

Modeling Individual and Group Actions in Meetings: a Two-Layer HMM Framework

Dong Zhang, Daniel Gatica-Perez, Samy Bengio, Iain McCowan, and Guillaume Lathoud
IDIAP Research Institute, Martigny, Switzerland
{zhang, gatica, bengio, mccowan, lathoud}@idiap.ch

Abstract

We address the problem of recognizing sequences of human interaction patterns in meetings, with the goal of structuring them in semantic terms. The investigated patterns are inherently group-based (defined by the individual activities of meeting participants, and their interplay), and multimodal (as captured by cameras and microphones). By defining a proper set of individual actions, group actions can be modeled as a two-layer process, one that models basic individual activities from low-level audio-visual features, and another one that models the interactions. We propose a two-layer Hidden Markov Model (HMM) framework that implements such concept in a principled manner, and that has advantages over previous works. First, by decomposing the problem hierarchically, learning is performed on low-dimensional observation spaces, which results in simpler models. Second, our framework is easier to interpret, as both individual and group actions have a clear meaning, and thus easier to improve. Third, different HMM models can be used in each layer, to better reflect the nature of each subproblem. Our framework is general and extensible, and we illustrate it with a set of eight group actions, using a public five-hour meeting corpus. Experiments and comparison with a single-layer HMM baseline system show its validity.

1 Introduction

Devising computational frameworks to automatically infer human behavior from sensors constitutes an open problem in many domains. Moving beyond the person-centered paradigm [23], recent work has started to explore multi-person scenarios, where not only individual but also group actions or interactions become relevant [6, 9, 20, 1].

Group activity plays a key role in meetings [24, 16], and this is documented by a significant amount of work in social psychology [14]. Viewed as a whole, a group shares information, engages in discussions, and makes decisions, proceeding through diverse communication phases both in single meetings and during the course of a long-term teamwork [14]. Recognizing group actions is therefore useful for browsing and retrieval purposes [24, 12], e.g., to structure a

meeting into a sequence of high-level items.

Interaction in meetings is inherently group-based [14] and multimodal [10]. In the first place, we can view a meeting as a continuous sequence of mutually exclusive group actions taken from an exhaustive set [12, 4]. Each of these group actions involves multiple simultaneous participants, and is thus implicitly constrained by the actions of the individuals. In the second place, as the principal modality in meetings, speech has recently been studied in the context of interaction modeling [8, 25, 4]. However, work analyzing the benefits of modeling individual and group actions using multiple modalities has been limited [1, 12, 13, 21], despite the fact that actions in meetings, both at the individual (e.g., note-taking or talking), and at the group level (e.g. dictating) are often defined by the joint occurrence of specific audio and visual patterns.

In this paper, we present a two-layer HMM framework for group action recognition in meetings. The fundamental idea is that, by defining an adequate set of individual actions, we can decompose the group action recognition problem into two levels, from individual to group actions. Both layers use ergodic HMMs or extensions. The goal of the lower layer is to recognize individual actions of participants using low-level audio-visual (AV) features. The output of this layer provides the input to the second layer, which models interactions. Individual actions naturally constitute the link between the low-level audio-visual features and high-level group actions. Similar to continuous speech recognition, we perform group action recognition directly on the data sequence, deriving the segmentation of group actions in the process. Our approach is general, extensible, and brings improvement over previous work, which reflects on the results obtained on a public meeting corpus, for a set of eight group actions based on multimodal turn-taking patterns.

The paper is organized as follows. Section 2 reviews related work. Section 3 introduces our approach. Section 4 and Section 5 describe the meeting data, and the feature extraction process. Experiments and discussion are presented in Section 6. Conclusions are drawn in Section 7.

2 Related Work

Current approaches to automatic activity recognition define models for specific activities that suit the goal in a particular domain, and use statistical methods for recognition. Predominately, the recognition of individual actions [23], or interaction involving few people [20, 9] has been investigated using visual features [23, 20, 9], although some work on the speech community can also be categorized as interaction recognition [8, 25]. To our knowledge however, little work has been conducted on recognition of group-based, multi-modal actions from multiple audio-visual streams captured by cameras and microphones [1, 12, 13]. [1] described automatic discovery of “influence” in a lounge room where people played interactive debating games. Our previous work presented different sequence models to recognize turn-taking patterns in a formal meeting room scenario, where people discuss around a table and use a white-board and a projector screen [13, 12].

