf{f}_{k+1} &= \mathbf{F}\mathbf{f}_k + \mathbf{G}\mathbf{u}_k \\ \mathbf{z}_k &= \mathbf{H}\mathbf{f}_k + \mathbf{r}_k, \end{aligned} \quad (6)$$

where

$$\begin{aligned} \mathbf{f}_k &= \begin{bmatrix} \mathbf{e}_k \\ \Delta\mathbf{e}_k \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} \mathbf{I}_{3 \times 3} & \mathbf{I}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \end{bmatrix} \\ \mathbf{G} &= \mathbf{I}_{6 \times 6}, \quad \mathbf{H} = \begin{bmatrix} \mathbf{I}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} \end{bmatrix}. \end{aligned} \quad (7)$$

The 3×1 vector $\Delta\mathbf{e}_k$ denotes the first differences of the Euler angles. The model noise \mathbf{u}_k and measurement noise \mathbf{r}_k are assumed to be uncorrelated, zero-mean white Gaussian processes. The output of the Kalman filter gives the final feature vector for the head gestures.

2.2 Extraction of Prosody Features

The prosodic speech events can be described by the temporal variations of loudness/intensity and pitch, as well as pauses between phrases, phoneme duration, timing, and rhythm. Among these, the most expressive one is the pitch, which is the rate of vocal-fold cycling. In this study, pitch frequency, V , and speech intensity, I , are considered as prosody features.

The pitch contour is extracted at a rate of 100 Hz from the speech signal using the autocorrelation method, as described in [26]. The mean of all pitch contours over all active utterances is removed to emphasize local variations [27], and then, the resulting mean-removed contours are low-pass filtered to reduce discontinuities. The regions between utterances without a valid pitch are filled with zero mean unit variance Gaussian noise. The intensity features are also extracted over the active utterances. The squared sound intensities are weighted with a 32 ms Kaiser-20 window, and the speech signal intensity is calculated as the sum of these weighted samples. The 32 ms window is shifted by 10 ms for each frame to extract intensity values at 100 Hz frame rate. The intensity features are also mean removed over active utterances and between-utterance regions are filled with zero mean unit variance Gaussian noise. The first-order derivative, ΔV_k , of the postprocessed pitch frequency at frame k is calculated using the following regression formula:

$$\Delta V_k = \frac{\sum_{i=1}^2 i(V_{k+i} - V_{k-i})}{2 \sum_{i=1}^2 i^2}. \quad (8)$$

Finally, the pitch frequency, its derivative and the intensity are concatenated to form the 3D prosody feature vector \mathbf{f}_k^p at frame k :

$$\mathbf{f}_k^p = [V_k \Delta V_k I_k]^T. \quad (9)$$

3 HEAD GESTURE-PROSODY PATTERN ANALYSIS

In this section, we propose a two-stage HMM-based unsupervised analysis framework, where the first stage aims to separately extract elementary gesture and prosody patterns for a speaker, and the second stage determines a correlation model between these head gesture and prosody patterns. In the first stage analysis, recurring elementary

gesture and prosody patterns are determined separately by unsupervised temporal clustering of individual gesture and prosody feature streams, respectively. The extracted elementary prosody and gesture patterns are analogous to phonemes and visemes in the speech and lip motion modeling. However, the elementary gesture and prosody patterns are not well established as in the case of phonemes and visemes, since the nature and strength of head gesture and prosody patterns may vary from person to person and in time. Hence, there is a need for unsupervised stage I analysis in order to extract these patterns for each speaker. Furthermore, the joint recurring nature of these patterns are also not well established as in the case of phoneme-viseme association; hence, the need for stage II analysis for joint modeling of correlations between head gesture and prosody patterns. In order to find a mapping between prosody and gesture patterns, unsupervised temporal segmentation of joint gesture and prosody pattern labels is performed, which defines the correlation between gesture and prosody pattern streams and relates co-occurring head gesture and prosody patterns.

We note that if a multistream HMM structure were directly employed for joint analysis of gesture and prosody feature streams, as commonly used for event detection [18], instead of the proposed two-stage analysis, the resulting joint gesture-prosody feature segments would not necessarily correspond to *independent* meaningful elementary gesture and prosody patterns. As a result, the synthesized gesture sequence might contain poorly defined gestural elements, which would degrade the quality of prosody-driven head gesture animation.

3.1 Stage I: Extraction of Elementary Head Gesture and Prosody Patterns

The first stage analysis defines recurrent elementary head gesture and prosody patterns separately using unsupervised temporal clustering over individual feature streams. The gesture and prosody feature streams F^g and F^p are separately used to train two HMM structures Λ_g and Λ_p , which capture recurrent head gesture segments ε^g and prosody segments ε^p . For ease of notation, we use a generic notation to represent the HMM structure that is identical for the gesture and prosody streams. The HMM structure Λ , which is used for unsupervised temporal segmentation, has M parallel branches and N states, as shown in Fig. 2. In the HMM structure Λ , observation probability densities are modeled by a single Gaussian with diagonal covariance for both gesture and prosody streams. The states labeled as s_s and s_e are nonemitting start and end states of the parallel HMM structure. Fig. 2 clearly illustrates that the parallel HMM Λ is composed of M parallel left-to-right HMMs, $\{\lambda_1, \lambda_2, \dots, \lambda_M\}$, where each λ_m is composed of N states, $\{s_{m,1}, s_{m,2}, \dots, s_{m,N}\}$. The state transition matrix A_{λ_m} of each λ_m is associated with a subdiagonal matrix of A_Λ . The feature stream is a sequence of feature vectors, $F = \{f_1, f_2, \dots, f_T\}$, where f_t denotes the feature vector at frame t . Unsupervised temporal segmentation using HMM model Λ yields L number of segments $\varepsilon = \{\varepsilon_1, \varepsilon_2, \dots, \varepsilon_L\}$. The l th temporal segment is associated with the following sequence of feature vectors:

$$\varepsilon_l = \{f_{t_l}, f_{t_l+1}, \dots, f_{t_{l+1}-1}\} \quad l = 1, 2, \dots, L, \quad (10)$$

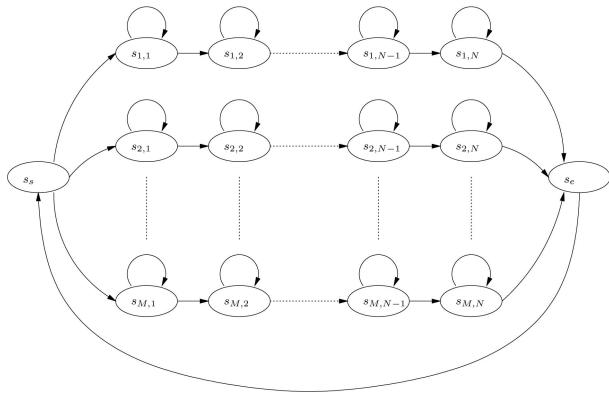


Fig. 2. Parallel HMM structure.

where f_{t_1} is the first feature vector f_1 , and $f_{t_{L+1}-1}$ is the last feature vector f_T .

The segmentation of the feature stream is performed using Viterbi decoding to maximize the probability of model match, which is the probability of feature sequence F given the trained parallel HMM Λ :

$$\begin{aligned} P(F|\Lambda) &= \max_{t_l, m_l} \prod_{l=1}^L P(\{f_{t_l}, f_{t_l+1}, \dots, f_{t_{l+1}-1}\} | \lambda_{m_l}) \\ &= \max_{\varepsilon_l, m_l} \prod_{l=1}^L P(\varepsilon_l | \lambda_{m_l}), \end{aligned} \quad (11)$$

where ε_l is the l th temporal segment, which is modeled by the m_l th branch of the parallel HMM Λ . One can show that λ_{m_l} is the best match for the feature sequence ε_l , that is

$$m_l = \operatorname{argmax}_m P(\varepsilon_l | \lambda_m). \quad (12)$$

Since, the temporal segment ε_l from frame t_l to $(t_{l+1}-1)$ is associated with segment label m_l , we define the sequence of frame labels based on this association as

$$\ell_t = m_l \quad \text{for } t = t_l, t_l + 1, \dots, t_{l+1} - 1, \quad (13)$$

where ℓ_t is the label of the t th frame, and we have a label sequence $\ell = \{\ell_1, \ell_2, \dots, \ell_T\}$ corresponding to the feature sequence F . The first stage analysis extracts the frame label sequences ℓ^g and ℓ^p given the head gesture and prosody feature streams F^g and F^p . Although mapping the gesture and prosody features to discrete frame labels, the mismatch between the frame rates of gesture and prosody is eliminated by downsampling the frame rate of prosody label stream to the rate of gesture label stream.

The parallel HMM structure has two important parameters to set before training the model Λ . The first parameter is the number of states in each branch, N . It should be selected by considering the minimum duration of temporal patterns. Selecting a small N may hamper modeling long-term statistics for each branch of the parallel HMM. The extreme case $N = 1$ reduces to K-Means unsupervised clustering. We select the number of states in each branch of the head gesture HMM Λ_g as $N_{\Lambda_g} = 10$, corresponding to the minimum gesture pattern duration of 10 frames ($\frac{1}{3}$ sec assuming 30 video frames/sec). Note that the gesture patterns can be longer than 10 frames since the HMM structure allows self-state transitions. On the other hand, the prosody patterns are expected to follow smooth

pitch frequency movements over several syllables. Considering the average syllable duration and smoothness of the pitch contours, we set $N_{\Lambda_p} = 5$ in each branch of the prosody HMM model Λ_p .

