



OPEN

# GameVibe: a multimodal affective game corpus

DATA DESCRIPTOR

Matthew Barthet<sup>1,2</sup>✉, Maria Kaselimi<sup>1,2</sup>, Kosmas Pinitas<sup>1,2</sup>, Konstantinos Makantasis<sup>1,2</sup>, Antonios Liapis<sup>1,2</sup> & Georgios N. Yannakakis<sup>1,2</sup>

As online video and streaming platforms continue to grow, affective computing research has undergone a shift towards more complex studies involving multiple modalities. However, there is still a lack of readily available datasets with high-quality audiovisual stimuli. In this paper, we present GameVibe, a novel affect corpus which consists of multimodal audiovisual stimuli, including in-game behavioural observations and third-person affect traces for viewer engagement. The corpus consists of videos from a diverse set of publicly available gameplay sessions across 30 games, with particular attention to ensure high-quality stimuli with good audiovisual and gameplay diversity. Furthermore, we present an analysis on the reliability of the annotators in terms of inter-annotator agreement.

## Background & Summary

Affective Computing (AC) is the interdisciplinary field that refers to the study of human emotions and the development of tools and technologies that learn, interpret and perceive these emotions<sup>1</sup>. The availability of large-scale corpora comprising affect manifestations elicited through appropriate stimuli is critical for AC. However, manifestations of affect are often highly subjective and difficult to assess; each individual's interpretation of a stimulus is influenced by factors such as preferences, memory, biases, systematic errors, and expectations<sup>2,3</sup>. Furthermore, emotions are a dynamic phenomenon and participants' reactions to similar stimuli may shift over time. Therefore, as AC advances and deep learning methods require an increasing amount of data and scale, it becomes important to design and implement experimental protocols that maximise the reliability of collected labels of large-scale affect corpora.

Video games are a hugely popular type of digital media, offering a unique form of human-computer-interaction (HCI) due to their rich content and interactive nature<sup>4</sup>. The global video game industry has seen substantial expansion and is already among the fastest growing subsectors within the broader entertainment sector<sup>5</sup>. Thus, using games as a test-bed to reveal the intricacies of the HCI loop is one of the most promising ways to analyse human behaviour and experience, at scale. Studying the behaviour of players and the experience of play can, in turn, be used to improve the technical aspects of game development, and to contribute to more engaging and personalised game experiences<sup>6,7</sup>.

The literature already contains a growing number of diverse affect corpora that vary in terms of modalities of user input considered, annotation methods, affect stimuli, and number of participants among other factors. Some notable examples of modalities recorded include audiovisual data<sup>8,9</sup>, physiological signals<sup>10</sup>, and facial expression data<sup>11</sup>. *First-person* annotation protocols ask the annotator to label their own experience, and have been employed in several affect corpora, including MazeBall<sup>12</sup>, PED<sup>13</sup>, FUNii<sup>14</sup>, and MUMBAI<sup>9</sup>. When employing a *third-person* protocol, instead, we ask annotators to label the experience of another person, such as in the RECOLA<sup>15</sup>, LIRIS-ACCEDE<sup>16</sup>, AFF-Wild<sup>10</sup>, AffectNet<sup>11</sup>, and SEWA<sup>17</sup> corpora. The provided labels can be *discrete* (e.g. categories, scales) such as in GAME-ON<sup>8</sup> and BIRAFFE2<sup>18</sup> or *continuous* traces as in RAGA<sup>19</sup> and MUMBAI<sup>9</sup>.

Contemporary affect corpora are gradually deviating from the controlled setting of in-lab experiments<sup>15,20,21</sup> to real-life (or *in-the-wild*) scenarios<sup>11,18,22</sup> in an attempt to elicit more natural user behaviours and manifestations of affect. Hiring annotators from crowdsourcing platforms has also been gaining interest in an attempt for AC to scale. Untrained crowd workers in uncontrolled settings, however, can be unreliable, resulting in corpora of questionable, or even limited, value<sup>23</sup>. It is fair to say that all affect corpora are characterised by subjectivity<sup>24–26</sup>, since consensus among annotators is not necessarily required. Whilst some studies attempt to mitigate

<sup>1</sup>Institute of Digital Games, University of Malta, Msida, Malta. <sup>2</sup>These authors contributed equally: Matthew Barthet, Maria Kaselimi, Kosmas Pinitas, Konstantinos Makantasis, Antonios Liapis, Georgios N. Yannakakis. ✉e-mail: [matthew.barthet@um.edu.mt](mailto:matthew.barthet@um.edu.mt)

this issue of variability (e.g. by using mood induction<sup>27</sup>) for better performance in downstream tasks, doing so may hamper the generalizability and reliability of results when deployed in uncontrolled real-life scenarios. Motivated by these issues, we seek to provide an affect corpus with validated annotator reliability and sufficient diversity in the stimuli to promote generalizability in downstream tasks. We aim to accomplish this by following a quality assurance protocol<sup>23</sup>, and collect stimuli which present rich contextual variety within a single domain to promote deeper research into the generalizability of affect predictions.

Regarding games as affect stimuli, the players' (or viewers') perceptions of a gameplay video as stimulus are inextricably linked to the game genre, the form of interfacing, the game's objective, the number of players, and potential social aspects<sup>4</sup>. We specifically focus on the First-Person Shooter (FPS) genre of games in this paper, for many reasons. FPS games have been popular among players since the 1990s<sup>28</sup>. Moreover, FPS games usually have high-quality graphics and audio, and a rich variability of stimuli due to the vast number of FPS games developed over the years. FPS games tend to be highly stimulating for both players and viewers, and share many fundamental gameplay elements whilst simultaneously employing very different styles (in terms of both art and gameplay). In addition, the FPS genre has been hugely popular on live-streaming platforms such as Twitch for several years. Indicatively, during December 2023, 4 of the top 10 live-streamed categories were FPS games, with over 220 million viewer hours watched in just 30 days (<https://www.twitchmetrics.net/games/viewership>, accessed on 16/12/2023). We argue that collectively these factors make FPS games the ideal genre for studying HCI, especially for understanding viewer engagement. This makes research in FPS affect modelling highly useful for content creators, streamers, and game designers.