Regarding statistical models, most of the existing work has used Hidden Markov Models (HMMs) [22, 23] and extensions, including coupled HMMs, input-output HMMs, multi-stream HMMs, and asynchronous HMMs . (see [18] for a recent review of models). Although the basic HMM, a discrete state-space model with an efficient learning algorithm, works well for temporally correlated sequential data, it is challenged by a large number of parameters, and the risk of over-fitting when learned from limited data [19]. This situation might occur in the case of multimodal group action recognition where, in the simplest case, possibly large vectors of AV features from each participant are concatenated to define the observation space [12, 13].

The above problem is general, and has been addressed using hierarchical representations [26, 4, 19]. In [26], an approach for unsupervised discovery of multilevel video structures using hierarchical HMMs was proposed, in the context of sports videos. In this model, the higher-level structure elements usually correspond to semantic events, while the lower-level states represents variations occurring within the same event. In [4], two methods for meeting structuring from audio were presented, using multilevel Dynamic Bayesian Networks (DBNs). However, in both [26, 4], the low-level actions have no obvious interpretation, and the number of low-level actions is a model parameter learned during training, or set by hand, which makes the structure of the models difficult to interpret. In [19], layered HMMs were proposed to model multimodal office activities at various time granularities. The lowest layer captured video, audio, keyboard and mouse activity features; the middle layer classifies AV features into basic events like “speech”, “music”, “one person”, “nobody”, etc. Finally, the highest layer uses outputs of previous layers to recognize office activities with longer temporal extent.

The solution we present to the problem of group action recognition is novel. Unlike previous work [12, 13], our framework explicitly models actions at different semantic levels (from individual to group level) at the same time scale. This layered structure coincides with the structure of meetings as modeled in social psychology, that is, that meetings comprise individual actions and interactions [14]. Our goal -modeling group activity- is also different than that of [19]. Since the two HMM layers are trained independently, our framework is easy to interpret and improve at different levels, and we have studied a number of models suitable for multimodal data. Overall, our work has a number of advantages, as described in the next section.

3 Group Action Recognition

In this section, we first introduce our computational framework. We then apply it to a specific set of individual and group actions. Finally, we describe some specific details.

3.1 Framework Overview

In our framework, we distinguish group actions (which belong to the whole set of participants) from individual actions (belonging to specific persons). Our ultimate goal is the recognition of group activity, and so individual actions should act as the bridge between group actions and low-level features, thus decomposing the problem in stages. The definition of both action sets is thus clearly intertwined.

Let *I-HMM* denote the lower recognition layer (individual action), and *G-HMM* denote the upper layer (group action). *I-HMM* receives as input AV features extracted from each participant, and outputs recognition results, either as soft or hard decisions (see section 3.3). In turn, *G-HMM* receives as input the output from *I-HMM*, and a set of *group features*, directly extracted from the raw streams, which are not associated to any particular individual. In our framework, each layer is trained independently, and can be substituted by any of the HMM variants that might capture better the characteristics of the data, more specifically asynchrony [2], or different noise conditions [5] between the audio and visual streams. Our approach is summarized in Fig.1.

Compared with a single-layer HMM, our approach has the following advantages: (1) a single-layer HMM is defined on a possibly large observation space, which might face the problem of over-fitting with limited training data. It is important to notice that the amount of training data becomes an issue in meetings where data labeling is not a cheap task. In contrast, the layers in our approach are defined over small-dimensional observation spaces, resulting in more stable performance in cases of limited amount of training data. (2) The *I-HMMs* are person-independent, and in practice can be trained with much more data from different persons, as each meeting provides multiple individual streams of training data. Better generalization performance

-
1. (Audio-Visual Feature Extraction)
 - 1-1. extract individual-level AV features
 - 1-2. extract group-level AV features
 2. (Individual Action Recognition)
 - 2-1. given individual features for each person, train *I-HMM* selecting best model by cross-validation
 - 2-2. output individual action recognition results
 3. (Group Action Recognition)
 - 3-1. construct a feature space by concatenating individual action results and group-level features
 - 3-2. train *G-HMM* selecting best model by cross-validation
 - 3-3. output group action recognition results
-

Figure 1: Two-layer HMM framework

can then be expected. (3) The *G-HMMs* are less sensitive to variations in the low-level features because their observations are the outputs of the individual action recognizers, which are expected to be well trained. (4) The two layers are trained independently. Thus, we can explore different HMM combination systems. In particular, we can replace the baseline *I-HMMs* with models that are more suitable for multi-modal asynchronous data sequences, with the goal of gaining understanding of the nature of the data (see Section 3.3). The framework is thus easier to interpret and to improve separately at each level. (5) Our framework is general and extensible to recognize new group actions defined in the future.

