

PAGAN: Video Affect Annotation Made Easy

David Melhart
Institute of Digital Games
University of Malta
 Msida, Malta
 david.melhart@um.edu.mt

Antonios Liapis
Institute of Digital Games
University of Malta
 Msida, Malta
 antonios.liapis@um.edu.mt

Georgios N. Yannakakis
Institute of Digital Games
University of Malta
 Msida, Malta
 georgios.yannakakis@um.edu.mt

Abstract—How could we gather affect annotations in a rapid, unobtrusive, and accessible fashion? How could we still make sure that these annotations are reliable enough for data-hungry affect modelling methods? This paper addresses these questions by introducing PAGAN, an accessible, general-purpose, online platform for crowdsourcing affect labels in videos. The design of PAGAN overcomes the accessibility limitations of existing annotation tools, which often require advanced technical skills or even the on-site involvement of the researcher. Such limitations often yield affective corpora that are restricted in size, scope and use, as the applicability of modern data-demanding machine learning methods is rather limited. The description of PAGAN is accompanied by an exploratory study which compares the reliability of three continuous annotation tools currently supported by the platform. Our key results reveal higher inter-rater agreement when annotation traces are processed in a relative manner and collected via unbounded labelling.

Index Terms—Affective computing, human-computer interaction, affect annotation, crowdsourcing, video

I. INTRODUCTION

The question of how to best collect many reliable and valid affect annotations is getting increasingly important in affective computing. Video-based annotation—as a popular approach in affective computing—requires participants to watch a set of videos and annotate their content, which is a cumbersome and costly process. Given the ever increasing use of data-hungry affect modelling techniques, however, the need for larger and more reliable affective corpora is growing.

Meanwhile, the majority of current frameworks for both discrete and continuous video affect annotation pose a number of limitations. Tools such as FeelTrace [1], ANNEMO [2], AffectButton [3], GTrace [4], CARMA [5], AffectRank [6], and RankTrace [7] often require local installation and calibration. Not only does this require the presence of a researcher while conducting the study, but it often necessitates knowledge of a programming language as well. This low level of accessibility constrains the widespread use of such annotation tools (i.e. in the wild) and, in turn, results in affective datasets of limited size and use. This limitation is even more severe for newly emerging fields, such as game user research, without large established corpora, where data collection is a necessity.

To address the above limitations this paper introduces a *general-purpose, online* video annotation platform, namely the *Platform for Audiovisual General-purpose ANnotation* (PA-

GAN)¹. The platform is publicly available and free to use for research purposes. PAGAN provides researchers with an easy and *accessible* way to crowdsource affect annotations for their videos. In contrast to other popular annotation tools, PAGAN does not require a local installation and is designed to help researchers organise and disseminate their research projects to a large pool of participants. Inspired by [8], the whole annotation process is done through a web interface operating on any modern web browser. Outsourcing the labelling task is as simple as sharing the corresponding project link.

PAGAN currently features three one-dimensional affect labelling techniques representing different methods for measuring the ground truth of affect: *GTrace* [4], *BTrace* (a modified version of AffectRank [6]) and *RankTrace* [7]. In addition to the detailed description of the platform, the paper offers an exploratory study which examines the reliability of the three annotation methods. The results of this study reveal higher degrees of inter-rater agreement when traces are processed in a relative manner and collected via unbounded labelling.

II. BACKGROUND

PAGAN is centred on *dimensional* and, primarily, *continuous* affect annotation. To motivate this focus, this section presents the theoretical background of categorical versus dimensional emotional representations, and time-discrete versus time-continuous affect annotation techniques.

A. Categorical vs. Dimensional Representation of Emotions

Theories of emotions are generally represented in two main ways: as *dimensions* or as *categories*. The former focuses on emotions as emerging sentiments, which are functions of simple affective dimensions [9]–[11]. The latter promotes an understanding, in which basic emotions are distinct from one-another in function and manifestation [12], [13]. Today, both schools of thought have contemporary continuation, with some frameworks aiming to reconcile the two viewpoints [14].

Categorical emotion representation is largely inspired by the work of Ekman [12] and is based on the assumption that humans elicit distinct emotions, which are inherent to the human psyche and universally understood. While normative studies have confirmed the generality of these frameworks to an extent [15], [16], putting these theories into practice also brings about

¹<http://pagan.institutedigitalgames.com/>

some conceptual limitations. The underlying assumption of clear division between basic emotional responses is challenged by a criterion bias when categorising fuzzy responses, and the subjective evaluation of emotions based on contextual cues highlights the relative nature of emotional appraisal [17] and calls the universality of these frameworks into question.

