

# Seeing and Hearing Egocentric Actions: How Much Can We Learn?

Alejandro Cartas<sup>1</sup>, Jordi Luque<sup>2</sup>, Petia Radeva<sup>1</sup>, Carlos Segura<sup>2</sup> and Mariella Dimiccoli<sup>3</sup>

<sup>1</sup>University of Barcelona, <sup>2</sup>Telefonica I+D, Research, Spain, <sup>3</sup>Institut de Robòtica i Informàtica Industrial (CSIC-UPC)

## Abstract

Our interaction with the world is an inherently multimodal experience. However, the understanding of human-to-object interactions has historically been addressed focusing on a single modality. In particular, a limited number of works have considered to integrate the visual and audio modalities for this purpose. In this work, we propose a multimodal approach for egocentric action recognition in a kitchen environment that relies on audio and visual information. Our model combines a sparse temporal sampling strategy with a late fusion of audio, spatial, and temporal streams. Experimental results on the EPIC-Kitchens dataset show that multimodal integration leads to better performance than unimodal approaches. In particular, we achieved a 5.18% improvement over the state of the art on verb classification.

## 1. Introduction

The ability to integrate multisensory information is a fundamental feature of the human brain that allows efficient interaction with the environment [14]. To mimic this human characteristic is crucial for autonomous robotics to reduce ambiguity about sensory environment and to form robust and meaningful representations. Beside artificial agents, multimodal integration has been used in traditionally vision-based tasks such as scene classification [3, 57], social behaviour analysis [1, 22, 36] and activity recognition [2, 35]. While scene classification [3, 57] and social behaviour analysis [1, 22, 36] have been approached by integrating mostly audio and visual features, comparative little attention has been paid to audiovisual integration for activity recognition [2, 35, 48]. Indeed, most existing approaches for the latter aimed at combining different inertial sensors such as accelerometers, gyroscopes and magnetic field sensors [34] or inertial cues with audio or depth cues [7, 10].

In this work, we focus on a particular case of activity recognition, that is the recognition of activities involving

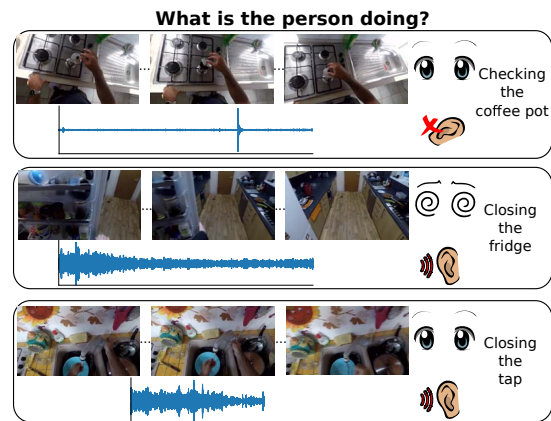


Figure 1. Audio and vision are complementary sources of information for recognizing egocentric object interactions. A limited number of interactions do not have an associated audio signal (top), but in most cases, auditory sources provide valuable information in situations such as the occlusion of the hands and objects (middle), and in some others they just strengthen the visual information (bottom).

object manipulations. More specifically, the goal is to identify actions of the type *verb+noun*, i.e. *pouring+jam* performed in a kitchen environment. Recently, promising results on this task based solely on audio features have been shown in [9]. Motivated by this work, we proposed a framework to integrate visual and audio features. Fig. 1 shows how audio cues are crucial to identify the activity being performed specially when visual information is ambiguous from an egocentric perspective due to self-occlusions. Our proposal aims at exploiting the complementarity of visual and audio features to obtain robust multimodal representations.

In particular, our contributions can be summarized as follows: (i) We provide an extensive evaluation and comparison with published methods of the proposed multimodal architecture on the EPIC-Kitchens dataset [12] (ii) In addition to action performance, we provide for the first time a detailed results on the *object* and *verb* components. The rest of the paper is organized as follows. The next section dis-

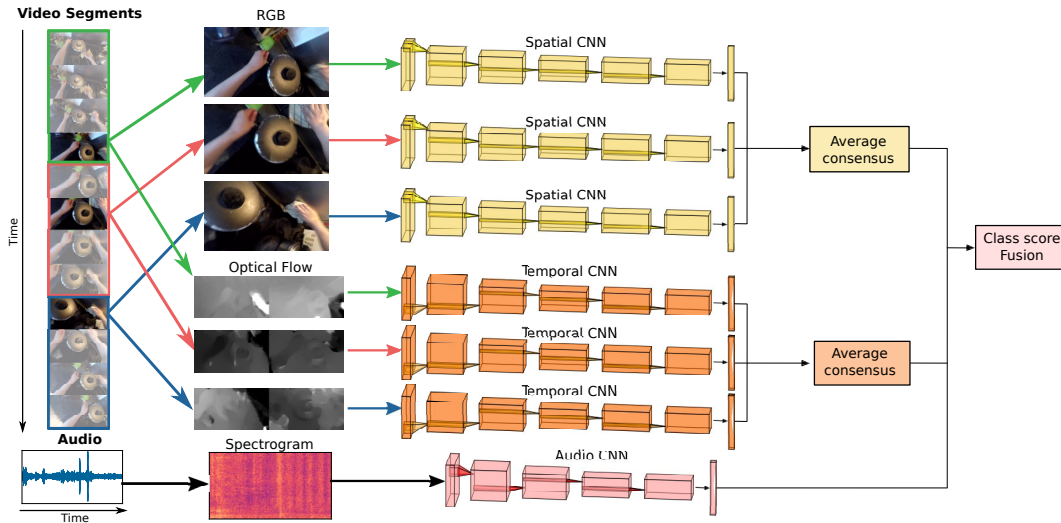


Figure 2. Pipeline of our proposed approach. A video is divided into  $K = 3$  time segments shown in green, red, and blue colors. Then, RGB and optical flow frames are sparsely sampled from each time segment to be processed in their respective spatial and temporal streams. At the end of each stream, the average consensus of the softmax scores is computed. A spectrogram is calculated from the raw audio signal and processed in its audio stream. The class scores of each stream are joined together using late fusion.

cusses related work. Section 3 introduces the proposed approach, while section 4 and section 5 detail the experimental setup and discuss the results, respectively. Finally, section 6 concludes the paper by summarizing the main findings.

