BAH DATASET FOR AMBIVALENCE/HESITANCY RECOGNITION IN VIDEOS FOR BEHAVIOURAL CHANGE

Anonymous authorsPaper under double-blind review

000 001 002

003

004

005

006 007 008

009

010 011 012

013

015

016

017

018

019

021

023

025

026

027

028

029

031

034

037

040

041

043

044

045

047

048

051

052

ABSTRACT

Recognizing complex emotions linked to ambivalence and hesitancy (A/H) can play a critical role in the personalization and effectiveness of digital behaviour change interventions. These subtle and conflicting emotions are manifested by a discord between multiple modalities, such as facial and vocal expressions, and body language. Although experts can be trained to identify A/H, integrating them into digital interventions is costly and less effective. Automatic learning systems provide a cost-effective alternative that can adapt to individual users, and operate seamlessly within real-time, and resource-limited environments. However, there are currently no datasets available for the design of machine/deep learning models to recognize A/H. This paper introduces a first Behavioural Ambivalence/Hesitancy (BAH) dataset collected for subject-based multimodal recognition of A/H in videos. It contains videos from 224 participants captured across nine provinces in Canada, with different age, and ethnicity. Through our web platform, we recruited participants to answer seven questions, some of which were designed to elicit A/H while recording themselves via webcam with microphone. BAH contains 1,118 videos for a total duration of 8.26 hours with 1.5 hours of A/H. Our behavioural team annotated timestamp segments to indicate where A/H occurs, and provide frame- and video-level annotations with the A/H cues. Video transcripts and their timestamps are also included, along with cropped and aligned faces in each frame, and a variety of participants meta-data. Additionally, this paper provides preliminary benchmarking results using baseline models trained on BAH for frame- and video-level recognition with mono- and multi-modal setups. It also includes results on models for zero-shot prediction, and for personalization using unsupervised domain adaptation. The limited performance of baseline models highlights the challenges of recognizing A/H in real-world videos. The data, code, and pretrained weights are publicly available.

1 Introduction

Emotion recognition plays a growing role in a range of health-related domains (Siddiqi et al., 2024), including disease prevention (Jin, 2024), diagnosis (Jiang et al., 2024; Maki et al., 2013), treatment monitoring (Dhuheir et al., 2021; Pepa et al., 2021; Suraj et al., 2022), and digital health promotion (Arabian et al., 2023; Subramanian et al., 2022), by supporting adaptive and responsive interventions (Liu et al., 2024b; Sinha et al., 2020). Emotion recognition technologies can support behaviour change interventions (Guo et al., 2024) by identifying affective states relevant to motivation, adherence, and engagement. Health-related behaviour change focuses on strategies to support individuals in adopting and maintaining healthy behaviours to prevent or manage chronic diseases, reduce early mortality, and improve mental health and well-being (Davidson & Scholz, 2020). Achieving and maintaining long-term behaviour change is a complex process (McDonald et al., 2002; Michie et al., 2013a;b), and often includes overcoming ambivalence and hesitancy (A/H). In face-to-face interactions, healthcare providers (e.g., clinicians, therapists) often identify A/H through a combination of speech and non-verbal cues (e.g., facial expressions and tone) (Heisel

Figure 1: Examples of body language cues used by annotators to identify the occurrence of A/H: "looking away," and "changing posture."

& Mongrain, 2004; Labbé et al., 2022; Miller & Rose, 2015). Identifying such complex emotion enables an adaptive, personalized, and more effective interventions.

A/H is understood as the simultaneous experience of desires for change (e.g., knowing a behaviour is beneficial for health) and against change (Armitage & Conner, 2000; Hohman et al., 2016) (e.g., building a new habit takes effort and time). This concept is considered to be at the core of starting and maintaining behaviour change (Conner & Armitage, 2008; Conner & Sparks, 2002; Manuel & Moyers, 2016; Miller & Rose, 2015). Individuals often find themselves somewhere between fully accepting and completely refusing an object or participating in a certain behaviour. However, accurately recognizing A/H, a subtle and multimodal emotion, is a challenging problem in personalized digital health (eHealth) interventions. Currently, there are no reliable and unobtrusive methods of assessing A/H in eHealth interventions. Therefore, designing robust automated methods to assist healthcare providers with A/H recognition capability is a key element for successful digital interventions.

Recent research on machine learning (ML) in emotion recognition focuses mainly on seven basic discrete emotions, e.g., 'Happy', 'Sad', and 'Surprised' (Belharbi et al., 2024a; Liu et al., 2024a; Xue et al., 2022). Other models in the literature predict ordinal levels, including pain and stress estimation (Aslam et al., 2024; Chaptoukaev et al., 2023; Zeeshan et al., 2024; Nasimzada et al., 2024), or continuous predictions such as valence-arousal (Dong et al., 2024; Praveen & Alam, 2024a; Praveen et al., 2023; 2021). However, real-world scenarios present more complex cases of emotions. Recently, there has been an increased interest in designing robust affect models for compound emotions, a case where a mixture of basic emotions is manifested (Kollias, 2023; Richet et al., 2024). In particular, compound emotions commonly occur in daily interactions. However, they are more difficult to discern as they are subtle, ambiguous, and resemble basic emotions. A/H recognition is related to such a task where intention and attitudes are conflicted or in a in-between state, between willingness and resistance (MacDonald, 2015), or positive and negative affect (Armitage & Conner, 2000). This can manifest in how individuals express themselves and can be recognized (Hayashi et al., 2023) in their facial expression, tone, verbal, and body language (Figure 1). As a result, A/H exhibits a multimodal nature that comes as the result of subtle interconnection between different cues. Unfortunately, such discord is extremely difficult to spot; a task that requires human training. This is a tedious and expensive procedure, leading to ineffective and less scalable eHealth interventions under limited resources. Assisting healthcare providers with automatic, reliable, and inconspicuous tools to help them recognize A/H can have a major impact in improving eHealth interventions. Although A/H is a common topic in behavioural science (Conner & Armitage, 2008; Hohman et al., 2016; Manuel & Moyers, 2016), it remains unexplored in the ML community, and as such, in the design of eHealth components that would allow to personalized behaviour change interventions. A possible reason is the lack of the necessary and specialized data for training and evaluation of ML models.

To address this limitation, we introduce in this work a first Behavioural Ambivalence/Hesitancy (BAH) dataset collected for subject-based multimodal recognition of A/H in videos. Through a collaboration with our behavioural science team, we have collected a large video dataset from 9 provinces in Canada. A data capture protocol is set in place to recruit diverse participants, including the development of a web-platform for video capturing, a dedicated storage server, and a specific annotation protocol. Our behavioural team designed seven questions to elicit responses regarding behaviours and to identify possible instances where participants are displaying A/H. Via



Figure 2: BAH dataset collection and annotation procedure. First, a participant access our web platform. They go through initial test/calibration to ensure the quality of the data. An avatar guides them throughout the entire process. Seven questions are presented to the participant. They are recorded while answering them. Once the data is captured, it is transferred by the Administrator to our local server. It is then annotated at several levels by an expert to determine when A/H occurs.

our web-platform, participants are presented these questions and asked to record themselves while answering via their device camera with a microphone. Participants are guided in the platform by an avatar throughout the entire data capture session. The dataset is composed of 224 participants. This amounts to a total of 1,118 videos (\sim 8.26 hours) where 638 videos contain A/H (\sim 1.50 hours). This amounts to 714,005 total frames where 131,103 contain A/H. Our behavioural analysis experts annotated the data at video- and frame-level to assess when A/H occurs. In addition, the video cues used by the annotators are reported such as facial expressions, body language, audio and language in addition to highlighting where there is inconsistency between the modalities. The BAH dataset is made public and it is provided with the raw videos with audio, cropped and aligned faces, detailed annotation/cues for video- and frame-level, audio transcript/timestamps/language, and participants meta-data such as age, ethnicity and more.

Our main contributions are summarized as follows. (1) A novel video dataset named BAH is proposed for subject-based multimodal recognition of A/H based on visual, audio, text and other relevant information. BAH is annotated by behavioural science experts at the video- and frame-levels. It can be used to develop and evaluate ML models for classification task, and build insight A/H for behaviour change interventions. The dataset is subject-based. It allows building different personalization learning methods such as domain adaptation techniques. (2) Preliminary benchmarking results for baseline models on BAH for frame- and video-based emotion recognition. Results allowed exploring the impact of key factors, including the impact of using temporal context, multimodal information, and feature fusion. Baseline results are also shown for other tasks – zero-shot prediction and personalization through subject-based domain adaptation. Our code and dataset is public.

2 The bah Dataset

2.1 Dataset Collection and Annotation

Capture. The BAH dataset contains Q&A videos. Its is constructed by collecting samples from participants over the age of 18 across Canada. Data collection and annotation process is presented in Figure 2. To proceed with the data collection, we developed "Automatic Expression Recognition" (AER) web-based platform (www.aerstudy.ca) where participants could record their responses to specific questions using their own computers or devices with camera and microphone. Users receive secure credentials to access the data collection platform, or they can create their own account. Participants first complete a brief survey to provide demographic information and indicate consent preferences (e.g., inclusion in secure datasets, challenges, or publications). They are then redirected to the AER platform, where they test their camera and microphone and choose an avatar to interact during data capture. The avatar guides them through seven questions. The session takes approximately 30 minutes. Participants are recruited and compensated via Prolific company (www.prolific.com), which also ensures population diversity and allows submission processing.

Participants answer seven questions designed by our behavioural team (Table 1), each one intended to elicit neutral, positive, negative, ambivalent, willing, resistant, and hesitant answers. Once the

Question no.	Response	Prompt
1	Neutral	Tell us about an activity you commonly do after waking up.
2	Positive	Talk about an activity that brings you joy, for example, a hobby. Tell us why.
3	Negative	Talk about an activity you dislike doing, for example, a chore or something you find boring or annoying. Tell us why.
4	Ambivalent	Tell us about something you enjoy doing but wish you stopped doing (like a guilty pleasure) or something you don't do but wish you did.
5	Willing	Tell us about an activity you are almost always willing to do, for example with friends, at work, at home.
6	Resistant	Tell us about something people around you do, but that you would not be willing to do, for example, with friends, at work, at home.
7	Hesitant	Tell us about something you could have done already but haven't done yet, for example, something you are procrastinating or haven't made up your mind about.

Table 1: The 7 questions (prompts) designed by our experts to create our videos for BAH dataset. To avoid influencing the participants answers, they are only shown prompts without indicating the expected emotion/response.

question is presented, the recording of the participant response starts. Skipping questions is allowed. At the end of each question, the participant has the option to rate their emotional response using a Likert-like 5-point scale. This self-rating is only employed for our analysis and does not serve as annotation. The order of the questions is randomized. In addition, participants are not aware what each questions is expected to illicit as emotion. During this capture procedure, several information is gathered including contact information of the participant, their demographics, consent, video recordings, survey responses, and software usage data (such as the time spent on each question). The participants' data is systematically downloaded and transferred to local secured server storage by the team for annotation and further analysis.

The study obtained human ethics approval from the two collaborating universities following all standard ethical practices. The dataset was collected between September 2024 and April 2025 in batches. This allowed us to adjust the targeted population (regarding participants' sex and Canadian province of residency) to ensure the dataset diversity.

Annotation. Three annotators were trained in expression recognition, specialized in identifying A/H, and in the annotation process of audio-visual data. A two-stage process was used: first, a global-level annotation determined the presence of A/H in each video; then, a frame-level annotation identified the precise segments where A/H occurred, specifying the start and end times (i.e., onset and offset) of each instance. Annotators also provided certainty ratings, and for some segments, indicated the cues that supported their judgment. To identify A/H, annotators tracked expressions across well-established modalities, (facial expressions, body language, audio, and language) and flagged cases where inconsistencies between modalities were observed. We do not include an "apex" annotation, as ambivalence and hesitancy do not reliably exhibit a peak moment of maximum intensity. Instead, they tend to manifest as sustained or fluctuating states, making the concept of an apex incompatible with their typical temporal structure. The videos were annotated following a codebook created specifically for the study. Videos were annotated using the ELAN 8 (archive.mpi.nl/tla/elan) software (Figure 2).

The annotation process followed a structured training protocol supported by a detailed training manual. Annotators first received a conceptual introduction to A/H, followed by hands-on training in using the ELAN annotation software. Practical application was conducted using a standardized set of videos from the dataset. This phase also introduced annotators to the codebook, emphasizing the cue list, with examples spanning facial, vocal, verbal, and bodily expressions. Annotators received feedback and additional sessions were provided when further alignment was needed. Only after this training phase, and a final assessment, did annotators proceed to independent annotation.

To promote consistency, annotators were instructed to flag cases of uncertainty or complexity. These cases were discussed collaboratively, often through co-annotation. A consistent lead annotator facilitated resolution efforts, ensuring that decisions reflected a shared interpretation. In parallel, a comprehensive annotation protocol guided how videos were managed, accessed, and annotated. Annotators followed standardized procedures: (1) watch the video without taking notes to understand

the participant and context; (2) re-watch the video to identify A/H segments and record start and end times; (3) reassess and refine selected segments; (4) identify and assign cues using the codebook; and (5) if needed, watch the video without audio or visual elements to isolate specific signals. Annotators were also encouraged to consult other videos from the same participant to establish expressive baselines in ambiguous cases.

The presence of A/H is assigned a single label (1), while its absence is assigned the label 0. Each video has a global- and frame-level label which can be used to train and evaluate ML models. The provided cues can also be used for interpretability aspect as well as to build insights on how people express A/H. The dataset is structured subject-wise which can be also useful for personalization training scenarios.

2.2 Dataset Variability

The dataset is designed to approximate the demographic distribution of sex and provincial representation in Canada. The BAH dataset is composed of 224 participant across Canada from nine provinces where 31.2% of participants is from British Columbia followed by Alberta with 20.5% and Quebec with 16.5% . All participants agreed to be part of this dataset. However, 50 participants (22.3%)¹ did not consent to be in publications while only seven participants (3.1%) did not consent to be part of challenges. The recorded videos contain both English and French languages. Each participant can record up to seven videos where 96 participants have recorded the full seven videos. We obtained an average of \sim 5 videos/participant where each participant has an average of \sim 2.84 videos with A/H which is equivalent to \sim 585 frames of A/H (or \sim 24.15 seconds of A/H). The dataset amounts a total of 1,118 videos (\sim 8.26 hours) where 638 videos contain A/H (\sim 1.50 hours). This amounts to 714,005 total frames where 131,103 contain A/H. Since captured videos represent answers to questions, they are relatively short. BAH dataset has an average video duration of 26.58 \pm 16.36 (seconds) with a minimum and maximum duration of 3 and 96 seconds.

An important characteristic of this dataset is the duration of the A/H segments in videos. BAH counts a total of 376 videos with multiple A/H segments and 259 videos with only one A/H segment. In total, there are 1,274 A/H segments. In particular, the duration of segment varies but it is brief with an average of 4.25 ± 2.47 seconds which is equivalent to 102.92 ± 59.16 frames. The minimum and maximum A/H segment is 0.01 seconds (1 frame), and 23.8 seconds (572 frames), respectively.

In terms of participants age, the dataset covers a large range from 18 to 66 years old. In particular, 37.1% of the participants covers the range 25-34 years, followed by the range of 35-44 years with 25.9%, then the range of 18-24 with 21.9%. In terms of sex, 59.8% are male, while 39.3 are female. As for ethnicity variation, White comes with 52.2% of the participants, followed by Asian with 21.0%, and Mixed with 11.6%, then Black with 10.3%. Large part of the participants are not students (65.2%) which limits common issues in recruit bias.

The public BAH dataset contains the row videos, detailed A/H annotation at video- and frame-level, cues, and per participant demographic information including age, birth country, Canada province where the participant lives, ethnicity, ethnicity simplified, sex, student status, consent to use recordings in publications. More details about the dataset diversity are provided in the appendix.