Alternatively, emotions can be represented through affective dimensions which typically follow Russell’s *Circumplex Model of Emotions* [11] or the *Pleasure-Arousal-Dominance model* [10]. Many contemporary annotation tools [1], [4], [7], [18] use one of these models for annotating one or more affective dimensions. The main limitation of these frameworks is that they cannot describe complex and self-reflexive emotions without expert interpretation, which could reintroduce biases to the observations. However, this simplicity also results in high *face validity* [19], reducing guesswork and *criterion bias* of the annotator [20] even in the case of fuzzy responses, which otherwise would be hard to categorise.

PAGAN focuses on time-continuous annotation to capture the temporal dynamics of affective experiences. Because this task often involves identifying fuzzy transitions between affective responses, it relies on a dimensional representation of emotion. The choice of this focus is also motivated by the relatively low cognitive load of one-dimensional labelling compared to evaluating the manifestations and transitions of multiple distinct emotional categories.

B. Discrete vs. Continuous Annotation Methods

Traditional surveys, such as the Self-Assessment Manikin [18], were developed to measure fixed scales with discrete items. While computer interfaces allowed for the development of time-continuous annotation tools, traditional surveys and digital tools for discrete affect annotation [3], [6] are still prevalent. Although these approaches capture less of the temporal dynamics of the experience [21], compartmentalising annotations could help reduce the noise of the labels and yield higher inter-rater agreement. Yannakakis and Martinez compared the nominal and ordinal representation of discrete affect annotations (AffectRank) with continuous bounded ratings (FeelTrace) [6]. Their study found that a nominal representation yields higher inter-rater agreement compared to treating a continuous trace as interval data.

Treating continuous annotation traces as interval data and processing them in an absolute fashion remains the prominent method of many studies [21], [22]. As this methodology necessitates that the trace is bounded to ensure a common scale among raters, interval processing of annotation traces provides data in a form that can be analysed via a wide array of statistical and machine learning approaches. However, there are serious caveats in representing inherently subjective experiences in an absolute fashion. Supported by the *adaptation level theory* [23], *habituation* [24], the *somatic-marker hypothesis* [25], and numerous studies within affective computing [26], [27] it appears that subjects experience stimuli in relation to their prior emotional and physiological states, experiences, and

memories. Thus, any annotation task is subject to a number of *anchoring* [28], *framing* [29] and *recency* [30] effects.

PAGAN aims to overcome the above limitations by featuring unbounded continuous annotation (via RankTrace). With this labelling protocol, the data is no longer structured along the same scale, which makes processing traces as absolute values problematic. However, this method overcomes the limitations of interval processing that arise from the discrepancy between the players’ cognitive evaluation processes and absolute scales. The relative processing of unbounded annotation traces has been shown to correlate with physiological signals [7], and predictive models based on these traces have been shown to generalise better [31].

C. Annotation Tools

In recent years, tools for affect labelling have diversified. Earlier examples such as FeelTrace [1] in 2000, AffectButton [3] in 2009, and AffectRank [6] in 2014 aim to capture a complex phenomenon by measuring two or three affective dimensions at once. Yet recent studies [4], [5], [7] focus on one-dimensional labelling. The shift away from multi-dimensional labelling can be explained by the increased cognitive load induced by these methods that comes with more complex tasks. Increased cognitive load can undermine the strengths of dimensional emotion representation [4], as one emotional axis can take precedence over the other—which could impact face validity. PAGAN implements three variations of one-dimensional affect labelling techniques, representing different methods for measuring the ground truth of affect: *GTrace*, as bounded and continuous; *BTrace* (binary trace), as real-time discrete; and *RankTrace*, as unbounded and continuous annotation techniques.

GTrace [4] was created as a bounded, continuous annotation tool and quickly became popular for affect labelling in human-computer interaction and affective computing [32]–[35]. *GTrace* has a limited memory and displays only the last few annotation values. PAGAN implements its own version of *GTrace* based on the description of the tool in [4] and [32].

Due to concerns regarding traces that are processed as intervals, *AffectRank* [6] was introduced as a real-time rank-based discrete labelling tool. As a two-dimensional annotation tool, the main improvement of *AffectRank* over *FeelTrace* was the focus on recording ordinal changes instead of absolute values. *BTrace* in PAGAN is inspired by *AffectRank* as it measures ordinal change, but is much simpler: instead of the 8 annotation options in *AffectRank*, *BTrace* focuses on one affective dimension and two nominal labels: positive vs. negative change. A major limitation of both *AffectRank* and *BTrace*, however, is the discrete nature of the provided labels which limits the resolution of the collected ground truth data.

To cater for the subjects’ relative judgement models, *RankTrace* was introduced for unbounded and relative annotation [7]. In *RankTrace* the annotator effectively draws a graph of their experience (Fig. 1) which acts as the annotators’ point of reference. *RankTrace* produces continuous and unbounded traces which can be observed as ordinal changes and be



Fig. 1. RankTrace interface in the PAGAN platform. *Apex Legends* (Electronic Arts, 2019) gameplay footage. No copyright infringement intended.