The second parameter is the number of temporal patterns, M . Since the number of head gesture and prosody patterns is speaker dependent, we propose a selection of M by using two fitness measures. The first fitness measure α , which is inversely related to in-class variance, is defined as the frame average of the log-probability of model match:

$$\alpha = \frac{1}{T} \log(P(F|\Lambda)). \quad (14)$$

The α measure is expected to saturate with an increasing number of parallel branches in Λ , since the training database is expected to contain a limited number of temporal patterns. However, small variations within temporal patterns are also expected; hence, the number of branches M can be more than the actual number of temporal patterns in the training corpus. Consequently, the second fitness measure, which is the average statistical separation between two similar temporal patterns, increases with the decreasing number of temporal patterns. The second fitness measure β is considered as the average statistical separation between two similar temporal patterns, and it is defined as

$$\beta = \frac{1}{T} \sum_{l=1}^L \log \left(\frac{P(\varepsilon_l | \lambda_{m_l})}{P(\varepsilon_l | \lambda_{m_l^*})} \right), \quad (15)$$

where $\lambda_{m_l^*}$ is the second best match for the temporal segment ε_l , that is

$$m_l^* = \operatorname{argmax}_{\forall m \neq m_l} P(\varepsilon_l | \lambda_m). \quad (16)$$

In general, the α measure increases with the number of patterns M , whereas β measure decreases. Hence, a good value for M can be selected such that β is high enough, whereas α reaches a certain value.

3.2 Stage II: Joint Modeling of Prosody-Gesture Patterns

In the second stage, unsupervised segmentation of the joint gesture-prosody label stream is performed to detect recurrent joint label patterns. Note that this task is similar to the task of stage I, except in the second stage, we have a multistream discrete observation sequence. For this task, the parallel HMM structure in Fig. 2 is used with discrete multistream HMM branches. In multistream HMMs, all streams share the same state transition structure however emission probabilities are determined independently for each stream.

The joint gesture-prosody frame label stream, denoted by ℓ^{gp} , is defined such that for every frame k , $\ell_k^{gp} = [\ell_k^g, \ell_k^p]^T$. We represent the discrete multistream parallel HMM structure by Γ_{gp} and its m th branch by γ_m^{gp} . The discrete HMM Γ_{gp} is trained over the joint gesture-prosody label stream. Each branch γ_m^{gp} , associated with a joint gesture-prosody temporal label pattern, is then described by

$$\gamma_m^{gp} = (\mathbf{A}_{\gamma_m^{gp}}, [\mathbf{B}_{\gamma_m^g} \mathbf{B}_{\gamma_m^p}], \boldsymbol{\Pi}_{\gamma_m^{gp}}), \quad (17)$$

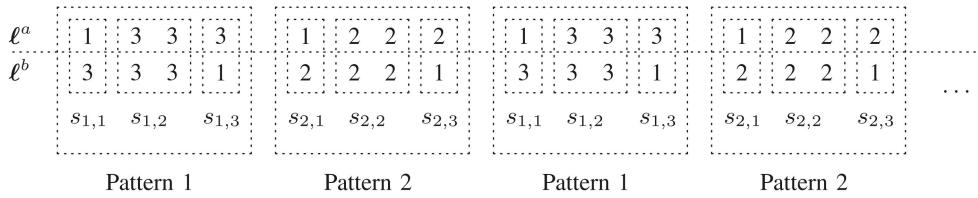


Fig. 3. Example for unsupervised joint label segmentation.

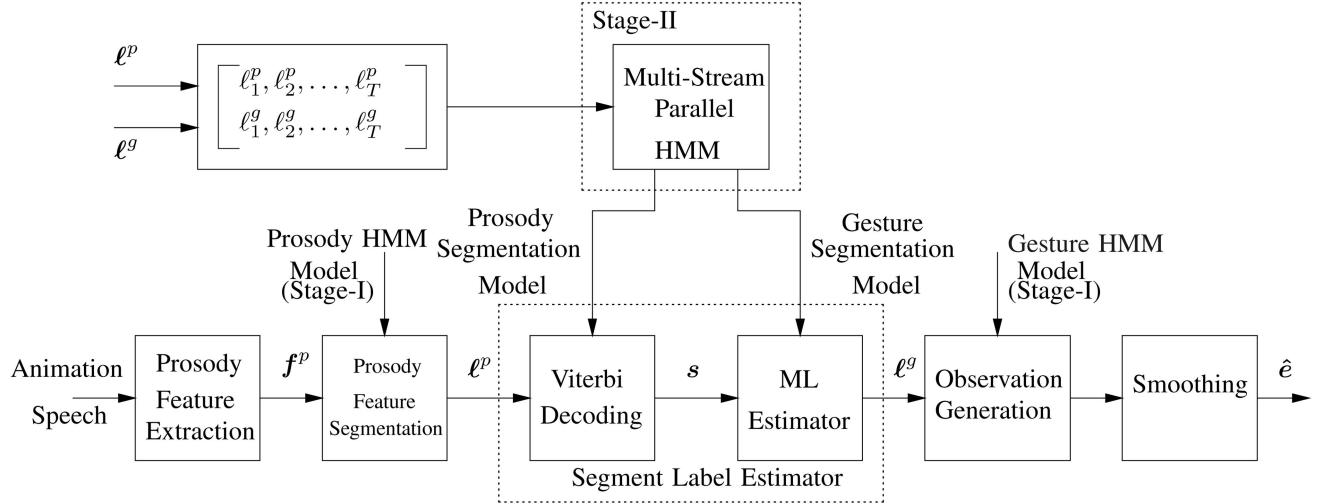


Fig. 4. The proposed prosody-driven gesture synthesis system.

where $A_{\gamma_m^{gp}}$ denotes a state transition matrix, $B_{\gamma_m^g}$ and $B_{\gamma_m^p}$ are discrete observation probability distributions for gesture and prosody label streams, and $\Pi_{\gamma_m^{gp}}$ is an initial state probability matrix. The distributions $B_{\gamma_m^g}$ and $B_{\gamma_m^p}$ define the probability of observing a gesture-prosody label at state s and frame k , given by

$$P(\ell_k^{gp}|s) = P(\ell_k^g|s)^{\kappa_g} P(\ell_k^p|s)^{\kappa_p}, \quad (18)$$

where the exponents, κ_g and κ_p , are the stream weights, which may be set to unity.

For the purpose of synthesis, each multistream discrete HMM branch, γ_m^{gp} , can be split into two individual single-stream discrete HMM models $\gamma_m^g = (A_{\gamma_m^{gp}}, B_{\gamma_m^g}, \Pi_{\gamma_m^{gp}})$ and $\gamma_m^p = (A_{\gamma_m^{gp}}, B_{\gamma_m^p}, \Pi_{\gamma_m^{gp}})$, respectively, for gesture and prosody streams. These single stream HMM models share the same state transition and initial state probability matrices, but their discrete observation probability distributions are different. The individual observation distributions are then given by $P(\ell_k^g|s)$ and $P(\ell_k^p|s)$ for gesture and prosody models, respectively.

Unsupervised temporal segmentation of joint label streams is demonstrated by the following example, which also illustrates how the asynchrony between gestures and prosody is handled in our scheme.

Example. Let us have two label streams ℓ^a and ℓ^b , where each label can assume values 1, 2, or 3. When temporal segmentation of the joint label stream is performed using the HMM structure Γ_{gp} with $M = 2$ patterns and $N = 3$ number of states for each pattern, we obtain the result shown in Fig. 3. One can observe that the recurrent joint label patterns are captured and the asynchrony between individual label streams is modeled by the first and the last states of the HMM branches.

The number of states $N_{\Gamma_{gp}}$ for each branch of Γ_{gp} should be selected according to the number of head gesture and prosody patterns determined by the stage I analysis, since Γ_{gp} models the recurrent joint gesture-prosody label pairs. Similarly, the number of branches $M_{\Gamma_{gp}}$ in Γ_{gp} should be selected by considering the two fitness measures α and β , as defined in (14) and (15). The selection of $N_{\Gamma_{gp}}$ and $M_{\Gamma_{gp}}$ is further discussed in Section 5.

4 PROSODY-DRIVEN GESTURE SYNTHESIS

In this section, we address prosody-driven gesture synthesis using the proposed gesture-prosody pattern model. A detailed block diagram of the proposed prosody-driven gesture synthesis system is shown in Fig. 4. The system takes speech as input and produces a sequence of head gesture features, that is, Euler angle vectors, which are naturally correlated with the input speech. The details of the subblocks are described in the following.

4.1.1 Prosody Feature Extraction

The prosody features, F^p , are extracted from the input speech signal, as described in Section 2.2.

4.1.2 Prosody Feature Segmentation

Temporal segmentation of prosody feature sequence F^p is performed using the HMM model Λ_p , which is trained in the stage I analysis in Section 3.1. During the temporal segmentation, the conditional probability $P(F^p|\Lambda_p)$ is maximized using Viterbi decoding to extract the temporal prosody segment sequence, ε^p , and the sequence of prosody frame labels, ℓ^p .

4.1.3 Gesture Segment Label Estimation

The aim of this step is to predict the sequence of gesture frame labels, ℓ^g , given the prosody frame labels ℓ^p . To this effect, temporal segmentation of the prosody frame labels, ℓ^p is performed using the HMM model Γ_p , which is extracted by splitting the jointly trained gesture-prosody HMM model Γ_{gp} . As a result of this temporal prosody label segmentation, a state sequence $s^p = \{s_1^p, s_2^p, \dots, s_K^p\}$ associated with $\ell^p = \{\ell_1^p, \ell_2^p, \dots, \ell_K^p\}$ is extracted. Then, the gesture frame label sequence ℓ^g is predicted by maximizing the probability of observing gesture label on the state sequence path s^p over the gesture HMM model Γ_g , such that

$$\ell_k^g = \arg \max_m P(m|s_k^p, \Gamma_g), \quad (19)$$

where k is the frame index, m runs over all possible M gesture patterns, and the conditional probability $P(m|s_k^p, \Gamma_g)$ is defined by the discrete observation probability distribution $B_{\gamma_m^g}$.

4.1.4 Generation of Euler Angles

This step computes the gesture segment sequence ε^g , consisting of the Euler angle features, given the gesture frame label sequence ℓ^g . First, we find the segment frame boundaries, $\{t_l\}_{l=1}^L$, by merging the same gesture frame labels in the sequence ℓ^g . Then, the Euler angle features for the l th segment, $\varepsilon_l^g = \{f_{t_l}^g, f_{t_l+1}^g, \dots, f_{t_{l+1}-1}^g\}$, are generated from the HMM $\lambda_{t_l}^g$, which is the t_l -th branch of the parallel HMM model Λ_g (computed in stage I).

