



Shukla, A., Gullapuram, S. S., Katti, H., Yadati, K., Kankanhalli, M. and Subramanian, R. (2017) Affect Recognition in Ads with Application to Computational Advertising. In: 2017 ACM on Multimedia Conference (MM '17), Mountain View, CA, USA, 23-27 Oct 2017, pp. 1148-1156. ISBN 9781450349062.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

© Association for Computer Machinery 2017. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in Proceedings of the 2017 ACM on Multimedia Conference (MM '17), Mountain View, CA, USA, 23-27 Oct 2017, pp. 1148-1156. ISBN 9781450349062. <https://doi.org/10.1145/3123266.3123444>.

<http://eprints.gla.ac.uk/150558/>

Deposited on: 27 October 2017

Affect Recognition in Ads with Application to Computational Advertising

Abhinav Shukla*
International Institute of Information
Technology
Hyderabad, India
abhinav.shukla@research.iiit.ac.in

Shruti Shriya Gullapuram*
International Institute of Information
Technology
Hyderabad, India
shruti.gullapuram@students.iiit.ac.in

Harish Katti
Centre for Neuroscience, Indian
Institute of Science
Bangalore, India
harish2006@gmail.com

Karthik Yadati
Delft University of Technology
Delft, Netherlands
n.k.yadati@tudelft.nl

Mohan Kankanhalli
School of Computing, National
University of Singapore
Singapore, Singapore
mohan@comp.nus.edu.sg

Ramanathan Subramanian
School of Computing Science,
University of Glasgow
Singapore, Singapore
ramanathan.subramanian@glasgow.
ac.uk

ABSTRACT

Advertisements (ads) often include strongly emotional content to leave a lasting impression on the viewer. This work (i) compiles an affective ad dataset capable of evoking coherent emotions across users, as determined from the affective opinions of five experts and 14 annotators; (ii) explores the efficacy of convolutional neural network (CNN) features for encoding emotions, and observes that CNN features outperform low-level audio-visual emotion descriptors [9] upon extensive experimentation; and (iii) demonstrates how enhanced affect prediction facilitates computational advertising, and leads to better viewing experience while watching an online video stream embedded with ads based on a study involving 17 users. We model ad emotions based on subjective human opinions as well as objective multimodal features, and show how effectively modeling ad emotions can positively impact a real-life application.

CCS CONCEPTS

• **Human-centered computing** → **HCI theory, concepts and models**; *User centered design*;

KEYWORDS

Affect Recognition, Advertisements, Human and Computational Perception, Convolutional Neural Networks (CNNs), Multimodal, Computational Advertising

1 INTRODUCTION

Advertising is a huge and profitable industry and advertisers intend to portray their products or services as not only useful, but also highly desirable and rewarding. Emotions are critical for conveying an effective message to viewers, and have been found to mediate consumer attitudes towards brands [10, 11, 19]. Similar objectives are at play in messages for public health and safety, where certain life choices are portrayed as beneficial and improving one’s quality of life, while others are portrayed as harmful and potentially fatal. The ability to objectively quantify advertisements (ads) in terms of emotional content therefore has a wide variety of

applications— *e.g.*, inserting the right type of ads at optimal temporal points within a video stream can beneficially impact both advertisers and consumers in video streaming websites such as YouTube [26, 27]. Subjective experience of pleasantness (valence) and emotional intensity (arousal) are important affective dimensions [21], and both modulate emotional responses to ads in distinct ways [6]. Affective content has also been shown to modulate recall of key concepts and episodes in movies [22] and video ads [27].

However, affect characterization in ads is a non-trivial problem as with other stimuli such as music and movie clips examined by prior works [1, 9, 14, 25]. Given that human emotional perception is subjective and the detection of specific emotions such as *joy*, *sorrow* and *disgust* is relatively hard, popular affect recognition (AR) works represent emotions along the valence and arousal dimensions [8, 21]. Also, methods have been devised to estimate emotions in a *content-centric* or *user-centric* manner. *Content-centric* methods estimate the emotion evoked by a stimulus by examining audio, visual and textual cues; a popular example is [9]. Of late, *user-centered* approaches that model the stimulus-evoked emotion via physiological changes induced in the viewer have gained in popularity [1, 14, 24]. While enabling a fine-grained examination of emotional perception, which is a transient phenomenon, user-centered methods nevertheless suffer from subjectivity limitations.

This work expressly investigates the modeling of emotions conveyed by ads, and employs subjective human opinions and objective multimedia features to this end. Firstly, upon carefully compiling a diverse set of 100 ads, we examine the efficacy of this ad dataset to coherently evoke emotions across viewers. To this end, we compare the affective opinions of five experts and 14 novice annotators, and find that the two groups are highly concordant. Secondly, we explore the utility of Convolutional Neural Networks (CNNs) for encoding audio-visual emotional features. As the compiled ad dataset is relatively small and insufficient for CNN training, we employ *domain adaptation* to transfer affective knowledge gained from the LIRIS-ACCEDE movie dataset [3] for modeling ad emotions. Extensive experimentation confirms that the synthesized CNN descriptors outperform popular audio-visual features proposed in [9] especially for valence recognition. Thirdly, we show how accurate encoding of the ad emotions can facilitate optimized insertion of

ads into streaming video, used for income generation by online websites such as YouTube. A user study with 17 viewers confirms that the insertion of emotionally relevant ads within the streamed video can maximize viewer experience. In summary, we make the following research contributions:

1. Ours is one of the few works to examine AR in ads, and the only work to characterize ad emotions in terms of subjective human opinions and objective audio-visual features.
2. We explore the utility of CNNs for encoding ad emotions. We show the effectiveness of a new CNN, *AdAffectNet* (AAN) generated by fine-tuning the *Places205* CNN architecture [28] for AR. For fine-tuning and domain adaptation, we have employed the extensively annotated LIRIS-ACCEDE movie dataset [3]. Extensive experiments reveal that the AAN features outperform emotional audio-visual descriptors proposed in [9], and the best AR performance is achieved with multi-task learning which exploits audio-visual similarities among emotionally homogeneous ads.
3. We show how improved affect modeling with the CNN features can facilitate the CAVVA ad-in-video insertion strategy [27]. Our AAN model can especially predict ad valence better than the baseline [9], which positively impacts viewer experience.

