Multi-modal Active Learning From Human Data: A Deep Reinforcement Learning Approach

Ognjen (Oggi) Rudovic¹, Meiru Zhang², Björn Schuller² and Rosalind W. Picard¹

¹Massachusetts Institute of Technology, USA, ²Imperial College London, UK

{orudovic,roz}@mit.edu, {meiru.zhang18, bjoern.schuller}@imperial.ac.uk

ABSTRACT

Human behavior expression and experience are inherently multimodal, and characterized by vast individual and contextual heterogeneity. To achieve meaningful human-computer and human-robot interactions, multi-modal models of the user's states (e.g., engagement) are therefore needed. Most of the existing works that try to build classifiers for the user's states assume that the data to train the models are fully labeled. Nevertheless, data labeling is costly and tedious, and also prone to subjective interpretations by the human coders. This is even more pronounced when the data are multi-modal (e.g., some users are more expressive with their facial expressions, some with their voice). Thus, building models that can accurately estimate the user's states during an interaction is challenging. To tackle this, we propose a novel multi-modal active learning (AL) approach that uses the notion of deep reinforcement learning (RL) to find an optimal policy for active selection of the user's data, needed to train the target (modality-specific) models. We investigate different strategies for multi-modal data fusion, and show that the proposed model-level fusion coupled with RL outperforms the feature-level and modality-specific models, and the naïve AL strategies such as random sampling, and the standard heuristics such as uncertainty sampling. We show the benefits of this approach on the task of engagement estimation from real-world child-robot interactions during an autism therapy. Importantly, we show that the proposed multi-modal AL approach can be used to efficiently personalize the engagement classifiers to the target user using a small amount of actively selected user's data.

ACM Reference Format:

Ognjen (Oggi) Rudovic¹, Meiru Zhang², Björn Schuller² and Rosalind W. Picard¹. 2019. Multi-modal Active Learning From Human Data: A Deep Reinforcement Learning Approach. In *2019 International Conference on Multimodal Interaction (ICMI '19)*, October 14–18, 2019, Suzhou, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3340555.3353742>

1 INTRODUCTION

Human behavior is inherently multi-modal, and individuals use eye gaze, hand gestures, facial expressions, body posture, and tone of voice along with speech to convey engagement and regulate social interactions [12]. A large body of work in human-computer interaction (HCI) and human-robot interaction (HRI) explored the use of various affective and social cues to facilitate and assess the user engagement [16, 36]. Most of these works can be divided into those that detect the presence of a set of the engagement cues or interaction events [9, 17, 29], or use supervised classifiers trained with social, physiological, or task-based interaction features [6, 8, 33]. A detailed overview of related multi-modal approaches can be found in [2, 28]. The majority of these approaches adopt either feature-level (e.g., by simple concatenation of the multi-modal input features) or model-level (e.g., by combining multiple classifiers trained for each data modality) fusion [4]. While this can improve the estimation of target outcomes (e.g., the user engagement) when compared to the single-modality classifiers, these methods usually adopt "one-size-fits-all" learning approach where the trained models are applied to new users without the adaptation to the user. Consequently, their performance is usually limited when the users' data are highly heterogeneous (e.g., due to the differences in facial expressions/body gestures as a result of individual engagement styles).

Recently, several works proposed models for personalized estimation of engagement in HRI. For instance, [31] proposed a multi-modal deep learning for engagement estimation that combines body, face, audio and autonomic physiology data of children with autism during the therapy sessions with a humanoid robot. However, most of multi-modal works in HCI and HRI are fully supervised [4, 20, 23], i.e., they assume that the data used to train the models are fully labeled. Obtaining these labels is expensive, especially when dealing with a large amount of audio-visual recordings. Therefore,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICMI '19, October 14–18, 2019, Suzhou, China

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6860-5/19/10...\$15.00

<https://doi.org/10.1145/3340555.3353742>

Figure 1: Overview of the proposed multi-modal AL approach. The input are the recordings of the child-robot interactions during an autism therapy (we used the camera placed behind the robot). The image frames are first processed using open-source tools (openFace [5] and openPose [7]) to obtain the facial and body cues. Likewise, the audio recordings are processed using the openSMILE toolkit [13]. We also used the data collected by the E4 wristband [1] on the child’s hand (providing autonomic physiology data such as galvanic skin conductance and body temperature, as well as the accelerometer data). These data are fed as input to the modality-specific engagement classifiers: the LSTM models followed by fully-connected layers (fcL). The classification of the engagement levels (low, medium, high) is performed by applying majority voting to the classifiers’ outputs. These are also fed into the Q-function of the RL policy for active data-selection (also modeled using an LSTM cell and fcL for the action selection: to ask or not ask for the label). If the label is requested, the input multi-modal data is stored in a data pool for labelling by the human expert. These data are then used to train the data-selection policy and target classifiers. During inference of the data of a new child, the actively-selected data are used to personalize the engagement classifiers.

there is a need for methods that can automatically select the most informative instances that need to be labeled in order to train the target estimation models. More importantly, to improve the generalization of these models, we need to select the data of a new user that can be used to personalize the target models. Yet, how to select the most informative instances of the multi-modal data from the target user is an open research problem that has not been investigated much.