Note that the segment duration for the l th segment is extended as $d_l = (t_{l+1} + \Delta - (t_l - \Delta))$ frames, where Δ is the number of overlapping frames at the segment boundaries to smooth segment-to-segment transitions. The state sequence s_l^g or equivalently the state occupancy duration for the l th segment is calculated using the diagonal terms of the d_l -step state transition matrix of the HMM $\lambda_{t_l}^g$. Having the state sequence s_l^g and the continuous observation probability $P(f^g|s_l^g)$, which are modeled using a Gaussian distribution, the Euler angle features are generated along the state sequence associated with the distribution $P(f^g|s_l^g)$. The segment boundaries have $2\Delta + 1$ number of frame overlaps, where the overlapped and averaged features generate smoother segment-to-segment transitions.

4.1.5 Smoothing of Euler Angles

As the final step of the gesture synthesis, the Euler angles are smoothed using median filtering followed by a Gaussian low pass filter to remove motion jerkiness. The median filtering is performed over 11 visual frames, and the Gaussian smoothing is performed over 15 visual frames. Fig. 5 depicts the samples generated from the HMM and outputs of the median and Gaussian filters. The figure clearly shows that the median filter removes jitters within a state and the Gaussian low pass filter smooths the state-to-state transitions.

There are two main advantages of using HMMs for gesture synthesis. The first is the random variations in the synthesized gesture patterns for each segment. This variation yields more natural looking synthesis results than using a fixed gesture dictionary, since humans produce slightly varying gestures at different occasions for the same

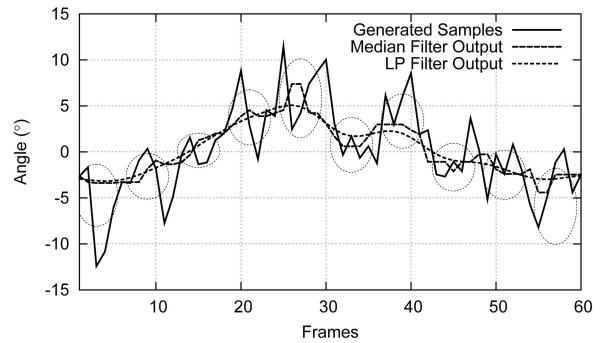


Fig. 5. The effect of filtering in the synthesis of Euler angle θ . The dashed circles represent the states of a single branch HMM model. The vertical position and size of each circle are adjusted considering the mean and variance of the Euler angles associated with each state.

semantics. The second advantage is generating gestures with varying duration in accordance with the prosody of the speaker.

5 EVALUATION AND RESULTS

In this section, we present experimental results and the evaluation of the proposed system. Section 5.1 describes the audio-visual database, which is used in the experimental evaluation to generate objective and subjective results. The evaluation of the gesture-prosody pattern analysis is presented in Section 5.2, and the objective and subjective performance results for synthesis are presented in Section 5.3. Speaker dependency of the prosody-driven head gesture synthesis system is evaluated in Section 5.4.

5.1 Database and Experimental Setup

We have conducted experiments using the MVGL-MASAL gesture-speech database. The database includes four recordings of two subjects telling stories in Turkish. The subjects are instructed to tell stories to children audience. All gestures are spontaneous within this context. Each story lasts approximately seven minutes. The audio-visual data is synchronously captured from the stereo camera and the sound card. The stereo video includes only upper body gestures with 30 frames per second, whereas the audio is recorded with 16-kHz sampling rate and 16 bits per sample. The detailed specification of the stereo camera can be found in [28]. The performance of the proposed analysis and synthesis system is evaluated in detail on the recordings of the first speaker, whereas the recordings of the second speaker are used to investigate the speaker dependency problem. For the first speaker, the database is partitioned into two parts such that three stories are used for the training of the models, and one story is used for testing. For objective evaluation of the synthesis, the Euler angles extracted from the test sequence are considered as the ground truth for the synthesized head motion.

5.2 Analysis Results

The head gesture and prosody correlation analysis includes unsupervised temporal segmentation of the individual feature streams as well as the joint gesture-prosody label stream. The objective and subjective evaluation of these tasks are presented in the following.

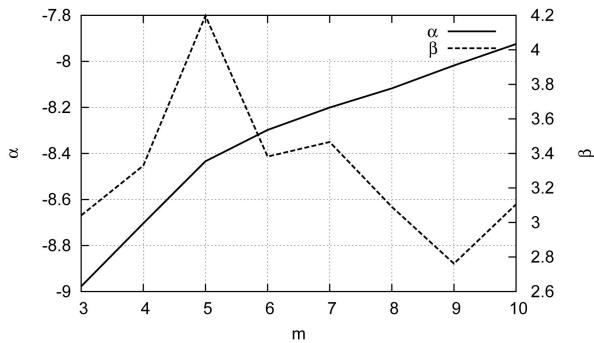


Fig. 6. The α and β fitness measures for varying number of head gesture patterns.

5.2.1 Segmentation of Head Gesture Patterns

The parallel HMM Λ_g is trained with features extracted from the training video using Expectation-Maximization (EM) algorithm. The resulting HMM structure provides a probabilistic cluster model for unsupervised segmentation of head gestures into recurring elementary patterns.

The number of branches, or equivalently the number of gesture patterns, M_{Λ_g} is a critical model parameter. In order to set M_{Λ_g} , the two fitness measures α and β , as respectively defined in (14) and (15), are calculated for varying number of gesture patterns and plotted in Fig. 6. The α value, which measures the probability of model match, increases with an increasing number of patterns as expected. Note that β measures the statistical separation between patterns. A good value for M_{Λ_g} is such that β is high enough, whereas α reaches a certain value. Therefore, M_{Λ_g} can be selected in the vicinity where α and β curves (normalized with their minimum and maximum values) intersect. In Fig. 6, β reaches the maximum at $M_{\Lambda_g} = 5$, which is near the intersection point. Hence, we set the number of gesture patterns M_{Λ_g} to 5.

Consequently, when the training head gesture sequence is segmented using Λ_g , the segments belonging to the same gestural patterns are observed to be visually alike. The mean Euler angle vectors and the typical thumbnails for the five gesture patterns are depicted in Fig. 7.

5.2.2 Segmentation of Prosody Patterns

The speech prosody feature sequence is extracted from the audio part of the training database. As defined in stage *I*, the HMM model Λ_p is trained with prosodic features to obtain unsupervised temporal segmentation of the audio stream.

The two fitness measures α and β are calculated for varying number of prosody patterns using HMM model Λ_p and plotted in Fig. 8. The α value, which measures the probability of model match, increases and the β value, which measures statistical separation between patterns, decreases with increasing number of patterns as expected. The number of prosody patterns M_{Λ_p} can thus be set to a value in the vicinity where α and β curves intersect. Hence, we select M_{Λ_p} as 5 in our experiments.

The means and standard deviations of the normalized intensity and pitch frequency trajectories for the five prosody patterns are depicted in Fig. 9. Note that the first pitch trajectory (upper left) is associated with the no-pitch segments that we filled with zero mean and unit variance

Gaussian noise. The noise filling is necessary for successful modeling of those segments with continuous density HMMs. The other four prosody patterns can be classified using the prosodic transcription conventions introduced by the American English Tones and Break Indices (ToBI) standard [29]. The two prosody patterns on the upper right are both falling boundary tones (L percent); the pattern on the lower left is a falling boundary tone, which makes a peak before the last syllable (HL percent), and the pattern on the lower right is a rising-falling boundary tone, which rises within the last syllable (LHL percent). We should note that these prosody patterns are obtained using unsupervised clustering over the training database, and they do not define a complete prosodic transcription convention for Turkish.

5.2.3 Segmentation of Joint Gesture-Prosody Patterns

In the first stage analysis, we obtain two independent HMM structures, Λ_g and Λ_p , respectively, for recurrent head gesture and prosody patterns. We then extract two independent and parallel streams of head gesture and prosody pattern labels via temporal segmentation using these HMM structures. In the second stage, the discrete multistream HMM structure Γ_{gp} is trained using EM over the joint gesture-prosody pattern label stream to perform unsupervised segmentation. The number of states for each branch of Γ_{gp} is selected as $N_{\Gamma_{gp}} = 4$ to model possible label pair transitions. These four states model four different gesture-prosody label pair combinations within a joint gesture-prosody label pattern. Note that the extreme case, $N_{\Gamma_{gp}} = 1$, can only model a single co-occurrence pattern of gesture-prosody labels.

The two fitness measures α and β for Γ_{gp} , and also the number of gesture patterns in Λ_g , are considered for selection of the number of joint gesture-prosody label patterns $M_{\Gamma_{gp}}$. The number of joint patterns $M_{\Gamma_{gp}}$ is expected to be larger than or equal to the number of gesture patterns M_{Λ_g} , since in a robust synthesis process all the gesture patterns need to be generated for some temporal prosody label pattern. Hence, for the selection of $M_{\Gamma_{gp}}$, we present the two fitness measures α and β together with the normalized euclidean distance measure ϵ_n , as defined in (20), for varying number of joint gesture-prosody label patterns in Fig. 10. The parameter $M_{\Gamma_{gp}}$ is selected as 6, since this value, which is near the intersection of α and β curves, is greater than M_{Λ_g} , and the distance ϵ_n has a minimum at $M_{\Gamma_{gp}} = 6$.

Observation probability distributions of the joint multistream HMM are plotted in Fig. 11. It can be seen that each branch of the HMM structure Γ_g models a temporal sequence of identical elementary gesture patterns. That is, in each of the six classes, a distribution of prosody patterns co-occurs with a single elementary gesture pattern. Note that this association between temporal prosody label patterns and a single gesture pattern is very beneficial to obtain smooth prosody-driven head gesture animations. Furthermore, boundaries of the prosody patterns within the co-occurring gesture pattern are determined by the state transition probabilities of Γ_p , and hence, the asynchrony problem is handled through the learned statistics of the joint multistream HMM Γ_{gp} .

