

Investigating the Influence of Personal Memories on Video-Induced Emotions

Dudzik, B.; Hung, H.S.; Neerincx, M.A.; Broekens, Joost

DOI

[10.1145/3340631.3394842](https://doi.org/10.1145/3340631.3394842)

Publication date

2020

Document Version

Final published version

Published in

Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization

Citation (APA)

Dudzik, B., Hung, H. S., Neerincx, M. A., & Broekens, J. (2020). Investigating the Influence of Personal Memories on Video-Induced Emotions. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization* (pp. 53-61) <https://doi.org/10.1145/3340631.3394842>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Investigating the Influence of Personal Memories on Video-Induced Emotions

Bernd Dudzik*

Delft University of Technology
Delft, South Holland, The Netherlands
B.J.W.Dudzik@tudelft.nl

Mark Neerincx

Delft University of Technology
Delft, South Holland, The Netherlands
M.A.Neerincx@tudelft.nl

Hayley Hung

Delft University of Technology
Delft, South Holland, The Netherlands
H.Hung@tudelft.nl

Joost Broekens

Leiden University
Leiden, South Holland, The Netherlands
D.J.Broekens@tudelft.nl

ABSTRACT

This paper contributes to the automatic estimation of the subjective emotional experience that audio-visual media content induces in individual viewers, e.g. to support affect-based recommendations. Making accurate predictions of these responses is a challenging task because of their highly person-dependent and situation-specific nature. Findings from psychology indicate that an important driver for the emotional impact of media is the triggering of personal memories in observers. However, existing research on automated predictions focuses on the isolated analysis of audiovisual content, ignoring such contextual influences. In a series of empirical investigations, we (1) quantify the impact of associated personal memories on viewers' emotional responses to music videos in-the-wild and (2) assess the potential value of information about triggered memories for personalizing automatic predictions in this setting. Our findings indicate that the occurrence of memories intensifies emotional responses to videos. Moreover, information about viewers' memory response explains more variation in video-induced emotions than either the identity of videos or relevant viewer-characteristics (e.g. personality or mood). We discuss the implications of these results for existing approaches to automated predictions and describe ways for progress towards developing memory-sensitive alternatives.

CCS CONCEPTS

• **Human-centered computing** → **User models.**

KEYWORDS

Video-Induced Emotions; Video Affective Content Analysis

ACM Reference Format:

Bernd Dudzik, Hayley Hung, Mark Neerincx, and Joost Broekens. 2020. Investigating the Influence of Personal Memories on Video-Induced Emotions. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20)*, July 14–17, 2020, Genoa, Italy. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3340631.3394842>

*This is the corresponding author

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

UMAP '20, July 14–17, 2020, Genoa, Italy

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6861-2/20/07...\$15.00

<https://doi.org/10.1145/3340631.3394842>

1 INTRODUCTION

Research on *Video Affective Content Analysis (VACA)* strives to enable technologies to automatically estimate the emotional responses videos induce in their viewers [2], e.g. to support affect-based recommendations [32, 36]. A fundamental challenge for this undertaking is that emotional experiences are highly subjective, expressing a dynamic relationship between individuals' ongoing personal needs and the perceived ability of their current situation to meet them [27]. Therefore, how people experience media stimuli depends on who they are and under what circumstances they encounter the stimuli (see, e.g. the findings of age-related differences in [19] or the existence of mood-effects [31]).

Throughout the existence of VACA research, the majority of technological investigations have avoided dealing with the issue of subjectivity. Instead, it has focused on the de-contextualized analysis of videos' audiovisual content to estimate emotional responses elicited in a majority of viewers [2]. As such, the emotional impact estimated by these technologies is not the subjective experience of a particular person viewing a video, but rather the expected response across a population of viewers, and independent of their viewing situation. However, without the capacity to reflect variations within and across individuals' impressions of the same video, the practical value of this approach for applications seems limited.

Only recently, research has started to openly address the issue of modelling individual affective experiences by exploring the usefulness of viewer- and situation-specific information in predictions, e.g. personality traits and cultural background [30], or the social setting in which viewing takes place [37]. Nevertheless, systematic research into such context-sensitive predictions remains scarce [2], and many essential drivers for human affective experiences remain unexplored and unaccounted for in computational models of video-induced emotions.

One significant influence that research has not yet touched upon is the recollection of personal memories [10]. Not only can memories about one's past have a significant emotional impact in their own right, but auditory and visual material can readily trigger them in an audience [18]. Moreover, many patterns of media creation and consumption revolve specifically around this ability to serve as cues for emotionally significant memories, e.g. the taking of photos as mementoes for the future, or the listening to music from a specific period in ones' past for the sake of reminiscence and nostalgia.

In this article, we present several empirical investigations to quantify the importance of accounting for viewers' memories when predicting video-induced emotions. In particular, we present the following contributions:

- (1) the collection of a dataset about viewers' recollection processes in emotional responses to music videos,
- (2) a series of analyses to quantify the influence of personal memories on video-induced emotions,
- (3) and a comparison of the impact on affective response prediction accuracy between information about *a*) the memories triggered in viewers, *b*) the eliciting video, and *c*) relevant viewer-characteristics (demographics, personality, and mood).

While the analyses that we present here focus on self-reported data, we discuss the relevance of our findings for existing computational approaches to predict video-induced emotions. In particular, we outline opportunities that future technological research could explore to account for the influence of personal memories in predictions made by automated systems.

2 BACKGROUND AND RELATED WORK

2.1 Affective Memory Processes and Media-Induced Emotions

While individuals can intentionally recall moments from their personal history, stimuli in the external environment can also remind them of these moments involuntarily [4]. Audiovisual media appear to be potent triggers for personal memories in observers [3, 23], and these can evoke strong emotions upon recollection [18]. The mechanisms that determine when external stimuli evoke memories from the past in a person are not entirely understood. Two discovered conditions are (1) the existence of some form of semantic or perceptual association between the triggering stimulus and the recollected content [22], and (2) a state of low attentional engagement (see, e.g. [34]). Once recollection takes place, memories can elicit strong emotional responses, as is demonstrated by their frequent use in emotion induction procedures [26]. Moreover, empirical evidence points to a direct connection between the affect associated with the memories triggered by a media stimulus – i.e. how one feels about the memory – to that stimulus' emotional impact [1]. Together these findings indicate that recollection of personal memories can drive viewers' experiences.