With this in mind, we introduce *GameVibe*, a novel multimodal affect corpus of viewer engagement for FPS game videos. This corpus consists of 2 hours of high-quality audio and visual data from 30 different FPS games extracted and curated from publicly available "Let's Play" videos on YouTube. Data collection involved 20 annotators, consisting of trained researchers, as well as postgraduate and undergraduate students. Affect traces were provided in the form of unbounded, time continuous signals using the RankTrace annotation tool<sup>6</sup> on the PAGAN platform<sup>29</sup>. As part of the corpus, we provide the raw videos used as stimuli during data collection, as well as latent representations extracted using pretrained foundation models for visuals (VideoMAE<sup>30</sup> and MVD<sup>31</sup>), and audio (BEATS<sup>32</sup>) for use in downstream tasks. Furthermore, we include quality assurance data on each annotator to give a better understanding of the reliability and validity of the annotations in the dataset. Finally, we test the reliability and validity of the dataset in terms of inter-annotator agreement.

Our in-lab data collection protocol and annotator quality assurance tools ensure that the *GameVibe* corpus consists of annotations which reliably approximate the ground truth. This corpus can be applied to a variety of downstream tasks, such as modelling engagement using the provided latent representations extracted from pre-trained models for vision and audio. Initial studies using *GameVibe* have shown that models can reliably predict viewer engagement in unseen clips when trained on clips of the same game<sup>33</sup>, which can be useful for game developers to assess the quality of new content. Another downstream task which benefits multimodal affect modelling research tackles the generalizability of affect models<sup>34</sup> through the diversity in stimuli of this corpus. Our generalizability study on building engagement models which predict labels in unseen games using few-shot learning has shown promising results<sup>35</sup>, and marks further progress towards the ultimate goal of general models of affect. The release of this corpus is expected to contribute both in engagement modelling and multimodal interaction research more broadly than these initial studies.

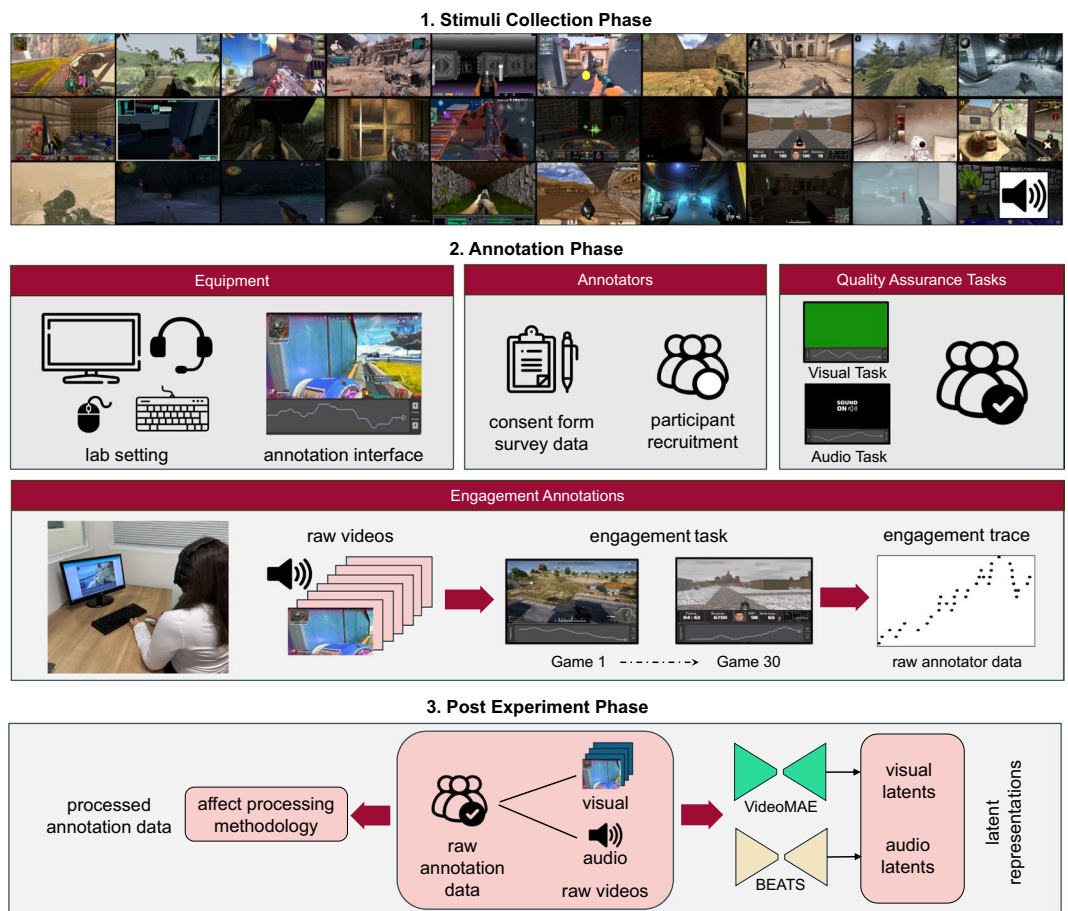
## Methods

Measuring the experience of digital game enjoyment and accurately identifying which game elements engage players are important goals for both the study of user experience in games and the development of better games<sup>4</sup>. In this section, we introduce the FPS affect corpus *GameVibe*, which was solicited to provide a rich, multimodal dataset of elicitors for viewer engagement. The collected dataset includes: (a) synchronised frames and audio per game, (b) extracted latent representations from audiovisual data, (c) engagement annotation traces per participant (in raw & processed form), (d) participants' replies to demographic surveys (anonymised data). The novelty of the dataset lies in the fact that engagement annotations are user-specific; by exploiting the different opinions and perspectives of different users for the same game, *GameVibe* enables the training of fair affect models that are robust to new users and new games.

