

Affect2MM: Affective Analysis of Multimedia Content Using Emotion Causality

Trisha Mittal¹, Puneet Mathur¹, Aniket Bera¹, Dinesh Manocha^{1,2}

¹Department of Computer Science, University of Maryland, College Park, USA

²Department of Electrical and Computer Engineering, University of Maryland, College Park, USA

{trisha, puneetm, bera, dmanocha}@umd.edu

Project URL: <https://gamma.umd.edu/affect2mm>

Abstract

We present *Affect2MM*, a learning method for time-series emotion prediction for multimedia content. Our goal is to automatically capture the varying emotions depicted by characters in real-life human-centric situations and behaviors. We use the ideas from emotion causation theories to computationally model and determine the emotional state evoked in clips of movies. *Affect2MM* explicitly models the temporal causality using attention-based methods and Granger causality. We use a variety of components like facial features of actors involved, scene understanding, visual aesthetics, action/situation description, and movie script to obtain an affective-rich representation to understand and perceive the scene. We use an LSTM-based learning model for emotion perception. To evaluate our method, we analyze and compare our performance on three datasets, *SENDv1*, *MovieGraphs*, and the *LIRIS-ACCEDE* dataset, and observe an average of 10 – 15% increase in the performance over *SOTA* methods for all three datasets.

1. Introduction

In affective computing, perceiving the emotions conveyed in images and videos has found applications in digital content management [30, 73], digital marketing [42, 25, 80], education [17, 2], and healthcare [14]. Such applications have resulted in automated ranking systems, indexing systems [72, 75, 32], and more personalized movie recommendation systems [49]. Affective analysis of movies has been a problem of interest in the community [58, 29], along similar lines. In our work, we explore the problem of affective analysis of movies with the goal of understanding the emotions that the movies invoke in the audience.

There has been a growing interest [51] in dynamically modeling the emotions over time (‘time series emotion recognition’ among the affective computing community. This underlying problem uses temporally continuous data (facial features, speech feature, or other modality features) from multimedia content as input and predicts the emotion labels at multiple timestamps (clips) of the input.

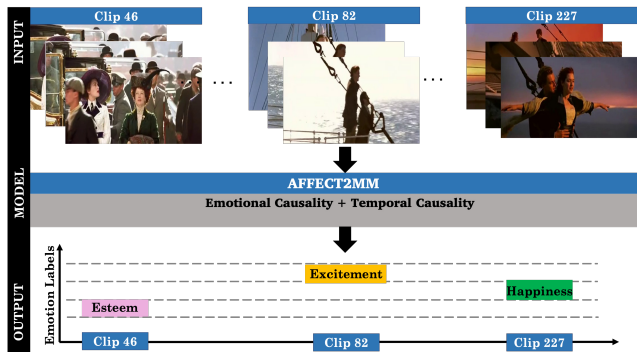


Figure 1: **Time-Series Emotion Perception Model:** We present *Affect2MM*, a learning model for time-series emotion perception for movies. We input a multimedia content in the form of multiple clips and predict emotion labels for each clip. *Affect2MM* is based on the theory of emotion causation and also borrows idea for temporal causality. We show some example clip-frames from the movie ‘Titanic’, a part of the *MovieGraphs* Dataset and corresponding emotion labels.

To aid in solving this time-series problem, several time-series emotion datasets have been proposed [44, 25, 69, 35, 6, 51]. While these datasets focus more on single-person emotional narratives recorded in controlled settings, multimedia datasets (movie databases) like *LIRIS-ACCEDE* [8] and *MovieGraphs* [71] (annotated for per-clip emotion labels) are also being explored for time-series emotion perception tasks.

There have been various efforts to understand how humans reason and interpret emotions resulting in various theories of emotion causation based on physiological, neurological, and cognitive frameworks. One such theory is the ‘emotional causality’ [15] that has been developed from the Causal Theory of Perception [27] and Conceptual Metaphor Theory [4, 48, 37]. ‘Emotional Causality’ refers to the understanding that an experience of emotion is embedded in a chain of events comprising of an (a) outer event; (b) an emotional state; and (c) a physiological response. Few works have explored such emotional causality for emotion perception in multimedia tasks.

Movies, as a time-series multimedia content, model mul-

multiple human-centric situations and are temporally very long, but coherent sequences. To be able to reason about emotions invoked at various clips of the movie, it is important to develop a causal understanding of the story. Generic methods of handling such temporality include recurrent neural networks [38, 63], attention-mechanisms [12], graph modeling [23], and statistical methods like Granger causality [21]. Explicit modeling of causality in the context of time-series emotion perception has been relatively unexplored.

Emotion labels have been explored extensively, both as discrete [36] and continuous [46], in affective analysis. The Valence-Arousal-Dominance (VAD) model [46] is used for representing emotions in a continuous space on a 3D plane with independent axes for valence, arousal, and dominance values. The Valence axis indicates how pleasant (vs. unpleasant) the emotion is; the Arousal axis indicates how high (or low) the physiological intensity of the emotion is, and the dominance axis indicates how much the emotion is tied to the assertion of high (vs. low) social status. A combination of 3 values picked from each axis represents a categorical emotion like ‘angry’ or ‘sad’, much like how an (x, y, z) point represents a physical location in 3-D Euclidean space. Various transformations [45, 24] can be used to map discrete emotion labels to the VAD space. In this work, we work with both continuous emotion labels and discrete emotion labels.

Main Contributions: The following are the novel contributions of our work.

1. We present Affect2MM, a learning-based method for capturing the dynamics of emotion over time. Affect2MM aligns with the psychological theory of “emotional causality” to better model the emotions evoked by each clip of a movie.
2. To better model the temporal causality in movies for long-range multimedia content like movies, we use attention methods and Granger causality to explicitly model the temporal causality (between clips in movies). Our approach can be used for predicting both continuous emotion labels (valence and arousal) and also discrete class labels.

We evaluate our method on two movie datasets, MovieGraphs [71] and the LIRIS-ACCEDE [8] dataset. To showcase our method’s generalizability, we also evaluate and compare our method on the SENDv1 [51] dataset, a single-person emotional narratives dataset.

2. Related Work

In this section, we summarize prior work done in related domains. We first look into available literature in affective analysis of multimedia content and various applications in Section 2.1. In Section 2.2, we discuss the visual affective representations that have previously been explored for related tasks. We also discuss conventional methods to model

temporality in Section 2.3. In Section 2.4, we discuss “emotional causality” and other theories of emotion, as suggested in Psychology literature and the need to align computation models with these theories.