2.3 ETHICAL CONSIDERATION, DATASET ACCESSIBILITY AND INTENDED USES

The collected data of human participants follows tightly ethical considerations. The project to collect BAH data was approved by ethical committees from both collaborating universities. Once recruited, participants have access to the full consent form prior to accessing the data capture platform and starting their data capture procedure. They are provided with details of the study, as well as a list of the potential risks and benefits of participating in the study. They are be instructed to read the consent form thoroughly and they are provided with a clear and simple video that summarizes the consent form. Participants are then able to decide the type of access they want the researcher to have to their audiovisual data, including if they want their images to be used for publications and presentations. In addition, these options are presented again at the end of the data capture procedure, just in case they change their mind around their participation in the study or the use of their data after they have finished recording their responses. At the end of the study, participants receive, via email,

¹The list of these participants is provided within the shared files of the BAH dataset.

a copy of the consent form that includes their choices about data usage and the contact information for the team should they have any further questions. Note that participants are given numerical codes for anonymity.

Following the guidelines of the funding agency , the BAH dataset is made public with open credentialed access for research purposes. To access the dataset, users are required to fill in a request form and sign an End-User License Agreement (EULA) as commonly done to ensure dataset security. Upon access approval, the user will receive a link to download the full dataset, including row videos, detailed annotation, cues, participants' meta-data, cropped-and-aligned-faces, frames, audio transcripts. BAH uses a proprietary license for research purposes. The dataset is hosted in a secured server as it is intended for long-term availability. Our public code is under an open-source license (BSD-3-Clause license). The code website will be used as a permanent page for the dataset that will reflect any future updates. Despite all our precautions, our dataset may still could be misused. We consider a thorough review of requests before granting data access. Reviewers can directly download the BAH dataset via the link provided in the appendix. Please read Sec.A in appendix before proceeding to download the dataset.

Our primary goal of building BAH dataset is to make public a first and unique dataset for A/H recognition in videos. Given the content of the dataset, its multi-modal aspect, and the provided annotation, it can be used to train and evaluate ML models for A/H recognition in videos at frame-and/or video-level with different learning scenarios. Since data is subject-based, it can also be used for personalization using domain adaptation, for instance. The provided cues used by annotators can also be used for interpretability learning, and further analysis to get more insight on our understanding of A/H in human behaviours. Such understanding and recognition of A/H can be leveraged in downstream tasks such as behavioural change, interventions and recommendations in clinics or via automated systems such as virtual trainers/assistants.

2.4 EXPERIMENTAL PROTOCOL

Dataset split. The dataset is divided randomly based on participants into 3 sets: train, validation and test set. The train and validation sets amounts to 3/4 of the total participants, while 1/4 goes to the test set. Videos of one participants belong to one and one set only. The details of each set is presented in Table 2. The split files are provided along with the dataset files. They contain the split in terms of videos and frames ready to use. Note that the dataset is highly imbalanced as depicted in Table 3, especially at frame level where only 18.15% contains A/H. This factor should be accounted for during training and evaluation. The dataset can be used for training at video- and/or frame-level. The participant identifiers are provided in the splits allowing subject-based learning scenarios.

Data subsets	Train	Validation	Test	Total
Number participants	143	25	56	224
Number participants with A/H	111	22	54	187
Number videos	616	110	392	1118
Number videos with A/H	340	59	239	638
Number frames	375,060	82,201	256,744	714,005
Number frames with A/H	65,282	15,157	50,664	131,103
Total duration (hour)	4.34	0.95	2.97	8.26
Total duration with A/H (hour)	0.74	0.17	0.58	1.50

Table 2: BAH dataset split into train, validation, and test sets.

Data subsets	Train (%)	Validation (%)	Test (%)	Total (%)
Participants with A/H	77.62	88.00	96.42	83.48
Videos with A/H	55.19	53.63	60.96	57.06
Frames with A/H	17.40	18.43	19.73	18.36
Duration with A/H	17.05	17.89	19.52	18.15

Table 3: Imbalance rate of BAH dataset split across train, validation, and test sets: (Total # items with A/H)/(Total # items).

Evaluation metrics. We refer here to the positive class as the class with label 1 indicating the presence of A/H, while negative class is the class 0 indicating the absence of A/H. To account

for the imbalance in BAH dataset, we use adapted standard evaluation metrics: - F1 score of the positive class. - Weighted F1 (WF1) score which is a weighted average of F1 of the positive and negative class. - Average F1 (AVGF1) score which is the unweighted mean of F1 of the positive and negative class. - Average precision score (AP) of the positive class which accounts for the performance sensitivity to the model's confidence. For AP score, a threshold list between 0 and 1 is used with a step of 0.001. Evaluation code of all measures is provided along with the public code of this dataset. We report results on more relevant metrics in the main paper, while results on all metrics are included in the appendix.

3 BASELINE RESULTS

This section provides preliminary results of different baseline models on our BAH dataset. In particular, we provide performance of models for the supervised frame-level classification task. We consider a 2-class classification problem where each frame is annotated, and models predict two outputs: one for the positive class (presence of A/H), and a second for the absence of A/H. Supervised video-level classification performance is included in the appendix. In addition, the results of other tasks – zero-shot prediction, and personalization through unsupervised domain adaptation – are also included.

We initially focus on the impact of using single vs multimodal learning for frame-level classification. Then, the performance of different individual modalities are explored, along with their multimodal fusion. In addition, we investigate the impact of temporal modeling and context vs single frame learning. In the following, we present the pre-processing of the three different used modalities: visual (facial), audio (vocal), and text transcripts (textual), and describe the baseline models used in each case.

3.1 Pre-processing of Modalities

1) Visual. All frames from each video are extracted, and for each frame, faces are located using RetinaFace model (Deng et al., 2019), cropped, then aligned. The face with the highest score is stored in case of multiple faces are detected in a frame. Faces are resized to 256×256 and stored as RGB images with a file name that maintains the order of frames. The video frame rate is 24 FPS. 2) Audio. We follow standard procedure to process audio data (Praveen & Alam, 2024b; Richet et al., 2024; Zhang et al., 2023). For audio modality, we first convert videos to single audio channels (mono) with a 16k sampling rate into way format. The log melspectrograms features are extracted using Vggish model (Hershey et al., 2017)(github.com/harritaylor/torchvggish). A hope of 1/FPS of the raw video is used to extract the spectrograms to synchronize audio with other modalities. 3) Text. The collected data captures the audio of participants. We consider audio transcripts as an extra modality that can help recognizing A/H since text is a significant cue used by annotators. To this end, we transcribe the audio of each video, and detect the language using Whisper model (Radford et al., 2023) (Whisper large-v3 multilingual: huggingface.co/openai/whisper-large-v3). We provide the timestamp of each transcript. Word-level features are then extracted using BERT Base Uncased model (Devlin et al., 2019)(pypi.org/project/pytorch-pretrained-bert/). A word may span more than one frame. To synchronize with other modalities, a word-level features is repeated per its timestamp for all the frames that correspond to the word.

Note that both cropped and aligned faces, and video transcripts are shipped with the shared public BAH dataset. Researchers can choose to use them or build their own.

3.2 PRE-TRAINING OF VISUAL BACKBONE

For audio and text modality, features are extracted offline and stored as described above. For visual modality, we explore different architectures including CNN- and ViT-based(Dosovitskiy et al., 2021). In particular, we explore ResNet family including ResNet18, 34, 50, 101, and 152 (He et al., 2016). For ViT family, we consider a recent model designed for basic emotion recognition, APViT (Xue et al., 2022). First, we pre-train each model on basic emotion recognition task including these emotions: "Anger", "Disgust", "Fear", "Happiness", "Sadness", "Surprise". To this end, we collected a large mixed dataset composed of 3 public common datasets for emotion recognition using images: RAF-DB (Li et al., 2017), and AffectNet (Mollahosseini et al., 2019), and

	Without context		With context (TCN	
Backbone	AVGF1	AP	AVGF1	AP
APViT (Xue et al., 2022)	0.5289	0.2221	0.5227	0.2217
ResNet18 (He et al., 2016)	0.4513	0.2064	0.5066	0.1842
ResNet34 (He et al., 2016)	0.4922	0.2107	0.4989	0.1683
ResNet50 (He et al., 2016)	0.4525	0.2124	0.5183	0.2068
ResNet101 (He et al., 2016)	0.5301	0.2280	0.5370	0.2191
ResNet152 (He et al., 2016)	0.5076	0.2096	0.5454	0.2674

Table 4: Visual modality performance on test set of BAH at frame-level classification: impact of architecture and context.

Aff-wild2 (Kollias & Zafeiriou, 2019). This amounts to more than 0.54 millions training images. Models are trained for basic emotion classification for 60 epochs with a batch size of 1,424 samples using 4 parallel NVIDIA A100 GPUs with 40 GB of memory. Standard cross-entropy loss and Stochastic Gradient descent (SGD) are used for training. Once pretrained, each model is further fine tuned on our BAH train set for A/H recognition. To account for class imbalance, we perform under-sampling of negative class over train set. This is achieved by randomly sampling the negative class samples to be the same as the positive class. The weights of both models will be made public. The backbones of each model are used later for feature extraction of visual modality.

3.3 IMPORTANCE OF CONTEXTUAL LEARNING

In this section, we aim to answer the question whether context modeling can help better A/H recognition. This is particularly interesting since A/H does not occur instantly, but within a context. Text and audio modality already capture context in their features making using them to answer this question less efficient. However, we can obtain frame features with and without any context. Therefore, we consider visual modality to answer our question. To this end, we study different visual backbones.

In the case of modeling without context, models simply train on independent frames without considering any context or dependency between them. Inference is done in the same way. In the case of modeling with context, both training and inference leverages temporal dependency between frames. To this end, a window of adjacent frames is fed to the model. We then use temporal convolutional network (TCN) (Bai et al., 2018) after the visual backbone to capture relations between frames embeddings. The window length defines the extent of the context. Table 4 shows the obtained results. Regardless of the context, we observe a very low performance over AP highlighting the difficulty of recognizing A/H based on images alone. In particular AP is below 0.2674. On the other hand, AVGF1 reaches 0.5454. Although it seems high, it is not significant as the

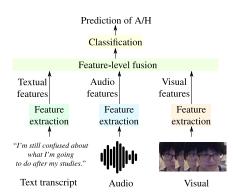


Figure 3: Multimodal model used to produce baseline performance (Richet et al., 2024).

negative class dominates making it easy to score high values. As a reference, predicting every frame as negative class yields an AVGF1 of 0.4452. We note that overall, using context boost performance of all metrics across all architectures. This is expected as A/H does not usually occur at a single frame but withing a context. This makes its recognition from a single frame challenging. We recall that the average A/H segments spans 102 frames (or 4.25 seconds). Future works should account for temporal context for better performance. However, as we will show in the ablations in the appendix, very large context could be counter-productive. Note that large ResNet models seems to yield better performance overall. Unless mentioned otherwise, all our next experiments will use ResNet152.

3.4 MULTIMODAL BASELINES

Since A/H is multimodal by nature, we explore the impact of using different modalities, including visual, audio, and transcript using the model presented in Figure 3 (Richet et al., 2024). Results are reported in Table 5. Using visual modality alone yields better performance compared to audio or text

modalities. However, looking to AP, visual modality provides more certain predictions followed by audio, then text modality. Combining a pair of modalities improves furthermore the performance to 0.5756 for AVGF1 with the case of visual and text and 0.2809 for AP. Combing the three modalities slightly reduces performance. This may suggest that better and more adapted fusion techniques are needed to recognized conflicts between modalities.

Table 6 shows the impact of using different feature fusing techniques including simple concatenation (CAN) (Zhang et al., 2023), co-attention (LFAN) (Zhang et al., 2023), transformer-based fusion (MT) (Waligora et al., 2024), and cross-attention fusion (JMT) (Waligora et al., 2024). We observe that the way of leveraging the interaction between the three modalities is a key factor. Although simple feature concatenation seems to yield the best performance, future works should pursue more adapted methods to A/H. Ambivalence and hesitancy are usually expressed as a conflict between willingness and resistance. This can be perceived through a parallel con-

Modalities	AVGF1	AP
Visual	0.5454	0.2674
Audio	0.5372	0.2520
Text	0.5309	0.2047
Visual + Audio	0.5636	0.2818
Visual + Text	0.5756	0.2809
Audio + Text	0.5644	0.2459
Visual + Audio + Text	0.5557	0.2416

Table 5: Multimodal models performance on test set of BAH at frame-level classification. For visual modality, ResNet152 backbone is used.

flict between modalities and or within modalities. For instance, a participant could say a sentence to convey a meaning but their facial expression, body behaviour, or tone may carry a contradictory emotion. Understanding such subtly and interconnection between different cues in different modalities could play an important role in designing robust methods for A/H recognition in videos.

We believe our new and unique dataset has brought a new challenging research direction to better understand complex and subtle human emotions that is A/H. Given the multimodal nature of A/H, our BAH dataset provides an essential and valuable toolkit for the research community to design and evaluate their methods. Important key and critical downstream tasks could potentially benefits from these methods including but not limited to clinical interviews, interventions, behavioural changes, and automated assistants such as online trainers. Our preliminary results suggest that leveraging context, multimodality, and their fusion could lead to better A/H recognition performance.

Fusion type	AVGF1	AP
LFAN (Zhang et al., 2023) (cvprw,2023)	0.5557	0.2416
CAN (Zhang et al., 2023) (cvprw,2023)	0.5708	0.2559
MT (Waligora et al., 2024) (cvprw,2024)	0.4834	0.2108
JMT (Waligora et al., 2024) (cvprw,2024)	0.5350	0.2335

Table 6: Feature fusion performance on test set of BAH at frame-level classification.

4 CONCLUSION

This work introduces a new and unique multimodal and subject-based video dataset, BAH, for A/H recognition in videos. BAH contains 224 participants across 9 provinces in Canada. Recruited participants answer 7 designed questions to elicit A/H while recording themselves via webcam and microphone via our web-platform. The dataset amounts to 1,118 videos for a total duration of 8.26 hours with 1.5 hours of A/H. It was annotated by our behavioural team at video- and frame-level. Our initial benchmarking yielded limited performance highlighting the difficulty of A/H recognition. Our results showed also that leveraging context, multimodality, and adapted feature fusion is a first good direction to design robust models. Our dataset and code are made public.

The following appendix contains related work, more detailed and relevant statistics about the datasets and its diversity, dataset limitations, implementation details, and additional results. Our code is included in the supplementary materials.