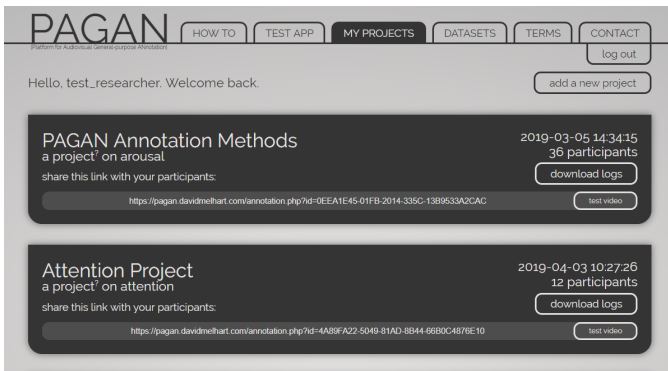


Fig. 2. Project summaries on the administration interface.

processed in a relative manner [7], [31]. Because RankTrace is unbounded, it lets the annotator react to the situation compared to previous experiences instead of forcing them to evaluate the stimuli in an absolute manner. In addition to GTrace and BTrace, the current PAGAN framework features a version of the RankTrace annotation method.

III. PAGAN PLATFORM DESCRIPTION

This section provides a description of the PAGAN platform, its user interface and general usage. The user interface of PAGAN consists of two separate sections. One is a web interface for researchers to prepare the annotation task (Section III-A) and the other is an interface for annotation by end-users (Section III-B). Section III-C details the three annotation methods incorporated currently in PAGAN and used in the evaluation study of Section IV.

A. Administration Interface

Researchers access and create their projects through a dedicated page. Each user has a secure login with a username and a password. After login, the researcher accesses their *project*



Fig. 3. Project creation screen on the researcher interface.

summaries (Fig. 2). Here, they can create new projects, view the progress of their ongoing studies, and access their corresponding annotation logs. Each project has a corresponding link, which is meant to be shared with potential participants. The annotation application can also be run in *test mode* from here, in which case the annotation logs are not saved.

The project creation screen can be seen in Fig. 3. Projects are highly customisable to accommodate different research needs. The project title identifies the study on the project summary page and displayed to the participants as part of the welcome message. The *annotation target* is the label for the *y* axis of the annotator (see Fig. 1). The project can be sourced from one or more uploaded videos or YouTube² links. The videos can be loaded either randomly or in sequence. If *endless mode* is selected, PAGAN rotates the videos indefinitely, allowing a participant to complete all tasks multiple times. In case of a randomised video order, there is an option to limit the number of videos a participant has to annotate. The videos can be played with or without sound; if videos are played with sound, PAGAN reminds participants to turn on their speakers or headphones. The researcher can optionally add information or instructions viewed before and/or after the annotation tasks to help integrate the platform into the larger research project.

²<https://www.youtube.com/>

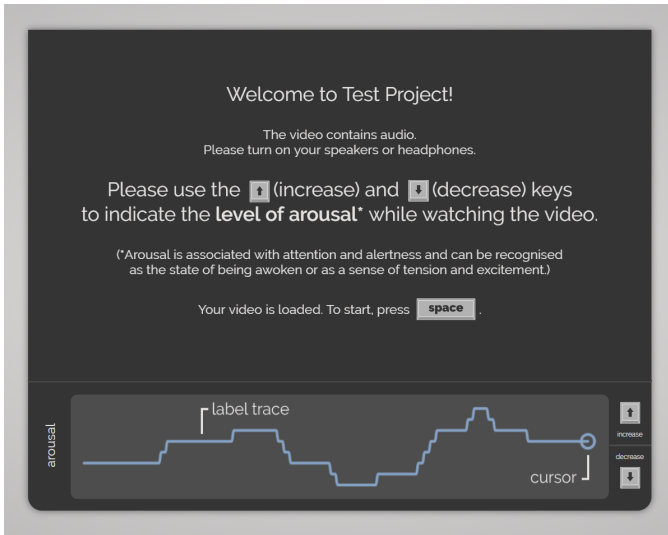


Fig. 4. Welcome message displayed to participants before annotation starts.

Finally, a survey link can be included, which is displayed to the participant at the end of the annotation session.

B. Annotator Interface

The annotator application is a separate interface from the researcher site and meant to be used by the participants of the study. The interface is designed to display only the necessary information, thus eliminating potential distractions. Upon navigating to the project link (see Section III-A), the participant is greeted by a welcome message which concisely explains the annotation procedure and provides some information about the annotation target (Fig. 4). After the video is loaded, the participant can start the annotation process (Fig. 1) at their leisure. The design of PAGAN eliminates the use of a computer mouse in favour of the more readily-available keyboard. The annotation is performed with the *up* and *down* keys on the keyboard and the session can be paused by pressing *space*. To minimise the amount of sessions with insufficient annotation, the system only logs a session as “completed” if at least 25% is seen and pauses if the browser tab is out of focus (i.e. if the participant leaves the annotation interface open but switches to a different tab or window).