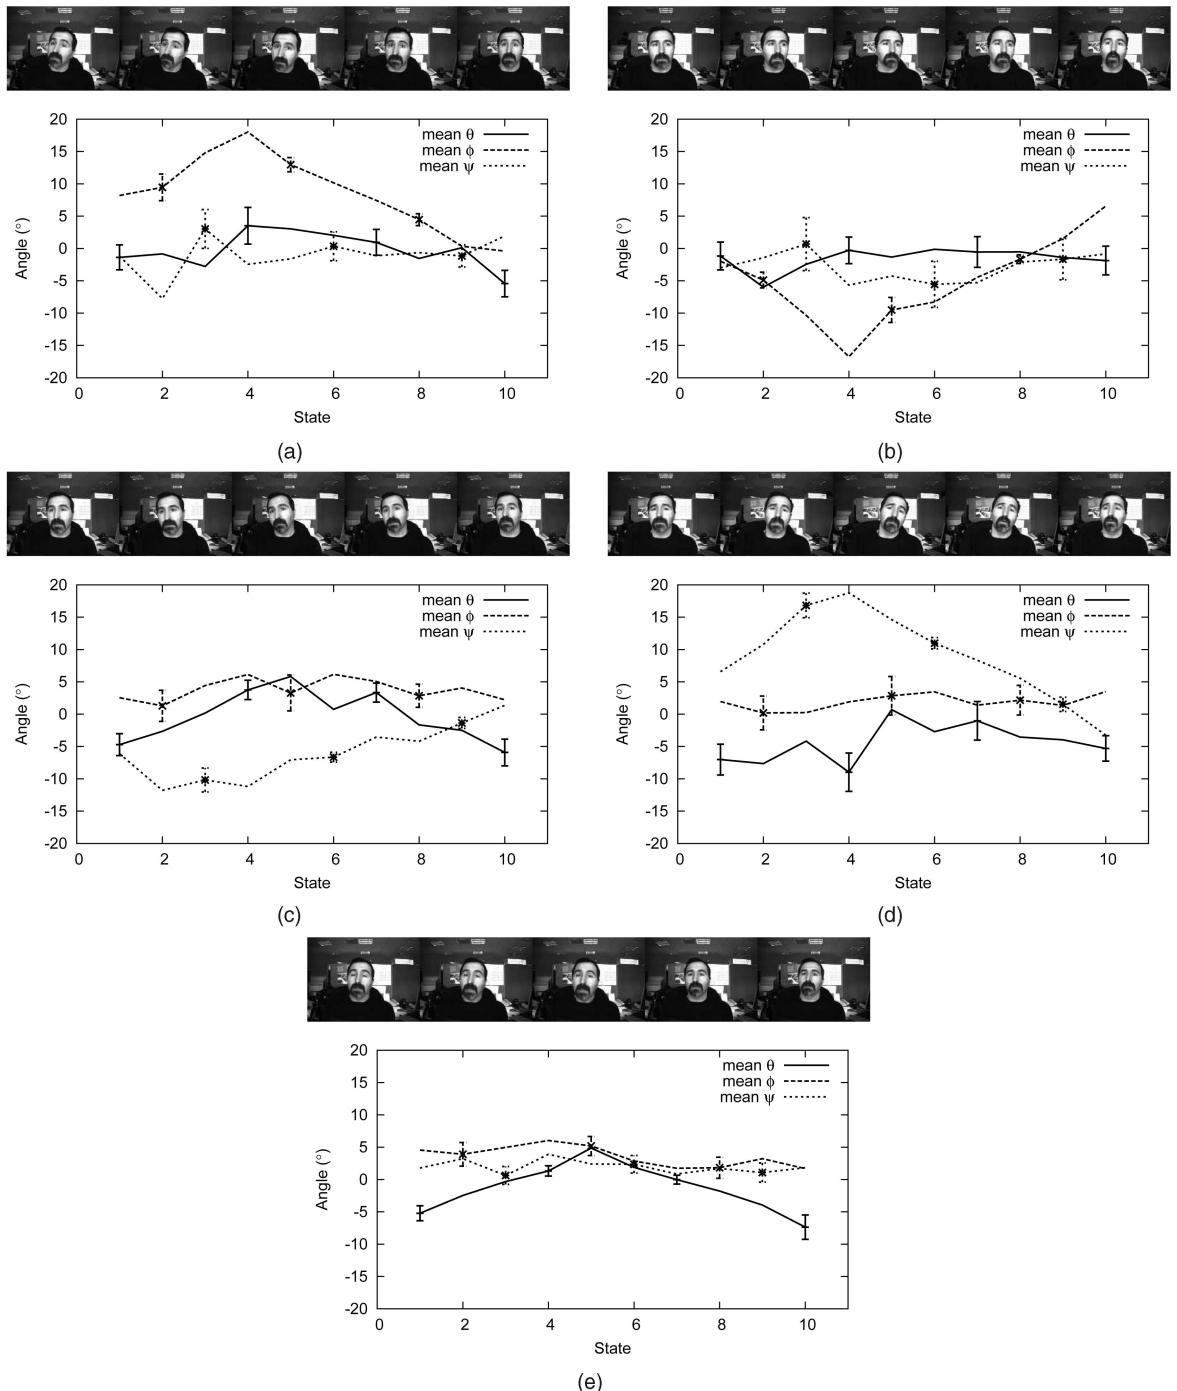


Fig. 7. The mean Euler angles with standard deviations and typical thumbnails for the five gesture patterns: (a) Turn Left, (b) Turn Right, (c) Tilt Left, (d) Tilt Right, and (e) Nod.

5.3 Synthesis Results

Prosody-driven head gesture synthesis generates an Euler angle sequence, which is naturally correlated to a given test speech signal. The details of the synthesis process is given in Section 4. In this section, we present objective and subjective evaluations of the prosody-driven head gesture synthesis process. The evaluations are performed over the test database, which is defined in Section 5.1.

The objective evaluations compare the difference between original and synthesized Euler angles. Furthermore, A-B comparison type subjective evaluations are performed

using the talking head avatar of *Momentum Inc.* [30], where the Euler angles that we deliver are used to drive head gestures/motion of the speech-driven talking head animation. The subjective tests are used to measure opinions on the naturalness of the synthesized head gestures using the speech-driven talking head animations.

We have adopted the Input-Output HMM (IOHMM) structure [8], [31] as a possible alternative scheme for the joint analysis of gesture and prosody label streams at the second stage. In that case, the IOHMM structure replaces the HMM Γ_{gp} to predict gesture segment labels from

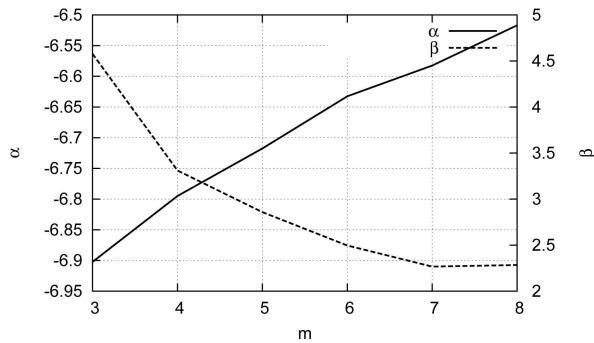


Fig. 8. The α and β fitness measures for varying number of prosody patterns.

prosody labels. The states in the IOHMM are fully connected, and the number of states is selected to be the same as the number of states in the Γ_{gp} model, which is 24. The IOHMM implementation of the Torch Machine Learning Library [32] is used in our experiments.

5.3.1 Objective Results

The objective evaluations compare the distance between original and synthesized Euler angles. In our evaluations, we have used three different distance measures. Let the original and synthesized Euler angles at frame k are represented with e_k and \hat{e}_k , respectively. The first distance measure ϵ_n is a normalized euclidean distance measure, which penalize Euler angles in wrong directions [6]:

$$\epsilon_n = \frac{\sum_{k=1}^K (\hat{e}_k - e_k)^T (\hat{e}_k - e_k)}{\sum_{k=1}^K (\hat{e}_k + e_k)^T (\hat{e}_k + e_k)}. \quad (20)$$

The second measure ϵ_m is the Mahalanobis distance, which is the euclidean distance weighted with the inverse covariance matrix, Σ^{-1} , of the original Euler angles e_k :

$$\epsilon_m = \frac{1}{K} \sum_{k=1}^K \sqrt{(\hat{e}_k - e_k)^T \Sigma^{-1} (\hat{e}_k - e_k)}. \quad (21)$$

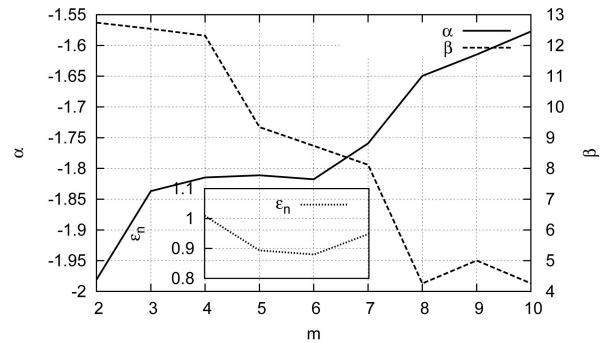


Fig. 10. The α and β fitness measures and the normalized euclidean distance measure ϵ_n for varying number of joint gesture-prosody label patterns.

The third distance measure is the standard euclidean distance, $\epsilon_e = \frac{1}{K} \sum_{k=1}^K \sqrt{(\hat{e}_k - e_k)^T (\hat{e}_k - e_k)}$.

The original Euler angles are extracted from the visual part of the test database to be used as the ground truth in the objective evaluations. Two sets of synthesized Euler angles are generated using the audio part of the test database. The first set is generated with the proposed head gesture synthesis system based on the Γ_{gp} model. The second set is generated by replacing the second stage joint gesture-prosody correlation model Γ_{gp} by IOHMM. The error measure statistics for the three distance measures ϵ_n , ϵ_m , and ϵ_e between the original and synthesized Euler angles are collected over synthesis trials repeated a hundred times. The mean and standard deviation of the distance measures are given in Table 1. Note that all three distance measures favor the proposed joint gesture-prosody correlation model Γ_{gp} .