From here on, Sec 2 discusses related work, while Sec 3 presents the compiled ad dataset and related statistics. Section 4 describes our AAN model and associated AR experiments. Sec 5 presents a user study to evaluate an ad-insertion strategy [27] with different affect encoding methods, while Sec 6 concludes the paper.

2 RELATED WORK

To highlight the novelty of our work, we briefly review prior works examining (i) AR, and (ii) the impact of ad-evoked emotions.

2.1 Affect Recognition

Many approaches have been devised to infer the emotions evoked by multimedia stimuli in a *content-centric* or *user-centric* manner. Content-centric approaches [9, 25] predict the elicited emotion by examining audio-visual cues in the analyzed stimuli. In contrast, user-centric AR methods [1, 14, 23] predict the stimulus-evoked emotion by measuring physiological changes in users (or content consumers). Nevertheless, both content and user centric methods require labels identifying stimulus emotion, and these labels are compiled from reliable annotators whose affective opinions are generally *acceptable*, given human subjectivity in emotion perception.

2.2 Emotional impact of ads

A number of works have studied the impact of ad-induced emotions on user behavior [10, 11, 19]. Holbrook and Batra [10] remark that emotions induced by ads strongly modulate users’ brand attitude. Pham *et al.* [19] conclude that ad-evoked feelings impact users both explicitly and implicitly. Ad-evoked emotions are found to *change* user attitude towards (especially hedonistic) products.

While a body of works have examined the correlation between ad emotions and user behavior, very few works have exploited these findings for developing targeted advertising mechanisms. The only work that incorporates emotional information for modeling context

Table 1: Summary statistics for quadrant-wise ads.

Quadrant	Mean length (s)	Mean asl	Mean val	Mean eng
H asl, H val	48.16	2.17	1.02	2.50
L asl, H val	44.18	1.37	0.91	2.23
L asl, L val	60.24	1.76	-0.76	2.47
H asl, L val	64.16	3.01	-1.16	2.56

in advertising is CAVVA [27], where arousal and valence evoked by video scenes are estimated via [9] to identify optimal ads and corresponding insertion points which maximize user engagement.

2.3 Analysis of related work

Examination of the literature reveals that (1) AR studies are hampered by the subjectivity in emotion perception, and a control dataset that can coherently evoke emotions across users is necessary for effectively learning content or physiology-based emotion predictors; (2) Despite the well-known impact of ad emotions on user behavior, there has hardly been any attempt to incorporate emotion-related findings in a computational advertising framework.

In this regard, we present the first work to compile a control set of affective ads, which elicit concordant affective opinions from experts and naive users. Also, we synthesize CNN-based emotion descriptors which are found to outperform audio visual features proposed in [9]. We also perform a user study and show how better affect encoding can facilitate the ad-insertion framework [27] to improve viewing experience. Details pertaining to our ad dataset are presented below.

3 ADVERTISEMENT DATASET

Defining *valence* as the feeling of *pleasantness/unpleasantness*, *arousal* as the *intensity of emotional feeling* and *engagement* as the *level of interest* while viewing an audio-visual stimulus, five experts carefully compiled a dataset of 100, roughly 1-minute long commercial advertisements (ads) which are used in this work. These ads are publicly available on online video websites and found to be uniformly distributed over the arousal–valence plane defined by Greenwald *et al.* [8] (Fig. 1). An ad was chosen if there was consensus among all five experts on its valence and arousal labels (defined as either *high* (H)/*low* (L)). The high valence ads typically involved product promotions, while low valence ads were social messages depicting the ill effects of smoking, alcohol and drug abuse, *etc.* Labels provided by experts were considered as *ground-truth*, and used for all recognition experiments in this work.

To evaluate the effectiveness of the ads as control stimuli, *i.e.*, examine how consistently they evoke emotions across viewers, 14 novice users rated the ads for valence (val), arousal (asl) and engagement (eng) upon familiarization. All ads were rated on a 5-point scale, which ranged from -2 (*very unpleasant*) to 2 (*very pleasant*) for val, 0 (*calm*) to 4 (*highly aroused*) for asl and 0 (*boring*) to 4 (*highly engaging*) for eng. Table 1 presents summary statistics for ads over the four quadrants. Evidently, low val ads are longer and are perceived as more arousing than high val ads suggesting that they evoked stronger emotional feelings among viewers.

We also computed inter-rater agreement in terms of the (i) Krippendorff’s α and (ii) Cohen’s κ measures. The α coefficient is applicable when multiple raters code data with ordinal scores– we

obtained $\alpha = 0.60, 0.37$ and 0.23 for val, asl and eng implying valence impressions were most consistent across raters. We then computed the κ agreement between annotator and ground-truth labels to determine concordance between the novice and expert groups. To this end, we thresholded each rater’s asl, val scores by their mean rating to assign H/L labels for each ad, and compared these labels with the ground truth. This procedure revealed a mean agreement of 0.84 for val and 0.67 for asl across raters. We also computed κ between the annotator and expert *populations* by thresholding the mean asl, val score per ad across raters against the grand mean– this method returned $\kappa = 0.94$ for val and 0.67 for asl¹. Clearly, there is good-to-excellent agreement between annotators and experts on affective impressions with considerably higher concordance for val. The observed concordance between the independent expert and annotator groups affirms that the compiled 100 ads are suitable control stimuli for affective studies.

Another desirable property of a control affective dataset is the independence of the asl and val dimensions. To this end, we (i) examined scatter plots of the annotator ratings, and (ii) computed correlations amongst those ratings. The scatter plot of the mean asl, val annotator ratings, and the distribution of asl and val ratings are presented in Fig. 1. The scatter plot is color-coded based on expert labels, and interestingly is different from the classical ‘C’ shape observed with images [17], music videos [14] and movie clips [1] owing to the difficulty in evoking medium asl/val but strong val/asl responses. The distributions of asl and val ratings are also roughly uniform resulting in Gaussian fits with large variance, with modes observed at median scale values of 2 and 0 respectively. A close examination of the scatter plot reveals that a number of ads are rated as moderate asl, but high/low val. This can be attributed to the fact that ads are constrained to effectively convey a positive or negative message to viewers, which is not typically true of images or movie scenes. Finally, Wilcoxon rank sum tests on annotator ratings revealed significantly different asl ratings for high and low asl ads ($p < 0.00005$), and distinctive val scores for high and low valence ads ($p < 0.000001$), consistent with expectation.

