Dyadic Speech-based Affect Recognition using DAMI-P2C Parent-child Multimodal Interaction Dataset

Huili Chen
MIT Media Lab
hchen25@media.mit.edu

Yue Zhang
MIT Media Lab
yuefw@mit.edu

Felix Weninger
Nuance Communications
felix@weninger.de

Rosalind Picard
MIT Media Lab
picard@media.mit.edu

Cynthia Breazeal
MIT Media Lab
cynthiab@media.mit.edu

Hae Won Park
MIT Media Lab
haewon@media.mit.edu

ABSTRACT

Automatic speech-based affect recognition of individuals in dyadic conversation is a challenging task, in part because of its heavy reliance on manual pre-processing. Traditional approaches frequently require hand-crafted speech features and segmentation of speaker turns. In this work, we design end-to-end deep learning methods to recognize each person's affective expression in an audio stream with two speakers, automatically discovering features and time regions relevant to the target speaker's affect. We integrate a local attention mechanism into the end-to-end architecture and compare the performance of three attention implementations – one mean pooling and two weighted pooling methods. Our results show that the proposed weighted-pooling attention solutions are able to learn to focus on the regions containing target speaker's affective information and successfully extract the individual's valence and arousal intensity. Here we introduce and use a “dyadic affect in multimodal interaction - parent to child” (DAMI-P2C) dataset collected in a study of 34 families, where a parent and a child (3-7 years old) engage in reading storybooks together. In contrast to existing public datasets for affect recognition, each instance for both speakers in the DAMI-P2C dataset is annotated for the perceived affect by three labelers. To encourage more research on the challenging task of multi-speaker affect sensing, we make the annotated DAMI-P2C dataset publicly available¹, including acoustic features of the dyads' raw audios, affect annotations, and a diverse set of developmental, social, and demographic profiles of each dyad.

CCS CONCEPTS

• **Affective computing and interaction;** • **Speech behaviours in social interaction;** • **Multimodal datasets and validation;**

KEYWORDS

dyadic affect dataset; speech affect recognition; parent-child interaction; end-to-end learning

¹The database is available on <https://forms.gle/FXJAPRUGkLwbW8NN9>.

This work is licensed under a Creative Commons Attribution International 4.0 License.

ICMI '20, October 25–29, 2020, Virtual event, Netherlands

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7581-8/20/10.

<https://doi.org/10.1145/3382507.3418842>

ACM Reference Format:

Huili Chen, Yue Zhang, Felix Weninger, Rosalind Picard, Cynthia Breazeal, and Hae Won Park. 2020. Dyadic Speech-based Affect Recognition using DAMI-P2C Parent-child Multimodal Interaction Dataset. In *Proceedings of the 2020 International Conference on Multimodal Interaction (ICMI '20)*, October 25–29, 2020, Virtual event, Netherlands. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3382507.3418842>

1 INTRODUCTION

1.1 Learning technologies for parent-child dyadic interactions

High-quality, social, responsive, and facilitative interactions between parents and their children are crucial for young children's language development [35, 48]. A major problem faced by many children, particularly those from low socioeconomic status (SES) families, is limited exposure to rich language during adult-child conversations at home. For instance, studies have reported that in low-SES families, parent-to-child conversations tend to be less frequent and of shorter duration [17, 36], and have fewer open-ended questions and discussions (e.g., parents negotiating with their children) [17, 18, 23, 36] compared higher-SES families. Such parental participation gap² is gaining significant attention, as the active participation of parents (or other adults children interact with daily) is crucial for children's educational success [29, 31]. It should be noted that the interest level of parents in their children's education does not vary across different SES groups [30]; however, the lack of access to parental education and guidance contributes to the participation gap. As noted by Hoover et al. [19], a well guided parental involvement in children's education has a significant impact on children's cognitive development and literacy skills.

Therefore, there is urgency in figuring out how to enrich social and conversational interactions between parents and children, particularly those from lower-SES households. Despite the increasing number of educational apps designed to support parental participation [26, 47], few artificial intelligence (AI) technologies have been designed to support multimodal and reciprocal adult-child interaction, for example, to proactively facilitate dialogic storytelling between two interlocutors in the here and now – where parent and child not only read but actively converse about the story, asking and answering questions, commenting on the narrative, and so on [3].

To achieve this objective, AI-enabled learning technologies (e.g., robot learning companions [4, 43]) need to have the ability to accurately recognize affective expressions of individuals in the

parent-child dyads. Children’s affective states have been shown to be crucial for their learning during their interaction with learning technologies [5], and affect-aware learning technologies promote children’s learning more effectively than affect-blind ones [14, 32]. Such affect sensing ability offers opportunities for learning technologies to provide real-time personalized interactions and interventions that optimize the learning experience for each family.

1.2 State-of-the-art affect sensing models

Current speech affect sensing models focus primarily on single-person affect detection and assume that an audio recording containing multiple speakers’ utterances is manually partitioned into homogeneous segments according to the speaker’s identity. Thus, speaker diarization has been an inevitable preprocessing step for these models, and as a result, modeling affect of individuals in dyads has been treated as a two-step problem. As far as we know, there is currently no approach that offers speaker diarization and affect recognition in a single model.

The traditional feature designing approach for speech-based affect recognition also creates additional barriers and requires much manual effort. After a large number of acoustic descriptors, such as pitch and energy, are extracted on a frame level from raw speech signals, feature designing techniques, e.g., statistical functionals, are applied to obtain a subset of most salient features for classification or regression with machine learning algorithms [11]. The feature attributes for model training are frequently selected through extensive and carefully prepared experiments and their effectiveness heavily relies on the implemented pattern recognition model [11]. Typically, a researcher invests a lot of manual effort in these steps, with a common belief being that the more effort, the greater the generalizability of the results.

As an increasingly popular alternative approach that can reduce manual effort, end-to-end deep learning automatically crafts a feature set and explores the most salient representations related to the task of interest by jointly training the representation learning and pattern recognition processes [24, 28, 41, 49, 50]. For example, Li et al. [24] have applied convolutional neural networks (CNNs) combined with attention pooling directly to the spectrograms extracted from speech utterances to jointly model temporal and frequency domain information and learn a final emotional representation. Instead of using CNNs, Mirsamadi et. al [28] have designed a deep recurrent neural network (RNN) system that learns both emotionally relevant features and their temporal aggregation from either raw audio signals or low-level descriptors of speech.