To address this, the approach proposed here uses the notion of AL [34]. Central to the AL framework is the query strategy used to decide when to request a label for target data. The most commonly used query strategies include uncertainty sampling, entropy, or query-by-committee [34]. Furthermore, more advanced query strategies have been proposed to adapt deep network classifiers based on the uncertainty of the network output (e.g., [22, 24, 38]). Yet, the candidate query strategies still must be specified by a human. More importantly, there is not one strategy that works the best for all users. Instead of using the heuristic strategies, recent (deep) AL approaches (e.g., [11, 15, 25, 37, 39]) have adopted a data-driven approach that learns a model-free AL off-line policy using RL [35]. For instance, [39] proposed a model where an agent makes a decision whether to request a label or make a prediction. The agent receives a reward related to its decision: a positive reward is given for correct predictions, and negative rewards for incorrect predictions or label requests. This can be achieved by the Q-function modeled using the notion of deep RL [15, 26]. The main goal of these approaches is to adapt the prediction model to new tasks, using a minimum number of queries. Our work is a

generalization of the RL framework for AL [15] to multi-modal data, where instead of dealing with a single agent-environment, the agent deals simultaneously with multiple environments (i.e., data modalities).

Note that RL has previously been applied in the tasks of multi-modal learning. For instance, [30] used RL to enhance the machine translation from different data modalities. Likewise, [21] used RL for image question answering, an inherently multi-modal learning problem. Also, in the context of visual dialogues, RL has been used with multi-modal learning [40]. However, none of these works explored RL for AL from multimodal data. The most related approach to ours is the multi-view AL framework [27]. It uses the standard heuristic AL strategies to select data from multiple views. While different views can be seen as different modalities of the same phenomenon, like facial and body gestures of human behaviour, to our knowledge no previous work has attempted multi-modal AL from the real-world human-interaction data and using RL to learn an optimal data selection policy. Moreover, the model personalization using such approach has not been explored before.

To tackle the challenges of learning from multi-modal human data (as typically encountered in HCI and HRI), in this work we formulate a novel multi-modal AL approach. We show that this approach can be used to personalize the data-modality-specific classifiers to the target user using a small amount of labeled data of the user, which are automatically selected using the newly proposed multi-modal AL strategy. The main contributions of this work are: (i) We propose a novel approach for multi-modal AL using RL for training a policy for active data-selection. (ii) We propose a

novel personalization strategy based on the actively-selected multi-modal data of the target user. (iii) We show on a highly challenging dataset of child-robot interactions during an autism therapy that the proposed approach leads to large improvements in estimation of engagement (low, medium, high) from the multi-modal data, when compared to non-personalized models, and heuristic AL strategies. The outline of the proposed approach is depicted in Fig. 1. Compared to traditional supervised classifiers for multi-modal data, our approach provides an efficient mechanism for actively selecting the most relevant data for training the target engagement classifier, thus, minimizing the human data-labelling efforts.

2 PRELIMINARIES

Problem Statement and Notation

In our learning setting, we use multi-modal data recordings of child-robot interactions during an autism therapy [31], as described in Sec. 4. Formally, we denote our dataset as $\mathcal{D} = \{D^{(1)}, \dots, D^{(m)}, \dots, D^{(M)}\}$, where $m = 1, \dots, M$ denotes the data modality (e.g., face, body, etc.). This dataset comprises video recordings of target interactions of C children (later split into training and test child-independent partitions). We assume a single recording per child, which may vary in duration. Each recording is represented with a set of multi-modal features $\{X, Y\}$, where $X = [x_1, \dots, x_N]$, with $x_i \in \mathcal{R}^{T \times M \times D_M}$ containing the collection of the multi-modal features extracted every 60 ms from a sliding window of 1 second duration (30 image frames), resulting in $T = 10$ temporally correlated multi-modal feature representations (the dimension of modality $m = 1, \dots, M$ is D_m). The features x_i of each 1 second interval are associated with the target engagement label $y_i = \{0, 1, 2\}$ (see Sec. 4 for details). Note that N can vary per child, i.e., per recording. Given these data, we address it as a multi-class multi-modal sequence classification problem, where our goal is two-fold: (i) to predict the target label given the input features extracted from the sliding window within the recording, and (ii) to actively select the data of each child so that our engagement estimation model can be personalized to the target child.

The Base Classification Model

As the base model for the engagement classifiers (φ) and also to implement the Q-function in the RL component of our approach, we use the Long Short-Term Memory (LSTM) [19] model, which enables long-range learning of time-feature dependencies. This has shown great success in tasks such as action recognition [3, 10] and speech analysis [14, 18], among others. Each LSTM cell has hidden states augmented with nonlinear mechanisms that allow the network state to propagate without modification, be updated, or be reset, using simple learned gating functions. More formally, a basic