2.2 Video-Affective Content Analysis

Video-Affective Content Analysis is an approach to predicting the emotional response of viewers to audiovisual media content. Research on this topic comprises roughly three components: the type of emotional response that forms the target of the prediction, the granularity at which responses' to video are to be predicted, and the sources of information that form the basis for predictions.

The types of emotional responses of primary concern in VACA research take two different forms [2]. The first is the expected emotion for a video (What feelings does it evoke in the majority of viewers?). The second, is the induced emotion (What feelings does a video evoke in a particular viewer?). Regarding the granularity of predictions, research efforts comprise of those undertaking a global analysis – which attempts to predict the emotional response for an

entire video –, or those conducting a continuous analysis – which consists of predictions for smaller windows (potentially down to the frame-level) [2]. Finally, existing VACA approaches are distinguishable by the information that they use to make their predictions [35]. *Direct VACA* focuses on the analysis of the audiovisual signals comprising the content of a video. In contrast, *Implicit VACA* relies on the automatic analysis of viewers' behaviour during exposure to a video clip, e.g. facial expressions (e.g. [24]) or physiological response (e.g. [21]).

Overall, the majority of existing research on predicting video-induced emotions focuses on expected emotions at a global level, using *Direct VACA* [2, 35]. Because this task focuses on assigning a single expected emotion to a video, most of the existing VACA research (either implicitly or explicitly) assumes homogeneous emotional responses across viewers. Video corpora used in VACA research are specially designed to adhere to this low-variation assumption by filtering out videos that elicit diverse responses (e.g. [21, 24]). Researchers in the VACA community are aware that technologies built on this assumption cannot reflect the natural variation in viewers' responses, and that doing so requires increasingly context-sensitive approaches [2, 31, 35]. However, explorations in this regard are only recently gaining traction. Initial efforts have touched on the value of information about personality and cultural background for personalizing predictions [15, 30], or the role of viewers' overall affective mood [31]. In principle, the combination of behavioural and physiological signals of viewers in combination with features describing the video content (e.g. [32]) can also be considered as a form of contextualized prediction. However, this response information is only available when viewers' are already exposed to a video, potentially ruling out some primary use-cases for undertaking VACA in the first place, e.g. intelligent recommendation.

In summary, few prior works investigate contextual features for predictions of individual viewers' emotional responses. Existing explorations revolve primarily around addressing viewer-specific differences, with some consensus among scholars for the importance of basic demographic factors, personality traits, and mood as features. Incorporating information about viewers' cognitive processes in general, and the effects of recollecting associated personal memories, in particular, has not been explored so far.

2.3 Representations of Emotions

There exist a wide variety of schemes for describing and classifying emotional experiences in psychology, many of which have been used in modelling video-induced emotions [2]. Broadly, these consist of two distinct groups: categorical and dimensional schemes. Categorical schemes typically build on psychological theories that assume a set of discrete emotional states, often with unique physiological response components attached to them (e.g. facial expressions, or the patterns of activity in the autonomous nervous system [12]). In contrast, dimensional schemes describe affect as points in continuous space with a set of orthogonal dimensions (e.g. [29]). Each dimension is supposed to capture an essential quality for discriminating between affective states.

Overall, there is no consensus about the merit of any particular type of scheme among researchers engaging in modelling video-induced emotions. However, categorical approaches may lack the necessary nuance for describing affective media experiences [2].

For this reason, we have used the dimensional PAD-framework [25] to measure video-induced emotions in the dataset collected for this research (see Section 3 below). It is prominent in psychological studies and also extensively used in VACA research (e.g. [21]). PAD characterizes affective experiences along three orthogonal axes, each with a positive and negative polarity: *pleasure* (*P*) (valence, is an experience positive or negative, enjoyable or unpleasant?), *arousal* (*A*) (Does it involve a high or low degree of bodily activation or alertness?), and *dominance* (*D*) (To what degree am I in control of the experienced situation?).

3 DATA COLLECTION

In this section, we provide a detailed description of the relevant elements of a dataset that we have collected to investigate the influence of recollection processes in emotional responses to videos.

3.1 Selected Video Stimuli

We used music video segments that have previously been evaluated for their affective impact in the *DEAP* dataset [21]. We opted for this choice, because (1) prior research demonstrated the potency of musical material for triggering personal recollections [18], and (2) each stimulus in this list was rated by multiple viewers using the PAD-framework. This second property provided us with information about the expected distribution of emotional responses from viewers, which we used for balancing the distribution of selected videos in the design of our study.

From the total of 150 video segments for which the validation study for the *DEAP* corpus collected ratings, we selected a subset of 42 videos based on their variation. In particular, we selected an equal amount of stimuli per affective dimension that possess either a high- or a low- degree of variation in their emotional responses. For example, we balanced videos where different viewers show very similar pleasure-responses with other videos where this is not the case. We opted for this scheme because we hypothesized that situation- and viewer-specific influences are more likely present in high variation stimuli. Consequently, these differences in variation might also reflect variation in the occurrence and influence of personal memories. It is important to note that for the *DEAP* corpus, only videos with a narrow distribution and pronounced average emotional responses were kept in their experiments [21].

3.2 Participants

We recruited 300 individuals via the crowd-sourcing platform *Amazon Mechanical Turk* and compensated each for their participation with 6 *USD*. We did not constrain our recruitment efforts to a geographic region or particular nationalities. However, we requested that individuals have command of the English language and that they participate in an environment that allows them to pay attention for the entire duration of the task. Additionally, we enforced restrictions on the age of participants in our online survey so that they were between 25 and 46 years of age. We implemented this to ensure that the release dates of music videos used in our study fall into a period in viewers' lifespan that cognitive psychology knows as the *reminiscence bump* [6]. Empirical findings indicate that memories made during this period of early adulthood (i.e. between age 15 to 30) remain particularly accessible throughout people's lives.

We expect this measure to maximize the possibility of viewers possessing accessible memories associated with our videos.

All subjects that we recruited gave their informed consent before entering the study itself, both regarding the tasks involved and the usage of their data. For a detailed overview about the demographics of the participating crowd-workers see the relevant section in *Table 1*. Because of the small number of participants in our dataset that are not from either the United States of America or The Republic of India (i.e. the *OTHER*-category in *Table 1*), we exclude all data belonging to these participants from the analysis reported throughout this article. This filtering reduces the total amount of individuals for which data is available to 288 unique viewers.

