Affective Computing for Large-scale Heterogeneous Multimedia Data: A Survey

SICHENG ZHAO, University of California, Berkeley, USA

SHANGFEI WANG, University of Science and Technology of China, China

MOHAMMAD SOLEYMANI, University of Southern California, USA

DHIRAJ JOSHI, IBM Research AI, USA

QIANG JI, Rensselaer Polytechnic Institute, USA

The wide popularity of digital photography and social networks has generated a rapidly growing volume of multimedia data (i.e., images, music, and videos), resulting in a great demand for managing, retrieving, and understanding these data. Affective computing (AC) of these data can help to understand human behaviors and enable wide applications. In this article, we survey the state-of-the-art AC technologies comprehensively for large-scale heterogeneous multimedia data. We begin this survey by introducing the typical emotion representation models from psychology that are widely employed in AC. We briefly describe the available datasets for evaluating AC algorithms. We then summarize and compare the representative methods on AC of different multimedia types, i.e., images, music, videos, and multimodal data, with the focus on both hand-crafted features-based methods and deep learning methods. Finally, we discuss some challenges and future directions for multimedia affective computing.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Information systems** → **Sentiment analysis**; • **Human-centered computing** → **Human computer interaction (HCI)**;

Additional Key Words and Phrases: Affective computing, emotion recognition, sentiment analysis, large-scale multimedia

ACM Reference format:

Sicheng Zhao, Shangfei Wang, Mohammad Soleymani, Dhiraj Joshi, and Qiang Ji. 2019. Affective Computing for Large-scale Heterogeneous Multimedia Data: A Survey. *ACM Trans. Multimedia Comput. Commun. Appl.* 15, 3s, Article 93 (December 2019), 32 pages.

<https://doi.org/10.1145/3363560>

This work was supported by Berkeley DeepDrive, the National Natural Science Foundation of China (Nos. 61701273, 91748129), and the National Key R&D Program of China (Grant No. 2017YFC011300). The work of MS is supported in part by the U.S. Army. Any opinion, content, or information presented does not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

Authors' addresses: S. Zhao, University of California, Berkeley, Berkeley, CA, 94720; email: schzhao@gmail.com; S. Wang (corresponding author), University of Science and Technology of China, Hefei, 230027, China; email: sfwang@ustc.edu.cn; M. Soleymani (corresponding author), University of Southern California, Playa Vista, CA, 90094; email: soleymani@ict.usc.edu; D. Joshi, IBM Research AI, Yorktown Heights, NY, 10598; email: djoshi@us.ibm.com; Q. Ji, Rensselaer Polytechnic Institute, Troy, NY, 12180; email: qji@ecse.rpi.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

1551-6857/2019/12-ART93 \$15.00

<https://doi.org/10.1145/3363560>

1 INTRODUCTION

Users are increasingly recording their daily activities, sharing interesting experiences, and expressing personal viewpoints using mobile devices on social networks, such as Twitter,¹ Facebook,² and Weibo,³ and so on. As a result, a rapidly growing volume of multimedia data (i.e., images, music, and videos) has been generated, as shown in Figure 1, which results in a great demand for the management, retrieval, and understanding of these data. Most existing work on multimedia analysis focuses on the cognitive aspects, i.e., understanding the objective content, such as object detection in images [44], speaker recognition in speech [47], and action recognition in videos [48]. Since what people feel has a direct influence on their decision making, affective computing (AC) of these multimedia data is of significant importance and has attracted increasing attention [22, 146, 154, 156, 164]. For example, companies would like to know how customers evaluate their products and can thus improve their services [53]; depression and anxiety detection from social media can help understand psychological distress and thus potentially prevent suicidal actions [106].

While the sentiment analysis in text [91] has long been a standard task, AC from other modalities, such as image and video, has just begun to be considered recently. In this article, we aim to review the existing AC technologies comprehensively for large-scale heterogeneous multimedia data, including images, music, videos, and multimodal data.

Affective computing of multimedia (ACM) aims to recognize the emotions that are expected to be evoked in viewers by a given stimuli. Similar to other supervised learning tasks, ACM is typically composed of three steps: data collection and annotation, feature extraction, and mapping learning between features and emotions [157]. One main challenge for ACM is the affective gap, i.e., “the lack of coincidence between the features and the expected affective state in which the user is brought by perceiving the signal” [45]. In the early stage, various handcrafted features were designed to bridge this gap with traditional machine learning algorithms, while more recently researchers have focused on end-to-end deep learning from raw multimedia data to recognize emotions. Existing ACM methods mainly assign the dominant (average) emotion category (DEC) to an input stimuli, based on the assumption that different viewers have similar reactions to the same stimuli. We can usually formulate this task as a single-label learning problem.

However, emotions are influenced by subjective and contextual factors, such as the educational background, cultural diversity, and social interaction [95, 141, 168]. As a result, different viewers may react differently to the same stimuli, which creates the subjective perception challenge. Therefore, the perception inconsistency makes it insufficient to simply predict the DEC for the highly subjective variable. As stated in Reference [168], we can perform two kinds of ACM tasks to deal with the subjectivity challenge: predicting personalized emotion perception for each viewer and assigning multiple emotion labels for each stimuli. For the latter one, we can either assign multiple labels to each stimuli with equal importance using multi-label learning methods, or predict the emotion distributions, which try to learn the degrees of each emotion [141].

In this article, we concentrate on surveying the existing methods on ACM and analyzing potential research trends. Section 2 introduces the widely used emotion representation models from psychology. Section 3 summarizes the existing available datasets for evaluating ACM tasks. Section 4, Section 5, Section 6, and Section 7 survey the representative methods on AC of images, music, videos, and multimodal data, respectively, including both handcrafted features-based methods and deep learning methods. Section 8 provides some suggestions for future research, followed by conclusion in Section 9.

¹<https://twitter.com>.

²<https://www.facebook.com>.

³<https://www.weibo.com>.

Fig. 1. Different multimedia data that are widely used to express emotions.

To the best of our knowledge, this article is among the first that provide a comprehensive survey of affective computing of multimedia data from different modalities. Previous surveys mainly focus on a single modality, such as images [56, 161], speech [35], music [60, 145], video [11, 131], and multimodal data [116]. From this survey, readers can more easily compare the correlations and differences among different AC settings. We believe that this will be instrumental in generating novel research ideas.

2 EMOTION MODELS FROM PSYCHOLOGY

There are two dominant emotion representation models deployed by psychologists: categorical emotions (CE), and dimensional emotion space (DES). CE models classify emotions into a few basic categories, such as *happiness*, *anger*, and so on. Some commonly used models include Ekman's six basic emotions [34] and Mikels's eight emotions [77]. When classifying emotions into *positive* and *negative* (polarity) [160, 163], sometimes including *neutral*, "emotion" is called "sentiment." However, sentiment is usually defined as an attitude held toward an object [116]. Emotions are usually represented by DES models as continuous coordinate points in a 3D or 2D Cartesian space, such as valence-arousal-dominance (VAD) [104] and activity-temperature-weight [66]. VAD is the most widely used DES model, where valence represents the pleasantness ranging from positive to negative, arousal represents the intensity of emotion ranging from excited to calm, and dominance represents the degree of control ranging from controlled to in control. Dominance is difficult to measure and is often omitted, leading to the commonly used two dimensional VA space [45].

The relationship between CE and DES and the transformation from one to the other are studied in Reference [124]. For example, positive valence relates to a happy state, while negative valence relates to a sad or angry state. CE models are easier for users to understand and label, but the limited set of categories may not well reflect the subtlety and complexity of emotions. DES can better describe detailed emotions with subtle differences flexibly, but it is difficult for users to distinguish the absolute continuous values, which may also be problematic. CE and DES are mainly employed in classification and regression tasks, respectively, with discrete and continuous emotion labels. If we discretize DES into several constant values, then we can also use it for classification [66]. Ranking based labeling can be applied to ease DEC comprehension difficulties in raters.

Although less explored in this context, one of the most well-known theories that explains the development of emotional experience is appraisal theory. According to this theory, cognitive evaluation or appraisal of a situation or content in case of multimedia results in emergence of

Table 1. Representative Emotion Models Employed in ACM

Model	Ref	Type	Emotion states/dimensions
Ekman	[34]	CE	happiness, sadness, anger, disgust, fear, surprise
Mikels	[77]	CE	amusement, anger, awe, contentment, disgust, excitement, fear, sadness
Plutchik	[97]	CE	(×3 scales) anger, anticipation, disgust, joy, sadness, surprise, fear, trust
Clusters	[50]	CE	29 discrete labels are consistently grouped into 5 clusters at a similar distance level
Sentiment		CE	positive, negative, (or neutral)
VA(D)	[104]	DES	valence-arousal(-dominance)
ATW	[66]	DES	activity-temperature-weight

emotions [88, 102]. According to Ortony, Clore, and Collins (OCC) [88], emotions are experienced following a scenario composed of a series of phases. First, there is a perception of an event, object, or an action. Then, there is an evaluation of events, objects, or action according to personal wishes or norms. Finally, perception and evaluation result in a specific emotion or emotions arising. Certain appraisal dimensions such as novelty and complexity can be labeled and detected from content. For example, Soleymani automatically recognized image novelty and complexity that are related to interestingness. There are also domain-specific emotion taxonomy and scales. Geneva Emotional Music Scale [153] is a music-specific emotion model for describing emotions induced by music. It consists of a hierarchical structure with 45 emotions, 9 emotional categories, and 3 superfactors that can describe emotion in music.

Another relevant concept worth mentioning is that emotion in response to multimedia can be expected, induced, or perceived emotion. Expected emotion is the emotion that the multimedia creator intends to make people feel, perceived emotion is what people perceive as being expressed, while induced/felt emotion is the actual emotion that is felt by a viewer. Discussing the difference or correlation of various emotion models is out of the scope of this article. The typical emotion models that have been widely used in ACM are listed in Table 1.

3 DATASETS

3.1 Datasets for AC of Images

The early datasets for AC of images mainly come from the psychology community with small-scale images. The International Affective Picture System (**IAPS**) is an image set that is widely used in psychology to evoke emotions [62]. Each image that depicts complex scenes is associated with the mean and standard deviation (STD) of VAD ratings in a nine-point scale by about 100 college students. The **IAPSa** dataset is selected from IAPS with 246 images [77] that are labeled by 20 undergraduate students. The **Abstract** dataset consists of 279 abstract paintings without contextual content. Approximately 230 people peer rated these paintings. The Artistic dataset (**ArtPhoto**) includes 806 artistic photographs from a photo sharing site [71] with emotions determined by the artist uploading the photos. The Geneva affective picture database (**GAPEd**) is composed of 520 negative, 121 positive, and 89 neutral images [26]. Besides, these images are also rated with valence and arousal values, ranging from 0 to 100 points. There are 500 abstract paintings in both **MART** and **devArt** datasets, which are collected from the Museum of Modern and Contemporary Art of Trento and Rovereto [3], and the “DeviantArt” online social network [3], respectively.

Recent datasets, especially the large-scale ones, are constructed using images from social networks. The Tweet dataset (**Tweet**) consists of 470 and 113 tweets for positive and negative sentiments, respectively [18]. The **FlickrCC** dataset includes about 500K Flickr creative common (CC) images that are generated based on 1,553 adjective noun pairs (ANPs) [18]. The images are mapped

Table 2. Released and Freely Available Datasets for AC of Images

Dataset	Ref	# Images	Type	# Annotators	Emotion model	Labeling	Labels
IAPS	[62]	1,182	natural	≈100 (half f)	VAD	annotation	average
IAPSa	[77]	246	natural	20 (10f, 10m)	Mikels	annotation	dominant
Abstract	[71]	279	abstract	≈230	Mikels	annotation	dominant
ArtPhoto	[71]	806	artistic	–	Mikels	keyword	dominant
GAPED	[26]	730	natural	60	Sentiment, VA	annotation	dominant, average
MART	[3]	500	abstract	25 (11f, 14m)	Sentiment	annotation	dominant
devArt	[3]	500	abstract	60 (27f, 33m)	Sentiment	annotation	dominant
Tweet	[18]	603	social	9	Sentiment	annotation	dominant
FlickrCC	[18]	≈500,000	social	–	Plutchik	keyword	dominant
Flickr	[144]	301,903	social	6,735	Ekman	keyword	dominant
Emotion6	[95]	1,980	social	432	Ekman+neutral	annotation	distribution
FI	[149]	23,308	social	225	Mikels	annotation	dominant
IESN	[168]	1,012,901	social	118,035	Mikels, VAD	keyword	personalized
FlickrLDL	[141]	10,700	social	11	Mikels	annotation	distribution
TwitterLDL	[141]	10,045	social	8	Mikels	annotation	distribution

Where “Ref” is short for reference, “# Images” and “# Annotators,” respectively, represent the total number of images and annotators (f: female, m: male), “Labeling” represents the method to obtain labels, such as human annotation (Annotation) and keyword searching (Keyword), and “Labels” means the detailed labels in the dataset, such as dominant emotion category (Dominant), average dimension values (Average), personalized emotion (Personalized), and emotion distribution (Distribution).

to the Plutchnik’s Wheel of Emotions with 8 basic emotions, each with 3 scales. The **Flickr** dataset contains about 300K images [144] with the emotion category defined by the synonym word list that has the most same words as the adjective words of an image’s tags and comments. The **FI** dataset consists of 23,308 images that are collected from Flickr and Instagram by searching the emotion keywords [149] and labeled by 225 Amazon Mechanical Turk (MTurk) workers. The number of images in each Mikels emotion category is larger than 1K. The **Emotion6** dataset [95] consists of 1,980 images collected from Flickr with 330 images for each Ekman’s emotion category. Each image was scored by 15 MTurk workers to obtain the discrete emotion distribution information. The **IESN** dataset that is constructed for personalized emotion prediction [168] contains about 1M images from Flickr. Lexicon-based methods and VAD averaging [134] are used to segment the text of metadata from uploaders for expected emotions and comments from viewers for personalized emotions. There are 7,723 active users with more than 50 involved images. We can also easily obtain the DEC and emotion distribution for each image. **FlickrLDL** and **TwitterLDL** datasets [141] are constructed for discrete emotion distribution learning. The former one is a subset of FlickrCC, which are labeled by 11 viewers. The latter one consists of 10,045 images that are collected by searching various sentiment key words from Twitter and labeled by 8 viewers. These datasets are summarized in Table 2.

3.2 Datasets for AC of Music

A notable benchmark for music recognition is music mood classification (AMC) task, organized by the annual Music Information Retrieval Evaluation eXchange⁴ (MIREX) [32]. In MIREX mood classification task, initially 600 songs were shared with the participants. Starting from 2013, 1,438 30-second excerpts from Korean pop songs have been added to MIREX. MIREX benchmark aims

⁴<http://www.music-ir.org/mirex/wiki/>.