LSTM cell can be described with the following equations:

$$\begin{aligned} \hat{g}^f, \hat{g}^i, \hat{g}^x, \hat{c}_t &= W_x \cdot x_t + W_h \cdot h_{t-1} + b \\ g^f &= \sigma(\hat{g}^f), g^i = \sigma(\hat{g}^i), g^x = \sigma(\hat{g}^x) \\ c_t &= g^f \odot c_{t-1} + g^i \odot \tanh(\hat{c}_t), \quad h_t = g^x \odot \tanh(c_t), \end{aligned} \quad (1)$$

where $\hat{g}^f, \hat{g}^i, \hat{g}^x$ are the forget gates, input gates, and output gates respectively, $\hat{c}_t$ is the candidate cell state, and c_t is a new LSTM cell state. W_x and W_h are the weights mapping from the observation (x_t) and hidden state (h_{t-1}), respectively, to the gates and candidate cell state, and b is the bias vector. $\odot$ represents element-wise multiplication; $\sigma(\cdot)$ and $\tanh(\cdot)$ are the sigmoid and hyperbolic tangent functions respectively [39]. To model the window of $T = 10$ time steps in (x, y) data pairs, we feed the feature vectors from $\mathbf{x} \in [x_1, \dots, x_t, \dots, x_T]$ to the temporally unrolled LSTM cell. Then, their output-state values h_t are passed through fully connected layers fcL_t , $t = 1, \dots, T$ (we use a rectified linear unit – ReLU), and average across the time steps. Finally, a sigmoid layer followed by the softmax is applied to this output to obtain the sequence label y^* .

RL for Data-selection Policy Learning

RL [35] is a framework that can be used to learn an optimal policy π for actively selecting data samples during model training and adaptation. This is a form of Markov Decision Process (MDP), which allows the learning of a policy that can dynamically select instances that are most informative [15]. More specifically, given an input feature vector ($\mathbf{x}_i$), we first compute a state-vector (s_i) that is then passed to the trained policy (Q-function), which outputs an action (a_i). During training of the Q-function, the goal is to maximize the action-value function $Q^*(s_i, a_i)$. This is at the heart of RL, and it specifies the expected sum of discounted future rewards for taking action a_i in state s_i and acting optimally from then on as:

$$a_i = \pi^*(s_i) = \arg \max_{a_i} Q^*(s_i, a_i); \quad (2)$$

The optimal Q function is given by the Bellman equation:

$$Q^*(s_i, a_i) = \mathbb{E}_{s_{i+1}} [r_i + \gamma \max_{a_{i+1}} Q^*(s_{i+1}, a_{i+1}) | s_i, a_i], \quad (3)$$

where $\mathbb{E}_{s_{i+1}}$ indicates an expected value over the distribution of possible next states s_{i+1} , r_i is the reward at the current state s_i (derived from the input features $\mathbf{x}_i$), and γ is a discount factor, which incentivizes the model to seek reward in fewer time steps. The tuples (a_i, s_i, r_i, s_{i+1}) represent a MDP, and they are used to train the Q-function in Eq. (3).

3 METHODOLOGY

We propose a novel approach for multi-modal AL that provides an optimal policy for the active data-selection. These

data are used consequently to re-train the classification models for estimation of the target output (in our case, the engagement level) from fixed-sized video segments (1 second long). The proposed approach consists of two sequential processes: (i) the training of the classifiers for the target output, and (ii) the learning of the Q-function of the RL model for active data selection. These two are performed in a loop, where first the target classifiers are trained using the data of each modality ($m = 1, \dots, M$), thus, M classifiers are trained in parallel. Then, based on their outputs, we perform the model-level fusion to obtain the target label for the input. This label is used, along with the input features and the true label, to train the Q-function for active data selection. During the models' training, this process is repeated by re-training the classifiers and the Q-function until both models have converged (or for a fixed number of iterations, i.e., training episodes). During inference of new data (in our case, the recording of the interactions between the new child and the robot), the learned group policy (i.e., the Q-function) is first used to select the data samples that the model is uncertain about. Consequently, these are then used to personalize the target classifiers by additionally training each classifier using the actively selected data of the target child. The personalized classifiers are then used to obtain the engagement estimates on the remaining data of the target child.

Engagement Classifiers

For each data modality ($m = 1, \dots, M$) we train a separate LSTM model, resulting in the following ensemble of the models: $\varphi = \{\varphi^{(1)}, \varphi^{(2)}, \dots, \varphi^{(M)}\}$. Each model is trained using the actively-selected samples from the target modality, where the number of possible samples is defined using a pre-defined ($\mathcal{B}$) for active learning. The number of hidden states in each LSTM was set to 64 (found on a separate validation set), and the size of the fcL was then set to 64×3 , where the network output was 1-hot encoded (three engagement levels). Specifically, in the dataset used, we have $M = 4$ data modalities, comprising of the FACE ($m = 1$), BODY ($m = 2$), autonomic physiology (A-PHYS) ($m = 3$) and AUDIO ($m = 4$) modality. Thus, for the target task, we trained four LSTM models. From each data modality, we used the feature representations extracted using the open-source codes for face and body processing from image data (openFace [5] and openPose [7]), as done in [31]. These provide the locations of characteristic facial points, facial action units activations (0/1) and their intensities (0-5), and locations of 18 body joints and their confidence. As features of A-PHYS, we used the data collected with the E4 [1] wristband on a child's hand, providing the galvanic skin response, heart-rate and body temperature, along with the 3D accelerometer data encoding the child's movements. From the audio recordings, we used 24 low-level descriptors provided by the openSMILE [13], an

open-source toolkit for audio-feature extraction. The feature dimension per modality was: 257D (FACE), 70D (BODY), 27D (A-PHYS), and 24D (Audio), thus, 378 features in total. For more details about the feature extraction, see [31].