3.3 Procedure and Apparatus

We developed an online application which crowd-workers could access through their web browser. It guided them through the entire data collection study, presented them with survey elements to acquire self-reports, and created audiovisual recordings using their webcams during their exposure to the videos.

After providing their informed consent for participation in the study, crowd-workers filled out an initial survey with background information about themselves (*Viewer-specific Self-Reports*). Then each person was provided with a random selection of 7 music videos from within our pool of 42 candidates. Acquisition of data about viewers' responses to each video had the following structure: first, they were watching the video, during which we recorded their upper-body behaviour with their device's webcam. Immediately after the video finished playing, the application collected self-reports about their experience (*Response-specific Self-reports*). This required participants to start by rating the emotional impact that the video had on them (its *Induced Emotion*). Then they were requested to report whether the video had evoked memories in them (i.e. whether any *Recollection* had occurred at all during exposure). If this was the case, we required participants to fill in a detailed survey to describe each of these memories in more detail, including the feelings that they associated with the memory (i.e. *Memory-Associated Affect*). After responding to all selected videos in this way, they had completed the data collection, and we debriefed them.

3.4 Self-Report Measures

Here we provide a detailed description of all self-report measures that we have collected. Descriptive statistics can be found in *Table 1* for viewer-specific, and in *Table 2* for response-specific measures.

3.4.1 Viewer-specific Measures.

Demographics: In previous studies, demographic information significantly accounted for variation in viewers' emotional responses (e.g. age [19]). We capture self-reports of the following basic features: participants' age in years, their gender, and their nationality.

Personality: We collected data about our participants' personality traits in terms of the HEXACO scheme. It is a framework that aims to account for a wide variety of individual differences across peoples' behaviours by differentiating between them with a set of stable personality traits. In the HEXACO scheme these traits are defined by six orthogonal dimensions: (1) Honesty-Humility (H), (2) Emotionality (E), (3) eXtraversion (X), (4) Agreeableness (A), (5) Conscientiousness (C), and (6) Openness to experience (O). In

Table 1: Overview of Viewer-Specific Measures

DEMOGRAPHICS						
	USA		India		Other	
	N_f	N_m	N_f	N_m	N_f	N_m
Gender	125	115	11	37	2	10
	M	SD	M	SD	M	SD
Age	33.57	6.01	30.54	4.91	32.83	4.91
PERSONALITY						
	USA		India		Other	
	M	SD	M	SD	M	SD
Honesty	2.78	0.72	1.95	0.55	2.56	0.55
Emotional	1.95	0.8	1.85	0.62	2.04	0.62
Extraver.	2.54	0.78	2.32	0.67	2.56	0.67
Agreeabl.	2.08	0.66	2.05	0.68	1.75	0.68
Conscien.	2.68	0.71	2.36	0.62	2.46	0.62
Openness	2.78	0.7	2.69	0.53	2.67	0.53
MOOD						
	USA		India		Other	
	M	SD	M	SD	M	SD
Pleasure	0.42	0.4	0.43	0.44	0.36	0.44
Arousal	-0.14	0.77	0.05	0.82	-0.38	0.82
Dominance	0.34	0.49	0.51	0.47	0.46	0.47

Measures taken once per viewer: $N = 300$

Table 2: Overview of Response-specific Measures

INDUCED EMOTION						
	USA		India		Other	
	M	SD	M	SD	M	SD
Pleasure	0.18	0.52	0.35	0.54	0.22	0.51
Arousal	-0.17	0.78	0.32	0.73	-0.23	0.78
Dominance	0.12	0.58	0.3	0.64	0.23	0.58
MEMORY ASSOCIATED-AFFECT						
	USA		India		Other	
	M	SD	M	SD	M	SD
Pleasure	0.31	0.53	0.48	0.55	0.35	0.61
Arousal	0.01	0.79	0.38	0.71	0.07	0.76
Dominance	0.28	0.57	0.42	0.62	0.36	0.57

Measures taken once per response: $N = 2098$

our study, we assessed viewers' HEXACO scores using the Brief HEXACO Inventory (BHI) [8]. Because its design is specifically aiming for brevity (it consists of only 24-items) without sacrificing validity, it is well suited for crowd-sourcing scenarios. Scores are continuous values in the range of $[1 - 5]$.

Mood: Before we exposed participants to any videos, they provided affective ratings for their overall mood for the same day. Findings show that mood has a significant influence on the emotions that videos induce in viewers [31]. Mood ratings in our corpus take the form of pleasure-, arousal- and dominance-ratings on a

continuous scale in the interval of $[-1, +1]$. We obtain them from participants with the *AffectButton* instrument – an interactive widget displaying an iconic facial expression which changes in response to mouse or touch interaction. Users can select the facial expression best fitting the affective judgment that they need to provide (see [5] for a detailed description and a validation study). We opted for this instrument because it allowed viewers without knowledge of the underlying psychological framework to provide quick and implicit PAD-ratings through choosing a face.

3.4.2 Response-specific Measures.

Induced Emotions: We also capture viewers' ratings for their emotional response to a video with the *AffectButton* instrument. Consequently, they are on a continuous scale for pleasure, arousal and dominance bounded by the $[-1, +1]$ interval. See Table 2 for relevant details about their distribution.

Memory-Associated Affect: Likewise, we measured the feelings that viewers associate with recollected memory content with the *AffectButton* instrument. Our dataset contains a total of 944 instances in which participants' recollected at least one memory in response to a music video. In principle, participants could report as many memories as were triggered in them by each video. However, only about 6% of recollections involved more than 2 memories (a total of 53 instances). In these cases, we decided to choose the most intense memory involved to represent their collective emotional meaning. This choice is motivated by empirical findings from psychology research, highlighting the dominance of emotionally intense parts of an event in retrospective summary judgments of emotional meaning (for example, [9]). To make this selection, we rank all memories involved in a given multi-memory recollection based on the following measure for the affective intensity (I) of their associated affect: $I = \sqrt{p^2 + ((a + 1)/2)^2 + d^2}$, where p , a , and d are the pleasure, arousal and dominant components of a particular rating. Note that this formula interprets negative arousal values as low-intensity experiences. This choice is motivated by the layout of the *AffectButton* instrument, which maps maximum negative arousal to neutral faces [5]. After calculation, we retain only the highest-scoring memory for further modelling activities. When we use *Memory-Associated Affect* in the remainder of this manuscript (including Table 2), we refer to the ratings selected in this way.