To our knowledge, the current end-to-end approaches are, however, designed for single-person affect recognition, and have not been applied to multi-speaker affect recognition. In a multi-speaker scenario, it is important to learn which time regions hold important information for affect prediction of a specific speaker. For instance, in an audio where a parent is talking with a child, an affect recognition model for the child’s arousal level needs to intelligently identify time regions when the child is speaking and the segments within those time regions most indicative of her arousal level. This means that an end-to-end model needs to have an attention mechanism that can learn the varying importance of the representations at different times as well as focus on the target speaker instead of

simply averaging them across an entire audio that may contain utterances from both speakers. .

1.3 Our approach and contributions

In this work, we integrate an attention mechanism into an end-to-end framework as an implicit speaker diarizer with the aim to automatically focus on the target speaker in dyads when modeling that speaker’s affect.

Since its inception, the attention mechanism has been widely used in various speech-related tasks including speech recognition and natural language processing. Additionally, more recent state-of-the-art work on emotion recognition started to implement attention mechanisms into end-to-end learning models in various ways and have increased the predictive power of the models [20, 27, 28, 51]. For example, Mirsamadi [28] has augmented deep neural networks (DNNs) with a temporal attention mechanism after RNNs to extract the most emotionally-salient parts of speech signals for emotion prediction. Another prior work added attention layers to CNNs for emotion recognition [24].

Here, we apply weighted pooling with local attention to study a novel problem in the speech affect field – dyadic affect recognition. To the best of our knowledge, this is the first work that is applying attention mechanisms in DNNs to model dyadic affect predictions. In fact, there has been very little research effort in recognizing the affect of both participants in a dyadic interaction. One of the main reasons for the lack of research in this field is that there is no existing dyad dataset with proper affect annotations. Even the most popular dyadic datasets such as IEMOCAP [2] does not provide the labels for both speakers’ emotions. To help advance research in this area, we present a new dataset, “dyadic affect in multimodal interaction – parent to child” (DAMI-P2C), which captures the affect (arousal and valence) of both parent and child in natural story-reading dyads. This is the first dataset that provides annotations for both speakers’ affective expressions for each instance.

In summary, the main contributions of this work are: (1) A rich dataset of parent-child dyads with affect annotations and other social and developmental profiles of the dyads, (2) The design and development of an end-to-end learning framework augmented with an attention mechanism for multi-speaker affect recognition, and (3) A comprehensive analysis with experimental results to provide a competitive affect recognition baseline for future multi-speaker affect recognition research.

2 DYAD INTERACTION DATA COLLECTION

2.1 Collection protocol

We invited families with children between the ages of 3-7 years to come to our lab space for a two-session data collection. The collection protocol consisted of two 45-minute in-lab sessions where a parent and their child read stories together for around 20 minutes, and the parent filled out surveys for the remaining 25 minutes. During the in-lab sessions, the parent and child sat next to each other as shown in Figure 1. Families that completed both sessions were given \$75 as their compensation.

A digitized version of our storybook corpus on a touchscreen tablet was used for the sessions. The storybook corpus consists of 30 storybooks recommended by early childhood education experts

Table 1: Gender identity, age range and language proficiency of the participant families in the dyadic dataset. English language learner is denoted as ELL. Native/bilingual speaker is denoted as N/B. The average ages for parents and children are 39.70 ± 5.47 and 5.49 ± 1.37 , respectively.

	parent (age)		child (age)	
Female	25	(38.78 $\pm$ 4.70)	13	(5.20 $\pm$ 1.96)
Male	9	(42.25 $\pm$ 6.90)	21	(5.65 $\pm$ 0.96)
N/B	28	(39.54 $\pm$ 5.15)	29	(5.58 $\pm$ 1.46)
ELL	6	(40.33 $\pm$ 7.15)	5	(5.00 $\pm$ 0.64)

Figure 1: An example scene of a parent-child storytime interaction in the lab setting. Annotators viewed audio-visual recordings similar to this one, except showing the faces, when coding parent’s and child’s affect attributes.

and teachers. Each story lasts from 3 to 15 minutes. Stories shorter than five minutes were categorized as short stories and the rest as long stories. During the story reading session, the parent and child could select any books they wanted to read from the corpus. They were encouraged to read stories together in a way that they would normally do at home, having conversations around the stories and making the activity fun and interactive.

Thirty-four families were recruited in our data collection from the Greater Boston area with their full consent (Table. 1). The dyad was always a pair of one parent and one child. Two parents did not report their children’s age, and four parents did not report their own. Three families withdrew from the data collection after the first session for reasons not related to the protocol. In sum, DAMI-P2C dataset consists of two sessions of 31 families and one session of three families.

2.2 Measurements

During the in-lab parent-child co-reading sessions, audio-visual recordings were captured using one microphone and four cameras installed in the story reading station (Figure 1). Four cameras were used to capture different angles of the dyadic interaction (i.e., frontal view, birds-eye view, parent-centered view, and child-centered view). The audio recordings were sent to a professional transcription service to obtain textual annotation of recorded speech.

In addition, we collected participant families’ demographic information, socio-economic status, home literacy environment, parenting style [22], a parental theory of mind assessment [34], and a child’s temperament and behavior questionnaire [33] completed by the parent.² The three families who only completed the first in-lab session did not fill out the surveys.

²The DAMI-P2C dataset is being released in phases, with the first phase having the acoustic features of raw audios, affect annotations, and the developmental, social and

2.3 Data annotation

We recruited three trained annotators with a psychology or education background to annotate the audio-visual recordings of the families’ co-reading interactions. The affect of both parent and child were annotated separately in terms of valence and arousal. While watching the recordings, the coders gave ratings every five non-overlapping seconds on a five-point ordinal scale [-2,2], with two corresponding to cases when the target person showed clear signs of high arousal/positive valence and -2 when the person showed clear signs of low arousal/negative valence.³ A five second window was selected as the fragment interval of target audio-visual recordings for the annotation to produce continuous quality scales, a threshold consistent with prior work on affect detection [38]. When annotating the recordings, annotators were instructed to judge whether a given fragment contained the story-related dyadic interaction and filter out those that did not. For those five-second fragments in which speakers paused and only made utterances in a small time interval, i.e., 1 second, we kept them as valid data instances, since a short period of silence frequently happens naturally in human conversations. In total, 16,593 five-second fragments have been annotated with 488.03 ± 123.25 fragments from each family on average, shown in Table 2.

