



Published in final edited form as:

*IEEE Trans Games*. 2020 June ; 12(2): 213–218. doi:10.1109/tg.2018.2877325.

## A Mobile Game for Automatic Emotion-Labeling of Images

**Haik Kalantarian\*** [Postdoctoral Scholar],

Department of Pediatrics and Biomedical Data Science at Stanford University.

**Khaled Jedoui\*** [Undergraduate],

Department of Mathematics at Stanford University.

**Peter Washington** [Graduate student],

Department of Bioengineering at Stanford University.

**Dennis P. Wall** [Associate Professor]

Departments of Pediatrics and Biomedical Data Science at Stanford University.

### Abstract

In this paper, we describe challenges in the development of a mobile charades-style game for delivery of social training to children with Autism Spectrum Disorder (ASD). Providing real-time feedback and adapting game difficulty in response to the child's performance necessitates the integration of emotion classifiers into the system. Due to the limited performance of existing emotion recognition platforms for children with ASD, we propose a novel technique to automatically extract emotion-labeled frames from video acquired from game sessions, which we hypothesize can be used to train new emotion classifiers to overcome these limitations. Our technique, which uses probability scores from three different classifiers and meta information from game sessions, correctly identified 83% of frames compared to a baseline of 51.6% from the best emotion classification API evaluated in our work.

### Keywords

autism; emotion; crowdsourcing; mobile

## I. INTRODUCTION

Autism Spectrum Disorder (ASD) is a developmental disorder characterized by deficits in social communication and the presence of repetitive behaviors and interests [1]. The prevalence of this condition has increased in recent years, rising from an estimated 1-in-68 in 2010 to 1-in-59 in 2014 [1]. Though there is no cure for autism, multiple studies have demonstrated that Applied Behavioral Analysis (ABA) therapy can improve developmental progress and social acuity if applied consistently from a young age [2]. The application of ABA therapy is customized to suit the child's deficits and needs, often including a method of teaching called Discrete Trial Training (DTT) [3] [2].

---

\*H. Kalantarian and K. Jedoui contributed equally to this work.

A discrete trial is a unit of instruction delivered by the teacher to the child that lasts between five and twenty seconds, consisting of a prompt, response, reinforcement, and brief pause before the next trial [4]. Two areas in which DTT has been shown to be effective include 1) teaching new discriminations: the recognition of various cues often presented using flashcards, and 2) imitation: the ability to provide a response identical to the queue [4]. As difficulties in emotion recognition and expression are hallmarks of Autism Spectrum Disorder [5], [6], discrete trial training can be an integral component of a treatment program designed to address these deficits.

Caring for a child with autism can pose a significant financial burden on families [7]. This is in part due to interventions that require long hours of 1-on-1 therapy administered by trained specialists in increasingly short supply due to substantial growth in the incidence of this condition [1]. Alternatives that can ameliorate some of these challenges could be derived from digital and mobile tools. We are developing *Guess What?*: a mobile charades-style game that can potentially deliver a form of discrete trial training to children at home through a social activity shared between the child, who must interpret and act out various emotive prompts shown on the screen, and the parent, who is tasked with guessing the emotion.

Presently, the caregiver is tasked with fulfilling the *reinforcement step* of discrete trial training based on the fidelity of the child's imitation. Automatic image-based emotion recognition algorithms can supplement the reinforcement process by detecting if the child is emoting the correct prompt. This functionality can facilitate the development of additional game features that are integral aspects of ABA therapy: (1) adapting the prompts to target the child's unique deficits [2], and (2) providing the appropriate visual feedback to guide the child toward the correct behavior without diminishing the challenge.

Most commercial emotion classifiers are trained on large databases of labeled images such as the CIFAR-100, ImageNet [8], Cohn-Kanade Database [9] and Belfast-Induced Natural Emotion Databases [10]. While these datasets contain thousands of images, children are significantly underrepresented in these sources. Thus, classifiers trained on these databases are not optimized for vision-based autism research. This motivates the development of new approaches for scalable aggregation of emotive frames from children that can be used to design future classifiers and augment existing ones.