2.1. Affective Analysis of Multimedia Content

Various approaches have explored understanding emotions evoked by multimedia content. Chen et al. [12], Ali et al. [1], Wei et al. [74] have performed affective multi-class classification on images collected from popular websites. Pilli et al. [57], Hussain et al. [26], and Zhang et al. [82] studied predicting sentiments in image advertisements. Vedula et al. [70] extended this idea and developed an advertisement recommendation system using sentiments in advertisement content. Understanding the relationship between emotional responses to content has been learned by recording viewer responses in many controlled use studies. Based on such studies, researchers have used facial responses [42, 31, 47], facial electromyography (EMG) [47], electroencephalogram (EEG), pupillary response and gaze [67, 64], smile [43, 67, 78]. Similarly, Philippot [56] and Gross and Levenson [22] were the first ones to propose a small dataset of clips from films with participants responses to them in controlled lab settings. McDuff et al. [42] have recognized the constraint of collecting such data in controlled settings and have proposed collecting large-scale viewer data using webcams. Other movie-based datasets with some affective annotations that were used are HUMAINE [16], FilmStim [60], DEAP [34], MAHNOB-HCI [64], EMDb [11], MovieGraphs [71] and LIRIS-ACCEDE [8]. In our work, we evaluate our method on some of these datasets.

2.2. Visual Affective Rich Representation

Affective analysis of multimedia from the viewer’s response (face and body posture reactions, EEG, and ECG signals) is not scalable due to lack of data. Subsequent efforts are being made to perform the same analysis using cues directly from the content: images/videos/movies. Wei et al. [74], and Panda et al. [53] report that general visual feature extractions used for standard vision tasks (object recognition) do not scale up in terms of performance for emotion-related tasks. Scene and place descriptors extracted from the image/frames have been explored [1] to understand the affective component better. Researchers [58, 29, 7] have also focused on the visual aesthetics of the multimedia content to understand how they affect the evocation of emotions from viewers. Zhao et al. [84] have also used background music to analyze the affective state. In our work, we present a total of 6 features that we believe can help in a better understanding of multimedia content.

2.3. Modeling Temporality in Time-Series Models

While emotional causality guides us to study the emotional state of a clip, it becomes important to keep track

of and model the temporality in emotions in long multimedia content like movies. While recurrent neural network architectures are inherently designed to keep track of such dependencies, explicitly modeling temporal causality has been a norm. Attention-based methods and their many variants [13, 12] have been used before to help models learn important parts to “attend to” in time series data. Furthermore, most commonly, causality is frequently studied using graph structures and modeling [23, 20]. More conventionally, statistical methods like Granger causality [21] have been used to quantify the dependence of past events into the future time series. More recently [62, 66], there have been multiple attempts to incorporate similar behaviors in neural networks.

2.4. Theory of Emotions: A Look Into Psychology

Some of the major theories of emotion that reason causation of emotion can be grouped into physiological, neurological, and cognitive theories. Physiological theories like the James-Lange Theory [28] and the Cannon-Bard Theory [10] suggest an external stimulus leads to a physiological reaction, and the emotional reaction is dependent on how the physical reaction is interpreted. Schachter-Singer Theory [59] is based on cognition propose, a two-factor theory according to which a stimulus leads to a physiological response that is then cognitively interpreted and labeled, resulting in an emotion. Many more such theories attempt to understand how humans think about emotional states, also called “affective cognition” by [52]. They also provide a taxonomy of inferences within affective cognition and also provide a model of the intuitive theory of emotion modeled as a Bayesian network. Another term scholars have used to understand this domain is “emotional causality” [15].

3. Background and Overview

We formally state our problem in Section 3.1. We then give a brief overview of co-attention mechanism (Section 3.2) and Granger causality (Section 3.3).

3.1. Problem Statement

We consider multimedia content, which can be any image, video, or audio in the form of multiple clips. Each clip, \mathcal{C} , is a short sequence of frames that contains information arising from multiple modalities such as facial expressions of the actors, their speech, the transcribed dialogues, the visual aesthetics of the scene, the description of the scene, and so on. Our goal is to predict the emotion label of each clip using the information contained in the corresponding frames. More formally,

Problem 3.1. *Given a set of clips spanning T time-steps, $\mathcal{C}^{1:T} = \{f_1^{1:T}, f_2^{1:T}, \dots, f_p^{1:T}\}$, where $f_i^{1:T}$ denotes the i^{th} feature for each clip, we predict the emotion label, y , for each clip, denoted by $y^{1:T}$.*

Our formulation allows y to be general in that it may represent either categorical emotion labels (“happy”, “sad”, ...)

or continuous real-valued labels, depending on the dataset. In this work, y can be either represent 26 categorical labels or one of the 2 real-valued labels - valence and arousal.

3.2. Co-attention Mechanism

Attention mechanisms [5] in neural networks have a long history in both the NLP and vision communities. The broad idea is to allow the learning model to only attend to relevant parts of the input rather than encoding the entire input. While intuitive, modeling it requires the computation of attention scores, a weighted matrix depicting the dependence of all outputs on each of the inputs. These attention scores are then used for making inferences. There have been many variants like self-attention [12], hierarchical and nested attention [79], and attention flow [61]. One such variant is the co-attention mechanism [77].

The co-attention mechanism calculates the shallow semantic similarity between the two inputs and uses that as a reference. Given two inputs ($\{u_p^t\}_{t=1}^{t=T}, \{u_q^t\}_{t=1}^{t=T}$), where T is the timestep, co-attention aligns them by constructing a soft-alignment matrix S . Each (i, j) entry of the matrix S is the multiplication of the tanh activation for both the inputs.

$$S_{i,j} = \tanh(w_p u_p^i) \cdot \tanh(w_q u_q^j)$$

We feed these inputs through a single-layer neural network followed by a softmax function to generate the attention distribution α [83].

$$\begin{aligned} z &= \tanh(w_p u_p^i \oplus w_q u_q^j) \\ \alpha &= \text{softmax}(w_\alpha z) \end{aligned} \quad (1)$$

We use \oplus to denote the concatenation of the two features. Based on the attention distribution, the attended/relevant parts can be obtained as follows:

$$\hat{u}_q^j = \sum_i \alpha_{i,j} \cdot u_p^i$$

3.3. Granger Causality (GC)

Granger causality [21] is a way of quantifying the extent to which the past activity of some time-series data is predictive of an other time-series data. In our case, we use the term “time-series data” to refer to an element of $\mathcal{C}^{1:T} = \{f_1^{1:T}, f_2^{1:T}, \dots, f_p^{1:T}\}$. The p features may include facial action units of the actors involved, the audio features, embeddings of the text/transcript, scene descriptions, action or situation descriptions, visual aesthetics of the clips, etc. In other words, the sequence of, say, facial features across T clips, $f_1^{1:T}$ is a time-series modality data. We use GC to reason the causality between different modalities connected across time. If there exists a causality between $f_1^{1:T}$ and $f_2^{1:T}$, then we say that f_1 Granger-causes f_2 .

In our approach, we first explore the existence of Granger causality between time-series modalities in temporal clips. If GC is found to exist between two modalities, then we show that it can be used to improve the accuracy of any general emotion recognition system.