Pearson correlations among the asl, val and eng dimensions from the annotator ratings as shown in Table 2. As each rater assesses multiple ads, we corrected for multiple comparisons by limiting the false discovery rate to within 5% via the procedure described in [4]. Bold values in Table 2 denote correlations found to be significant ($p < 0.05$) over all raters. Rating correlations reveal that the asl and val dimensions are only weakly correlated, while a positive and significant correlation is noted between asl and eng in line with prior studies [1, 14]. Analyses from here on will focus on the asl and val attributes given the strong connection between asl and eng. The user study (Sec 5) will focus on user engagement/experience.

Overall, examination of ad labels reveals that (i) Our ads constitute a control dataset for affective studies as asl and val ratings are largely uncorrelated; (ii) Also, different from the ‘C’-shape characterizing the asl-val relationship for other stimulus types, asl and val ratings are uniformly distributed for the ad stimuli, and (iii) There is considerable concordance between the expert and annotator populations on affective labels, implying that the selected




Figure 1: (left) Scatter plot of mean asl, val ratings color-coded with expert labels. (middle) Asl and (right) Val rating distribution with Gaussian pdf overlay (view under zoom).

ads effectively evoke coherent emotions across viewers. The next section describes content-centric ad AR.

Table 2: Mean correlations between self-rated attributes. Significant correlations ($p < 0.05$) are denoted in bold.

	asl	val	eng
Arousal	1	-0.19	0.36
Valence		1	0.11
Engagement			1

4 COMPUTATIONAL MODEL

This section describes the datasets and the computational model employed for recognizing emotions conveyed by our ads.

4.1 Datasets

Due to subjective variance in emotion perception, careful affective labeling is imperative for effectively learning content-centric [9, 25] or user-centric [1, 14] affective correlates, which is why we analyze ads that evoked perfect consensus among experts. Of late, convolutional neural networks (CNNs) have become extremely popular for visual [16] and audio [12] recognition, but these models require huge training data. Given the small size of our ad dataset, we fine-tune the pre-trained *Places205* [28] model using the affective LIRIS-ACCEDE movie dataset [3], and employ this fine-tuned model for encoding ad emotions– a process known as *domain adaptation* in machine learning literature.

To learn deep features for modeling ad affect, we employed the *Places205* CNN [28] intended for image classification. *Places205* is trained using the Places-205 dataset comprising 2.5 million images and 205 scene categories. The Places-205 dataset contains a wide variety of scenes with varying illumination, viewpoint and field of view, and we hypothesized a strong relationship between scene perspective, lighting and the scene mood. LIRIS-ACCEDE contains asl, val ratings for ≈ 10 s long movie snippets, whereas our ads are about a minute-long (ranging from 30–120 s).

4.2 FC7 Feature Extraction via CNNs

For ad AR, we represent the *visual* modality using *key-frame* images, and the *audio* modality using *spectrograms*. We fine-tune *Places205* via the LIRIS-ACCEDE [3] dataset to synthesize *AdAffectNet* (AAN), and use the fully connected layer (fc7) AAN descriptors for our analysis.

4.2.1 Keyframes as Visual Descriptors. From each video in the ad and LIRIS-ACCEDE datasets, we sample one *key frame* every three seconds– this enables extraction of a continuous video profile for AR. This process generated 1791 key-frames for our 100 ads.

¹Chance agreement corresponds to a κ value of 0.




Figure 2: Exemplar spectrograms for varied emotional ads. x denotes time (0-10 s), while y denotes spectral magnitude observed at each time instant. High and low frequency densities are respectively shown in red and green shades.

4.2.2 Spectrograms as Audio Descriptors. Spectrograms (SGs) are visual representations of the audio frequency spectrum, and have been successfully employed for AR from speech and music [2]. Specifically, transforming the audio content to a spectrogram image allows for audio classification to be treated as a visual recognition problem. We extract spectrograms over the 10 s long LIRIS-ACCEDE clips, and consistently from 10 s ad segments. This process generates 610 spectrograms for our ad dataset. Following [2], we combine multiple tracks to obtain a single spectrogram (as opposed to two for stereo). Each spectrogram is generated using a 40 ms window short time Fourier transform (STFT), with 20 ms overlap. Fig. 2 shows three exemplar spectrograms indicative of emotional ad content. Note greater densities of high frequencies in high asl ads, and such intense scenes are often characterized by sharp frequency changes.

4.2.3 CNN Training. We use the Caffe [13] deep learning framework for fine-tuning *Places205*, with a momentum of 0.9, weight decay of 0.0005, and a base learning rate of 0.0001 reduced by $\frac{1}{10}^{th}$ every 20000 iterations. We totally train four binary classification AAN networks this way to recognize high and low asl/val from audio/visual features. To fine-tune *Places205*, we use only the top and bottom 1/3rd LIRIS-ACCEDE videos in terms of asl and val rankings under the assumption that descriptors learned for the extreme-rated clips will effectively represent affective concepts. 4096-dimensional *fc7* outputs extracted from the four AAN networks are used for ad AR.

4.3 AR with audio-visual features

We will mainly compare our AAN network based AR framework against the algorithm of Hanjalic and Xu [9] in this work. Even after a decade since it was proposed, this algorithm remains one of the most popular AR baselines as noted from recent works such as [1, 14]. In [9], asl and val are modeled via low-level descriptors describing motion activity, colorfulness, shot change frequency, voice pitch and sound energy in the scene. These hand-crafted features are interpretable, and employed to estimate time-continuous asl and val levels conveyed by the scene. Table 3 summarizes the audio-visual features used in our AR experiments, and the proportion of positive audio/video frames for val and asl in our ad dataset.