To obtain the engagement estimate by the proposed multi-modal approach, we perform the model-level fusion by combining the target predictions of the modality-specific LSTMs using the majority vote as:

$$\hat{y}_i \leftarrow \text{majorityVote}(\varphi^{(1)}(\mathbf{x}_i^{(1)}), \varphi^{(2)}(\mathbf{x}_i^{(2)}), \varphi^{(3)}(\mathbf{x}_i^{(3)}), \varphi^{(4)}(\mathbf{x}_i^{(4)})), \quad (4)$$

where in the case of ties, we select the most confident estimate, based on the soft-max outputs of each classifier. We also tried other model-fusion schemes, such as confidence-based weighting, and majority vote on the three most confident estimates, but the basic majority vote performed the best, and was used in the experiments reported.

Group-policy Learning for Active Data-selection

We first use the training data (i.e., recordings of the children in the training set) to learn the (initial) group-policy π_g for making the decision whether to query or not the label for the input features ($\mathbf{x}$). For this, we implement the Q-function using the LSTM model, which receives in its input the states (s) and outputs actions (a), as described in Sec. 2.

States and Actions. We approximate the Q-function using the LSTM model with 32 hidden units. This is followed by a ReLU layer with the weight and bias parameters $\{W_l, b_l\}$, and softmax in the output, i.e., actions. The actions in our model are binary (ask/ do not ask for label), and are 1-hot encoded using a 2D output vector a_i for input states s_i . For instance, if a label is requested, $a_i = [1, 0]$; otherwise, $a_i = [0, 1]$ and no data label is provided by an oracle (e.g. the human expert). For training the Q-function, the design of the input states is critical. Typically, the raw input features (in our case, $\mathbf{x}_i$, are considered to be the states (s_i) of the model [39]. While this is feasible, in the case of multiple modalities, this can lead to a large state-space, which in turn can easily lead to the over-fitting of the LSTM used to implement the Q-function. Here, we propose a different approach. Instead of using the input content ($\mathbf{x}_i$) directly, we use the output of the engagement classifiers to form the states of the Q-function. Namely, the sigmoid layer of each target classifier outputs the probabilities for each class label in y , i.e., $\varphi^{(m)} \rightarrow (p_1^{(m)}, p_2^{(m)}, p_3^{(m)})$. To obtain the "content-free"¹ states for our RL model, for the target input $\mathbf{x}_i$, we concatenate these estimated probabilities from each classifier to form the state vector s_i . Furthermore, we augment this state vector by also adding the overall confidence of each classifier. This is computed as

¹Note that this can also be termed as "meta-weights", as in the previously proposed meta-learning AL frameworks, e.g., [24].

Algorithm 1: Multi-modal Q-learning (MMQL)

Input: Dataset $D = \{D^{(1)}, D^{(2)}, \dots, D^{(M)}\}$, models $\{\varphi_0 = \{\varphi_0^{(1)}, \varphi_0^{(2)}, \dots, \varphi_0^{(M)}\}, Q_0^\pi\} \leftarrow \text{rand}, \mathcal{B}$

Output: Optimized models $\{\varphi_*, Q_*^\pi\}$

```

for  $e := 1$  to  $N$  do
   $D_0^\ell \leftarrow \emptyset$ , shuffle  $D$ ;
   $\varphi_e \leftarrow \varphi_{e-1}$ ,  $Q_e^\pi \leftarrow Q_{e-1}^\pi$ ;
  for  $i := 1$  to  $|D|$  do
     $\hat{y}_i \leftarrow \text{majorityVote}(\varphi_e(\mathbf{x}_i))$ ;
     $s_i \leftarrow \mathbf{x}_i$  # construct a new state;
     $a_i = \arg \max_a Q_e^\pi(s_i, a)$  # make a decision;
    if  $a_i == 1$  then
       $D_e^\ell \leftarrow D_e^\ell \cup (\mathbf{x}_i, y_i)$  # ask for label;
    end
     $r_i \leftarrow R(a_i, \hat{y}_i, y_i)$  # compute reward;
    if  $|D_e^\ell| == \mathcal{B}$  then
      store  $(s_i, a_i, r_i, \text{end}) \rightarrow \mathcal{M}$ ;
      break;
    end
     $s_{i+1} \leftarrow \mathbf{x}_{i+1}$  # construct a new state;
    store  $(s_i, a_i, r_i, s_{i+1}) \rightarrow \mathcal{M}$ ;
    update  $Q_e^\pi$  using a batch from  $\mathcal{M}$ ;
  end
  update models  $\varphi_e$  using  $D_e^\ell$ ;
end
Return:  $\varphi_* \leftarrow \varphi_e$ ,  $Q_*^\pi \leftarrow Q_e^\pi$ 

```

$C^{(m)} = 1 - \sum_{i=1}^{|y|} p_i^{(m)} \log(p_i^{(m)})$, where the sum on the right-hand side of the equation is the entropy of the classifier (we bound it to $[0, 1]$). Finally, in the case of $M = 4$, this results in 16D state-vectors s_i . As we show in our experiments, such the state representation leads to overall better results and learning of the Q-function than when the raw high-dimensional input features are used to represent the states (in our case, $\dim(\mathbf{x}) = 378$).