3.2 Definition of Actions

As an implementation of the proposed framework, we define a set of group actions and individual actions in this section. On one hand, a group of $N_G = 8$ group actions are defined based on multi-modal turn-taking patterns. The list is defined in Table 1. These group actions are multimodal, and commonly found in meetings. For modeling purposes, they are assumed to define a partition (i.e., the action set is non-overlapping and exhaustive). This set is richer compared to the one that we defined in [13], as it includes simultaneous occurrence of actions, like “*monologue+note-taking*” which could occur during real situations, like dictating or minute-taking.

On the other hand, we define a small set of $N_I = 3$ multimodal individual actions which, as stated earlier, will help bridge the gap between group actions and low-level AV features. The list appears in Table 2. While the list of potentially interesting individual actions in meetings is large, our ultimate goal is recognition of the group-level actions defined in Table 1. It is interesting to note that, although at first glance one would not think of “*speaking*” or “*writing*” as multimodal, joint sound and visual patterns do occur in these cases and are useful in recognition, as the results in later sections confirm.

Finally, meeting rooms can be equipped with white-

Table 1: Description of group actions

Action	Description
Discussion	most participants engaged in conversations
Monologue	one participant speaking continuously without interruption
Monologue+Note-taking	one participant speaking continuously others taking notes.
Note-taking	most participants taking notes
Presentation	one participant presenting using the projector screen
Presentation+Note-taking	one participant presenting using projector screen, others taking notes
White-board	one participant speaking using the white-board
White-board+Note-taking	one participant speaking using white-board, others taking notes

Table 2: Description of individual actions

Action	Description
Speaking	one participant speaking
Writing	one participant taking notes
Idle	one participant neither speaking nor writing

boards or projector screens which are shared by the group. Extracting features from these group devices also helps recognize group actions. They constitute the group features described in the previous subsection. Their detailed description will be presented in section 5.

The logical relations between individual actions, group actions, and group features are summarized in Table 3. The group actions can be seen as combinations of individual actions plus states of group devices. For example, “*presentation + note-taking*” can be decomposed into “*speaking*” by one individual, with more than one “*writing*” participant, while the group device of *projector screen* is in use. Needless to say, our approach is not rule-based, but Table 3 is useful to conceptually relate the two layers.

3.3 Implementing the Two-layer Framework

In this section, we present some details about the architecture of our framework.

In the first place, we investigate three models for the lower-layer *I-HMM*, each of which attempts to model specific properties of the data (for space reasons, we refer to each model only briefly, please refer to the original references for details):

Early Integration, where a basic HMM [22] is trained on combined AV features. This method involves aligning and synchronizing AV features to form one concatenated set of features which is then treated as a single stream of data.

Audio-Visual Multi-Stream, which combines the audio-only and visual-only streams. Each stream is modeled independently. The final classification is based on the fusion of the outputs of both modalities by estimating their joint

Table 3: Relationships between group actions, individual actions and group features. Symbol “★” indicates that the white-board or projector screen are in use when the corresponding group action takes place. Symbol “/” indicates that the number of participants for the corresponding action is not certain.

Group Actions	Individual Actions			Group Features	
	speaking	writing	idle	white-board	projector screen
discussion	>2	/	/		
monologue	1	0	/		
monologue+note-taking	1	>=1	/		
note-taking	0	>2	0		
presentation	1	0	/		★
presentation+note-taking	1	>=1	/		★
white-board	1	0	/	★	
white-board+note-taking	1	>=1	/	★	

occurrence [5].

Audio-Visual Asynchronous, which also combines audio-only and visual-only streams, by learning the joint distribution of pairs of sequences when these sequences are not synchronized and are not of the same length or rate [2].

In the second place, a mechanism to link the two HMM layers has to be specified. There are two approaches to do so, based on different *I-HMM* outputs. Let $a^t = (a_1^t, \dots, a_{N_I}^t) \in \mathbb{R}^{N_I}$ denote a vector in a continuous space of dimension equal to the number of individual actions, which indicates the degree of confidence in the recognition of each individual action at time t for a sequence $x_1^t = x_1, x_2, \dots, x_t$. In the first approach, the individual action model with the highest probability outputs a value of 1, while all other models output a zero value. The vector a^t generated in this way is used as input to *G-HMM*. We refer to it as *hard decision*. The second approach directly outputs the probability p_k^t for each individual action model $M_k, k = 1, \dots, N_I$, as input feature vector to *G-HMM*, $a_k^t = p_k^t$ for all k . We refer to it as *soft decision*.