4.4 Experiments and Results

We first provide a brief description of the classifiers used and settings employed for our binary AR experiments, where the objective

Table 3: Extracted features and +ve class proportions (in %) for the audio and visual modalities.

Attribute	Valence/Arousal		
	Audio	Video	aud+vid (A+V)
AAN Features	4096D AAN FC7 features obtained with 10s SGs.	4096D AAN FC7 features by extracted from keyframes sampled every 3 seconds.	8192D AAN FC7 features with SGs + keyframes over 10s intervals.
Hanjalic [9] Features	Per-second sound energy and pitch statistics [9].	Per-second shot change frequency and motion statistics [9].	Concatenation of audio-visual features.
val/asl +ve class prop (%)	43.8/51.9	43.4/51.6	43.8/51.9

is to assign a binary (H/L) label for asl and val evoked by each ad, using the extracted *fc7*/low-level audio visual features. Experimental results will be discussed thereafter.

Classifiers: We employed the Linear Discriminant Analysis (LDA), linear SVM (LSVM) and Radial Basis SVM (RSVM) classifiers for AR. LDA and LSVM attempt to separate H/L labeled training data with a hyperplane, while RSVM is a non-linear classifier which separates H and L classes, linearly inseparable in the input space, via transformation onto a high-dimensional feature space.

In addition to the above *single-task learning* methods which do not exploit the underlying structure of the input data, we also explored the use of *multi-task learning* (MTL) for AR. When posed with the learning of multiple *related* tasks, MTL seeks to jointly learn a set of task-specific classifiers on modeling task relationships, which is highly beneficial when learning with few examples. Among the MTL methods available as part of the MALSAR package [29], we employed the sparse graph-regularized MTL (SR-MTL) where *a-priori* knowledge regarding task-relatedness is modeled in the form of a graph R . Given tasks $t = 1..T$, with X_t denoting training data for task t and Y_t their labels, SR-MTL jointly learns a weight matrix $W = [W_1..W_T]$ such that the objective function $\sum_{t=1}^T \|W_t^T X_t - Y_t\|_F^2 + \alpha \|WR\|_F^2 + \beta \|W\|_1 + \gamma \|W\|_F^2$ is minimized. Here, α, β, γ are regularization parameters, while $\|\cdot\|_F$ and $\|\cdot\|_1$ denote the matrix Frobenius (ℓ_2) and ℓ_1 -norms respectively.

MTL is particularly suited for dimensional AR, and one can expect similarities in terms of audio-visual content among high val or high asl ads. We exploit underlying similarities by modeling each asl-val quadrant as a task (*i.e.*, all H asl, H val ads will have identical task labels). Also, quadrants with same asl/val labels are deemed as related tasks, while those with dissimilar labels are considered unrelated. The graph R guides the learning of W_t 's, as shown in the three examples in Fig.3, where SR-MTL is fed with the specified features computed over the final 30 s of all ads. Darker shades denote salient MTL weights. Shot change frequency is found to be a key predictor of asl in [9], and one can notice salient weights for H asl H val ads in particular. The attributable reason is that our H asl H val ads involve frequent shot changes to maintain emotional intensity, while the mood of our H asl L val ads² is strongly influenced by semantics. Likewise, pitch amplitude is deemed as a key val predictor, and salient weights can be consistently seen over the 30 s temporal window for HV ads. Finally, more salient weights for H val ads with the motion activity feature implies that our positive val ads involve accentuated motion.

²which depict topics like drug and alcohol abuse, and overspeeding.




Figure 3: Learned MTL weights for the four quadrants (tasks) when fed with the specified low-level features computed over the final 30 s of the 100 ads.

Metrics and Experimental Settings: We used the F1-score (F1), defined as the harmonic mean of precision and recall as our performance metric, given our unbalanced dataset (Table 3). Apart from unimodal (audio (A) or visual (V)) fc7 features, we also employed feature fusion and probabilistic decision fusion of the unimodal outputs. Feature fusion (A+V) involved concatenation of fc7 A and V features over 10 s windows (see Table 3), while the W_{est} technique [15] is employed for decision fusion (DF). In DF, the test label is computed as $\sum_{i=1}^2 \alpha_i^* t_i p_i$, where i indexes the A,V modalities, p_i 's denote posterior A,V classifier probabilities and $\{\alpha_i^*\}$ are the optimal weights maximizing test F1-score, and determined via a 2D grid search. If F_i denotes the training F1-score for the i^{th} modality, then $t_i = \alpha_i F_i / \sum_{i=1}^2 \alpha_i F_i$ for given α_i .

As the Hanjalic (Han) algorithm [9] uses audio plus visual features to model asl and val, we only consider (feature and decision) fusion performance in this case. As we evaluate AR performance on a small dataset, we present AR results over 10 repetitions of 5-fold cross validation (CV). CV is typically used to overcome the *overfitting* problem on small datasets, and the optimal classifier parameters (including regularization parameters for MTL) are determined from the range $[10^{-3}, 10^3]$ via an inner five-fold CV on the training set. Finally, in order to examine the temporal variance in AR performance, we present F1-scores obtained over (a) all ad frames ('All'), (b) last three frames (L3) and (c) last frame (L)³.

4.4.1 Results Overview. Table 4 presents the asl, val F1-scores under the various settings. The highest F1 over all the considered temporal windows, achieved with the single-task and multi-task classifiers, and via feature/decision fusion are denoted in bold. Based on the observed results, we make the following claims.

Focusing on **unimodal fc7 features and single-task classifiers**, val (peak F1 = 0.79) is generally recognized better than asl (peak F1 = 0.68) and especially with video features. A and V fc7 features perform comparably for asl. Much higher asl, val recognition scores are achievable with the **MTL classifier** (F1 of 0.96 for val and 0.94 for asl) due to its ability to exploit the underlying similarities among audio and visual features among similarly labeled ads. MTL F1-scores are consistently higher with V features for both asl and val.