4 THE IMPACT OF MEMORY PROCESSES ON VIDEO-INDUCED EMOTIONS

In this section we present two empirical investigations to quantify the impact of personal memories on viewers' emotional responses during video-viewing scenarios in-the-wild.

4.1 Exp. 1: Video-Induced Emotions Differ when Personal Memories are Recollected

In this first experiment, we investigate whether there are differences in the emotional responses of viewers' to a video (i.e. their Induced Pleasure, Induced Arousal, or Induced Dominance), depending on whether it made them recollect personal memories or not. Moreover, we explore the degree to which these differences depend on the identity of eliciting videos.

Method and Approach: We use linear mixed regression models in our analysis to account for the repeated measures of responses from the same viewers in our data collection. A separate model was fitted for each affective dimension of viewers' induced emotions: *Induced Pleasure*, *Induced Arousal*, and *Induced Dominance*. The fixed effects included in these models are the *identity of the video (VID)* shown to viewers, the *Occurrence of Recollection (REC)*, as well as their two-way interaction term. *REC* is a binary variable denoting whether a viewer's response involved the recollection of a video or not. *VID*, on the other hand, is a factor with 42 levels, each denoting the identity of the particular music video that we showed to participants. Additionally, we specify viewers' identity as a random effect in the models, thereby accounting for the dependence among their responses due to repeated measures.

Results: We conducted an analysis of variance for the fixed effects in the separate models (Table 3). To account for the multiple comparisons between the same set of dependent variables with each of the three independent variables, we have applied Bonferroni corrections to all statistical tests. Results show that the occurrence of recollection has a significant effect on video-induced emotion across all dimensions. This finding indicates that there is a difference in video-induced emotions when viewers' responses also involved the recollection of personal memories. To investigate the direction of these differences, we compared the means of responses involving recollections to those without them. Results indicate that when recollections are present ratings for induced emotions are higher across all affective dimensions (induced pleasure: $M\Delta = +.16$, $t(2043.7) = 7.05$, $p = < .001$; induced arousal: $M\Delta = +.17$, $t(2033) = 5.09$, $p = < .001$; induced dominance: $M\Delta = +.19$, $t(2058.2) = 7.46$, $p = < .001$). Moreover, the absence of a significant interaction-effect reveals that the magnitudes of these differences are comparable across videos. Finally, we investigated the increase in explained variance contributed by each of the fixed effects in the model to gain insights into their relative explanatory power. For this purpose we use a measure for the explained variance of the fixed effects in linear mixed models – Marginal R^2 (R_m^2) [28]. Consequently, the measure of ΔR_m^2 in Table 3 captures the additional variance explained by a model that includes the predictors for a given fixed effect, compared to one that does not. A comparison of this metric across models for the different affective dimensions reveals that video identity is the effect explaining the highest amount of unique variance ($Avg\Delta R_m^2 = .13$). While the effects of occurring recollections are significant, they explain only a rather small amount of unique variation ($Avg\Delta R_m^2 = .017$).

4.2 Exp. 2: Memory-Associated Affect Predicts Video-Induced Emotions

Psychological findings indicate that the affect associated with the memory content triggered by media stimuli correlates with the emotions that these stimuli induce in them [1]. Here, we explore the existence of this relationship for those instances in our dataset, where viewers' have recollected memory content in response to music videos. In particular, we assess whether their feelings towards the memory content evoked by a particular video (i.e. *Memory-associated Affect*) are indicative of their emotional response to it.

Method and Approach: Like in the previous experiment we fit three separate linear mixed regression models for our analysis, each targeting one of the affective dimensions of viewers induced emotions (i.e. *Induced Pleasure*, *Induced Arousal*, and *Induced Dominance*). All models include fixed effects for (1) video identity *VID* (again a factor with 42 levels), and (2) the memory-associated affect *MA*, which consists of ratings for pleasure, arousal and dominance as predictors. Additionally, we include all two-way interactions between the predictors of the specified *VID* and *MA*-effects. All affective ratings for induced emotion and memory-associated affect are continuous numerical variables and constrained to the interval $[-1, 1]$. Before introducing them to the model, they are standardized by subtracting their mean and dividing by their standard deviation. We include participants' identity as a random effect to account for the repeated-measures design of our data collection.

Results: We conduct an analysis of variance for the fixed effects in the specified models, the results of which we present in Table 4). We applied Bonferroni correction when testing for significance to account for multiple comparisons between dependent and independent variables across models. Significant effects exist both for memory-associated affect and video identity across all models for induced emotion. Moreover, we find significant two-way interactions between video identity and memory-associated affect. This finding indicates video-specific variations in the strength of the relationship between the emotions induced in viewers and their feelings towards recollected memories. Finally, we compare the changes in uniquely explained variance for each of the fixed effects in the specified models. This reveals that memory-associated affect makes the strongest contribution ($Avg\Delta R_m^2 = .306$), going beyond the unique share of video identity ($Avg\Delta R_m^2 = .04$). Moreover, the interactions between these effects explain a relatively large share of additional variance in viewers' responses ($Avg\Delta R_m^2 = .1$).

4.3 Discussion

Our first experiment demonstrates that peoples' experience of a video significantly differs depending on whether it triggers personal memories or not. This finding points to a recollection-specific bias, causing videos that trigger memories to display induced emotions with heightened levels of induced pleasure, arousal and dominance. This finding confirms the results of existing psychological research involving music-evoked recollections (e.g. [18], and shows that such biases are also present for video material the unconstrained scenarios captured by our dataset. More generally, our results highlight the existence of systematic influences on the emotions induced by videos that one cannot feasibly attribute to their content alone but instead result from effects in the situation under which viewing takes place. Naturally, technologies for emotion prediction that solely rely on the analysis of the audiovisual signals comprising this content – as is the dominant approach in VACA – cannot account for this kind of influence.

In our second experiment, we have discovered a strong relationship between the affect that a viewer associates with a video-triggered memory and the triggering video's emotional impact on him/her. This effect explains an amount of variation that goes significantly above and beyond what one can reasonably attribute to video-specific influences alone, further underlining the

Table 3: Exp. 1 – Effect of Occurring Recollections on Video-Induced Emotion.