C. Annotation Methods

This section presents the annotation techniques included in the PAGAN framework: RankTrace, GTrace, and BTrace.

1) *RankTrace*: The implementation of RankTrace closely follows the original by Lopes et al. [7] (Fig. 5a). The only major distinction to their version is the exponential acceleration of the annotator cursor when a control key is held down. This change was made because the original version of RankTrace uses a wheel interface where the magnitude of change can be controlled easier by the participant. As the annotation trace displays the entire history, the participant has sufficient visual feedback which acts as a reference (anchoring) point [27] for the subjective evaluation of the experience.

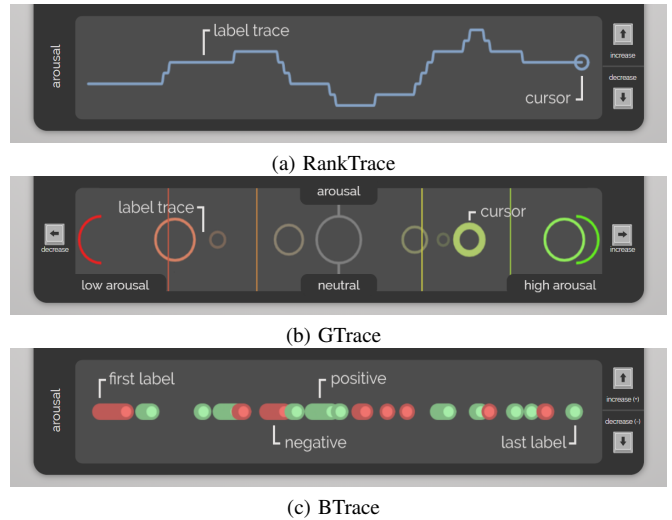


Fig. 5. Interfaces for the annotation methods included in this study.

2) *GTrace*: Similarly to how the tool is used in Baveye et al. [32], the user interface is moved under the video; vertical lines are added as an allusion to a traditional 7-item scale to provide a visual aid for the absolute evaluation of the trace (Fig. 5b). Similarly to RankTrace, the movement of the cursor is accelerated when a key is held down as the original implementation used a mouse cursor allowing for higher speed while retaining precision. When the participant stops the cursor, it leaves a mark which slowly fades, providing limited memory of previous positions, to which the participant can compare new labels. The limited memory from the fading mark differs from both BTrace and RankTrace which display the full history of the session.

3) *BTrace*: Binary Trace (BTrace) is a new annotation tool introduced in this paper which is largely based on AffectRank [6]. BTrace is designed as a simple alternative to relative annotation in a discrete manner, using two nominal categories: +1 as increase (or positive change) and -1 as decrease (or negative change). In that regard, it could be viewed as an one-dimensional version of AffectRank. The design of the tool, however, is based on the benefits reference points have on the reliability of the obtained annotation labels [7], [27] and thus it displays the full history of the annotation session as red and green blobs (see Fig. 5c).

IV. EXAMPLE STUDY

This section presents a small-scale exploratory study conducted with the PAGAN platform. The goal of this study is two-fold. First, we present the usage of the system in a real-world scenario; and second, we examine the effectiveness of relative annotation methods compared to absolute affect labelling. This study focuses on the perceived arousal level of different videos with emotional content. Our two relative annotation methods are RankTrace and BTrace, and our absolute method is a variant of GTrace (see Sec. III-C and Fig. 5).

TABLE I
SUMMARY OF PARTICIPANTS’ CHOICE OF MOST AND LEAST INTERESTING VIDEO AND MOST AND LEAST INTUITIVE ANNOTATION TOOL

Most / Least Interesting Video		
Apex	GoT	SEMAINE
11 / 6	15 / 1	3 / 22
Most / Least Intuitive Tool		
RankTrace	GTrace	BTrace
16 / 4	8 / 8	5 / 17

TABLE II
NUMBER OF ANNOTATION TRACES FOR EACH VIDEO AND ANNOTATION TYPE, THE AVERAGE NUMBER OF SAMPLES ARE SHOWN IN BRACKETS

Annotation Type	Apex	GoT	SEMAINE
RankTrace	11 (229)	10 (514)	10 (425)
GTrace	11 (429)	12 (358)	8 (133)
BTrace	9 (302)	8 (120)	14 (85)