The agreement of the three annotators was measured using the intra-class correlation (ICC) [42] type (3,1) for average fixed raters. The ICC is commonly used in behavioral sciences to assess the annotators’ agreement and score ranges from 0-100%. The average ICC and its confidence interval among the annotators for each of our labels are presented in Table 3. According to the commonly-cited cutoffs for qualitative inter-rater reliability (IRR) based on ICC values [7, 16], IRR is good if $.60 \leq ICC < .75$, and is excellent if $.75 \leq ICC \leq 1.0$. Given this evaluation criteria, the annotations for the four affect attributes all achieve either good or excellent qualities. After recordings were coded independently by the annotators, we took the average of a scale’s ratings from the three annotators for each recording fragment as its final score. The ratings were standardized to the zero mean and unit standard deviation for each annotator in order to eliminate individual judgement biases of the three annotators in the use of the rating scale. Both the average standardized ratings and the three independent raw ratings are available in the DAMI-P2C dataset.

The label distribution for each target attribute after the individual ratings are averaged and standardized is depicted in Figure 2, and the pairwise Spearman correlation coefficients between the four affect attributes are reported in Table 3. The correlation results showed that speaker valence and arousal are moderately positively correlated with each other for both parent and child. Additionally, parent’s and child’s arousals are negatively correlated with each other. The latter may be explained by the conversational interactions in which the parent and the child took turns being a speaker and listener.

demographic profiles of each dyad. We plan to add in the raw audio-visual recordings with faces unobstructed for facial affect analysis in later phases.

³Coding manuals for valence/arousal of both parent and child can be accessed via the link (<https://drive.google.com/file/d/1wSedbtZcik8f21AvHU4tKSMOPFU1b5z/view?usp=sharing>).

Table 2: Statistics of the number of valence/arousal annotations collected across families. The first, second and third quartiles are denoted as Q1, Q2 and Q3, respectively

Mean $\pm$ SD	Min	Max	Q1	Q2	Q3
488.03 $\pm$ 123.23	235	724	438.25	482.00	576.75

Table 3: Left: Pairwise Pearson correlation coefficients between child’s arousal (CA), parent’s arousal (PA), child’s valence (CV), and parent’s valence (PV). Right: The average and confidence interval (CI) of intra-class correlation for the four affect attributes.

	CA	PA	CV	PV	Label	ICC score	ICC CI
CA	1				CA	0.84	[0.84, 0.85]
PA	-0.27	1			PA	0.83	[0.82, 0.83]
CV	0.32	-0.01	1		CV	0.64	[0.63, 0.65]
PV	-0.19	0.43	0.14	1	PV	0.61	[0.60, 0.62]

Figure 2: Histograms of the distribution for each affect attribute label after three individual ratings are averaged and standardized.

3 METHODOLOGY

3.1 End-to-end speech emotion recognition

In deep learning, end-to-end architecture refers to the process of learning the direct mapping between the input and final output using a neural network model as a black box and training it according to the global objective function. State-of-the-art end-to-end models can be broadly categorized into 1) spectrogram (e.g., [15]) and 2) raw wave form (e.g., [40]) based on the input signal type. Different combinations of neural network models, e.g., convolutional and recurrent, have been explored as end-to-end emotion recognition systems [53].

3.2 Low-level descriptors

Low-level descriptors (LLDs) are a set of acoustic features computed from short frames of typically 20 to 50 milliseconds. The most relevant feature groups are prosodic, spectral, cepstral, and voice quality features⁴. The traditional approach to feature extraction is to apply statistical functionals to the LLDs over the duration of the utterance, and then concatenate the aggregated values into a long

feature vector at the utterance level to form high-level statistical functions. Instead, we compute 65 LLDs from each frame of the audio samples using the openSMILE 2.0 toolkit [11], and extract temporal information by feeding the frame-level LLDs into the end-to-end learning framework, which systematically crafts the feature set and may save much manual effort.

3.3 Model design

Our end-to-end architecture for parent-child affect recognition consists of convolutional, recurrent, and deep neural networks (CRDNN), as illustrated in Figure 3. Frame-level LLDs of the raw wave form audios are used as the input signals. As the first module in the architecture, the CNN is used to extract high-level emotional features from LLDs, treating LLDs as two-dimensional images with time and frequency dimensions and reducing the feature engineering. Then, the RNN is used to extract the contextual and time-series information of the CNN output, analyzing the sequence feature of the audio signals. With the spectral and temporal modeling handled by the CNN and RNN layers, an attention layer is used to select regions of speech signals vital for the task of interest. Lastly, the output from the attention layer is fed into a fully-connected (FC) layer, which maps the attention output, i.e., utterance representation, to another space easier for the task of continuous affect recognition.

Input. For each five-second audio instance, a set of LLD features is extracted at the frame level, and each LLD is normalized by its global mean and standard deviation calculated across the entire dataset. Then, the input is a sequence of frames $\mathbf{x} = \{x_1, x_2, \dots, x_L\}$ where x_t is a set of LLDs at frame t and L is the total number of frames in the five-second audio.

CNN. The CNN structure consists of a single one-dimensional convolutional layer with filters of 1×8 , and takes $\mathbf{x}$ as the input. Then, it is followed by an overlapping max-pooling layer with pooling size=(1,3) on the frequency axis, and an element-wise rectified linear unit (ReLU). Then, the batch normalization layer [21] is applied to the CNN output to speed up training. The final output is denoted as $\mathbf{f} = \{f_1, f_2, \dots, f_L\}$.

RNN. The sequence of higher-level representations from CNNs ($\mathbf{f}$) is then directed into two bidirectional gated recurrent unit (GRU) layers each having 128 hidden cells to learn a sequential pattern of the input signal data. GRU is chosen for its efficacy of temporal summarization and less demand on training data size in comparison with long short-term memory (LSTM) networks [6], as well as its effectiveness in modeling individual affect shown in prior work [52]. The output sequence of the last recurrent layer is represented as $\mathbf{h} = \{h_1, h_2, \dots, h_L\}$.

Attention layer. For the models without any attention mechanism, the RNN output is directed into a temporal mean-pooling layer to transform into one single vector $\mathbf{z}$ for the FC module. For the models with an attention mechanism, an attention layer is applied to transform the RNN output into $\mathbf{z}$. More details on the attention mechanism are in Section 3.4.

FC. As a regression module, the FC layer with input size of 128, hidden units of 512, and output size of 1 maps the pooling layer’s output, $\mathbf{z}$, to a final continuous affect score.

⁴The full list of LLDs computed in this work can be accessed via link (<https://drive.google.com/file/d/1Bq6rWLVUggQwyzaJ4c9QU1epOBalCHdr/view?usp=sharing>).