In this section, we provide a detailed description of the process we followed to build *GameVibe*. The core phases, as illustrated in Fig. 1, are as follows:

- **Design phase:** we outline the objectives, tools, methodologies, and parameters of the study.
- **Stimuli collection phase:** we choose and curate video content from diverse FPS games.
- **Annotation phase:** before collecting annotations of viewers' emotional states, we recruit participants, procure and prepare the equipment, and develop informed consent forms to uphold ethical standards. In addition, we conduct Quality Assurance tests to measure participants' reliability. We detail all the above below.
- **Post-experiment phase:** we assess the quality of the dataset, process and analyse the collected data and document our findings. This ensures the dataset's efficacy and utility for subsequent research endeavours.

**Design phase.** This phase involves defining the scope and purpose of the dataset, determining the specific types of data to be collected, establishing data sources and acquisition methods, and devising a structure for organizing and storing the data. Additionally, factors such as data format, granularity, quality, and potential biases were considered. Beyond research standards, collaboration with domain experts from the game industry is essential to understand requirements and ensure that the dataset aligns with intended use cases. In this work, we draw from our experiences during our long-term collaboration with *Ubisoft's Massive Entertainment*, with whom



**Fig. 1** High-level overview of the experimental protocol used for data collection in this study.

we co-designed an effective experimental protocol for collecting annotations and training affect models in *Tom Clancy's The Division 2*<sup>26</sup>. Following similar practices, we ensure that our dataset closely aligns with both research and industry goals. Furthermore, privacy and ethical considerations were carefully weighed to safeguard sensitive information and uphold ethical standards. In this study, we exploit already tested and reliable tools—based on self-reporting—to measure a game's engagement level as perceived by different annotators. We also rely on well-established methodologies for gathering emotion-related datasets, specifically the PAGAN data collection framework<sup>29</sup> and the RankTrace annotation tool<sup>6</sup>. Furthermore, our core objective was to solicit reliable engagement annotations across a wide variety of contexts to allow for further research into the generalizability of affect models across multiple annotators and stimuli.

**Stimuli collection phase.** As discussed in Background & Summary, we chose 30 FPS games as affect stimuli in order to encompass the widest possible range of audiovisual styles and design aspects (including modes of gameplay and winning conditions). In terms of *game style*, we distinguish between games featuring a “realistic” vs. “stylized” art style. In terms of *game era*, we pick games of both “modern” and “retro” styles (Fig. 2). For example, a game such as *Overwatch 2* (game 23 in Fig. 2) falls under the modern, stylised art style, whereas *Wolfenstein* (game 28 in Fig. 2) would fall under the retro gameplay with a realistic art style. For distinguishing across *game design* patterns, we made an effort to select games with different game modes or “sub-genres” within the FPS space, such as “Battle Royale” (e.g. *PUBG*), “death match” (e.g. *Counter Strike*), and “single-player” (e.g. *Doom*) games. The distribution of these design patterns is visualised in Fig. 2, where we can see single player games make up 67% of the stimuli selected, with death match and Battle Royale making up 20% and 13% respectively. We hope that the richness in variety and multimodality of the dataset will empower affect models to generalise better to unseen games when predicting viewer engagement, and help us understand how game design and graphical style impact human responses to these stimuli.

Within each game, when picking the videos to be included in the dataset, we ensured that the content consisted of primarily gameplay, meaning that there are no videos with more than 15 seconds of non-gameplay footage (e.g. menu screens, cut scenes or transition animations). Furthermore, we made sure to only use videos with audio consisting solely of in-game sounds and that did not contain any player or streamer commentary. Based on these criteria, we acquired the videos for each of our chosen FPS games from “Let's Play” videos on YouTube, ensuring the videos were at least 5 minutes long in order to have enough content to extract multiple short videos



post-survey questionnaire (see below) included personal data, we preserved anonymity by using IDs rather than annotator names and ensured that such questions could not be used to identify the participants.

**Survey data.** At the end of the annotation experiment, participants were required to fill in a short survey on their demographic information (age range, ethnicity, gender, handedness, education level), familiarity with video games, familiarity with FPS games, familiarity with video annotation as well as their favourite game. Most participants were between 25 and 35 years old (42%), while 32% were between 18 and 25 years and 21% were between 35 and 45 years old; 1 participant was over 45 years old. Almost all participants (97.4%) were Caucasian, and the vast majority (89.5%) were right-handed. Furthermore, 75% of participants were male, 20% were female and 5% identified as non-binary. In terms of familiarity with games or video annotation, most participants were very familiar with games (average Likert score 4.25 out of 5) and slightly less so with FPS games (average Likert score 3.55 out of 5). Most participants were not very familiar with video annotation of affect data (average Likert score 2.65 out of 5), although some outliers existed (6 annotators rated their familiarity at 4 or 5 on a 5-point Likert scale). Finally, participants mentioned many different favourite games, with 7 out of 20 participants mentioning games which were first-person or third-person shooters.

**Annotation protocol and tools.** As mentioned, each participant was assigned to a session and asked to annotate their engagement for a randomised sequence of 30 videos. The random video order was imposed to minimise carry-over effects between stimuli, as has been established in the literature<sup>37</sup>. Annotation tasks were carried out using the PAGAN annotation platform<sup>29</sup> using the RankTrace annotation interface<sup>6</sup>. RankTrace allows for annotation in a continuous and unbounded fashion, and has been used extensively to collect reliable ground truths in an ordinal fashion<sup>3</sup> (see equipment in Fig. 1).


Before annotating gameplay videos, participants were tasked with annotating two Quality Assurance (QA) tasks, which were set up as independent PAGAN projects and linked together. Such QA tasks measure the ability to annotate a simple objective task where the ground truth signal is known. Such QA tasks have been proven to reliably predict annotator reliability in a previous study which used the first two sessions of this dataset<sup>23</sup> as well as other stimuli<sup>38</sup>. The first task is a visual task which requires the participant to annotate their perceived changes in brightness of a video showing a green screen, inspired by a previous study<sup>39</sup>. The second task is an auditory task, which requires participants to annotate their perceived changes in pitch whilst listening to an oscillating sound wave. These two stimuli are included in the dataset along with their respective annotations to provide a measure of our annotators' reliability on simple objective tasks, and by extension a better insight into their reliability in our engagement study.