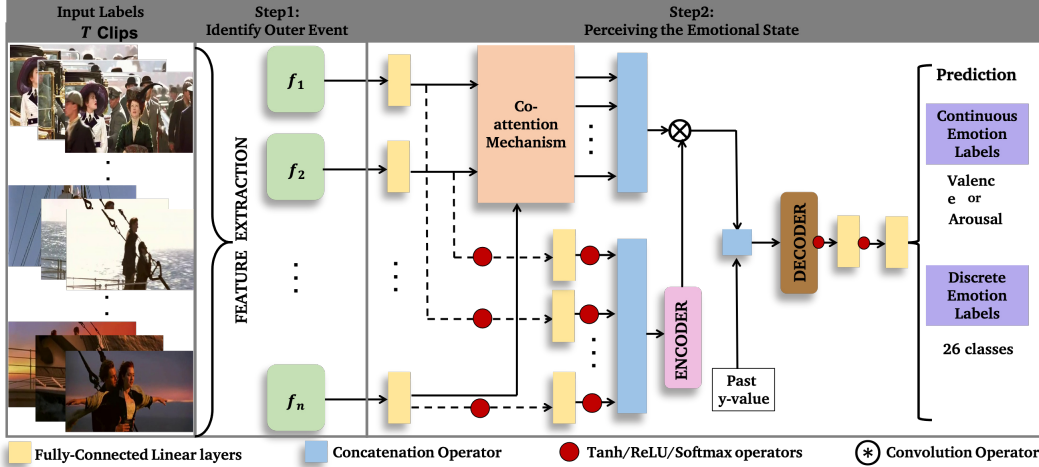


Figure 2: **Affect2MM**: We use the components of Emotional Causality to infer the emotional state depicted in the multimedia content. Given the input T clips, we first extract features, $f_1 \dots f_n$, to build the affective-rich representation and help identify the outer event/cause. We then pass these feature representations through a co-attention mechanism, cLSTM-based encoder, and decoder to perceive the emotional state in either continuous values (valence/arousal) or discrete labels. The networks consist of fully-connected layers (yellow), concat layers (blue), and non-linearity (red dots).

Existence of GC in LSTMs: The most widely-used model to estimate GC in linear systems is the Vector Autoregressive (VAR) model [41, 40]. However, solutions for the linear models do not scale to non-linear time series data [68, 62]. Tank et al. [66] address non-linear GC in neural network architectures like Multi-Layer Perceptron (MLP) and LSTMs by introducing a new LSTM architecture called component-wise LSTM (cLSTM), which models each time-series modality through independent LSTM networks. More formally, consider input

$$X_{1:T} = \Pi \left(\begin{bmatrix} f_{1,1} & f_{1,2} & \dots & f_{1,T} \\ f_{2,1} & f_{2,2} & \dots & f_{2,T} \\ \vdots & \vdots & \ddots & \vdots \\ f_{p,1} & f_{p,2} & \dots & f_{p,T} \end{bmatrix} \right) \quad (2)$$

where the j^{th} row of $X_{1:T}$ is denoted as $X^{(j)}$ and linear transformation Π is applied row-wise (more details in Section 4.3). To check the existence of GC, we begin by passing $X_{1:T}$ to the cLSTM which is made up of p separate LSTMs, $L^{(j)}$, $j = 1, 2, \dots, p$, each operating on the input $X_{1:T}$. In our system, this stage is executed using Equation 4.3. Then, following the theoretical framework put forth in Tank et al. [66], we check for GC by solving the following optimization problem using line search [3]:

$$W^{*(j)} = \arg \min_{W^{(j)}} \sum_t (x_{i,t} - L^{(j)}(X_{1:T}))^2 + \lambda \sum_{k=1}^p \left\| W_{:k}^{(j)} \right\|_2, \quad (3)$$

If $\left\| W_{:k}^{*(j)} \right\|_2 = 0$, then the cLSTM identifies the k^{th} time series modality to be *Granger non-causal* to the j^{th} time-series modality. Otherwise, it is Granger-causal. In Equation 3, $W^{(j)}$ denotes the weights of the LSTM $L^{(j)}$ and $W_{:k}^{(j)}$ denotes the k^{th} column of $W^{(j)}$.

4. Affect2MM: Our Approach

We first give an overview of our method and some notations used in Section 4.1. This is followed by a description of each of the individual components of Affect2MM in Section 4.2, 4.3, 4.3, and 4.3.

4.1. Overview

We present an overview of our time-series emotion prediction model for multimedia content, Affect2MM in Figure 2. Our approach draws on the theory of Emotional Causality to infer the emotional state that is conveyed in the multimedia content. The theory of Emotion Causality [15] consists of the following main events:

1. **Identifying Outer Event:** This stage refers to a stimulus that is contained in the multimedia content that causes an emotion in the consumer of that multimedia content. Such emotion-causing stimuli in multimedia commonly include affective cues such as facial expressions and speech of the actors, but often also include the visual aesthetic features as well as the context of the multimedia content.

We work along these lines and extract these affective cues and visual aesthetic features from the multimedia content and build an affective-rich representation, described further in Section 4.2.

2. **Perceiving the Emotional State:** This event refers to the formation of an emotional state in the consumer of the multimedia content upon receiving the stimuli contained in the latter. We develop a novel co-attention-based deep neural network that predicts a perceived emotion conveyed through multimedia.

Table 1: **Dataset Information:** We summarise details of the three datasets we used and evaluated Affect2MM on. We provide details about the dataset splits, features used and the evaluation metrics. We also mention emotion labels (C:continuous and D:discrete) available and used for training.

Dataset	Train/Val/Test	Affective Features						Labels (C/D)	Evaluation Metrics
		Facial	Audio	Textual	VA	Scene	Situation		
SEND [52]	144/39/41 videos	✓	✓	✓	×	×	×	Valence (C)	CCC
MovieGraphs [71]	33/7/10 movies	✓	✓	✓	✓	✓	✓	26 classes (D)	Accuracy
LIRIS-ACCEDE [8]	44/10/12 movies	✓	✓	×	✓	✓	×	Valence/Arousal (C)	MSE

Lastly, the theory discussed thus far assumes a notion of causality between the two events. That is, the “outer event” *causes* the “perceived emotional state”. In order to computationally model this causality, we investigate the causality between the affective cues using Granger causality [21]. In the following sections, we describe the computational details of the two events of Emotion Causality.

4.2. Building the Affective-Rich Representation

Our goal here is to build a representation of features that are capable of inciting a perceived emotional state in the audience. We extract features that can contribute to an affective-rich representation from the input $\mathcal{C}^{1:T}$ clips of the multimedia content. For each clip, we extract at most p feature representations, including but not limited to– the facial action units of the actors involved, the audio features, embeddings of the text/transcript, scene descriptions, action or situation descriptions and visual aesthetics of the clips,

$$f_i^{1:T} = \mathcal{F}_i(\mathcal{C}^{1:T}) \quad (4)$$

where $f_i^{1:T} \in \mathbb{R}^{p \times |\text{feature size}|_i \times T}$ for $i = 1, 2, \dots, p$ is the i^{th} feature representation obtained using the extractor, \mathcal{F}_i . We describe the feature extractor \mathcal{F} in Section 5.2. Affect2MM can work with any subset of the 6 features representations mentioned. This is also shown in our results.