Figure 3: End-to-end convolutional recurrent deep neural networks (CRDNN) architecture design. (a): baseline CRDNN framework without a mean pooling layer (Section 3.4.1). (b): CRDNN framework with a weighted pooling using the RNN output (Section 3.4.2). (c): CRDNN framework with a weighted pooling using the CNN output (Section 3.4.3).

Dropout. To understand the impact of dropout on model performance, we implement two models with different dropout layer functions for each CRDNN architecture. The first model has one dropout on the CNN layer and one on the RNN. The second model has an additional dropout on the attention layer in addition to the first two dropouts. We do not implement models without CNN and RNN dropout, as prior work suggested applying dropout to the CNN and RNN as a common practice to increase the model’s generalizability [13, 44]. A 10% dropout is applied to each CNN, RNN, and FC layer during training to prevent overfitting and increase model generalizability.

3.4 Attention mechanism

We implement the following attention mechanisms. In each architecture, the attention mechanism is added as the layer directly after the second RNN layer and before the FC layer, and the obtained result from the attention mechanism z is fed into the FC layer for affect regression.

3.4.1 Mean-pooling over time without local attention. For the baseline model without local attention, we perform a mean-pooling over time on the RNN outputs (z), computed as follows:

$$z = \frac{\sum_{t=1}^L h_t}{|L|}$$

3.4.2 Weighted-pooling with local attention using RNN output. Instead of the mean pooling over time, we compute a weighted sum of the RNN outputs with the following equation:

$$z = \sum_{t=1}^L \alpha_t h_t$$

where weight α_i is the weight of the RNN output h_i . The equation to compute α_t is as follows:

$$\alpha_t = \frac{\exp(w^T h_t)}{\sum_{i=1}^L \exp(w^T h_i)}$$

At each time frame t , the inner product between the learnable parameter w and the RNN output h_t is interpreted as a score for the contribution of frame t to the final utterance-level representation of the affect. Then, the result is fed into a softmax function to obtain a set of final weights α_t which sum to unity.

3.4.3 Weighted-pooling with local attention using the CNN output. Instead of learning to weight the frames based on the signal’s temporal dynamic from the RNN output, we design a new weighted-pooling layer that allows the model to weight the frames based on speaker characteristics, which are mostly present in the input features or lower layers of the network. This weighted pooling with the CNN output is similar to the one with the RNN output except that the weight of the RNN output α_t is calculated using the CNN output (f_t) instead of the RNN output (h_t), as follows:

$$\alpha_t = \frac{\exp(w^T f_t)}{\sum_{i=1}^L \exp(w^T f_i)}$$

4 EXPERIMENTS

In this study, we focused on the audio modality in the collected audio-visual dataset described in Section 2. The dataset has four affect attribute categories: parent’s arousal and valence, child’s arousal and valence. All four categories were used to train the deep models.

4.1 Evaluation settings

To evaluate the different deep learning architectures, we used a leave one speaker group out cross validation (LOGO-CV). We partitioned 34 dyads into five non-overlapping groups. The evaluation setting was subject-independent; therefore, data instances (five-second audio fragments) from each dyad only appeared in one of the five groups. When partitioning the dataset, each group’s size was additionally balanced by sorting dyads based on their instance sizes and assigning the largest dyad in the unassigned dyad list to each group iteratively. For each iteration, we selected one as the development (dev) set, one as the test set and the rest as the train set. The experiment repeated until each group had been used as dev and test sets once, and the experiment was repeated five times in total. In total, each experiment was trained on a training set of 9955.8 ± 327 instances, a dev set of 3318.6 ± 255 instances and a test set of 3318.6 ± 255 instances. We plan to publish the split with the dataset for other researchers to use the same split for model comparisons. To evaluate model performance, the Spearman’s rank correlation coefficient (ρ) was selected for its ability to assess monotonic relationships (whether linear or not) without assuming normal distributions of variables, given that the four affect annotations did not have normal distributions. Note that ρ ranges from $[-1, 1]$; however, we report them in %, i.e., $[-100, 100]$.

4.2 CRDNN Model training

Our CRDNN models were implemented using TensorFlow, and different models were trained separately for each speaker and each affect attribute. The training of the proposed architectures was conducted using the Adam optimization algorithm with an initial learning rate of 0.0001 and a batch size of 20. For the loss function, we used the concordance correlation coefficient (ccc), which measures agreement as a departure from perfect linearity ($y = \hat{y}$). An early stop was implemented to avoid overfitting issues. Specifically, the training would stop if the model's performance on the dev dataset did not improve after 15 epochs.

Table 4: Model performance on predicting child's arousal

Methods	DROP(ALL) ρ (mean $\pm$ SD)	DROP(CR) ρ (mean $\pm$ SD)
ATT(NO)	52.1 $\pm$ 6.6%	53.0 $\pm$ 7.2%
ATT(R)	53.5 $\pm$ 6.8%	54.7 $\pm$ 5.9%** [†]
ATT(C)	54.0 $\pm$ 7.8%*	51.9 $\pm$ 9.0%

Table 5: Model performance on predicting parent's arousal.

Methods	DROP(ALL) ρ (mean $\pm$ SD)	DROP(CR) ρ (mean $\pm$ SD)
ATT(NO)	53.6 $\pm$ 9.8%	53.5 $\pm$ 9.9%
ATT(R)	53.7 $\pm$ 11.7%	53.8 $\pm$ 10.0%
ATT(C)	53.1 $\pm$ 9.9%	53.6 $\pm$ 9.6%

Table 6: Model performance on predicting child's valence.

Methods	DROP(ALL) ρ (mean $\pm$ SD)	DROP(CR) ρ (mean $\pm$ SD)
ATT(NO)	22.1 $\pm$ 6.4%*	21.4 $\pm$ 4.2%
ATT(R)	19.8 $\pm$ 5.2%	19.5 $\pm$ 5.8%
ATT(C)	20.4 $\pm$ 5.9%	21.7 $\pm$ 4.2%

Table 7: Model performance on predicting parent's valence.