Effect	df_n	Induced Pleasure					Induced Arousal				Induced Dominance				$Avg\Delta R_m^2$
		df_d	F	p	ΔR_m^2		df_d	F	p	ΔR_m^2	df_d	F	p	ΔR_m^2	
<i>VID</i>	41	1937.48	10.27	<.001*	.16		1864.78	6.43	<.001*	.1	1878.43	8.74	<.001*	.13	.13
<i>REC</i>	1	1683.29	47.1	<.001*	.02		1977.8	30.96	<.001*	.01	1942.93	38.48	<.001*	.02	.017
<i>VID * REC</i>	41	1957.52	1.29	.106	.02		1880.93	1.28	.11	.02	1895.85	1.28	.115	.02	.02

* Result below Bonferroni corrected critical value for significance: $\alpha_{adj} = .017$

Table 4: Exp. 2 – Effect of Memory-associated Affect on Video-Induced Emotion.

Effect	df_n	Induced Pleasure					Induced Arousal				Induced Dominance				$Avg\Delta R_m^2$
		df_d	F	p	ΔR_m^2		df_d	F	p	ΔR_m^2	df_d	F	p	ΔR_m^2	
<i>VID</i>	41	877.07	2.64	<.001*	.05		865.28	1.36	.067	.03	885.22	1.86	.001*	.04	.04
<i>MA</i>	3	889.39	228.92	<.001*	.34		895.33	141.62	<.001*	.29	878.35	161.02	<.001*	.29	.306
<i>VID * MA</i>	123	738.16	1.4	.005*	.08		733.37	1.34	.014*	.11	738.99	1.63	<.001*	.11	.1

* Result below Bonferroni corrected critical value for significance: $\alpha_{adj} = .017$

importance of accounting for memories when modelling individuals' experiences. Moreover, this relationship varies in intensity across different videos. A possible explanation for this is that the memories evoked by some videos are more engaging as a target for viewers than their audiovisual content. Exposure to these videos could create conditions where viewers' attention is more likely to drift inwards, thereby increasing memories' emotional impact. Such mind-wandering phenomena can occur across different forms of media consumption, e.g. reading, and its emergence depends on individuals' availability of attentional resources [14]. Being able to identify videos that are less impacted by memory-associated affect automatically could be a valuable effort for computational research because it is likely that responses to these are more firmly grounded in a video's audiovisual content. Such a grounding matches the stimulus-centric modelling assumption of direct VACA, and consequently applying it to these videos might improve results.

Together, both studies demonstrate the scope and depth of the role that personal memories play in viewers' experiences. Moreover, they point towards the potential that information about individuals' recollection processes could hold as context for predictions of video-induced emotions in technological systems.

5 USING MEMORIES TO PERSONALIZE PREDICTIONS OF INDUCED EMOTIONS

We found that memory-associated affect strongly correlates with video-induced emotion. Therefore, in this section, we explore to what extent the occurrence of recollections and memory-associated affect influence the accuracy of predicting video-induced emotion. We assess the relative contribution of these memory-related features for personalizing predictions compared to information about viewers' demographics, their personality traits, and their overall mood at the time of exposure to the video.

5.1 Exp. 3: Occurrence of Recollections

In this experiment, we explore the increase in predictive power provided solely by information about the occurrence of recollections in response to videos.

Method and Approach: We constructed a linear mixed regression model for each affective dimension of the emotions induced in viewers: *Induced Pleasure*, *Induced Arousal*, and *Induced Dominance*. The fixed effect Occurrence of Recollections (*REC*) is a factor with two levels, capturing the presence or absence of personal memories as part of the response to a video. A fixed-effect for the identity of the video (*VID*), consisting of a factor with 42 levels, captures information that can be feasibly provided by the video content itself. Additionally, we define the following fixed effects for various characteristics collected as viewer-specific measures (see Table 1):

- (1) Demographics (*DEM*) includes a continuous predictor for viewers' age, one 2-level factor representing viewers' nationality (USA/India), and another 2-level factor for their gender (male/ female).
- (2) Personality (*PER*) comprises five continuous predictors, one for each of the HEXACO personality traits
- (3) Mood (*MOOD*) includes three predictors, one each for viewers' self-reported pleasure, arousal, and dominance

We standardized the continuous predictors in all of the specified fixed effects, as well as target variables, by subtracting their respective mean and dividing by their standard deviation before introducing them into the regression model. Finally, all models include participants' identity as a random effect to account for the repeated measures of responses.

Results: The results for an analysis of variance of the fixed effects specified in our regression models are presented in Table 5. All tests for significance are Bonferroni corrected to account for multiple comparisons between the predictors and the target variables across the different models. Our findings indicate that information about recollection occurrences contributes to the accuracy of models above and beyond the other included sources. However, the amount of unique variation explained by it is rather small ($Avg\Delta R_m^2 = .016$). A look at the remaining effects in the models reveals that video identity provides the biggest insights into viewers' responses across all affective dimensions ($Avg\Delta R_m^2 = .126$). Note that the results for these effects are different from those in

Table 5: Exp. 3 – Comparisons including the Effect of Occurring Recollections on Video-Induced Emotion

Effect	df_n	Induced Pleasure					Induced Arousal				Induced Dominance				$Avg\Delta R_m^2$
		df_d	F	p	ΔR_m^2		df_d	F	p	ΔR_m^2	df_d	F	p	ΔR_m^2	
VID	41	1915.96	9.8	<.001*	.159		1852	6.1	<.001*	.095	1852.99	8.34	<.001*	.126	.126
DEM	3	277.28	8.88	<.001*	.012		277.06	12.39	<.001*	.026	277.07	3.18	.025	.007	.012
PERS	6	275.43	1.38	.222	.004		275.48	1.39	.22	.007	275.48	1.85	.09	.008	.006
MOOD	3	275.2	4.6	.004*	.007		275.29	8.01	<.001*	.021	275.29	6.21	<.001*	.015	.014
REC	1	1629.06	45.58	<.001*	.018		1889.19	29.8	<.001*	.013	1886.63	38.23	<.001*	.017	.016

* Result below Bonferroni corrected critical value for significance: $\alpha_{adj} = .017$

Table 6: Exp. 4 – Comparisons including the Effect of Memory-associated Affect on Video-Induced Emotion