Table 3. Released and Freely Available Datasets for Music Emotion Recognition

Dataset	Ref	# Songs	Type	# Annotators	Emotion model	Labeling	Labels
MIREX mood	[32]	2,038	western and kpop	2–3	Clusters	annotation	dominant, distribution
CAL500	[128]	500	western	>2	–	annotation	dominant
Soundtracks	[33]	110	instrumental	110	self-defined, VA	annotation	distribution
MoodSwings	[121]	240	western	Unknown	VA	annotation	distribution
DEAP	[61]	120	western	14–16	VAD	annotation	average
AMG1608	[23]	1,608	western	15	VA	annotation	distribution
DEAM	[6]	1,802	diverse	5–10	VA	annotation	average, distribution
PMemo	[155]	794	western	10	VA	annotation	distribution

Where “# Songs” and “# Annotators,” respectively, represent the total number of songs and annotators per song, “Labeling” represents the method to obtain labels, such as human annotation (Annotation), and “Labels” means the detailed labels in the dataset, such as dominant emotion category (Dominant), average dimension values (Average), personalized emotion (Personalized), and emotion distribution (Distribution).

to automatically classify songs into five emotion clusters derived from cluster analysis of online tags. MIREX mood challenge emotional representation has been debated in the literature due to its data-driven origin rather than psychology of emotion. For example, in Reference [65], semantic and acoustic overlaps have been found between clusters. MIREX mood challenge considers only one label for the whole song and disregards the dynamic time evolving nature of music.

Computer Audition Lab 500 (CAL500) is a dataset of 500 popular songs that is labeled by multiple tags including emotions [128]. The dataset is labeled in the lab by 66 labelers. Soundtrack datasets [33] for music and emotion is developed by Eerola and Vuoskoski and contains instrumental music from soundtracks of 60 movies. The expert annotators selected songs based on five basic discrete categories (*anger, fear, sadness, happiness, and tenderness*) and dimensional VA representation of emotions. Although not developed with music content analysis in mind, the Database for Emotion Analysis using Physiological Signals, or DEAP, dataset [61] also includes valence, arousal, and dominance ratings for 120 one-minute music video clips of western pop music. Each video clip is annotated by 14–16 participants who were asked to report their felt valence, arousal, and dominance on a nine-point scale. AMG1608 [23] is another music dataset that contains arousal and valence ratings for 1,608 Western songs in different genres and is annotated through MTurk.

Music datasets with emotion labels usually consider one emotion label per song (static). MoodSwings dataset [121] was the first to annotate music dynamically over time. MoodSwings was developed by Schmidt et al. and includes 240 15-second excerpts of western pop songs with per-second valence and arousal labels, collected on MTurk. The MediaEval “Emotion in Music” challenge was organized between 2013 and 2015 in MediaEval Multimedia Evaluation initiative.⁵ MediaEval is a community-driven benchmarking campaign dedicated to evaluating algorithms for social and human-centered multimedia access and retrieval [63]. Unlike MIREX, “Emotion in Music” task focused on dynamic emotion recognition in music tracking arousal and valence over time [6, 115]. The data from MediaEval tasks were compiled in MediaEval Database for Emotional Analysis in Music (DEAM), which is the largest available dataset with dynamic annotations, at 2 Hz, with valence and arousal annotations for 1,802 songs and song excerpts licensed under Creative Commons license. PMemo is a dataset of 794 songs with dynamic and static arousal and valence annotations in addition to electrodermal responses from 10 participants [155].

These datasets are summarized in Table 3. For a more detailed review of available music datasets with emotional labels, we refer the readers to Reference [89].

⁵<http://www.multimediaeval.org>.

3.3 Datasets for AC of Videos

The target of video affective content computing is to recognize the emotions evoked by videos. In this field, it is necessary to construct a large benchmark dataset with precise emotional tags. However, the majority of existing research evaluates their proposed methods on their own collected datasets. The scarce video resources in those self-collected datasets, combined with the copyright restrictions, result in limited accessibility for other researchers to reproduce existing work. Therefore, it is beneficial to summarize some publicly available datasets in this field. In general, publicly available datasets can be classified into two types: datasets consisting of videos only, such as movie clips or user generated videos, and datasets including both videos and audience information.

3.3.1 Datasets Consisting of Videos Only. The LIRIS-ACCEDE dataset [30] is one of the largest datasets in this area. Because it is collected under Creative Commons licenses, there are no copyright issues. The LIRIS-ACCEDE dataset is a living database in development. To fulfill requirements of different tasks, new data, features, and tags are included. The LIRIS-ACCEDE dataset includes the Discrete LIRIS-ACCEDE collection and the Continuous LIRIS-ACCEDE collection, which was used for the MediaEval Emotional Impact of Movies tasks from 2015 to 2018.

The Discrete LIRIS-ACCEDE collection [13] includes 9,800 clips derived from 40 feature films and 120 short films. Specifically, the majority of the 160 films are collected from the video platform VODO. The duration of all 160 films is about 73.6 hours in total. All of the 9,800 video clips last about 27 hours in total, and the duration of each clip is between 8 and 12 seconds, which is long enough to induce emotions in viewers to feel emotions. In this collection, all the 9,800 video clips are labeled by values of valence and arousal.

The Continuous LIRIS-ACCEDE collection [12] differs from the Discrete LIRIS-ACCEDE collection in annotation type. Roughly speaking, the annotations for movie clips in the Discrete LIRIS-ACCEDE collection are global. It means that a whole 8-to-12-second video clip is represented by a single value of valence and arousal. This annotation type limits the possibility for tracking emotions. To address this issue, 30 longer films are selected from the 160 films mentioned above. The total duration of all the selected films is about 7.4 hours. There are emotional annotations according to valence and arousal of each second of the films in the collection.

The MediaEval Affective Impact of Movies collections between 2015 and 2018 are used for the MediaEval affective Impact of Movies tasks in each corresponding year. Specifically, the MediaEval 2015 Affective Impact of Movies [112] includes two sub-tasks: affect detection and violence detection. The Discrete LIRIS-ACCEDE collection was used as the development set. And 1,100 additional video clips were extracted from 39 new movies and included. Indeed, all the new collected data were shared under Creative Commons licenses. In addition, three values were used to label the 10,900 video clips: a binary signal representing the presence of violence, a class tag of the excerpt for felt arousal, and an annotation for felt valence.

The MediaEval 2016 Affective Impact of Movies Task [27] also includes two sub-tasks: Global emotion prediction and Continuous emotion prediction. The Discrete LIRIS-ACCEDE collection and the Continuous LIRIS-ACCEDE collection were used as the development sets for the first and second sub-tasks, respectively. In addition, 49 new movies were chosen as the test sets. 1,200 short video clips from the new movies were extracted for the first task, and 10 long movies were selected for the second task. For the first sub-task, the tags include scores of valence and arousal for each whole movie clip. And for the second sub-task, scores of valence and arousal for each second of the movies are evaluated.

The MediaEval 2017 Affective Impact of Movies Task [28] is focused on long movies for two sub-tasks: valence/arousal prediction and Fear prediction. The Continuous LIRIS-ACCEDE collection was selected as the development set, and an additional 14 new movies were collected as the test set.

The annotations contain a valence value and an arousal value. In addition, there is a binary value to represent whether the segment is supposed to induce fear or not for each 10-second segment.

The MediaEval 2018 Affective Impact of Movies task [29] is also dedicated to valence/arousal prediction and fear prediction. The Continuous LIRIS-ACCEDE collection and the test set of the MediaEval 2017 Emotional Impact of Movies task were used as the development set. In addition, 12 other movies selected from the set of the 160 movies mentioned in the Discrete LIRIS-ACCEDE part were used as test set. Specifically, for the first sub-task, there are annotations containing valence and arousal values for each second of the movies; and the beginning and ending times of each sequence in movies that induce fear are recorded for the second sub-task.

The VideoEmotion dataset [55] is a well-designed user-generated video collection. It contains 1,101 videos downloaded from web platforms, such as YouTube and Flickr. The annotations of the videos in this dataset are based on Plutchik's wheel of emotions [97].

Both the YF-E6 Dataset and the VideoStory-P14 Dataset are introduced in Reference [136]. To collect the YF-E6 emotion dataset, six basic emotion types are used as keywords to search videos on YouTube and Flickr. There are 3K videos collected in the YF-E6 dataset totally. Then there were 10 annotators performing the labeling tasks. Only when all tags for a video clip were more than 50 percent consistent, the video clip was added to the dataset. Finally, the dataset includes 1,637 videos labeled with six basic emotion types. The VideoStory-P14 Dataset is based on the VideoStory dataset. Similar to the VideoEmotion Dataset, the keywords in Plutchik's Wheel of Emotions were used for the search process of the construction of the VideoStory dataset. Finally, there are 626 videos in the videoStory-P14 dataset with each having a unique emotion tag.

3.3.2 Datasets Including Both Videos and Audience's Reactions. The DEAP dataset [61] includes the electroencephalography (EEG) and peripheral physiological signals that are collected from 32 participants during watching 40 one-minute-long excerpts of music videos. In addition, frontal face videos collected from 22 of the 32 participants are gathered. Annotators labeled each video according to the level of like/dislike, familiarity, arousal, valence, and dominance. Though the DEAP dataset is publicly available, it should be noted that it does not include the actual videos because of the licensing issues, but the links of videos are provided.

The MAHNOB-HCI [119] is a multimodal dataset including multi-class information recorded in response to video affective stimuli. Particularly, speeches, face videos, and eye gazes are recorded. In addition, two experiments were conducted to record both peripheral and central nervous system physiological signals from 27 subjects. In the first experiment, subjects were assigned to report their emotional responses to 20 affective induced videos, including the level of arousal, valence and dominance, and predictability as well as emotion categories. In the second experiment, the participants evaluated whether they agreed with the displayed labels after watching short videos and images. The dataset is available for academic use through a web-interface.

The DECAF dataset [1] consists of Infra-red facial video signals, Electrocardiogram (ECG), Magnetoencephalogram (MEG), horizontal Electrooculogram (hEOG), and Trapezius Electromyogram (tEMG), recorded from 30 participants watching 36 movie clips and 40 one-minute music videos, which are derived from the DEAP dataset [61]. The subjective feedback is based on valence, arousal, and dominance space. In addition, time-continuous emotion annotations for movie clips are also included in the dataset.

The AMIGOS dataset [24] includes multi-class affective data, individual and groups of viewers' responses to both short and long videos. The EEG, ECG, GSR, frontal, and full body video were recorded in two experimental settings, i.e., 40 participants watching 16 short emotional clips and 4 long clips. The duration of each selected short videos is between 51 and 150 seconds, and the duration of each long excerpt is about 20 minutes. Finally, participants annotated the affective level of valence, arousal, control, familiarity, liking, and basic emotions.

Table 4. Released and Freely Available Datasets for Video Emotion Recognition

Dataset	Ref	#Clips	Hours	Type	# Annotators	Emotion model	Labeling	Labels
Discrete LIRIS-ACCEDE	[13]	9,800	26.9	film	–	VA	annotation	dominant
Continuous LIRIS-ACCEDE	[12]	30	7.4	film	10 (7f, 3m)	VA	annotation	average
MediaEval 2015	[112]	1,100	–	film	–	3 discrete VA values	annotation	dominant
MediaEval 2016	[27]	1,210	–	film	–	VA	annotation	distribution, average
MediaEval 2017	[28]	14	8	film	–	VA, fear	annotation	average
MediaEval 2018	[29]	12	9	film	–	VA, fear	annotation	average
VideoEmotion	[55]	1,101	32.7	user-generated	10 (5f, 5m)	Plutchik	annotation	dominant
YF-E6	[136]	1,637	50.9	user-generated	10(5f, 5m)	Emkan	annotation	dominant
VideoStory-P14	[136]	626	–	user-generated	–	Plutchik	keyword	dominant
DEAP	[61]	120	2	music video	–	VAD	annotation	personalized
MAHNOB-HCI	[119]	–	–	multiple types	–	VAD, Ekman+neutral	annotation	personalized
DECAF	[1]	76	–	music video/movies	–	VAD	annotation	personalized
AMIGOS	[24]	20	–	movies collection	–	VAD, Ekman	annotation	personalized
ASCERTAIN	[123]	36	–	movies collection	58 (21f, 37m)	VA	annotation	personalized

Where “#Clips” and “Hours,” respectively, represent the total number and hours of video clips, “Type” means the genre of the videos in the dataset, “Emotion Model” represents the labeling type, “Labeling” represents the method to obtain labels, such as human annotation (Annotation) and keyword searching (Keyword), and “Labels” means the detailed labels in the dataset, such as dominant emotion category (Dominant), average dimension values (Average), personalized emotion (Personalized), and emotion distribution (Distribution).

Big-five personality scales and affective self-ratings of 58 users together with their EEG, ECG, GSR, and facial activity data were included in the ASCERTAIN dataset [123]. The number of videos used as the stimulus is 36 and the length of each video clip is between 51 and 128 seconds. It is the first physiological dataset that is useful for both affective and personality recognition.

The publicly available datasets for video affective content analysis are summarized in Table 4.

3.4 Datasets for AC of Multimodal Data

In addition to audiovisual content and viewers’ reactions, other modalities, such as language, also contain significant information for affective understanding of multimedia content.

Visual sentiment is the sentiment associated with the concepts depicted in images. Two datasets were developed through mining images associated with the adjective-noun pair (ANP) representations that have affective significance [17]. ANPs in Reference [17] were generated by first using seed terms from Plutchik’s Wheel of Emotion [97] to query Flickr⁶ and YouTube.⁷ After mining the tags associated with visual content on YouTube and Flickr, adjective and noun candidates were identified through part-of-speech tagging. Then adjective and nouns were paired to create ANP candidates that were filtered by sentiment strength, named entities, and popularity. The Visual Sentiment Ontology (VSO), Reference [17]⁸, are the results of this process. Sentibank resulted in the creation of a set of photo-tweet sentiment dataset, with both visual and textual data with polarity labels, collected on Amazon Mechanical Turk.⁹ This work was later extended to form a multilingual ANP set and its dataset, in References [25, 57],¹⁰ contains 15,630 ANPs from

⁶<https://www.flickr.com>.

⁷<https://www.youtube.com/>.

⁸<https://visual-sentiment-ontology.appspot.com>.

⁹http://www.ee.columbia.edu/in/dvmm/vso/download/twitter_dataset.html.

¹⁰<http://mvso.cs.columbia.edu>.

Table 5. Released and Freely Available Datasets for Multimodal Multimodal Emotion Recognition

Dataset	Ref	#Samples	Modalities	Type	Emotion model	Labeling	Labels
SentiBank tweet	[17]	603	images, text	images	Sentiment	annotation	dominant
MVSO	[57]	7.36M	image, metadata	photos	Sentiment	automatic	average
CMU-MOSEI	[152]	23,500	video, audio, text	YouTube videos	Sentiment	annotation	average
MELD	[98]	13,000	video, audio, text	TV series	Sentiment, Disc.	annotation	dominant
COGNIMUSE	[178]	3.5h	video, audio, text	movies	VA	annotation, self-report	dominant
VR	[68]	73	video, audio	VR videos	VA	self-report	average

Disc. for MELD corresponds to six Ekman emotions in addition to neutral. “Labeling” represents the method to obtain labels, such as human annotation (annotation), self-reported felt emotion, and keyword searching (keyword), “Labels” means the detailed labels in the dataset, such as dominant emotion category (dominant), average dimension values (average), personalized emotion (personalized), and emotion distribution (distribution).