A. Collected Data

The collected data consists of 108 annotated videos from 36 participants. Participants were found through the social and academic network of the researchers, while subsequent parties were added through snowball sampling by participants sharing the project link. The average age of the participants is 29 years old and 69% identified as male, 24% identified as female, one subject identified as queer and one did not want to identify themselves. The majority of the participants were avid gamers, with 59% playing more than once a week. Each participant was asked to annotate three videos with different but emotionally evocative content: (a) recorded gameplay from *Apex Legends* (Electronic Arts, 2019) (Apex), a popular Battle Royale-style game; (b) the Season 8 trailer of the TV series *Game of Thrones* (HBO, 2019) (GoT); (c) a conversation between a human participant and “Spike”, the angry virtual agent in the *SEMAINE* database [36]. All videos are approximately 2 minutes long. Each video was assigned a random annotation type, discussed in Section III-C. The order of videos was also randomised.

Participants were asked to name the most and least interesting of the three videos and the most and least intuitive of the three annotation tools, effectively ranking them. The results of their preferences are summarised in Table I. The GoT trailer was the most popular (only one participant rated it as the least interesting), while the video from the SEMAINE database was by far the least liked (it collected 81% of “least interesting” votes). In terms of usability, participants ranked RankTrace the most intuitive (as it received 55% of “most intuitive” votes), GTrace second, and BTrace the least intuitive.

B. Methodology

To measure the reliability of the different annotation techniques over the different videos, we observe the inter-rater agreement between participants. Inspired by Yannakakis and Martinez [6], we measure the inter-rater agreement with the *Krippendorff’s* α coefficient [37], which is a robust metric of the degree of agreement corrected for chance between any

number of observers and any type of data. *Krippendorff’s* $\alpha = 1 - D_o/D_e$, where D_e denotes the expected and D_o the observed disagreements between annotations.

Krippendorff’s α is adjusted to the level of measurement of the observations through the weighing of the expected and observed coincidences (see [38] for a complete explanation). This robustness allows for a fair comparison between different annotation methods. *Krippendorff’s* α has an upper bound of 1, which indicates absolute agreement, while 0 signifies no agreement or pure chance. At *Krippendorff’s* $\alpha < 0$, disagreements between annotators are systematic and go beyond chance-based levels.

To allow for a comparison between discrete and continuous annotation and smooth out some of the surface differences between individual traces, we compartmentalise the signals into equal length time-windows. This method of preprocessing is often used in affective computing to preprocess time-continuous signals [6], [7], [31]. We clean the dataset of traces which either had extremely few samples from annotation (less than 3) or where viewing time was less than a minute. This cleanup process removed 15% of traces, and the final datasets comprise of 92 traces. Table II shows the number of traces and samples in each dataset and annotation method. In this study 3-second time windows are considered without any overlap. Potentially the 3-second processing provides approximately 40 windows per participant. As some participants did not complete the full annotation task, this number can vary. However, to maximise the sample sizes, we decided to keep these traces as *Krippendorff’s* α can be applied to data with missing observations as well.

As BTrace already encodes perceived change, similarly to AffectRank [6], we compute the value of time windows as the *sum of annotation values* (Σ_A) within each window, adding values in case of increase and subtracting them in case of decrease. For RankTrace and GTrace, we consider both an absolute and a relative metric [31]: the mean value (μ_A) and average gradient (Δ_A) of time-windows based on the min-max normalised traces. We consider the mean value an absolute metric because it denotes the general level of the participant’s response in a given time-window. In contrast, the average gradient of a time-window considers the amount and direction of the change that happened, as it is computed from the differences of adjacent datapoints of the trace [7], [31]. The calculation of *Krippendorff’s* α is adjusted to the observed metric. When the annotation trace is processed into a relative metric (Σ_A, Δ_A), we compare annotation values as ordinal variables. When the annotation trace is processed into an absolute metric (μ_A), we compare annotation values as interval variables.

C. Results

This section presents the results of the statistical analysis and an interpretation of the results. The calculated inter-rater agreement based on *Krippendorff’s* α scores are displayed in Table III. For the purpose of comparisons of RankTrace and GTrace, we use the highest α value between Δ_A and μ_A .

TABLE III
KRIPPENDORFF’S α ACROSS ANNOTATION TRACES PROCESSED AS
3-SECOND TIME WINDOWS

Annotation		Video		
Tool	Processing	Apex	GoT	SEMAINE
RankTrace	Δ_A	0.2025	0.1760	-0.0043
	μ_A	0.1542	-0.0227	0.0147
GTrace	Δ_A	0.1517	0.1224	-0.0347
	μ_A	0.1857	0.0549	0.0926
BTrace	Σ_A	0.3193	0.0973	0.0249

The highest α values for RankTrace are 0.20 for Apex and 0.18 for GoT, which are higher than the highest α values for GTrace (0.19 and 0.12 respectively). For both GoT and Apex videos, the highest α values are found with Δ_A in three of the four instances examined (except for annotations with GTrace on the Apex dataset), which is further evidence that processing time-windows of GTrace ratings through a relative measure yields more consistent results. Interestingly, both GTrace and RankTrace have a higher α value with μ_A for the SEMAINE video (with GTrace having superior performance), although generally these values are very low and any inter-rater agreement could be chance-based. The general findings from these comparisons are in line with a growing body of research promoting the relative collection and processing of affective annotation traces [27], [31], [39].