Effect	df_n	Induced Pleasure					Induced Arousal				Induced Dominance				$Avg\Delta R_m^2$
		df_d	F	p	ΔR_m^2		df_d	F	p	ΔR_m^2	df_d	F	p	ΔR_m^2	
VID	41	1921.56	6.67	<.001*	.092		1853.22	4.75	<.001*	.064	1848.95	5.47	<.001*	.071	.075
DEM	3	277.51	9.18	<.001*	.011		277.3	11.16	<.001*	.024	277.28	3.07	.028	.006	.013
PERS	6	275.67	1.91	.08	.005		275.7	0.8	.568	.004	275.69	2.28	.037	.008	.005
MOOD	3	277.3	2.86	.037	.004		276.82	6.01	.001*	.016	276.78	4.7	.003*	.012	.014
MEM	8	1933.19	49.13	<.001*	.134		1919.4	35.56	<.001*	.104	1916.07	42.19	<.001*	.115	.117

* Result below Bonferroni corrected critical value for significance: $\alpha_{adj} = .017$

experiment 2 (see Table 4) because models for the current analysis include viewer-characteristics as additional predictors and do not contain interaction effects. Moreover, the analysis reveals that demographics explain a small amount of variation ($Avg\Delta R_m^2 = .012$) in induced pleasure and arousal, but do not contribute to predictions of dominance. Mood consistently provides a degree of information about emotional responses comparable to that of occurring recollection ($Avg\Delta R_m^2 = .014$). The average overall fit for the fixed effects specified in the models was modest ($AvgR_m^2 = .2$).

5.2 Exp. 4: Memory-Associated Affect

In this experiment, we assess the added benefits for personalizing predictions gained by information about the affect associated with a recollected memory.

Method and Approach: We specify separate linear mixed-effects regression models for each of the affective dimensions of viewers' induced emotions. Each includes the fixed effects specified in Exp. 4 for viewers' demographics (DEM), their personality traits (PERS), and their mood (MOOD), as well as the effect of video identity (VID). In addition to these, all models include a fixed effect for memory-associated affect (MEM), which is a factor with nine levels. It denotes whether a response either (1) involves no recollection, or (2) the octant of the affective rating associated with the recollected memory in the three-dimensional PAD-space (e.g. a memory associated with positive values for pleasure, arousal, and dominance would be assigned to the octant $P_+A_+D_+$). All models include a random effect for viewers' identity to account for the repeated measures in our design. This coding allows models to predict both responses involving memories, and those that do not.

Results: Table 5 lists the results for an analysis of variance of the fixed effects specified in our regression models. All tests for significance are Bonferroni corrected to account for multiple comparisons

between the predictors and the target variables across the different models. Our findings indicate that information about recollection occurrences contributes to the accuracy of models above and beyond the other included sources. However, the amount of unique variation explained by it is rather small ($Avg\Delta R_m^2 = .016$). A look at the remaining effects in the models reveals that video identity is the best predictor for all affective dimensions ($Avg\Delta R_m^2 = .126$). Note that the results for these effects are different from those in experiment 2 (see Table 4) because models for the current analysis include viewer-characteristics as additional predictors and do not contain interaction effects. A look at the effects for viewer-characteristics reveals that demographics explain a small amount of variation ($Avg\Delta R_m^2 = .012$) in induced pleasure and arousal but do not contribute to predictions of dominance. In contrast to this, mood consistently provides a degree of information about emotional responses comparable to that of occurring recollection ($Avg\Delta R_m^2 = .014$), thereby confirming earlier findings for mood-dependency in media responses [31]. Contrary to the results of prior research [15], personality as a whole provided no significant insights into viewers' emotional responses. The average overall fit for the fixed effects specified in the models was modest ($AvgR_m^2 = .2$).

5.3 Discussion

While the occurrence of recollections offers only minor contributions to predictions about viewers' emotional responses to videos, memory-associated affect emerged as the strongest predictor. Moreover, these contributions go above and beyond those offered by viewer-characteristics, or the identity of the eliciting music video.

Providing intelligent applications with information about viewers' recollection processes at prediction time in the dynamic fashion assumed in our analyses is a challenging problem because it requires technology that is both able to meaningfully estimate (1) when recollections do occur and (2) what affect viewers' associate with the

evoked memories. However, despite these substantial obstacles, recollection and memory-associated affect are the most important source of information for predicting video-induced emotions, above and beyond the video itself, personality, mood, and demographics. Our research shows that not addressing the influence of memories will always limit the accuracy of automatic predictions of video-induced affect. In and of itself, this is not a problem, but something to be aware of when developing such systems.

Finally, the overall modest fit across our models in both analyses (occurrence of recollections: $AvgR_m^2 = .2$; memory-associated affect: $AvgR_m^2 = .3$) points towards significant room for improvement by incorporating additional viewer- and situation-specific features. Empirical research from cognitive psychology has described a wide variety of contextual influences on acting on human cognitive-affective processing that could be explored in computational models of video-induced emotions, e.g. the presence of enduring goals and values [27]. An essential step in this direction is the careful development of datasets of emotional responses that systematically capture relevant contextual attributes [11, 31]. Importantly, the overall limited insights offered by static viewer-characteristics in our analyses underline that such efforts should focus increasingly on dynamic attributes of viewers, their cognitions, and the situations in which they take place.

6 LIMITATIONS

There are notable limitations to how our empirical findings generalize to other types of media material or a different viewership. First, our data collection involves only responses to a particular format of media content, i.e. music videos. It is plausible that the connection between personal memories and emotional impact is less profound for other content formats. For example, in feature films, empathy with the portrayed characters in the narrative is a critical aspect [17] that could overshadow the influence of any personal memories. It is also important to point out that we have purposefully selected both participants and content to increase the chance for recollection. While this should not have an impact on the validity of our findings regarding the effects of memories, our data may display an inflated rate of their occurrence. However, a realistic understanding of the conditions under which videos trigger memories in members of the general population requires a more diverse content-participant mixture.

7 TOWARDS MODELING MEMORIES IN AUTOMATIC PREDICTIONS

Attempting to provide applications with information about viewers' recollections offers numerous opportunities for empirical and technological exploration. For example, there exists no direct computational research modelling the evocative properties of videos or of the situations in which people view them. Similar to existing work estimating the likelihood of media content to be remembered (e.g. [20]), modelling videos' capacity to trigger memories could be explored based on their audiovisual features. Additionally, Given that involuntary memory recollection has been connected to attentional drifting, work on predicting such mind-wandering from multimodal data can serve as a starting point for modelling viewers' pensiveness in a situation. In this setting, researchers have

successfully used measures of an individual's physiology or overt behaviour to detect when attention is turning inwards and away from video content to other thoughts [33]. Likewise, existing work from ubiquitous computing, sentiment analysis, and cognitive modelling can form the basis for predicting the affect viewers' associate with their memories (see [10] for a detailed discussion). Such efforts could centre around personal data that has been either collected implicitly by rich ubiquitous sensing (e.g. through lifelogging [16]), or provided explicitly as social media content (e.g. comments in response to media [7], or as entries in smart journals [13]).