12 major languages and 7.37M images [58]. My Reaction When (MRW) dataset contains 50,107 video-sentence pairs crawled from social media, depicting physical or emotional reactions to the situations described in sentences [120]. The GIFs are sourced from Giphy.¹¹ Even though there are no emotional labels, the language and visual associations are mainly based on sentiment, which makes this dataset an interesting resource for affective content analysis.

CMU Multimodal Opinion Sentiment and Emotion Intensity (CMU-MOSEI) is a collection of multiple datasets for multimodal sentiment analysis and emotion recognition. This collection includes more than 23,500 sentence utterance videos from more than 1K people from YouTube [152].¹² All the videos are transcribed and aligned with audiovisual modalities. A multimodal multi-party dataset for emotion recognition in conversation (MELD) was primarily developed for emotion recognition in multiparty interaction purposes [98]. MELD contains visual, audio, and textual modalities and includes 13K utterances from 1,433 dialogues from the TV series *Friends*, with each utterance labeled with emotion and sentiment.

COGNIMUSE is a collection of videos annotated with sensory and semantic saliency, events, cross-media semantics, and emotions [178]. A subset of 3.5 hours extracted from movies, including textual modality, are annotated on arousal and valence. Li et al. collected a dataset of 360-degree virtual reality videos that can elicit different emotions [68]. Even though the dataset consists of 73 short videos, on average 183 seconds long, it is one of the first datasets of its kind whose content understanding stays limited. These multimodal datasets are summarized in Table 5.

4 AFFECTIVE COMPUTING OF IMAGES

In the early stages, AC researchers mainly worked on designing handcrafted features to bridge the affective gap. Recently, with the advent of deep learning, especially convolutional neural networks (CNNs), current methods have shifted to an end-to-end deep representation learning. Motivated by the fact that the perception of image emotions may be dependent on different types of features [171], some methods employ fusion strategies to jointly consider multiple features. In this section, we summarize and compare these methods. Please note that here we classify the directly extracted CNN features based on pre-trained deep models into handcrafted features category.

4.1 Handcrafted Features-based Methods for AC of Images

Low-level Features are difficult to be understood by viewers. These features are often directly derived from other computer vision tasks. Some widely extracted features include GIST, HOG2x2,

¹¹<https://giphy.com/>.

¹²<https://github.com/A2Zadeh/CMU-MultimodalSDK>.

self-similarity, and geometric context color histogram features as in Reference [94], because of their individual power and distinct description of visual phenomena in a scene perspective.

Compared with the above generic features, some specific features derived from art theory and psychology have been designed. For example, Machajdik and Hanbury [71] extracted elements-of-art features, including *color* and *texture*. The MPEG-7 visual descriptors are employed in Reference [66], which include four color-related ideas and two texture-related ideas. How shape features in natural images influence emotions is investigated in Reference [70] by modeling the concepts of roundness-angularity and simplicity-complexity. Sartori et al. [103] designed two kinds of visual features to represent different color combinations based on Itten's color wheel.

Mid-level Features contain more semantics, are more easily interpreted by viewers than low-level features, and thus are more relevant to emotions. Patterson and Hays [94] proposed to detect 102 attributes in five different categories, including materials, surface properties, functions or affordances, spatial envelop attributes, and object presence. Besides these attributes, eigenfaces that may contribute to facial images are also incorporated in Reference [150]. More recently, in Reference [100], SIFT features are first extracted as basic features, which are fed into bag-of-visual-words (BoVW) to represent the multi-scale blocks. Another mid-level representation is the latent topic distribution estimated by probabilistic latent semantic analysis.

Harmonious composition is essential in an artwork. Several compositional features, such as low depth of field, are designed to analyze such characteristics of an image [71]. Based on the fact that figure-ground relationships, color patterns, shapes, and their diverse combinations are often jointly employed by artists to express emotions in their artworks, Wang et al. [133] proposed to extract interpretable aesthetic features. Inspired by principles-of-art, Zhao et al. [162] designed corresponding mid-level features, including *balance*, *emphasis*, *harmony*, *variety*, *gradation*, and *movement*. For example, Itten's color contrasts and the rate of focused attention are employed to measure *emphasis*.

High-level Features that represent the semantic content contained in images can be easily understood by viewers. We can also well recognize the conveyed emotions in images through these semantics. In the early years, simple semantic content including faces and skins contained in images are extracted in Reference [71]. For the images that contain faces, facial expressions may directly determine the emotions. Yang et al. [142] extracted eight kinds of facial expressions as high-level features. They built compositional features of local Haar appearances by a minimum error-based optimization strategy, which are embedded into an improved AdaBoost algorithm. For the images detected without faces, the expressions are simply set as *neutral*. Finally, they generated an eight-dimensional vector with each element representing the number of corresponding facial expressions.

More recently, the semantic concepts are described by adjective noun pairs (ANPs) [18, 21], which are detected by SentiBank [18] or DeepSentiBank [21]. The advantages of ANP are that it turns a neutral noun into an ANP with strong emotions and makes the concepts more detectable, compared to nouns and adjectives, individually. A 1,200-dimensional vector representing the probability of the ANPs can form a feature vector.

Table 6 summarizes the above-mentioned handcrafted features at different levels for AC of images. Some recent methods also extracted CNN features from pre-trained deep models, such as AlexNet [157, 158] and VGGNet [141].

To map the extracted handcrafted features to emotions, **Machine Learning Methods** are commonly employed. Some typical learning models include Naive Bayes (NB) [71, 133], support vector machine (SVM) [70, 150, 162], K nearest neighbor (KNN) [66], sparse learning (SL) [67, 103], logistic regression (LR) [150], multiple instance learning (MIL) [100], and matrix completion (MC) [3]

Table 6. Summary of the Handcrafted Features at Different Levels for AC of Images

Feature	Ref	Level	Short description	# Feat
LOW_C	[94]	low	GIST, HOG2x2, self-similarity, and geometric context color histogram features	17,032
Elements	[71]	low	color: mean saturation, brightness and hue, emotional coordinates, colorfulness, color names, Itten contrast, Wang's semantic descriptions of colors, area statistics; texture: Tamura, Wavelet and gray-level co-occurrence matrix	97
MPEG-7	[66]	low	color: layout, structure, scalable color, dominant color; texture: edge histogram, texture browsing	≈200
Shape	[70]	low	line segments, continuous lines, angles, curves	219
IttenColor	[103]	low	color co-occurrence features and patch-based color-combination features	16,485
Attributes	[94]	mid	scene attributes	102
Sentributes	[150]	mid	scene attributes, eigenfaces	109
Composition	[71]	mid	level of detail, low depth of field, dynamics, rule of thirds	45
Aesthetics	[133]	mid	figure-ground relationship, color pattern, shape, composition	13
Principles	[162]	mid	principles-of-art: balance, contrast, harmony, variety, gradation, movement	165
BoVW	[100]	mid	bag-of-visual-words on SIFT, latent topics	330
FS	[71]	high	number of faces and skin pixels, size of the biggest face, amount of skin w.r.t. the size of faces	4
ANP	[18]	high	semantic concepts based on adjective noun pairs	1,200
Expressions	[142]	high	automatically assessed facial expressions (anger, contempt, disgust, fear, happiness, sadness, surprise, neutral)	8

"# Feat" indicates the dimension of each feature.

for emotion classification, support vector regression (SVR) [70, 162] for emotion regression, and multi-graph learning (MGL) [171] for emotion retrieval.

Instead of assigning the DEC to an image, some recent methods began to focus on the perception subjectivity challenge, i.e., predicting personalized emotions for each viewer or learning emotion distributions for each image. The personalized emotion perceptions of a specified user after viewing an image is predicted in References [166, 168], associated with online social networks. They considered different types of factors that may contribute to emotion recognition, including the images' visual content, the social context related to the corresponding users, the emotions' temporal evolution, and the images' location information. To jointly model these factors, they proposed rolling multi-task hypergraph learning (RMTHG), which can also easily handle the data incompleteness issue.

Generally, the distribution learning task can be formulated as a regression problem, which slightly differs for different distribution categories (i.e., discrete or continuous). For example, if emotion is represented by CE, then the regression problem targets predicting the discrete probability of each emotion category with the sum equal to 1; if we represent emotion based on DES, then the regression problem is typically transformed to the prediction of the parameters of specified continuous probability distributions. For the latter one, we usually need to first determine the form of continuous distributions, such as exponential distribution and Gaussian distribution. Some representative learning methods for emotion distribution learning of discrete emotions include shared sparse learning (SSL) [169], weighted multimodal SSL (WMMSSL) [157, 159], augmented conditional probability neural network (ACPNN) [141], and weighted multi-model CPNN (WMMCPNN) [158]. Both SSL and WMMSSL can only model one test image each time, which is computationally inefficient. After the parameters are learned, ACPNN and WMMCPNN can easily predict the emotion distributions of a test image. Based on the assumption that the VA emotion labels can be well modeled by a mixture of two bidimensional Gaussian mixture models (GMMs),

Table 7. Representative Work on AC of Images Using Handcrafted Features

Ref	Feature	Fusion	Learning	Dataset	Task	Result
[71]	Elements, Composition, FS	early	NB	IAPSa, Abstract, ArtPhoto	cla	0.471, 0.357, 0.495
[66]	MPEG-7	–	KNN	unreleased	cla	0.827
[70]	Shape, Elements	early	SVM, SVR	IAPSa; IAPS	cla; reg	0.314; V-1.350, A-0.912
[67]	Segmented objects	–	SL	IAPS, ArtPhoto	cla	0.612, 0.610
[150]	Sentributes	–	SVM, LR	Tweet	cla	0.824
[133]	Aesthetics	–	NB	Abstract, ArtPhoto	cla	0.726, 0.631
[162]	Principles	–	SVM, SVR	IAPSa, Abstract, ArtPhoto; IAPS	cla; reg	0.635, 0.605, 0.669; V-1.270, A-0.820
[171]	LOW_C, Elements, Attributes, Principles, ANP, Expressions	graph	MGL	IAPSa, Abstract, ArtPhoto, GAPEd, Tweet	ret	0.773, 0.735, 0.658, 0.811, 0.701
[103]	IttenColor	–	SL	MART, devArt	cla	0.751, 0.745
[100]	BoVW	–	MIL	IAPSa, Abstract, ArtPhoto	cla	0.699, 0.636, 0.707
[3]	IttenColor	–	MC	MART, devArt	cla	0.728, 0.761
[168]	GIST, Elements, Attributes, Principles, ANP, Expressions	graph	RMTHG	IESN	cla_p	0.582
[169]	GIST, Elements, Principles	–	SSL	Abstract	dis_d	0.134
[157]	GIST, Elements, Attributes, Principles, ANP, deep features from AlexNet	weighted	WMSSL	Abstract, Emotion6, IESN	dis_d	0.482, 0.479, 0.478
[141]	ANP, VGG16	–	ACPNN	Abstract, Emotion6, FlickrLDL, TwitterLDL	dis_d	0.480, 0.506, 0.469, 0.555
[158]	GIST, Elements, Attributes, Principles, ANP, AlexNet	weighted	WMCPNN	Abstract, Emotion6, IESN	dis_d	0.461, 0.464, 0.470
[167]	GIST, Elements, Attributes, Principles, ANP, AlexNet	–	MTSSR	IESN	dis_c	0.436

Where “Fusion” indicates the fusion strategy of different features, “Cla, Reg, Ret, cla_p, dis_d, dis_c” in the task column are short for classification, regression, retrieval, personalized classification, discrete distribution learning, continuous distribution learning (the Same Below), respectively, “Result” is the reported best accuracy for classification, mean squared error for regression, discounted cumulative gain for retrieval, F1 for personalized classification, and KL divergence for distribution learning (the First Line [169] is the result on sum of squared difference) on the corresponding datasets.

Zhao et al. [167] proposed to learn continuous emotion distributions in VA space by multi-task shared sparse regression (MTSSR). Specifically, the parameters of GMMs are regressed, including the mean vector and covariance matrix of the two Gaussian components as well as the mixing coefficients.

Table 7 summarizes some representative work based on handcrafted features. Generally, high-level features (such as ANP) can achieve better recognition performance for images with rich semantics, mid-level features (such as Principles) are more effective for artistic photos, while low-level features (such as Elements) perform better for abstract paintings.

4.2 Deep Learning-based Methods for AC of Images

To deal with the situation where images are weakly labeled, a potentially cleaner subset of the training instances are selected progressively [148]. First, they trained an initial CNN model based on the training data. Second, they selected the training samples with distinct sentiment scores between the two classes with a high probability based on the prediction score of the trained model on the training data itself. Finally, the pre-trained AlexNet on ImageNet is fine-tuned to classify emotions into 8 categories by changing the last layer of the CNN from 1K to 8 [149]. Besides using

the fully connected layer as classifier, they also trained an SVM classifier based on the extracted features from the second to the last layer of the pre-trained AlexNet model.

Multi-level deep representations (MldrNet) are learned in Reference [101] for image emotion classification. They segmented the input image into three levels of patches, which are input to three different CNN models, including Alexnet, aesthetics CNN (ACNN), and texture CNN (TCNN). The fused features are fed into multiple instance learning (MIL) to obtain the emotion labels. Zhu et al. [175] proposed to integrate the different levels of features by a Bidirectional GRU model (BiGRU) to exploit their dependencies based on MldrNet. They generated two features from the Bi-GRU model and concatenated them as the final feature representations. To enforce the feature vectors extracted from each pair of images from the same category to be close enough, and those from different categories to be far away, they proposed to jointly optimize a contrastive loss together with the traditional cross-entropy loss.

More recently, Yang et al. [138] employed deep metric learning to explore the correlation of emotional labels with the same polarity, and proposed a multi-task deep framework to optimize both retrieval and classification tasks. By considering the relations among emotional categories in the Mikels' wheel, they jointly optimized a novel sentiment constraint with the cross-entropy loss. Extending triplet constraints to a hierarchical structure, the sentiment constraint employs a sentiment vector based on the texture information from the convolutional layer to measure the difference between affective images. In References [105, 139], Yang et al. proposed a weakly supervised coupled convolutional neural network to exploit the discriminability of localized regions for emotion classification. Based on the image-level labels, a sentiment map is first detected in one branch with the cross spatial pooling strategy. And then the holistic and localized information are jointly combined in the other branch to conduct a classification task. The detected sentiment map can easily explain which regions of an image determine the emotions.

The above deep methods mainly focused on the dominant emotion prediction. There are also some works on emotion distribution learning based on deep models. The very first work is a mixed bag of emotions, which trains a deep CNN regressor (CNNR) for each emotion category in Emotion6 [95] based on the AlexNet architecture. They changed the number of output nodes to 1 to predict a real value for each emotion category and replaced the Softmax loss with Euclidean loss. To ensure the sum of different probabilities to be 1, they normalized the predicted probabilities of all emotion categories. However, CNNR has some limitations. First, the predicted probability cannot be guaranteed to be non-negative. Second, the probability correlations among different emotions are ignored, since the regressor for each emotion category is trained independently. In Reference [140], Yang et al. designed a multi-task deep framework based on VGG16 by jointly optimizing the cross-entropy loss for emotion classification and Kullback-Leibler (KL) divergence loss for emotion distribution learning. To match the single emotion dataset to emotion distribution learning settings, they transformed each single label to emotion distribution with emotion distances computed on Mikels' wheel [166, 168]. By extending the size of training samples, this method achieves the state-of-the-art performance for discrete emotion distribution learning.