Based on Table III, it seems that BTrace achieves the highest inter-rater agreement on the Apex dataset, while showing lower reliability on GoT and SEMAINE videos. As the compartmentalised binary labels denote the rough amount of perceived change in a time-window (but not its magnitude), the possibility of relatively high inter-rater agreement is not surprising. However, results on the GoT and SEMAINE videos show the unreliability of this method. A possible reason for the high variance in the inter-rater agreement is the low face validity of the method. BTrace collected 59% of the “least intuitive” votes among the three annotation methods. Therefore, despite its potential robustness in certain cases, BTrace has shown to be the least reliable and intuitive to use.

An unexpected finding of this analysis is the overall low inter-rater agreement of all methods on the chosen SEMAINE video, which was also ranked as the least interesting by participants. A plausible explanation of the results is a connection between the context and intensity of the affective content and the reliability of the annotation traces. While games and trailers are designed to elicit arousal, the slow pace of SEMAINE videos can be unappealing by comparison. The differences in inter-rater agreements between the Apex and GoT datasets also point towards the role of context in emotion elicitation. While the GoT video is authored to elicit high arousal, the Apex footage presents a more organic scenario with relative calm periods and high-octane action. Especially for frequent videogame players, who have personal experiences with the dynamics of shooter games, this video is easier to interpret and the affective high-points are easier to recognise. This is also supported by a recent study of Jaiswal et al. [40], who

also observed an effect between the context of the annotation task and the quality of labels.

V. DISCUSSION

This paper presented an online platform for crowdsourcing affect annotations, providing researchers with an accessible tool for labelling any kind of audiovisual content. A companion study showcased the usability of the platform, highlighted the reliability of the supported annotation techniques, and compared bounded, unbounded and binary annotations of arousal. Results showed that an unbounded relative annotation method which includes the entire history of labels as reference points is more intuitive to use. Moreover, the study included three videos indicative of different sources of arousal: a game, a TV series trailer, and a dialogue with a virtual agent. Our analysis reveals low inter-rater agreement on the SEMAINE database video, which raises the question on whether more engaging forms of emotion elicitation such as games would offer more reliable benchmarks for affective computing research.

The main limitation of the user study of Section IV is the preliminary nature of the analysis. Arguably with a more thorough pruning, regularisation of the annotation traces [41], quality control of the labels using gold-standards [42], or a strict selection process for the included participants [43], higher inter-rater agreement can be achieved. Moreover, the exploratory nature of this study assessed how different types of videos can be annotated; in a more concise study the set of videos should likely be both larger and more consistent in terms of subject matter. However, as the main focus of this paper was the introduction of the PAGAN platform, these explorations were out of scope of the current study.

There is a large number of features that can be incorporated into the PAGAN platform, such as the support for more flexible research protocols through participant uploads. In the future, PAGAN can be extended with data preprocessing, analysis, and visualisation tools, providing researchers a toolbox for not just data collection but preliminary analysis as well. Such a toolbox could include automatic processing of traces into time windows, outlier detection and pruning; statistical summary and analysis in terms of inter-rater agreement. Machine learning support can be integrated with PAGAN as well, either to preprocess and format data for other software, such as the *Preference Learning Toolbox* [44], [45] or as a light-weight predictive modelling module in PAGAN itself.

VI. CONCLUSION

This paper presented a highly customisable and accessible online platform to aid affective computing researchers and practitioners in the crowdsourcing process of video annotation tasks. In a companion study, we demonstrated the reliability of the supported annotation techniques and showed the strength of relative annotation processing. Our key findings advocate the use of relative, continuous, and unbounded annotation techniques and the use of videogames as active elicitors of emotional responses.

ACKNOWLEDGEMENTS

This paper is funded, in part, by the H2020 project Com N Play Science (project no: 787476).