Reward. Another important aspect of RL is the design of the reward function. Given the input multi-modal feature vectors $\mathbf{x}_i$, the active learner chooses an action a_i of either requesting the true label y_i or not. If the label is requested, the model receives a negative reward to reflect that obtaining labels is costly. On the other hand, if no label is requested, the model receives positive reward if the estimation is correct; otherwise, it receives negative reward. This is encoded by

Algorithm 2: Personalized Engagement Estimation

Input: New data $D = \{D^{(1)}, D^{(2)}, \dots, D^{(K)}\}$, models $\{\varphi_* = \{\varphi_*^{(1)}, \varphi_*^{(2)}, \dots, \varphi_*^{(M)}\}, Q_*^\pi\}, \mathcal{B}$

Output: labeled data in D , adapted models φ_*

```

 $D^\ell \leftarrow \emptyset$ , shuffle  $D$ 
for  $i := 1$  to  $|D|$  do
   $\hat{y}_i \leftarrow \text{majorityVote}(\varphi_*(\mathbf{x}_i))$ ;
   $s_i \leftarrow \mathbf{x}_i$  # construct a new state;
   $a_i = \arg \max_a Q_*^\pi(s_i, a)$  # make a decision;
  if  $a_i == 1$  then
     $D^\ell \leftarrow D^\ell \cup (\mathbf{x}_i, y_i)$  # ask for label;
     $D \leftarrow D \setminus \mathbf{x}_i$ ;
  end
  if  $|D^\ell| == \mathcal{B}$  then
    break;
  end
end
update models  $\varphi_*$  using  $D^\ell$ ;
make estimates  $\hat{Y} \leftarrow \text{majorityVote}(\varphi_*(\mathbf{X} \in D))$ ;
Return:  $\hat{Y} \in D$ ,  $Y \in D^\ell$ ,  $\varphi_*$ 

```

the following RL reward function:

$$r_i(a_i, \hat{y}_i, y_i) = \begin{cases} r_{req} = -0.05, & \text{if } a_i = 1 \\ r_{cor} = +1, & \text{if } a_i = 0 \wedge \hat{y}_i = y_i \\ r_{inc} = -1, & \text{if } a_i = 0 \wedge \hat{y}_i \neq y_i \end{cases}, \quad (5)$$

where y_i is the target label obtained by the majority vote from the modality-specific engagement classifiers (Sec. 3). This reward is critical for training the Q-function that we use to learn the data-selection policy. Note that this type of reward has previously been proposed in [39], however, it has not been used in the context of multi-modal AL.

Optimization. Given the space-action pairs, and the rewards, the parameters of the Q-function (modeled with a single fcL) are optimized by minimizing the Bellman loss on the training data:

$$L_B^{(i)}(\Theta) = [Q_\Theta(s_i, a_i) - (r_i + \gamma \max_{a_{i+1}} Q_\Theta(s_{i+1}, a_{i+1}))]^2, \quad (6)$$

which encourages the model to improve its estimate of the expected reward at each training iteration (we set $\gamma = 0.9$). This is performed sequentially over $i = 1, \dots, N$ data samples (from the experience replay buffer), and over 100 training episodes. Note that for training the taggers, we use the standard softmax loss. The overall loss minimization is performed alternatively between the Q-learning and the classification loss, using Adam optimizer with the learning rate 0.001. The

training of this new AL approach, named the Multi-modal Q-learning (MMQL), is summarized in Alg. 1.

Personalized Estimation of Engagement

The learned group-policy for active data selection can be applied to multi-modal recordings of previously unseen children. However, the trained multi-modal engagement classifiers may be suboptimal due to the highly diverse styles of engagement expressions across the children with autism, in terms of their facial expressions, head movements, body gestures and positions, among others. These may vary from child to child not only in their appearance but also dynamics during the engagement episodes. To account for these individual differences, we adapt the proposed MMQL to each child. We do so by additionally training (fine-tuning) the modality-specific engagement classifiers using the actively selected data samples from the target child. Specifically, we start with the group-level engagement classifiers to obtain the initial engagement estimates, but also use the learned Q-function to select difficult data that need be expert-labeled. This is performed in an off-line manner: the requested samples of the target child are first annotated by an expert, and then used to personalize the engagement classifier. The main premise here is that with a small number of human-labeled videos, the engagement classifiers can easily be optimized for the target child as:

$$\varphi_*^{new} = \max \langle \varphi_*^{(1)}, \varphi_*^{(2)}, \dots, \varphi_*^{(M)} \rangle_{D^\ell}, \quad (7)$$

where D^ℓ are the child data-samples actively selected using the group-level Q-function, and under the $\mathcal{B}$. This is described in detail in Alg. 2.

4 EXPERIMENTS

Data and Features. To evaluate the proposed approach, we used the cross-cultural dataset of children with ASC attending a single session (on average, 25 mins long) of a robot-assisted autism therapy [32]. During the therapy, an experienced educational therapist worked on teaching the children socio-emotional skills, focusing on recognition and imitation of behavioral expressions as shown by neurotypical population. To this end, the NAO robot was used to demonstrate examples of these expressions. The data comprises highly synchronized audio-visual and autonomic physiological recordings of 17/18 children, ages 3-13, with Japanese/European cultural background, respectively. All the children have a prior medical diagnosis of ASC, varying in its severity. The audio-visual recordings were annotated by the human experts on a continuous scale [-1,1]. We discretized these annotations by binning the average continuous engagement score within 1 sec intervals into: low [-1,0.5], medium (0.5,0.8], and high (0.8,1] engagement. The multi-modal features were obtained as described in Sec. 3. We split the children into training (20)

and test (14)² at random, and used their recordings for the evaluation of the models presented here. Table 1 shows the highly imbalanced nature of the data of the test children.