In soft decision, the probability p_k^t of model M_k given a sequence x_1^t , is computed in the following way. Let us define the forward variable $\alpha(i, t) \stackrel{\text{def}}{=} P(x_1^t, q_t = i)$, which is the probability of having generated the sequence x_1^t and being in the state i at time t in standard Baum-Welch algorithm [22]. For multi-stream HMM and asynchronous HMM, $\alpha(i, t)$ can be replaced by corresponding factor [5, 2]. Given that the probabilities of all states sum up to one, $\sum_{j=1}^{N_S} P(q_t = j) = 1$ (N_S is the number of all states for all models), the probability $P(q_t = i | x_1^t)$ of state i given a sequence x_1^t is:

$$P(q_t = i | x_1^t) = \frac{P(q_t = i, x_1^t)}{P(x_1^t)} = \frac{P(q_t = i, x_1^t)}{\sum_{j=1}^{N_S} P(q_t = j, x_1^t)} \quad (1)$$

$$= \frac{\alpha(i, t)}{\sum_{j=1}^{N_S} \alpha(j, t)}. \quad (2)$$

With this, the probability p_k^t of model M_k given a sequence

x_1^t is then computed as:

$$p_k^t = \sum_{i \in M_k} P(q_t = i | x_1^t) = \sum_{i \in M_k} \frac{\alpha(i, t)}{\sum_{j=1}^{N_S} \alpha(j, t)} \quad (3)$$

where i is the state in model M_k , which is a subset of the states of all models, and N_S is the total number of all states. The probability p_k^t of model M_k is the sum of the probabilities of all states in model M_k . We concatenate the individual recognition vectors from all participants, together with the group-level features, into a $(N_I \times N_P + N_{GF})$ -dimensional vector (where N_P is the number of participants, and N_{GF} is the dimension of the group features) as observations to *G-HMM* for group action recognition.

4 Meeting Database

We used the public meeting corpus we first described in [12], which was collected in a meeting room equipped with synchronized multi-channel audio and video recorders¹. The sensors include three fixed cameras and twelve microphones [15]. Two cameras have an upper-body, frontal view of two participants including part of the table. A third wide-view camera captures the projector screen and white-board. The multi-camera meeting room and visual feature extraction is shown in Fig.2. Audio was recorded using lapel microphones for all participants, and an eight-microphone array place in the center of the table. The corpus consists of 59 short meetings at five-minute average duration, with four participants per meeting. The group action structure was scripted before recording, so part of the group actions labels we define were already available as part of the public corpus. However, we needed to relabel the rest of the group actions (e.g. *monologues* into either *monologues* or *monologues+note-taking*), and to label the entire corpus in terms of individual actions. All ground-truth was produced using *Anvil*, a publicly available video annotation tool².

¹<http://mmm.idiap.ch/>

²<http://www.dfki.de/kipp/anvil/>

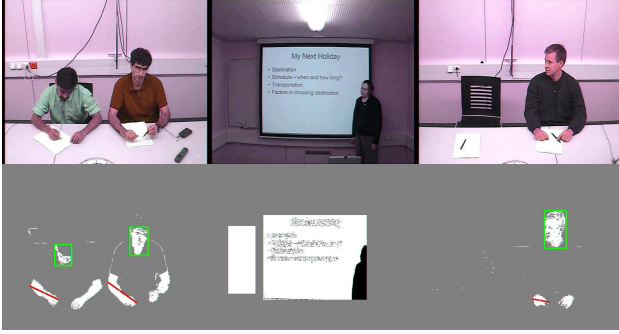


Figure 2: Multi-camera meeting room and visual feature extraction

5 Multi-Modal Feature Extraction

In this section, we describe the process to extract the two types of AV features used in this work: person-specific AV features and group-level AV features. The former are extracted from individual participants. The latter are extracted from the whiteboard and projector screen regions.

5.1 Person-Specific AV Features

Person-specific visual features were extracted from the cameras that have a close view of the participants. Person-specific audio features were extracted from the lapel microphones attached to each person, and from the microphone array. The complete set of features is listed in Table 4.