Methods	DROP(ALL) ρ (mean $\pm$ SD)	DROP(CR) ρ (mean $\pm$ SD)
ATT(NO)	28.1 $\pm$ 16.1%*	31.6 $\pm$ 14.4%
ATT(R)	26.1 $\pm$ 17.3%	30.4 $\pm$ 14.3%
ATT(C)	26.7 $\pm$ 14.6%	33.7 $\pm$ 12.4%** [†]

5 RESULTS AND DISCUSSION

To compare the performance of different CRDNN architectures, we trained and evaluated 24 models in total with six models per learning task. Models are categorized into three attention types: (1) ATT(NO) denotes no attention layer models with a mean-pooling layer (Section 3.4.1), (2) ATT(R) denotes models with a weighted-pooling layer using the RNN output (Section 3.4.2), and (3) ATT(C) denotes models with a weighted-pooling layer using the CNN output (Section 3.4.3). For each attention implementations, we trained two models, which are the DROP(ALL) model that has dropouts on its CNN, RNN and attention layers, and the DROP(CR) model that has dropouts on its CNN and RNN layers but not on its attention layer. Results showing Spearman correlation coefficient means and

standard deviations obtained on the six models for child's arousal, parent's arousal, child's valence, and parent's valence are reported in Tables 4, 5, 6 and 7, respectively.

When comparing the performance between two models, we used a z-test to evaluate the statistical significance of the difference in their Spearman correlation coefficients (ρ). The sample size in the z-test was the total dataset size ($N = 16,593$), as every data instance was used as the test instance once in the cross-validation [8]. In each learning task, we evaluated the performance difference between the ATT(NO) model and the most effective attention-augmented model within the same dropout configuration, and the z-test results with $p < 0.05$ significance level marked with * and $p < 0.001$ as ** in the tables. Additionally, we evaluated the performance difference between the most effective ATT(NO) model and attention-augmented model across both DROP(ALL) and DROP(CR) configurations in each task, and the overall most effective models with the significant z-test results are marked with [†] in the tables.

Figure 4: The attention-augmented model, ATT(R)+DROP(CR), automatically learned attention weights for child's and parent's arousal predictions; results are shown here for two audio wave files from the test set. The time region for parent's and child's utterances are highlighted in blue and yellow, respectively. The top, middle and bottom subplot are the amplitude of the raw audio signals and attention weights for parent's arousal model, and weights for child's arousal model, respectively.

5.1 Model performance across learning tasks

5.1.1 Comparisons between arousal and valence. All six models achieved significantly higher ρ scores on child's arousal and parent's arousal than on child's valence and parent's valence, according to the z-test results ($p < 0.05$). The most effective models for child's arousal and parent's arousal achieved 54.7 $\pm$ 5.9% and 53.8 $\pm$ 10.0%, respectively; conversely, the most effective models for valence achieved only 22.1 $\pm$ 6.4% and 33.7 $\pm$ 12.4%. This performance difference between predicting a speaker's arousal and valence is consistent with findings from prior studies in affect recognition research [49, 51], which have repeatedly shown that valence is much more difficult to predict from speech compared to arousal.

5.1.2 Comparisons between parent's and child's valence. All six models achieved significantly higher ρ scores when predicting parent's valence than predicting child's valence according to the z-test results ($p < 0.05$). The ρ scores across the six models are 29.4 $\pm$ 3.0% and 20.8 $\pm$ 1.1% on average for parent's and child's valence, respectively. This finding suggests that it was easier to predict parent's valence than child's valence, and such difference could be explained

by how parents and children interacted when reading stories together. Based on our observation, parents in our study were more verbally expressive and made more utterances than children did, making it easier for models to pick up their valence from their verbal behaviors. Additionally, children in our study were diverse in their personalities and temperaments, and several of them tended to exhibit their valence predominantly through body gesture and facial expression without making utterances, particularly when they were listening to their parents reading stories. This would make it more difficult for a single-modality speech model to predict children's valence. The finding creates a higher motivation to integrate children's developmental profiles (e.g., temperament) as an additional modality into future dyadic affect model design.

5.2 Effect of attention mechanism on model performance

5.2.1 Overall performances of attention-augmented models. In child's arousal and parent's valence tasks, the most effective attention-augmented model achieved significantly higher ρ scores than the most effective attention-free model, as analyzed by the z-test (child's arousal: $p < 0.001$; parent's valence: $p < 0.001$). In the other two tasks, the most effective attention-augmented and attention-free models achieved comparable performances with their z-test results not being statistically significant (parent's arousal: $p > 0.05$; child's valence: $p > 0.05$). This result showed that integrating a weighted attention layer into a CRDNN model led to its improved performance on predicting child's arousal and parent's valence in parent-child interactions, indicating the overall effectiveness of using an attention mechanism in dyadic affect recognition despite that statistically significant performance gains may not be evident in all affect learning tasks.

5.2.2 Parent's and child's arousal. A comparison of model performance between parent's arousal and child's arousal tasks showed that integrating a weighted attention layer into a CRDNN model led to a larger performance improvement on child's arousal than parent's arousal as seen in Tables 4 and 5. For child's arousal, the most effective attention-augmented model ATT(R)+DROP(CR) achieved 54.7%, showing a statically significant increase of 1.7% over the most effective attention-free model ($\rho = 53.0\%$), according to the z-test ($p < 0.001$). In addition, the most effective attention-augmented models in the two dropout configurations, that is, ATT(C)+DROP(ALL) and ATT(R)+DROP(CR), significantly outperformed their attention-free counterparts in child's arousal task ($p < 0.05$), respectively. This finding further strengthens the claim that having a weighted-pooling attention layer would improve a CRDNN model's prediction on child's arousal. Conversely, in the parent's arousal task, the performance difference between the most effective attention-augmented and attention-free models was not statistically significant in each dropout configuration.

To further investigate how the weighted-pooling attention layer in the proposed framework was trained, we illustrate the learned attention weights of ATT(R)+DROP(CR) model in parent's and child's arousal tasks for two wave audio files from the test set along with the audio files' speech amplitudes, seen in Figure 4. Generally speaking, as shown in the weight distributions of both child's and parent's arousal models in each audio, one speaker's

high amplitude was not picked up by the other speaker's model and vice versa. Additionally, the model identified the correct speaker, i.e., implicitly learned speaker diarization, with a speaker's weight scoring higher in the parts of signals having a high amplitude of the target speaker's utterance, a signal of speaker's arousal.

The weighted-pooling layer's ability to identify regions of signals relevant to individual affect led to a statistically significant improved performance on predicting child's arousal but not parent's arousal. It could be partially explained by the dyadic interaction styles in our study. In most dyads in our dataset, the parent made more utterances than the child did because the parent read stories and led the story-related discussions. As a result, it became comparatively less crucial for models to accurately pick up the specific parts of signals relevant to parent's utterances, as averaging entire audio signals could sufficiently reflect the parent's arousal. However, each five-second instance might have only short time regions of signals relevant to a child's arousal, with the rest containing noisy frames (e.g., parent's utterances). In the child's arousal task, identifying the correct time regions and filtering out the noisy regions thus became more crucial for prediction, and the importance of the weighted-pooling layer became more evident. From the model design perspective, it was possible that the RNN layer already learned to focus on the correct speaker, particularly when one speaker made more utterances than the other in an audio; therefore, the hidden state from the RNN layer used to train the attention layer did not contain much information from the speaker who made fewer utterances, and averaging the hidden states was sufficient for predicting the arousal of the dominant speaker, i.e., parent's arousal.