## 2. Related Work

**Activity recognition** The literature on activity recognition is vast and spans several decades [24, 39]. In the following, we will limit the discussion to egocentric object interactions recognition methods. Early approaches aimed at modelling spatio-temporal features through probabilistic models [16, 17], and temporal or spatio-temporal pyramids [33, 38]. Later on, multi-stream neural network based approaches were proposed [32, 47]. Typically, each stream treats a different modality (motion, RGB) or models different cues (hands, objects, etc). Attention-based mechanisms were proposed in [50, 51]. The key idea underlying these approaches is learning to attend regions containing objects correlated with the activity under consideration. Recently, Baradel et al. [6] build on the output of an object detector to perform a spatio-temporal reasoning about the action being performed.

Two relevant works have been focused on acoustic scene and activity classification from data collected in the wild [26, 29]. Hwang and Lee [26] clustered thirteen acoustic classes based on Mel-frequency cepstrum coefficient (MFCC) features from data gathered using mobile phones. They used as a classification method a  $k$ -nearest neighborhood on hand-crafted histograms from the same spectral features. A closer work to ours is presented in [29]. They

proposed a model to classify 15 different one-minute home activities of daily living. A shallow network was trained using the encodings generated by pre-trained VGG-11 network [25] and an oversampling strategy using a large audio set [18].

**Multisensory integration** Most multimodal approaches for activity recognition through wearable sensors build on the combination of the several inertial features captured by accelerometers, gyroscopes and magnetic sensors [34]. Inertial cues are often integrated with visual information [55]. Audiovisual integration has been extensively used in the context of smart room scene analysis [44] and event detection in surveillance videos [11]. More recently, the integration of audio and visual features has been successfully used for detecting human-to-human interactions [5], for apparent personality trait recognition [22], for video description generation [27], and for multispeaker tracking [40]. Although several works that combine audiovisual sources have been reported in the context first-person action recognition challenges [15, 20, 21], they provide few details about their models. In [30, 31] are proposed attention mechanisms for action recognition using audio as a modality branch. However, the use of audio-visual cues for object interaction recognition is still very limited and previous works only reported results on the full interaction (*action*) and not its components (*verb* and *noun*).

The three main model-agnostic approaches largely used in the literature for combining the audio-visual features are early, late, and hybrid fusion [4, 28]. In the first approach, a model learns from the multimodal features after joining

them [27, 37, 40]. In the second approach, different unimodal predictors are trained and the final decision is made by combining their output [22]. The last approach combines the output from early and late fusion predictors [11, 44].

### 3. Proposed approach

In this section we describe our multimodal approach for action recognition<sup>1</sup>, that is summarized in Fig. 2. We first present our vision-based recognition approach that uses Temporal Segments Network (TSN)[52] for the spatial (RGB) and temporal (optical flow) visual modalities. Then, we present our audio-based recognition approach that uses two different convolutional neural networks (CNNs): VGG-11[46] and a custom network based on [41]. Finally, we detail our fusion strategy to integrate the different modalities.

**Vision** The visual spatial and temporal input modalities are RGB and optical flow frames calculated using [56]. Each visual modality was trained as a TSN stream. On the TSN model, the frames of a video are grouped into  $K$  sequential segments of equal size. Similarly to [52], we decided to set  $K = 3$  as originally presented. Simultaneously from each segment, a frame is sparsely sampled and processed by a CNN. In our case, we used as backbone networks a ResNet-18 and ResNet-50 [23] in our experiments. Then, a consensus of the scores from each processed frame is done. We used as a consensus function the average of the softmax scores. This model is an extension of [45], but it learns long-range temporal structure of the action in the video.

**Audio** The audio modality uses as input the spectrogram of the raw audio signal from the video. The spectrogram is calculated as follows. First, when the video has multiple audio channels, we join them by obtaining their mean. Then, we compute the short-time Fourier transform (STFT) from this signal using a sampling frequency of 16 KHz. The STFT uses a Hamming window of length equal to 30 ms with 50% time overlapping. The signal spectrogram is calculated as the logarithm value of the squared magnitude of its STFT. The final step consists in normalizing all the input spectrograms. The spectrogram has a resulting dimension size for the frequency of 331. As in [9], we only consider the first four seconds of the audio spectrogram. When it has less than four seconds duration then a zero padding is applied. This constraint results in a time dimension size of 248 for the input spectrogram.

A single spectrogram covers a larger time window than the visual input frames. Therefore, our model only needs one CNN to process the audio modality. Nonetheless, for

<sup>1</sup>Code at: [http://github.com/gorayni/seeing\\_and\\_hearing](http://github.com/gorayni/seeing_and_hearing)

Layer type	Output size	#Filters	Kernel size	Dilation
Conv2D	331×248	64	11×7	9×4
max-pool				
Conv2D	166×124	64	6×4	9×4
Conv2D	166×124	32	6×4	9×4
Conv2D	166×124	16	6×4	9×4
max-pool				
Dense	256	-	-	-
Dense	256	-	-	-

Table 1. Architecture of the traditional dilated network for audio classification.

longer video durations a long short-term memory (LSTM) could be added as in [43]. Our backbone audio CNN models are a VGG-11[46] network and a proposed smaller CNN based on [42]. We call the latter traditional dilated network and show its architecture on Table 1. This network was adapted to spectrograms with bigger sizes by using dilation convolutions [54].