4.3. Perceiving the Emotional State

We use the features, $f_1^{1:T}, f_2^{1:T} \dots f_p^{1:T}$, generated from the first event to predict perceived emotions in the consumer of the multimedia content. Our approach consists of a deep neural network that uses co-attention to learn and, ultimately, to be able to focus on the useful elements of the features at different time instances. The key intuition here is that the relevance of features varies during the length of the multimedia. For example, the scene-setting and the visual aesthetics may stand out towards the beginning of a movie, as the viewers are not acquainted with the actors, rather are trying to build the scene in their minds. But later, the facial expressions and the speech of the actors may develop a stronger presence as the viewers get to know the actors and their stories. Co-attention helps capture the time-varying nature of the pairwise interdependent affective-rich features, implicitly handling transitivity amongst related modalities.

We begin by simultaneously encoding the feature representation (Eqn. 4.3) using recurrent neural network ar-

chitectures (cLSTMs in our case) and computing the co-attention between pairs of features (Eqn. 7). The results from these operations are convolved to obtain a final context vector (Eqn. 8). The final perceived emotion label is computed by passing the context vector through an LSTM decoder, followed by a combination of linear and non-linear matrix multiplication operations (Eqn. 10).

cLSTM Encoder: We encode the p time-series features using a cLSTM. As described in Section 3.3, we first compute $X_{1:T}$ using Equation 2. More formally, the transformation, $\Pi(f_i^{1:T})$ is given by,

$$\Pi(\cdot) = \text{softmax}(\phi(\cdot))$$

Then, the transformed feature representation of each time-series is computed as,

$$x_i^{1:T} = \Pi(f_i^{1:T}), \quad (5)$$

where ϕ is a linear operation with suitable weights. We can then obtain $X^{1:T}$ by row-wise stacking $x_i^{1:T}$ as follows,

$$X^{1:T} = x_1^{1:T} \oslash x_2^{1:T} \oslash \dots \oslash x_p^{1:T},$$

where \oslash is a row-wise stacking operator. The stacked inputs are encoded using the cLSTM encoder as defined in Eq. 3.

$$h_{\text{enc}} = \text{cLSTM}(X^{1:T}) \quad (6)$$

Co-attention Scores: We learn the interactions between the different features $f_1^{1:T}, f_2^{1:T} \dots f_p^{1:T}$ by aligning and combining modalities pairwise using Eq. 1. We obtain m values of α_k , where $k = 1, 2, \dots, m$ and $m = \binom{p}{2}$, corresponding to each pairwise co-attention operation.

$$\alpha_k = \text{Co-attention}(\phi(f_{k_1}^{1:T}), \phi(f_{k_2}^{1:T})) \quad (7)$$

where k_1, k_2 are indices denoting one of the p feature vectors, $\phi(\cdot)$ are linear layer operators with appropriate weights and \oplus is the concatenation operator. We obtain a final α as,

$$\alpha = \alpha_1 \oplus \alpha_2 \oplus \dots \oplus \alpha_m$$

Decoder: Finally, once we have computed h_{enc} and α , we can obtain the ‘context vector’ d , by convolving the attention weights (α) with the encoded feature representation, h_{enc} ,

$$d = h_{enc} \otimes \alpha \quad (8)$$

The idea is to retain more of the corresponding values of h_{enc} where the attention weights α are high and less information from the parts where the attention weights are low. Finally, the decoder uses the ‘context vector’, d and the past emotion values y' , concatenated together. We simply concatenate these two vectors and feed the merged vector to the decoder. This returns \hat{y} the predicted labels.

$$h_{dec} = \text{LSTM}(d \oplus y') \quad (9)$$

$$\hat{y} = \phi(\text{ReLU}(\phi(h_{dec}))) \quad (10)$$

Vector Auto Regressive (VAR) training of Shared cLSTM Encoder: The cLSTM encoder is shared in a multitask learning fashion to regress future values of input multimodal time-series data $\{f_1^{1:T}, f_2^{1:T}, \dots, f_p^{1:T}\}$ through vector autoregressive training [66] as shown in Equation 3. The VAR training can be viewed as a secondary task to the primary emotion prediction task, involving shared encoder layers. The group lasso penalty applied to the columns of the $W_{:k}^{(j)}$ matrix forces the cLSTM to predict the future values of k^{th} modality without relying on the past values of j^{th} modality. The ridge regularization penalty (r), parameter for non-smooth regularization (l) and the learning rate of VAR training determine the sparsity of the Granger Causal relationship matrix [66] to mitigate the problem of multicollinearity amongst the multivariate time series input.

5. Implementation Details

We give an overview of the datasets (SENDv1, LIRIS-ACCEDE, and MovieGraphs) used for evaluating Affect2MM in Section 5.1. In Section 5.2, we mention the features extracted for training. Finally, to aid reproducibility of our work, we list the training hyperparameters in Section 5.3.

5.1. Datasets

Here we discuss details of the three datasets we used to evaluate and benchmark Affect2MM. For further readability, we have summarized these details in Table 1.

SENDv1 Dataset: The dataset consists of video clips of people recounting important and emotional life stories unscripted. The videos have been recorded in a face-centered setting with no background. Valence ratings collected are provided for every 0.5 seconds of the video.

Evaluation Metrics: Most previous works in this dataset have reported the Concordance Correlation Coefficient (CCC) [39] for validation and test splits along with the standard deviation. The CCC captures the expected discrepancy between the two vectors, compared to the

expected discrepancy if the two vectors were uncorrelated and can be calculated as follows:

$$\text{CCC}(Y, \hat{Y}) = \frac{2 \text{Corr}(Y, \hat{Y}) \sigma_Y \sigma_{\hat{Y}}}{\sigma_Y^2 \sigma_{\hat{Y}}^2 + (\mu_Y - \mu_{\hat{Y}})^2}$$

where $\text{Corr}(Y, \hat{Y})$ is the Pearson correlation between the groundtruth (Y) and predicted valence values (\hat{Y}) for T clips/timestamps of a video, and the μ and σ are the mean and standard deviation predictions.

LIRIS-ACCEDE Dataset: The dataset contains videos from a set of 160 professionally made and amateur movies. The movies are in various languages including English, Italian, Spanish, and French.