5.2.3 Parent's and child's valence. For parent's valence, the most effective model, ATT(C)+DROP(CR), achieved a ρ value of 33.7%, outperforming the most effective attention-free model ($\rho = 31.6\%$) by 2.7% with a statistically significant z-test result ($p < 0.05$). For child's valence, the performance difference between the most effective attention-augmented and attention-free models was not significant. We do not show learned attention weights associated with raw speech amplitudes for valence given that speech amplitude is known to have little association with valence (unlike arousal, which often changes with speech amplitude [10]).

In both valence tasks, the weighted-pooling layer using the CNN output, ATT(C), yielded significantly higher ρ scores than the weighted-pooling layer using the RNN output, ATT(R). This finding was, however, not evident in the two arousal tasks, and evidenced that lower-level speaker's speech characteristics directly extracted from the CNN layer contained information more crucial for the weighted pooling layer to pick up time regions vital to speaker valence than the temporal domain context information processed by the RNN layer. Overall, this finding suggests that, when predicting individual affect in dyads, it is important not to assume that an attention mechanism could work equally effectively for both valence and arousal tasks and the effectiveness of an attention mechanisms depends on the affect prediction task (either valence or arousal).

5.3 Effect of dropout on model performance

In four learning tasks, the performance difference between DROP(ALL) and DROP(CR) models was not consistently observed.

In a few cases, DROP(CR) models outperformed DROP(ALL) models. For example, the performance of ATT(R) model on child's arousal improved from 53.5% to 54.7% when using DROP(CR) instead of DROP(ALL), a difference that is statistically significant ($p < 0.05$). A similar trend was observed for parent's valence where ATT(C)+DROP(CR) and ATT(R)+DROP(CR) achieved 33.7% and 30.4%, respectively, while with DROP(ALL) only 26.7% and 26.1% was achieved. These two differences are also both statistically significant ($p < 0.001$). Only in the case of using the ATT(C) model for child's arousal, having a dropout function on the model's attention layer significantly improved its performance with an increase from 51.9% to 54.0% ($p < 0.001$). This finding suggests that the dropout function on a CRDNN model's attention layer does not necessarily improve its performance, and may even impede its prediction quality in some affect recognition tasks.

In three of four learning tasks, DROP(CR) models had smaller variability in the correlation scores (ρ) during a five-fold cross validation compared to DROP(ALL) models. Specifically, the average standard deviation scores of ρ values dropped from 5.8% to 4.7%, 16.0% to 13.7%, 10.4% to 9.8% when the dropout function on the attention layer was removed from models prediction of child's valence, parent's valence, parent's arousal, respectively. This effect was more evident in the attention-augmented models than in the attention-free models. This finding suggests that having a dropout function on the attention layer could potentially reduce the robustness of a CRDNN model and lead to greater variability in its performance on unseen datasets. Overall, a dropout function on the attention layer is not recommended for future model design, particularly when the attention layer is implemented using weighted pooling.

6 CONCLUSION

Results on the new DAMI-P2C dataset show that an end-to-end learning framework can achieve competitive performance in predicting the affect of individuals in a dyadic interaction. With a weighted pooling layer added between the RNN and FC layers as the model's attention layer, the end-to-end system significantly improved its performance in terms of the Spearman correlation coefficient on predicting child and parent arousal levels. The most effective attention mechanism was different across the four learning tasks: The weighted-pooling layer using the RNN output led to a larger performance increase for child's arousal, whereas the weighted-pooling layer using the CNN output was more effective for parent's valence. Lastly, we found that the attention-augmented models are able to focus on the target speaker, providing implicit diarization of who speaks when.

We are open-sourcing this dataset of parent-child dyadic co-reading interaction, which contains LLDs of the 34 dyads' audios, arousal and valence annotations, and a diverse set of developmental, social and demographic profiles of each dyad. Participant families came from diverse social and cultural backgrounds, allowing for the exploration of building affect sensing systems personalized to individual variations in affect across dyads. The dataset is a unique contribution to the study of modeling multi-speaker speech affect in human-human interactions, and we hope that the end-to-end deep learning systems designed in this work provide a competitive affect

recognition baseline for future research in this field. Lastly, we also plan to open-source the video and tablet-based interaction log data in future, and this multimodal dataset would further increase the applicability of affect detectors to challenging real-world cases (e.g., handling technical difficulties such as background noise and poor lighting) [1] as well as recognize a greater variety of affective states (e.g., mind wandering [45, 46]).

7 LIMITATIONS AND FUTURE WORK

This work has several limitations. Since we did not directly compare the attention-augmented models trained on speaker-blind data with models trained on speaker-diarized data, our experimental results do not quantify how accurately the models are able to diarize speakers. Instead, our findings suggest only that the attention-augmented models' higher performances on predicting child's arousal and parent's valence could be attributed to their attention mechanisms that implicitly diarize speakers when picking up parts of signals vital for the learning task of interest. This indication provides an increased motivation for such direct comparisons between speaker-blind and speaker-diarized models in future work. Secondly, our system takes frame-level acoustic features as its input, but it can be extended to directly take the raw waveform by adding a time-convolutional layer [40], which can learn an acoustic model and transform the raw time-domain waveform to a frame-level feature. This extension would further reduce our system's reliance on offline pre-processing and enlarge an intelligent agent's capability of sensing and reacting to human affect in real-time interactions.

Our dataset provides many opportunities to analyze and model multi-speaker affect in parent-child co-reading interactions. For example, an end-to-end neural speaker diarization [12] or multi-task learning [51] can be integrated into our current affect prediction systems to jointly learn speaker diarization and individual speaker's valence and arousal. Further, the demographic and developmental profiles can be integrated into the current system to model multi-speaker affect in a personalized or culture-sensitive manner [37, 39], as individual differences modulate affective expressions suggested by a theoretical foundation for affect detection [9] and cultural differences were empirically found to exist in some aspects of emotions, particularly emotional arousal level [25]. Computational models that integrated a child's cultural background as an additional input modality were also found more effective in modeling children's affect and engagement than models that did not [37, 39]. Thus, personalized or culture-sensitive affect models may further improve model prediction performance, and help this technology to facilitate better parent-child dyadic experiences in the future.