Concerning recognition with **single-task classifiers and fused fc7 features**, comparable or better F1 scores are achieved with multimodal approaches. In general, better recognition is achieved

via decision fusion as compared to feature fusion⁴. For val, the best fusion performance (0.75 with feature fusion and RSVM classifier) is superior compared to A-based (F1 = 0.66), but inferior compared to V-based (F1 = 0.79) recognition. Contrastingly for asl, fusion F1-score (0.75 with DF) considerably outperforms unimodal methods (0.68 with A, and 0.67 with V). Focusing on the **MTL classifier**, MTL F1-scores in the **A+V FC7 + MTL** condition are considerably higher than single-task F1-scores⁵ analogous to the unimodal case, even though the best F1-scores are still less than those achieved with video fc7 features + MTL.

Comparing A+V **fc7 vs Han** features, fc7 descriptors clearly outperform Han features with both single and multi-task approaches. The difference in performance is prominent for val, while comparable recognition is achieved with both features for asl. RSVM produces the best F1-scores for both asl and val among **single-task classifiers** with unimodal and multimodal approaches. However, the linear **MTL** model considerably outperforms all single-task methods with both fc7 and Han features. These observations suggest that while the H and L asl/val features for *all ads* are difficult to linearly classify *per se*, exploiting underlying similarities among **quadrant-specific ads** enables better linear separability.

Relatively small σ values are observed for the 'All' condition with the adopted five-fold CV procedure in Table 4, especially with fc7 features suggesting that the AR results are **minimally impacted** by **overfitting**. Examining **temporal windows** considered for AR, higher σ 's are observed for the L3 and L cases, which denote model performance on the terminal ad frames. Surprisingly, one can note a general degradation in asl recognition for the L3 and L conditions with A/V features, while val F1-scores are more consistent. Also, a sharp degradation in performance is noted with MTL for the L3 and L conditions. Three inferences can be made from the above observations, namely, (1) Greater heterogeneity in the ad content towards endings is highlighted by the large variance with fusion and MTL-based approaches; (2) Fusion models synthesized with Han features appear to be more prone to overfitting, given the generally larger σ values seen with the corresponding models; (3) That asl recognition is lower in the L3 and L conditions highlights the limitation of using a **single** asl/val label (as opposed to dynamic labeling) over time. Generally lower F1-scores achieved for asl with all methods suggests that asl is a more transient phenomenon as

³This equates to estimating the ad asl/val typically over the terminal 30/10 s, when one would expect the conveyed emotion to be strongest.

⁴The relative efficacy of feature or decision fusion depends on the specific problem and features on hand.

⁵We are not aware of MTL-based decision fusion methods.

Table 4: Ad AR from content analysis. F1 scores are presented in the form $\mu \pm \sigma$.

Method	Valence			Arousal		
	F1 (all)	F1 (L3)	F1 (L)	F1 (all)	F1 (L3)	F1 (L)
Audio FC7 + LDA	0.61±0.04	0.62±0.10	0.55±0.18	0.65±0.04	0.59±0.10	0.53±0.19
Audio FC7 + LSVM	0.60±0.04	0.60±0.09	0.55±0.19	0.63±0.04	0.57±0.09	0.50±0.18
Audio FC7 + RSVM	0.64±0.04	0.66±0.08	0.62±0.17	0.68±0.04	0.60±0.10	0.53±0.19
Video FC7 + LDA	0.69±0.02	0.79±0.08	0.77±0.13	0.63±0.03	0.58±0.10	0.57±0.18
Video FC7 + LSVM	0.69±0.02	0.74±0.08	0.70±0.15	0.62±0.02	0.57±0.09	0.52±0.17
Video FC7 + RSVM	0.72±0.02	0.79±0.07	0.74±0.15	0.67±0.02	0.62±0.10	0.58±0.19
Audio FC7 + MTL	0.85±0.02	0.83±0.10	0.78±0.20	0.78±0.03	0.62±0.14	0.45±0.16
Video FC7 + MTL	0.96±0.01	0.94±0.07	0.82±0.25	0.94±0.01	0.87±0.12	0.63±0.29
A+V FC7 + LDA	0.70±0.04	0.66±0.08	0.49±0.18	0.60±0.04	0.52±0.10	0.51±0.18
A+V FC7 + LSVM	0.71±0.04	0.66±0.07	0.49±0.19	0.56±0.04	0.49±0.10	0.47±0.19
A+V FC7 + RSVM	0.75±0.04	0.70±0.07	0.55±0.17	0.63±0.04	0.56±0.11	0.49±0.19
A+V Han + LDA	0.59±0.09	0.63±0.08	0.64±0.12	0.54±0.09	0.50±0.10	0.58±0.08
A+V Han + LSVM	0.62±0.09	0.62±0.10	0.65±0.11	0.55±0.10	0.51±0.11	0.57±0.09
A+V Han + RSVM	0.65±0.09	0.62±0.11	0.62±0.12	0.59±0.12	0.58±0.11	0.56±0.10
A+V FC7 LDA DF	0.60±0.04	0.66±0.04	0.70±0.19	0.59±0.02	0.60±0.07	0.57±0.15
A+V FC7 LSVM DF	0.65±0.02	0.66±0.04	0.65±0.08	0.60±0.04	0.63±0.10	0.53±0.13
A+V FC7 RSVM DF	0.72±0.04	0.70±0.04	0.70±0.12	0.69±0.06	0.75±0.07	0.70±0.07
A+V Han LDA DF	0.58±0.09	0.58±0.09	0.61±0.09	0.59±0.06	0.59±0.07	0.61±0.08
A+V Han LSVM DF	0.59±0.10	0.59±0.09	0.60±0.10	0.61±0.05	0.61±0.08	0.60±0.09
A+V Han RSVM DF	0.60±0.08	0.56±0.10	0.58±0.09	0.58±0.09	0.56±0.06	0.58±0.09
A+V FC7 + MTL	0.89±0.03	0.88±0.11	0.77±0.26	0.87±0.03	0.68±0.17	0.46±0.20
A+V Han + MTL	0.77±0.04	0.79±0.07	0.74±0.15	0.78±0.04	0.73±0.11	0.58±0.22

compared to val, and that coherency between val features and labels is sustainable over time.