5.3.2 Subjective Results

Subjective A-B comparisons are performed using the speech-driven talking head animations to measure opinions on the naturalness of the synthesized head gestures. The subjects are asked to evaluate the naturalness of the speech-driven synthesized head gestures for an A-B test pair on a

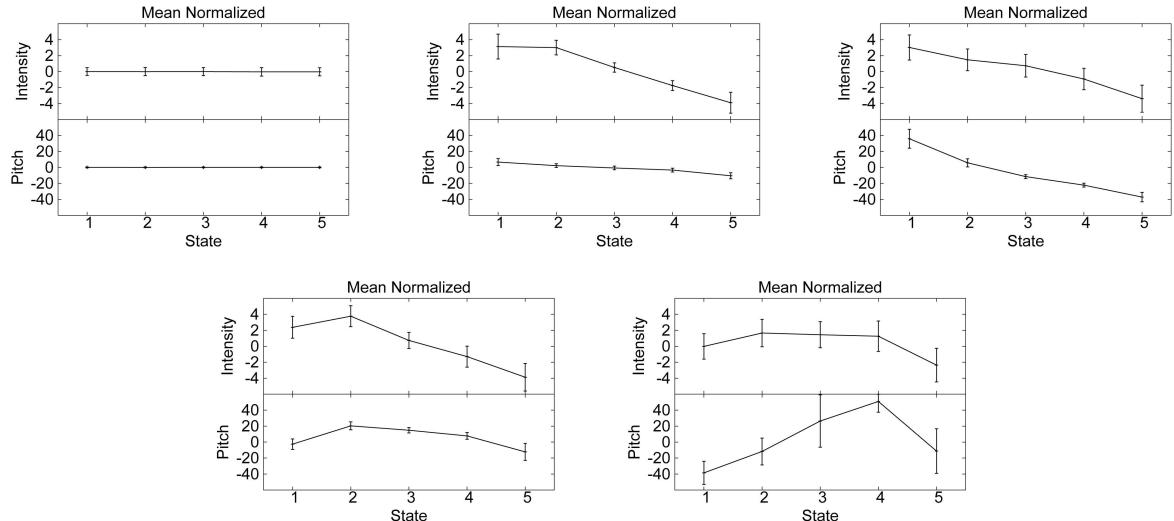


Fig. 9. The means and standard deviations of the normalized intensity (dB) and pitch frequency (Hz) trajectories for the five prosody patterns.

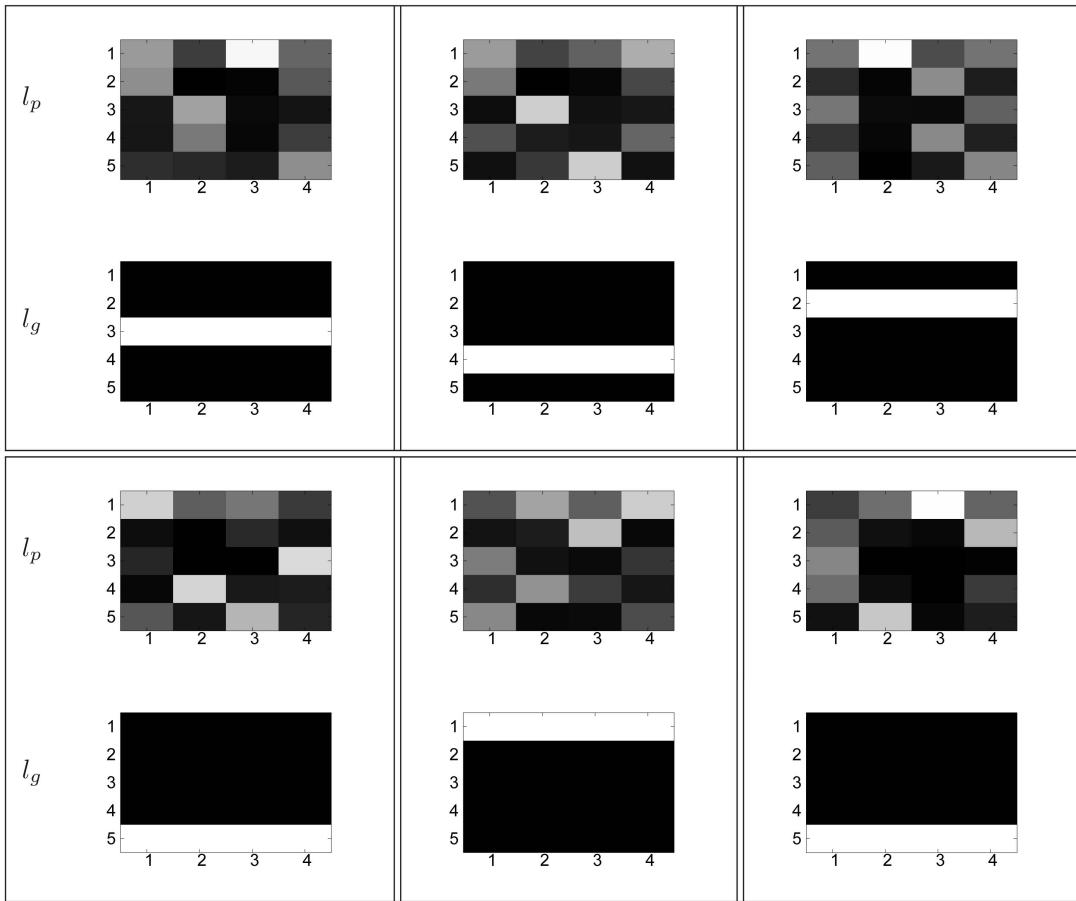


Fig. 11. Observation probability distributions of the joint multistream HMM. The discrete prosody and gesture labels are on the y -axis, whereas states are on the x -axis. Dark and white regions represent low- and high-probability values, respectively. Note that in each of the six classes, a distribution of prosody patterns co-occurs with a single elementary gesture pattern.

scale of $(-2, -1, 0, 1, 2)$, where the scale corresponds to A much better, A better, no preference, B better, and B much better.

The whole test database is manually partitioned into meaningful 15 segments, where each segment is approximately 12 secs. For each evaluation eight segments out of 15 are randomly selected. Three sets of A-B comparison pairs, each including these eight segments, are considered for the speech-driven talking head animations using the original and two sets of synthesized Euler angles. Furthermore, three random startup A-B test pairs and another three test pairs with identical synthesis algorithms are also included to the subjective test set. Hence, the total number of A-B pairs in a test is 30. Apart from the three random start-up A-B pairs, all the pairs are randomized across conditions and pairwise. The subjective tests are performed over 15 subjects. The average preference scores for the three

comparison sets are presented in Table 2. Note that the scores of the three random start-up pairs are ignored in calculating the final preference scores. As expected, the subjective A-B comparisons indicate a preference for the talking head animations with the original Euler angles. On the other hand, animations synthesized with the proposed joint gesture-prosody correlation model Γ_{gp} are preferred over animations generated using the IOHMM correlation model with an average preference score of -0.56 . Also note that the preference for the animations with the original Euler angles is stronger in the case of IOHMM driven animations as compared to the proposed Γ_{gp} driven animations. This is expected, since the output and transition probabilities in the IOHMM structure are conditional directly on the input sequence, whereas in the joint multistream HMM, the output gesture patterns are affected by the states only and not directly by the input. Hence, the

TABLE 1

The Mean and Standard Deviation of the Distance Measures between the Original and the Two Sets of Synthesized Euler Angles from the Proposed Γ_{gp} and IOHMM Models

Model	Γ_{gp}	IOHMM
$[\mu_{\epsilon_n}, \sigma_{\epsilon_n}]$	[0.817896, 0.010981]	[0.890652, 0.012861]
$[\mu_{\epsilon_m}, \sigma_{\epsilon_m}]$	[1.946043, 0.020871]	[2.290755, 0.073021]
$[\mu_{\epsilon_e}, \sigma_{\epsilon_e}]$	[13.694374, 0.158944]	[16.427827, 0.525493]

TABLE 2
The Subjective A-B Comparison Results

A-B pair	Preference Score
Original - Γ_{gp}	-0.23
Original - IOHMM	-0.83
Γ_{gp} - IOHMM	-0.56
Identical pairs	0.04

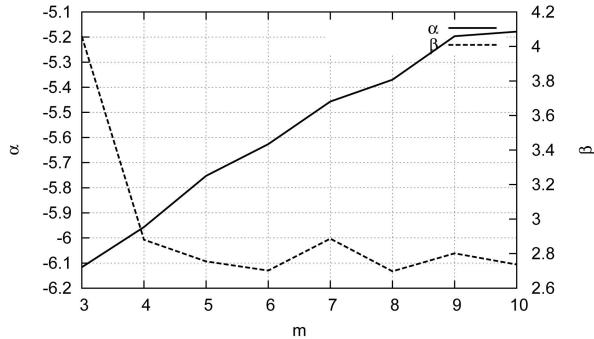


Fig. 12. The α and β fitness measures for varying number of head gesture patterns of the second speaker.

use of parallel multistream HMM in the second stage is more robust to any noise in the input stream.

Samples of the audio-visual sequences for the prosody-driven talking head animations are available online [33]. These samples are selected to demonstrate three possible related applications. The first one is the speaker dependent prosody-driven gesture synthesis application, where gesture-prosody correlation model of a speaker is used to animate the same speaker with her/his speech. The second application is a head gesture transplant, where gesture-prosody correlation model of speaker A is used to animate speaker B from speaker A 's speech. Furthermore, the prosody transplant is considered as the third application, where gesture-prosody correlation model of speaker A is used to animate speaker A from speaker B 's speech. In the demonstration of the prosody transplant, we used speech input from audio-book recordings in English, where the gesture-prosody correlation model is performed over the story telling recordings in Turkish. Although one should expect differences in prosody patterns across different languages, the naturalness of the animations is observed to be acceptable. We also note that the talking speed of these two speakers are different, where the native Turkish speaker has a faster rate than the native English speaker. As expected from the proposed correlation model, we observe slower head gesture animations for the native English speaker.

5.4 Speaker Dependency

The proposed analysis method is capable of providing personalized elementary head gesture and prosody patterns and a personalized prosody to gesture mapping model. To demonstrate this, we have repeated the experiments with a second speaker. The second speaker is also instructed to tell the same four stories to an audience of children. The system is then trained using the recordings of the speaker. At the end of the two-stage analysis for modeling gesture-prosody correlation, we have observed that the resulting elementary patterns for both prosody and head gestures significantly differ from those of the first speaker.