The representative deep learning-based methods are summarized in Table 8. The deep representation features generally perform better than the handcrafted ones, which are intuitively designed for specific domains based on several small-scale datasets. However, how the deep features correlate to specific emotions is unclear.

5 AFFECTIVE COMPUTING OF MUSIC

Music emotion recognition (MER) strives to identify emotion expressed by music and subsequently predict listeners' felt emotion from acoustic content and music metadata, e.g., lyrics, genre. Emotional understanding of music has applications in music recommendation and is particularly useful

Table 8. Representative Work on Deep Learning-based AC of Images

Ref	Base net	Pre	#Feat	Cla	Loss	Dataset	Task	Result
[148]	self-defined	no	24	–	–	FlickrCC	cla	0.781
[149]	AlexNet	yes	4,096	SVM	–	FI, IAPSa, Abstract, ArtPhoto	cla	0.583, 0.872, 0.776, 0.737
[101]	AlexNet, ACNN, TCNN	yes	4,096, 256, 4,096	MIL	–	FI, IAPSa, Abstract, ArtPhoto, MART	cla	0.652, 0.889, 0.825, 0.834, 0.764
[175]	self-defined	no	512	–	contrastive	FI, IAPSa, ArtPhoto	cla	0.730, 0.902, 0.855
[138]	GoogleNet-Inception	yes	1,024	–	sentiment	FI, IAPSa, Abstract, ArtPhoto	cla; ret	0.676, 0.442, 0.382, 0.400; 0.780, 0.819, 0.788, 0.704
[105, 139]	ResNet-101	yes	2,048	–	–	FI, Tweet	cla	0.701, 0.814
[95]	AlexNet	yes	4,096	–	Euclidean	Emotion6	dis_d	0.480
[140]	VGG16	yes	4,096	–	KL	Emotion6, FlickrLDL, TwitterLDL	dis_d	0.420, 0.530, 0.530

Where “Pre” Indicates Whether the Network Is Pre-trained Using ImageNet, “# Feat” Indicates the Dimension of Last Feature Mapping Layer Before the Emotion Output Layer, “Cla” Indicates the Classifier Used After the Last Feature Mapping with Default Softmax, “Loss” Indicates the Loss Objectives (Besides the Common Cross-entropy Loss for Classification), and “Result” is the Reported Best Accuracy for Classification, Discounted Cumulative Gain for Retrieval, and KL Divergence for Distribution Learning on the Corresponding Datasets.

Table 9. Musical Dimensions and Acoustic Features Describing Them

Musical dimension	Acoustic features
Melody	Pitch
Harmony	chromagram, chromagram peak, key, mode, key clarity, harmonic, change, chords
Rhythm	tempo, beat histograms, rhythm regularity, rhythm strength, onset rate
Dynamics and loudness	RMS energy, loudness, timbral width
Timbre	MFCC, spectral shapres (centroid, shape, spread, skewness, kurtosis, contrast, and flatness), brightness, rolloff frequency, zero crossing rate, spectral contrast, auditory modulation features, inharmonicity, roughness, dissonance, odd to even harmonic ratio [4]
Musical form	Similarity Matrix (similarity between all possible frames) [89]
Texture	attack slope, attack time

for producing music retrieval. An analysis of search queries from creative professionals showed that 80% contain emotional terms, showing emotions prominence in that field [52]. A growing number of works have tried to address emotional understanding of music from acoustic content and metadata (see References [60, 145] for earlier reviews on this topic).

Earlier work on emotion recognition from music relied on extracting acoustic features similar to the ones used in speech analysis, such as audio energy and formants. Acoustic features describe attributes related to musical dimensions. Musical dimensions include melody, harmony, rhythm, dynamics, timbre (tone color), expressive techniques, musical texture, and musical form [89], as shown in Table 9. Some also add energy as a musical feature, which is important for MER [145]. Melody is a linear succession of tones and can be captured by features representing key, pitch, and tonality. Among others, chroma is often used to represent melodic features [145]. Harmony is how the combination of various pitches are processed during hearing. Understanding harmony involves chords or multiple notes played together. Examples of acoustic features capturing harmony include chromagram, key, mode, and chords [4]. Rhythm consists of repeated patterns of musical sounds, i.e., notes and pulses that can be described in terms of tempo and meter. Higher tempo songs often induce higher arousal and fluent rhythm is associated with higher valence and

firm rhythm is associated with sad songs [145]. Mid-level acoustic features, such as onset rate, tempo, and beat histogram, can represent rhythmic characteristics of music. Dynamics of music involve the variation in softness or loudness of notes, which include change of loudness (contrast) and emphasis on individual sounds (accent) [89]. Dynamics of music can be captured by changes in acoustic features related to energy such as root mean square (RMS) energy. Timbre is the perceived sound quality of musical notes. Timbre is what differentiates different voices and instruments playing the same sound. Acoustic features capturing timbre, such as MFCC and spectrum shape, describe sound quality [143]. Acoustic features describing timbre include MFCC, spectral features (centroid, contract, flatness), and zero crossing rate [4]. Expressive techniques are the way a musical piece is played, including tempo and articulation [89]. Acoustic features, such as tempo, attack slope, and time, can be used to describe this dimension. Musical texture is how rhythmic, melodic, and harmonic features are combined in music production [89]. It is related to the range of tones played at the same time. Musical form describes how a song is structured, such as introduction verse and chorus [89]. Energy, whose dynamics are described in music dynamic features, is strongly associated with arousal perception.

There are a number of toolboxes available for extracting acoustic features from music that can be used for music emotion recognition. Music Analysis, Retrieval and Synthesis for Audio Signals (Marsyas) [129] is an open source framework developed in C++ that supports extracting a large range of acoustic features with music information retrieval applications in mind, including time-domain zero-crossings, spectral centroid, rolloff, flux, and Mel-Frequency Cepstral Coefficients (MFCC), and so on. MIRToolbox [64] is an open source toolbox implemented in MATLAB for music information retrieval applications. MIRToolbox offers the ability to extract a comprehensive set of acoustic features at different levels, including features related to tonality, rhythm, and structures. Speech and music interpretation by large-space extraction or OpenSMILE [36, 37] is an open source software developed in C++ with the ability to extract a large number of acoustic features for speech and music analysis in real-time. LibROSA [75] is a Python package for music and audio analysis. It is mainly developed with music information retrieval application in mind and supports importing from different audio sources and extracting musical features such as onsets chroma and tempo in addition to the low-level acoustic features. ESSENTIA [16] is an open source library developed in C++ with Python interface that is developed for audio analysis. ESSENTIA contains an extensive collection of algorithms supporting audio input/output functionality, standard digital signal processing blocks, statistical characterization of data, and a large set of spectral, temporal, tonal, and high-level music descriptors.

Music emotion recognition either attempts to classify songs, excerpt into categories (classification), or estimate their expressed emotions on continuous dimensions (regression). The choice of machine learning model in music emotion recognition depends on the emotional representation used. Mood clusters [32], dimensional representations such as arousal, tension, valence, as well as music-specific emotion representation, can be used. An analysis of the methods proposed for MediaEval “Music in Emotion” task submissions revealed that using deep learning accounted for the superior performance for emotion recognition much more than the choice features [6]. Recent methods for emotion recognition in music rely on deep learning and often use spectrogram features that are converted to images [5]. Aljanaki and Soleymani proposed learning musically meaningful mid-level perceptual features that can describe emotions in music [5]. They demonstrated that perceptual features such as melodiousness, modality, rhythmic complexity, and dissonance can describe a large portion of emotional variance in music both in dimensional representation and MIREX clusters. They also trained a deep convolutional neural network to recognize these mid-level attributes. There have also been works attempting to use lyrics in addition to acoustic

content for recognizing emotion in music [73]. However, lyrics are copyrighted and not easily available, which hinders further work in this direction.

6 AFFECTIVE COMPUTING OF VIDEOS

Currently, the features used in affective video content analysis are mainly from two categories [107, 109]. One is considering the stimulus of video content and extracting the features reflecting the emotions conveyed by the video content itself; the other is extracting features from the viewers. Features extracted from the video content are content-based features, and features formed from the signals of the viewers' responses are viewer-related features.

6.1 Content-related Features

Generally speaking, the video content is composed of a series of ordered frames as well as corresponding audio signals. Therefore, it is natural to extract features from these two modalities. The audiovisual features can further be divided into low-level and mid-level, according to their ability to describe the semantics of video content.

6.1.1 Low-level Features. Commonly, the low-level features are directly computed from the raw visual and audio content, and usually carry no semantic information. As for visual content, color, lighting, and tempo are important elements that can endow the video with strong emotional rendering and further give viewers direct visual stimuli. In many cases, computations are conducted over each frame of the video, and the average values of the computational results of the overall video are considered as visual features. Specifically, the color-related features often contain the histogram and variance of color [20, 107, 177], the proportions of color [82, 86], the number of white frame and fades [177], the grayness [20], darkness ratio, color energy [132, 170], brightness ratio and saturation [85, 86], and so on. In addition, the differences of dark and light can be reflected by the lighting key, which is used to evoke emotions in video and draw the attention of viewers by creating an emotional atmosphere [82]. As for the tempo-related features, properties of shot can reinforce the expression of video, such as shot change rate and shot length variance [20, 85, 87], according to movie grammar. To better take advantage of the temporal information of the video, the motion vectors have been computed as features in Reference [173]. Since the optical flow can characterize the influence of camera motions, the histogram of optimal flow matrix (HOF) has been computed as features in Reference [147]. Additionally, Yi and Wang [147] traced motion key points at multiple spatial scales and computed the mean motion magnitude of each frame as features.

To represent audio content, pitch, zero crossing rates (ZCR), Mel frequency cepstrum coefficients (MFCC), and energy are the most popular features [2, 43, 82, 132, 170]. In particular, the MFCC [14, 49, 74, 147, 176, 177] and its $\Delta MFCC$ are used to characterize emotions in video clips frequently; while the derivatives and statistics (min, max, mean) of MFCC or $\Delta MFCC$ are also explored widely. As for pitch, Reference [87] shows that pitch of sound is associated closely with some emotions, such as anger with higher pitch and sadness with lower standard deviation of pitch. Similar situations can also occur in the energy [86, 87]. For example, the total energy of anger or happiness is higher than the counterpart of unexciting emotions. ZCR [86] is used to separate different types of audio signals, such as music, environmental sound, and speech of human. Besides these frequently related features, audio flatness [177], spectral flux [177], delta spectrum magnitude, harmony [86, 111, 177], band energy ratio, spectral centroid [49, 177], and spectral contrast [86] are also utilized.

Evidently, the aforementioned features are mostly handcrafted. With the emergence of deep learning, features can be automatically learned through deep neural networks. Some pre-trained convolutional neural networks (CNNs) are used to learn static representations from every frame

or some selected key frames, while a Long-Short Term Memory (LSTM) is exploited to capture dynamic representations existing in videos.

For instance, in Reference [136], an AlexNet with seven fully connected layers trained on 2,600 ImageNet classes is used to learn features. A Convolutional Auto-Encoder (CAE) is designed to ensure the CNNs can extract the visual features effectively in Reference [92]. Ben-Ahmed and Huet [14] first used the pre-trained ResNet-152 to extract feature vectors. And then these vectors are fed into an LSTM according to their temporal order to extract high-order representations. Pre-trained model SoundNet is utilized to learn audio features. Because the expressive emotions of video are induced and communicated by the protagonist in video in many cases, the features of protagonist are extracted from the key frame by a pre-trained CNN and used in video affective analysis in References [176, 177]. In addition to the protagonist, other objects in each frame of video also give insights into emotional expression of video. For example, in Reference [110], Shukla et al. removed the non-gaze regions from video frames (Eye ROI) and built the coarse-grained scene structure remaining gist information by Gaussian filter with variance. After the operations above in Reference [110], the next video affective analysis may pay more attention to important information and reduce unnecessary noise.

6.1.2 Mid-level Features. Unlike low-level features, mid-level features often contain semantic information. For example, EmoBase10 feature depicting audio cues is computed in Reference [147]. Hu et al. [49] proposed a method of combining the audio and visual features to model contextual structures of the key frames selected from video. This can produce a kind of so-called multi-instance sparse coding (MI-SC) for next analysis. In addition, the lexical features [81] are extracted from the dialogues of speakers by using a natural language toolkit. These features can reflect the emotional changes in videos and can also represent a certain emotional expression in overall videos. Muszynski et al. [81] used aesthetic movie highlights related to occurrences of meaningful movie scenes to define some experts. These features produced by experts are more knowledgeable and abstract for video affective analysis, especially movies. HHTC features, which are computed on the basis of combination of Huang Transform in visual-audio and cross-correlation features, are proposed in Reference [85].

6.2 Viewer-related Features

Besides the content-related features, viewers' facial expressions and changes of physiological signals evoked by content of videos are the most common sources for extracting viewer-related features. McDuff and Soleymani [74] coded the facial actions of viewers for further affective video analysis. Among various physiological signals, electrocardiography (ECG), galvanic skin response (GSR), and electroencephalography (EEG) are the most frequent ones, and their statistical measures, such as mean, median, spectral power bands, and so on, are often recommended as features. Wang et al. [132] used EEG signals to construct a new EEG feature with the assistance of the relationship among video content by exploiting canonical correlation analysis (CCA). In Reference [40], some viewer-related features are extracted from the whole pupil dilation ratio time-series without the differences among pupil diameter in human eyes, such as its average and derivation for global features as well as the four spectral power bands for local features.

In addition to viewers' responses mentioned above, the comments or other textual information produced by viewers can also reflect their attitudes or emotional reactions toward the videos. In light of this, it is reasonable to consider users' textual comments or other textual information to extract features. In Reference [82], the "sentiment analysis" module using Unigrams and Bigrams is built to learn comment-related features of the collected data according to the YouTube link provided by the DEAP dataset.

6.3 Machine Learning Methods

After feature extracting, a classifier or a regressor is used to obtain emotional analysis results. For classification, there are several frequently used classifiers, including support vector machines (SVM) [2, 42, 83, 109, 130, 147], Naive Bayes (NB) [42, 83], Linear Discriminant Analysis [110], logistic regression (LR) [137], ensemble learning [2], and so on.

Recent works show that the SVM-based methods are very popular for affective video content analysis due to its simplicity, max-margin training property, and use of kernels [131]. For example, Yi and Wang's work [147] demonstrated that linear SVM is more suitable for classification than RBM, MLP, and LR. In Reference [109], LDA, linear SVM (LSVM), and Radial Basis SVM (RSVM) classifiers are employed in emotion recognition experiments, and the RSVM obtained the best F1 scores. In Reference [42], both Naive Bayes and SVM are used as classifiers in unimodal and multimodal conditions. In the unimodal experimental condition, NB is not better than SVM. And the fusion results showed that SVM is much better than NB in multimodal situations. However, SVM also has its shortages, such as the difficulty of selecting suitable kernel functions. Indeed, SVM is not always the best choice. In Reference [2], the results demonstrated that ensemble learning outperforms SVM in terms of classification accuracy. Ensemble learning has acquired a lot of attention in many fields because of its accuracy, simplicity, and robustness. In addition, in Reference [137], LR is adopted as the classifier for its effectiveness and simplicity. In fact, LR is used frequently in many transfer learning tasks.