When participants were ready to annotate gameplay videos, they were provided the following definition of viewer engagement: "A high level of engagement is associated with a feeling of tension, excitement, and readiness. A low level of engagement is associated with boredom, low interest, and disassociation with the game". We emphasised that participants annotated in a *first-person* manner, i.e. they annotated their own engagement as a viewer and not their perceived engagement of the player. Based on this definition, participants were instructed to scroll up on the mouse wheel whenever they felt an increase in their engagement, down when they felt a decrease, and to remain idle (i.e. no scrolling) if their engagement remained unchanged. Since RankTrace annotations are unbounded, mouse scroll wheel offers an ideal—and intuitive—interface to collect such data.

**Collected annotation data.** As shown in Fig. 3, the dataset consists of 4 sessions of annotated footage, producing a corpus of gameplay videos of a total duration of 2 hours. Each of the 4 sessions contains 30 different 1-minute videos (i.e. one video from each of the 30 game titles), which are presented to the annotator in random order. This means that for each session we collect 150 annotation traces, with every video being annotated by 5 participants. In total, this amounts to 600 engagement annotation traces for 120 gameplay videos (of 30 different game titles), annotated by 20 participants. The data are contained in the dataset repository<sup>40</sup>, detailed under Data Records.

**Quality assurance and reproducibility of the GameVibe affect annotations.** The dataset was created with high-quality videos from numerous FPS games, ordered randomly. The random order is commonly used to prevent bias from ordering effects<sup>41</sup>, while the varied short video stimuli minimise user fatigue. The same room, equipment and light conditions were kept during the process for consistency of the annotation experience and data collection instruments (i.e. mouse wheel). Instructions were provided to the annotators to avoid gross errors during the annotation process. Finally, the dataset includes data from two QA tasks completed by annotators prior to the engagement study. Performance of annotators in such QA tests could be used in future studies to filter inconsistent annotators; however, in this paper, we include all annotators in the Technical Validation below.

**Post-experiment phase. Audiovisual stimuli data processing.** In this study, we exploit large-scale pre-trained models to extract latent representations in visual data using (a) Video Masked Autoencoders (VMA)<sup>30</sup> and (b) Masked Video Distillation (MVD)<sup>31</sup> models for visual representation learning. VMA is a model for robust representation learning for video data by utilising a masked autoencoder architecture<sup>30</sup>. MVD provides efficient knowledge distillation by leveraging the principles of masked attention mechanisms to distil knowledge from a teacher network to a student network<sup>31</sup>. For audio data, we leverage the BEATS model proposed by Chen *et al.*<sup>32</sup>, which serves as a self-supervised architecture designed for audio processing tasks. Given that FPS games have a high representation fidelity (regardless of art style) with the real world, we hypothesised that the above pretrained models can identify and classify different visual elements or sounds present within the audiovisual stimuli across different FPS games.



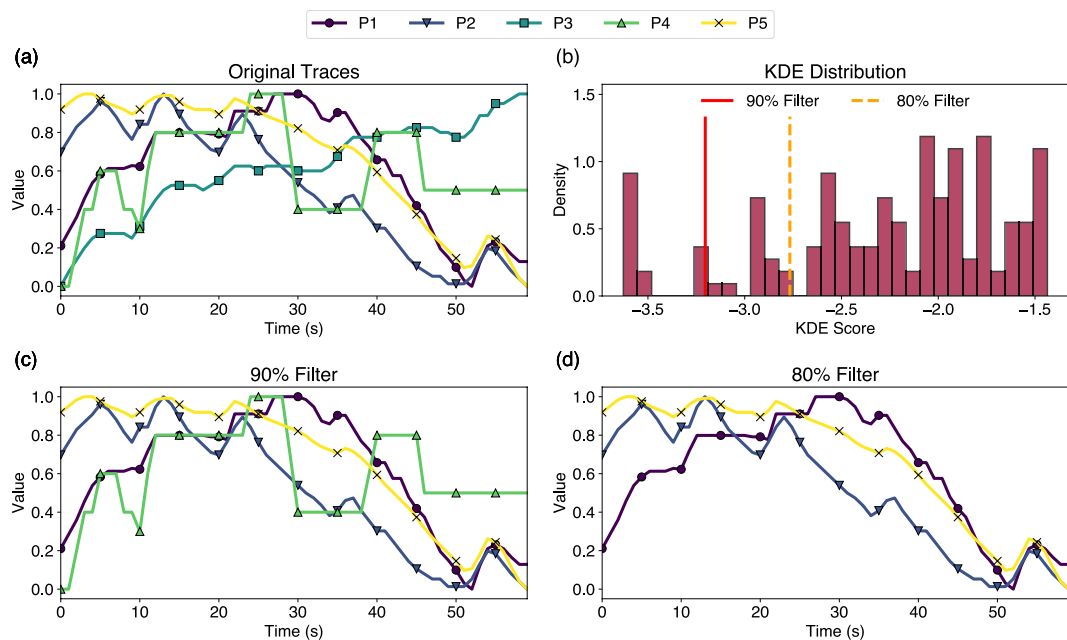
**Fig. 3** GameVibe's post-processing pipeline, visualising the transformed annotations and the output files of each stage.

**Affect processing methodology.** The raw discrete values that indicate changes in engagement are generated from each PAGAN project and require a number of pre-processing steps before they are suitable for use in downstream tasks. Raw data from PAGAN is a list of values paired with a timestamp for when that value was entered by the participant (after scrolling on the mouse wheel), stored in a CSV file for each session. This data is converted into a time continuous signal by interpolating the values into a series of 1-second time windows (see Step 1 in Fig. 3). We also ensure that the annotator has made at least one input to the signal to be considered a valid signal; flat signals (with no changes) are discarded from the dataset, however no such instances were recorded. Finally, we combine the data into a single Python *pickle* file, detailed under Data Records.