**Audio-visual fusion** In our experiments we used two late fusion methods. The first method is the weighted sum of the class scores from each stream. The second method uses a network with two fully connected (FC) layers. Its input vector is calculated by concatenating the outputs of the penultimate FC layers from each stream. During training, the weights of each modality CNN stream are kept frozen.

## 4. Experimental Setup

The main objective of our experiments was to measure the performance of our proposed multimodal approach on an egocentric object interaction recognition task from videos. More specifically, the task consists on predicting what a person is doing (*verb*) using a specific object (*noun*). Both classifications can be trained and evaluated separately or combined as a single *action* classification. Therefore, our secondary objective was to determine the contribution of audio and visual information on each type of classification (noun, verb, action). Contrary to previous works [15, 20, 21], we did not make any assumption on which classification type the audio source would perform better. We provide further details of our experimental setup in the following subsections. In section 4.1, we describe the used dataset and its data partition. The specific evaluation metrics and baselines are presented in section 4.2. Finally, the implementation details of our model are detailed in section 4.3.

### 4.1. Dataset

We carried out our experiments on the EPIC Kitchens dataset [12]. Each video segment in the dataset shows a participant doing one specific cooking related action in a kitchen environment. Some examples of their labels are “cut potato” or “wash cup”. The EPIC Kitchens dataset

		Top-1 Accuracy			Top-5 Accuracy			Avg Class Precision			Avg Class Recall				
		VERB	NOUN	ACTION	VERB	NOUN	ACTION	VERB	NOUN	ACTION	VERB	NOUN	ACTION		
Chance/Random		11.38	01.58	00.43	47.58	07.74	02.12	01.00	00.34	00.06	01.00	00.34	00.07		
Largest class		20.19	04.11	02.10	66.93	18.38	07.54	00.21	00.01	00.00	01.05	00.36	00.07		
ACTION	Audio	VGG-11	29.44	05.99	<b>05.84</b>	70.49	16.90	<b>17.41</b>	00.92	00.51	<b>00.20</b>	01.86	00.79	<b>00.45</b>	
		Traditional Dilated	32.16	06.86	<b>10.77</b>	72.22	22.34	<b>25.86</b>	01.64	00.84	<b>03.04</b>	02.31	01.90	<b>04.58</b>	
	Vision	TSN Flow	36.70	14.82	<b>22.54</b>	75.98	37.25	<b>41.52</b>	02.94	03.46	<b>05.98</b>	03.89	03.06	<b>07.04</b>	
		TSN RGB	34.06	13.01	<b>32.80</b>	76.71	63.90	<b>59.06</b>	04.17	13.45	<b>15.37</b>	04.83	11.16	<b>19.62</b>	
		RGB+Flow	37.12	30.15	<b>36.19</b>	79.56	62.02	<b>59.92</b>	04.96	11.94	<b>15.14</b>	04.19	09.21	<b>17.84</b>	
	Multimodal	Flow+Audio	Weighted	38.04	13.03	<b>25.35</b>	77.42	35.81	<b>45.06</b>	03.89	04.52	<b>09.25</b>	03.59	03.21	<b>09.32</b>
			FC	37.65	27.03	<b>34.93</b>	79.52	59.41	<b>60.69</b>	04.03	13.88	<b>16.08</b>	03.66	09.49	<b>18.50</b>
		RGB+Audio	Weighted	39.95	27.45	<b>36.78</b>	80.47	59.17	<b>60.38</b>	03.30	11.03	<b>15.51</b>	04.02	08.26	<b>16.96</b>
			FC	39.66	14.44	<b>26.03</b>	77.53	39.15	<b>46.16</b>	03.60	04.78	<b>09.30</b>	03.97	03.39	<b>09.29</b>
		RGB+Flow+Audio	Weighted	38.80	30.15	<b>35.50</b>	80.31	63.44	<b>61.71</b>	04.22	14.97	<b>16.71</b>	04.20	10.91	<b>19.63</b>
			FC	40.06	29.68	<b>36.92</b>	<b>80.82</b>	61.69	<b>61.14</b>	03.16	12.47	<b>15.56</b>	04.12	09.47	<b>17.76</b>
		Flow+Audio	Weighted	41.58	21.26	<b>27.34</b>	79.63	48.33	<b>47.84</b>	05.81	07.16	<b>12.63</b>	05.08	05.69	<b>14.01</b>
			FC	40.85	36.39	<b>35.94</b>	76.84	70.05	<b>61.31</b>	<b>09.03</b>	17.13	<b>16.40</b>	07.11	12.88	<b>19.57</b>
		RGB+Flow+Audio	Weighted	42.56	<b>36.81</b>	<b>40.15</b>	77.06	<b>70.38</b>	<b>64.19</b>	08.48	<b>18.08</b>	<b>19.21</b>	<b>07.55</b>	<b>12.93</b>	<b>22.68</b>
			FC												
	VERB+NOUN	Audio	VGG-11	34.48	09.51	03.56	74.50	26.63	12.17	05.26	01.32	00.28	04.04	01.32	01.72
			Traditional Dilated	34.82	15.44	06.26	74.72	36.96	17.83	04.53	05.77	01.12	03.88	04.95	01.39
		Vision	TSN Flow	49.08	22.72	13.54	81.60	46.32	30.77	10.80	08.81	02.53	07.12	04.97	02.23
TSN RGB			50.65	54.01	32.51	88.63	80.87	59.72	25.96	38.83	16.23	<b>19.36</b>	<b>34.43</b>	18.94	
RGB+Flow			55.47	52.82	32.76	88.48	78.01	58.13	28.94	39.82	13.53	14.25	<b>27.81</b>	14.22	
Multimodal		Flow+Audio	Weighted	50.06	26.26	15.13	81.02	51.45	33.49	11.72	11.32	02.99	06.61	06.42	02.42
			FC	53.51	53.11	32.21	87.35	79.63	57.82	26.57	38.89	13.67	13.07	28.98	14.94
		RGB+Audio	Weighted	56.27	51.09	32.27	87.24	77.15	55.96	25.06	37.17	11.81	11.34	24.28	12.05
			FC	51.40	26.39	15.62	81.57	51.54	34.24	11.86	12.57	03.37	06.98	06.48	02.72
		RGB+Flow+Audio	Weighted	54.24	54.90	33.84	88.19	80.89	59.72	<b>31.21</b>	38.96	15.03	15.07	31.74	16.73
			FC	56.65	53.90	33.86	87.70	79.47	58.37	25.75	<b>40.87</b>	13.36	12.58	28.17	14.02
		Flow+Audio	Weighted	52.00	33.02	20.22	82.24	58.50	39.37	09.72	18.58	05.80	08.11	16.63	06.58
			FC	55.41	55.08	35.21	87.15	79.78	60.27	20.49	38.47	15.35	14.28	33.55	17.64
		RGB+Flow+Audio	Weighted	<b>60.21</b>	<b>56.14</b>	38.55	<b>89.07</b>	<b>80.96</b>	62.97	27.05	39.89	17.14	19.09	33.85	19.28
			FC												