Performance Metrics and Models. We report the average accuracy (ACC) and F-1 score of the models in the task of the 3-level engagement classification. The reported results are obtained by evaluating the group-level models and models adapted to each test child (Alg. 2). The latter was repeated 10 times by random shuffling of the test child data, and the average results are reported. For training/adaptation of the models, we varied the for active data selection as $\mathcal{B} \in \{5, 10, 20, 50, 100\}$. The number of episodes during training was set to 100, and LSTMs were trained for 10 epochs in each episode (and during the classifier adaptation to the test child). To evaluate different model settings, we start with the uni-modal models (thus, trained/tested using only $M = 1$ data modality – FACE, BODY, A-PHYS or AUDIO) and the multi-modal approach with the feature level fusion (i.e., by concatenating the input features from modalities $M = 1, \dots, 4$). We compare these models to the proposed model-level fusion in the MMQL approach. As the baselines, we show the performance achieved with alternative data-selection strategies: random sampling (RND) and the most common AL heuristic for data selection – uncertainty sampling (UNC). In the case of multi-modal learning, the UNC scores for each sample were computed as the sum of the classifiers’ entropy (Sec.3).

Results. Tables 2&3 show the summary of the performances achieved by the compared models. We first note that the classifiers trained on the training children have generalized poorly on the test children. This is expected especially when using data of children with autism, who exhibit very different engagement patterns during their interactions. As can be seen from the numbers on the right-hand side (obtained after the personalization of the classifiers using actively selected data of the target child), the models’ performances largely increase. This evidences the importance of the model personalization using the data of the target child. Overall, among the uni-modal models, the BODY modality achieves the best performance, followed by the FACE, A-PHYS, and AUDIO modality, as also noted in [31]. However, both multi-modal versions of the models (feature- and model-level fusion) bring gains in the performance, with the model-level fusion performing the best on average.

Comparing the MMQL models with the states based on the data content (cont=1) and those constructed from the classifiers outputs (cont=0), we note that there is no large difference in the performance for most models. Yet, the F-1 score of the MMQL with model-level fusion achieves a larger improvement (4.3%). On the other hand, we note that the MMQL approach with cont=0 achieves a similar performance

²The data of one child were discarded because of severe face occlusions.

Table 1: The distribution of the engagement levels for each test child. The data samples are obtained by applying a 1-second sliding window, with a shift of 20 ms, to the original recordings of the target children from Japan (C1) and Serbia (C2).

Test child ID	1(C2)	2(C1)	3(C1)	4(C1)	5(C2)	6(C2)	7(C1)	8(C2)	9(C1)	10(C1)	11(C1)	12(C1)	13(C1)	14(C2)
# of class 0	697	0	72	31	541	2281	3903	478	125	0	188	20	0	1134
# of class 1	683	0	107	72	241	1026	1004	772	318	0	101	204	107	44
# of class 2	0	2539	2095	1452	0	0	898	87	549	1689	1524	5686	2107	0

Table 2: ACC [%] of the models before (*left*) and after (*right*) adaptation to test children using different $\mathcal{B}$ (5, 10, 20, 50 and 100) for active data-selection (averaged across the children and $\mathcal{B}$).

Model	FACE	BODY	A-PHYS	AUDIO	FEATURE-F	MODEL-F
MMQL (cont=0)	49±10 / 78± 13	61± 21 / 81± 13	36± 10 / 77± 14	49± 21 / 76± 17	61± 20/ 81± 14	57± 14 / 82± 13
MMQL (cont=1)	51± 14 / 77± 15	60± 22 / 81± 14	31± 12 / 79± 15	49± 22 / 74± 18	58± 23 / 80± 14	55± 21 / 82± 14
UNC	48± 14 / 69± 14	63± 22 / 74± 16	19± 12 / 53± 12	50± 26 / 70± 21	59± 23 / 72± 18	46± 12 / 73± 15
RND	48± 13 / 70± 15	63± 24 / 73± 17	22± 12 / 57± 14	45± 12 / 71± 21	59± 18 / 74± 16	50± 12 / 75± 16

Table 3: F-1 [%] of the models before (*left*) and after (*right*) adaptation to test children using different $\mathcal{B}$ (5, 10, 20, 50 and 100) for active data-selection (averaged across the 14 children and $\mathcal{B}=5$).

Model	FACE	BODY	A-PHYS	AUDIO	FEATURE-F	MODEL-F
MMQL (cont=0)	29± 7 / 42± 6	34± 8 / 52± 8	23± 6 / 44± 5	24± 7 / 36± 3	33± 7 / 49± 5	32± 6 / 53± 7
MMQL (cont=1)	31± 8 / 43± 7	33± 9 / 52± 9	20± 7 / 46± 6	24± 8 / 35± 3	33± 9 / 49± 5	28± 7 / 53± 11
UNC	30± 9 / 36± 7	36± 11 / 42± 10	14± 9 / 30± 5	24± 10 / 32± 3	33± 9 / 38± 6	28± 7 / 38± 6
RND	30± 9 / 37± 7	35± 12 / 42± 10	16± 9 / 31± 6	24± 6 / 33± 4	33± 8 / 39± 8	30± 7 / 39± 6

Figure 2: The performance of the feature- and model-level fusion with different active data-selection strategies.