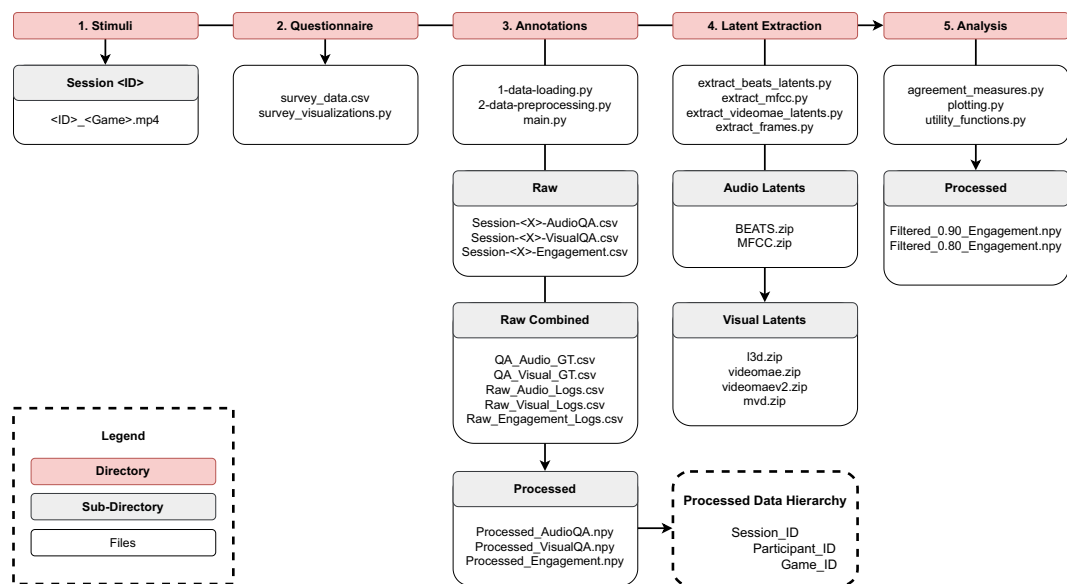
As RankTrace<sup>6</sup> was used to record engagement using a mouse scroll wheel, the raw data collected are unbounded (i.e., have no minimum or maximum values). In order to facilitate comparison across annotators, we include a normalised version of our dataset where engagement is normalised to [0,1] via Min-Max normalisation on a per-trace basis (see Step 2 in Fig. 3). In our processing library, we also include the option to perform moving average smoothing on the signals in the dataset, or to change the size of the time window of interpolated signals. We selected a 1-second resampling rate, considering the trade-off between increasing data volume and detecting variations in engagement. Larger time windows or moving average filters would result in a smoother signal at the cost of signal accuracy and resolution. We offer the raw annotation traces from PAGAN in the dataset (see below) for further experiments with other processing methods. The balance of the collected data is highly dependent on said processing methods, and therefore we do not conduct an in-depth study on the engagement distributions for each stimulus. As an indication of the label balance, the corpus consists of 15.7% increases, 13.4% decreases and 70.9% stable engagement levels when the raw data is processed into ordinal signals<sup>42</sup> using 1 second time windows.

Since this study handles subjective annotations, where no concrete ground truth can ever be identified, two strategies are common to establish a consensus. One strategy is to embrace the potentially high variance in the dataset and use this information to identify regions of high vs low agreement<sup>43</sup>, or even model it as part of the training process to improve performance<sup>44</sup>. Another strategy is to empirically analyse the data and use established techniques such as outlier detection<sup>45</sup> to remove problematic annotators, or factor in annotator consensus<sup>46</sup>. We follow the latter strategy, detecting and removing outliers in the dataset to improve the quality of the ground truth signal derived in any downstream tasks (see Step 3 in Fig. 3).

For filtering outliers in each video in a session (see Step 4 in Fig. 4), we use Dynamic Time Warping<sup>47</sup> (DTW) to create a distance matrix of the normalised annotation signals from each participant. We chose DTW as a distance measure due to its proven use in similar studies<sup>48</sup> and its ability to focus on comparing the shape of the signals (i.e. an ordinal distance measure) whilst also factoring out time shifts between participants due to different lags in their annotation<sup>49</sup>. For example, for the video of *Wolfenstein 3D* in Session 1, we calculate the pairwise DTW distances between all possible pairs of annotators and assemble them into a  $5 \times 5$  DTW distance matrix. We use the minimum distance to the nearest participant to judge whether a participant is a singular outlier (and should therefore be removed), or part of a cluster of two or more annotators (and should therefore remain in the dataset). This process involves setting a threshold to determine what will be considered an in/outlier. Rather than using a static threshold which would require hand-picking for each session, we use the non-parametric Kernel Density Estimation (KDE) for outlier removal to create a threshold, which is visualised in Fig. 4. This involves calculating all the pairwise KDE scores between signals for all videos in a session and sorting them in ascending order. We then test two filters, a strict filter which only includes signals with a nearest-neighbour distance within the best 80% of KDE scores, and a more relaxed filter which includes signals with a nearest-neighbour distance within the best 90% of KDE scores. Remaining signals are considered outliers and are removed: examples of removed outliers for *Wolfenstein 3D* are shown in Fig. 4.



**Fig. 4** An example of outlier detection using the game *Wolfenstein 3D* from Session 1. (a) depicts the unfiltered traces from five annotators. The frequency distribution can be seen in (b), depicting the signals' KDE score used to create the outlier filters. The red and orange lines depict the 90% filter and 80% filter thresholds, respectively. Annotators to the right of both filters are considered inliers and are not removed. (c) shows the first filter applied removing one outlier (P3), and (d) shows the strictest filter applied removing two outliers (P3, P4).



**Fig. 5** Diagram overview of the dataset and its structure.

### Data Records

The dataset and files described in this section are publicly available and can be found on our Open Science Framework repository<sup>40</sup>. We provide annotation data in CSV format and the respective video stimuli in *mp4* format: both formats can be processed in most software packages or programming language. We provide Python files for data processing and feature extraction. The `GameVibe_README.md` file details the organisation of the dataset, explaining the structure, naming convention, and specific contents of each file.