Evaluation Metrics: Consistent with prior methods we report the Mean Squared Error (MSE) for the test splits. Given predicted valence value, \hat{y} and true valence value y for T clips of a movie, we compute the MSE as follows:

$$\text{MSE} = \frac{1}{T} \sum_{i=1}^N (y_i - \hat{y}_i)^2$$

MovieGraphs Dataset: This dataset provides detailed graph-based annotations of social situations depicted in movie clips for 51 popular movies. Each graph consists of several types of nodes to capture the emotional and physical attributes of actors, their relationships, and the interactions between them. The dataset was collected and manually annotated using crowd-sourcing methods. We process all clips in every movie and extract the available ‘emotional’ attributes and group them in 26 discrete emotion labels as described by Kosti et al. [36]. Details on creating processed discrete 26 labels are listed in Section A.3.

Evaluation Metrics: Because the labels in MovieGraphs are discrete, detailing 26 emotion classes, we report and compare against the Top-1 accuracy.

5.2. Feature Extraction

To build up an affective rich representation of the multimedia content, we work with a total of 6 features: facial, audio, textual, visual aesthetics, scene, and situation descriptors. We summarize the features available for training Affect2MM in Table 1. An in-depth discussion of these feature extraction steps specific to the datasets has been added in Section A.1.

5.3. Training Hyperparameters

All our results were generated on an NVIDIA GeForce GTX1080 Ti GPU. Hyper-parameters for our model were tuned on the validation set to find the best configurations. We used RMSprop for optimizing our models with a batch size of 1. We experimented with the range of our model’s hyperparameters such as: number of hidden layers (1, 2, 3), size of hidden layers (cLSTM, LSTM Decoder, Dense), dropout {0.2, 0.3, 0.4, 0.5, 0.6}, hidden dimension {64, 128, 256, 512}, and embedding dimension

{64, 128, 256, 512}. The ridge penalty (r) and non-smooth regularization parameter (l) of VAR training of the cLSTM was kept constant at $1e^{-4}$ and 0.001, respectively. The learning rate of both the tasks - emotion prediction and VAR were in this range - $\{1e^{-5}, 1e^{-4}, 1e^{-3}, 1e^{-2}\}$. More specific details on model specific training hyperparameters are added in Section A.2.

6. Results

We first discuss our quantitative results in Section 6.1, where we compare the performance of our method with SOTA methods on the three datasets. We then go over the ablation experiments performed in Section 6.2. Finally, in Section 6.3, we present some qualitative results for Affect2MM.

6.1. Quantitative Results

We compare Affect2MM with SOTA methods on the three datasets.

SENDv1 Dataset: For this dataset we summarise the results in Table 2. We provide CCC score for every combination of the three features/modalities used and also the combined results. Two prior works [52, 50] have experimented with various network architectures and ideas and reported CCC values. We list and compare our performance from their best-performing methods in Table 2.

MovieGraphs Dataset: The dataset has not been previously tested for any similar task. We trained a recently proposed SOTA method, EmotionNet [74], for affective analysis of web images on MovieGraphs dataset and compared our performance with this approach. We summarize this result in Table 3.

LIRIS-ACCEDE Dataset: For this dataset we summarise the results in Table 4. To be consistent with prior methods, we report Means Squared Error (MSE) and Pearsons Correlation Coefficient (PCC) for our method. We compare against 7 existing SOTA methods evaluated on the same dataset. Some of these listed methods were a part of the MediaEval2018 Challenge.

6.2. Ablation Experiments

We perform a small ablation study to analyze the importance of each of the two components we incorporate for modeling the temporal causality in Affect2MM. We report the performance without co-attention between the features and also without Granger causality. These results are summarised in Table 5. As reported, we see a performance improvement of about 4 – 5% across all datasets and another 2 – 3% with the addition of Granger causality.

6.3. Qualitative Results

Time-Series Emotion Predictions: We present qualitative results for Affect2MM in Figures 3 and 4. In Figure 3(a), we show the valence values (range $\in [0, 1]$) learned for a single sample from the test set for SENDv1 Dataset.

Table 2: **Evaluation on SENDv1 Dataset:** for time-series emotion perception of Affect2MM. We report CCC values.

Method	Modalities							
	FA		AT		FT		FAT	
	Val	Test	Val	Test	Val	Test	Val	Test
LSTM [51]	.100	.160	.800	.090	.280	.400	.140	.150
VRNN [51]	.140	.170	.320	.350	.240	.300	.170	.240
SFT [50]	.150	.160	.080	.080	.320	.350	.120	.140
MFT [50]	.060	.080	.360	.330	.400	.360	.420	.440
B3-MFN [50]	0.220	.090	.370	.330	.330	.310	.340	.280
Human [51]	-	-	-	-	-	-	.470	.500
Ours	.557	.582	.592	.601	.556	.567	.599	.597

Table 3: **Evaluation on MovieGraphs:** for time-series emotion perception of Affect2MM for Affective Analysis of Multimedia. We report top-1 accuracy for comparisons.

Method	Validation (Top-1 Acc)	Test (Top-1 Acc)
EmotionNet [74]	35.60	27.90
Ours	39.88	30.58

Table 4: **Evaluation on LIRIS-ACCEDE (MediaEval2018 Dataset):** for time-series emotion perception of Affect2MM for Affective Analysis of Multimedia. We report MSE for comparisons (lower the better).

Method	Valence (MSE)	Arousal (MSE)
CERTH-ITI [7]	0.117	0.138
THUHCSI [29]	0.092	0.140
Quan et al. [58]	0.115	0.171
Yi et al. [81]	0.090	0.136
GLA [65]	0.084	0.133
Ko et al. [33]	0.102	0.149
Zhao et al [84]	0.071	0.137
Ours	0.068	0.128

Table 5: **Ablation Studies:** We perform ablation experiments to understand the importance of co-attention and Granger Causality for modeling temporal causality.

Experiment	SENDv1 (CCC)	MG (Acc)	LIRIS-ACCEDE	
			Valence (MSE)	Arousal (MSE)
Affect2MM w/o (co-attn.& GC)	.570	28.290	0.122	0.143
Affect2MM w/o GC	.585	29.450	0.095	0.135
Affect2MM	.597	30.580	0.068	0.128

In Figure 3(b), we show the predicted valence and arousal values for a sample from LIRIS-ACCEDE Dataset (range $\in [-1, 1]$) along with groundtruth labels. In Figure 3(c), we plot the 27×27 confusion matrix for all the test points in MovieGraphs dataset. The horizontal axis represents the predicted class labels while the vertical axis represents the true labels. The discrete classes are in alphabetical order with 27th class as ‘None’ (Section A.3). Brighter colors on the diagonal indicate more correct classifications.

Interpreting GC matrix: We visualize the Granger causality plots in Figure 4 ((a), (b) and (c)) for all the three datasets. The dimensions of the GC matrix is $|\text{features} \times \text{features}|$. Hence, it is $|3 \times 3|$ for SENDv1




Figure 3: **Qualitative Plots:** (a) We show the valence values learned for a single sample on the test set in the SENDv1 Dataset, and in (b), we show both valence and arousal values from the LIRIS-ACCEDE Dataset. (c) We show the confusion matrix for the MovieGraphs dataset. Brighter colors on the diagonal indicate more number of correct classifications.