4.4.2 *Discussion.* As ads are inherently emotional and have great influencing/monetizing capacity [10, 19], the ability to infer ad emotions and make optimal ad insertions within video streams would be highly advantageous for multimedia systems. Therefore, it is surprising that very few works [18, 27] have attempted to mine visual and emotional content in ads. In this regard, our work expressly sets out to model the emotion conveyed by 100 ads based on subjective human opinions and objective audio-visual features.

We carefully curated a small but diverse set of ads based on consensus among experts, and examined if those ads could coherently evoke emotions across viewers by acquiring asl and val ratings from 14 annotators. A good-to-excellent agreement on asl and val impressions is noted between the expert and annotator groups. Also, annotator ad ratings are found to be uniformly distributed over the asl-val plane, with only a weak negative correlation noted between asl and val ratings.

As the compiled ads are found to constitute a control affective dataset, we modeled the emotion conveyed by these ads in terms of audio-visual features. Specifically, we extracted fc7 layer outputs from the *AdAffectNet* CNNs fed with key frames and spectrograms for video and audio-based affect modeling. While CNNs have been previously used for video and audio-based AR [3, 7], the modeled scenes are only a few second long snippets. In contrast, we have explored the validity of CNN-based AR for full-length ads, some of which are over a minute long, in this work.

Obtained AR results confirm that the synthesized fc7 features are effective predictors of asl and val. They outperform the audio-visual

features proposed by Hanjalic and Xu [9] with both single and multi-task classifiers. In particular, while fc7 features are considerably better for val, Han features provide competitive performance for asl. Optimal AR is achieved with the MTL classifier, which is able to effectively exploit the underlying similarities among emotionally homogeneous ads in terms of audio visual content. Nevertheless, a significant drop in recognition performance is generally noted for the terminal ad portion with most methods, and especially for asl, implying that (a) asl is a more transient phenomenon as compared to val, and there is less coherence between the employed AV features and asl labels towards the ad endings, and (b) the use of a single asl/val over the entire ad duration may be inappropriate, and one may need to acquire time-varying labels for affective studies. The next section will evaluate whether the superior AR achieved with AAN fc7 features translates to optimized ad insertion in a computational advertising task via a user study.

5 COMPUTATIONAL ADVERTISING- USER STUDY

Given the superior AR achieved by our *AdAffectNet* CNN features, we hypothesized that this should in turn enable optimized selection and insertion of affective ads within streamed video content, as discussed in the CAVVA ad insertion framework [27]. Video-in-video advertising is complex, as it aims to strike a balance between (a) maximizing ad impact, and (b) minimally disrupting (or ideally, enhancing) viewing experience while watching a program video into which the ads are embedded. Also, while ad insertion strategies have modeled ad-video relevance in terms of low-level visual context [18] and high-level emotional context [27], their performance has not been compared against human context assessment.

Table 5: Summary of program video statistics.

Name	Scene length (s)	Manual Rating	
		Valence	Arousal
coh	127±46	0.08±1.18	1.53±0.58
ipoh	110±44	0.03±1.04	1.97±0.49
friends	119±69	1.08±0.37	2.15±0.65

To this end, we performed a user study to evaluate whether the ad insertion framework formulated in [27], which employs the affect estimation methodology of Hanjalic and Xu [9], would benefit from better affect prediction via our deep AAN features. Better estimation of the affect induced by the video content and candidate ads can enable optimized selection of ads and corresponding insertion points. Specifically, we compared video program sequences generated via the CAVVA framework by estimating arousal (asl) and valence (val) scores for the ads and video scenes via (a) the baseline method of Hanjalic and Xu [9], (b) our deep AAN model and (c) human annotators.

5.1 Dataset

For the user study, we chose 28 ads (from the 100 used in this work) and three program videos. The program videos were scenes from a television sitcom *Friends* (**friends**) and two movies *The Pursuit of Happyness* (**ipoh**) and *Children of Heaven* (**coh**), with predominantly social themes and situations invoking high-to-low val and asl. Summary statistics of the three program videos are presented in Table 5. Each program video was segmented into 8 scenes, and the average scene length was 118 seconds. We obtained val and asl scores for the video scenes and 28 ads using (a) normalized softmax class probabilities [5] output by our AAN model, with video and audio fc7 features respectively used for val and asl estimation (b) the baseline model (Han) [9] and (c) ratings from three experts (Manual). We then inserted ads into each program video based on method-specific affect scores with the optimization strategy described in [27], and obtained 9 unique **video program** sequences (mean length 19.4 min) comprising the inserted ads. Exactly 5 ads were inserted in each program video, and 21 of the 28 chosen ads were cumulatively inserted at least once onto the 9 video programs (upon being selected via any of the three methods), with an average insertion frequency of 2.14.

5.2 Experiment and Questionnaire Design

We recruited a total of 17 users (5 female, mean age 20.5 years) to evaluate the video program sequences. Each user saw one exemplar sequence corresponding to the three affect prediction strategies. We followed a randomized 3×3 Latin square design so that all nine video programs were covered with three users.

Our user evaluation was in two parts; In line with the twin goals underlying seamless ad insertion within streaming video, we evaluated whether the ad insertion strategy resulted in (a) increased brand recall, and (b) minimal disturbance and improved viewing experience. We performed recall evaluation by measuring the impact of the ad insertion strategy on *immediate* and *day-after* recall. These *objective* measures quantified the impact of ad insertion on short-term (immediate) and long-term (day-after) memory of viewers, on viewing the video programs. Specifically, we measured the

proportion of (i) inserted ads that were recalled correctly (*Correct* recall), (ii) inserted ads that were not recalled (*Forgotten*) and (iii) non-inserted ads incorrectly recalled as viewed, perhaps owing to their inherent salience (*Incorrect* recall). For those ads that were inserted into program sequences and were correctly recalled, we also assessed whether viewers perceived them to be contextually appropriate with respect to program content.

The viewer was provided with a key-frame visual from each of the 28 ads, as well as a response sheet for every video program sequence. In addition to the recall related questions, we asked viewers to indicate whether they perceived the correctly recalled ads as being inserted at an appropriate position in the video stream (*Good insertion*)⁶. All recall and insertion quality-related responses were acquired from users as binary values. We pooled responses from viewers after they had watched video sequences generated via deep AAN, Han and manual affective scores for analyses.