However, all the classifiers mentioned above are not able to capture the temporal information. Some other methods try to use temporal information. For example, Gui et al. [40] combined SVM and LSTM to predict emotion labels. Specifically, global features and sequence features are proposed to represent the pupillary response signals. Then an SVM classifier is trained with the global features and an LSTM classifier is trained with the sequence features. Finally, a decision fusion strategy is proposed to combine these two classifiers.

A regressor is needed when mapping the extracted features to the continuous dimensional emotion space. Recently, one of the most popular regression methods is support vector regression (SVR) [12, 79, 177]. For example, in Reference [12], video features such as audio, color, and aesthetic are fed into SVR in the SVR-Standard experiment. And in the SVR-Transfer learning experiment, the pre-trained CNN is treated as a feature extractor. The CNN's outputs are used as the input to the SVR. The experimental results showed that the SVR-Transfer learning outperforms other methods. Indeed, the various kernel functions in SVR provide a stronger adaptability.

6.4 Data Fusion

In total, there are two fusion strategies for multimodal information: feature-level fusion and decision-level fusion. Feature-level fusion means that the multimodal features are combined and then used as the input of a classifier or a regressor. Decision-level fusion fuses several results of different classifiers, and the final results are computed according to the fusion methods.

One way of feature-level fusion is implemented by feature accumulation or concatenation [14, 147, 177]. In Reference [14], two feature vectors for visual and audio data are averaged as the global genre representations. In Reference [177], multi-class features are concatenated to generate a high-dimensional joint representation. Some machine learning methods are also employed to learn joint features [41, 43, 90, 135]. In Reference [41], a two-branch network is used to combine the visual and audio features. The outputs of the two-branch network are then fed into a classifier, and the experiment results showed that the joint features outperform other methods. In Reference [43], the low-level audio-visual features and fMRI-derived features are fed into multimodal DBM to learn joint representations. The target of this method is to learn the relation between audio-visual features and fMRI-derived features. In Reference [135], PCA is used to learn the multimodal joint

features. In Reference [132], canonical correlation analysis (CCA) is used to construct a new video feature space with the help of EEG features and a new EEG feature space with the assistance of video content, so only one modality is needed to predict emotion during the testing process.

By combining the results of different classifiers, decision-level fusion strategy is able to achieve better results [2, 40, 42, 82, 83, 111]. In Reference [2], linear fusion and SVM-based fusion techniques are explored to combine outputs of several classifiers. Specifically, the output of each classifier has its own weight in linear fusion. The final result is the weighted sum of all outputs. In SVM-based fusion, the outputs of unimodal classifiers are concatenated together. And then the higher-level representations for each video clip are fed into a fusion SVM to predict the emotion. Based on these results, linear fusion is better than SVM-based fusion. In References [42, 83], linear fusion is also used to fuse the outputs of multiple classifiers. The differences among these linear fusion methods depend on the distribution of weights. Representative methods on AC of videos are summarized in Table 10.

6.5 Deep Learning Methods

Traditionally, video emotional recognition includes two steps, i.e., feature extraction step and regression or classification step. Because of the lack of consensus on the most relevant emotional features, we may not be able to extract the best features for the problem at hand. As a result, this two-step mode has hampered the development of affective video content analysis. To solve this problem, some methods based on end-to-end training frameworks are proposed. Khorrami et al. [59] combined CNN and RNN to recognize the emotional information of videos. According to their method, a CNN is trained using frame facial images sampled from videos to extract features. Then the features are fed into an RNN to perform continuous emotion recognition. In Reference [51], a single network using ConvLSTM is proposed, where videos are input to the network and the predicted emotional information is output directly. In fact, due to the complexity of CNNs and RNNs, the training of these frameworks needs large amounts of data. However, in video affective content analysis, the samples in existing datasets are usually limited. This is the reason why end-to-end methods are still less common compared to the traditional two-step methods, despite their influential potential.

7 AFFECTIVE COMPUTING OF MULTIMODAL DATA

In this section, we survey the works that analyze multimodal data beyond audiovisual content. Most of the existing works on affective understanding of multimedia rely on one modality, even when additional modalities are available, for example in videos [55]. Earlier work on emotional understanding of multimedia used handcrafted features from different modalities that are fused at feature or decision levels [8, 15, 46, 117, 126]. The more recent works mainly use deep learning models [55, 93].

Language is a commonly used modality in addition to vision and audio. There is a large body of work on text-based sentiment analysis [91]. Sentiment analysis from text is well-established and is deployed at scale in industry on a broad set of applications involving opinion mining [116]. With the shift toward an increasingly multimodal social web, multimodal sentiment analysis is becoming more relevant. For example, vloggers post their opinions on YouTube, and photos commonly accompany user posts on Instagram and Twitter. Analyzing text for emotion recognition requires representing terms by features. Lexically based approaches are one of the most popular methods for text-based emotion recognition. They involve using knowledge of words' affect for estimating document or content's affect. Linguistic Inquiry and Word Count (LIWC) is a well-known lexical tool that matches the terms in a document with its dictionary and generates scores along different dimensions including affective and cognitive constructs such as "present focus" and "positive

Table 10. Representative Work on AC of Videos Using Kinds of Features

Ref	Feature	Fusion	Learning	Dataset	Task	Result
[111]	Mel frequency spectral; MFCC, Chroma, and their derivatives and statistics; Audio compressibility; Harmonicity; Shot frequency; HOF and statistics; Histogram of 3D HSV and statistics; Video compressibility; Histogram of facial area	decision	LSTM	Dataset described by Malandrakis [72]	reg	$P_a : 0.84 \pm 0.06$ $MSE_a : 0.08 \pm 0.04$ $P_v : 0.50 \pm 0.14$ $MSE_v : 0.21 \pm 0.06$
[40]	Average, standard deviation, and four spectral power bands of pupil dilation ratio time-series	decision	SVM, LSTM	MAHNOB-HCI	cla	$Acc_a : 0.730, Acc_v : 0.780$
[14]	Audio and visual deep features from pretrained model	feature	CNN, LSTM, SVM	PMIT	cla	$MAP : 0.02122$
[2]	MFCC; Color values; HoG; Dense trajectory descriptor; CNN-learned features	decision	CNN, SVM, Ensemble	DEAP	cla	$Acc : 0.81, 0.49$
[79]	HHTC features	feature	SVR	Discrete LIRIS-ACCEDE	reg	$MSE_a : 0.294, MSE_v : 0.290$
[42]	Statistical measures (such as mean, median, skewness kurtosis) for EEG data, power spectral features, ECG, GSR, Face/Head-pose	decision	SVM/NB	music excerpts [80]	cla	F1 (v : 0.59, 0.58, a : 0.60, 0.57)
[41]	Time-span visual and visual features	feature	CNN Opensmile toolbox	music excerpts [80]	cla	$MSE_a : 0.082$ $MSE_v : 0.071$
[43]	tempo; pitch; zero cross; roll off; MFCCs; Saturation; Color heat; Shot length feature; General preferences; Visual excitement; Motion feature; fMRI feature	feature	DBM SVM	TRECVID	cla	–
[177]	Colorfulness; MFCC; CNN-learned features from the keyframes containing protagonist	decision	CNN, SVM, SVR	LIRIS-ACCEDE, PMSZU	cla/reg	–
[173]	Multi-frame motion vectors	decision	CNN	SumMe, TVsum, Continuous LIRIS-ACCEDE	reg	–
[49]	The median of the L values in Luv space; means and variances of components in HSV space; texture feature; mean and standard deviation of motions between frames in a shot; MFCC; Spectral power; mean and variance of the spectral centroids; Time domain zero crossings rate; Multi-instance sparse coding	feature	SVM	Musk1, Musk2, Elephant, Fox, Tiger	cla	$Acc : 0.911, 0.906, 0.885, 0.627, 0.868$
[82]	Lighting key; Color; Motion vectors; ZCR; energy; MFCC; pitch; Textual features	decision	SMO, Naive Bayes	DEAP	cla	F1 : 0.849 0.811 $Acc : 0.911 0.883$
[20]	Key lighting; Grayness; Fast motion; Shot change rate; Shot length variation; MFCC; CNN-learned features; power spectral density; EEG; ECG; respiration; galvanic skin resistance	feature	SVM	DEAP	cla	$Acc_v : 0.7 0.7 0.7125, Acc_a : 0.6876 0.7 0.8$ F1 (A : 0.664 0.687 0.789)
[83]	MFCC; ZCR; energy; pitch, color histograms; lighting key; motion vector	decision	SVM, Naive Bayes	DEAP	cla	F1 : 0.869, 0.846 $Acc : 0.925, 0.897$
[109]	CNN feature, low-level audio visual features, EEG	decision	LDA, LSVM, RSVM	Dataset introduced by the authors	cla	–
[147]	MKT; ConvNets feature; EmoLarge; IS13; MFCC; EmoBase10; DSIFT; HSH	decision	SVM, LR, RBM, MLP	MediaEval 2015, 2016 Affective Impact of Movies	cla	$Acc_a : 0.574, Acc_v : 0.462$
[110]	CNN feature	–	SVM, LDA	dataset in Reference [108]	cla	–
[12]	CNN feature	–	SVR	LIRIS-ACCEDE	reg	$MSE_a : 0.021, MSE_v : 0.027$

Where P_a , P_v , MSE_a , MSE_v , Acc_a , Acc_v , Acc , and MAP Indicate the Pearson Correlation Coefficients of Arousal and Valence, the Mean Sum Error of Arousal and Valence, the Accuracy of Arousal and Valence, the Average Accuracy, and Mean Average Precision, Respectively. “Statistics” means (min, max, mean).

emotion” [54]. The terms in each category or selected by experts are extensively validated on different content. AffectNet is another notable lexical resource that includes a semantic network of 10K items with representations for “pleasantness,” “attention,” “sensitivity,” and “aptitude” [19]. The continuous representations can be mapped to 24 distinct emotions. DepecheMood is a lexicon created through a data-driven method mining a news website annotated with its particular set of discrete emotions, namely, “afraid,” “amused,” “anger,” “annoyed,” “don’t care,” “happy,” and “inspired” [122]. DepecheMood is extended to DepecheMood++ by including Italian [7].

The more recent development in text-based affective analysis are models powered by deep learning. Leveraging large-scale data, deep neural networks are able to learn representations that are relevant for affective analysis in language. Word embeddings are one of the most common representations used to represent language. Word embeddings, such as Word2Vec [78] or GloVe [96], learn language context of the word by learning a representation (a vector) that can capture semantic and syntactic similarities. More recently, representation learning models that can encode the whole sequence of terms (sentences, documents) showed impressive performance in different language understanding tasks, including sentiment and emotional analysis. Bidirectional Encoder Representations from Transformers (BERT) [31] is a method for learning a language model that can be trained on large amounts of data in an unsupervised manner. This pre-trained model is very effective in representing a sequence of terms as a fixed-length representation (vector). BERT architecture is a multi-layer bidirectional Transformer network that encodes the whole sequence at once. BERT representation achieves state-of-the-art results in multiple natural language understanding tasks.

The audiovisual features that are used for multimodal understanding of affect are similar to the ones discussed in previous sections. The main technique between multimodal models lies in methods for multimodal fusion. Multimodal methods involve extracting features from multiple modalities, e.g., audiovisual, and training joint or separate machine learning models for fusion [9]. Multimodal fusion can be done in model-based and model-agnostic ways. The model-agnostic fusion methods do not rely on a specific classification or regression method and include feature-level, decision-level, or hybrid fusion techniques. Model-based methods address multimodal fusion in model construction. Examples of model-based fusion methods include Multiple Kernel Learning (MKL) [38], graphical models, such as Conditional Random Fields [10], and neural networks [84, 99].

Pang et al. [93] used Deep Boltzmann Machine (DBM) to learn a joint representation across text, vision, and audio to recognize expected emotions from social media videos. Each modality is separately encoded with stacking multiple Restricted Boltzmann Machines (RBM) and pathways are merged to a joint representation layer. The model was evaluated for recognizing eight emotion categories for 1,101 videos from Reference [55]. Muszynski et al. [81] studied perceived vs induced emotion in movies. To this end, they collected additional labels on a subset of LIRIS-ACCEDE dataset [13]. They found that perceived and induced emotions do not always agree. Using multimodal Deep Belief Networks (DBN), they could demonstrate that fusion of electrodermal responses with audiovisual content features improves the overall accuracy for emotion recognition [81].

In Reference [111], authors performed regression to estimate intended arousal and valence levels (as judged by experts in Reference [72]). LSTM recurrent neural networks are used for unimodal regressions and fused via early and late fusion for audiovisual estimation, with late fusion achieving the best results. Tarvainen et al. [125] performed an in-depth analysis on how emotions are constructed in movies. They identified scene type as a major factor in emotions in movies. They then used content features to recognize emotions along three dimensions of hedonic tone (valence), energetic arousal (awake–tired), and tense arousal (tense–calm).

Bilinear fusion is a method that is proposed to model inter- and intra- modality interaction among modalities by performing outer product between unimodal embeddings [69]. Zadeh

et al. [151] extended this to a Tensor Fusion Network to model intra-modality and inter-modality dynamics in multimodal sentiment analysis. The tensor fusion network includes modality embedding sub-networks, a tensor fusion layer modeling the unimodal, bimodal, and trimodal interactions using a three-fold Cartesian product from modality embeddings along with a final sentiment inference sub-network conditioned on the tensor fusion layer. The main drawback of such methods is the increase in the dimensionality of the resulting multimodal representation.

8 FUTURE DIRECTIONS

Although remarkable progress has been made on affective computing of multimedia (ACM) data, there are still several open issues and directions that can boost the performance of ACM.

Multimedia Content Understanding. As emotions may be directly evoked by the multimedia content in viewers, accurately understanding what is contained in multimedia data can significantly improve the performance of ACM. Sometimes it is even necessary to analyze the subtle details. For example, we may feel “amused” by a video with a laughing baby; but if the laugh is from a negative character, then it is more possible for us to feel “angry.” In such cases, besides the common property, such as “laugh,” we may need to further recognize the identity, such as “a lovely baby” and “an evil antagonist.”

Multimedia Summarization. Emotions can play a vital role in selection of multimedia for creation of summaries or highlights. This is an important application in entertainment and sports industries (e.g., movie trailers, sports highlights). There has been some recent work in this direction where affect information from audio visual cues has led to the successful creation of video summaries [76, 113]. In particular, work reported in Reference [113] used audiovisual emotions in part to create an AI trailer for a 20th Century Fox film in 2016. Similarly, AI Highlights, described in Reference [76], hinged on audiovisual emotional cues and have successfully been employed to create the official highlights at Wimbledon and US Open since 2017. This is a very promising direction for affective multimedia computing, which can have a direct impact on real-world media applications.