Figure 3: The performance per test child of the MMQL (model-level fusion, cont=0) before (in green) and after (in blue) the personalization of the engagement classifiers. The results are shown for $\mathcal{B}=10$. The star in front of the child ID indicates that the performance difference is statistically significant ($p = 0.05$)

to when the content is used. Thus, this simpler model is preferable in practice. In the rest of the experiments, we show the performance of the MMQL (cont=0) only. Compared to the baselines, we note that the proposed MMQL largely outperforms these base strategies for active data-selection, under the same constraints. This evidences that the proposed is able to learn a more efficient data-selection policy. By comparing the ACC and F-1 scores, we note that the proposed is able to improve the classification of each engagement level, while the RND/UNC strategies tend to overfit the majority class.

Similar observations can be made from Fig. 2, showing the performance of personalized models after the adaptation using different $\mathcal{B}$. MMQL consistently outperforms RND and UNC sampling strategies, which we attribute again to the superior performance of the data selection strategy attained by the RL Q-function. This trend is even more pronounced for larger $\mathcal{B}$, evidencing that the proposed Q-function consistently selects the more informative data samples, that are used to adapt the modality-specific classifiers to the target child. This holds for both, the feature- and model-level fusion within the proposed MMQL approach.

Fig. 3 shows the performance of the MMQL (model-fusion, cont=0) approach per child before and after the model personalization using data actively selected with the proposed RL approach. Note that even with 5 samples only, the engagement classification improves largely and for almost all test children. However, because of the highly imbalanced nature of these data (see Table 3), F-1 scores are relatively low, yet, the improvements due to the active adaptation of the target classifiers are evident. One of the challenges when working with such imbalanced data is that the classifiers tend to overfit the majority class as most of the active samples come from that class. While this is much more pronounced in the

RND/UNC selection strategies, it is also one of the bottlenecks of the current approach since the target classifiers are updated offline. We plan to address this in our future work.

5 CONCLUSIONS

We proposed a novel active learning approach for multi-modal data of human behaviour. Instead of using heuristic strategies for active data-selection (such as the uncertainty sampling), our approach uses the notion of deep RL for active selection of the most informative data samples for the model adaptation to the target user. We showed the effectiveness of this approach on a highly challenging multi-modal dataset of child-robot interactions during an autism therapy. Specifically, we showed that the learned data-selection policy can generalize well to new children by being able to select their data samples that allow the pre-trained engagement classifiers to adapt quickly to the target child. We showed that this multi-modal model personalization can largely improve the performance of the engagement estimation for each test child using only a few expert-labeled data of the target child.

ACKNOWLEDGEMENTS

This work was funded by EC H2020 Marie Curie Action – Individual Fellowship no. 701236 (EngageMe), EC H2020 no. 688835 (DE-ENIGMA), and ACLEW (Analyzing Child Language Experiences Around the World), project no. HJ-253479.

REFERENCES

- [1] 2015. *Empatica E4*: <https://www.empatica.com/en-eu/research/e4/>.
- [2] Xavier Alameda-Pineda, Elisa Ricci, and Nicu Sebe. 2019. Multimodal behavior analysis in the wild: An introduction. In *Multimodal Behavior Analysis in the Wild*. Elsevier, 1–8.
- [3] Moez Baccouche, Franck Mamalet, Christian Wolf, Christophe Garcia, and Atilla Baskurt. 2011. Sequential deep learning for human action recognition. In *International Workshop on Human Behavior Understanding*. Springer, 29–39.