REFERENCES

- H. Arabian, T. Abdulbaki Alshirbaji, R. Schmid, V. Wagner-Hartl, J. Chase, and K. Moeller. Harnessing wearable devices for emotional intelligence: Therapeutic applications in digital health. *Sensors*, 23(19):8092, 2023.
- C. J Armitage and M. Conner. Attitudinal ambivalence: A test of three key hypotheses. *Personality and Social Psychology Bulletin*, 26(11):1421–1432, 2000.
- M.H. Aslam, M.O. Zeeshan, S. Belharbi, M. Pedersoli, Al.L. Koerich, S. Bacon, and E. Granger.
 Distilling privileged multimodal information for expression recognition using optimal transport.
 In *International Conference on Automatic Face and Gesture Recognition (FG)*, 2024.
 - S. Bai, J.Z. Kolter, and V. Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *CoRR*, abs/1803.01271, 2018.
- S. Belharbi, M. Pedersoli, A. L. Koerich, S. Bacon, and E. Granger. Guided interpretable facial expression recognition via spatial action unit cues. In *International Conference on Automatic Face and Gesture Recognition (FG)*, 2024a.
 - S. Belharbi, M. Pedersoli, A.L. Koerich, S. Bacon, and E. Granger. Spatial action unit cues for interpretable deep facial expression recognition. In *AI and Digital Health Symposium*, 2024b.
 - M. Bishay, P. Palasek, S. Priebe, and I. Patras. Schinet: Automatic estimation of symptoms of schizophrenia from facial behaviour analysis. *IEEE Transactions on Affective Computing*, 12(4): 949–961, 2019.
 - J. Bonnard, A. Dapogny, F. Dhombres, and K. Bailly. Privileged attribution constrained deep networks for facial expression recognition. In *ICPR*, 2022.
 - E.H. Bradley, L.A. Curry, and K.J. Devers. Qualitative data analysis for health services research: developing taxonomy, themes, and theory. *Health services research*, 42(4):1758–1772, 2007.
 - C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J.N. Chang, S. Lee, and S.S. Narayanan. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42:335–359, 2008.
 - H. Chaptoukaev, V. Strizhkova, M. Panariello, B. Dalpaos, A. Reka, V. Manera, S. Thummler, E. ISMAILOVA, N. Evans, F. Bremond, M. Todisco, M. A Zuluaga, and L. M. Ferrari. StressID: a multimodal dataset for stress identification. In *NeurIPS*, 2023.
 - J. Choe, S.J. Oh, S. Chun, S. Lee, Z. Akata, and H. Shim. Evaluation for weakly supervised object localization: Protocol, metrics, and datasets. *TPAMI*, pp. 1–1, 2022.
- Y. Chu, L. Liao, Z. Zhou, C.-W. Ngo, and R. Hong. Towards multimodal emotional support conversation systems. *CoRR*, abs/2408.03650, 2024.
 - M. Conner and C.J. Armitage. Attitudinal ambivalence. 2008.
- M. Conner and P. Sparks. Ambivalence and attitudes. *European review of social psychology*, 12(1): 37–70, 2002.
- K.W. Davidson and U. Scholz. Understanding and predicting health behaviour change: a contemporary view through the lenses of meta-reviews. *Health psychology review*, 14(1):1–5, 2020.
- M. De-la Torre, E. Granger, P. VW Radtke, R. Sabourin, and D.O. Gorodnichy. Partially-supervised learning from facial trajectories for face recognition in video surveillance. *Information fusion*, 24: 31–53, 2015.
 - J. Deng, J. Guo, Y. Zhou, J. Yu, I. Kotsia, and S. Zafeiriou. Retinaface: Single-stage dense face localisation in the wild. *CoRR*, abs/1905.00641, 2019.
 - J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pp. 4171–4186, 2019.

- M. Dhuheir, A. Albaseer, E. Baccour, A. Erbad, M. Abdallah, and M. Hamdi. Emotion recognition for healthcare surveillance systems using neural networks: A survey. In 2021 International Wireless Communications and Mobile Computing (IWCMC), pp. 681–687, 2021.
- L. Dong, X. Wang, S. Setlur, V. Govindaraju, and I. Nwogu. Ig3d: Integrating 3d face representations in facial expression inference. *CoRR*, abs/2408.16907, 2024.
 - A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- Y. Fan, J. Lam, and V. Li. Facial action unit intensity estimation via semantic correspondence learning with dynamic graph convolution. In *AAAI*, 2020.
 - R. Guo, H. Guo, L. Wang, M. Chen, D. Yang, and B. Li. Development and application of emotion recognition technology—a systematic literature review. *BMC psychology*, 12(1):95, 2024.
 - X. Guo, B. Zhu, L.F. Polanía, C. Boncelet, and K.E. Barner. Group-level emotion recognition using hybrid deep models based on faces, scenes, skeletons and visual attentions. In *Proceedings of the 20th ACM international conference on multimodal interaction*, pp. 635–639, 2018.
 - J.A. Hall, J.A. Harrigan, and R. Rosenthal. Nonverbal behavior in clinician—patient interaction. *Applied and preventive psychology*, 4(1):21–37, 1995.
- J. Han, L. Xie, J. Liu, and X. Li. Personalized broad learning system for facial expression. *Multi-media Tools and Applications*, 2020.
- D. Hayashi, S.D.L. Carvalho, P.A.B. Ribeiro, R.C.M. Rodrigues, T.M. São-João, K. Lavoie, S. Bacon, and M. E. Cornelio. Methods to assess ambivalence towards food and diet: a scoping review. *Journal of Human Nutrition and Dietetics*, 36(5):2010–2025, 2023.
 - K. He, X. Zhang, S.g Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
 - Z. He, Z. Li, F. Yang, L. Wang, J. Li, C. Zhou, and J. Pan. Advances in multimodal emotion recognition based on brain–computer interfaces. *Brain sciences*, 10(10):687, 2020.
 - M.J. Heisel and M. Mongrain. Facial expressions and ambivalence: Looking for conflict in all the right faces. *Journal of Nonverbal Behavior*, 28:35–52, 2004.
 - S. Hershey, S. Chaudhuri, D.P.W. Ellis, J.F. Gemmeke, A. Jansen, R.C. Moore, M. Plakal, D. Platt, R.A. Saurous, B. Seybold, M. Slaney, R.J. Weiss, and K.W. Wilson. Cnn architectures for large-scale audio classification. In *ICASSP*, 2017.
 - Z.P. Hohman, W.D. Crano, and E.M. Niedbala. Attitude ambivalence, social norms, and behavioral intentions: Developing effective antitobacco persuasive communications. *Psychology of Addictive Behaviors*, 30(2):209, 2016.
 - S. Hornstein, K. Zantvoort, U. Lueken, B. Funk, and K. Hilbert. Personalization strategies in digital mental health interventions: a systematic review and conceptual framework for depressive symptoms. *Frontiers in digital health*, 5:1170002, 2023.
 - G. M. Jacob and B. Stenger. Facial action unit detection with transformers. In CVPR, 2021.
- D. Jiang, L. Yan, and F. Mayrand. Emotion expressions and cognitive impairments in the elderly: review of the contactless detection approach. *Frontiers in Digital Health*, 6:1335289, 2024.
- H. Jin. A comparative analysis of single and multi-modality-based emotion recognition for disease prevention. In *International Conference on Artificial Intelligence and Communication (ICAIC)*, volume 185, pp. 323, 2024.
 - R. Khanna, N. Robinson, M. O'Donnell, H. Eyre, and E. Smith. Affective computing in psychotherapy. *Advances in Psychiatry and Behavioral Health*, 2(1):95–105, Sep 2022.

D. Kollias. Multi-label compound expression recognition: C-expr database & network. In *CVPR*, 2023.

- D. Kollias and S. Zafeiriou. Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arcface. *CoRR*, 2019.
- D. Kollias, P. Tzirakis, A. S. Cowen, S. Zafeiriou, I. Kotsia, E. Granger, M. Pedersoli, S. L. Bacon, A. Baird, C. Gagne, C. Shao, G. Hu, S. Belharbi, and M. H. Aslam. Advancements in Affective and Behavior Analysis: The 8th ABAW Workshop and Competition. 3 2025. doi: 10.6084/m9. figshare.28524563.v4.

- J. Kossaifi, R. Walecki, Y. Panagakis, J. Shen, M. Schmitt, F. Ringeval, J. Han, V. Pandit, A. Toisoul, B. Schuller, et al. Sewa db: A rich database for audio-visual emotion and sentiment research in the wild. *TPAMI*, 43(3):1022–1040, 2019.
- K. Kraack. A multimodal emotion recognition system: Integrating facial expressions, body movement, speech, and spoken language. *CoRR*, abs/2412.17907, 2024.
 - S. Labbé, I. Colmegna, V. Valerio, V.G. Boucher, S. Peláez, A.I. Dragomir, C. Laurin, E.M. Hazel, S.L. Bacon, and K.L. Lavoie. Training physicians in motivational communication to address influenza vaccine hesitation: a proof-of-concept study. *Vaccines*, 10(2):143, 2022.
 - I. Lee, E. Lee, and S.B. Yoo. Latent-ofer: Detect, mask, and reconstruct with latent vectors for occluded facial expression recognition. In *ICCV*, 2023.
 - S. Li and W. Deng. Deep emotion transfer network for cross-database facial expression recognition. In *ICPR*, 2018.
 - S. Li, W. Deng, and J. Du. Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild. In *CVPR*, 2017.
 - J. Liang, D. Hu, and J. Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *ICML*, 2020.
 - B. Lin, Y. Ye, B. Zhu, J. Cui, M. Ning, P. Jin, and L. Yuan. Video-llava: Learning united visual representation by alignment before projection, 2024.
 - C. Liu, X. Zhang, X. Liu, T. Zhang, L. Meng, Y. Liu, Y. Deng, and W. Jiang. Facial expression recognition based on multi-modal features for videos in the wild. In *CVPR*, 2023a.
 - D. Liu, H. Zhang, and P. Zhou. Video-based facial expression recognition using graph convolutional networks. In *ICPR*, 2021a.
 - H. Liu, R. An, Z. Zhang, B. Ma, W. Zhang, Y. Song, Y. Hu, W. Chen, and Y. Ding. Norface: Improving facial expression analysis by identity normalization. *ECCV*, 2024a.
- X. Liu, L. Jin, X. Han, J. Lu, J. You, and L. Kong. Identity-aware facial expression recognition in compressed video. In *ICPR*, 2021b.
 - Y. Liu, W. Wang, Ch. Feng, H. Zhang, Z. Chen, and Y. Zhan. Expression snippet transformer for robust video-based facial expression recognition. *Pattern Recognition*, 138:109368, 2023b.
 - Y. Liu, Y. Zhang, and Y. Wang. Application of deep learning-based image processing in emotion recognition and psychological therapy. *Traitement du Signal*, 41(6):2923, 2024b.
- H. Lokhande, C. Garware, T. Kudale, and R. Kumar. Personalized well-being interventions (pwis):
 A new frontier in mental health. In *Affective Computing for Social Good: Enhancing Well-being, Empathy, and Equity*, pp. 183–200. 2024.
- 645 I. Loshchilov and F. Hutter. SGDR: stochastic gradient descent with warm restarts. In *ICLR*, 2017.
 - C. Luo, S. Song, W. Xie, L. Shen, and H. Gunes. Learning multi-dimensional edge feature-based AU relation graph for facial action unit recognition. In *IJCAI*, 2022.

- N.E. MacDonald. Vaccine hesitancy: Definition, scope and determinants. *Vaccine*, 33(34):4161–4164, 2015.
- Y. Maki, H. Yoshida, T. Yamaguchi, and H. Yamaguchi. Relative preservation of the recognition of positive facial expression "happiness" in alzheimer disease. *International Psychogeriatrics*, 25 (1):105–110, 2013.
 - J.K. Manuel and T.B. Moyers. The role of ambivalence in behavior change. *Addiction*, 111(11): 1910–1912, November 2016.
 - J. Mao, R. Xu, X. Yin, Y. Chang, B. Nie, A. Huang, and Y. Wang. Poster++: A simpler and stronger facial expression recognition network. *Pattern Recognition*, pp. 110951, 2024.
 - H.P. McDonald, A.X. Garg, and R.B. Haynes. Interventions to enhance patient adherence to medication prescriptions: scientific review. *Jama*, 288(22):2868–2879, 2002.
 - S. Michie, M. Richardson, M. Johnston, C. Abraham, J. Francis, W. Hardeman, M.P. Eccles, J. Cane, and C.E. Wood. The behavior change technique taxonomy (v1) of 93 hierarchically clustered techniques: building an international consensus for the reporting of behavior change interventions. *Annals of behavioral medicine*, 46(1):81–95, 2013a.
 - S. Michie, R. West, and B. Spring. Moving from theory to practice and back in social and health psychology. 2013b.
 - W.R. Miller and G.S. Rose. Motivational interviewing and decisional balance: contrasting responses to client ambivalence. *Behavioural and cognitive psychotherapy*, 43(2):129–141, 2015.
 - J.A. Miranda Calero, L. Gutiérrez-Martín, E. Rituerto-González, E. Romero-Perales, J.M. Lanza-Gutiérrez, C. Peláez-Moreno, and C. López-Ongil. Wemac: Women and emotion multi-modal affective computing dataset. *Scientific data*, 11(1):1182, 2024.
 - A. Mollahosseini, B. Hassani, and M. H. Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Trans. Affect. Comput.*, 10(1):18–31, 2019.
 - S. Murtaza, S. Belharbi, M. Pedersoli, and E. Granger. A realistic protocol for evaluation of weakly supervised object localization. In *WACV*, 2025.
 - J. Nasimzada, J. Kleesiek, K. Herrmann, A. Roitberg, and C. Seibold. Towards synthetic data generation for improved pain recognition in videos under patient constraints. *CoRR*, abs/2409.16382, 2024.
 - A. O'Donnell, M. Addison, L. Spencer, H. Zurhold, M. Rosenkranz, R. McGovern, E. Gilvarry, M.-S. Martens, U. Verthein, and E. Kaner. Which individual, social and environmental influences shape key phases in the amphetamine type stimulant use trajectory? a systematic narrative review and thematic synthesis of the qualitative literature. *Addiction*, 114(1):24–47, 2019.
 - C. Ortiz, T. López-Cuadrado, A. Ayuso-Álvarez, C. Rodríguez-Blázquez, and I. Galán. Co-occurrence of behavioural risk factors for non-communicable diseases and mortality risk in spain: a population-based cohort study. *BMJ open*, 15(1):e093037, 2025.
 - M. Pantic and L.J.M. Rothkrantz. Toward an affect-sensitive multimodal human-computer interaction. *Proceedings of the IEEE*, 91(9):1370–1390, 2003.
 - L. Pepa, L. Spalazzi, M. Capecci, and M. G. Ceravolo. Automatic emotion recognition in clinical scenario: a systematic review of methods. *IEEE Transactions on Affective Computing*, 14(2): 1675–1695, 2021.
 - S. Poria, D. Hazarika, N. Majumder, G. Naik, E. Cambria, and R. Mihalcea. MELD: A multimodal multi-party dataset for emotion recognition in conversations. In *Conference of the Association for Computational Linguistics*, pp. 527–536, 2019.
 - R. G. Praveen and Ja. Alam. Inconsistency-aware cross-attention for audio-visual fusion in dimensional emotion recognition. *CoRR*, abs/405.12853, 2024a.

- R. G. Praveen, E. Granger, and P. Cardinal. Cross attentional audio-visual fusion for dimensional emotion recognition. In *International Conference on Automatic Face and Gesture Recognition* (FG 2021), 2021.
- R. G. Praveen, P. Cardinal, and E. Granger. Audio-visual fusion for emotion recognition in the valence-arousal space using joint cross-attention. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 5(3):360–373, 2023.
- R.G. Praveen and J. Alam. Recursive joint cross-modal attention for multimodal fusion in dimensional emotion recognition. In *ABAW workshop at CVPR*, 2024b.
- A. Radford, J. Wook Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever. Robust speech recognition via large-scale weak supervision. In *ICLM*, 2023.
- N. Richet, S. Belharbi, H. Aslam, O. Zeeshan, A. L. Koerich, M. Pedersoli, S. Bacon, and E. Granger. Textualized and feature-based models for compound multimodal emotion recognition in the wild. In *ABAW workshop at ECCV*, 2024.
- F. Ringeval, A. Sonderegger, J. Sauer, and D. Lalanne. Introducing the recola multimodal corpus of remote collaborative and affective interactions. In *International conference and workshops on automatic face and gesture recognition (FG)*, pp. 1–8, 2013.
- E. Sánchez-Lozano, G. Tzimiropoulos, and M. F. Valstar. Joint action unit localisation and intensity estimation through heatmap regression. In *BMVC*, 2018.
- D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and rkhs-based statistics in hypothesis testing. *The annals of statistics*, pp. 2263–2291, 2013.
- M. Sharafi, E. Ollivier, M.O. Zeeshan, S. Belharbi, M. Pedersoli, A.L. Koerich, S. Bacon, and E. Granger. Disentangled source-free personalization for facial expression recognition with neutral target data. In *International Conference on Automatic Face and Gesture Recognition (FG)*, 2025.
- M. H. Siddiqi, I. Ahmad, Y. Alhwaiti, and F. Khan. Facial expression recognition for healthcare monitoring systems using neural random forest. *IEEE Journal of Biomedical and Health Informatics*, 2024.
- S. Sinha, S.Kr. Mishra, and S. Bilgaiyan. Emotion analysis to provide counseling to students fighting from depression and anxiety by using cctv surveillance. In *Machine Learning and Information Processing: Proceedings of ICMLIP* 2019, pp. 81–94, 2020.
- B. Subramanian, J. Kim, M. Maray, and A. Paul. Digital twin model: A real-time emotion recognition system for personalized healthcare. *IEEE Access*, 10:81155–81165, 2022.
- A. Suraj, A. S Kaushik, and K. Bai. Patient monitoring using emotion recognition. *International Journal for Research in Applied Science and Engineering Technology*, 10(8):1197–1201, 2022.
- I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *ICML*, 2013.
- A.D. Vairamani. Advancements in Multimodal Emotion Recognition: Integrating Facial Expressions and Physiological Signals, pp. 217–240. Cham, 2024.
 - P. Waligora, H. Aslam, O. Zeeshan, S. Belharbi, A. L. Koerich, M. Pedersoli, S. Bacon, and E. Granger. Joint multimodal transformer for emotion recognition in the wild. In *ABAW workshop at CVPR*, 2024.
 - S. Walter, S. Gruss, H. Ehleiter, J. Tan, H.C. Traue, P. Werner, A. Al-Hamadi, S. Crawcour, A.O. Andrade, and G. Moreira da Silva. The biovid heat pain database data for the advancement and systematic validation of an automated pain recognition system. In *International Conference on Cybernetics (CYBCO)*, pp. 128–131, 2013.
- J. Wang and L. Kawka. Gimefive: Towards interpretable facial emotion classification. *CoRR*, abs/2402.15662, 2024.