Table 2. Classification performance for the equally stratified *action* data split. The *action* part show the results of training one classifier on the action labels, whereas the *verb+noun* part show the results of independently training two classifiers over the verb and noun labels. The scores in gray color were calculated based on the respective *action* or *verb+noun* classifiers.




Figure 3. Accuracy difference of the *noun* split for the unweighted test predictions that changed with respect to RGB+Flow and RGB+Flow+Audio.

includes 432 videos recorded from a first-person perspective by 32 participants in their own kitchens while cooking/preparing something. Each video was divided into segments in which the person is doing one specific *action* (a *verb* plus a *noun*). The total number of verbs and nouns categories in dataset is 125 and 352, correspondingly.

For comparison purposes, we considered two data partitions derived from the labeled data of the EPIC Kitchen Challenge:

**Home made partition** In this partition, all the participants were considered for the training, validation, and test splits. The data proportions for the validation and test splits were 10% and 15%, accordingly. Since the resulting distribution of action classes is highly unbalanced, the data split was done as follows. At least one sample of each action category was put in the training split. If the category had

at least two samples, one of them went to the test split. We also report the results obtained on the EPIC Kitchens Challenge board from the models trained on this partition.

**Unseen verb partition** The second data partition was the one proposed [6]. This partition only used the *verb* classes from the labeled data. Moreover, the training and test splits are on the participants 01-25 and 26-31, respectively. Therefore, the test set is only composed of unseen kitchens. In order to train our methods, we created a randomly stratified validation split with 10% of data from the training split.

## 4.2. Evaluation metrics

Following [12], we measured the classification performance using aggregate and per-class metrics. As aggregate metrics for measuring the classification performance we used the top-1 and top-5 accuracy, whereas as per-class





	Method	Top-1 Accuracy	Top-5 Accuracy	Avg. Class Precision	Avg. Class Recall
	Chance/Random	11.75	48.87	00.99	00.98
	Largest class	21.27	69.44	00.31	01.41
Audio	Traditional Dilated	30.51	74.19	04.71	03.60
	VGG-11	33.27	74.13	05.72	04.08
Vision	ResNet-18 [23]†	32.05	-	-	-
	I3D ResNet-18 [8]†	34.20	-	-	-
	TSN ResNet-18 RGB	34.69	77.13	08.38	05.08
	ORN [6]†	40.89	-	-	-
	TSN ResNet-18 Flow	44.48	77.88	08.15	06.18
	RGB+Flow	43.36	78.77	09.98	05.58
Multimodal	RGB+Audio VGG-11	41.09	80.10	08.82	04.76
	Flow+Audio VGG-11	45.75	80.34	08.41	05.57
	RGB+Flow+Audio VGG-11	45.86	80.72	09.33	05.25
Weighted	RGB+Audio VGG-11	40.83	79.42	09.45	04.85
	Flow+Audio VGG-11	43.74	79.80	09.68	05.31
	RGB+Flow+Audio VGG-11	<b>46.07</b>	<b>80.76</b>	09.34	05.30
FC	RGB+Audio VGG-11	42.08	79.44	<b>12.51</b>	06.55
	Flow+Audio VGG-11	44.50	78.68	08.17	06.86
	RGB+Flow+Audio VGG-11	45.92	80.33	11.01	<b>07.31</b>

Table 3. Classification performance results on the comparison *verb* data split. The results marked with † were originally reported in [6].




Figure 6. Accuracy difference of the *action* split for the unweighted test predictions that changed with respect to RGB+Flow and RGB+Flow+Audio.

## 5. Results

**Noun** The performance results for the *noun* classification on the home made partition are shown in the lower part of Table 2. The best accuracy score was achieved by the weighted multimodal combination that improved the visual baseline by 1.24%. These results also indicate that the separated or combined unweighted fusion of the optical flow and audio decreases the top-1 and top-5 accuracy of the task. This effect can also be seen on the higher number of classes that decreased their accuracy on Fig. 3. The most misclassified pair of objects are *spoon-knife*, *spoon-fork*, *plate-bowl*, *tap-sponge*, and *knife-fork*, as shown in Fig. 5.

**Verb** The *verb* classification results on the home made partition are also presented in the lower part of Table 2. Each visual method combination boosted their respective performance by adding the unweighted audio score. The multimodality combination is greater than the best visual

method by 4.74%. According to Fig. 4, the multimodality helps to disambiguate *turn-on* and *turn-off* verbs, but fails on verbs that lack of sound like *scoop* or *adjust*. Fig. 8 shows that the most misclassified pair of verbs are *take-put*, *put-open*, *take-close*, *take-open*, and *put-close*.