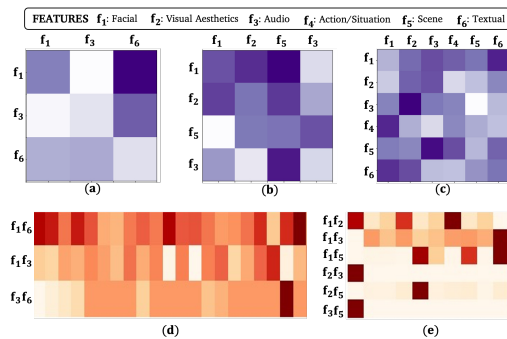


Figure 4: **GC Matrices and Co-attention Weights:** (Top) We show the GC matrices for all three datasets [51, 8, 71]. Darker shades indicate existence of a Granger causality between corresponding features. (Bottom) We plot normalized co-attention weights for one sample each from the (d) SENDv1 and (e) LIRIS-ACCEDE datasets with time-stamps on the x -axis and pairs of features on the y -axis. Darker shades imply a higher co-attention between the pair of the features and more relevance to predictions of the current and the following clips.

dataset, $|4 \times 4|$ for LIRIS-ACCEDE dataset and $|6 \times 6|$ for Moviegraphs dataset. To interpret the GC matrix, we read the features appearing on the horizontal axis and query whether they Granger-cause (or non-cause) the features on the vertical axis. Darker shades of purple indicate causality while lighter shades indicate non-causality. So for example, in (a) f_3 Granger-non-causes f_1 for SENDv1 dataset, and in (b), f_5 Granger-causes f_1 with a higher certainty.

Causality Qualitative Analysis: Here we further analyse co-attention and GC qualitatively. We present the analysis and reasoning of GC and co-attention for one such video from SENDv1 dataset (ID123vid1, timestamp 1:52 in supplementary video), for which the corresponding GC matrix and co-attention weights are shown in Figures 4(a) and (d).

Interpreting Co-attention Weights: We visualize the co-attention weights for one sample each from SENDv1 and LIRIS-ACCEDE in Figure 4 ((d), (e)). We plot the clips/timestamps on the x -axis and pairs of all the features on the y -axis. For every clip, darker shades of red indicate a higher co-attention weight between a pair of modalities.

1. **Text GC Face:** In Figure 4a, the text modality strongly

Granger-causes the facial modality because at 2:20, when the subject says, “... he had a sailor’s mouth ...”, the subject then moments later motions to his face and changes his facial expression to mimic a sailor face.

2. **Text GC Speech:** In Figure 4a, the text Granger-causes speech which is most clearly evidenced when at 2:05, the subject says, “... feel really lucky, really blessed ...”, then subject then audibly sighs with relief.

3. **Speech correlation with text:** In Figure 4d, we see a weak correlation between speech and text in the beginning because when the subject is remembering the good things about his grandfather, the text contains positive stories about the grandfather while the tone in his speech is still sad. However, the correlation between text and speech is stronger at the end of the video when the subject mentions the word, “death”, and then takes a long pause signifying great sadness and loss.

7. Conclusion, Limitations and Future Work

We present Affect2MM, a learning method for time-series emotion prediction for multimedia content. We use Emotional Causality via co-attention and Granger causality for modeling temporal dependencies in the input data and building an affective-rich representation to understand and perceive the scene. We evaluated our method on three benchmark datasets and achieved state-of-the-art results.

There are some limitations with our approach. We need to make the feature extraction process online. Furthermore, our approach does not currently work with scenes containing single actors. In the future, we would like to extend our approach to solve Step 3 of the Emotional Causality Theory to predict the physiological response of the viewer to the movie. We would also like to explore causality and correlation of modalities in single-actor scenes to automate analysis of debates and group discussions. Finally, we will incorporate the ideas presented in this paper with recommendation systems

8. Acknowledgements

This work was supported in part by ARO Grants W911NF1910069, W911NF1910315, W911NF2110026 and Intel.