- Y. Wang, S. Yan, Y. Liu, W. Song, J. Liu, Y. Chang, X. Mai, X. Hu, W. Zhang, and Z. Gan. A survey on facial expression recognition of static and dynamic emotions. *CoRR*, abs/2408.15777, 2024.
- Z. Wu and J. Cui. La-net: Landmark-aware learning for reliable facial expression recognition under label noise. In *ICCV*, 2023.
 - F. Xue, Q. Wang, and G. Guo. Transfer: Learning relation-aware facial expression representations with transformers. In *ICCV*, 2021.
 - F. Xue, Q. Wang, Z. Tan, Z. Ma, and G. Guo. Vision transformer with attentive pooling for robust facial expression recognition. *IEEE Transactions on Affective Computing*, 2022.
 - J. Yan, P. Li, C. Du, K. Zhu, X. Zhou, Y. Liu, and J. Wei. Multimodal emotion recognition based on facial expressions, speech, and body gestures. *Electronics* (2079-9292), 13(18), 2024.
 - S. Yang, J. Van de Weijer, L. Herranz, and S. Jui. Exploiting the intrinsic neighborhood structure for source-free domain adaptation. *NeurIPS*, 2021.
 - M.O. Zeeshan, M.H. Aslam, S. Belharbi, A.L. Koerich, M. Pedersoli, S. Bacon, and E. Granger. Subject-based domain adaptation for facial expression recognition. In *International Conference on Automatic Face and Gesture Recognition (FG)*, 2024.
 - Muhammad Osama Zeeshan, Marco Pedersoli, Alessandro Lameiras Koerich, and Eric Grange. Progressive multi-source domain adaptation for personalized facial expression recognition. *CoRR*, 2025.
 - D. Zeng, Z. Lin, X. Yan, Y. Liu, F. Wang, and B. Tang. Face2exp: Combating data biases for facial expression recognition. In *CVPR*, 2022.
 - S. Zhang, Z. Zhao, and C. Guan. Multimodal continuous emotion recognition: A technical report for abaw5. In *ABAW workshop at CVPR*, 2023.
 - Y. Zhang, R. Zhao, W. Dong, B.-G. Hu, and Q. Ji. Bilateral ordinal relevance multi-instance regression for facial action unit intensity estimation. In *CVPR*, 2018.
 - S. Zhao, G. Jia, J. Yang, G. Ding, and K. Keutzer. Emotion recognition from multiple modalities: Fundamentals and methodologies. *IEEE Signal Processing Magazine*, 38(6):59–73, 2021.
 - C. Zheng, M. Mendieta, and C. Chen. POSTER: A pyramid cross-fusion transformer network for facial expression recognition. In *ICCVw*, 2023.

Appendix

Table of Contents

A	To Reviewers: Please Read Before Downloading BAH Dataset	16
В	Related Work	17
C	BAH Dataset Limitations	
D	BAH Dataset File Structure	20
E	BAH Data Collection Web-platform	20
F	BAH Dataset Diversity	21
G	BAH Dataset Annotation Codebook	26
Н	Baseline Results	31
	H.1 Supervised Learning Cases (Followup from Main paper)	31
	H.2 Zero-shot Inference: Multimodal Large Language Models (M-LLMs)	32
	H.3 Personalization using Domain Adaptation	34

A TO REVIEWERS: PLEASE READ BEFORE DOWNLOADING BAH DATASET

We setup two strategies to directly download BAH for review purposes only:

- 1. Via our newly installed private and anonymous server at our university through this link https://142.137.245.13/index.php/s/MyY2GyzBwjNXFLq. For the first time usage, please accept the security certificate on your internet browser before proceeding we are working to fix that. If this fails, please use the second option.
- 2. Via Google drive link: https://drive.google.com/drive/folders/1VyR9G0fb4ez_RIZtFr3kaGmnndfKJBoI?usp=drive_link.

Use this password to unzip the file "BAH_DB-shared-public-no-comp-ICLR2026.zip": @ICRL_oSY5QhGTHH5ckAf3qKCF_2026_Brazil
For access issues, please reach out to the ICLR organizers so we can help you.

To ensure that we are compliant with our ethical requirements we are asking reviewers to read the "0_Read_before_downloading_data.pdf" (which is in the root of the files pointed by the download link of BAH dataset) before downloading the dataset. If you agree to the terms, please proceed to the dataset download. Here is the content of the file "0_Read_before_downloading_data.pdf":

```
To be consistent with the ethical requirements of the BAH dataset, by downloading the BAH dataset, you are agreeing to the following terms:

o Purpose of the access: You are accessing the BAH dataset as a part of a blind review process for the ICLR 2026 purposes.

o Use of the data: You will not redistribute, republish, or disseminate the BAH dataset.

o Duration of the access: The dataset must be securely destroyed after the review process has been completed.

If you have any concerns regarding these terms, please contact the ICLR 2026 organisers. "
```

The download link provides access to the dataset itself for review purposes. Additionally, we included other materials such as anonymized EULA, and request forms. We also include a presentation to our data collection platform www.aerstudy.ca.

867 868

864

866

RELATED WORK В

870 871

This section provides works in affective computing related to behavioural science.

872 873 874

875

876

877

878

879

880

881

882

883

884

885

886

887

889

890

891

892 893

894

895

896

897

898

a) Affect Recognition using Machine Learning:

Basic Emotions. An important line for ML research in affective computing is discrete emotion recognition in facial image (facial Expression Recognition – FER) (Bonnard et al., 2022; Liu et al., 2024a; Kollias et al., 2025; Lee et al., 2023; Mao et al., 2024; Wang et al., 2024; Wu & Cui, 2023; Xue et al., 2021; Zeng et al., 2022; Zheng et al., 2023). This usually involves classifying facial images into one of seven or eight basic emotions, such as 'Happy', 'Sad', and 'Surprised'. Other works focus on videos (Liu et al., 2023a; 2021a;b; 2023b) as well. There has also a recent interest in designing robust FER methods that are interpretable (Belharbi et al., 2024a;b; Wang & Kawka, 2024; Xue et al., 2022). They typically produce a heat map that points to relevant regions used by a model to perform a prediction. This is usually formulated as an attention map or a Class-Activation Map (CAM) (Choe et al., 2022; Murtaza et al., 2025). Other work aims to predict ordinal levels (i.e, ordered labels), including pain and stress estimation (Aslam et al., 2024; Chaptoukaev et al., 2023; Zeeshan et al., 2024; Nasimzada et al., 2024); a task that can be extremely useful in healthcare applications. Some datasets such as BioVid (Walter et al., 2013) rely on advanced and expensive modalities such as bio-signals to predict pain for instance. Dimension recognition of emotions, typically aims to estimate continuous valence and arousal values linked to emotions (Dong et al., 2024; Praveen & Alam, 2024a; Praveen et al., 2023; 2021). Finally, another task related in emotion recognition is Action Units (AUs) detection (Jacob & Stenger, 2021; Luo et al., 2022). It aims to predicting active AUs in the face under a multi-label classification framework. Other works go further to estimate the intensity of AUs (Fan et al., 2020; Zhang et al., 2018), or both (Sánchez-Lozano et al., 2018), a much more challenging task.

Compound Emotions. Real-world scenarios often present complex emotions that combine basic ones. There has been recent interest in building affective computing models to predict compound emotions, a case where a mixture of basic emotions are expressed (Kollias, 2023; Richet et al., 2024). These are show in several practical real-world application since such complex emotions occur in daily interactions. However, they are more difficult to recognize as they are subtle, ambiguous, and resemble basic emotions. A recent specialized video-based dataset named C-EXPR-DB (Kol-

899 900 901

Dataset	Affect	Modalities	Subject-based	Num. of participants	Num. of samples	Environment	Annotation
RAF-DB (Li et al., 2017)	Basic/compound emotions	Images	No	-	15,339 images	Wild	Image-level
AffectNet (Mollahosseini et al., 2019)	Basic emotions	Images	No	-	450k images	Wild	Image label
Aff-wild2 (Kollias & Zafeiriou, 2019)	Basic emotions, Valence/Arousal, Action Units	Video, audio	No	-	564 videos	Wild	Frame-level
MELD (Poria et al., 2019)	Basic emotions	Video, audio	No	-	13000 utterances	Actors/TV-show	Frame-level
C-EXPR-DB (Kollias, 2023)	Compound emotions	Video, audio	No	-	400 videos	Wild	Frame-level
UNBC-McMaster (Kollias & Zafeiriou, 2019)	Pain estimation	Frames	Yes	25	200 videos	Lab	Frame-level
BioVid (Walter et al., 2013)	Pain estimation	Frames, biomedical signals (GSR, ECG, and EMG at trapezius muscle)	Yes	90	18017 samples	Lab	Frame-level
RECOLA (Ringeval et al., 2013)	Apparent Emotional Reaction Recognition	video, audio, physiology (electrocardiogram, and electrodermal activity)	Yes	46	46 videos	Lab	Frame-level
SEWA (Kossaifi et al., 2019)	Apparent Emotional Reaction Recognition	video, audio	Yes	398	1,990 videos	Wild	Frame-level
WEMAC (Miranda Calero et al., 2024)	Discrete, dimensional emotions	Physiology (blood volume pulse, galvanic skin response, and skin temperature), audio	Yes	100	100 records	Lab	Self-reported
StressID (Chaptoukaev et al., 2023)	Stress	EDA, ECG, Respiration, Face video, Speech	Yes	65	587 videos	Lab	Frame-level
SchiNet (Bishay et al., 2019)	Estimation of Symptoms of Schizophrenia	video	Yes	91	91 videos	Wild	Video-level
MESC (Chu et al., 2024)	Emotional Support Conversation	video, audio, text	Yes	-	1,019 dialogues	Wild	Utterance-level
IEMOCAP (Busso et al., 2008)	Improvisations of scripted scenarios for basic emotions	video, audio, text	Yes	10 actors	-	Lab/Actors	Frame-level
BAH (ours)	Ambivalence/Hesitancy	Video, audio, transcript	Yes	224	1,118 videos	Wild	Video-level, Frame-level, A/H cues

914 915 916

Table 7: Common affective computing datasets for emotion modeling in health contexts.

lias, 2023) has been constructed for the design/evaluation of models. The dataset accounts for the difficulty of the task as different modalities are required to better recognize compound emotions.

Despite the recent progress in affect modelling, Ambivalence/Hesitancy recognition is still unexplored in ML. A possible reason is the lack of specialized dataset for training and evaluation of ML models. As it is implicated in healthcare and interventions, A/H is a common topic in behavioural science (Hohman et al., 2016; Manuel & Moyers, 2016). A/H recognition is related to compound emotion recognition task where intention and attitudes are conflicted or in a in-between state, between willingness and resistance (MacDonald, 2015), or positive and negative affect (Armitage & Conner, 2000). This can manifest in how an individual expresses them self and can be recognized (Hayashi et al., 2023) in their facial expression, tone, verbal, and body language. As a result, A/H exhibits a multimodal nature that comes as the result of subtle interconnection between different cues. Unfortunately, such discord is extremely difficult to spot; a task that requires human training. This is a tedious and expensive procedure, leading to ineffective and less scalable eHealth interventions under limited resources. Assisting healthcare providers with automatic, reliable and inconspicuous tools to help them recognize A/H can have a major impact in improving eHealth interventions.

Our BAH dataset fills in the gap in the literature, and to provide an important resource to design/evaluate ML models for A/H recognition task. It is a video Q&A dataset from which we extract audiovisual information with transcripts, offering multiple modalities. The dataset is fully annotated by behaviour science experts at video- and frame-level. In addition, cues used by annotators to recognize A/H at each segment are provided. This includes facial and vocal expressions, body language, language in addition to highlighting where there is inconsistency between the modalities. As shown in Table 7, our BAH dataset is competitive compared to existing affective computing datasets in terms of modalities, number and diversity of participants, and annotations. While no dataset matches the specific focus on A/H in digital health interventions, datasets like MESC (Chu et al., 2024), SchiNet (Bishay et al., 2019), and IEMOCAP (Busso et al., 2008) contain videos from interviews with psychological relevance. Therefore, BAH provides an important asset for the ML community to begin research in A/H recognition.

b) Behavioural Science:

Health-Related Behaviour Change and Non-Communicable Diseases. High-risk health behaviours, such as tobacco use, physical inactivity, unhealthy diets, and harmful alcohol consumption, are responsible for the vast majority of non-communicable diseases (NCDs), which include cardio-vascular disease, type 2 diabetes, cancer, and chronic respiratory illnesses. According to the World Health Organization (WHO) (Ortiz et al., 2025), NCDs account for approximately 74% of global deaths, and these outcomes are disproportionately influenced by modifiable behavioural factors. Evidence suggests that around 80% of chronic disease risk is attributable to these high-risk behaviours.

Consequently, health-related behaviour change has become a primary target for preventive and therapeutic interventions. Traditional methods, such as motivational interviewing (MI) and cognitive behavioural therapy (CBT), rely on face-to-face clinical interviews, which remain foundational to behavioural health practice (O'Donnell et al., 2019). These interactions provide unique opportunities for clinicians to detect ambivalence, hesitancy, and other complex affective states, often through subtle verbal and nonverbal cues (Hall et al., 1995). Despite the growing shift toward digital platforms, clinical interviews remain the gold standard for eliciting meaningful emotional and cognitive responses, insights that are essential to tailoring behaviour change strategies. Efforts to change health behaviours over the long term are inherently complex. Individuals often experience ambivalence and hesitancy, understood as fluctuating between intention and resistance, when attempting to adopt healthier lifestyles. In traditional healthcare contexts, providers rely on both verbal communication and non-verbal cues (e.g., tone, gestures, facial expressions) to recognize and address such motivational conflicts. This in-person interaction allows for nuanced support that can adapt to a patient's readiness for change (Davidson & Scholz, 2020). The purpose of developing multimodal A/H recognition systems is to capture and replicate this nuanced understanding of patient behaviour within digital health interventions, thereby supporting clinicians and scaling behavioural health care.