**Action** The results of our experiments on the home made partition are presented in the upper part of Table 2. They show that using multimodal information outperforms single audio or visual classification when training separately the *verb* and *noun* classifiers. Even though the best top-1 accuracy was achieved when considering the *verb* and *noun* labels as a single action classification problem, all other performance metrics were lower than when considering them separately. The upper part of Table 2 shows that the accuracy of *noun* is diminished when adding audio scores to the *action* classification. The categories that changed their prediction on the multimodality setting for the action classifier

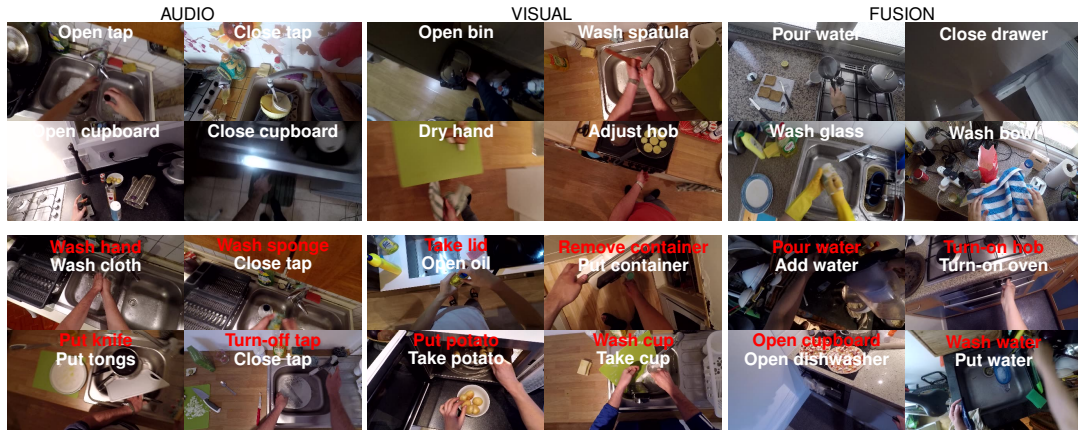


Figure 7. Qualitative results for the unweighted multimodal *action* classification experiment. The top and bottom rows show true and false positive prediction, respectively. The columns indicate when the audio, vision, or their fusion scores were the final multimodal decision. The true and false labels are shown in white and red colors.

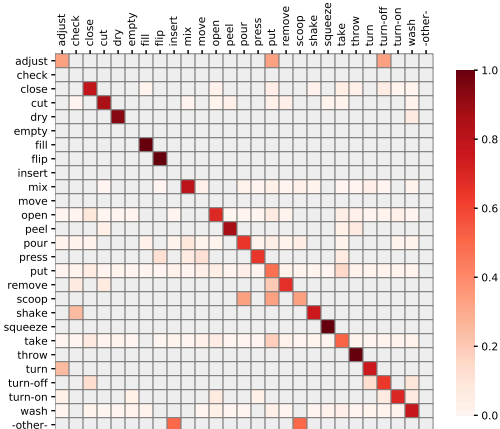


Figure 8. Normalized verb confusion matrix with an accuracy of 62.43%. Only the categories with more than 100 samples are shown.

are presented in Fig. 6.

Some qualitative results of true and false positive predictions are shown in Fig. 7. The action predictions made by the audio modality rely more on the *verb* than *noun* classification. For instance, it correctly predicts the verb *wash* in Fig. 7, but falsely predicts the noun *hand* instead of *cloth*. The visual modality obtained similar accuracy on *verb* and *noun* classification. Their predictions fail in cases where the actions occur in similar contexts, for example, the actions *take cup* and *wash cup* can occur in the sink, as illustrated in in Fig. 7. The combination of audio and visual modalities help to disambiguate actions where the objects are occluded such as *open drawer* and *close drawer*, as shown in first row of the fusion column in Fig. 7. Their complementarity also helps on the classification of actions with the same *verb*, but different *noun*, such as the case of *washing glass* and *bowl*,

shown in the second row of the fusion column in Fig. 7.

**Comparison with other methods** Table 4 shows the results obtained on the EPIC Kitchen Challenge board using the models trained on the home made partition. These results indicate that directly training over the *action* performs better than the combination of *verb+noun*. Our multimodal models obtained better scores than the challenge baseline and has similar results as previous works [50]. Additionally, the results obtained on the unseen participants (S2) test split are in the top-ten ranking of the first challenge. We can also observe that the unweighted addition as fusion method for *noun* diminishes the aggregate and per-class performance, not only for our method but also in the baseline results [12].

The results on the comparison data split originally presented on [6] are shown in Table 3. Our method obtained an improvement in accuracy of 5.18% with respect to the ORN method [6]. This test on unseen kitchens showed that methods that only rely on RGB underperform optical flow methods. Likewise, they also showed that audio classification methods can have similar performance to visual methods. The addition of auditory sources increased the performance of the best visual method by 3.47%.

**Comparison of audio CNN architectures** The VGG-11 and the Traditional Dilated network have similar classification performance. The results on Table 2 indicate that the Traditional Dilated network has better results on seen test users, but the results on Table 3 shows that the VGG-11 network outperforms the Traditional Dilated network on unseen test users.