8 ACKNOWLEDGEMENT

This work was supported by the MIT Integrated Learning Initiative (MITili) grant and partly by the Intel Graduate Fellowship.

REFERENCES

- [1] Nigel Bosch, Huili Chen, Sidney D'Mello, Ryan Baker, and Valerie Shute. 2015. Accuracy vs. Availability Heuristic in Multimodal Affect Detection in the Wild. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction* (Seattle, Washington, USA) (ICMI'15). Association for Computing Machinery, New York, NY, USA, 267–274. <https://doi.org/10.1145/2818346.2820739>

- [2] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette Chang, Sungbok Lee, and Shrikanth Narayanan. 2008. IEMOCAP: Interactive emotional dyadic motion capture database. *Language Resources and Evaluation* 42, 4 (2008), 335–359.
- [3] Angela Chang and Cynthia Breazeal. 2011. TinkRBook: shared reading interfaces for storytelling. In *Proc. of International Conference on Interaction Design and Children*. ACM, Ann Arbor, Michigan, 145–148.
- [4] Huili Chen, Hae Won Park, and Cynthia Breazeal. 2020. Teaching and learning with children: Impact of reciprocal peer learning with a social robot on children's learning and emotive engagement. *Computers & Education* 150 (2020), 103836. <https://doi.org/10.1016/j.compedu.2020.103836>
- [5] Huili Chen, Hae Won Park, Xiajie Zhang, and Cynthia Breazeal. 2020. Impact of Interaction Context on the Student Affect-Learning Relationship in Child-Robot Interaction. In *Proc. of the ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, Cambridge, United Kingdom, 389–397.
- [6] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *Proc. of Advances in Neural Information Processing Systems, Workshop on Deep Learning and Representation Learning*. Neural Information Processing Systems Foundation, Inc., Montréal, Canada, 1–9.
- [7] D. Cicchetti. 1994. Guidelines, Criteria, and Rules of Thumb for Evaluating Normed and Standardized Assessment Instruments in Psychology. *Psychological Assessment* 6 (1994), 284–290.
- [8] Thomas Dietterich. 1998. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation* 10, 7 (1998), 1895–1923.
- [9] Sidney K. D'Mello, Nigel Bosch, and Huili Chen. 2018. *Multimodal-Multisensor Affect Detection*. Association for Computing Machinery and Morgan & Claypool, 167–202. <https://doi.org/10.1145/3107990.3107998>
- [10] Florian Eyben. 2015. *Real-time Speech and Music Classification by Large Audio Feature Space Extraction*. Springer International Publishing, 1–338 pages.
- [11] Florian Eyben, Felix Weninger, Florian Gross, and Björn Schuller. 2013. Recent developments in openSMILE, the Munich open-source multimedia feature extractor. In *Proc. of ACM Multimedia*. ACM, Barcelona, Spain, 835–838.
- [12] Yusuke Fujita, Naoyuki Kanda, Shota Horiguchi, Yawen Xue, Kenji Nagamatsu, and Shinji Watanabe. 2019. End-to-End Neural Speaker Diarization with Self-attention. In *Proc. of ASRU*. IEEE, Singapore, Singapore.
- [13] Yarin Gal and Zoubin Ghahramani. 2015. A Theoretically Grounded Application of Dropout in Recurrent Neural Networks. [arXiv:1512.05287 \[stat.ML\]](https://arxiv.org/abs/1512.05287)
- [14] 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 *Proc. of AAAI*. AAAI Press, Phoenix, Arizona, 3951–3957.
- [15] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *Proc. of ICASSP*. IEEE, Vancouver, Canada, 6645–6649.
- [16] Kevin A Hallgren. 2012. Computing Inter-Rater Reliability for Observational Data: An Overview and Tutorial. *Tutorials in quantitative methods for psychology* 8 1 (2012), 23–34.
- [17] Betty Hart and Todd R Risley. 1995. *Meaningful differences in the everyday experience of young American children*. Paul H Brookes Publishing.
- [18] Erika Hoff. 2003. The specificity of environmental influence: Socioeconomic status affects early vocabulary development via maternal speech. *Child Development* 74, 5 (2003), 1368–1378.
- [19] Kathleen V Hoover-Dempsey and Howard M Sandler. 1995. Parental involvement in children's education: Why does it make a difference? *Teachers College Record* 97, 2 (1995), 310–331.
- [20] Che-Wei Huang and Shrikanth Shri Narayanan. 2017. Deep convolutional recurrent neural network with attention mechanism for robust speech emotion recognition. In *Proc. of International Conference on Multimedia and Expo (ICME)*. IEEE, Hong Kong, China, 583–588.
- [21] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proc. of ICMI*, Vol. 37. JMLR.org, Lille, France, 448–456.
- [22] Randy W Kamphaus and Cecil R Reynolds. 2006. *PRQ: Parenting relationship questionnaire manual*. NCS Pearson, Minneapolis, MN.
- [23] Valerie W Lawrence and Elizabeth F Shipley. 1996. Parental speech to middle- and working-class children from two racial groups in three settings. *Applied Psycholinguistics* 17, 2 (1996), 233–255.
- [24] Pengcheng Li, Yan Song, Ian Vince McLoughlin, Wu Guo, and Li-Rong Dai. 2018. An attention pooling based representation learning method for speech emotion recognition. In *Proc. of INTERSPEECH*. ISCA, Hyderabad, India, 3087–3091.
- [25] Nangyeon Lim. 2016. Cultural differences in emotion: differences in emotional arousal level between the East and the West. *Integrative Medicine Research* 5, 2 (2016), 105 – 109. <https://doi.org/10.1016/j.imr.2016.03.004>
- [26] Katrina McNab and Ruth Fielding-Barnsley. 2013. Digital Texts, iPads, and Families: An Examination of Families' Shared Reading Behaviours. *International Journal of Learning: Annual Review* 20 (2013), 53–62.
- [27] Hao Meng, Tianhao Yan, Fei Yuan, and Hongwei Wei. 2019. Speech emotion recognition from 3D log-mel spectrograms with deep learning network. *IEEE Access* 7 (2019), 125868–125881.
- [28] Seyedmahdad Mirsamadi, Emad Barsoum, and Cha Zhang. 2017. Automatic speech emotion recognition using recurrent neural networks with local attention. In *Proc. of ICASSP*. IEEE, New Orleans, LA, 2227–2231.
- [29] Susan B Neuman and Donna C Celano. 2012. Worlds Apart: One City, Two Libraries, and Ten Years of Watching Inequality Grow. *American Educator* 36, 3 (2012), 13.
- [30] S. B. Neuman and D. C. Celano. 2012. Worlds apart: One city, two libraries, and ten years of watching inequality grow. *Am. Educ. Fall* (2012), 13–23.
- [31] Susan B Neuman and Donna C Celano. 2015. *Giving our children a fighting chance: Poverty, literacy, and the development of information capital*. Teachers College Press.
- [32] 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 *Proc. of AAAI*. AAAI Press, Honolulu, Hawaii, 687–694.
- [33] Samuel P Putnam and Mary K Rothbart. 2006. Development of short and very short forms of the Children's Behavior Questionnaire. *Journal of Personality Assessment* 87, 1 (2006), 102–112.
- [34] Katherine Rice and Elizabeth Redcay. 2015. Spontaneous mentalizing captures variability in the cortical thickness of social brain regions. *Social Cognitive and Affective Neuroscience* 10, 3 (2015), 327–334.
- [35] Sarah Roseberry, Kathy Hirsh-Pasek, and Roberta M Golinkoff. 2014. Skype me! Socially contingent interactions help toddlers learn language. *Child Development* 85, 3 (2014), 956–970.
- [36] Meredith L Rowe. 2008. Child-directed speech: Relation to socioeconomic status, knowledge of child development and child vocabulary skill. *Journal of Child Language* 35, 1 (2008), 185–205.
- [37] Ognjen Rudovic, Jaeryoung Lee, Miles Dai, Björn Schuller, and Rosalind W. Picard. 2018. Personalized machine learning for robot perception of affect and engagement in autism therapy. *Science Robotics* 3, 19 (2018). <https://doi.org/10.1126/scirobotics.aao6760>
- [38] O. Rudovic, H. W. Park, J. Busche, B. Schuller, C. Breazeal, and R. W. Picard. 2019. Personalized Estimation of Engagement From Videos Using Active Learning With Deep Reinforcement Learning. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 217–226.
- [39] O. Rudovic, Y. Utsumi, J. Lee, J. Hernandez, E. C. Ferrer, B. Schuller, and R. W. Picard. 2018. CultureNet: A Deep Learning Approach for Engagement Intensity Estimation from Face Images of Children with Autism. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 339–346.
- [40] Tara N Sainath, Ron J Weiss, Andrew Senior, Kevin W Wilson, and Oriol Vinyals. 2015. Learning the speech front-end with raw waveform CLDNNs. In *Proc. of INTERSPEECH*. ISCA, Dresden, Germany, 1–5.
- [41] Mousmita Sarma, Pegah Ghahremani, Daniel Povey, Nagendra Kumar Goel, Kandarpa Kumar Sarma, and Najim Dehak. 2018. Emotion Identification from Raw Speech Signals Using DNNs. In *Proc. of INTERSPEECH*. ISCA, Hyderabad, India, 3097–3101.
- [42] Patrick E Shrout and Joseph L Fleiss. 1979. Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin* 86, 2 (1979), 420–428.
- [43] Samuel Spaulding, Huili Chen, Safinah Ali, Michael Kulinski, and Cynthia Breazeal. 2018. A Social Robot System for Modeling Children's Word Pronunciation: Socially Interactive Agents Track. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Stockholm, Sweden) (AAMAS'18). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1658–1666.
- [44] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- [45] Angela Stewart, Nigel Bosch, Huili Chen, Patrick Donnelly, and Sidney D'Mello. 2017. Face Forward: Detecting Mind Wandering from Video During Narrative Film Comprehension. In *Artificial Intelligence in Education*, Elisabeth André, Ryan Baker, Xiangen Hu, Ma. Mercedes T. Rodrigo, and Benedict du Boulay (Eds.). Springer International Publishing, Cham, 359–370.
- [46] Angela Stewart, Nigel Bosch, Huili Chen, Patrick J. Donnelly, and Sidney K. D'Mello. 2016. Where's Your Mind At? Video-Based Mind Wandering Detection During Film Viewing. In *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization* (Halifax, Nova Scotia, Canada) (UMAP'16). Association for Computing Machinery, New York, NY, USA, 295–296. <https://doi.org/10.1145/2930238.2930266>
- [47] Lori Takeuchi and Reed Stevens. 2011. The new coviewing: Designing for learning through joint media engagement. The Joan Ganz Cooney Center at Sesame Workshop.
- [48] Catherine S Tamis-LeMonda, Marc H Bornstein, and Lisa Baumwell. 2001. Maternal responsiveness and children's achievement of language milestones. *Child Development* 72, 3 (2001), 748–767.