Multimodal Cues and the Detection of Complex Emotions. Identifying complex emotional states such as ambivalence, resistance, or hesitancy is crucial for tailoring behavioural interventions. Research in psychology and human-computer interaction has shown that complex emotional states,

 such as ambivalence, uncertainty, or defensiveness, are communicated through a combination of facial expressions, body posture, vocal tone, speech patterns, and physiological responses (Guo et al., 2018; Pantic & Rothkrantz, 2003). In digital contexts, however, the absence of physical presence makes this task more difficult. Recent research in psychology and computer science has focused on the use of multimodal cues, such as facial expressions, voice tone, body posture, and physiological responses, as proxies for emotional and motivational states (Kraack, 2024; Yan et al., 2024). These cues can reveal underlying emotional conflict or uncertainty that might not be captured by self-report alone. Studies have shown that combining multiple input channels (e.g., audio-visual data) can enhance the accuracy of emotion recognition systems. For instance, multimodal datasets are being used to train models that detect affective states like confusion, frustration, and mixed emotions, which are highly relevant in contexts such as education, mental health, and behaviour change. By incorporating these data streams, researchers can better approximate the nuanced human capacity for reading emotions, paving the way for emotionally aware systems (He et al., 2020; Zhao et al., 2021).

Affective Computing and Personalized Digital Health Interventions Affective computing, a subfield of artificial intelligence (AI) focused on recognizing, interpreting, and responding to human emotions, holds promise for advancing personalized digital health interventions. By leveraging emotion-aware algorithms, digital platforms can better understand users' psychological readiness and tailor support accordingly (Lokhande et al., 2024; Vairamani, 2024). For example, interventions that dynamically respond to detected signs of resistance or disengagement may improve user retention and behavioural outcomes. Incorporating affective computing into digital health technologies also allows dynamic tailoring of content based on users' real-time affect, responsive dialogue, mimicking the adaptability found in face-to-face interactions. Recent advancements in conversational agents, voice analysis, and facial expression recognition have made it possible for digital interventions to adapt content delivery based on real-time emotional assessments (Khanna et al., 2022). This not only improves user engagement but also enhances intervention effectiveness by ensuring messages are delivered in an emotionally congruent and contextually appropriate manner (Hornstein et al., 2023).

C BAH DATASET LIMITATIONS

While BAH dataset offers a novel contribution to emotion recognition for digital behaviour change, several limitations should be considered.

- 1) Data collection constraints. The web-based platform occasionally experienced technical issues, preventing some participants from completing all seven videos. As participants used their own devices in home settings, video and audio quality varied significantly despite clear instructions and testing. Response length was participant-determined, leading to high variability in content. Some environmental noise or visual distractions (e.g., background conversations, movement) were present in a subset of recordings.
- 2) Participant representation. Although participants were recruited from nine Canadian provinces with diverse age and ethnic backgrounds, individuals from under-resourced areas or without reliable

```
1 BAH_DB-shared-public/
2 — BAH_dataset_documentation.pdf
3 — bah-video.csv
4 — bah-video.json
5 — cropped-aligned-faces
6 — meta_data.yml
7 — readme.md
8 — split
9 — split-frames
10 — transcription
11 — video_annotation_transcript.yaml
12 — Videos
```

Figure 4: File structure of the shared BAH dataset.

internet access were likely underrepresented. Gender identity was collected but not used in sampling, and no data on socioeconomic status was recorded. Digital literacy and access may have biased participation toward more tech-savvy individuals.

3) Multimodal and data balance issues. The expressiveness of cues (facial, vocal, bodily, verbal) varied widely by participant, complicating consistent multimodal analysis. Though the dataset is balanced at the video level, frame-level imbalance exists (fewer A/H frames than non-A/H). Training strategies that account for class imbalance should be considered.

D BAH DATASET FILE STRUCTURE

Figure 4 shows the file structure of the shared BAH dataset. The file "BAH_dataset_documentation.pdf" contains the detailed documentation about all files/directory, including annotation structure.

E BAH DATA COLLECTION WEB-PLATFORM

Alongside the dataset files, we include a slide presentation of our "Automatic Expression Recognition" (AER) web-based platform (www.aerstudy.ca). The presentation is in the file "AER-web-platform.pdf". Figure 5 shows an example of the platform.

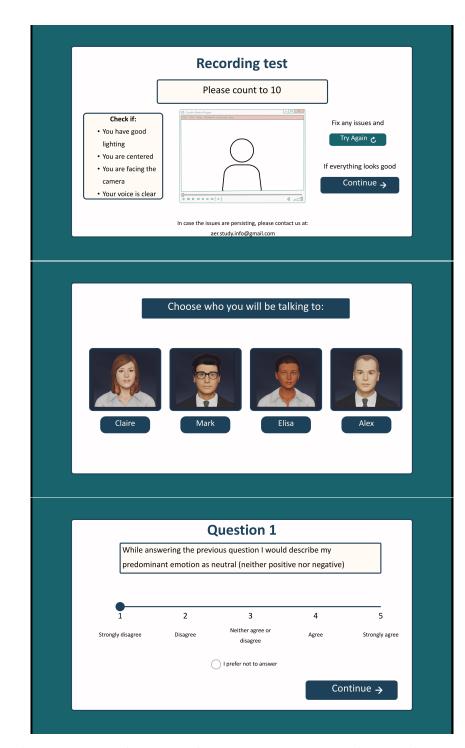


Figure 5: Examples taken from the platform to present our "Automatic Expression Recognition" (AER) web-based platform (www.aerstudy.ca).

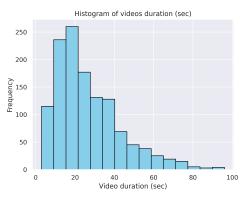
F BAH DATASET DIVERSITY

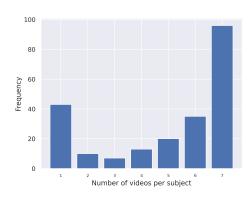
This section includes more statistics about BAH dataset to highlight its diversity. Figure 11 shows a general overview via a nutrition label. Overall, BAH dataset has significant diversity. It covers different Canadian provinces, age range, ethnicities, and male/female presence. It has a large number

of videos (1,118) where 638 videos contain A/H. Most asked questions elicited A/H, especially question-4 (Ambivalent). In addition, since we have less control over the participants, and their environment, the dataset is considered in-the-wild. On top of video and audio modality, we provide audio transcript which has shown to be an important modality for A/H recognition. BAH is fully annotated at video- and frame-level. Moreover, annotators report the used cues to recognize A/H at each segment. All these properties make our dataset a realistic and relevant asset to design ML model for the task of A/H recognition in videos.

We include the following general information:

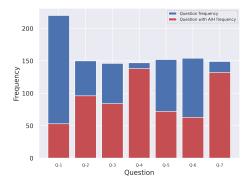
- Videos durations distribution (Figure 6a).
- Videos per participants distribution. (Figure 6b).
- Questions and A/H distribution (Figure 6c).
- A/H segments duration (Figure 7).





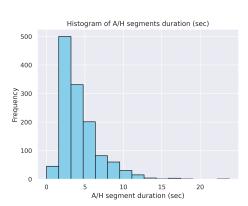
(a) Videos duration histogram.

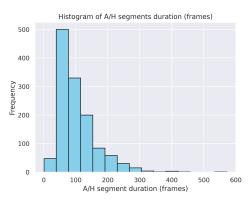
(b) Distribution of number of videos per participant.



(c) Distribution over 7 questions: Num. videos perquestion (blue), Num. videos with A/H (red).

Figure 6: Video duration (a), and videos/participant (b), and question distribution (c) over BAH dataset.





- (a) Distribution of A/H segment duration in seconds.
- (b) Distribution of A/H segment duration in frames.

Figure 7: Distribution of A/H segment duration in seconds (a), and frames (b) over BAH dataset.

In addition, more demographics statistics are included as well:

- Participants' age distribution (Figure 8).
- Participants' age range distribution (Figure 9a).
- Distribution of Canada provinces where participants live (Figure 9b).
- Participants' simplified ethnicity distribution (Figure 10a).
- Participants' student-status distribution (Figure 10b).
- Participants' consent to use their data in challenges distribution (Figure 10c).
- Participants' consent to use their data in publications distribution (Figure 10d).

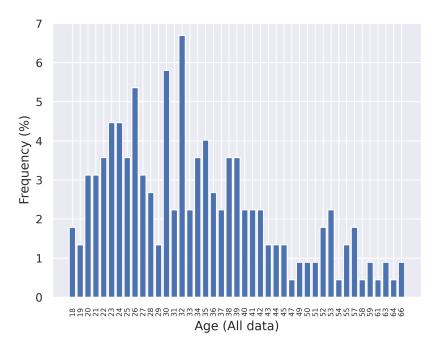
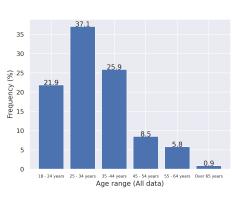
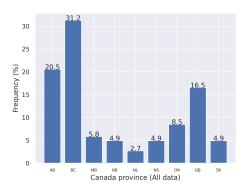


Figure 8: Participants' age distribution in BAH dataset.

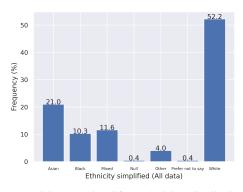


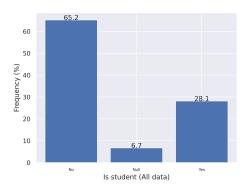


(a) Participants' age range distribution.

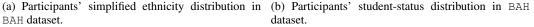
(b) Distribution of Canada provinces where participants live.

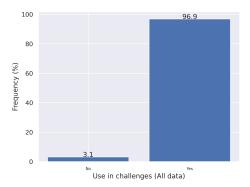
Figure 9: Participants' age range (a), and where the provinces where they live (b) over BAH dataset. Name of provinces: 'Manitoba (MB)', 'Alberta (AB)', 'Nova Scotia (NS)', 'Newfoundland and Labrador (NL)', 'Saskatchewan (SK)', 'New Brunswick (NB)', 'Ontario (ON)', 'Quebec (QC)', 'British Columbia (BC)'.

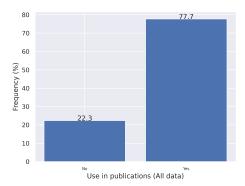




BAH dataset.







(c) Participants' consent to use their data in challenges distribution in BAH dataset.

(d) Participants' consent to use their data in publications distribution in BAH dataset.

Figure 10: Distribution of participants' simplified ethnicity (a), their student-status (b), their consent to use their data in challenges (c), and publications (d) over BAH dataset.

1296	
1297	Dati Detect Foote
1298	BAH Dataset Facts
1299	Dataset BAH (Behavioural Ambivalence/Hesitancy – A/H)
1300	Nature of Dataset A Dataset for Ambivalence/Hesitancy recognition in
1301	videos for participants recruited in Canada Participants Country Canada
1302	Number of provinces in Canada 9
1303	Provinces in Canada 'Manitoba (MB)', 'Alberta (AB)', 'Nova Scotia (NS)',
1304	'Newfoundland and Labrador (NL)', 'Saskatchewan (SK)', 'New Brunswick
1305	(NB)', 'Ontario (ON)', 'Quebec (QC)', 'British Columbia (BC)'. Number of participants 224
1306	Number of videos 1,118 where 638 videos contains A/H
1307	Video length 26.58 ± 16.36 (seconds) with a minimum and maximum
1308	duration of 3 and 96 seconds Total duration 8.26 hours where A/H duration is 1.5 hours
1309	Total number of frames 714,005 where 131,103 frames contains A/H
1310	Total number of A/H video segments 1,274
1311	Length A/H video segment 4.25 ± 2.47 seconds or 102.92 ± 59.16
1312	frames. The minimum and maximum A/H segment is 0.01 seconds (1
1313	frame), and 23.8 seconds (572 frames) Data capture web-platform www.aerstudy.ca
1314	
1315	Motivation
1316	Summary Behavioural Ambivalence/Hesitancy
1317	(BAH) is a dataset collected for subject-based multimodal recognition of
1318	A/H in videos. It contains videos from 224 participants captured across 9 provinces in Canada, with different age, and ethnicity. Through our web
1319	platform, we recruited participants to answer 7 questions, some of which
1320	were designed to elicit A/H while recording themselves via webcam with
1321	microphone. BAH amounts to 1,118 videos for a total duration of 8.26
1322	hours with 1.5 hours of A/H. Our behavioural team annotated timestamp
1323	segments to indicate where A/H occurs, and provide frame- and video-
1324	level annotations with the A/H cues. Video transcripts and their timestamps are also included, along with cropped and aligned faces in each frame, and
1325	a variety of participants meta-data.
1326	Original Authors Redacted for anonymity reasons.
1327	Metadata
1328	URL Redacted for anonymity reasons.
1329	Keywords Ambivalence/hesitancy, eHealth, digital health intervention,
1330	video, Deep Learning, Benchmark
1331	Available participants meta-data Age, birth country,
1332	Canada province where the participant lives, ethnicity, ethnicity simplified,
1333	sex, student status, consent to use recordings in publications Video format *.mp4
1334	Ethical Review Redacted for anonymity reasons.
1335	License Custom - for research purposes only.
1336	How to request the data? Fill in this form, sign, and upload the EULA -
1337	Redacted for anonymity reasons.
1338	First release 2025
1339	Annotation
1340	Annotators 3 experts in behavioural science
1341	Video- and frame-level Label "1" for presence of A/H, "0", its absence
1342	Cues provided by annotators for each A/H segment Facial expressions,
1343	body language, audio and language in addition to highlighting where there
1344	is inconsistency between the modalities
1345	Data size
1346	All files zipped (*.zip) 8.3 GB
12/17	0.0 GB

Figure 11: A data card styled (nutrition label) for ${\tt BAH}$ dataset.

G BAH DATASET ANNOTATION CODEBOOK

This section contains relevant information regarding our designed annotation codebook for A/H recognition. We provide the definitions of A/H and the types of cues (Table 8), as well as a more detailed description of the most relevant cues in each modality used to detect A/H, which include facial cues (Table 9), language cues (Table 10), audio cues (Table 11), body language cues (Table 12), and cross-modal inconsistency cues (Table 13). The codebook is a working document that continues to evolve in response to relevant insights emerging from expert annotations and contributions from behaviour change experts. Updates on the codebook will be made available and communicated upon request. This iterative approach aligns with established qualitative research practices, where coding frameworks are refined throughout the analysis process to better reflect the complexity and richness of the data (Bradley et al., 2007).

Term	Definition
Ambivalence/Hesitancy	The simultaneous presence of competing positive and negative feelings, ideas, thoughts, or emotions towards one same object or goal. A state in which a person has not entirely made up their mind about doing something; when they aren't fully decided on how to act (towards a behaviour or object; not necessarily the goal behaviour; excluding towards language or answering questions)
Facial Cues	Different motions of the muscles in the face. Facial expressions commonly occur around the mouth and eyes, including changes in a person's gaze. They can be used to assess a person's emotional state.
Language Cues	Includes verbal/speech-based expressions of ambivalence or hesitancy. Some common verbal expressions can include the use of 'I want to but', 'mmmm', among others.
Audio Cues	Changes in a person's non-verbal language, such as changes in tone, speed and pitch.
Body Cues	Non-verbal signals that include gestures, body posture and movements. Some of the cues that can be annotated as body language are hand movements, head tilts, shoulders shrugging and sighs (chest movement).
Cross-modal inconsistency Cues	Simultaneous incompatibility between two or more modalities or different types of cues. For example, this could be represented by someone saying 'yes' while shaking their head side to side.

Table 8: BAH dataset annotation codebook: definitions.