References

- [1] Afsheen Rafaqat Ali, Usman Shahid, Mohsen Ali, and Jeffrey Ho. High-level concepts for affective understanding of images. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 679–687. IEEE, 2017. 2
- [2] Fehaid Alqahtani and Naeem Ramzan. Comparison and efficacy of synergistic intelligent tutoring systems with human physiological response. *Sensors*, 19(3):460, 2019. 1
- [3] Larry Armijo. Minimization of functions having lipschitz continuous first partial derivatives. *Pacific J. Math.*, 16(1):1–3, 1966. 4
- [4] Angeliki Athanasiadou and Elzbieta Tabakowska. *Speaking of emotions: Conceptualisation and expression*, volume 10. Walter de Gruyter, 2010. 1
- [5] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate, 2016. 3
- [6] Pablo Barros, Nikhil Churamani, Egor Lakomkin, Henrique Siqueira, Alexander Sutherland, and Stefan Wermtner. The omg-emotion behavior dataset. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2018. 1
- [7] Elissavet Batziou, Emmanouil Michail, Konstantinos Avgerinakis, Stefanos Vrochidis, Ioannis Patras, and Ioannis Kompatsiaris. Visual and audio analysis of movies video for emotion detection@ emotional impact of movies task mediaeval 2018. In *MediaEval*, 2018. 2, 7
- [8] Yoann Baveye, Emmanuel Dellandrea, Christel Chamaret, and Liming Chen. Liris-accede: A video database for affective content analysis. *IEEE Transactions on Affective Computing*, 6(1):43–55, 2015. 1, 2, 5, 8
- [9] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *International Conference on Computer Vision*, 2017. 12
- [10] Walter B Cannon. The james-lange theory of emotions: A critical examination and an alternative theory. *The American journal of psychology*, 39(1/4):106–124, 1927. 3
- [11] Sandra Carvalho, Jorge Leite, Santiago Galdo-Álvarez, and Oscar F Gonçalves. The emotional movie database (emdb): A self-report and psychophysiological study. *Applied psychophysiology and biofeedback*, 37(4):279–294, 2012. 2
- [12] Tao Chen, Damian Borth, Trevor Darrell, and Shih-Fu Chang. DeepSentibank: Visual sentiment concept classification with deep convolutional neural networks. *arXiv preprint arXiv:1410.8586*, 2014. 2, 3
- [13] Jianpeng Cheng, Li Dong, and Mirella Lapata. Long short-term memory-networks for machine reading, 2016. 3
- [14] Jeffrey F Cohn, Tomas Simon Kruez, Iain Matthews, Ying Yang, Minh Hoai Nguyen, Margara Tejera Padilla, Feng Zhou, and Fernando De la Torre. Detecting depression from facial actions and vocal prosody. In *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*, pages 1–7. IEEE, 2009. 1
- [15] Maarten Coëgnarts and Peter Kravanja. Perceiving causality in character perception: A metaphorical study of causation in film. *Metaphor and Symbol*, 31(2):91–107, 2016. 1, 3, 4
- [16] Ellen Douglas-Cowie, Roddy Cowie, Ian Sneddon, Cate Cox, Orla Lowry, Margaret Mcrorie, Jean-Claude Martin, Laurence Devillers, Sarkis Abrilian, Anton Batliner, et al. The humane database: Addressing the collection and annotation of naturalistic and induced emotional data. In *International conference on affective computing and intelligent interaction*, pages 488–500. Springer, 2007. 2
- [17] Andrew Downs and Paul Strand. Effectiveness of emotion recognition training for young children with developmental delays. *Journal of Early and Intensive Behavior Intervention*, 5(1):75, 2008. 1
- [18] Rosenberg Ekman. *What the face reveals: sic and applied studies of spontaneous expression using the Facial Action Coding System (FACS)*. Oxford University Press, USA, 1997. 12
- [19] Florian Eyben, Felix Weninger, Florian Gross, and Björn Schuller. Recent developments in opensmile, the munich open-source multimedia feature extractor. In *Proceedings of the 21st ACM international conference on Multimedia*, pages 835–838, 2013. 12
- [20] Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in Genetics*, 10:524, 2019. 3
- [21] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969. 2, 3, 5
- [22] James J Gross and Robert W Levenson. Emotion elicitation using films. *Cognition & emotion*, 9(1):87–108, 1995. 2
- [23] Ruocheng Guo, Lu Cheng, Jundong Li, P. Richard Hahn, and Huan Liu. A survey of learning causality with data: Problems and methods. *ACM Comput. Surv.*, 53(4), July 2020. 2, 3
- [24] Holger Hoffmann, Andreas Scheck, Timo Schuster, Stefan Walter, Kerstin Limbrecht, Harald C Traue, and Henrik Kessler. Mapping discrete emotions into the dimensional space: An empirical approach. In *2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 3316–3320. IEEE, 2012. 2
- [25] Zaeem Hussain, Mingda Zhang, X. Zhang, Keren Ye, C. Thomas, Zuha Agha, Nathan Ong, and Adriana Kovashka. Automatic understanding of image and video advertisements. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1100–1110, 2017. 1
- [26] Zaeem Hussain, Mingda Zhang, Xiaozhong Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, and Adriana Kovashka. Automatic understanding of image and video advertisements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1705–1715, 2017. 2
- [27] John Hyman. The causal theory of perception. *The Philosophical Quarterly (1950-)*, 42(168):277–296, 1992. 1
- [28] WILLIAM JAMES. II.—WHAT IS AN EMOTION ? *Mind*, os-IX(34):188–205, 04 1884. 3
- [29] Zitong Jin, Yuqi Yao, Ye Ma, and Mingxing Xu. Thuhcsi in mediaeval 2017 emotional impact of movies task. *MediaEval*, 17:13–17, 2017. 1, 2, 7
- [30] Dhiraj Joshi, Ritendra Datta, Elena Fedorovskaya, Xin Lu, Quang-Tuan Luong, James Z Wang, Jia Li, and Jiebo Luo. On aesthetics and emotions in scene images: A computational perspective. *Scene vision: making sense of what we see*, page 241, 2014. 1
- [31] Karim Sadik Kassam. *Assessment of emotional experience through facial expression*. Harvard University, 2010. 2

- [32] Kenta Kimura, Satoshi Haramizu, Kazue Sanada, and Akiko Oshida. Emotional state of being moved elicited by films: A comparison with several positive emotions. *Frontiers in Psychology*. 1
- [33] Tobey H Ko, Zhonglei Gu, Tiantian He, and Yang Liu. Towards learning emotional subspace. In *MediaEval*, 2018. 7
- [34] Sander Koelstra, Christian Muhl, Mohammad Soleymani, Jong-Seok Lee, Ashkan Yazdani, Touradj Ebrahimi, and Pun. Deap: A database for emotion analysis; using physiological signals. 2
- [35] Jean Kossaifi, Robert Walecki, Yannis Panagakis, Jie Shen, Maximilian Schmitt, Fabien Ringeval, Jing Han, Vedhas Pandit, Antoine Toisoul, Bjoern W Schuller, et al. Sewa db: A rich database for audio-visual emotion and sentiment research in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 1
- [36] Ronak Kostı, Jose M. Alvarez, Adria Recasens, and Agata Lapedriza. Emotic: Emotions in context dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017. 2, 6
- [37] Zoltán Kövecses. *Metaphor and emotion: Language, culture, and body in human feeling*. Cambridge University Press, 2003. 1
- [38] J Richard Landis and Gary G Koch. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174, 1977. 2
- [39] I Lawrence and Kuei Lin. A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, pages 255–268, 1989. 6
- [40] Aurelie C Lozano, Naoki Abe, Yan Liu, and Saharon Rosset. Grouped graphical granger modeling methods for temporal causal modeling. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 577–586, 2009. 4
- [41] Helmut Lütkepohl. *New introduction to multiple time series analysis*. Springer Science & Business Media, 2005. 4
- [42] Daniel McDuff, Rana El Kaliouby, Jeffrey F Cohn, and Rosalind W Picard. Predicting ad liking and purchase intent: Large-scale analysis of facial responses to ads. *IEEE Transactions on Affective Computing*, 6(3):223–235, 2014. 1, 2
- [43] Daniel McDuff, Rana El Kaliouby, David Demirdjian, and Rosalind Picard. Predicting online media effectiveness based on smile responses gathered over the internet. In *2013 10th IEEE international conference and workshops on automatic face and gesture recognition (FG)*, pages 1–7. IEEE, 2013. 2
- [44] Gary McKeown, Michel F Valstar, Roderick Cowie, and Maja Pantic. The semaine corpus of emotionally coloured character interactions. In *2010 IEEE International Conference on Multimedia and Expo*, pages 1079–1084. IEEE, 2010. 1
- [45] Albert Mehrabian. Analysis of the big-five personality factors in terms of the pad temperament model. *Australian journal of Psychology*, 48(2):86–92, 1996. 2
- [46] Albert Mehrabian and James A Russell. *An approach to environmental psychology*. the MIT Press, 1974. 2
- [47] Anca Cristina Micu and Joseph T Plummer. Measurable emotions: How television ads really work: Patterns of reactions to commercials can demonstrate advertising effectiveness. *Journal of Advertising Research*, 50(2):137–153, 2010. 2
- [48] Susanne Niemeier and René Dirven. *The language of emotions: Conceptualization, expression, and theoretical foundation*. John Benjamins Publishing, 1997. 1
- [49] Eva Oliveira, Pedro Martins, and Teresa Chambel. Ifelt: accessing movies through our emotions. In *Proceedings of the 9th European Conference on Interactive TV and Video*, pages 105–114, 2011. 1
- [50] D. Ong, Z. Wu, Z. Tan, M. Reddan, I. Kahhale, A. Mattek, and J. Zaki. Modeling emotion in complex stories: the stanford emotional narratives dataset. *IEEE Transactions on Affective Computing*, pages 1–1, 2019. 7
- [51] Desmond Ong, Zhengxuan Wu, Zhi-Xuan Tan, Marianne Reddan, Isabella Kahhale, Alison Mattek, and Jamil Zaki. Modeling emotion in complex stories: the stanford emotional narratives dataset. *IEEE Transactions on Affective Computing*, 2019. 1, 2, 7, 8, 12
- [52] Desmond C Ong, Jamil Zaki, and Noah D Goodman. Computational models of emotion inference in theory of mind: A review and roadmap. *Topics in cognitive science*, 11(2):338–357, 2019. 3, 5, 7
- [53] Rameswar Panda, Jianming Zhang, Haoxiang Li, Joon-Young Lee, Xin Lu, and Amit K Roy-Chowdhury. Contemplating visual emotions: Understanding and overcoming dataset bias. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 579–595, 2018. 2
- [54] Yilang Peng and JOHN B JEMMOTT III. Feast for the eyes: Effects of food perceptions and computer vision features on food photo popularity. *International Journal of Communication (19328036)*, 12, 2018. 12
- [55] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014. 12
- [56] Pierre Philippot. Inducing and assessing differentiated emotion-feeling states in the laboratory. *Cognition and emotion*, 7(2):171–193, 1993. 2
- [57] Stephen Pilli, Manasi Patwardhan, Niranjan Pedanekar, and Shirish Karande. Predicting sentiments in image advertisements using semantic relations among sentiment labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 408–409, 2020. 2
- [58] Khanh-An C Quan, Vinh-Tiep Nguyen, and Minh-Triet Tran. Frame-based evaluation with deep features to predict emotional impact of movies. In *MediaEval*, 2018. 1, 2, 7
- [59] Stanley Schachter and Jerome Singer. Cognitive, social, and physiological determinants of emotional state. *Psychological review*, 69(5):379, 1962. 3
- [60] Alexandre Schaefer, Frédéric Nils, Xavier Sanchez, and Pierre Philippot. Assessing the effectiveness of a large database of emotion-eliciting films: A new tool for emotion researchers. *Cognition and emotion*, 24(7):1153–1172, 2010. 2
- [61] Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. Bidirectional attention flow for machine comprehension, 2018. 3
- [62] Vikas Sindhwani, Ha Quang Minh, and Aurélie C. Lozano. Scalable matrix-valued kernel learning for high-dimensional nonlinear multivariate regression and granger causality. In *Proceedings of the Twenty-Ninth Conference on Uncertainty*