Contextual Knowledge Modeling. The contextual information of a viewer watching some multimedia is very important. Similar multimedia data under different contexts may evoke totally different emotions. For example, we may feel “happy” when listening to a song about love in a wedding; but if the same song is played when two lovers are departing, then it is more likely that we feel “sad.” The prior knowledge of viewers or multimedia data may also influence the emotion perceptions. An optimistic viewer and a pessimistic viewer may have totally different emotions about the same multimedia data.

Group Emotion Clustering. It is too generic to simply recognize the dominant emotion, while it is too specific to predict personalized emotion. It would make more sense to model emotions for groups or cliques of viewers with similar interests and backgrounds. Clustering different viewers into corresponding groups possibly based on the user profiles may provide a feasible solution to this problem.

New AC Setting Adaptation. Because of the domain shift [127], the deep learning models trained on one labeled source domain may not work well on the other unlabeled or sparsely labeled target domain, which results in the models’ low transferability to new domains. Exploring domain adaptation techniques that fit well on the AC tasks is worth investigating. One possible solution is to translate the source data to an intermediate domain that is indistinguishable from the target data while preserving the source labels [165, 172] using Generative Adversarial Networks [39, 174]. How to deal with some practical settings, such as multiple labeled source domains and emotion models’ homogeneity, is more challenging.

Regions-of-interest Selection. The contributions of different regions of given multimedia may vary to the emotion recognition. For example, the regions that contain the most important semantic information in images are more discriminative than background; some video frames are of no use to emotion recognition. Detecting and selecting the regions of interest may significantly improve the recognition performance as well as the computation efficiency.

Viewer-multimedia Interaction. Instead of direct analysis of the multimedia content or implicit consideration of viewers' physiological signals (such as facial expressions, Electroencephalogram signals, etc.), joint modeling of both multimedia content and viewers' responses may better bridge the affective gap and result in superior performances. We should also study how to deal with missing or corrupted data. For example, some physiological signals are unavailable during the data collection stage.

Affective Computing Applications. Although AC is claimed to be important in real-world applications, few practical systems have been developed due to the relatively low performance. With the availability of larger datasets and improvements in self-supervised and semi-supervised learning, we foresee the deployment of ACM in real-world applications. For example, in media analytics, the content-understanding methods will identify the emotional preferences of users and emotional nuances of social media content to better target advertising efforts; in fashion recommendation, intelligent costumer service, such as customer-multimedia interaction, can provide better experience to customers; in advertising, generating or curating multimedia that evokes strong emotions can attract more attention. We believe that an emotional artificial intelligence will become a significant component of mainstream multimedia applications.

Benchmark Dataset Construction. Existing studies on ACM mainly adopt small-scale datasets or construct relatively larger-scale ones using keyword searching strategy without annotation quality guaranteed. To advance the development of ACM, creating a large-scale and high-quality dataset is in urgent need. It has shown that there are three critical factors for dataset construction of ACM, i.e., the context of viewer response, personal variation among viewers, and the effectiveness and efficiency of corpus creation [118]. To include a large number of samples, we may exploit online systems and crowdsourcing platforms to recruit large numbers of viewers with a representative spread of backgrounds to annotate multimedia and provide contextual information on their emotional responses. Since emotion is a subjective variable, personalized emotion annotation would make more sense, from which we can obtain the dominant emotion and emotion distribution. Further, accurate understanding of multimedia content can boost the affective computing performance. Inferring emotional labels from social media users' interaction with data, e.g., likes, comments, in addition to their spontaneous responses, e.g., facial expression, where possible, will provide new avenues for enriching affective datasets.

9 CONCLUSION

In this article, we have surveyed affective computing (AC) methods for heterogeneous multimedia data. For each multimedia type, i.e., images, music, videos, and multimodal data, we summarized and compared available datasets, handcrafted features, machine learning methods, deep learning models, and experimental results. We also briefly introduced the commonly employed emotion models and outlined potential research directions in this area. Although deep learning-based AC methods have achieved remarkable progress in recent years, an efficient and robust AC method that is able to obtain high accuracy under unconstrained conditions is yet to be achieved. With the advent of deep understanding of emotion evocation in brain science, accurate emotion measurement in psychology, and novel deep learning network architectures in machine learning, affective computing of multimedia data will remain an active research topic for a long time.

REFERENCES

- [1] Mojtaba Khomami Abadi, Ramanathan Subramanian, Seyed Mostafa Kia, Paolo Avesani, Ioannis Patras, and Nicu Sebe. 2015. DECAF: MEG-based multimodal database for decoding affective physiological responses. *IEEE Trans. Affect. Comput.* 6, 3 (2015), 209–222.
- [2] Esra Acar, Frank Hopfgartner, and Sahin Albayrak. 2017. A comprehensive study on mid-level representation and ensemble learning for emotional analysis of video material. *Multim. Tools Appl.* 76, 9 (2017), 11809–11837.
- [3] Xavier Alameda-Pineda, Elisa Ricci, Yan Yan, and Nicu Sebe. 2016. Recognizing emotions from abstract paintings using non-linear matrix completion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5240–5248.
- [4] Anna Aljanaki. 2016. *Emotion in Music: Representation and Computational Modeling*. Ph.D. Dissertation. Utrecht University.
- [5] Anna Aljanaki and Mohammad Soleymani. 2018. A data-driven approach to mid-level perceptual musical feature modeling. In *Proceedings of the International Society for Music Information Retrieval Conference*.
- [6] Anna Aljanaki, Yi-Hsuan Yang, and Mohammad Soleymani. 2017. Developing a benchmark for emotional analysis of music. *PloS One* 12, 3 (2017), e0173392.
- [7] Oscar Araque, Lorenzo Gatti, Jacopo Staiano, and Marco Guerini. 2018. DepecheMood++: A bilingual emotion lexicon built through simple yet powerful techniques. Retrieved from: *Arxiv Preprint Arxiv:1810.03660* (2018).
- [8] Sutjipto Arifin and Peter Y. K. Cheung. 2008. Affective level video segmentation by utilizing the pleasure-arousal-dominance information. *IEEE Trans. Multim.* 10, 7 (2008), 1325–1341.
- [9] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2019. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 2 (2019), 423–443.
- [10] Tadas Baltrušaitis, Ntombikayise Banda, and Peter Robinson. 2013. Dimensional affect recognition using continuous conditional random fields. In *Proceedings of the IEEE International Conference and Workshops on Automatic Face and Gesture Recognition*. 1–8.
- [11] Yoann Baveye, Christel Chamaret, Emmanuel Dellandréa, and Liming Chen. 2018. Affective video content analysis: A multidisciplinary insight. *IEEE Trans. Affect. Comput.* 9, 4 (2018), 396–409.
- [12] Yoann Baveye, Emmanuel Dellandréa, Christel Chamaret, and Liming Chen. 2015. Deep learning vs. kernel methods: Performance for emotion prediction in videos. In *Proceedings of the International Conference on Affective Computing and Intelligent Interaction*. 77–83.
- [13] Yoann Baveye, Emmanuel Dellandrea, Christel Chamaret, and Liming Chen. 2015. Liris-accede: A video database for affective content analysis. *IEEE Trans. Affect. Comput.* 6, 1 (2015), 43–55.
- [14] Olfa Ben-Ahmed and Benoit Huet. 2018. Deep multimodal features for movie genre and interestingness prediction. In *Proceedings of the International Conference on Content-Based Multimedia Indexing*. 1–6.
- [15] Sergio Benini, Luca Canini, and Riccardo Leonardi. 2011. A connotative space for supporting movie affective recommendation. *IEEE Trans. Multim.* 13, 6 (2011), 1356–1370.
- [16] Dmitry Bogdanov, Nicolas Wack, Emilia Gómez, Sankalp Gulati, Perfecto Herrera, O. Mayor, Gerard Roma, Justin Salamon, J. R. Zapata, and Xavier Serra. 2013. ESSENTIA: An audio analysis library for music information retrieval. In *Proceedings of the International Society for Music Information Retrieval Conference*. 493–498.
- [17] Damian Borth, Tao Chen, Rongrong Ji, and Shih-Fu Chang. 2013. Sentibank: Large-scale ontology and classifiers for detecting sentiment and emotions in visual content. In *Proceedings of the ACM International Conference on Multimedia*. 459–460.
- [18] Damian Borth, Rongrong Ji, Tao Chen, Thomas Breuel, and Shih-Fu Chang. 2013. Large-scale visual sentiment ontology and detectors using adjective noun pairs. In *Proceedings of the ACM International Conference on Multimedia*. 223–232.
- [19] Erik Cambria, Thomas Mazzocco, Amir Hussain, and Chris Eckl. 2011. Sentic medoids: Organizing affective common sense knowledge in a multi-dimensional vector space. In *Advances in Neural Networks*, Vol. 6677. Springer, 601–610.
- [20] Mo Chen, Gong Cheng, and Lei Guo. 2018. Identifying affective levels on music video via completing the missing modality. *Multim. Tools Appl.* 77, 3 (2018), 3287–3302.
- [21] Tao Chen, Damian Borth, Trevor Darrell, and Shih-Fu Chang. 2014. DeepSentiBank: Visual sentiment concept classification with deep convolutional neural networks. *Comput. Sci.* (2014).
- [22] Tao Chen, Felix X. Yu, Jiawei Chen, Yin Cui, Yan-Ying Chen, and Shih-Fu Chang. 2014. Object-based visual sentiment concept analysis and application. In *Proceedings of the ACM International Conference on Multimedia*. 367–376.
- [23] Yu-An Chen, Yi-Hsuan Yang, Ju-Chiang Wang, and Homer Chen. 2015. The AMG1608 dataset for music emotion recognition. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*. 693–697.
- [24] Juan Abdon Miranda Correa, Mojtaba Khomami Abadi, Nicu Sebe, and Ioannis Patras. 2017. AMIGOS: A dataset for mood, personality, and affect research on individuals and GrOuPS. Retrieved from: *Arxiv Preprint Arxiv:1702.02510* (2017).

- [25] Vaidehi Dalmia, Hongyi Liu, and Shih-Fu Chang. 2016. Columbia MVSO image sentiment dataset. Retrieved from: *Arxiv Preprint Arxiv:1611.04455* (2016).
- [26] Elise S. Dan-Glauser and Klaus R. Scherer. 2011. The Geneva affective picture database (GAPED): A new 730-picture database focusing on valence and normative significance. *Behav. Res. Meth.* 43, 2 (2011), 468–477.
- [27] Emmanuel Dellandréa, Liming Chen, Yoann Baveye, Mats Viktor Sjöberg, Christel Chamaret et al. 2016. The MediaEval 2016 emotional impact of movies task. In *Proceedings of the CEUR Workshop*.
- [28] Emmanuel Dellandréa, Martijn Huigsloot, Liming Chen, Yoann Baveye, and Mats Sjöberg. 2017. The MediaEval 2017 emotional impact of movies task. In *Proceedings of the MediaEval Workshop*.
- [29] Emmanuel Dellandréa, Martijn Huigsloot, Liming Chen, Yoann Baveye, Zhongzhe Xiao, and Mats Sjöberg. 2018. The MediaEval 2018 emotional impact of movies task. In *Proceedings of the MediaEval Workshop*.
- [30] Emmanuel Dellandréa, Martijn Huigsloot, Liming Chen, Yoann Baveye, Zhongzhe Xiao, and Mats Sjöberg. 2019. Datasets column: Predicting the emotional impact of movies. *ACM SIGMultim. Rec.* 10, 4 (2019), 6.
- [31] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*. 4171–4186.
- [32] X. H. J. S. Downie, Cyril Laurier, and M. B. A. F. Ehmann. 2008. The 2007 MIREX audio mood classification task: Lessons learned. In *Proceedings of the International Society for Music Information Retrieval Conference*. 462–467.
- [33] Tuomas Eerola and Jonna K. Vuoskoski. 2011. A comparison of the discrete and dimensional models of emotion in music. *Psych. Music* 39, 1 (2011), 18–49.
- [34] Paul Ekman. 1992. An argument for basic emotions. *Cog. Emot.* 6, 3–4 (1992), 169–200.
- [35] Moataz El Ayadi, Mohamed S. Kamel, and Fakhri Karray. 2011. Survey on speech emotion recognition: Features, classification schemes, and databases. *Pattern Recog.* 44, 3 (2011), 572–587.
- [36] Florian Eyben, Felix Weninger, Florian Gross, and Björn Schuller. 2013. Recent developments in openSMILE, the Munich open-source multimedia feature extractor. In *Proceedings of the ACM International Conference on Multimedia*. 835–838.
- [37] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. OpenSMILE: The Munich versatile and fast open-source audio feature extractor. In *Proceedings of the ACM International Conference on Multimedia*. 1459–1462.
- [38] Mehmet Gönen and Ethem Alpaydin. 2011. Multiple kernel learning algorithms. *J. Mach. Learn. Res.* 12, July (2011), 2211–2268.
- [39] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Proceedings of the Conference on Advances in Neural Information Processing Systems*. 2672–2680.
- [40] Dongdong Gui, Sheng-hua Zhong, and Zhong Ming. 2018. Implicit affective video tagging using pupillary response. In *Proceedings of the International Conference on Multimedia Modeling*. 165–176.
- [41] Jie Guo, Bin Song, Peng Zhang, Mengdi Ma, Wenwen Luo et al. 2019. Affective video content analysis based on multimodal data fusion in heterogeneous networks. *Inform. Fus.* 51 (2019), 224–232.
- [42] Rishabh Gupta, Mojtaba Khomami Abadi, Jesús Alejandro Cárdenes Cabré, Fabio Morreale, Tiago H. Falk, and Nicu Sebe. 2016. A quality adaptive multimodal affect recognition system for user-centric multimedia indexing. In *Proceedings of the ACM International Conference on Multimedia Retrieval*. 317–320.
- [43] Junwei Han, Xiang Ji, Xintao Hu, Lei Guo, and Tianming Liu. 2015. Arousal recognition using audio-visual features and fMRI-based brain response. *IEEE Trans. Affect. Comput.* 6, 4 (2015), 337–347.
- [44] Junwei Han, Dingwen Zhang, Gong Cheng, Nian Liu, and Dong Xu. 2018. Advanced deep-learning techniques for salient and category-specific object detection: A survey. *IEEE Sig. Proc. Mag.* 35, 1 (2018), 84–100.
- [45] Alan Hanjalic. 2006. Extracting moods from pictures and sounds: Towards truly personalized TV. *IEEE Sig. Proc. Mag.* 23, 2 (2006), 90–100.
- [46] Alan Hanjalic and Li-Qun Xu. 2005. Affective video content representation and modeling. *IEEE Trans. Multim.* 7, 1 (2005), 143–154.
- [47] John H. L. Hansen and Taufiq Hasan. 2015. Speaker recognition by machines and humans: A tutorial review. *IEEE Sig. Proc. Mag.* 32, 6 (2015), 74–99.
- [48] Samitha Herath, Mehrtaash Harandi, and Fatih Porikli. 2017. Going deeper into action recognition: A survey. *Image Vis. Comput.* 60 (2017), 4–21.
- [49] Weiming Hu, Xinmiao Ding, Bing Li, Jianchao Wang, Yan Gao, Fangshi Wang, and Stephen Maybank. 2016. Multi-perspective cost-sensitive context-aware multi-instance sparse coding and its application to sensitive video recognition. *IEEE Trans. Multim.* 18, 1 (2016), 76–89.
- [50] Xiao Hu and J. Stephen Downie. 2007. Exploring mood metadata: Relationships with genre, artist, and usage metadata. In *Proceedings of the International Society for Music Information Retrieval Conference*. 67–72.