Facial cues	Definition	
Upper Region		
Close 2	A change in the frequency with which someone blinks or closing one's eyes for a longer period (e.g., either keeping them closed, or blinking for a long time). This excludes normal blinking, it is annotated when there is a difference compared to the participants baseline. Closing of both eyes; "blinking" (with both eyes).	
Close 1	Closing one eye at the time; includes winking. The duration of the wink is no relevant, it can be a quick wink or a longer one.	
Squint	Partially closing one or both eyes. Significant or identifiable changes or contractions in the muscles around the outer or inner corners of the eyes. It might involve some changes in the eyebrows, forehead and cheeks. Includes squinting eyes, muscles contracting around the eyes.	
Frown	To bring your eyebrows together (inner eyebrow) so that there are lines on your face above your eyes. Frown, forehead fold, small frown, tensed forehead, wrinkled forehead, furrowing brows, lowering inside corners of eyebrow	
Eyebrow	Lowering/raising external parts of eyebrow(s) (or full eyebrow(s)). [i.e., one of both eyebrows]	
Gaze	Changes in the direction of the gaze by moving the eyes. Moving eyes (not face to look down, up, to the side.	
Eye roll	Eye-rolling is a transitory gesture in which a person briefly turns their eyes upward, often in an arcing motion from one side to the other. The eyes do not se on anything in particular and go back to their previous position.	
Open	Opening the eyes, looks like an increase in awareness. Eyes look slightly bigger Engagement of the eyelids, contracting the eyelid muscles to make them look wider.	
Lower Region		
Smile	Ends of the mouth/lips curve up, often with the lips moving apart. Includes Smile, smirking, half a smirk, fake smile, raising both sides of the mouth, side smile, half smile	
Pout	Pushing one's lips or one's bottom lip forward; or turning the outer sides of the lips downwards. Pouting, pursed lips, "frowning" with one's mouth	
Lip press	Contracting or pressing lips without pushing them forward. Includes: pressed lips, pressing lips together, putting lips together. Excludes pressing lips to pout/purse.	
Hanging	Leaving one's mouth open for an extended time (e.g., hanging mouth, gaping mouth).	
Mouth	Any other movements of the mouth that (1) are not captured by smile, pout, or pressing lips, and (2) is not a result of the baseline speech patterns. Opening mouth, opened mouth, raising upper lip, rising one side of the mouth, taking the mouth corners back and lower them	
Wrinkle Chin	Moving the muscles around the chin to create identifiable lines, folds, ridges or furrows in the chin. Usually seen as a contraction of the chin muscles creating creases around or on the chin.	
Nose	Changes in the movement or looks of the nose. Includes significant movements on the nostrils, the tip of the nose, scrunching the nose, or any other muscle movement that would create a change in the nose.	

Table 9: BAH dataset annotation codebook: facial cues.

Language cues	Definition	
Filler sound	Sound made during a pause in speech signalling the person isn't done taking. Examples: "mmm", "umm", "hum", "emmm", "err", "uh", "ah"	
Filler word	Words used that do not contain substantive content, but are used as fillers to fill in space while the person thinks (or to signal they are not done talking, or that they are about to talk): "like", "you know", "I mean", "okay", "so", "actually", "basically"	
Hedging	Words/expressions used to express ambiguity about what one is saying (about to say or just said). Examples: "somewhat"; "I'm not an expert, but"; " right?"; " isn't it?"; "I do not know"; "all I know"; "I think"	
Correction	Corrects something they said. This focuses on the content of what is said, not on a syntax-based error, or speech error.	
Repetition	Emphasizing a phrase by repeating, or repetition of a word, might be related to trying to find the right word or expression	
Com-B Constructs		
Positive	Statement of positive feelings towards a behaviour or action.	
Negative	Statement of negative feelings towards a behaviour or action.	
Excuse	Statement where the participant mentions an excuse, a reason or justification for something that has happened or hasn't happened. It can also be an expression of regret for doing/not doing something. Use of 'but'. Shows avoidance or lack of responsibility	
Success	Statement of success with goal (focused on the behaviour)	
Fail	Statement of failure with goal (focused on the behaviour)	
Cap	Mentions having the capability to change their behaviour. Includes physical capability (e.g., balance, dexterity) or psychological capability (e.g., knowledge, skills, memory).	
No cap	Mentions NOT having the capability to change their behaviour. Includes physical capability (e.g., balance, dexterity) or psychological capability (e.g., knowledge, skills, memory).	
Mot	Mentions having motivation to change their behaviour. Includes reflective motivation (e.g., making plans, having positive attitudes/beliefs) or automatic motivation (e.g., desires, habit, feelings)	
No mot	Mentions NOT having motivation to change their behaviour OR motivation NOT to change their behaviour. Includes reflective motivation (e.g., making plans, having positive attitudes/beliefs) or automatic motivation (e.g., desires, habit, feelings).	
Opp	Mentions having opportunity to change their behaviour. Includes physical opportunities (e.g., access to financial resources, location, time) or social opportunity (e.g., support/encouragement from others; norms)	
No opp	Mentions NOT having opportunity to change their behaviour. Includes physical opportunities (e.g., access to financial resources, location, time) or social opportunity (e.g., support/encouragement from others; norms)	

Table 10: BAH dataset annotation codebook: language cues.

Audio cues	Definition
Pause	Briefly interrupting a sentence by having silent pauses in between words or ideas that differ from the usual pace of how the participants speaks. It includes silent pauses, paused speech or ideas.
Cut words	Ending speaking a word before completing the utterance of the word (e.g., say "exer" instead of "exercise"). Breaking the words or interrupting the words while they are being spoken. Might involve correcting syntax/speech
Slow	Reducing the speed of speaking. There is a perceptible change in the speed while someone is talking, making it slower or de-accelerated. It differs from paused speech or cutting off words since the words, phrases or ideas are not cut off or left in the middle, there are no significant silences in the answers. It can include elongating syllables or words. Speed change is determined in comparison to the person's own baseline.
Fast	Changes in the speed of the answers, making it faster. Information is given quickly, briskly or lively. Speed change is determined in comparison to the person's own baseline.
Volume	Changes in volume of speech. Differences in how loud or quiet an answer is shared, or there can be differences in the volume of specific words or syllables. Includes: Raising volume, lowering volume, high volume, low volume, and mumbling. Volume change is determined in comparison to the person's own baseline.
Shaky voice	When there is an rapid fluctuation or trembling of rhythm or tone (i.e., there is instability) to the way someone is speaking. It includes voice shaking, quivering in voice. Excludes case when shaking is due to laughing
Breath	Audible breath, inhaling or exhaling, it can be while the person is talking or before/after a phrase. It includes changes in the breathing rhythm, intensity or deepness of the person (compared to the baseline) that create a sound. Includes sigh, deep breath
Click	Quick sound made by pressing the tongue against the roof of the mouth or back of the teeth and snapping it downward. It often signals disapproval, unsureness or impatience. The sound resembles a "tsk" or "tsk-tsk."
Laugh	Engaging in laughter, or variations thereof (e.g., snicker, chortle, giggle)
Stuttering	Involuntary repetition of sounds while speaking. This can be seen as a disruption or blocking of the speech by prolongation sounds or by struggling to say a word or a part of a word. Even though the stutter might cut off a word or phrase it is different since the person will finish the word or idea. Includes stammer, stumble.

Table 11: BAH dataset annotation codebook: audio cues.

Body cues	Definition
Look away	Moving the orientation of the head away from the baseline position such that eyes or the gaze will look away. Includes the head facing down, head facing up, looking down, looking up, looking from side to side, lowering head, raising head.
Shake	Turning the head from side to side, it can be done with repetitive head movements or with a slight turn of the head to one or both sides. Includes shaking head "no". Rotation is on the horizontal plane
Tilt	Angling the head to the side without focusing on something else, and holding the position. Changing the position of the head so it is in a sloping position. It can be accompanied by changes in the gaze but not necessarily. Includes head tilting up and down, tilting head to the side, tilted head. Includes bobbling head.
Throw	Throwing the head in a rapid movement in a particular direction.
Sigh	Movements of the chest, shoulder or head that accompany a sigh or a deep breath. It includes long sigh, deep breath, sigh, big sigh. Noticeable bringing the chest or diaphragm muscles up and down. Change determined in comparison to the person's own baseline.
Nod	Moving the head up and down. Lowering and raising the head, it can be done by slight or clearly marked movements. Includes movements such as back and forward or a single small nod.
Shrug	Raising of the shoulders, it can be a momentary or slight rise or a longer movement where one or both shoulders is raised. It includes shrugging shoulders, shrugs
Hands	Movements or placement of the hands that differs from baseline
Posture	Movements in the overall positioning of the spine, body or arms (independent from the head). The changes are determined by each person's baseline. It includes movements like readjusting in the seat, sloughing, turning to the sides. Needs to involve more than just the head. Excludes shrugging.
Scratch	Movements in the hands and arms to scratch or caress another part of the body or face. It includes scratching head, scratching neck, scratching eyes, scratching chin
Restless	Rhythmic and repeated movements. Can be swaying, shaking, being jittery.

Table 12: BAH dataset annotation codebook: body cues.

Cross-modal inconsistency	Definition
FL	Face and language/speech do not match. E.g., looking uncomfortable while saying yes, looking annoyed or uncomfortable while saying they are happy, smiling while saying they are worried.
FA	Face and audio do not match. E.g., speaks in a sad, energetic tone while smiling.
FB	Face and body do not match. E.g., Nodding while looking afraid or concerned, showing disgust but leaning forward
LA	Language/speech and audio do not match. E.g., speaks in a sad, energetic tone while saying they are happy.
LB	Language/speech and body do not match. E.g., seems like they are about to say something but do not, nod is discrepant with verbal speech, shaking head while saying yes
AB	Body and language/speech do not match. E.g., unengaged tone while nodding (in agreement)

Table 13: BAH dataset annotation codebook: cross-modal inconsistency cues - occurring simultaneously.

H BASELINE RESULTS

H.1 SUPERVISED LEARNING CASES (FOLLOWUP FROM MAIN PAPER)

Training details. For both cases, pretraining visual models on RAF-DB (Li et al., 2017), AffectNet (Mollahosseini et al., 2019), and Aff-wild2 (Kollias & Zafeiriou, 2019), their finetuning, and final training with multimodal setup, we used a learning rate between 0.0009 to 0.001 with multiplying coefficient of 10. When training on BAH, and in the case of using context, we used a window size between 24 (1 second) to 2880 (2 minutes) with a step of 1 second (24 frames). In this case, we use a mini-batch size in $\{2,4,8,16,32\}$, where a sample in the mini-batch is a window of frames; and a single-GPU training. In all trainings, we used a weight decay of 0.0001. All our experiments were conducted on a server with 4 NVIDIA A100 GPUs with 40 GB of memory, AMD EPYC 7413 24-Core Processor, and 503GB of RAM. We present in Table 14 the computation time of the multimodal case.

Ablation over the window length. We conduct an ablation to study the impact of the context (window length) on the performance of recognizing A/H. To this end, we use a window length from 24 to 3264 with a step of 1 second (24 frames). Figure 12 shows the obtained results. By considering WF1, performance improves with the increase of the context where it can reach above 0.825. However, F1 does not necessarily improve with the increase of the context. On the other hand, AP prefers small context. Using a small context of few seconds could be a good compromise for all the metrics. Note that the average length of an A/H segment is around 4 seconds (96 frames).

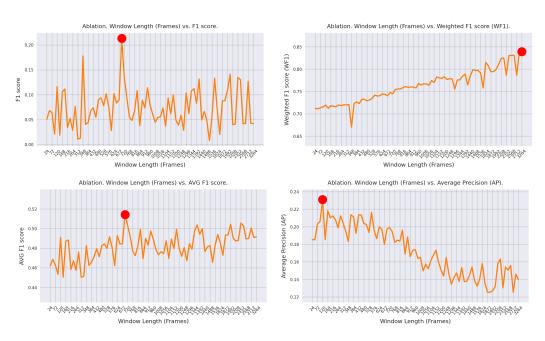


Figure 12: Impact of context (window) length on the performance of frame-level classification when using visual modality alone (ResNet18): F1, WF1, AVGF1, and AP. Best performance is indicated in red dot.

Video-level classification. To obtain video level predictions, we resort to a simple post-processing of frame-level predictions in the main paper. In particular, we follow (De-la Torre et al., 2015), where a sliding window averages the probability of the positive class at each frame. The case where high probability in a window suggests the presence of A/H in the video. In our case, we consider a context of 1 second (24 frame). The maximum probability across all windows is considered the video probability to have the positive class. The probability of the negative class is the complement of the probability of the positive one. We then proceed to measure the same performance metrics. Similar to the main paper, we report the performance at video-level for different cases: using visual backbone only without and with context (Table 15), multimodal (Table 16), and fusion (Table 17).

Case	Value
Train time 1 epoch	\sim 5mins
Inference time per-frame	~ 0.12 ms
Total n. params.	$\sim 223M$
N. learnable params.	$\sim 5 \mathrm{M}$
N. FLOPs	\sim 1.87 TFLOPs
N. MACs	\sim 938 GMACs

Table 14: Computation time, number of parameters, number of FLOPs/MACs for multimodal case with visual, audio and text (with ResNet152 for visual backbone). Visual backbone is frozen, while audio and text backbones are used to extract features offline and store them.

	Without context			With context (TCN)				
Backbone	F1	WF1	AVGF1	AP	F1	WF1	AVGF1	AP
APViT (Xue et al., 2022)	0.4069	0.6659	0.5947	0.5100	0.5154	0.4295	0.4531	0.3900
ResNet18 (He et al., 2016)	0.0000	0.5628	0.4081	0.2872	0.4263	0.6370	0.5739	0.4519
ResNet34 (He et al., 2016)	0.0000	0.5628	0.4081	0.3194	0.0000	0.5949	0.4165	0.2934
ResNet50 (He et al., 2016)	0.0394	0.5771	0.4293	0.6070	0.3793	0.7058	0.6079	0.4457
ResNet101 (He et al., 2016)	0.3481	0.6644	0.5774	0.4375	0.4712	0.2811	0.3381	0.4236
ResNet152 (He et al., 2016)	0.0765	0.5892	0.4483	0.4014	0.0000	0.5620	0.4075	0.3585

Table 15: Visual modality performance on test set of BAH at video-level classification: impact of architecture and context.

Similar to frame-level results, using context and multimodal yields better performance. In addition, simple concatenation using fusion yields the highest performance.

Modalities	F1	WF1	AVGF1	AP
Visual	0.0000	0.5620	0.4075	0.3585
Audio	0.0084	0.5379	0.3923	0.2293
Text	0.3458	0.5671	0.5007	0.2442
Visual + Audio	0.5200	0.5105	0.5131	0.5787
Visual + Text	0.1627	0.5946	0.4759	0.3227
Audio + Text	0.0861	0.5399	0.4039	0.2236
Visual + Audio + Text	0.0129	0.5175	0.3662	0.2123

Table 16: Multimodal models performance on test set of BAH at video-level classification. For visual modality, ResNet152 backbone is used.

Fusion type	F1	WF1	AVGF1	AP
LFAN (Zhang et al., 2023) (cvprw,2023)	0.0129	0.5175	0.3662	0.2123
CAN (Zhang et al., 2023) (cvprw,2023)	0.4906	0.6279	0.5868	0.3234
MT (Waligora et al., 2024) (cvprw,2024)	0.4800	0.1825	0.2643	0.3471
JMT (Waligora et al., 2024) (cvprw,2024)	0.4727	0.3014	0.3527	0.3068

Table 17: Feature fusion performance on test set of BAH at video-level classification.

H.2 ZERO-SHOT INFERENCE: MULTIMODAL LARGE LANGUAGE MODELS (M-LLMS)

Multimodal LLMs (M-LLMs) have gained significant attention in the affective computing space due to their ability to infer cross-modal dynamics across the visual, aural, and textual modalities. The problem of detecting ambivalence and hesitancy in videos is inherently multimodal as it requires also capturing the cross-modal inconsistency. To get out-of-the-box performance of existing SOTA M-LLM, we performed zero-shot inference using the 'Video-LLaVA-7B-hf' (Lin et al., 2024). Since the performance of an M-LLM or LLMs in general can be heavily influenced by the query prompt, we experiment with different variations of the prompts. Table 18 summarizes the different prompt variations used for zero-shot inference.