		Top-1 Accuracy			Top-5 Accuracy			Avg Class Precision			Avg Class Recall				
		VERB	NOUN	ACTION	VERB	NOUN	ACTION	VERB	NOUN	ACTION	VERB	NOUN	ACTION		
S1	Audio	Traditional Dilated (Verb+Noun)	35.11	10.65	03.95	75.33	28.63	13.01	15.43	06.19	01.75	11.03	06.64	01.26	
		Traditional Dilated (Action)	34.06	05.31	07.43	73.51	18.24	20.94	05.19	02.74	01.81	07.85	04.14	03.08	
		TSN BNInception (FUSION) [12]	48.23	36.71	20.54	84.09	62.32	39.79	47.26	35.42	10.46	22.33	30.53	08.83	
	Visual	TSN ResNet-50 (Verb+Noun) (FUSION)	55.08	38.59	24.38	86.36	64.16	45.37	43.69	38.59	14.91	28.63	32.10	12.12	
		TSN ResNet-50 (Action) (FUSION)	38.62	25.84	27.95	79.51	54.18	49.12	10.50	24.63	14.13	14.30	21.12	14.61	
		LSTA (two stream) [50]	59.55	38.35	30.33	85.77	61.49	49.97	42.72	36.19	14.46	38.12	36.19	17.76	
		3rd Place Challenge [49]	63.34	44.75	35.54	89.01	69.88	57.18	63.21	42.26	19.76	37.77	41.28	21.19	
		2nd Place Challenge [19]	64.14	47.65	35.75	87.64	70.66	54.65	43.64	40.52	18.95	38.31	45.29	21.13	
	1st Place Challenge [53]	<b>69.80</b>	<b>52.27</b>	<b>41.37</b>	<b>90.95</b>	<b>76.71</b>	<b>63.59</b>	<b>63.55</b>	<b>46.86</b>	<b>25.13</b>	<b>46.94</b>	<b>49.17</b>	<b>26.39</b>		
	Multimodal	Action	Ours	56.37	37.69	24.00	85.47	63.45	44.66	48.15	38.02	13.49	25.54	30.31	10.50
			Ours (Weighted)	56.44	<b>39.42</b>	25.26	85.87	<b>65.27</b>	46.27	<b>51.39</b>	<b>38.36</b>	14.88	26.66	32.88	11.90
			Ours (FC)	<b>58.88</b>	39.13	27.35	<b>87.15</b>	64.83	47.68	46.36	37.92	<b>16.63</b>	<b>38.13</b>	<b>34.90</b>	15.00
			Ours	40.80	22.29	28.83	81.04	50.59	49.68	11.43	23.00	15.89	13.40	17.90	14.18
	Multimodal	Verb+Noun	Ours (Weighted)	41.22	24.29	29.09	81.16	53.25	<b>50.57</b>	11.24	23.61	15.35	14.04	19.72	14.51
			Ours (FC)	44.64	30.64	<b>29.13</b>	76.41	59.39	49.71	19.90	32.28	16.51	21.99	25.28	<b>16.54</b>
			Ours	30.73	07.20	02.53	67.26	21.65	09.90	13.92	04.52	01.90	09.79	04.68	01.20
			Traditional Dilated (Action)	31.96	03.89	03.96	64.73	13.96	13.45	05.05	01.66	00.97	07.91	04.04	01.94
	S2	Audio	Traditional Dilated (Verb+Noun)	30.73	07.20	02.53	67.26	21.65	09.90	13.92	04.52	01.90	09.79	04.68	01.20
Traditional Dilated (Action)			31.96	03.89	03.96	64.73	13.96	13.45	05.05	01.66	00.97	07.91	04.04	01.94	
TSN BNInception (FUSION) [12]			39.40	22.70	10.89	74.29	45.72	25.26	22.54	15.33	05.60	13.06	17.52	05.81	
Visual		TSN ResNet-50 (Verb+Noun) (FUSION)	45.72	24.89	14.95	77.06	49.37	31.07	24.44	20.30	08.79	18.04	18.96	10.10	
		TSN ResNet-50 (Action) (FUSION)	36.63	18.06	17.14	75.28	42.03	34.65	11.38	10.89	07.27	12.96	14.38	10.78	
		LSTA (two stream) [50]	47.32	22.16	16.63	77.02	43.15	30.93	31.57	17.91	08.97	26.17	17.80	11.92	
		3rd Place Challenge [49]	49.37	27.11	20.25	77.50	51.96	37.56	31.09	21.06	09.18	18.73	21.88	14.23	
		2nd Place Challenge [19]	55.24	33.87	23.93	80.23	58.25	40.15	25.71	28.19	<b>15.72</b>	25.69	29.51	17.06	
1st Place Challenge [53]		<b>59.68</b>	<b>34.14</b>	<b>25.06</b>	<b>82.69</b>	<b>62.38</b>	<b>45.95</b>	<b>37.20</b>	<b>29.14</b>	15.44	<b>29.81</b>	<b>30.48</b>	<b>18.67</b>		
Multimodal		Action	Ours	46.88	25.16	14.58	77.13	48.69	31.00	<b>28.72</b>	<b>16.63</b>	08.93	17.25	17.83	08.55
			Ours Weighted	47.46	25.95	15.74	<b>77.16</b>	<b>50.12</b>	31.85	28.71	16.47	09.26	17.85	19.21	09.94
			Ours FC	<b>47.49</b>	<b>26.36</b>	15.98	76.68	49.37	31.75	24.64	20.61	<b>09.80</b>	<b>20.59</b>	<b>21.35</b>	10.03
			Ours	38.37	15.23	<b>18.40</b>	75.15	39.84	35.64	10.93	11.60	06.88	11.75	13.31	10.91
Multimodal		Verb+Noun	Ours (Weighted)	38.10	16.76	18.23	75.38	42.23	<b>35.68</b>	11.58	12.83	07.72	11.79	14.16	11.26
			Ours (FC)	40.87	20.38	17.65	69.27	45.82	33.73	15.72	15.35	09.61	17.34	16.95	<b>12.20</b>

Table 4. Performance comparison with EPIC Kitchens challenge baseline results. The results highlighted in **bold blue** are the best obtained by our method.