- [51] Jian Huang, Ya Li, Jianhua Tao, Zheng Lian, and Jiangyan Yi. 2018. End-to-end continuous emotion recognition from video using 3D ConvLSTM networks. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*. 6837–6841.
- [52] Charles Inskip, Andy Macfarlane, and Pauline Rafferty. 2012. Towards the disintermediation of creative music search: Analysing queries to determine important facets. *Int. J. Dig. Lib.* 12, 2–3 (2012), 137–147.
- [53] Bernard J. Jansen, Mimi Zhang, Kate Sobel, and Abdur Chowdury. 2009. Twitter power: Tweets as electronic word of mouth. *J. Assoc. Inform. Sci. Technol.* 60, 11 (2009), 2169–2188.
- [54] Audra E. Massey, Jeffrey H. Kahn, Renée M. Tobin, and Jennifer A. Anderson. 2007. Measuring emotional expression with the linguistic inquiry and word count. *JSTOR: Amer. J. Psychol.* 120, 2 (2007), 263–286.
- [55] Yu-Gang Jiang, Baohan Xu, and Xiangyang Xue. 2014. Predicting emotions in user-generated videos. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*.
- [56] Dhiraj Joshi, Ritendra Datta, Elena Fedorovskaya, Quang-Tuan Luong, James Z. Wang, Jia Li, and Jiebo Luo. 2011. Aesthetics and emotions in images. *IEEE Sig. Proc. Mag.* 28, 5 (2011), 94–115.
- [57] Brendan Jou, Tao Chen, Nikolaos Pappas, Miriam Redi, Mercan Topkara, and Shih-Fu Chang. 2015. Visual affect around the world: A large-scale multilingual visual sentiment ontology. In *Proceedings of the ACM International Conference on Multimedia*. 159–168.
- [58] Brendan Jou, Margaret Yuying Qian, and Shih-Fu Chang. 2016. SentiCart: Cartography and geo-contextualization for multilingual visual sentiment. In *Proceedings of the ACM International Conference on Multimedia Retrieval*. 389–392.
- [59] Pooya Khorrami, Tom Le Paine, Kevin Brady, Charlie Dagli, and Thomas S. Huang. 2016. How deep neural networks can improve emotion recognition on video data. In *Proceedings of the IEEE International Conference on Image Processing*. 619–623.
- [60] Youngmoo E. Kim, Erik M. Schmidt, Raymond Migneco, Brandon G. Morton, Patrick Richardson, Jeffrey Scott, Jacquelin A. Speck, and Douglas Turnbull. 2010. Music emotion recognition: A state-of-the-art review. In *Proceedings of the International Society for Music Information Retrieval Conference*, Vol. 86. 937–952.
- [61] Sander Koelstra, Christian Muhl, Mohammad Soleymani, Jong-Seok Lee, Ashkan Yazdani, Touradj Ebrahimi, Thierry Pun, Anton Nijholt, and Ioannis Patras. 2012. DEAP: A database for emotion analysis using physiological signals. *IEEE Trans. Affect. Comput.* 3, 1 (2012), 18–31.
- [62] Peter J. Lang, Margaret M. Bradley, and Bruce N. Cuthbert. 1997. International affective picture system (IAPS): Technical manual and affective ratings. *NIMH Center Study Emot. Attent.* (1997), 39–58.
- [63] Martha Larson, Mohammad Soleymani, Guillaume Gravier, Bogdan Ionescu, and Gareth J. F. Jones. 2017. The benchmarking initiative for multimedia evaluation: MediaEval 2016. *IEEE MultiM.* 24, 1 (2017), 93–96.
- [64] Olivier Lartillot, Petri Toivainen, and Tuomas Eerola. 2008. A Matlab toolbox for music information retrieval. In *Data Analysis, Machine Learning and Applications*. Springer, 261–268.
- [65] Cyril Laurier, Perfecto Herrera, M. Mandel, and D. Ellis. 2007. Audio music mood classification using support vector machine. In *MIREX task on Audio Mood Classification*. IEEE, 2–4.
- [66] Joonwhoan Lee and EunJong Park. 2011. Fuzzy similarity-based emotional classification of color images. *IEEE Trans. Multim.* 13, 5 (2011), 1031–1039.
- [67] Bing Li, Weihua Xiong, Weiming Hu, and Xinmiao Ding. 2012. Context-aware affective images classification based on bilayer sparse representation. In *Proceedings of the ACM International Conference on Multimedia*. 721–724.
- [68] Benjamin J. Li, Jeremy N. Bailenson, Adam Pines, Walter J. Greenleaf, and Leanne M. Williams. 2017. A public database of immersive VR videos with corresponding ratings of arousal, valence, and correlations between head movements and self report measures. *Front. Psych.* 8 (2017), 2116.
- [69] Tsung-Yu Lin, Aruni Roy Chowdhury, and Subhransu Maji. 2015. Bilinear CNN models for fine-grained visual recognition. In *Proceedings of the IEEE International Conference on Computer Vision*. 1449–1457.
- [70] Xin Lu, Poonam Suryanarayan, Reginald B. Adams Jr., Jia Li, Michelle G. Newman, and James Z. Wang. 2012. On shape and the computability of emotions. In *Proceedings of the ACM International Conference on Multimedia*. 229–238.
- [71] Jana Machajdik and Allan Hanbury. 2010. Affective image classification using features inspired by psychology and art theory. In *Proceedings of the ACM International Conference on Multimedia*. 83–92.
- [72] Nikos Malandrakis, Alexandros Potamianos, Georgios Evangelopoulos, and Athanasia Zlatintsi. 2011. A supervised approach to movie emotion tracking. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*. 2376–2379.
- [73] Ricardo Malheiro, Renato Panda, Paulo Gomes, and Rui Pedro Paiva. 2016. Emotionally relevant features for classification and regression of music lyrics. *IEEE Trans. Affect. Comput.* 9, 2 (2016), 240–254.
- [74] Daniel McDuff and Mohammad Soleymani. 2017. Large-scale affective content analysis: Combining media content features and facial reactions. In *Proceedings of the IEEE International Conference on Automatic Face & Gesture Recognition*. 339–345.

- [75] Brian McFee, Colin Raffel, Dawen Liang, Daniel P. W. Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. 2015. librosa: Audio and music signal analysis in Python. In *Proceedings of the Python in Science Conferences*, Vol. 8. 18–25.
- [76] Michele Merler, Khoi-Nguyen C. Mac, Dhiraj Joshi, Quoc-Bao Nguyen, Stephen Hammer, John Kent, Jinjun Xiong, Minh N. Do, John R. Smith, and Rogerio S. Feris. 2018. Automatic curation of sports highlights using multimodal excitement features. *IEEE Trans. Multim.* 21, 5 (2018), 1147–1160.
- [77] Joseph A. Mikels, Barbara L. Fredrickson, Gregory R. Larkin, Casey M. Lindberg, Sam J. Maglio, and Patricia A. Reuter-Lorenz. 2005. Emotional category data on images from the International Affective Picture System. *Behav. Res. Meth.* 37, 4 (2005), 626–630.
- [78] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of the Conference on Advances in Neural Information Processing Systems*. 3111–3119.
- [79] Shasha Mo, Jianwei Niu, Yiming Su, and Sajal K. Das. 2018. A novel feature set for video emotion recognition. *Neurocomputing* 291 (2018), 11–20.
- [80] Fabio Morreale, Raul Masu, Antonella De Angeli et al. 2013. Robin: An algorithmic composer for interactive scenarios. In *Proceedings of the Sound and Music Computing Conference*. 207–212.
- [81] Michal Muszynski, Leimin Tian, Catherine Lai, Johanna Moore, Theodoros Kostoulas, Patrizia Lombardo, Thierry Pun, and Guillaume Chanel. 2019. Recognizing induced emotions of movie audiences from multimodal information. *IEEE Trans. Affect. Comput.* (2019).
- [82] Shahla Nemati and Ahmad Reza Naghsh-Nilchi. 2016. Incorporating social media comments in affective video retrieval. *J. Inform. Sci.* 42, 4 (2016), 524–538.
- [83] Shahla Nemati and Ahmad Reza Naghsh-Nilchi. 2017. An evidential data fusion method for affective music video retrieval. *Intell. Data Anal.* 21, 2 (2017), 427–441.
- [84] Mihalís A. Nicolaou, Hatice Gunes, and Maja Pantic. 2011. Continuous prediction of spontaneous affect from multiple cues and modalities in valence-arousal space. *IEEE Trans. Affect. Comput.* 2, 2 (2011), 92–105.
- [85] Jianwei Niu, Yiming Su, Shasha Mo, and Zeyu Zhu. 2017. A novel affective visualization system for videos based on acoustic and visual features. In *Proceedings of the International Conference on Multimedia Modeling*. 15–27.
- [86] Jianwei Niu, Shihao Wang, Yiming Su, and Song Guo. 2017. Temporal factor-aware video affective analysis and recommendation for cyber-based social media. *IEEE Tran. Emerg. Topics Comput.* 5 (2017), 412–424.
- [87] Jianwei Niu, Xiaoke Zhao, and Muhammad Ali Abdul Aziz. 2016. A novel affect-based model of similarity measure of videos. *Neurocomputing* 173 (2016), 339–345.
- [88] Andrew Ortony, Gerald L. Clore, and Allan Collins. 1988. *The Cognitive Structure of Emotions*. Cambridge University Press.
- [89] Renato Eduardo Silva Panda. 2019. *Emotion-based Analysis and Classification of Audio Music*. Ph.D. Dissertation. Univesity of Coimbra, Coimbra, Portugal.
- [90] Yagya Raj Pandeya and Lee Joonwhoan. 2019. Music-video emotion analysis using late fusion of multimodal. *DEStech Trans. Comput. Sci. Eng.* (2019).
- [91] Bo Pang, Lillian Lee et al. 2008. Opinion mining and sentiment analysis. *Found. Trends® Inform. Retr.* 2, 1–2 (2008), 1–135.
- [92] Lei Pang and Chong-Wah Ngo. 2015. Mutlimodal learning with deep Boltzmann machine for emotion prediction in user generated videos. In *Proceedings of the ACM International Conference on Multimedia Retrieval*. 619–622.
- [93] Lei Pang, Shiai Zhu, and Chong Wah Ngo. 2015. Deep multimodal learning for affective analysis and retrieval. *IEEE Trans. Multim.* 17, 11 (2015), 2008–2020.
- [94] Genevieve Patterson and James Hays. 2012. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2751–2758.
- [95] Kuan-Chuan Peng, Amir Sadovnik, Andrew Gallagher, and Tsuhan Chen. 2015. A mixed bag of emotions: Model, predict, and transfer emotion distributions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 860–868.
- [96] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 1532–1543.
- [97] Robert Plutchik. 1980. *Emotion: A Psychoevolutionary Synthesis*. Harpercollins College Division.
- [98] Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. MELD: A multimodal multi-party dataset for emotion recognition in conversations. In *Proceedings of the Meeting of the Association for Computational Linguistics*. 527–536.
- [99] Shyam Sundar Rajagopalan, Louis-Philippe Morency, Tadas Baltrušaitis, and Roland Goecke. 2016. Extending long short-term memory for multi-view structured learning. In *Proceedings of the European Conference on Computer Vision*. 338–353.

- [100] Tianrong Rao, Min Xu, Huiying Liu, Jinqiao Wang, and Ian Burnett. 2016. Multi-scale blocks based image emotion classification using multiple instance learning. In *Proceedings of the IEEE International Conference on Image Processing*. 634–638.
- [101] Tianrong Rao, Min Xu, and Dong Xu. 2016. Learning multi-level deep representations for image emotion classification. Retrieved from: *Arxiv Preprint Arxiv:1611.07145* (2016).
- [102] David Sander, Didier Grandjean, and Klaus R. Scherer. 2005. A systems approach to appraisal mechanisms in emotion. *Neural Netw.* 18, 4 (2005), 317–352.
- [103] Andreza Sartori, Dubravko Culibrk, Yan Yan, and Nicu Sebe. 2015. Who's afraid of itten: Using the art theory of color combination to analyze emotions in abstract paintings. In *Proceedings of the ACM International Conference on Multimedia*. 311–320.
- [104] Harold Schlosberg. 1954. Three dimensions of emotion. *Psych. Rev.* 61, 2 (1954), 81.
- [105] Dongyu She, Jufeng Yang, Ming-Ming Cheng, Yu-Kun Lai, Paul L. Rosin, and Liang Wang. 2019. WSCNet: Weakly supervised coupled networks for visual sentiment classification and detection. *IEEE Trans. Multim.* (2019).
- [106] Guangyao Shen, Jia Jia, Liqiang Nie, Fuli Feng, Cunjun Zhang, Tianrui Hu, Tat-Seng Chua, and Wenwu Zhu. 2017. Depression detection via harvesting social media: A multimodal dictionary learning solution. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 3838–3844.
- [107] Abhinav Shukla. 2018. *Multimodal Emotion Recognition from Advertisements with Application to Computational Advertising*. Ph.D. Dissertation. International Institute of Information Technology Hyderabad.
- [108] Abhinav Shukla, Shruti Shriya Gullapuram, Harish Katti, Karthik Yadati, Mohan Kankanhalli, and Ramanathan Subramanian. 2017. Affect recognition in ads with application to computational advertising. In *Proceedings of the ACM International Conference on Multimedia*. 1148–1156.
- [109] Abhinav Shukla, Shruti Shriya Gullapuram, Harish Katti, Karthik Yadati, Mohan Kankanhalli, and Ramanathan Subramanian. 2017. Evaluating content-centric vs. user-centric ad affect recognition. In *Proceedings of the ACM International Conference on Multimodal Interaction*. 402–410.
- [110] Abhinav Shukla, Harish Katti, Mohan Kankanhalli, and Ramanathan Subramanian. 2018. Looking beyond a clever narrative: Visual context and attention are primary drivers of affect in video advertisements. In *Proceedings of the ACM International Conference on Multimodal Interaction*. 210–219.
- [111] Sarath Sivaprasad, Tanmayee Joshi, Rishabh Agrawal, and Niranjan Pedaneekar. 2018. Multimodal continuous prediction of emotions in movies using long short-term memory networks. In *Proceedings of the ACM International Conference on Multimedia Retrieval*. 413–419.
- [112] Mats Sjöberg, Yoann Baveye, Hanli Wang, Vu Lam Quang, Bogdan Ionescu, Emmanuel Dellandréa, Markus Schedl, Claire-Hélène Demarty, and Liming Chen. 2015. The MediaEval 2015 affective impact of movies task. In *Proceedings of the MediaEval Conference*.
- [113] John R. Smith, Dhiraj Joshi, Benoit Huet, Winston Hsu, and Jozef Cota. 2017. Harnessing A.I. for augmenting creativity: Application to movie trailer creation. In *Proceedings of the ACM International Conference on Multimedia*. 1799–1808.
- [114] Mohammad Soleymani. 2015. The quest for visual interest. In *Proceedings of the ACM International Conference on Multimedia*. 919–922.
- [115] Mohammad Soleymani, Micheal N. Caro, Erik M. Schmidt, Cheng-Ya Sha, and Yi-Hsuan Yang. 2013. 1000 Songs for emotional analysis of music. In *Proceedings of the ACM International Workshop on Crowdsourcing for Multimedia*. 1–6.
- [116] Mohammad Soleymani, David Garcia, Brendan Jou, Björn Schuller, Shih-Fu Chang, and Maja Pantic. 2017. A survey of multimodal sentiment analysis. *Image Vis. Comput.* 65 (2017), 3–14.
- [117] Mohammad Soleymani, Joep J. M. Kierkels, Guillaume Chanel, and Thierry Pun. 2009. A Bayesian framework for video affective representation. In *Proceedings of the International Conference on Affective Computing and Intelligent Interaction and Workshops*. 1–7.
- [118] Mohammad Soleymani, Martha Larson, Thierry Pun, and Alan Hanjalic. 2014. Corpus development for affective video indexing. *IEEE Trans. Multim.* 16, 4 (2014), 1075–1089.
- [119] Mohammad Soleymani, Jeroen Lichtenauer, Thierry Pun, and Maja Pantic. 2012. A multimodal database for affect recognition and implicit tagging. *IEEE Trans. Affect. Comput.* 3, 1 (2012), 42–55.
- [120] Yale Song and Mohammad Soleymani. 2019. Polysemous visual-semantic embedding for cross-modal retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1979–1988.
- [121] Jacquelin A. Speck, Erik M. Schmidt, Brandon G. Morton, and Youngmoo E. Kim. 2011. A comparative study of collaborative vs. traditional musical mood annotation. In *Proceedings of the International Society for Music Information Retrieval Conference*, Vol. 104. 549–554.
- [122] Jacopo Staiano and Marco Guerini. 2014. Depeche mood: A lexicon for emotion analysis from crowd annotated news. In *Proceedings of the Meeting of the Association for Computational Linguistics*. 427–433.