	Prompt	
Cimalo	'Classify the emotion in the video as either 'Non-Ambivalent'	
Simple	or 'Ambivalent'.' Respond with only one word: '	
	'Definition: Ambivalence is the state of having contradictory or	
Definition 1	conflicting feelings or attitudes towards something or someone	
Demittion 1	simultaneously. Classify the emotion in the video as either	
	'Non-Ambivalent' or 'Ambivalent'. Respond with only one word: '	
	'Definition: Ambivalence and Hesitancy is understood as the	
Definition 2	simultaneous experience of desires for change and against change.	
Delimition 2	Classify the emotion in the video as either 'Non-Ambivalent' or	
	'Ambivalent'. Respond with only one word: '	
	'Video transcript: {transcript}. Definition: Ambivalence is the	
	state of having contradictory or conflicting feelings or attitudes	
Transcript + Def 1	towards something or someone simultaneously. Classify the	
	emotion in the video as either 'Non-Ambivalent' or 'Ambivalent'.	
	Respond with only one word: '	
	'Video transcript: {transcript}. Definition: Ambivalence and	
	Hesitancy understood as the simultaneous experience of desires	
Transcript + Def 2	for change and against change. Classify the emotion in the video	
1	as either 'Non-Ambivalent' or 'Ambivalent'.	
	Respond with only one word:'	

Table 18: Summary of prompt variations for zero-shot inference.

H.2.1 Frame-Level Prediction

For frame-level prediction, we adopt a segment-wise strategy, where the entire video is divided into 8-frame chunks and passed through the model using a sliding window. This way, the model sees all the frames in each video. A single prediction is obtained for the window, which is replicated for the segment to match the total number of frame labels in each video. The model's output, 'Non-Ambivalent' or 'Ambivalent', is mapped to 0 and 1 respectively to match the ground truth.

Prompt	F1	WF1
Simple	0.0000	0.7148
Definition Only 1	0.1360	0.6889
Definition Only 2	0.3296	0.0652
Transcript + Def 1	0.3604	0.5057
Transcript + Def 2	0.3293	0.0915

Table 19: Frame Level Prediction using M-LLM.

Table 19 shows the results obtained for frame-level predictions using different prompts. The simple prompt where no context about the definition of A/H or the transcript is provided, the model predicts everything to be 'Non-Ambivalent', resulting in a higher WF1 but F1 score of zero. Adding just the definition of A/H to the prompt increases the performance as it provides the model with a little bit of context about about to look for in the video. The best results for frame-level prediction are obtained using the 'Transcript + Def 1' prompt, where the actual transcript of the video is also provided, along with a straightforward definition of A/H.

In the segment-wise approach applied with a sliding window of 8 frames, the model essentially sees every frame, but this approach limits the context window to be 1/3 of a second, which may not be enough to capture the temporal dependencies in the visual modality. We investigate the effect of various lengths context windows on the overall performance.

We selected the best-performing query prompt from the first experiment (Table 19) to perform the ablation on the context window size. Table 20 shows the results with different lengths of context window for the visual modality. 24 frames represent a one-second context window. Increasing the context window size to 120 frames (5 seconds) only marginally improves the overall performance of

Context Window	F1	WF1
24 Frames	0.3591	0.5048
48 Frames	0.3613	0.5051
80 Frames	0.3591	0.5040
120 Frames	0.3631	0.5053
192 Frames	0.3618	0.5095

Table 20: Performance comparison with increasing size context window for frame-level prediction.

the model, and it plateaus at 120 frames and then starts to drop which is an indicator that the visual encoder starts losing information with a longer context window.

H.2.2 VIDEO-LEVEL PREDICTION

For video-level prediction, the entire video is fed to the model, and the transcript is embedded in the prompt. The model selects 8 uniformly spaced frames from the video and predicts a single output. Similar to frame-level predictions, the model's output is mapped to 0 and 1, and the performance metrics are calculated.

Prompt	F1	WF1
Simple	0.0000	0.2191
Definition Only 1	0.1836	0.3111
Definition Only 2	0.7575	0.4618
Transcript + Def 1	0.7233	0.6787
Transcript + Def 2	0.7504	0.4766

Table 21: Video Level Prediction using M-LLM.

Table 21 presents the video-level prediction results. Similar to frame-level predictions, the 'simple' prompt without any context on the definition or the transcript performs the worst and predicts all samples to be 'Non-Ambivalent'. A similar trend is also observed here, i.e., adding the definition and the transcript substantially affects the model performance.

H.2.3 ANALYSIS AND DISCUSSION

The performance of M-LLM with zero-shot inference is substantially influenced by the query prompt. As observed from tables 19 and 21, simply asking the model to predict emotion based on the visual modality only performs the worst, whereas adding only the definition of A/H in the query prompt helps the model better identify the positive(A/H) class. Best results in all cases are obtained with the introduction of the text transcript of the video in the query prompt. We conjecture that this happens for two reasons: i) the textual modality serves a significant role in the identification of the A/H class, and ii) the current M-LLMs' performance is heavily reliant on the textual modality. This coheres with the overall structure of traditional M-LLMs that are built upon well-trained LLMs with the addition of a visual encoder like ViT, which is used to encode the visual information that is fed to the LLM for downstream tasks. Intuitively, the performance should increase with careful fine-tuning on the BAH dataset.

Further, the idea of textualizing the aural and visual modalities explored in (Richet et al., 2024) can be well-suited for a task like this where the audio and visual modalities essentially summarize the cues detected in the corresponding modalities. Particularly for tasks like subtle emotion recognition or the detection of A/H, where cross-modal inconsistency has to be considered. Textualizing the aural and visual modalities can be done to adequately exploit the reasoning abilities of SOTA LLMs.

H.3 Personalization using Domain Adaptation

Domain adaptation (DA) (Han et al., 2020; Li & Deng, 2018) has emerged as a promising approach for personalized expression recognition, where the model is trained on diverse labeled source data

to generalize to unlabeled target domains representing individual users. Recent research emphasizes on subject-based domain adaptation (Zeeshan et al., 2024; 2025; Sharafi et al., 2025), where each individual is defined as a distinct domain. DA will be employed to personalize ML models, by considering each participant in the test set as a separate target domain.

Experimental Protocol. For personalized in BAH, we adopt the standard protocol from prior work (Zeeshan et al., 2024; Sharafi et al., 2025), which involves partitioning the data of each target individual into train, validation, and test sets. Given the class imbalance in the BAH dataset, we ensure a balanced representation of positive and negative samples within each split. Sample counts per class and split for each subject are detailed in Table 23. Note that participants 25 (82609) and 55 (82864) were excluded due to a lack of positive samples. We establish the following baseline methods to evaluate the effectiveness of personalized BAH recognition: Source-only: The model trained on the source data is directly evaluated on the target individual test set without any adaptation. This assesses the generalization capability of the source model. Unsupervised Domain Adaptation (UDA): Source data with labels is utilized to adapt the model to each target individual using unlabeled data. This explores the potential of leveraging source knowledge for personalization in the absence of target labels. Source Free Unsupervised Domain Adaptation (SFUDA): Adaptation is performed solely using the unlabeled data from the target individual, without access to the source data. This examines the feasibility of personalization when the source data is unavailable. Oracle: The model is fine-tuned using the labeled data from the target individual during training. This provides an upper-bound performance, representing a fully supervised model.

Visual backbone. We employ a ViT-based model for personalization, leveraging its superior performance over ResNet-based architectures for visual tasks without contextual information. In all our experiments, we utilize a ViT-based model pre-trained on the source data.

H.3.1 UNSUPERVISED DOMAIN ADAPTATION:

We investigated two unsupervised domain adaptation (UDA) approaches for personalized BAH recognition: (i) a discrepancy-based method using Maximum Mean Discrepancy (MMD) (Sejdinovic et al., 2013) to minimize the domain gap and improve performance on the target subject, and (ii) a subject-based (Zeeshan et al., 2024) method using self-supervision that trains the model by generating pseudo-labels for the target domain, followed by reducing the domain shift using MMD that aligns source and target.

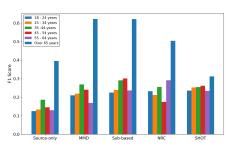
Implementation detail. In UDA experiments, we optimize our model using Stochastic Gradient Descent (SGD) (Sutskever et al., 2013) with a learning rate of 2e-4, momentum of 0.9, weight decay of 5e-4, and a cosine annealing scheduler (Loshchilov & Hutter, 2017) with a minimum learning rate of 2e-5. We set the batch size to 128 and run each target adaptation for 10 epochs. For the subject-based method, we introduce a hyperparameter $\gamma_3=0.01$ to weight the target loss, computed using pseudo-labels generated by the Augmented Confident Pseudo-Label (ACPL) technique (Zeeshan et al., 2024). This weighting is essential for mitigating noise in the pseudo-labels, in conjunction with a confidence threshold of 0.95 that is updated every 4 epochs.

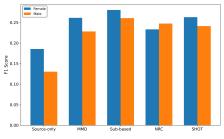
Methods	F1	WF1	AP
Source-only	0.1547 ± 0.1608	0.6814 ± 0.1687	0.2462 ± 0.1665
UDA (MMD) (Sejdinovic et al., 2013)	0.2418 ± 0.1513	0.6494 ± 0.1484	0.2608 ± 0.1685
UDA (Sub-Based) (Zeeshan et al., 2024) (fg,2024)	0.2674 ± 0.1475	0.6461 ± 0.1534	0.2673 ± 0.1642
SFUDA (SHOT) (Liang et al., 2020) (icml,2020)	0.2487 ± 0.1189	0.6537 ± 0.1142	0.25311 ± 0.0970
SFUDA (NRC) (Yang et al., 2021) (neurips,2021)	0.2400 ± 0.1158	0.6370 ± 0.1221	0.2401 ± 0.1367
Oracle	0.3699 ± 0.2048	0.7394 ± 0.0829	0.3932 ± 0.2171

Table 22: Performance of UDA and SFUDA with Source-only and Oracle on BAH.

H.3.2 SOURCE FREE UNSUPERVISED DOMAIN ADAPTATION

Two source-free unsupervised domain adaptation (SFUDA) approaches were explored for personalized BAH recognition: (i) a representation learning strategy inspired by hypothesis transfer (Liang et al., 2020), where information maximization was used to adapt the model to the target domain, and target-specific prototypes guided pseudo-labelling for class-level alignment, and (ii) a





- (a) Comparison between different age groups on DA methods.
- (b) Comparison between different DA methods based on participant Sex.

Figure 13: Comparison between different age groups and sex on DA methods.

neighbourhood-based method (Yang et al., 2021) in which label consistency was encouraged among target features and their reciprocal nearest neighbours, while expanded neighbourhoods were used to aggregate local structure and reduce the impact of noisy supervision through self-regularization and affinity-weighted loss.

Experimental protocol. We optimize the model using SGD with a learning rate of 1e-4, momentum of 0.9, and weight decay of 1e-3. The model is trained for 30 epochs with a batch size of 64. For NRC-based adaptation, we maintain memory banks of target features and predictions to retrieve K=3 nearest neighbors and M=2 expanded neighbors.

H.3.3 RESULT AND ANALYSIS

The average performance across 54 target participants, along with the standard deviation, is presented in Table 22. All reported results are based on evaluations of the respective target test sets. Our analysis demonstrates the effectiveness of domain adaptation for personalized detection of the A/H class in the BAH dataset. All tested methods surpass the Source-only baseline in F1 and AP positive-class metrics. Notably, Sub-based achieves the highest F1 (0.2674) and AP (0.2674), outperforming other domain adaptation techniques. While MMD and SHOT show moderate improvements, they lag behind Sub-based, highlighting the benefit of pseudo-labeling for enhancing minority-class recall and precision. Even NRC, despite some sensitivity to domain shifts, exceeds the Source-only performance. The Oracle upper bound (F1: 0.3699) underscores the substantial potential for further advancements in positive-class detection within this context. We observed a decrease in WF1 for leading positive-class detection methods like Sub-based (0.6461) and MMD (0.6494), relative to the Source-only baseline (0.6814), which represents a justifiable trade-off. The inherent bias of WF1 towards the majority negative class (comprising 80% of the data) penalizes approaches that prioritize minority-class recall. Even slight degradations in negative-class performance disproportionately impact the overall AP score. For example, Sub-based method emphasizes enhancing positive class identification, likely incurs a cost in precision or recall on the more frequent negative class, a necessary compromise to effectively detect the A/H class.

Sex-wise analysis. In the Figure 13b, we can observe that the *female* category generally exhibits higher values across most methods compared to the *male*. Specifically, *female* shows the highest values in **Sub-based** (0.28), **SHOT** (0.26), and **MMD** (0.26). It can also be noted that the number of female subjects (24) is less than the male subjects (29), as shown in Table 23.

Age-wise analysis. Figure 13a illustrates the varying impact of age on DA methods. The *Over 65 years* age group generally exhibits higher performance across all methods. However, it can be noted that it contains fewer subjects (2) as shown in Table 23, suggesting that the observed performance trends might not be fully representative of the broader population in this age range. Generally, **MMD** tends to increase with age, reaching its highest value in the 35 - 44 years group (0.269). **Subbased** generally shows a trend of increasing values with age, peaking in the 45 - 54 years age group (0.301). **NRC** shows more variability across age groups, with the highest value in the 55 - 64 years group (0.291) and the lowest in the 45 - 54 years group (0.175). **SHOT** shows relatively consistent performance across age groups, with a slight peak in the 45 - 54 years group (0.261). **Source-only** has the lowest values across all age groups, with a slight increase in the 35 - 44 years group (0.186).

 Individual performance. Tables 24 26, 25, show the individual performance of each subject on DA methods compared to source-only and oracle baselines, using F1, WF1, and AP metrics. Notably, 13 subjects achieve an F1 score of 0 in the source-only setting but exhibit higher WF1 scores. This indicates that performance gains for these subjects rely heavily on the dominant neutral class and do not reflect accurate classification of the A/H classes. The high WF1 in source-only is likely due to the dataset imbalance (80% neutral samples), where even predicting all samples as neutral yields a high WF1. In contrast, DA methods aim to improve the prediction of the A/H classes, leading to higher F1 and AP scores, even if it means a slight decrease in WF1 due to the reduced accuracy on the neutral class. This trade-off aligns with the primary goal of recognizing the A/H classes in the dataset. Specifically, MMD improved the F1 score for 43 subjects, Sub-based improved the F1 score for 44 subjects, SHOT for 40, and NRC for 38, compared to source-only. For AP, MMD improved performance for 30 subjects, Sub-based for 38, SHOT for 34, and NRC for 27, compared to source-only.