## 5.1. Discussion

The experimental results show that our model improves the top-1 accuracy by 3.61% in average for *verb*, *noun*, and *action* classification on our home made data partition. Likewise, the results suggest that audiovisual multimodality benefits the classification of *verb*, and consequently *action*, more than for the classification of *noun*. Furthermore, although multimodality improves the aggregate performance metrics and the avg. class precision for *noun* classification, the unweighted fusion decreases their value as observed in Table 2. This might be as a consequence of three main reasons. First, an interacting object can produce several sounds and not being described by one in particular. For instance, the sounds of a fridge being opened and closed are characteristically different. Second, rather than describing an object, sounds are better suited for describing the materials they are made of. For example, water and milk are liquids and their emitting sound while being poured is indistinguishable. Third, objects lack of a time dimension, but doing an action and making a sound involve time. Nonetheless, not all actions produce any sound, like checking the coffee pot. The verb confusion matrix in Fig. 8 shows that the multimodality classification on verbs fails on categories that does not produce any sound and that are more visually abstract, like *checking* the heat. Additionally, harder audiovisual classes are *empty*, *flip*, and *squeeze*, as seen in Fig. 8 and Fig. 4. The most discriminative input source for the

*noun*, *verb*, and *action* comes from the RGB frames as observed from Tables 2 and 4. However their performance decreases when the test is performed on images from unseen persons and the optical flow achieves a higher accuracy, as seen in Table 3.

## 6. Conclusions

We presented a multimodal approach for egocentric action classification and we validated it on the EPIC Kitchens dataset. Our approach combines audiovisual input sources. Specifically, its audio input is the spectrogram extracted from the raw audio of the video, while its visual inputs are RGB and optical flow frames. We tested and analyzed our approach for classifying each separate category (*verb* and *noun*) or merged (*action*). The obtained results show that our model improves the top-1 accuracy by 3.61% in average. Additionally, the results suggest that multimodal information is specially beneficial for the *verb* recognition problem. Indeed, our multimodal approach outperformed the state of the art methods on *verb* classification by 5.18% accuracy.

**Acknowledgment** This work was partially funded by TIN2018-095232-B-C21, 2017 SGR 1742, Nestore, Validithi, 20141510 (La MaratoTV3) and CERCA Programme/Generalitat de Catalunya. A.C. is supported by a doctoral fellowship from the Mexican Council of Science and Technology (CONACYT) (grant-no. 366596). We acknowledge the support of NVIDIA Corporation with the donation of Titan Xp GPUs.



## References

- [1] Sharifa Alghowinem, Roland Goecke, Michael Wagner, Julien Epps, Matthew Hyett, Gordon Parker, and Michael Breakspear. Multimodal depression detection: fusion analysis of paralinguistic, head pose and eye gaze behaviors. *IEEE Transactions on Affective Computing*, 2016. 1
- [2] R. Arandjelović and A. Zisserman. Look, listen and learn. In *IEEE International Conference on Computer Vision*, 2017. 1
- [3] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. In *Advances in Neural Information Processing Systems*, 2016. 1
- [4] T. Baltruaitis, C. Ahuja, and L. Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):423–443, Feb 2019. 2
- [5] Sophia Bano, Tamas Suveges, Jianguo Zhang, and Stephen J Mckenna. Multimodal egocentric analysis of focused interactions. *IEEE Access*, 6:37493–37505, 2018. 2
- [6] Fabien Baradel, Natalia Neverova, Christian Wolf, Julien Mille, and Greg Mori. Object level visual reasoning in videos. In *European Conference on Computer Vision (ECCV)*, June 2018. 2, 4, 5, 6, 7
- [7] Andreas Bulling, Ulf Blanke, and Bernt Schiele. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys (CSUR)*, 46(3):33, 2014. 1
- [8] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 6
- [9] Alejandro Cartas, Jordi Luque, Petia Radeva, Carlos Segura, and Mariella Dimiccoli. How much does audio matter to recognize egocentric object interactions? *arXiv preprint arXiv:1906.00634*, 2019. 1, 3, 5
- [10] Chen Chen, Roozbeh Jafari, and Nasser Kehtarnavaz. A survey of depth and inertial sensor fusion for human action recognition. *Multimedia Tools and Applications*, 76(3):4405–4425, 2017. 1
- [11] Marco Cristani, Manuele Bicego, and Vittorio Murino. Audio-visual event recognition in surveillance video sequences. *IEEE Transactions on Multimedia*, 9(2):257–267, 2007. 2, 3
- [12] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, 2018. 1, 3, 4, 5, 7, 8
- [13] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, June 2009. 5
- [14] Marc O Ernst and Heinrich H Bülthoff. Merging the senses into a robust percept. *Trends in cognitive sciences*, 8(4):162–169, 2004. 1
- [15] Bernard Ghanem Fabian Caba Heilbron, Victor Escorcia and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–970, 2015. 2, 3
- [16] Alireza Fathi, Ali Farhadi, and James M Rehg. Understanding egocentric activities. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 407–414, 2011. 2
- [17] Alireza Fathi, Yin Li, and James Rehg. Learning to recognize daily actions using gaze. *European Conference in Computer Vision (ECCV)*, pages 314–327, 2012. 2
- [18] Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *Proc. IEEE ICASSP 2017*, New Orleans, LA, 2017. 2
- [19] Deepti Ghadiyaram, Matt Feiszli, Du Tran, Xueting Yan, Heng Wang, and Dhruv Mahajan. Large-scale weakly-supervised pre-training for video action recognition. *CoRR*, abs/1905.00561, 2019. 8
- [20] Bernard Ghanem, Juan Carlos Niebles, Cees Snoek, Fabian Caba Heilbron, Humam Alwassel, Victor Escorcia, Ranjay Krishna, Shyamal Buch, and Cuong Duc Dao. The activitynet large-scale activity recognition challenge 2018 summary. *CoRR*, abs/1808.03766, 2018. 2, 3
- [21] Bernard Ghanem, Juan Carlos Niebles, Cees Snoek, Fabian Caba Heilbron, Humam Alwassel, Ranjay Krishna, Victor Escorcia, Kenji Hata, and Shyamal Buch. Activitynet challenge 2017 summary. *CoRR*, abs/1710.08011, 2017. 2, 3
- [22] Yağmur Güçlütürk, Umur Güçlü, Marcel AJ van Gerven, and Rob van Lier. Deep impression: Audiovisual deep residual networks for multimodal apparent personality trait recognition. In *European Conference on Computer Vision*, pages 349–358. Springer, 2016. 1, 2, 3
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 3, 6
- [24] Samitha Herath, Mehrtash Harandi, and Fatih Porikli. Going deeper into action recognition: A survey. *Image and vision computing*, 60:4–21, 2017. 2
- [25] Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron Weiss, and Kevin Wilson. Cnn architectures for large-scale audio classification. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017. 2
- [26] Kyuwoong Hwang and Soo-Young Lee. Environmental audio scene and activity recognition through mobile-based crowdsourcing. *Consumer Electronics, IEEE Transactions on*, 58:700–705, 05 2012. 2
- [27] Qin Jin and Junwei Liang. Video description generation using audio and visual cues. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval*, pages 239–242. ACM, 2016. 2, 3
- [28] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video