8 SUMMARY AND CONCLUSION

In this article, we have presented two contributions relevant to predictions of video-induced emotions: (1) two empirical investigations exploring the effects that being reminded of personal memories by a video has on induced emotions, and (2) two additional experiments in which we explore the relative value of access to features describing viewers' recollection processes for understanding variations in their emotional responses.

The findings of our first set of experiments show that the presence of associated personal memories coincides with a stronger emotional impact on them, independently of the video that is being viewed. This indicates that recollections are a ubiquitous influence on viewers' subjective experience of video material. Moreover, when memories are triggered, induced emotions are often similar to the affect that viewers associate with what has been remembered. However, the degree of this similarity varies across videos, showing that viewers' experience of some videos is more strongly influenced by their memories than that of others. As a consequence, one goal for technological research should be to automatically detect the importance of memories for experiencing a particular video.

In our second set of experiments, we found that both the occurrence of and affect-associated with personal memories explain variation in viewers' emotional responses to videos above and beyond the video itself and relevant viewer-characteristics. These results indicate that providing this information to computational models holds significant potential for predictions of subjective viewing experiences. Moreover, the negligible contribution of static viewer-characteristics to predictions of induced emotions (e.g. personality traits), highlights the necessity of access to such highly situation-specific information. Consequently, without accounting for dynamic influences like personal memories in computational models, accurate predictions of video-induced emotions in real-life applications will remain out of reach. This is a challenging endeavor, but we have outlined several lines of existing technological research that can form a starting point for exploring automatic predictions of when memories occur, and how they impact a viewers' experience. As such, progress seems difficult, but possible.

ACKNOWLEDGMENTS

This work has been partially supported by the 4TU research center Humans & Technology (H&T) project (Systems for Smart Social Spaces for Living Well: S4) and by the Netherlands Organization for Scientific Research (NWO) under the MINGLE project number 639.022.606.