While increased ad recall reflects a key desired effect of a successful ad insertion strategy, ads that are out of sync with the video program flow may disrupt viewer experience. In some cases, this disruptiveness may indirectly contribute to the recall, but would adversely impact viewing experience. So, mere recall alone does not indicate optimal ad insertion, relevance of the ad to the program or an enhanced viewer experience. To address these issues, we defined a second set of *subjective* experience evaluation measures and asked users to provide ratings on a Likert scale of 0–4 for the following questions, with 4 implying *best* and 0 denoting *worst*:

1. Whether advertisements were uniformly distributed across the video program?
2. Whether the inserted ads blended well with the flow of the video program?
3. Whether the inserted ads had a content and mood similar to surrounding program?
4. What was the overall viewer experience while each video program?

Each participant filled the recall and experience-related questionnaires (provided in supplementary material) after watching each video program. They also filled in the day-after recall questionnaire, a day after completing the experiment.

5.3 Results and Discussion

We evaluate the effectiveness of affective scores obtained from (a) our AAN-fc7 features, (b) Han features [9] employed in CAVVA [27], and (c) human assessments, for modeling contextual relevance based on user recall and experience responses.

Fig. 4 depicts user recall and experience-related measures as obtained with the three affect measurement approaches. Focusing on recall, although ad insertion via Han method [9] results in higher immediate and day-after ad recall, lower incorrect recall and forgottenness ($p < 0.05$ in all cases), video programs generated with AAN-based affective scores are found to maximize user experience ($p < 0.05$ for insertion uniformity and ad relevance, and $p < 0.1$ for non-disruptive ad insertion). The significance of these effects was assessed by comparing the proportion mean distributions for each question and method, via independent t -tests. Ads placed via the AAN model and correctly recalled by users were ‘well inserted’

⁶This was one way of inferring if the ad placements facilitated their recall.




Figure 4: Summary of user study results in terms of recall and user experience-related measures.

(difference with respect to manual scores significant at $p < 0.05$) based on responses compiled immediately upon viewing.

Notably, ad insertions via AAN-based affective scores were opined to be (i) ‘uniformly distributed’ across the streamed video, and (ii) ‘most relevant’ in terms of emotional context with the video (Fig.4). These observations imply that our AAN model is more accurate than Han in capturing the mood of the ads and video scenes. In contrast, while the Han method achieves the best recall from viewers, it also scores the least with respect to insertion-point distribution and relevance, implying an adverse impact on viewing experience.

To examine *how affective attributes influenced ad recall*, we correlated the ad recall measures with their (manually assigned) mean val, asl ratings. A meaningful relationship was noted between *estimated ad valence* and the *forgottenness rate* (Pearson $\rho = 0.45, p < 0.05$), indicating that positive val ads tend to be forgotten more easily. This observation agrees with the prior findings of Rimmel *et al.* [20], who discovered that recall performance was maximum for negative valence images in a memory study.

Surprisingly, ad insertions based on manual affective scores resulted in lowest recall and highest forgottenness among the three methods, while performing second best with respect to experience measures. This can be partly attributed to the higher val ratings observed for the selected ads based on manual scores ($\mu_{val} = 0.6$ for 12 unique inserted ads) as compared to ads selected based on AAN-based val estimates ($\mu_{val} = 0.3$ over 11 unique inserted ads). Ads selected based on Han val estimates had the lowest mean val ($\mu_{val} = 0.23$ over 12 unique inserted ads), suggesting that more low val ads were selected via the Han approach, resulting in least forgottenness. Nevertheless, viewers forgot nearly half the ads immediately and most ads a day later with all the considered methods. This reveals the scope for improving the ad-placement strategy by placing specific emphasis on ad retention.

To examine the *relationship between affective scores* estimated by the three methods *and the inserted ads*, we first examined if there was any relationship between manual ratings and computationally estimated scores. We found a significant correlation between Han-predicted and manual asl scores (Pearson $\rho = 0.4, p < 0.05$), but only a weakly significant correlation between manual ratings and AAN-based asl scores (Pearson $\rho = 0.24, p = 0.09$) on considering the 24 program video scenes and 28 ads (52 scores in total). Conversely, manual val ratings correlated significantly with our AAN model (Pearson $\rho = 0.45, p < 0.001$), but only weakly with Han estimates (Pearson $\rho = 0.25, p = 0.08$).

The CAVVA optimization framework [27] has two components—one for selection of ad insertion points into the program video, and

another for selecting the set of ads to be inserted. Asl scores only play a role in the choice of insertion points, whereas valence scores influence both components. Our results suggest that accurate val prediction, as accomplished by our AAN model, plays a critical role in enhancing the subjective user experience. Although this improved experience seems to come at the expense of ad recall, we note that the Han method results in a disruptive experience despite high recall, and hence solely emphasizing on recall may not necessarily lead to the optimal ad placement strategy.

6 CONCLUSIONS AND FUTURE WORK

This work discusses affect prediction from ads, and the utility of better ad affect estimation is demonstrated via a computational advertising application. A curative set of 100 diverse ads is compiled based on expert consensus, and its effectiveness as an affective dataset is examined based on ratings acquired from 14 raters. Dataset suitability is confirmed by (1) excellent agreement between the expert and annotator groups, and (2) uniform distribution of the asl and val ratings with minimal correlation between them.

AAN-based audio-visual features are then proposed for encoding ad affect, and are found to significantly outperform features proposed by Hanjalic and Xu [9] for val recognition. Best results with the AAN features are achieved with the MTL classifier, which effectively exploits the underlying audio-visual similarities among emotionally homogeneous ads. Finally, a study involving 17 users confirms that better modeling of ad emotions facilitates insertion of contextually relevant ads onto a streamed video. Specifically, the proposed AAN model is able to estimate ad valence better than the baseline [9], resulting in enhanced viewing experience.

Future work will focus on the design of other informative (*e.g.*, recurrent neural network-based) multimedia features for modeling affect. Another interesting line of inquiry is affect prediction via user physiological measurements similar to [1, 14]. Finally, efficient approaches need to be designed for ad insertion within streamed video, so as to maximize both ad recall and viewing experience.