- classification with convolutional neural networks. In *Proceedings of International Computer Vision and Pattern Recognition (CVPR 2014)*, 2014. 2
- [29] Dawei Liang and Edison Thomaz. Audio-based activities of daily living (adl) recognition with large-scale acoustic embeddings from online videos. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 3(1):17:1–17:18, Mar. 2019. 2
- [30] Xiang Long, Chuang Gan, Gerard de Melo, Jiajun Wu, Xiao Liu, and Shilei Wen. Attention clusters: Purely attention based local feature integration for video classification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2
- [31] Xiang Long, Chuang Gan, Gerard Melo, Xiao Liu, Yandong Li, Fu Li, and Shilei Wen. Multimodal keyless attention fusion for video classification, 2018. 2
- [32] Minghuang Ma, Haoqi Fan, and Kris M. Kitani. Going deeper into first-person activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [33] Tomas McCandless and Kristen Grauman. Object-centric spatio-temporal pyramids for egocentric activity recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*, volume 2, page 3, 2013. 2
- [34] Francisco Ordóñez and Daniel Roggen. Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1):115, 2016. 1, 2
- [35] Andrew Owens, Phillip Isola, Josh H. McDermott, Antonio Torralba, Edward H. Adelson, and William T. Freeman. Visually indicated sounds. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2405–2413, 2016. 1
- [36] Maja Pantic, Alex Pentland, Anton Nijholt, and Thomas S Huang. Human computing and machine understanding of human behavior: A survey. In *Artificial Intelligence for Human Computing*, pages 47–71. Springer, 2007. 1
- [37] A. Pieropan, G. Salvi, K. Pauwels, and H. Kjellström. Audio-visual classification and detection of human manipulation actions. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3045–3052, Sep. 2014. 3
- [38] Hamed Pirsiavash and Deva Ramanan. Detecting activities of daily living in first-person camera views. In *Proceedings of the IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 2847–2854. IEEE, 2012. 2
- [39] Ronald Poppe. A survey on vision-based human action recognition. *Image and vision computing*, 28(6):976–990, 2010. 2
- [40] Xinyuan Qian, Alessio Brutti, Oswald Lanz, Maurizio Omologo, and Andrea Cavallaro. Multi-speaker tracking from an audio-visual sensing device. *IEEE Transactions on Multimedia*, 2019. 2, 3
- [41] Tara Sainath and Carolina Parada. Convolutional neural networks for small-footprint keyword spotting. In *Interspeech*, 2015. 3
- [42] Tara Sainath and Carolina Parada. Convolutional neural networks for small-footprint keyword spotting. In *Interspeech*, 2015. 3
- [43] Tara Sainath, Oriol Vinyals, Andrew Senior, and Hasim Sak. Convolutional, long short-term memory, fully connected deep neural networks. pages 4580–4584, 04 2015. 3
- [44] Shankar T Shivappa, Mohan M Trivedi, and Bhaskar D Rao. Hierarchical audio-visual cue integration framework for activity analysis in intelligent meeting rooms. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 107–114. IEEE, 2009. 2, 3
- [45] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems (NIPS)*, pages 568–576, 2014. 3
- [46] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 3, 5
- [47] Suriya Singh, Chetan Arora, and C. V. Jawahar. First person action recognition using deep learned descriptors. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [48] Mahesh Subedar, Ranganath Krishnan, Paulo Lopez Meyer, Omesh Tickoo, and Jonathan Huang. Uncertainty aware audiovisual activity recognition using deep bayesian variational inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 103–108, 2019. 1
- [49] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. FBK-HUPBA submission to the epic-kitchens 2019 action recognition challenge. *CoRR*, abs/1906.08960, 2019. 8
- [50] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. Lsta: Long short-term attention for egocentric action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9954–9963, 2019. 2, 7, 8
- [51] Swathikiran Sudhakaran and Oswald Lanz. Attention is all we need: Nailing down object-centric attention for egocentric activity recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2018. 2
- [52] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European Conference on Computer Vision (ECCV)*, 2016. 3
- [53] Xiaohan Wang, Yu Wu, Linchao Zhu, and Yi Yang. Baiduuts submission to the epic-kitchens action recognition challenge 2019. *CoRR*, abs/1906.09383, 2019. 8
- [54] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2016. 3
- [55] Haibin Yu, Guoxiong Pan, Mian Pan, Chong Li, Wenyan Jia, Li Zhang, and Mingui Sun. A hierarchical deep fusion framework for egocentric activity recognition using a wearable hybrid sensor system. *Sensors*, 19(3):546, 2019. 2
- [56] C. Zach, T. Pock, and H. Bischof. A duality based approach for realtime tv-l1 optical flow. In Fred A. Hamprecht, Christoph Schnörr, and Bernd Jähne, editors, *Pattern Recognition*, pages 214–223, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg. 3

- [57] Tong Zhang and C-C Jay Kuo. Audio content analysis for on-line audiovisual data segmentation and classification. *IEEE Transactions on speech and audio processing*, 9(4):441–457, 2001. 1