1998	Tar Sub	ID	Sex	Age Group	Train A/H	Train N	Val A/H	Val N	Test A/H	Test N
1999	P1	82557	Male	18-24	240	3383	241	275	120	915
2000	P2		Female	25-34	644	3850	320	320	241	1043
2001	P3	82711	Male	25-34	3174	1402	326	326	876	432
2002	P4	82565	Male	35-44	316	1827	153	153	117	496
2003	P5	82615	Male	35-44	194	1906	150	150	86	514
2004	P6	82594	Male	55-64	165	2424	165	203	83	657
2005	P7 P8		Female Female	25-34 25-34	733 160	1109 1996	131 154	131 154	217 78	310 538
2006	P9		Female	25-34 35-44	439	1740	155	155	76 149	336 474
2007	P10		Female	45-54	134	888	72	72	52	240
2008	P11		Female	35-44	447	2270	194	194	160	617
2009	P12	82564	Female	35-44	433	2779	229	229	166	752
	P13	82568	Male	18-24	235	1282	108	108	86	348
2010	P14	82569	Male	45-54	454	4911	383	383	209	1324
2011	P15	82573	Male	25-34	251	1607	132	132	96	435
2012	P16	82578	Male	35-44	69	1417	70	142	35	390
2013	P17 P18	82583 82585	Male Male	35-44 25-34	243 162	4041 2148	243 163	369 167	122 81	1102 580
2014	P18 P19		Female	25-34 25-34	156	3376	157	347	78	932
2015	P20	82591	Male	18-24	211	969	84	84	74	263
2016	P21		Female	25-34	185	2718	185	229	92	738
2017	P22	82593	Male	45-54	108	1440	109	111	54	389
2018	P23		Female	25-34	410	4448	347	347	189	1200
2019	P24		Female	55-64	644	3877	322	322	242	1050
2020	P25	82609	Male	35-44	0	1320	0	87	0	352
2021	P26 P27		Female Female	35-44 45-54	653 736	4572 3104	373 274	373 274	256 253	1237 844
2022	P28	82637	Male	35-44	215	5432	215	215	107	1507
2023	P29	82649	Male	18-24	937	3646	327	327	316	994
2024	P30	82666	Female	55-64	262	3478	262	272	131	938
2025	P31	82671	Female	45-54	335	2187	180	180	129	592
2026	P32		Female	25-34	714	1247	140	140	214	347
2027	P33		Female	35-44	1644	1507	225	225	468	433
2028	P34		Female	25-34	3136	4477	543	543	920	1255
2029	P35 P36	82694 82696	Male Female	25-34 55-64	1386 434	2014 1075	242 107	242 107	408 135	564 296
	P37	82708	Male	25-34	341	895	88	88	107	246
2030	P38	82710	Male	35-44	496	1242	124	124	155	342
2031	P39	82714	Male	Over65	819	806	116	116	234	231
2032	P40	82722	Male	Over65	1789	1058	203	203	498	316
2033	P41	82729	No	18-24	329	4786	329	401	164	1298
2034	P42	82735	Male	25-34	448	6037	448	478	224	1629
2035	P43 P44		Female	18-24 18-24	136 94	4876 2115	137 95	488	68 47	1342 584
2036	P44 P45		Male Female	25-34	306	6867	93 307	219 717	47 153	384 1897
2037	P46		Female	18-24	235	2937	226	226	116	791
2038	P47	82755	Male	18-24	531	2911	245	245	194	789
2039	P48	82764	Male	25-34	30	1882	31	161	15	511
2040	P49	82775	Male	35-44	100	1317	101	101	50	355
2041	P50		Female	35-44	408	3974	313	313	181	1072
2042	P51		Female	35-44	17	1310	17	107	9	354
2043	P52	82801	Male	25-34	1047	3898	353	353	350	1063
2044	P53 P54	82851 82861	Male Female	18-24 18-24	168 174	3220 3352	168 175	316 327	84 87	885 921
2045	P55	82864	Male	18-24	0	1043	0	69	0	278
2046	P56	82883	Male	25-34	1045	4537	398	398	361	1234
2047	Total	_	29/24	_	29172	150930	11424	13538	10415	40958
4VT1	-									

Table 23: Frame counts of Ambivalence/Hesitancy (A/H) and Neutral (N) classes across training, validation, and test splits for each target subject, with age groups added from metadata. Participants lacking A/H samples are highlighted. 'P' represents Participant. Total number of participants according to Sex 'Male/Female'.

2052	Tou Cub	Course only	MMD	Cub Dogod	CHOT	NRC	Omasla
2053	P-1	Source-only			0.2205	0.2051	Oracle 0.2652
2054	P-1 P-2	0.0000 0.1022	0.2043	0.1900			
	P-2 P-3	0.1022	0.2846 0.1774	0.2812 0.1724	0.1933		0.3879 0.7345
2055	P-4	0.0484	0.1774	0.1724			0.7343
2056	P-5	0.2180	0.3007	0.3230		0.2002	
2057	P-6	0.0000	0.0442	0.0482		0.2392	
2058	P-7	0.0000	0.0328	0.1237		0.2392	
2059	P-8	0.0000	0.1328	0.2090		0.2316	
2060	P-9	0.0233	0.2393	0.4552			0.2899
2061	P-10	0.0000	0.4352	0.4332		0.2766	
2062	P-11	0.0585	0.3300	0.3279		0.3233	
2063	P-12	0.2857	0.3702	0.3601		0.3914	
2064	P-13	0.0000	0.1379	0.3175			0.3810
2065	P-14	0.0439	0.1197	0.2282			0.3992
2066	P-15	0.0000	0.1284	0.0917		0.3062	
	P-16	0.3214	0.3913	0.3361	0.2570	0.1410	0.3248
2067	P-17	0.0282	0.1518	0.1631	0.1600	0.1813	0.1777
2068	P-18	0.2028	0.1741	0.2000	0.0979	0.2052	0.3380
2069	P-19	0.0233	0.2033	0.2166	0.0253	0.2371	0.1807
2070	P-20	0.0000	0.3432	0.3575	0.3077	0.3489	0.4583
2071	P-21	0.2899	0.2615	0.2647	0.1439	0.1622	0.3121
2072	P-22	0.0000	0.2138	0.1667	0.2400	0.2991	0.1584
2073	P-23	0.1308	0.0655	0.0936		0.2529	0.2451
2074	P-24	0.1271	0.0758	0.2228	0.2704	0.3678	0.0599
2075	P-25						
2076	P-26	0.2770	0.2953	0.2873			0.0875
2077	P-27	0.3828	0.3921	0.3846			0.4702
2078	P-28	0.0321	0.0449	0.0867		0.1374	
2079	P-29	0.3710	0.3840	0.3810		0.4972	
2080	P-30 P-31	0.1111 0.2989	0.1325 0.4393	0.1769 0.4237		0.1352 0.1669	
2081	P-32	0.2989	0.4393	0.4237		0.1009	0.7216
2082	P-33	0.1957	0.4447	0.4300	0.5390		0.6606
2083	P-34	0.3936	0.5081	0.4800		0.2192	
	P-35	0.0783	0.2446	0.2531		0.2447	
2084	P-36	0.2793	0.4359	0.4200		0.1927	
2085	P-37	0.0000	0.0721	0.4713			0.7386
2086	P-38	0.2924	0.2835	0.3609	0.2659	0.1074	0.4847
2087	P-39	0.0467	0.4989	0.5000		0.3469	
2088	P-40	0.7431	0.7442	0.7421	0.4773	0.2772	0.8160
2089	P-41	0.0737	0.1388	0.1338	0.2060	0.1079	0.3354
2090	P-42	0.0216	0.1589	0.1572	0.1513	0.2407	0.0682
2091	P-43	0.0000	0.0103	0.0000		0.0919	
2092	P-44	0.2128	0.1569	0.1584		0.1277	
2093	P-45	0.1318	0.1298	0.1568		0.1156	
2094	P-46	0.2214	0.2018	0.2003		0.2179	
2095	P-47	0.2076	0.3479	0.3543		0.3706	
2096	P-48	0.0976	0.0803	0.0787			0.0887
2097	P-49	0.2597	0.1975	0.2621		0.2420	
2098	P-50	0.2054	0.2241	0.2250		0.2540	
2099	P-51	0.0000	0.2222	0.2222		0.1639	
	P-52	0.3418	0.3490	0.3486		0.3459	
2100	P-53 P-54	0.0800 0.1538	0.1536 0.1594	0.1475 0.1393		0.1102 0.1591	
2101	P-54 P-55		0.1394	0.1393	0.2113	0.1391	0.2030
2102	P-56	0.0000	0.2458	0.2720		0.2150	0.5606
2103	1-50	0.0000	0.2430	0.2720	0.5203	0.2130	0.5000

Table 24: Participant-wise F1 across different domain adaptation settings: source-only, UDA (MMD and Sub-Based), SFUDA (SHOT and NRC), and the oracle setting. Participants lacking A/H samples are highlighted. 'P' represents Participant.

2106	Tar Sub	Source-only	MMD	Sub-Rased	SHOT	NRC	Oracle
2107	P-1	0.8280	0.7279	0.7800	0.6749		0.8134
2108	P-2	0.7442	0.7180	0.6157		0.5913	0.7356
2109	P-3	0.1949	0.2830	0.2803			0.6767
2110	P-4	0.7188	0.6639	0.6600		0.6608	0.6540
2111	P-5	0.7759	0.7171	0.7322	0.5749		0.6478
	P-6	0.8342	0.8141	0.7900	0.5686	0.4830	0.8369
2112	P-7	0.4371	0.492	0.4700	0.4356	0.4834	0.6236
2113	P-8	0.8180	0.8337	0.8284	0.8181	0.7796	0.7845
2114	P-9	0.7076	0.6676	0.6700	0.5745	0.6670	0.8263
2115	P-10	0.7407	0.7434	0.7900	0.7208		0.766
2116	P-11	0.7144	0.7335	0.7316	0.4910	0.4897	0.777
2117	P-12	0.6937	0.5285	0.6470	0.5985	0.6461	0.7225
2118	P-13	0.7130	0.7219	0.7117	0.6650	0.6466	0.6487
2119	P-14	0.8032	0.7654	0.7703		0.8005	0.8202
2120	P-15	0.7373	0.7603	0.7502			0.7650
2121	P-16	0.9015	0.8831	0.8500			0.8465
2122	P-17	0.8491	0.8074	0.8200		0.7181	
2123	P-18	0.3837	0.4321	0.5800		0.7567	0.7001
2124	P-19	0.8844	0.7014	0.7200		0.6562	0.6817
2125	P-20	0.6789	0.6833	0.6800	0.5066		0.6461
	P-21	0.7602	0.7836	0.800		0.8537	
2126	P-22	0.8185	0.7526	0.800			0.8019
2127	P-23	0.7086	0.7616	0.7666		0.5677	
2128	P-24 P-25	0.6515	0.6913	0.6000	0.3074	0.4817	0.7363
2129	P-26	0.4858	0.4026	0.3106			0.7526
2130	P-27	0.3634	0.4020	0.4900		0.7520	0.7320
2131	P-28	0.8390	0.8101	0.4900		0.4927	0.9007
2132	P-29	0.5440	0.5242	0.5229	0.5818		0.6712
2133	P-30	0.8270	0.7610	0.7500		0.6173	0.8295
2134	P-31	0.7470	0.7049	0.6700		0.7959	
2135	P-32	0.4150	0.4659	0.4442		0.7279	
2136	P-33	0.4110	0.5165	0.5000		0.5514	
2137	P-34	0.4965	0.4446	0.4700	0.5243	0.5697	0.6742
2138	P-35	0.4589	0.5069	0.5103	0.5231	0.5038	0.6540
2139	P-36	0.6450	0.6791	0.6647	0.6131	0.7197	0.6832
2140	P-37	0.5713	0.5970	0.7185			0.8258
2141	P-38	0.4730	0.4597	0.4031			0.6380
2142	P-39	0.3382	0.5021	0.5000	0.4945		0.6589
	P-40	0.6443	0.6248	0.6072			0.7575
2143	P-41	0.8301	0.8103	0.8102		0.7577	
2144	P-42	0.7862	0.7613	0.7554		0.5631	0.8244
2145	P-43	0.9217	0.8829	0.8900		0.7857	0.9090
2146	P-44	0.7312	0.4289	0.4163		0.7937	0.5793
2147	P-45	0.7202	0.7202	0.7035			0.8174
2148	P-46 P-47	0.6241 0.7338	0.4856 0.4946	0.5000 0.3500		0.5005 0.6158	0.6369
2149	P-48	0.7536	0.4940	0.6073		0.0138	
2150	P-49	0.7556	0.0383	0.6900			0.7002
2151	P-50	0.7336	0.3983	0.5233			0.7749
2152	P-51	0.9629	0.4973	0.9304		0.9045	0.7749
2153	P-52	0.6466	0.9304	0.6500	0.6664		0.7480
2154	P-53	0.8622	0.6120	0.4527			0.8701
2155	P-54	0.8501	0.8054	0.8172	0.7319		0.7785
2156	P-55	-	-	-	-	-	-
	P-56	0.6753	0.6960	0.7000		0.7228	0.8077
2157							

Table 25: Participant-wise WF1 across different domain adaptation settings: source-only, UDA (MMD and Sub-Based), SFUDA (SHOT and NRC), and the oracle setting. Participants lacking A/H samples are highlighted. 'P' represents Participant.

2160					~~		
2160		Source-only					Oracle
2161	P-1	0.1210	0.1356	0.1321	0.1255		0.1818
2162	P-2	0.2425	0.2364	0.2010		0.3595	
2163	P-3	0.7141	0.7156	0.7200		0.3659	
2164	P-4	0.2263	0.2305	0.2300		0.1945	
2165	P-5	0.1121	0.1133	0.1135		0.1971	
2166	P-6	0.1012	0.1024	0.1072		0.1273	
2167	P-7	0.4696	0.4582	0.4713		0.3585	
2168	P-8	0.2103	0.2422	0.2508		0.3144	
2169	P-9	0.3859	0.3543	0.3600	0.3298		0.7228
	P-10	0.1707	0.2818	0.3098		0.2062	
2170	P-11	0.3277	0.3201	0.3146		0.2249	
2171	P-12	0.2124	0.2350	0.2400		0.3081	
2172	P-13	0.2060	0.2609	0.2554		0.2935	
2173	P-14	0.1226	0.1273	0.1932		0.1465	0.377
2174	P-15	0.2540	0.2606	0.2524		0.2214	
2175	P-16	0.3330	0.3419	0.3500		0.2987	
2176	P-17	0.1128	0.1118	0.1151		0.1944	
2177	P-18	0.1153	0.1141	0.1283		0.3041	
2178	P-19	0.1439	0.1405	0.1425		0.2903	0.132
2179	P-20	0.1777	0.2608	0.2647		0.2564	
2180	P-21	0.2363	0.2626	0.2613		0.2532	
	P-22 P-23	0.1163	0.1299	0.1337 0.1218		0.3965 0.1733	
2181	P-23 P-24	0.1120 0.1694	0.1137 0.1687	0.1218		0.1733	
2182	P-25	-	-	-	0.1033	-	-
2183	P-26	0.1927	0.1880	0.2130		0.2078	
2184	P-27	0.1327	0.3151	0.3349		0.3720	
2185	P-28	0.0493	0.0489	0.0547		0.2682	
2186	P-29	0.2553	0.2706	0.2700		0.4657	
2187	P-30	0.1572	0.1314	0.1419		0.2847	
2188	P-31	0.2982	0.2836	0.3007		0.2911	
2189	P-32	0.4454	0.4259	0.4569		0.2091	
2190	P-33	0.5795	0.5950	0.5873		0.4025	
2191	P-34	0.4147	0.4320	0.4407	0.4396	0.2138	0.6924
2192	P-35	0.4562	0.4385	0.4375	0.4166	0.2239	0.6593
2193	P-36	0.4630	0.5522	0.5478	0.4110	0.1744	0.7183
2194	P-37	0.6750	0.6536	0.6236	0.3856	0.2773	0.8139
	P-38	0.2802	0.2628	0.2722	0.258	0.1143	0.4661
2195	P-39	0.4912	0.4753	0.4728	0.5258	0.4601	0.7091
2196	P-40	0.6798	0.6724	0.6649	0.5816	0.3985	0.8876
2197	P-41	0.0993	0.1235	0.1240		0.2101	
2198	P-42	0.0959	0.1205	0.1182		0.1595	
2199	P-43	0.0432	0.0447	0.0448		0.0486	
2200	P-44	0.1382	0.1542	0.1495		0.0852	
2201	P-45	0.0801	0.0760	0.0838		0.0688	
2202	P-46	0.1459	0.1473	0.1457		0.2547	
2203	P-47	0.2451	0.2686	0.2639		0.3361	
2204	P-48	0.1192	0.0927	0.1124		0.2058	
2205	P-49	0.1400	0.1202	0.1602		0.1485	
2206	P-50	0.1293	0.1440	0.1432		0.1895	
	P-51	0.0714	0.5359	0.5339			0.6549
2207	P-52	0.2945	0.3046	0.3044		0.2602	
2208	P-53	0.1588	0.1070	0.2165		0.3122	
2209	P-54	0.1234	0.1173	0.1179	0.1//1	0.2725	0.2968
2210	P-55	0.2340	0.2655	0.2654	0.2252	0.3140	0.5042
2211	P-56	0.2340	0.2655	0.2654	0.2333	0.3140	0.3943

Table 26: Participant-wise AP across different domain adaptation settings: source-only, UDA (MMD and Sub-Based), SFUDA (SHOT and NRC), and the oracle setting. Participants lacking A/H samples are highlighted. 'P' represents Participant.