Person-specific visual features. For each video frame, the raw image is converted to a skin-color likelihood image, using a 5-component skin-color Gaussian mixture model (GMM). We use the chromatic color space, known to be less variant to the skin color of different people [27]. The chromatic colors are defined by a normalization process: $r = \frac{R}{R+G+B}$, $g = \frac{G}{R+G+B}$. Skin pixels were then classified based on thresholding of the skin likelihood. A morphological postprocessing step was performed to remove noise. The skin-color likelihood image is the input to a connected-component algorithm (flood filling) that extracts blobs. All blobs whose areas are smaller than a given threshold were removed. We use 2-D blob features to represent each participant in the meeting, assuming that the extracted blobs correspond to human faces and hands. First, we use a multi-view face detector to verify blobs corresponding to the face. The blob with the highest confidence output by the face detector is recognized as the face. Among the remaining blobs, the one that has the rightmost centroid horizontal position is identified as the right hand (we only extracted features from the right hands since the participants in the corpus are predominately right-handed). For each person, the detected face blob is represented by its vertical centroid position and eccentricity [23]. The hand blob is represented by its horizontal centroid position, eccentricity, and angle. Additionally, the motion magnitude for head and right hand

Table 4: Audio-visual feature list

		Description
		Person-Specific Features
speech relative pitch		
speech energy		
speech rate		
Visual	head vertical centroid	
	head eccentricity	
	right hand horizontal centroid	
	right hand angle	
	right hand eccentricity	
	head and hand motion	
Group Features	Audio	SRP-PHAT from white-board
		SRP-PHAT from projector screen
	Visual	mean difference from white-board
		mean difference from projector screen

are also extracted and summed into one single feature.

Person-specific audio features. Using the microphone array and the lapels, we extracted two types of person-specific audio features. On one hand, speech activity was estimated at four seated locations, from the microphone array waveforms. The seated locations were fixed 3-D vectors measured on-site. The speech activity measure was SRP-PHAT [3], which is a continuous, bounded value that indicates the activity at a particular location. On the other hand, three acoustic features were estimated from each lapel waveform: energy, pitch and speaking rate. We computed these features on speech segments, setting a value of zero on silence segments. Speech segments were detected using the microphone array, because it is well suited for multiparty speech. We used the SIFT algorithm [11] to extract pitch, and a combination of estimators [17] to extract speaking rate.

5.2 Group AV Features

Group AV features were extracted from the white-board and projector screen regions, and are listed in Table 4.

Group visual features. These were extracted from the camera that looks towards the white-board and projector screen area. We first get difference images between a reference background image and the image at each time, in the white-board and projector screen regions (see Fig.2). On these difference images, we use the average intensity over a grid of 16×16 blocks as features.

Group audio features. These are SRP-PHAT features extracted using the microphone array from two locations corresponding to the white-board and projector screen.

6 Experiments

In this section, we first describe the measures used to evaluate our results, and then present results for both individual action recognition and group action recognition.

6.1 Performance Measures

We use the *action error rate (AER)* and the *frame error rate (FER)* as measures to evaluate the results of group action recognition and individual action recognition, respectively.

AER is equivalent to the word error rate widely used in speech recognition, and is defined as the sum of insertion (Ins), deletion (Del), and substitution (Subs) errors, divided by the total number of actions in the ground-truth, $AER = \frac{Subs+Del+Ins}{TotalActions} \times 100\%$. For group action recognition, we have $N_G = 8$ possible actions which in many cases have no clear-cut temporal boundaries. Furthermore, at least five actions occur in each meeting in the corpus. We believe that *AER* is a thus good measure to evaluate group action recognition, as we are more interested in the recognition of the correct action sequence rather than the precise time alignment of the recognized action segments.

However, *AER* overlooks the time alignment between recognized and target action segments. For individual action recognition, there are only $N_I = 3$ possible actions. Furthermore, some streams (participants) in the corpus consist of only two individual actions (e.g., a person who talks only once during the course of a meeting). *AER* might not provide a meaningful assessment in such cases. In this view, we adopt *FER* as the performance measure for individual action recognition. *FER* is defined as one minus the ratio between the number of correctly recognized frames and the number of total frames, $FER = (1 - \frac{correctframes}{totalframes}) \times 100\%$. This measure reflects well the accuracy of the boundaries (begin and end time) of the recognized actions, compared to manually labeled action boundaries.

With limited number of training and testing actions, results are likely to vary due to the random initialization of the training procedure (based on Expectation-Maximization [22]). For this reason, and to assess consistency in the results, we report the mean and standard deviation (*STD*) for *AER* and *FER*, computed over 10 runs.