- [4] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2019. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 2 (2019), 423–443.
- [5] Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. 2016. Openface: an open source facial behavior analysis toolkit. In *IEEE Conference on Applications of Computer Vision*. 1–10.
- [6] Dan Bohus and Eric Horvitz. 2014. Managing human-robot engagement with forecasts and... um... hesitations. In *Proceedings of the 16th international conference on multimodal interaction*. ACM, 2–9.
- [7] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2017. Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [8] Ginevra Castellano, André Pereira, Iolanda Leite, Ana Paiva, and Peter W McOwan. 2009. Detecting user engagement with a robot companion using task and social interaction-based features. In *Proceedings of the 2009 international conference on Multimodal interfaces*. ACM, 119–126.
- [9] Cheng Chang, Cheng Zhang, Lei Chen, and Yang Liu. 2018. An Ensemble Model Using Face and Body Tracking for Engagement Detection. In *Proceedings of the International Conference on Multimodal Interaction*. ACM, 616–622.
- [10] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. 2015. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2625–2634.
- [11] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. 2016. RL2: Fast Reinforcement Learning via Slow Reinforcement Learning. *arXiv preprint arXiv:1611.02779* (2016).
- [12] Rana El Kaliouby, Rosalind Picard, and Simon Baron-Cohen. 2006. Affective computing and autism. *Annals of the New York Academy of Sciences* 1093, 1 (2006), 228–248.
- [13] Florian Eyben, Felix Weninger, Florian Gross, and Björn Schuller. 2013. Recent developments in opensmile, the munich open-source multimedia feature extractor. In *ACM International Conference on Multimedia*. 835–838.
- [14] Florian Eyben, Felix Weninger, Stefano Squartini, and Björn Schuller. 2013. Real-life voice activity detection with lstm recurrent neural networks and an application to hollywood movies. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 483–487.
- [15] Meng Fang, Yuan Li, and Trevor Cohn. 2017. Learning how to active learn: A deep reinforcement learning approach. *arXiv preprint arXiv:1708.02383* (2017).
- [16] Michael A Goodrich, Alan C Schultz, et al. 2008. Human–robot interaction: a survey. *Foundations and Trends® in Human–Computer Interaction* 1, 3 (2008), 203–275.
- [17] Goren Gordon, Samuel Spaulding, Jacqueline Kory Westlund, Jin Joo Lee, Luke Plummer, Marayna Martinez, Madhurima Das, and Cynthia Breazeal. 2016. Affective Personalization of a Social Robot Tutor for Children’s Second Language Skills. In *Thirtieth AAAI Conference on Artificial Intelligence*. <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11759>
- [18] Alex Graves and Navdeep Jaitly. 2014. Towards end-to-end speech recognition with recurrent neural networks. In *International Conference on Machine Learning*. 1764–1772.
- [19] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [20] Alejandro Jaimes and Nicu Sebe. 2007. Multimodal human–computer interaction: A survey. *Computer vision and image understanding* 108, 1-2 (2007), 116–134.
- [21] Ai-Wen Jiang, Bo Liu, and Ming-Wen Wang. 2017. Deep Multimodal Reinforcement Network with Contextually Guided Recurrent Attention for Image Question Answering. *Journal of Computer Science and Technology* 32, 4 (2017), 738–748.
- [22] Christoph Käding, Erik Rodner, Alexander Freytag, and Joachim Denzler. 2016. Active and Continuous Exploration with Deep Neural Networks and Expected Model Output Changes. *arXiv preprint arXiv:1612.06129* (2016).
- [23] Samira Ebrahimi Kahou, Xavier Bouthillier, Pascal Lamblin, Caglar Gulcehre, Vincent Michalski, Kishore Konda, Sébastien Jean, Pierre Froumenty, Yann Dauphin, Nicolas Boulanger-Lewandowski, et al. 2016. Emonets: Multimodal deep learning approaches for emotion recognition in video. *Journal on Multimodal User Interfaces* 10, 2 (2016), 99–111.
- [24] Ksenia Konyushkova, Raphael Sznitman, and Pascal Fua. 2017. Learning active learning from data. In *Advances in Neural Information Processing Systems*. 4228–4238.
- [25] Ming Liu, Wray Buntine, and Gholamreza Haffari. 2018. Learning to Actively Learn Neural Machine Translation. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*. 334–344.
- [26] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [27] Ion Muslea, Steven Minton, and Craig A Knoblock. 2006. Active learning with multiple views. *Journal of Artificial Intelligence Research* 27 (2006), 203–233.
- [28] Maja Pantic and Leon JM Rothkrantz. 2003. Toward an affect-sensitive multimodal human-computer interaction. *Proc. IEEE* 91, 9 (2003), 1370–1390.
- [29] Hae Won Park, Ishaan Grover, Samuel Spaulding, Louis Gomez, and Cynthia Breazeal. 2019. A Model-free Affective Reinforcement Learning Approach to Personalization of an Autonomous Social Robot Companion for Early Literacy Education. In *Proceedings of the Thirty Third AAAI Conference on Artificial Intelligence (AAAI’19)*. AAAI Press.
- [30] Xin Qian, Ziyi Zhong, and Jieli Zhou. 2018. Multimodal machine translation with reinforcement learning. *arXiv preprint arXiv:1805.02356* (2018).
- [31] Ognjen Rudovic, Jaeryoung Lee, Miles Dai, Björn Schuller, and Rosalind Picard. 2018. Personalized Machine Learning for Robot Perception of Affect and Engagement in Autism Therapy. *Science Robotics* (2018).
- [32] Ognjen Rudovic, Jaeryoung Lee, Lea Mascarell-Maricic, Björn W Schuller, and Rosalind W Picard. 2017. Measuring Engagement in Robot-Assisted Autism Therapy: A Cross-Cultural Study. *Frontiers in Robotics and AI* 4 (2017), 36.
- [33] Jyotirmay Sanghvi, Ginevra Castellano, Iolanda Leite, André Pereira, Peter W McOwan, and Ana Paiva. 2011. Automatic analysis of affective postures and body motion to detect engagement with a game companion. In *Proceedings of the 6th international conference on Human-robot interaction*. ACM, 305–312.
- [34] Burr Settles. 2010. Active learning literature survey. *University of Wisconsin, Madison* 52, 55-66 (2010), 11.
- [35] Richard S Sutton and Andrew G Barto. 1998. Reinforcement learning: An introduction. *MIT press Cambridge* (1998).
- [36] Christiana Tsiourti, Astrid Weiss, Katarzyna Wac, and Markus Vincze. 2019. Multimodal Integration of Emotional Signals from Voice, Body, and Context: Effects of (In) Congruence on Emotion Recognition and Attitudes Towards Robots. *International Journal of Social Robotics* (2019), 1–19.
- [37] Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharmashan Kumar, and Matt Botvinick. 2016. Learning to reinforcement learn. *arXiv preprint*

- arXiv:1611.05763* (2016).
- [38] Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. 2016. Cost-effective active learning for deep image classification. *IEEE Transactions on Circuits and Systems for Video Technology* (2016).
 - [39] Mark Woodward and Chelsea Finn. 2016. Active one-shot learning. *NIPS, Deep Reinforcement Learning Workshop* (2016).
 - [40] Jiaping Zhang, Tiancheng Zhao, and Zhou Yu. 2018. Multimodal hierarchical reinforcement learning policy for task-oriented visual dialog. *arXiv preprint arXiv:1805.03257* (2018).