ACKNOWLEDGMENTS

This research is supported by the National Research Foundation, Prime Ministers Office, Singapore under its International Research Centre in Singapore Funding Initiative.

REFERENCES

- [1] M. K. Abadi, R. Subramanian, S. M. Kia, P. Avesani, I. Patras, and N. Sebe. 2015. DECAF: MEG-Based Multimodal Database for Decoding Affective Physiological Responses. *IEEE Trans. Affective Computing* 6, 3 (2015), 209–222.

- [2] Yoann Baveye. 2015. *Automatic prediction of emotions induced by movies*. Theses. Ecole Centrale de Lyon.
- [3] Yoann Baveye, Emmanuel Dellandrea, Christel Chamaret, and Liming Chen. 2015. LIRIS-ACCEDE: A video database for affective content analysis. *IEEE Trans. Affective Computing* 6, 1 (2015), 43–55.
- [4] Yoav Benjamini and Yoel Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Royal Stat. Soc. Series B (Methodological)* 57, 1 (1995), 289–300.
- [5] Christopher M. Bishop. 2013. *Pattern Recognition and Machine Learning*. Vol. 53. Springer.
- [6] V Carter Broach, Thomas J Page, and R Dale Wilson. 1995. Television Programming and Its Influence on Viewers' Perceptions of Commercials: The Role of Program Arousal and Pleasantness. *Journal of Advertising* 24, 4 (1995), 45–54. <http://www.jstor.org/stable/4188988> \$delimiter="026E30F\$ \$nhttp://about.jstor.org/terms
- [7] Yin Fan, Xiangju Lu, Dian Li, and Yuanliu Liu. 2016. Video-based Emotion Recognition Using CNN-RNN and C3D Hybrid Networks. In *International Conference on Multimodal Interaction*. 445–450.
- [8] M. K. Greenwald, E. W. Cook, and P. J. Lang. 1989. Affective judgement and psychophysiological response: dimensional covariation in the evaluation of pictorial stimuli. *Journal of Psychophysiology* 3 (1989), 51–64.
- [9] Alan Hanjalic and Li-Quan Xu. 2005. Affective Video Content Representation. *IEEE Trans. Multimedia* 7, 1 (2005), 143–154.
- [10] Morris B Holbrook, Rajeev Batra, and Rajeev Batra. 1987. Assessing the Role of Emotions as Mediators of Consumer Responses to Advertising. *Journal of Consumer Research* 14, 3 (1987), 404–420. <https://doi.org/10.1002/job.305>
- [11] Morris B Holbrook and John O Shaughnessy. 1984. The role of emotion in advertising. *Psychology & Marketing* 1, 2 (1984), 45–64.
- [12] Zhengwei Huang, Ming Dong, Qirong Mao, and Yongzhao Zhan. 2014. Speech Emotion Recognition Using CNN. In *ACM Multimedia*. 801–804.
- [13] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. CAFFE: Convolutional Architecture for Fast Feature Embedding. In *ACM Int'l Conference on Multimedia*. 675–678.
- [14] Sander Koelstra, Christian Muehl, 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. Affective Computing* 3, 1 (2012), 18–31.
- [15] Sander Koelstra and Ioannis Patras. 2013. Fusion of facial expressions and EEG for implicit affective tagging. *Image and Vision Computing* 31, 2 (2013), 164–174.
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Neural Information Processing Systems*. 1097–1105.
- [17] Peter J. Lang, Margaret M. Bradley, and B. N. Cuthbert. 2008. *International affective picture system (IAPS): Affective ratings of pictures and instruction manual*. Technical Report A-8. The Center for Research in Psychophysiology, University of Florida, Gainesville, FL.
- [18] Tao Mei, Xian-Sheng Hua, Linjun Yang, and Shipeng Li. 2007. VideoSense: Towards Effective Online Video Advertising. In *ACM Int'l Conference on Multimedia*. 1075–1084.
- [19] Michel Tuan Pham, Maggie Geuens, and Patrick De Pelsmacker. 2013. The influence of ad-evoked feelings on brand evaluations: Empirical generalizations from consumer responses to more than 1000 {TV} commercials. *International Journal of Research in Marketing* 30, 4 (2013), 383 – 394.
- [20] U Rimmele, L Davachi, R Petrov, S Dougal, and E.A Phelps. 2011. Emotion enhances the subjective feeling of remembering, despite lower accuracy for contextual details. *Emotion* 11 (2011), 553–562.
- [21] James Russell. 1980. A circumplex model of affect. (1980), 1161–1178 pages.
- [22] Ramanathan Subramanian, Divya Shankar, Nicu Sebe, and David Melcher. 2014. Emotion modulates eye movement patterns and subsequent memory for the gist and details of movie scenes. *Journal of vision* 14, 3 (2014), 1–18.
- [23] Ramanathan Subramanian, Julia Wache, Mojtaba Abadi, Radu Vieriu, Stefan Winkler, and Nicu Sebe. 2016. Ascertain: Emotion and personality recognition using commercial sensors. *IEEE Transactions on Affective Computing* (2016).
- [24] Ramanathan Subramanian, Victoria Yanulevskaya, and Nicu Sebe. 2011. Can computers learn from humans to see better?: inferring scene semantics from viewers' eye movements. In *ACM international conference on Multimedia*. ACM, 33–42.
- [25] Hee Lin Wang and Loong-Fah Cheong. 2006. Affective understanding in film. *IEEE Trans. Circ. Syst. V. Tech.* 16, 6 (2006), 689–704.
- [26] Karthik Yadati. 2013. *Online Multimedia Advertising*. Master's thesis. National University of Singapore, Singapore.
- [27] Karthik Yadati, Harish Katti, and Mohan Kankanhalli. 2014. CAVVA: Computational affective video-in-video advertising. *IEEE Trans. Multimedia* 16, 1 (2014), 15–23.
- [28] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. 2014. Learning deep features for scene recognition using places database. In *Advances in neural information processing systems*. 487–495.
- [29] J. Zhou, J. Chen, and J. Ye. 2011. *MALSAR: Multi-Task Learning via Structural Regularization*. Arizona State University. <http://www.public.asu.edu/~jye02/Software/MALSAR>