Finally, we also use confusion matrices, whose rows and columns index the recognized and ground-truth actions, respectively; The element c_{ij} of the confusion matrix corresponds to either the percentage (for individual actions) or the instances (for group actions) of action j recognized as action i . The confusion matrix for group actions is based on *AER*, so there are substitution, insertion, and deletion errors. For individual actions, there are neither insertions nor deletions because the performance measure is *FER*.

6.2 Experiments Protocol

For both individual and group action recognition, we use 6-fold cross-validation on the training set to select the best parameters. For cross-validation, we split the training set into a training set and a validation set. For individual action recognition, we split the training streams according to

Table 5: Number of individual actions in different sets

Individual Action	Train	Test
Speaking	1088	897
Writing	363	390
Idle	1426	1349
Total	2877	2636

Table 6: Number of group actions in different sets

Group Action	Train	Test
discussion	48	49
monologue	26	26
monologue + note-taking	17	23
note-taking	6	3
presentation	6	9
presentation + note-taking	5	9
white-board	5	1
white-board + note-taking	11	19
Total	124	139

person identities. There are no people overlapping across training, validation, and test sets. For group actions, there is no overlap in participants between training and test sets. We intend the models to be person-independent, so they can be used for future meetings without need for re-training. After the best model parameters are selected, we train models on the whole training set and apply them on the test set.

From the 59 meetings, 30 are used as training data, and the remaining 29 are used for testing. The number of individual actions and group actions in the different data sets are summarized in Table 6-6. The number of individual actions is much larger than that of group actions. There are two reasons. First, for individual action recognition, there are four participants for each meeting. Therefore, there are $30 \times 4 = 120$ streams for training and $29 \times 4 = 116$ streams for testing. Second, the duration of individual actions is typically shorter than that of group actions.

6.3 Individual Action Recognition

Three HMM methods were tested for individual action recognition.

Early integration. Normal HMM trained on three feature sets: audio-only, visual-only. and audio-visual.

Audio-Visual multi-stream. Multi-stream HMM combining individual audio and visual streams.

Audio-Visual asynchronous: Asynchronous HMM combining individual audio and visual streams.

Multi-stream HMMs allow us to give different weights to different modalities. Following [13], two sets of weights were tested: (0.5, 0.5) and (0.8, 0.2) for audio and video, respectively. For asynchronous HMM, the allowed asynchrony ranges from 0.2 seconds to 2.2 seconds. Results are presented in Table 7 in terms of *FER* mean and standard

Table 7: Results of individual action recognition

Method	Features	FER (%)	STD
Early Integration	Visual	34.17	3.64
	Audio	23.48	2.70
	Audio-Visual	9.98	2.65
MS-HMM (0.5, 0.5)	Audio-Visual	14.12%	2.93
MS-HMM (0.8, 0.2)	Audio-Visual	8.58%	1.76
Asynchronous HMM	Audio-Visual	7.42%	1.13

Table 8: Confusion matrix of recognized individual actions; early-integration HMM, using AV features

	Speaking	Writing	Idle
Speaking	96.60%	1.62%	1.17%
Writing	0.16%	90.80%	17.33%
Idle	3.24%	7.58%	81.50%

deviation, obtained over 10 runs.

From Table 7, we observe that all methods using AV features got less than 15% FER, which is about 10% improvement over using audio-only features, and about 20% improvement over using visual-only features. Asynchronous HMM produced the best result. Given that the total number of frames is over 43,000, the improvement using asynchronous HMM over the other HMM methods is statistically significant with a confidence level above 99%, using a standard proportion test [7]. The improvement suggests that there exist asynchronous effects between the audio and visual modalities. Additionally, the multi-stream HMM with larger weight on audio outperformed the equal-weight scheme. This is not completely surprising given the fact the predominant role of audio in the defined actions.

The confusion matrix for early integration AV HMM, corresponding to a randomly chosen single run, is shown in Table 8. In summary, “speaking” is well detected, “writing” is generally well-detected but tends to get confused with “idle”, which in turn is the action with the highest FER. This is likely due to the catch-all role that this action plays. In practice, “idle” includes all other possible AV patterns, (e.g. pointing, laughing, etc.), which makes its modeling more difficult, compared with the other two well-defined actions.