REFERENCES

- [1] Hans Baumgartner. 1992. Remembrance of things past: Music, autobiographical memory, and emotion. *Advances in Consumer Research* 19 (1992), 613–620. arXiv:9301106587
- [2] Yoann Baveye, Christel Chamaret, Emmanuel Dellandrea, and Liming Chen. 2018. Affective video content analysis: A multidisciplinary insight. *IEEE Transactions on Affective Computing* 9, 4 (2018), 396–409. <https://doi.org/10.1109/TAFFC.2017.2661284>
- [3] Amy M. Belfi, Brett Karlan, and Daniel Tranel. 2016. Music evokes vivid autobiographical memories. *Memory* 24, 7 (aug 2016), 979–989. <https://doi.org/10.1080/09658211.2015.1061012>
- [4] Dorte Berntsen, Søren Rislov Staugaard, and Louise Maria Torp Sørensen. 2013. Why am I remembering this now? Predicting the occurrence of involuntary (spontaneous) episodic memories. *Journal of Experimental Psychology: General* 142, 2 (2013), 426–444. <https://doi.org/10.1037/a0029128>
- [5] Joost Broekens and Willem-Paul Brinkman. 2013. AffectButton: A method for reliable and valid affective self-report. *International Journal of Human-Computer Studies* 71, 6 (jun 2013), 641–667. <https://doi.org/10.1016/j.ijhcs.2013.02.003>
- [6] Martin A. Conway and Shamsul Haque. 1999. Overshadowing the Reminiscence Bump: Memories of a Struggle for Independence. *Journal of Adult Development* 6, 1 (1999), 35–44. <https://doi.org/10.1023/A:1021672208155>
- [7] Dan Cosley, Victoria Schwanda, S Tejaswi Peesapaty, Johnathon Schultz, and Jonathan Baxter. 2009. Experiences with a publicly deployed tool for reminiscing. In *Proc. First Int'l Workshop on Reminiscence Systems*. 31–36.
- [8] Reinout E. de Vries. 2013. The 24-item Brief HEXACO Inventory (BHI). *Journal of Research in Personality* 47, 6 (dec 2013), 871–880. <https://doi.org/10.1016/j.jrp.2013.09.003>
- [9] A. M. DO, A. V. RUPERT, and G. WOLFORD. 2008. Evaluations of pleasurable experiences: The peak-end rule. *Psychonomic Bulletin & Review* 15, 1 (feb 2008), 96–98. <https://doi.org/10.3758/PBR.15.1.96>
- [10] Bernd Dudzik, Hayley Hung, Mark Neerinx, and Joost Broekens. 2018. Artificial Empathic Memory. In *Proceedings of the 2018 Workshop on Understanding Subjective Attributes of Data, with the Focus on Evoked Emotions - EE-USAD '18*. ACM Press, New York, New York, USA, 1–8. <https://doi.org/10.1145/3267799.3267801>
- [11] Bernd Dudzik, Michel-Pierre Jansen, Franziska Burger, Frank Kaptein, Joost Broekens, Dirk K.J. Heylen, Hayley Hung, Mark A. Neerinx, and Khiet P. Truong. 2019. Context in Human Emotion Perception for Automatic Affect Detection: A Survey of Audiovisual Databases. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 206–212. <https://doi.org/10.1109/ACII.2019.8925446>
- [12] Paul Ekman. 1999. Basic emotions. , 45–60 pages. <https://doi.org/10.1002/0470013494.ch3>
- [13] Chris Elsdén, Abigail C. Durrant, and David S. Kirk. 2016. It's Just My History Isn't It?. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 2819–2831. <https://doi.org/10.1145/2858036.2858103>
- [14] Shi Feng, Sidney D'Mello, and Arthur C. Graesser. 2013. Mind wandering while reading easy and difficult texts. *Psychonomic Bulletin & Review* 20, 3 (jun 2013), 586–592. <https://doi.org/10.3758/s13423-012-0367-y>
- [15] Sharath Chandra Guntuku, Weisi Lin, Michael James Scott, and Gheorghita Ghinea. 2015. Modelling the influence of personality and culture on affect and enjoyment in multimedia. In *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 236–242. <https://doi.org/10.1109/ACII.2015.7344577>
- [16] Cathal Gurrin, Alan F. Smeaton, and Aiden R. Doherty. 2014. LifeLogging: Personal Big Data. *Foundations and Trends® in Information Retrieval* 8, 1 (2014), 1–125. <https://doi.org/10.1561/15000000033>
- [17] Juan-José Igartua. 2010. Identification with characters and narrative persuasion through fictional feature films. *Communications* 35, 4 (jan 2010), 347–373. <https://doi.org/10.1515/comm.2010.019>
- [18] Petr Janata, Stefan T. Tomic, and Sonja K. Rakowski. 2007. Characterisation of music-evoked autobiographical memories. *Memory* 15, 8 (nov 2007), 845–860. <https://doi.org/10.1080/09658210701734593>
- [19] Lisanne M. Jenkins and David G. Andrews. 2012. A New Set of Standardised Verbal and Non-verbal Contemporary Film Stimuli for the Elicitation of Emotions. *Brain Impairment* 13, 2 (sep 2012), 212–227. <https://doi.org/10.1017/BrImp.2012.18>
- [20] Aditya Khosla, Akhil S. Raju, Antonio Torralba, and Aude Oliva. 2015. Understanding and Predicting Image Memorability at a Large Scale. In *2015 IEEE International Conference on Computer Vision (ICCV)*, Vol. 2015 Inter. IEEE, 2390–2398. <https://doi.org/10.1109/ICCV.2015.275>
- [21] S. Koelstra, C. Muhl, M. Soleymani, Jong-Seok Lee, A. Yazdani, T. Ebrahimi, T. Pun, A. Nijholt, and I. Patras. 2012. DEAP: A Database for Emotion Analysis Using Physiological Signals. *IEEE Transactions on Affective Computing* 3, 1 (jan 2012), 18–31. <https://doi.org/10.1109/T-AFFC.2011.15>
- [22] John H. Mace. 2004. Involuntary autobiographical memories are highly dependent on abstract cuing: the Proustian view is incorrect. *Applied Cognitive Psychology* 18, 7 (nov 2004), 893–899. <https://doi.org/10.1002/acp.1020>
- [23] Daniel G. McDonald, Melanie A. Sarge, Shu-Fang Lin, James G. Collier, and Bridget Potocki. 2015. A Role for the Self: Media Content as Triggers for Involuntary Autobiographical Memories. *Communication Research* 42, 1 (feb 2015), 3–29. <https://doi.org/10.1177/0093650212464771>
- [24] Daniel McDuff and Mohammad Soleymani. 2017. Large-scale Affective Content Analysis: Combining Media Content Features and Facial Reactions. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*. IEEE, 339–345. <https://doi.org/10.1109/FG.2017.49>
- [25] Albert Mehrabian. 1996. Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in Temperament. *Current Psychology* 14, 4 (dec 1996), 261–292. <https://doi.org/10.1007/BF02686918>
- [26] Caitlin Mills and Sidney D'Mello. 2014. On the Validity of the Autobiographical Emotional Memory Task for Emotion Induction. *PLoS ONE* 9, 4 (apr 2014), e95837. <https://doi.org/10.1371/journal.pone.0095837>
- [27] Agnes Moors, Phoebe C. Ellsworth, Klaus R. Scherer, and Nico H. Frijda. 2013. Appraisal Theories of Emotion: State of the Art and Future Development. *Emotion Review* 5, 2 (apr 2013), 119–124. <https://doi.org/10.1177/1754073912468165>
- [28] Shinichi Nakagawa and Holger Schielzeth. 2013. A general and simple method for obtaining R² from generalized linear mixed-effects models. *Methods in Ecology and Evolution* 4, 2 (feb 2013), 133–142. <https://doi.org/10.1111/j.2041-210x.2012.00261.x>
- [29] James A. Russell. 1980. A circumplex model of affect. *Journal of Personality and Social Psychology* 39, 6 (dec 1980), 1161–1178. <https://doi.org/10.1037/h0077714>
- [30] Michael James Scott, Sharath Chandra Guntuku, Weisi Lin, and Gheorghita Ghinea. 2016. Do Personality and Culture Influence Perceived Video Quality and Enjoyment? *IEEE Transactions on Multimedia* 18, 9 (sep 2016), 1796–1807. <https://doi.org/10.1109/TMM.2016.2574623>
- [31] Mohammad Soleymani, Martha Larson, Thierry Pun, and Alan Hanjalic. 2014. Corpus Development for Affective Video Indexing. *IEEE Transactions on Multimedia* 16, 4 (jun 2014), 1075–1089. <https://doi.org/10.1109/TMM.2014.2305573> arXiv:1211.5492
- [32] Mohammad Soleymani, Maja Pantic, and Thierry Pun. 2012. Multimodal Emotion Recognition in Response to Videos. *IEEE Transactions on Affective Computing* 3, 2 (apr 2012), 211–223. <https://doi.org/10.1109/T-AFFC.2011.37>
- [33] Angela Stewart, Nigel Bosch, Huili Chen, Patrick Donnelly, and Sidney D'Mello. 2017. Face Forward: Detecting Mind Wandering from Video During Narrative Film Comprehension. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 10331 LNAI. Springer, Cham, 359–370. https://doi.org/10.1007/978-3-319-61425-0_30
- [34] Doménique van Gennip, Elise van den Hoven, and Panos Markopoulos. 2015. Things That Make Us Reminisce. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*, Vol. 1. ACM Press, New York, New York, USA, 3443–3452. <https://doi.org/10.1145/2702123.2702460>
- [35] Shangfei Wang and Qiang Ji. 2015. Video Affective Content Analysis: A Survey of State-of-the-Art Methods. *IEEE Transactions on Affective Computing* 6, 4 (oct 2015), 410–430. <https://doi.org/10.1109/TAFFC.2015.2432791>
- [36] Ashkan Yazdani, Evangelos Skodras, Nikolaos Fakotakis, and Touradj Ebrahimi. 2013. Multimedia content analysis for emotional characterization of music video clips. *EURASIP Journal on Image and Video Processing* 2013, 1 (dec 2013), 26. <https://doi.org/10.1186/1687-5281-2013-26>
- [37] Yi Zhu, Ingrid Heynderickx, and Judith A. Redi. 2015. Understanding the role of social context and user factors in video Quality of Experience. *Computers in Human Behavior* 49 (aug 2015), 412–426. <https://doi.org/10.1016/j.chb.2015.02.054>