The dataset is split into five directories as depicted in Fig. 5. The “*Stimuli*” directory contains the 1-minute game videos organised into subdirectories by session and named according to the game in each video. The “*Questionnaire*” directory contains participants' responses to the post-experiment questionnaire in CSV format, and a Python script to process the responses and create visualisations. The “*Annotations*” directory contains the annotation data returned from PAGAN and the Python scripts required to process them; this folder has several

Session 1					Session 2					Session 3					Session 4								
	P1	P2	P3	P4	P5		P6	P7	P8	P9	P10		P11	P12	P13	P14	P15		P16	P17	P18	P19	P20
P1		4.1	8.9	8.5	4.3	P6		10.8	10.9	12.5	10.3	P11		8.7	10.8	8.0	9.9	P16		7.3	9.4	8.9	9.4
P2	4.1		9.1	7.8	3.8	P7	10.8		6.4	8.2	10.3	P12	8.7		8.1	8.8	8.8	P17	7.3		12.2	12.8	11.3
P3	8.9	9.1		5.1	11.2	P8	10.9	6.4		9.8	10.2	P13	10.8	8.1		11.7	11.4	P18	9.4	12.2		10.5	8.6
P4	8.5	7.8	5.1		9.2	P9	12.5	8.2	9.8		9.6	P14	8.0	8.8	11.7		9.6	P19	8.9	12.8	10.5		6.3
P5	4.3	3.8	11.2	9.2		P10	10.3	10.3	10.2	9.6		P15	9.9	8.8	11.4	9.6		P20	9.4	11.3	8.6	6.3	
Avg.	6.4	6.2	8.6	7.6	7.1	Avg.	11.1	8.9	9.3	10.0	10.1	Avg.	9.4	8.6	10.5	9.5	9.9	Avg.	8.8	10.9	10.2	9.6	8.9

**Table 1.** DTW distance matrices between all participants in the same session, averaged across 30 video stimuli per session.

subdirectories of importance. In the “Raw” subdirectory, the raw annotations outputted by PAGAN for each session can be found. The “Raw Combined” subdirectory contains the raw data assembled into a single CSV for each task, and the ground truth signals for the QA test (see Methods). The “Processed” subdirectory contains the processed annotation data for the QA tasks and the Engagement task for all sessions in NumPy file format. These processed files consist of a Python dictionary containing the following hierarchy: data is sorted into sessions, which in turn are sorted into participants, and further sorted by game where the participants’ annotations can be found. The “Latent Extraction” subdirectory contains Python scripts to extract the latents from the audiovisual data found in the “Stimuli” directory. We also include examples of extracted latents for a 3-second time window that can be used for further analysis. Finally, the “Analysis” directory contains the scripts used to conduct the analysis in this paper, and includes a “Processed” subdirectory with annotation data with outliers removed (see Methods).

## Technical Validation

**Annotator quality assurance.** We first illustrate the reliability of our annotators by computing their average QA scores across our visual and auditory QA tasks (see Methods). Since both QA tasks involve an annotation task where the ground truth (screen brightness and audio frequency) is known in advance, we use the Signed Differential Agreement (SDA) metric<sup>39</sup> to measure similarity. We chose SDA as it has been proven effective in our previous QA study<sup>23</sup> and, due to its bounded nature between  $-1$  and  $1$ , is an intuitive metric for selecting a filter threshold. Across all participants and both QA tasks, the average SDA score was  $0.43 \pm 0.15$ , meaning that annotators correctly annotated 71.5% of all time windows among QA signals. If we used a QA filtering threshold value similar to previous work<sup>23</sup> (SDA = 0, or 50% accuracy), only 3 annotators (P6 and P7 from Session 2, P17 from Session 4) would be removed due to poor QA performance. This verifies that almost all our annotators understood the annotation process and could accurately produce annotations before moving on to the engagement task. For the purposes of the technical validation of this paper, however, no participants are removed through this QA filtering in order to assess all raw data included in the dataset.

**Inter-annotator agreement.** We use DTW (see Methods) to create distance matrices for each video in each session, in order to calculate the inter-annotator agreement between annotators’ normalised engagement traces on the same stimuli. We average those DTW values between pairs of participants on a per-session basis, since the same participant with the same identifier (e.g. P1) annotated every video in the same session. We report the average DTW distance across all stimuli (30) per session in Table 1, using different participant identifiers because different participants annotated different sessions. We observe that there are differences between annotators; only annotators in Session 1 were more in agreement. However, annotators are rarely consistently disagreeing with all other annotators; the most obvious instances of this are P6 (with DTW distances over 10 with all other annotators in Session 2) and P17 (with DTW distances over 11 with three other annotators in Session 4). We note that both annotators performed poorly in the QA tasks as well, and would have been removed via an SDA cut-off on the two QA tasks (see above). In general, there is a statistically significant ( $p < 0.05$ ) negative Pearson correlation between average DTW distance of each participant’s engagement traces with all others in the same session and SDA score averaged from the two QA tasks ( $r = -0.56$ ). This is a promising finding for the quality assurance offered by the additional QA tasks, as the closer the annotators are to the known ground truth in those two tasks (an SDA score near 1), the less likely they are to disagree with other annotators in the more challenging engagement annotation task.

**Outliers per participant.** We detect outliers on the engagement traces (see Methods) based on the KDE of all DTW matrices of all videos in the same session, and report results with two filters for outlier removal. When using the 90% filter (see Methods), 60 out of 600 annotations are removed (15 per session), whilst with the 80% filter 120 annotations are removed (30 per session). Interestingly, while P7 would be removed due to poor QA task performance, none of their 30 traces are removed as outliers for the 90% filter and only 2 traces are removed as outliers for the 80% filter. On the other hand, P17 has the most outliers of all 20 participants: 11 of their 30 traces are removed as outliers for the 90% filter and an extraordinary 16 out of 30 removed for the 80% filter. Other participants with many removed outliers are P4 (with 9 outliers for 80% filter) and P3, P13, P20 (each with 8 outliers for 80% filter). The average SDA score from the two QA tasks is not significantly correlated ( $p > 0.05$ ) with the number of outliers removed:  $r = -0.33$  for removed outliers per participant with the 90% filter and  $r = -0.25$  for outliers with the 80% filter.