6.4 Group Action Recognition

Using the outputs from *I-HMM* and the group-level features, concatenated as described in Section 3.3, we investigated the following cases for recognition of group actions:

Early integration, visual-only, soft decision. A normal HMM is trained using the combination of the results of the *I-HMM* trained on visual-only features, and the visual group features. The soft decision criteria is used.

Early integration, audio-only, soft decision. Same as above, but replacing visual-only by audio-only information.

Table 9: Results of group action recognition

Method	Features	AER	STD
Single-Layer	Visual-Only	48.20 %	3.78
	Audio-Only	36.70%	4.12
	Audio-Visual	23.74%	2.97
Two-Layer	Visual-Only	42.45%	2.85
	Audio-Only	32.37%	2.10
	A-HMM <i>hard decision</i>	17.85%	2.87
	A-HMM <i>soft decision</i>	15.11%	1.48

Asynchronous HMM (A-HMM), AV, soft decision. We use the asynchronous HMM for individual action layer and audio-visual features. The soft decision criteria is used.

Asynchronous HMM, AV, hard decision. Same as above, but changing the criteria to link two HMM layers.

As baseline methods for comparison, we use single-layer HMMs, using low-level audio-only, visual-only, and AV features as observations [12], and trained by cross-validation following the same experimental protocol. The results appear in Table 9, in terms of AER mean and standard deviation over 10 runs.

We observe from Table 9 that the use of AV features outperformed the use of single modalities for both single-layer HMM and two-layer HMM methods. This results supports the hypothesis that the group actions we defined are inherently multimodal. Furthermore, the two-layer HMM method using AV features improved the performance by over 8% compared to the AV single-layer HMM. Given the small number of group actions in the corpus, a standard proportion test indicates that the difference in performance between AV single-layer and two-layer HMMs using soft decision is significant at the 96% level. Additionally, the standard deviation for the two-layer approach is half the baseline’s, which suggests that our approach might be more robust to variations in initialization, given the fact that each HMM stage in our approach is trained using an observation space of relatively low dimension. Regarding hard vs. soft decision, soft decision produced a slightly better result, although not statistically significant. However, the standard deviation using soft-decision is again around half the corresponding to hard-decision. Overall, the soft decision two-layer HMM appears to be favored by the results.

The confusion matrix for AV, soft-decision, two-layer HMM for a randomly chosen single run is shown in Table 10. Empty cells represent zero values. We see that most substitution errors come from confusions between actions with and without note-taking. This might be mainly because several instances of “writing” could not be reliably detected as individual actions, as mentioned in the previous subsection. On the other hand, “discussion” and “note-taking” actions can be recognized reasonably well.

Table 10: Confusion matrix of recognized group actions for two-layer HMM frame using soft decision

	discussion	monologue	mono+note	note-taking	presentation	pres+note	white-board	wb+note	DEL
discussion	47								
monologue		21	4						1
mono+note	1	3	19						
note-taking				2					1
presentation					6	3			1
pres+note					1	5			
white-board	1	1			1		1	2	
wb+note						1		17	
INS		2	1						

7 Conclusions and Future Work

In this paper, meetings were defined as sequences of multimodal group actions. We addressed the problem of modeling and recognizing such group actions, proposing a two-layer HMM framework to decompose the group action recognition problem into two layers. The first layer maps low-level AV features into individual actions. The second layer uses results from the first layer as input to recognize group actions. Experiments on a public 59-meeting corpus demonstrate the effectiveness of the proposed framework to recognize a set of eight multimodal turn-taking actions, compared to a baseline, single-layer HMM system. We believe our methodology to be promising. In the short-term, we will explore its applicability to other sets of group actions.

Acknowledgments

This work was partly supported by the EU 5th FWP IST project M4 (IST-2001-34485), the EU 6th FWP IST Integrated Project AMI (FP6-506811, pub. AMI-11), and the Swiss NCCR IM2. We also thank Stephan Reiter (TUM) for providing part of the ground-truth for individual actions.