- in *Artificial Intelligence*, UAI'13, page 586–595, Arlington, Virginia, USA, 2013. AUAI Press. 3, 4
- [63] Mohammad Soleymani, Sadjad Asghari-Esfeden, Maja Pantic, and Yun Fu. Continuous emotion detection using eeg signals and facial expressions. In *2014 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2014. 2
- [64] Mohammad Soleymani, Maja Pantic, and Thierry Pun. Multimodal emotion recognition in response to videos. *IEEE transactions on affective computing*, 3(2):211–223, 2011. 2
- [65] Jennifer J Sun, Ting Liu, and Gautam Prasad. Gla in mediaeval 2018 emotional impact of movies task. *arXiv preprint arXiv:1911.12361*, 2019. 7
- [66] Alex Tank, Ian Covert, Nicholas Foti, Ali Shojaie, and Emily Fox. Neural granger causality for nonlinear time series, 2018. 3, 4, 6
- [67] Thales Teixeira, Rosalind Picard, and Rana El Kaliouby. Why, when, and how much to entertain consumers in advertisements? a web-based facial tracking field study. *Marketing Science*, 33(6):809–827, 2014. 2
- [68] T. Teräsvirta, D. Tjøstheim, and C. Granger. Modelling nonlinear economic time series. 2011. 4
- [69] George Trigeorgis, Fabien Ringeval, Raymond Brueckner, Erik Marchi, Mihalis A Nicolaou, Björn Schuller, and Stefanos Zafeiriou. Adieu features? end-to-end speech emotion recognition using a deep convolutional recurrent network. In *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5200–5204. IEEE, 2016. 1
- [70] Nikhita Vedula, Wei Sun, Hyunhwan Lee, Harsh Gupta, Mitsunori Ogihara, Joseph Johnson, Gang Ren, and Srinivasan Parthasarathy. Multimodal content analysis for effective advertisements on youtube. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 1123–1128. IEEE, 2017. 2
- [71] Paul Vicol, Makarand Tapaswi, Lluís Castrejon, and Sanja Fidler. Moviegraphs: Towards understanding human-centric situations from videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 5, 8
- [72] Hee Lin Wang and Loong-Fah Cheong. Affective understanding in film. *IEEE Transactions on circuits and systems for video technology*, 16(6):689–704, 2006. 1
- [73] Yilin Wang and Baoxin Li. Sentiment analysis for social media images. In *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*, pages 1584–1591. IEEE, 2015. 1
- [74] Zijun Wei, Jianming Zhang, Zhe Lin, Joon-Young Lee, Niranjan Balasubramanian, Minh Hoai, and Dimitris Samaras. Learning visual emotion representations from web data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13106–13115, 2020. 2, 7
- [75] Norbert Wiley. Emotion and film theory. *Studies in Symbolic Interaction*, 26:169–190, 2003. 1
- [76] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *European Conference on Computer Vision*. Springer, 2018. 12
- [77] Caiming Xiong, Victor Zhong, and Richard Socher. Dynamic coattention networks for question answering, 2018. 3
- [78] Songfan Yang, Mehran Kafai, Le An, and Bir Bhanu. Zapping index: using smile to measure advertisement zapping likelihood. *IEEE Transactions on Affective Computing*, 5(4):432–444, 2014. 2
- [79] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489, San Diego, California, June 2016. Association for Computational Linguistics. 3
- [80] Keren Ye and Adriana Kovashka. Advise: Symbolism and external knowledge for decoding advertisements. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 837–855, 2018. 1
- [81] Yun Yi, Hanli Wang, and Qinyu Li. Cnn features for emotional impact of movies task. In *MediaEval*, 2018. 7
- [82] Huaizheng Zhang, Yong Luo, Qiming Ai, Yonggang Wen, and Han Hu. Look, read and feel: Benchmarking ads understanding with multimodal multitask learning. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 430–438, 2020. 2
- [83] Qi Zhang, Jinlan Fu, X. Liu, and X. Huang. Adaptive coattention network for named entity recognition in tweets. In *AAAI*, 2018. 3
- [84] Yin Zhao, Longjun Cai, Chaoping Tu, Jie Zhang, and Wu Wei. Video affective impact prediction with multimodal fusion and long-short temporal context. 2019. 2, 7