**Outliers per game and video.** It is worthwhile to observe which game stimuli were more prone to inter-annotator disagreements, which would lead to more outliers removed through the above filtering process. After applying the 90% filter, outliers were removed from 26 out of the 30 games. The games with no outliers were

*Apex Legends*, *Wolfem Wolfenstein*, *Medal of Honour* (PS1) and *Superhot*. The games with the most outliers were *Outlaws*, *Counter-Strike 16*, and *Wolfenstein 3D* with 4 outliers each (out of 20 annotation traces across 4 videos of the same game). After applying the 80% filter, outliers were removed from every game except *Apex Legends*. The game with the most outliers when using the 80% filter was *Operation Body Count*, with 9 outliers removed (out of 20 traces). The consistency of *Apex Legends* across sessions and annotators is likely due to intertwined factors, such as its highly stylised graphics, strong emphasis on audio cues, and quick well-defined changes in tempo (e.g. alternating between fast-paced team fights and slow-paced periods of healing, looting, exploration, etc.). Finally, observing individual videos with the most outliers after applying the 90% filter, Session 1's *Outlaws* and Session 4's *Counter-Strike 16* had 3 outliers removed each (out of 5 traces each). When using the stricter 80% filter, Session 4's *Superhot* video had all of its annotations classified as outliers and removed, indicating that this video had particularly poor inter-rater agreement compared to others in that session. This aggressive filtering process could evidently prove problematic, as in the case of this video, no data remain for affect modelling. For this reason, the data records contain all raw and processed traces and Python scripts that would allow researchers to perform their own filtering process or apply different KDE thresholds depending on their goals.

### Code availability

The data have been deposited into an Open Science Framework repository<sup>40</sup> and is available for public use. The dataset can be managed, visualised and pre-processed using Python files.

Received: 24 June 2024; Accepted: 22 October 2024;

Published online: 29 November 2024

### References

- Picard, R. W., Vyzas, E. & Healey, J. Toward machine emotional intelligence: Analysis of affective physiological state. *IEEE transactions on pattern analysis and machine intelligence* **23**, 1175–1191 (2001).
- Sciutti, A., Barros, P., Castellano, G. & Nagai, Y. Affective shared perception. *Frontiers in Integrative Neuroscience* **16**, 1024267 (2022).
- Yannakakis, G. N., Cowie, R. & Busso, C. The ordinal nature of emotions: An emerging approach. *IEEE Transactions on Affective Computing* **12**, 16–35 (2018).
- Yannakakis, G. N. & Melhart, D. Affective Game Computing: A Survey. In *In Proceedings of the IEEE* (2023).
- Goh, E., Al-Tabbaa, O. & Khan, Z. Unravelling the complexity of the video game industry: An integrative framework and future research directions. *Telematics and Informatics Reports* 100100 (2023).
- Lopes, P., Yannakakis, G. N. & Liapis, A. RankTrace: Relative and unbounded affect annotation. In *Proceedings of the IEEE International Conference on Affective Computing and Intelligent Interaction*, 158–163 (2017).
- Kotsia, I., Zafeiriou, S. & Fotopoulos, S. Affective gaming: A comprehensive survey. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 663–670 (2013).
- Maman, L. *et al.* Game-on: A multimodal dataset for cohesion and group analysis. *IEEE Access* **8**, 124185–124203 (2020).
- Doyran, M. *et al.* Mumbai: multi-person, multimodal board game affect and interaction analysis dataset. *Journal on Multimodal User Interfaces* (2021).
- Kollias, D. Abaw: Valence-arousal estimation, expression recognition, action unit detection & multi-task learning challenges. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2328–2336 (2022).
- Mollahosseini, A., Hasani, B. & Mahoor, M. H. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing* **10**, 18–31 (2017).
- Yannakakis, G. N., Martinez, H. P. & Jhala, A. Towards affective camera control in games. *Transactions on User Modeling and User-Adapted Interaction* **20**, 313–340 (2010).
- Karpouzis, K., Yannakakis, G. N., Shaker, N. & Asteriadis, S. The platformer experience dataset. In *Proceedings of the IEEE International Conference on Affective Computing and Intelligent Interaction*, 712–718 (2015).
- Beaudoin-Gagnon, N. *et al.* The funii database: A physiological, behavioral, demographic and subjective video game database for affective gaming and player experience research. In *Proceedings of the IEEE International Conference on Affective Computing and Intelligent Interaction* (2019).
- Ringeval, F., Sonderegger, A., Sauer, J. & Lalanne, D. Introducing the recola multimodal corpus of remote collaborative and affective interactions. In *Proceedings of the 10th IEEE International Conference and workshops on automatic face and gesture recognition (FG)* (2013).
- Baveye, Y., Dellandrea, E., Chamaret, C. & Chen, L. Liris-accede: A video database for affective content analysis. *IEEE Transactions on Affective Computing* **6**, 43–55 (2015).
- Kossaiji, J. *et al.* Sewa db: A rich database for audio-visual emotion and sentiment research in the wild. *IEEE transactions on pattern analysis and machine intelligence* **43**, 1022–1040 (2019).
- Kutt, K. *et al.* Biraffe2, a multimodal dataset for emotion-based personalization in rich affective game environments. *Scientific Data* **9**, 274 (2022).
- Granato, M., Gadia, D., Maggiorini, D. & Ripamonti, L. A. An empirical study of players' emotions in vr racing games based on a dataset of physiological data. *Multimedia tools and applications* **79**, 33657–33686 (2020).
- Forgas, J. P., Bower, G. H. & Krantz, S. E. The influence of mood on perceptions of social interactions. *Journal of Experimental Social Psychology* **20**, 497–513 (1984).
- Pinilla, A., Tamayo, R. M. & Neira, J. How do induced affective states bias emotional contagion to faces? a three-dimensional model. *Frontiers in psychology* **11**, 97 (2020).
- Park, C. Y. *et al.* K-emocon, a multimodal sensor dataset for continuous emotion recognition in naturalistic conversations. *Scientific Data* **7**, 293 (2020).
- Barthet, M. *et al.* Knowing your annotator: Rapidly testing the reliability of affect annotation. In *Proceedings of the IEEE International Conference on Affective Computing and Intelligent Interaction Workshops (ACIIW)* (2023).
- Martnez-Miwa, C. A. & Castelán, M. On reliability of annotations in contextual emotion imagery. *Scientific Data* **10**, 538 (2023).
- Miranda-Correa, J. A., Abadi, M. K., Sebe, N. & Patras, I. Amigos: A dataset for affect, personality and mood research on individuals and groups. *IEEE Transactions on Affective Computing* **12**, 479–493 (2018).
- Kollias, D. *et al.* Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *International Journal of Computer Vision* **127**, 907–929 (2019).
- Lench, H. C., Flores, S. A. & Bench, S. W. Discrete emotions predict changes in cognition, judgment, experience, behavior, and physiology: a meta-analysis of experimental emotion elicitations. *Psychological bulletin* **137**, 834 (2011).
- Morris, S. W. bots and mods: Multiplayer fps games as co-creative media. In *Proceedings of DiGRA 2003 Conference: Level Up* (2003).
- Melhart, D., Liapis, A. & Yannakakis, G. N. PAGAN: Video affect annotation made easy. In *Proceedings of the IEEE International Conference on Affective Computing and Intelligent Interaction (ACII)*, 130–136 (2019).