In order to set the number of gesture patterns, M_{Λ_g} , and prosody patterns, M_{Λ_p} , the two fitness measures α and β are calculated for varying number of gesture and prosody patterns, respectively, which are plotted in Figs. 12 and 13. In both plots, the probability of model match, α , increases, and the statistical separation between patterns, β , decreases with an increasing number of patterns. The numbers of gesture and prosody patterns are selected in the vicinity

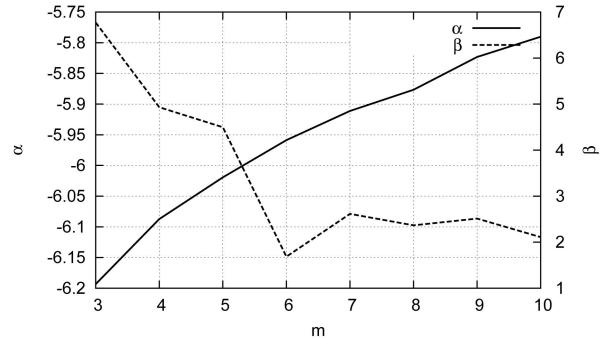


Fig. 13. The α and β fitness measures for varying number of prosody patterns of the second speaker.

where α and β curves intersect, as $M_{\Lambda_g} = 4$ and $M_{\Lambda_p} = 5$. The mean Euler angle vectors and typical thumbnails for the four gesture patterns are plotted in Fig. 14. Similarly, the means and standard deviations of the normalized intensity and pitch frequency trajectories for the five prosody patterns are plotted in Fig. 15. Note that the elementary gesture patterns for the second speaker is distinctively different than the ones for the first speaker (see Fig. 7 for comparison). Sample video streams of the typical elementary gesture patterns are available online in [33]. The elementary prosody patterns also differ for the second speaker. Three of the prosody patterns are falling boundary tones (L percent), and the other one is a rising boundary tone (H percent) for the second speaker.

At the second stage analysis, the joint gesture-prosody pattern label stream is segmented in an unsupervised manner using the discrete multistream HMM structure Γ_{gp} . The two fitness measures for Γ_{gp} are plotted in Fig. 16. In the unsupervised segmentation, the number of joint gesture-prosody patterns is set to $M_{\Gamma_{gp}} = 5$. As for the demonstration of synthesis results, to better emphasize speaker dependency, we have used the same audio-book recordings in English and the same face model to derive the head gesture animations for the two different speakers. A sample animation video is available online in [33], where the video stream resulting from the second speaker's gesture-prosody correlation model, is presented in parallel with the video stream generated from the first speaker's model for visual evaluation of the speaker dependency performance.

6 CONCLUSIONS

We proposed a new two-stage joint head gesture and speech prosody analysis framework. In the first stage of the analysis, elementary gesture and prosody patterns are extracted using unsupervised segmentation for a speaker, and in the second stage, a correlation model between head gesture and prosody patterns is developed. The proposed two-stage analysis framework offers the following advantages:

1. Meaningful elementary gesture and prosody patterns are defined for a speaker at the first stage.
2. A mapping between these elementary prosody and head gesture patterns is obtained with unsupervised segmentation of the joint gesture-prosody label stream.

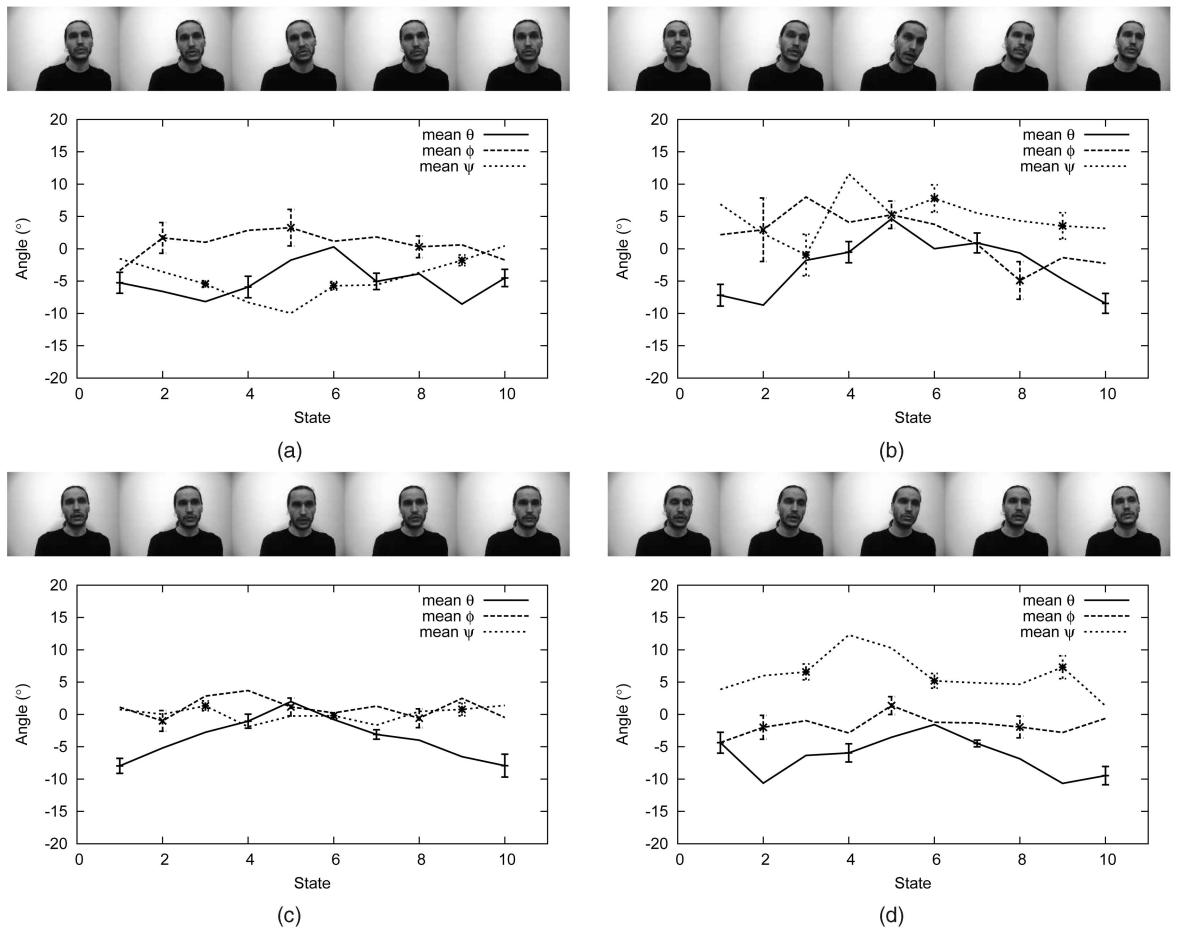


Fig. 14. The mean Euler angles with standard deviations and typical thumbnails from the second speaker for the four gesture patterns: (a) Tilt Left, (b) Nod with Tilt Right, (c) Nod, and (d) Tilt Right.

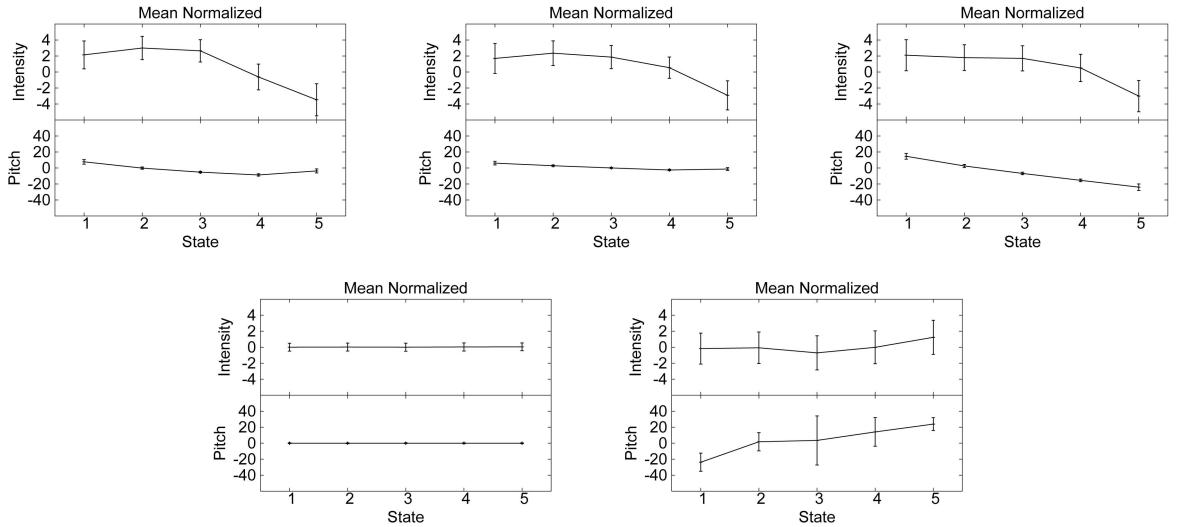


Fig. 15. The means and standard deviations of the normalized intensity (dB) and pitch frequency (Hz) trajectories for the five prosody patterns of the second speaker.

3. The HMM-based analysis and synthesis yields flexibility in modeling structural and durational variations within gestural and prosodic patterns.
4. Automatic generation of the elementary gesture patterns produces natural looking prosody-driven head gesture synthesis.