REFERENCES

- [1] R. Cowie, E. Douglas-Cowie, S. Savvidou*, E. McMahon, M. Sawey, and M. Schröder, “FEELTRACE’: An instrument for recording perceived emotion in real time,” in *ISCA tutorial and research workshop (ITRW) on speech and emotion*, 2000.
- [2] F. Ringeval, A. Sonderegger, J. Sauer, and D. Lalanne, “Introducing the recola multimodal corpus of remote collaborative and affective interactions,” in *Proceedings of the Automatic Face and Gesture Recognition Conference*. IEEE, 2013.
- [3] J. Broekens and W.-P. Brinkman, “Affectbutton: A method for reliable and valid affective self-report,” *Intl. Journal of Human-Computer Studies*, vol. 71, no. 6, pp. 641–667, 2013.
- [4] R. Cowie, M. Sawey, C. Doherty, J. Jaimovich, C. Fyans, and P. Stapleton, “Gtrace: General trace program compatible with EmotionML,” in *Proceedings of the Intl. Conference on Affective Computing and Intelligent Interaction*. IEEE, 2013, pp. 709–710.
- [5] J. M. Girard, “Carma: Software for continuous affect rating and media annotation,” *Journal of Open Research Software*, vol. 2, no. 1, 2014.
- [6] G. N. Yannakakis and H. P. Martinez, “Grounding truth via ordinal annotation,” in *Proceedings of the Intl. Conference on Affective Computing and Intelligent Interaction*. IEEE, 2015, pp. 574–580.
- [7] P. Lopes, G. N. Yannakakis, and A. Liapis, “Ranktrace: Relative and unbounded affect annotation,” in *Proceedings of the Intl. Conference on Affective Computing and Intelligent Interaction*. IEEE, 2017, pp. 158–163.
- [8] D. McDuff, R. Kaliouby, T. Senechal, M. Amr, J. Cohn, and R. Picard, “Affectiva-MIT facial expression dataset (AM-FED): Naturalistic and spontaneous facial expressions collected “in-the-wild”,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2013, pp. 881–888.
- [9] H. Schlosberg, “Three dimensions of emotion.” *Psychological review*, vol. 61, no. 2, p. 81, 1954.
- [10] A. Mehrabian, *Basic dimensions for a general psychological theory implications for personality, social, environmental, and developmental studies*. Cambridge, 1980.
- [11] J. A. Russell, “A circumplex model of affect.” *Journal of personality and social psychology*, vol. 39, no. 6, 1980.
- [12] P. Ekman, “An argument for basic emotions,” *Cognition & emotion*, vol. 6, no. 3-4, pp. 169–200, 1992.
- [13] R. S. Lazarus and B. N. Lazarus, *Passion and reason: Making sense of our emotions*. Oxford University Press, USA, 1996.
- [14] E. Cambria, A. Livingstone, and A. Hussain, “The hourglass of emotions,” in *Cognitive behavioural systems*. Springer, 2012, pp. 144–157.
- [15] J. Diehl-Schmid, C. Pohl, C. Ruprecht, S. Wagenpfeil, H. Foerstl, and A. Kurz, “The Ekman 60 faces test as a diagnostic instrument in frontotemporal dementia,” *Archives of Clinical Neuropsychology*, vol. 22, no. 4, pp. 459–464, 2007.
- [16] C. Westbury, J. Keith, B. B. Briesemeister, M. J. Hofmann, and A. M. Jacobs, “Avoid violence, rioting, and outrage; approach celebration, delight, and strength: Using large text corpora to compute valence, arousal, and the basic emotions,” *The Quarterly Journal of Experimental Psychology*, vol. 68, no. 8, pp. 1599–1622, 2015.
- [17] H. Aviezer, R. R. Hassin, J. Ryan, C. Grady, J. Susskind, A. Anderson, M. Moscovitch, and S. Bentin, “Angry, disgusted, or afraid? studies on the malleability of emotion perception,” *Psychological science*, vol. 19, no. 7, pp. 724–732, 2008.
- [18] J. D. Morris, “SAM: the self-assessment manikin. an efficient cross-cultural measurement of emotional response,” *Journal of advertising research*, vol. 35, no. 6, pp. 63–69, 1995.
- [19] B. Nevo, “Face validity revisited,” *Journal of Educational Measurement*, vol. 22, no. 4, pp. 287–293, 1985.
- [20] H. Martinez, G. Yannakakis, and J. Hallam, “Don’t classify ratings of affect; rank them!” *IEEE transactions on Affective Computing*, vol. 5, no. 3, pp. 314–326, 2014.
- [21] A. Metallinou and S. Narayanan, “Annotation and processing of continuous emotional attributes: Challenges and opportunities,” in *Proceedings of the automatic face and gesture recognition conference*. IEEE, 2013, pp. 1–8.
- [22] H. Gunes and B. Schuller, “Categorical and dimensional affect analysis in continuous input: Current trends and future directions,” *Image and Vision Computing*, vol. 31, no. 2, pp. 120–136, 2013.