30. Tong, Z., Song, Y., Wang, J. & Wang, L. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems* **35**, 10078–10093 (2022).
31. Wang, R. *et al.* Masked video distillation: Rethinking masked feature modeling for self-supervised video representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6312–6322 (2023).
32. Chen, S. *et al.* Beats: audio pre-training with acoustic tokenizers. In *Proceedings of the 40th International Conference on Machine Learning*, 5178–5193 (2023).
33. Pinitas, K. *et al.* Varying the context to advance affect modelling: A study on game engagement prediction. In *Proceedings of the IEEE International Conference on Affective Computing and Intelligent Interaction (ACII)* (2024).
34. Camilleri, E., Yannakakis, G. N. & Liapis, A. Towards general models of player affect. In *Proceedings of the 7th International Conference on Affective Computing and Intelligent Interaction (ACII)*, 333–339 (2017).
35. Pinitas, K., Makantasis, K. & G. N. Yannakakis. Across-game engagement modelling via few-shot learning. In *Proceedings of the First Workshop on Computer Vision for Videogames, European Conference on Computer Vision (ECCV)* (2024).
36. Pinitas, K. *et al.* Predicting player engagement in tom clancy's the division 2: A multimodal approach via pixels and gamepad actions. In *Proceedings of the 25th International Conference on Multimodal Interaction*, 488–497 (2023).
37. Sharma, K., Castellini, C., Van Den Broek, E. L., Albu-Schaeffer, A. & Schwenker, F. A dataset of continuous affect annotations and physiological signals for emotion analysis. *Scientific data* **6**, 196 (2019).
38. Burmania, A., Parthasarathy, S. & Busso, C. Increasing the reliability of crowdsourcing evaluations using online quality assessment. *IEEE Transactions on Affective Computing* **7**, 374–388 (2015).
39. Booth, B. M. & Narayanan, S. S. Fifty shades of green: Towards a robust measure of inter-annotator agreement for continuous signals. In *Proceedings of the International Conference on Multimodal Interaction*, 204–212 (2020).
40. Barthet, M. *et al.* Gamevibe: a multimodal affective game corpus, <https://doi.org/10.17605/OSF.IO/P4NGX> (2024).
41. D'Mello, S., Kappas, A. & Gratch, J. The affective computing approach to affect measurement. *Emotion Review* **10**, 174–183 (2018).
42. Yannakakis, G. N., Cowie, R. & Busso, C. The ordinal nature of emotions. In *Proceedings of the IEEE International Conference on Affective Computing and Intelligent Interaction*, 248–255 (2017).
43. Girard, J. M., Tie, Y. & Liebenthal, E. Dynamos: The dynamic affective movie clip database for subjectivity analysis. In *Proceedings of the IEEE International Conference on Affective Computing and Intelligent Interaction (ACII)* (2023).
44. Rizos, G. & Schuller, B. Modelling sample informativeness for deep affective computing. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3482–3486 (2019).
45. Boukerche, A., Zheng, L. & Alfandi, O. Outlier detection: Methods, models, and classification. *ACM Computing Surveys (CSUR)* **53**, 1–37 (2020).
46. Parthasarathy, S., Cowie, R. & Busso, C. Using agreement on direction of change to build rank-based emotion classifiers. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **24**, 2108–2121 (2016).
47. Sakoe, H. & Chiba, S. Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing* **26**, 43–49 (1978).
48. Melhart, D., Liapis, A. & Yannakakis, G. N. The arousal video game annotation (again) dataset. *IEEE Transactions on Affective Computing* **13**, 2171–2184 (2022).
49. Mariooryad, S. & Busso, C. Correcting time-continuous emotional labels by modeling the reaction lag of evaluators. *IEEE Transactions on Affective Computing* **6**, 97–108 (2014).

## Acknowledgements

This research is supported by the Foundation.AI (Foundation Modelling of Games via Artificial Intelligence) research project funded by the Ministry for Education, Sport, Youth, Research and Innovation in Malta, as well as by the Malta Council for Science and Technology through the SINO-MALTA Fund 2022, Project OPTiMaL. Makantasis was supported by Project ERICA (GA: REP-2023-36) financed by the Xjenza Malta (Malta Council for Science & Technology - MCST), for and on behalf of the Foundation for Science and Technology, through the FUSION: R&I Research Excellence Programme.

## Author contributions

G.N.Y. conceived the experiments, M.B., M.K. and K.P. conducted the experiments under the supervision of K.M., A.L. and G.N.Y. and M.B. and M.K. performed the analysis of results. All authors discussed the results, and all authors contributed equally to the final manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to M.B.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024