In addition to successful demonstration of speaker dependent speech-driven head gesture synthesis system, different applications, such as head gesture transplant and prosody transplant, are also demonstrated. After extracting a gesture-prosody correlation model for speaker A, head gesture transplant animates speaker B from speaker A's

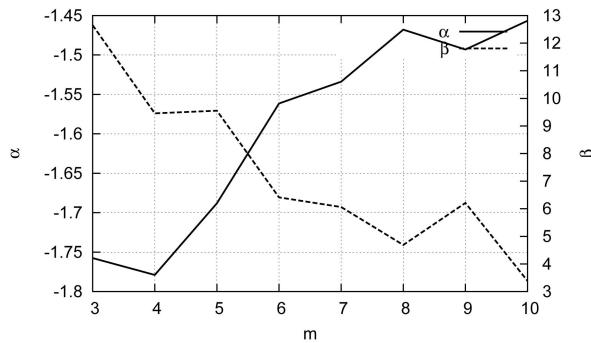


Fig. 16. The α and β fitness measures and the normalized euclidean distance measure ϵ_n for varying number of joint gesture-prosody label patterns for the second speaker.

speech, and prosody transplant animates speaker *A* from speaker *B*'s speech. In the prosody transplant demonstration, gesture-prosody correlation model is trained with audio-visual recordings in Turkish, and prosody-driven gesture synthesis is performed with speech input recordings in English. The naturalness of the prosody transplant is found to be acceptable. Also in this demonstration, we observe slower head gesture animations for the native English speaker whose talking speed is slower.

The proposed HMM based two-stage head gesture and speech prosody analysis system can be utilized to model the correlation between any other loosely correlated modalities, such as facial expressions and speech prosody, arm gestures and speech semantics, etc. Furthermore, the proposed speaker dependent speech-driven head gesture synthesis system can be tailored to model a speaker's emotion and mood. We also note that prosody patterns obtained using the proposed stage I analysis over a multispeaker phonetically rich Turkish (or any other language) training database can be used to define a complete ToBI-like prosodic transcription convention for Turkish (or any other language) intonation.

APPENDIX

RIGID MOTION PARAMETER ESTIMATION BY CONSTRAINED OPTIMIZATION

This appendix summarizes the method used for estimating the rotation matrix, R , and translation vector, t , which describe the rigid motion between the world point coordinate matrices W_k and W_{k_r} (see Section 2.1).

Let \mathbf{m}_k denote the mean of the column vectors in the matrix W_k such that

$$\mathbf{m}_k = \frac{1}{N} \sum_{i=1}^N \mathbf{w}_k^i \quad (22)$$

and \mathbf{m}_{k_r} be defined similarly. Then, the translation t between W_k and W_{k_r} is given by

$$\mathbf{t} = \mathbf{m}_k - \mathbf{m}_{k_r}. \quad (23)$$

Furthermore, let W'_{k_r} and W'_k represent the mean-removed coordinate matrices such that

$$W'_k = W_k - \mathbf{m}_k \mathbf{1}^T, \quad \text{and} \quad W'_{k_r} = W_{k_r} - \mathbf{m}_{k_r} \mathbf{1}^T. \quad (24)$$

Then, the rotation matrix R can be found by minimizing the cost function:

$$f(R) = \|E\|_F^2 = \text{tr}(EE^T), \quad (25)$$

where $\|\cdot\|_F$ and $\text{tr}(\cdot)$ denote the Frobenius-norm and the matrix trace, respectively, and

$$E = W'_k - RW'_{k_r}. \quad (26)$$

The minimization of the cost function $f(R)$, $f : \mathbb{R}^{3 \times 3} \rightarrow \mathbb{R}$, is a nonlinear optimization problem, under the unitary constraint $R^T R = I$, which can be solved by the algorithm proposed in [34], where Manton describes a modified Newton method for optimization on the complex Stiefel manifold that defines the space related with the unitary constraint.

We simplified this method to minimize the cost function $f(R)$ for a square and real matrix R subject to the constraint $R^T R = I$ as follows:

1. Choose the initial R such that $R^T R = I$.

For small rotations, R can be approximated in terms of a parameter vector $\mathbf{u} = [u_x, u_y, u_z]^T$ such that [35]

$$R \approx I + S = I + \begin{bmatrix} 0 & -u_z & u_y \\ u_z & 0 & -u_x \\ -u_y & u_x & 0 \end{bmatrix}. \quad (27)$$

Equating the residual defined in (26) to zero, we obtain the following equation to solve for S :

$$W'_k - W'_{k_r} = SW'_{k_r}, \quad (28)$$

which can be expressed in terms of \mathbf{u} as

$$\text{vec}(W'_k - W'_{k_r}) = Ku = \begin{bmatrix} K_1 \\ \vdots \\ K_N \end{bmatrix} \quad (29)$$

$$K_n = \begin{bmatrix} 0 & Z_n & -Y_n \\ -Z_n & 0 & X_n \\ -X_n & Y_n & 0 \end{bmatrix},$$

where operator $\text{vec}(\cdot)$ obtains a column vector by stacking the columns of the operand matrix, and each 3×3 submatrix K_n is constructed using the n th point (X_n, Y_n, Z_n) from W'_{k_r} . The least squares solution of (29) can then be used to find \mathbf{u} and to construct S . The initial guess for R can finally be obtained by projection onto the unitary space $R = \pi(I + S)$ (described in step 5 below).

2. Compute the derivative D_R and the Hessian H_R of f given by

$$D_R = -2EW'^T_{k_r}, \quad (30)$$

$$H_R = 2((W'_{k_r} W'^T_{k_r}) \otimes I_{3 \times 3}), \quad (31)$$

where \otimes denotes Kronecker product.

3. If $\sqrt{\text{tr}(D_R^T D_R - R^T D_R R^T D_R)} < \epsilon$, then stop.

4. Compute the Newton step size $Z := Z^{(cp)}$.

The Newton step size is defined as the value of Z , $Z \in \mathbb{R}^{3 \times 3}$, confined to the tangent space V , at which the quadratic approximation $g(Z)$ has its critical point:

$$g(Z) \approx f(R) + \text{tr}(Z^T D) + (1/2)\text{vec}(Z)^T H \text{vec}(Z), \quad (32)$$

where

$$D = D_R, \quad H = H_R - (1/2)[(R^T D_R + D_R^T R)^T \otimes I]. \quad (33)$$

The tangent space V is defined as a subset of $\mathbb{R}^{3 \times 3}$ such that $Z = RA$, where A is skew symmetric. The critical point $Z^{(cp)} \in V$, that is, the Newton step size, satisfies the following linear constraint:

$$\text{tr}(Z^T D) + [\text{vec}(Z)^T H] \text{vec}(Z^{(cp)}) = 0. \quad (34)$$

By writing Z as $Z = \sum_{i=1}^3 \alpha_i RA_i$, where A_i ($i = 1, 2, 3$) is an arbitrary basis for skew-symmetric matrix, the critical point $Z^{(cp)}$ can be found by solving the following linear equation for α_i ($i = 1, 2, 3$):

$$\frac{\partial g(Z)}{\partial \alpha_i} = \text{tr}((RA_i)^T D) + \text{vec}(RA_i)^T H \text{vec}(Z) = 0. \quad (35)$$

Note that the above equation can be put into a matrix form:

$$K^T H K \alpha = K^T \text{vec}(D), \quad (36)$$

where

$$K = (I \otimes R)[\text{vec}(A_1), \text{vec}(A_2), \text{vec}(A_3)] \quad (37)$$

$$\alpha = [\alpha_1, \alpha_2, \alpha_3]^T.$$

5. Set $R' := \pi(R + Z)$.

The projection $\pi(R)$, $\pi : \mathbb{R}^{3 \times 3} \rightarrow St$, onto the Stiefel manifold, $St = \{R \in \mathbb{R}^{3 \times 3} : R^T R = I\}$, is defined as

$$\pi(R) = \underset{Q \in St}{\text{argmin}} \|R - Q\|^2. \quad (38)$$

If the singular value decomposition of R is $U\Sigma V^T$, then the projection is simply given by [34]:

$$\pi(R) = UI_{3 \times 3}V^T. \quad (39)$$

6. If $f(R) \leq f(R')$, then abort.

7. Set $R := R'$. Go to step 2.

ACKNOWLEDGMENTS

The authors would like to thank *Momentum Inc.* for making the talking head avatar available and for their collaboration to build the MVGL-MASAL gesture-speech database. This work has been supported by the European FP6 Network of Excellence SIMILAR (<http://www.similar.cc>)

and by TUBITAK under project EEEAG-106E201 and COST2102 action). The work of M.E. Sargin was done while he worked at Koc University.