- [123] Ramanathan Subramanian, Julia Wache, Mojtaba Khomami Abadi, Radu L. Vieriu, Stefan Winkler, and Nicu Sebe. 2018. ASCERTAIN: Emotion and personality recognition using commercial sensors. *IEEE Trans. Affect. Comput.* 9, 2 (2018), 147–160.
- [124] Kai Sun, Junqing Yu, Yue Huang, and Xiaoqiang Hu. 2009. An improved valence-arousal emotion space for video affective content representation and recognition. In *Proceedings of the IEEE International Conference on Multimedia and Expo*. 566–569.
- [125] Jussi Tarvainen, Jorma Laaksonen, and Tapio Takala. 2018. Film mood and its quantitative determinants in different types of scenes. *IEEE Trans. Affect. Comput.* (2018).
- [126] René Marcelino Abritta Teixeira, Toshihiko Yamasaki, and Kiyoharu Aizawa. 2012. Determination of emotional content of video clips by low-level audiovisual features. *Multim. Tools Appl.* 61, 1 (2012), 21–49.
- [127] Antonio Torralba and Alexei A. Efros. 2011. Unbiased look at dataset bias. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1521–1528.
- [128] Douglas Turnbull, Luke Barrington, David Torres, and Gert Lanckriet. 2007. Towards musical query-by-semantic-description using the cal500 data set. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*. 439–446.
- [129] George Tzanetakis and Perry Cook. 2000. Marsyas: A framework for audio analysis. *Org. Sound* 4, 3 (2000), 169–175.
- [130] Shangfei Wang, Shiyu Chen, and Qiang Ji. 2019. Content-based video emotion tagging augmented by users’ multiple physiological responses. *IEEE Trans. Affect. Comput.* 10, 2 (2019), 155–166.
- [131] Shangfei Wang and Qiang Ji. 2015. Video affective content analysis: A survey of state-of-the-art methods. *IEEE Trans. Affect. Comput.* 6, 4 (2015), 410–430.
- [132] Shangfei Wang, Yachen Zhu, Lihua Yue, and Qiang Ji. 2015. Emotion recognition with the help of privileged information. *IEEE Trans. Auton. Ment. Dev.* 7, 3 (2015), 189–200.
- [133] Xiaohui Wang, Jia Jia, Jiaming Yin, and Lianhong Cai. 2013. Interpretable aesthetic features for affective image classification. In *Proceedings of the IEEE International Conference on Image Processing*. 3230–3234.
- [134] Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behav. Res. Meth.* 45, 4 (2013), 1191–1207.
- [135] Baixi Xing, Hui Zhang, Kejun Zhang, Lekai Zhang, Xinda Wu, Xiaoying Shi, Shanghai Yu, and Sanyuan Zhang. 2019. Exploiting EEG signals and audiovisual feature fusion for video emotion recognition. *IEEE Access* 7 (2019), 59844–59861.
- [136] Baohan Xu, Yanwei Fu, Yu-Gang Jiang, Boyang Li, and Leonid Sigal. 2016. Heterogeneous knowledge transfer in video emotion recognition, attribution, and summarization. *IEEE Trans. Affect. Comput.* 9, 2 (2016), 255–270.
- [137] Can Xu, Suleyman Cetintas, Kuang-Chih Lee, and Li-Jia Li. 2014. Visual sentiment prediction with deep convolutional neural networks. Retrieved from: *Arxiv Preprint Arxiv:1411.5731* (2014).
- [138] Jufeng Yang, Dongyu She, Yukun Lai, and Ming-Hsuan Yang. 2018. Retrieving and classifying affective Images via deep metric learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [139] Jufeng Yang, Dongyu She, Yu-Kun Lai, Paul L. Rosin, and Ming-Hsuan Yang. 2018. Weakly supervised coupled networks for visual sentiment analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7584–7592.
- [140] Jufeng Yang, Dongyu She, and Ming Sun. 2017. Joint image emotion classification and distribution learning via deep convolutional neural network. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 3266–3272.
- [141] Jufeng Yang, Ming Sun, and Xiaoxiao Sun. 2017. Learning visual sentiment distributions via augmented conditional probability neural network. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 224–230.
- [142] Peng Yang, Qingshan Liu, and Dimitris N. Metaxas. 2010. Exploring facial expressions with compositional features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2638–2644.
- [143] Xinyu Yang, Yizhuo Dong, and Juan Li. 2018. Review of data features-based music emotion recognition methods. *Multim. Syst.* 24, 4 (2018), 365–389.
- [144] Yang Yang, Jia Jia, Shumei Zhang, Boya Wu, Qicong Chen, Juanzi Li, Chunxiao Xing, and Jie Tang. 2014. How do your friends on social media disclose your emotions? In *Proceedings of the AAAI Conference on Artificial Intelligence*. 306–312.
- [145] Yi-Hsuan Yang and Homer H. Chen. 2012. Machine recognition of music emotion: A review. *ACM Trans. Intell. Syst. Technol.* 3, 3 (2012), 40.
- [146] Xingxu Yao, Dongyu She, Sicheng Zhao, Jie Liang, Yu-Kun Lai, and Jufeng Yang. 2019. Attention-aware polarity sensitive embedding for affective image retrieval. In *Proceedings of the IEEE International Conference on Computer Vision*.
- [147] Yun Yi and Hanli Wang. 2018. Multi-modal learning for affective content analysis in movies. *Multim. Tools Appl.* (2018), 1–20.

- [148] Quanzeng You, Jiebo Luo, Hailin Jin, and Jianchao Yang. 2015. Robust image sentiment analysis using progressively trained and domain transferred deep networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 381–388.
- [149] Quanzeng You, Jiebo Luo, Hailin Jin, and Jianchao Yang. 2016. Building a large scale dataset for image emotion recognition: The fine print and the benchmark. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 308–314.
- [150] Jianbo Yuan, Sean Mcdonough, Quanzeng You, and Jiebo Luo. 2013. Sentitribute: Image sentiment analysis from a mid-level perspective. In *Proceedings of the ACM International Workshop on Issues of Sentiment Discovery and Opinion Mining*. 10.
- [151] Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2017. Tensor fusion network for multimodal sentiment analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 1103–1114.
- [152] Amir Zadeh, Paul Pu Liang, Soujanya Poria, Prateek Vij, Erik Cambria, and Louis-Philippe Morency. 2018. Multi-attention recurrent network for human communication comprehension. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [153] Marcel Zentner, Didier Grandjean, and Klaus R. Scherer. 2008. Emotions evoked by the sound of music: Characterization, classification, and measurement. *Emotion* 8, 4 (2008), 494–521.
- [154] Chi Zhan, Dongyu She, Sicheng Zhao, Ming-Ming Cheng, and Jufeng Yang. 2019. Zero-shot emotion recognition via affective structural embedding. In *Proceedings of the IEEE International Conference on Computer Vision*.
- [155] Kejun Zhang, Hui Zhang, Simeng Li, Changyuan Yang, and Lingyun Sun. 2018. The PMemo dataset for music emotion recognition. In *Proceedings of the ACM International Conference on Multimedia Retrieval*. 135–142.
- [156] Yanhao Zhang, Lei Qin, Rongrong Ji, Sicheng Zhao, Qingming Huang, and Jiebo Luo. 2016. Exploring coherent motion patterns via structured trajectory learning for crowd mood modeling. *IEEE Trans. Circ. Syst. Vid. Technol.* 27, 3 (2016), 635–648.
- [157] Sicheng Zhao, Guiguang Ding, Yue Gao, and Jungong Han. 2017. Approximating discrete probability distribution of image emotions by multi-modal features fusion. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 4669–4675.
- [158] Sicheng Zhao, Guiguang Ding, Yue Gao, and Jungong Han. 2017. Learning visual emotion distributions via multi-modal features fusion. In *Proceedings of the ACM International Conference on Multimedia*. 369–377.
- [159] Sicheng Zhao, Guiguang Ding, Yue Gao, Xin Zhao, Youbao Tang, Jungong Han, Hongxun Yao, and Qingming Huang. 2018. Discrete probability distribution prediction of image emotions with shared sparse learning. *IEEE Trans. Affect. Comput.* (2018).
- [160] Sicheng Zhao, Guiguang Ding, Jungong Han, and Yue Gao. 2018. Personality-aware personalized emotion recognition from physiological signals. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- [161] Sicheng Zhao, Guiguang Ding, Qingming Huang, Tat-Seng Chua, Björn W. Schuller, and Kurt Keutzer. 2018. Affective image content analysis: A comprehensive survey. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 5534–5541.
- [162] Sicheng Zhao, Yue Gao, Xiaolei Jiang, Hongxun Yao, Tat-Seng Chua, and Xiaoshuai Sun. 2014. Exploring principles-of-art features for image emotion recognition. In *Proceedings of the ACM International Conference on Multimedia*. 47–56.
- [163] Sicheng Zhao, Amir Gholaminejad, Guiguang Ding, Yue Gao, Jungong Han, and Kurt Keutzer. 2019. Personalized emotion recognition by personality-aware high-order learning of physiological signals. *ACM Trans. Multim. Comput., Commun. Appl.* 15, 1s (2019), 14.
- [164] Sicheng Zhao, Zizhou Jia, Hui Chen, Leida Li, Guiguang Ding, and Kurt Keutzer. 2019. PDANet: Polarity-consistent deep attention network for fine-grained visual emotion regression. In *Proceedings of the ACM International Conference on Multimedia*.
- [165] Sicheng Zhao, Chuang Lin, Pengfei Xu, Sendong Zhao, Yuchen Guo, Ravi Krishna, Guiguang Ding, and Kurt Keutzer. 2019. CycleEmotionGAN: Emotional semantic consistency preserved cycleGAN for adapting image emotions. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 2620–2627.
- [166] Sicheng Zhao, Hongxun Yao, Yue Gao, Guiguang Ding, and Tat-Seng Chua. 2018. Predicting personalized image emotion perceptions in social networks. *IEEE Trans. Affect. Comput.* 9, 4 (2018), 526–540.
- [167] Sicheng Zhao, Hongxun Yao, Yue Gao, Rongrong Ji, and Guiguang Ding. 2017. Continuous probability distribution prediction of image emotions via multi-task shared sparse regression. *IEEE Trans. Multim.* 19, 3 (2017), 632–645.
- [168] Sicheng Zhao, Hongxun Yao, Yue Gao, Rongrong Ji, Wenlong Xie, Xiaolei Jiang, and Tat-Seng Chua. 2016. Predicting personalized emotion perceptions of social images. In *Proceedings of the ACM International Conference on Multimedia*. 1385–1394.

- [169] Sicheng Zhao, Hongxun Yao, Xiaolei Jiang, and Xiaoshuai Sun. 2015. Predicting discrete probability distribution of image emotions. In *Proceedings of the IEEE International Conference on Image Processing*. 2459–2463.
- [170] Sicheng Zhao, Hongxun Yao, Xiaoshuai Sun, Xiaolei Jiang, and Pengfei Xu. 2013. Flexible presentation of videos based on affective content analysis. In *Proceedings of the International Conference on Multimedia Modeling*. 368–379.
- [171] Sicheng Zhao, Hongxun Yao, You Yang, and Yanhao Zhang. 2014. Affective image retrieval via multi-graph learning. In *Proceedings of the ACM International Conference on Multimedia*. 1025–1028.
- [172] Sicheng Zhao, Xin Zhao, Guiguang Ding, and Kurt Keutzer. 2018. EmotionGAN: Unsupervised domain adaptation for learning discrete probability distributions of image emotions. In *Proceedings of the ACM International Conference on Multimedia*. 1319–1327.
- [173] Sheng-hua Zhong, Jiaxin Wu, and Jianmin Jiang. 2019. Video summarization via spatio-temporal deep architecture. *Neurocomputing* 332 (2019), 224–235.
- [174] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*. 2223–2232.
- [175] Xinge Zhu, Liang Li, Weigang Zhang, Tianrong Rao, Min Xu, Qingming Huang, and Dong Xu. 2017. Dependency exploitation: A unified CNN-RNN approach for visual emotion recognition. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 3595–3601.
- [176] Yingying Zhu, Zhengbo Jiang, Jianfeng Peng, and Sheng-hua Zhong. 2016. Video affective content analysis based on protagonist via convolutional neural network. In *Proceedings of the Pacific Rim Conference on Multimedia*. 170–180.
- [177] Yingying Zhu, Min Tong, Zhengbo Jiang, Shenghua Zhong, and Qi Tian. 2019. Hybrid feature-based analysis of video's affective content using protagonist detection. *Exp. Syst. Appl.* 128 (2019), 316–326.
- [178] Athanasia Zlatintsi, Petros Koutras, Georgios Evangelopoulos, Nikolaos Malandrakis, Niki Efthymiou, Katerina Pas-tra, Alexandros Potamianos, and Petros Maragos. 2017. COGNIMUSE: A multimodal video database annotated with saliency, events, semantics and emotion with application to summarization. *EURASIP J. Image Vid. Proc.* 1 (2017), 54.

Received August 2019; revised September 2019; accepted September 2019