References

- [1] S. Basu, T. Choudhury, B. Clarkson, and A. Pentland. Towards measuring human interactions in conversational settings. In *Proc. IEEE CVPR Workshop on Cues in Communication*, Kawai, Dec. 2001.
- [2] S. Bengio. An asynchronous hidden Markov model for audio-visual speech recognition. In S. Becker, S. Thrun, and K. Obermayer, editors, *Proc. NIPS 15*, 2003.
- [3] J. DiBiase, H. Silverman, and M. Brandstein. Robust localization in reverberant rooms. In M. Brandstein and D. Ward, editors, *Microphone Arrays*, chapter 8, pages 157–180. Springer, 2001.
- [4] A. Dielmann and S. Renals. Dynamic Bayesian networks for meeting structuring. in *Proc. IEEE ICASSP*, 2004.
- [5] S. Dupont and J. Luetttin. Audio-visual speech modeling for continuous speech recognition. *IEEE Transactions on Multimedia*, 2(3):141–151, September 2000.
- [6] A. Galata, N. Johnson, and D. Hogg. Learning behavior models of human activities. In *British Machine Vision Conference*, 1999.
- [7] D. Gibbon, R. Moore, and R. Winksi. Handbook of Standards and Resources for Spoken Language Systems. *Mouton de Gruyter*, 1997.
- [8] D. Hillard, M. Ostendorf, and E. Shriberg. Detection of agreement vs. disagreement in meetings: Training with unlabeled data. In *Proc. HLT-NAACL Conference*, Edmonton, May 2003.
- [9] S. Hongeng and R. Nevatia. Multi-agent event recognition. In *Proc. IEEE ICCV*, Vancouver, July 2001.
- [10] R. Krauss and C. Garlock and P. Bricker and L. McMahon". The role of audible and visible back-channel responses in interpersonal communication. *Journal of Personality and Social Psychology*, 35(7):523-529, 1977.
- [11] J. D. Markel. The SIFT algorithm for fundamental frequency estimation. *IEEE Transactions on Audio and Electroacoustics*, 20:367–377, 1972.
- [12] I. McCowan, S. Bengio, D. Gatica-Perez, G. Lathoud, F. Monay, D. Moore, P. Wellner, and H. Bourlard. Modeling human interactions in meetings. In *Proc. IEEE ICASSP*, Hong Kong, April 2003.
- [13] I. McCowan, D. Gatica-Perez, S. Bengio, G. Lathoud, M. Barnard, and D. Zhang. Automatic analysis of multimodal group actions in meetings. IDIAP-RR 27, IDIAP, Martigny, Switzerland, May 2003.
- [14] J. E. McGrath. *Groups: Interaction and Performance*. Prentice-Hall, 1984.
- [15] D. Moore. The IDIAP smart meeting room. IDIAP-COM 07, IDIAP, 2002.
- [16] N. Morgan, D. Baron, J. Edwards, D. Ellis, D. Gelbart, A. Janin, T. Pfau, E. Shriberg, and A. Stolcke. The meeting project at ICSI. In *Proc. HLT Conference*, San Diego, CA, March 2001.
- [17] N. Morgan and E. Fosler-Lussier. Combining multiple estimators of speaking rate. in *Proc. ICASSP*, 1998.
- [18] K. Murphy. Dynamic Bayesian networks: Representation, inference and learning. *Ph.D. dissertation, UC Berkeley*, 2002.
- [19] N. Oliver, E. Horvitz, and A. Garg. Layered representations for learning and inferring office activity from multiple sensory channels. In *Proc. ICMI*, October 2002.
- [20] N. Oliver, B. Rosario, and A. Pentland. A Bayesian computer vision system for modeling human interactions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8), August 2000.
- [21] E. Padilha and J. Carletta. Nonverbal behaviours improve a simulation of small group discussion. In *Proc. First Int. Nordic Symposium of Multi-modal Communication*, Copenhagen, Sep. 2003.
- [22] L. R. Rabiner and B.-H. Juang. *Fundamentals of Speech Recognition*. Prentice-Hall, 1993.
- [23] T. Starner and A. Pentland. Visual recognition of american sign language using HMMs. In *Proc. Int. Work. on AFGR*, Zurich, 1995.
- [24] A. Waibel, M. Bett, F. Metze, K. Ries, T. Schaaf, T. Schultz, H. Soltan, H. Yu, and K. Zechner. Advances in automatic meeting record creation and access. in *Proc. IEEE ICASSP*, May 1999.
- [25] B. Wrede and E. Shriberg. The relationship between dialogue acts and hot spots in meetings. In *Proc. ASRU*, Virgin Islands, Dec. 2003.
- [26] L. Xie, S.-F. Chang, A. Divakaran, and H. Sun. Unsupervised discovery of multilevel statistical video structures using hierarchical hidden Markov models. in *Proc. ICME*, July 2003.
- [27] J. Yang, L. Weier, and A. Waibel. Skin-color modeling and adaptation. in *Proc. ACCV*, 1998.