- [49] George Trigeorgis, Fabien Ringeval, Raymond Brueckner, Erik Marchi, Mihalis A Nicolaou, Björn Schuller, and Stefanos Zafeiriou. 2016. Adieu features? End-to-end speech emotion recognition using a deep convolutional recurrent network. In *Proc. of ICASSP*. IEEE, Shanghai, China, 5200–5204.
- [50] Panagiotis Tzirakis, George Trigeorgis, Mihalis A Nicolaou, Björn W Schuller, and Stefanos Zafeiriou. 2017. End-to-end multimodal emotion recognition using deep neural networks. *IEEE Journal of Selected Topics in Signal Processing* 11, 8 (2017), 1301–1309.
- [51] Zixing Zhang, Bingwen Wu, and Björn Schuller. 2019. Attention-augmented end-to-end multi-task learning for emotion prediction from speech. In *Proc. of ICASSP*. IEEE, Brighton, UK, 6705–6709.
- [52] Zixing Zhang, Bingwen Wu, and Björn Schuller. 2019. Attention-augmented End-to-end Multi-task Learning for Emotion Prediction from Speech. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, Vol. 2019-May. Institute of Electrical and Electronics Engineers Inc., 6705–6709. <https://doi.org/10.1109/ICASSP.2019.8682896>
- [53] Huijuan Zhao, Ning Ye, and Ruchuan Wang. 2018. A Survey on Automatic Emotion Recognition Using Audio Big Data and Deep Learning Architectures. In *Proc. of International Conference on Big Data Security on Cloud (BigDataSecurity), IEEE International Conference on High Performance and Smart Computing (HPSC) and IEEE International Conference on Intelligent Data and Security (IDS)*. IEEE, Omaha, NE, 139–142.