- [23] H. Helson, *Adaptation-level theory: an experimental and systematic approach to behavior*. Harper and Row: New York, 1964.
- [24] R. L. Solomon and J. D. Corbit, “An opponent-process theory of motivation: I. temporal dynamics of affect.” *Psychological Review*, vol. 81, no. 2, p. 119, 1974.
- [25] A. R. Damasio, *Descartes error: Emotion, rationality and the human brain*. New York: Putnam, 1994.
- [26] G. N. Yannakakis, R. Cowie, and C. Busso, “The ordinal nature of emotions,” in *Proceedings of the Intl. Conference on Affective Computing and Intelligent Interaction*. IEEE, 2017, pp. 248–255.
- [27] —, “The ordinal nature of emotions: An emerging approach,” *IEEE Transactions on Affective Computing (Early Access)*, 2018.
- [28] B. Seymour and S. M. McClure, “Anchors, scales and the relative coding of value in the brain,” *Current opinion in neurobiology*, vol. 18, no. 2, pp. 173–178, 2008.
- [29] A. Tversky and D. Kahneman, “The framing of decisions and the psychology of choice,” *Science*, vol. 211, no. 4481, pp. 453–458, 1981.
- [30] S. Erk, M. Kiefer, J. Grothe, A. P. Wunderlich, M. Spitzer, and H. Walter, “Emotional context modulates subsequent memory effect,” *Neuroimage*, vol. 18, no. 2, pp. 439–447, 2003.
- [31] E. Camilleri, G. N. Yannakakis, and A. Liapis, “Towards general models of player affect,” in *Proceedings of the Intl. Conference on Affective Computing and Intelligent Interaction*. IEEE, 2017, pp. 333–339.
- [32] Y. Baveye, E. Dellandréa, C. Chamaret, and L. Chen, “Deep learning vs. kernel methods: Performance for emotion prediction in videos,” in *Proceedings of the Intl. Conference on Affective Computing and Intelligent Interaction*. IEEE, 2015, pp. 77–83.
- [33] P. M. Müller, S. Amin, P. Verma, M. Andriluka, and A. Bulling, “Emotion recognition from embedded bodily expressions and speech during dyadic interactions,” in *2015 Intl. Conference on Affective Computing and Intelligent Interaction*. IEEE, 2015, pp. 663–669.
- [34] E. Dellandréa, L. Chen, Y. Baveye, M. V. Sjöberg, C. Chamaret *et al.*, “The mediaeval 2016 emotional impact of movies task,” in *CEUR Workshop Proceedings*, 2016.
- [35] S. Dhamija and T. E. Boulton, “Automated action units vs. expert raters: Face off,” in *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018, pp. 259–268.
- [36] G. McKeown, M. Valstar, R. Cowie, M. Pantic, and M. Schroder, “The SEMAINE database: Annotated multimodal records of emotionally colored conversations between a person and a limited agent,” *IEEE Transactions on Affective Computing*, vol. 3, no. 1, pp. 5–17, 2012.
- [37] K. Krippendorff, “Reliability in content analysis,” *Human communication research*, vol. 30, no. 3, pp. 411–433, 2004.
- [38] —, “Computing Krippendorff’s alpha-reliability,” https://repository.upenn.edu/asc_papers/43, 2011, accessed 15 April 2019.
- [39] G. N. Yannakakis and H. P. Martínez, “Ratings are overrated!” *Frontiers in ICT*, vol. 2, p. 13, 2015.
- [40] M. Jaiswal, Z. Aldeneh, C.-P. Bara, Y. Luo, M. Burzo, R. Mihalcea, and E. M. Provost, “Muse-ing on the impact of utterance ordering on crowdsourced emotion annotations,” *arXiv preprint arXiv:1903.11672*, 2019.
- [41] C. Wang, P. Lopes, T. Pun, and G. Chanel, “Towards a better gold standard: Denoising and modelling continuous emotion annotations based on feature agglomeration and outlier regularisation,” in *Proceedings of the Audio/Visual Emotion Challenge*. ACM, 2018, pp. 73–81.
- [42] A. Burmania, S. Parthasarathy, and C. Busso, “Increasing the reliability of crowdsourcing evaluations using online quality assessment,” *IEEE Transactions on Affective Computing*, vol. 7, no. 4, pp. 374–388, 2016.
- [43] M. Soleymani and M. Larson, “Crowdsourcing for affective annotation of video: Development of a viewer-reported boredom corpus,” in *Proceedings of the SIGIR Workshop on Crowdsourcing for Search Evaluation*, 2010.
- [44] V. E. Farrugia, H. P. Martínez, and G. N. Yannakakis, “The preference learning toolbox,” *arXiv preprint arXiv:1506.01709*, 2015.
- [45] E. Camilleri, G. N. Yannakakis, D. Melhart, and A. Liapis, “PyPLT: Python Preference Learning Toolbox,” in *Proceedings of the Intl. Conference on Affective Computing and Intelligent Interaction*. IEEE, 2019.