REFERENCES

- [1] T. Chen, "Audiovisual Speech Processing," *IEEE Signal Processing Magazine*, vol. 18, pp. 9-21, 2001.
- [2] S. Morishima, K. Aizawa, and H. Harashima, "An Intelligent Facial Image Coding Driven by Speech and Phoneme," *Proc. Int'l Conf. Acoustics, Speech, and Signal Processing (ICASSP '89)*, pp. 1795-1798, 1989.
- [3] C. Bregler, M. Covell, and M. Slaney, "Video Rewrite: Driving Visual Speech with Audio," *Proc. ACM SIGGRAPH '97*, pp. 353-360, 1997.
- [4] F. Huang and T. Chen, "Real-Time Lip-Synch Face Animation Driven by Human Voice," *Proc. IEEE Second Workshop Multimedia Signal Processing*, pp. 352-357, 1998.
- [5] E. Yamamoto, S. Nakamura, and K. ShiKano, "Lip Movement Synthesis from Speech Based on Hidden Markov Models," *Speech Comm.*, pp. 105-115, 1998.
- [6] M. Brand, "Voice Puppetry," *Proc. 26th Ann. Conf. Computer Graphics and Interactive Techniques*, pp. 21-28, 1999.
- [7] P.S. Alekseev and A.K. Katsaggelos, "Speech-to-Video Synthesis Using Facial Animation Parameters," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 14, no. 5, pp. 682-692, 2004.
- [8] Y. Li and H.-Y. Shum, "Learning Dynamic Audio-Visual Mapping with Inputoutput Hidden Markov Models," *IEEE Trans. Multimedia*, vol. 8, no. 3, pp. 542-549, 2006.
- [9] J. Xue, J. Borgstrom, J. Jiang, L. Bernstein, and A. Alwan, "Acoustically-Driven Talking Face Synthesis Using Dynamic Bayesian Networks," *Proc. Int'l Conf. Multimedia and Expo (ICME '06)*, pp. 1165-1168, 2006.
- [10] L. Valbonesi, R. Ansari, D. McNeill, F. Quek, S. Duncan, K.E. McCullough, and R. Bryll, "Multimodal Signal Analysis of Prosody and Hand Motion: Temporal Correlation of Speech and Gestures," *Proc. European Signal Processing Conf. (EUSIPCO '02)*, vol. 1, pp. 75-78, 2002.
- [11] K. Munhall, J.A. Jones, D.E. Callan, T. Kuratake, and E. Vatikiotis-Bateson, "Visual Prosody and Speech Intelligibility: Head Movement Improves Auditory Speech Perception," *Psychological Science*, vol. 15, no. 2, pp. 133-137, 2004.
- [12] F. Quek, D. McNeill, R. Ansari, X. Ma, R. Bryll, S. Duncan, and K. McCullough, "Gesture Cues for Conversational Interaction in Monocular Video," *Proc. Int'l Workshop Recognition, Analysis, and Tracking of Faces and Gestures in Real-Time Systems*, pp. 64-69, 1999.
- [13] T. Kuratake, K.G. Munhall, P.E. Rubin, E. Vatikiotis-Bateson, and H. Yehia, "Audio-Visual Synthesis of Talking Faces from Speech Production Correlates," *Proc. Sixth European Conf. Speech Comm. and Technology (EUROSPEECH '99)*, pp. 1279-1282, 1999.
- [14] H.P. Graf, E. Cosatto, V. Strom, and F.J. Huang, "Visual Prosody: Facial Movements Accompanying Speech," *Proc. IEEE Int'l Conf. Automatic Face and Gesture Recognition*, pp. 381-386, 2002.
- [15] E. Chuang and C. Bregler, "Mood Swings: Expressive Speech Animation," *ACM Trans. Graphics*, vol. 24, no. 2, pp. 331-347, 2005.
- [16] Z. Deng, C. Busso, S. Narayanan, and U. Neumann, "Audio-Based Head Motion Synthesis for Avatar-Based Telepresence Systems," *Proc. ACM SIGMM Workshop Effective Telepresence (ETP '04)*, pp. 24-30, 2004.
- [17] M.E. Sargin, F. Ofli, Y. Yasinnik, O. Aran, A. Karpov, S. Wilson, E. Erzin, Y. Yemez, and A.M. Tekalp, "Gesture-Speech Correlation Analysis and Speech Driven Gesture Synthesis," *Proc. Int'l Conf. Multimedia and Expo (ICME '06)*, 2006.
- [18] M. Naphade and T. Huang, "Discovering Recurrent Events in Video Using Unsupervised Methods," *Proc. Int'l Conf. Image Processing (ICIP '02)*, 2, pp. 13-16, 2002.
- [19] P. Viola and M. Jones, "Rapid Object Detection Using a Boosted Cascade of Simple Features," *Proc. IEEE Computer Vision and Pattern Recognition (CVPR '01)*, pp. 511-518, 2001.
- [20] R. Lienhart and J. Maydt, "An Extended Set of Haar-Like Features for Rapid Object Detection," *Proc. Int'l Conf. Image Processing (ICIP '02)*, vol. 1, pp. 900-903, 2002.
- [21] J.Y. Bouguet, *Pyramidal Implementation of the Lucas Kanade Feature Tracker*—description of the Algorithm, OpenCV Documents, Intel Corp., Microprocessor Research Labs, 1999.

- [22] M. Brown, D. Burschka, and G. Hager, "Advances in Computational Stereo," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no. 8, pp. 993-1008, Aug. 2003.
- [23] P. Fua, "Combining Stereo and Monocular Information to Compute Dense Depth Maps That Preserve Depth Discontinuities," *Proc. 12th Int'l Joint Conf. Artificial Intelligence*, pp. 1292-1298, 1997.
- [24] D. Varshalovich, A. Moskalev, and V. Khersonskii, "Description of Rotation in Terms of the Euler Angles," *Quantum Theory of Angular Momentum*, World Scientific, 1988.
- [25] K. Shoemake, "Animating Rotation with Quaternion Curves," *Proc. 12th Ann. Conf. Computer Graphics and Interactive Techniques*, pp. 245-254, 1985.
- [26] P. Boersma, "Accurate Short-Term Analysis of the Fundamental Frequency and the Harmonics-to-Noise Ratio of a Sampled Sound," *Proc. Inst. Phonetic Sciences*, vol. 17, pp. 97-110, 1993.
- [27] S. Ananthakrishnan and S. Narayanan, "An Automatic Prosody Recognizer Using a Coupled Multi-Stream Acoustic Model and a Syntactic-Prosodic Language Model," *Proc. IEEE Int'l Conf. Acoustics, Speech, and Signal Processing (ICASSP '05)*, vol. 1, 2005.
- [28] Point Grey Research Inc., <http://www.ptgrey.com/>, 2008.
- [29] K. Silverman, M. Beckman, J. Pitrelli, M. Ostendorf, C. Wightman, P. Price, J. Pierrehumbert, and J. Hirschberg, "Tobi: A Standard for Labeling English Prosody," *Proc. Int'l Conf. Spoken Language Processing (ICSLP '92)*, pp. 867-870, 1992.
- [30] Momentum Inc., Speech-Driven Talking Head Avatar, <http://www.momentum-dmt.com/>, 2008.
- [31] Y. Bengio and P. Frasconi, "Input-Output HMMs for Sequence Processing," *IEEE Trans. Neural Networks*, vol. 7, no. 5, pp. 1231-1249, 1996.
- [32] R. Collobert, S. Bengio, and J. Mariethoz, "Torch: A Modular Machine Learning Software Library," *IDIAP Research Report*, vol. 2, p. 46, 2002.
- [33] *Prosody-Driven Head Gesture Animation*, <http://mvgl.ku.edu.tr/prosodygesture/>, 2008.
- [34] J.H. Manton, "Optimisation Algorithms Exploiting Unitary Constraints," *IEEE Trans. Signal Processing*, vol. 50, no. 3, pp. 635-650, Mar. 2002.
- [35] D. Demirdjian and T. Darrell, "Motion Estimation from Disparity Images," *Proc. Eighth IEEE Int'l Conf. Computer Vision*, vol. 1, pp. 213-218, 2001.



Mehmet Emre Sargin received the BSc degree in electrical and electronics engineering from Middle East Technical University, Ankara, Turkey, in 2004, and the MSc degree in electrical and computer engineering from Koc University, Istanbul, in 2006. He is currently working toward the PhD degree in electrical and computer engineering at University of California, Santa Barbara. His research interests include computer vision, pattern recognition, machine learning, and bioimage informatics. He is the second prize winner in the ICASSP '07 student paper contest. He is a student member of the IEEE.



Yücel Yemez received the BS degree from Middle East Technical University, Ankara, Turkey, in 1989, and the MSc and PhD degrees from Bogazici University, Istanbul, respectively, in 1992 and 1997, all in electrical engineering. From 1997 to 2000, he was a postdoctoral researcher in the Image and Signal Processing Department, Télécom Paris (Ecole Nationale Supérieure des Télécommunications). Currently, he is an assistant professor in the Computer Engineering Department, Koc University, Istanbul. His current research interests include computer vision and 3D computer graphics. He is a member of the IEEE.



Engin Erzin received the BSc, MSc, and PhD degrees in electrical engineering from the Bilkent University, Ankara, in 1990, 1992, and 1995, respectively. From 1995 to 1996, he was a postdoctoral Fellow in the Signal Compression Laboratory, University of California, Santa Barbara. He joined Lucent Technologies in September 1996, and he was with Consumer Products for one year as a Member of Technical Staff of the Global Wireless Products Group.

From 1997 to 2001, he was with the Speech and Audio Technology Group, Network Wireless Systems. Since January 2001, he has been with the Electrical and Electronics Engineering Department and Computer Engineering Department, Koç University, Istanbul. His research interests include speech signal processing, pattern recognition, and adaptive signal processing. He is a senior member of the IEEE.



A. Murat Tekalp received the MS and PhD degrees in electrical, computer, and systems engineering from Rensselaer Polytechnic Institute (RPI), Troy, New York, in 1982 and 1984, respectively. He was with the Eastman Kodak Company, Rochester, New York, from December 1984 to June 1987 and with the University of Rochester from July 1987 to June 2005, where he was promoted to a distinguished university professor. Since June 2001, he has been a professor at Koc University, Istanbul. He is serving as the chairman of the Electronics and Informatics Group, Turkish Science and Technology Foundation (TUBITAK) and as an independent expert to review projects for the European Commission. His research interests include digital image and video processing, including video compression and streaming, motion-compensated video filtering for high-resolution, video segmentation, content-based video analysis and summarization, 3DTV/video processing and compression, multicamera surveillance video processing, and protection of digital content. He is the author of the book *Digital Video Processing* (Prentice Hall, 1995). He served as an associate editor for the *IEEE Transactions on Signal Processing* from 1990 to 1992 and the *IEEE Transactions on Image Processing* from 1994 to 1996, and the Kluwer journal *Multidimensional Systems and Signal Processing* from 1994 to 2002. He was an area editor for the Academic Press journal *Graphical Models and Image Processing* from 1995 to 1998. He was also on the editorial board of the Academic Press journal *Visual Communication and Image Representation* from 1995 to 2002. Currently, he is the editor in chief of the EURASIP journal *Signal Processing: Image Communication* (Elsevier). He holds seven US patents. His group contributed technology to the ISO/IEC MPEG-4 and MPEG-7 standards. He was named as distinguished lecturer by the IEEE Signal Processing Society in 1998 and awarded a Fulbright Senior Scholarship in 1999. He received the TUBITAK Science Award (highest scientific award in Turkey) in 2004. He chaired the IEEE Signal Processing Society Technical Committee on Image and Multidimensional Signal Processing from January 1996 December 1997. He was appointed as the special sessions chair for the 1995 IEEE International Conference on Image Processing, the technical program cochair for IEEE ICASSP 2000, the general chair of IEEE International Conference on Image Processing (ICIP), Rochester, New York, in 2002, and technical program cochair of EUSIPCO 2005, Antalya, Turkey. He is the founder and first chairman of the Rochester Chapter, IEEE Signal Processing Society. He was elected as the chair of the Rochester Section, IEEE from 1994 to 1995. He is a fellow of the IEEE.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.