The primary contributions of this paper are as follows:

- We present a preliminary version of our mobile charades-style game, *Guess What?* [11], as it is developed into a form of discrete trial training for teaching emotion recognition and expression.
- We describe the repurposing of the game for aggregation of emotive egocentric video for autism research.
- We propose and evaluate two automatic labeling algorithms, which use probability scores from existing emotion classifiers and contextual meta-information to extract labeled frames from videos derived from these game sessions.

This paper is organized as follows. In Section II, we briefly cover related work in this area. In Section III, we describe the game design. In Section IV, we present our algorithms. In Section V, we describe our experimental methods, followed by results in Section VI and concluding remarks in Section VII.

## II. RELATED WORK

Meta classifiers are systems which consider the output of multiple classifiers as features, used to build another classifier for final class label assignment [12]. This technique has been applied to a variety of problem domains, such as grammar correction [13], news video classification, [12], and autism identification [14]. In [15], Chaibi et al. propose an ensembleclassification approach to detect emotions. Similarly, Perikos et al. propose an ensemble-based method to detect emotion from textual data using bagging and boosting techniques [16]. Our approach leverages the success of these previous ensemble- based emotion recognition techniques by using classification scores from three classifiers combined with meta information from the game session to tag each frame with an emotion with a higher accuracy than any individual classifier could achieve.

In [17], Burmania et al. evaluate the ability of raters to correctly label data based on the inclusion of reference sets within the database with a pre-determined ground truth. Sessions in which raters are performing poorly are therefore paused, which increases the accuracy of aggregated frames. This idea has also been explored in earlier works such as that of Le et al. [18]. These approaches are conceptually similar to our method, though we do not rely on human raters and instead characterize the systematic biases of individual classifiers in the determination of the final class label assignment.

In [19], the authors use a deep-learning architecture to evaluate several manual labeling techniques to develop a framework in which scores from ten raters can be combined to generate a final label with highest accuracy. Similarly, Yu et al. [20] demonstrate that an ensemble of deep learning classifiers can significantly outperform a single classifier for facial emotion recognition. In contrast with these previous approaches, our method fuses classification confidence scores with game meta information rather than selecting the label with the maximum probability. By considering both per-class probabilities and *a priori* knowledge about the prompt shown at the time, labeling accuracy is significantly improved.

A novel approach for crowdsourcing labeled expressions is described in [21]. The authors propose an iPad puzzle game in which players are recorded through the device's front camera while periodically instructed to make various expressions. This work bears similarity to our approach, though we target a younger audience and feature a social interplay between caregiver and child. Another crowdsourcing approach by Tuite et al. [22] is the Meme Quiz: a game in which users are tasked with making an expression which the system attempts to recognize based on a continuously-expanding training dataset. Results demonstrated statistically significant increases in classification accuracy over time despite an increasing numbers of classes. Unlike this platform, our system does not contain any online learning functionality in its current form. However, repeated game sessions will provide larger datasets that can be used to train a more robust emotion classifier offline.

Aside from crowdsourcing data, various educational and therapeutic games have been proposed in recent years. An example is Recovery Rapids: a gamified implementation of Constraint-Induced Movement Therapy for stroke rehabilitation in which the authors adapt therapy to the individual user in real-time [23]. A similar work within the domain of autism research is described by Harrold et al. in [24]. Rather than acquiring labeled images, the authors propose an educational iPad game which targets deficits in emotion recognition and mimicry in children with developmental delay. During game-play, the child is given a voice-instruction to imitate a face shown on the screen, which is presented beside the child's face acquired from the device's front camera. Our game, while similar, requires a caregiver to drive the experience due to the targeted age of participants.

### III. GAME DESIGN

*Guess What?* is a mobile game available for Android and iOS platforms [11] designed to be a shared experience between the child, who attempts to enact the prompt shown on the screen through gestures and facial expressions, and the parent, who is tasked with guessing the word associated with the prompt during the 90 second game session. During each session, the parent holds the phone with the screen directed outward toward the child, who is recorded with the phone's front camera. This interplay, generally structured around the inversion-problem game described in [25], has the potential to provide a social, engaging, and educational experience for the child while providing structured video for researchers to develop a dataset of semi-labeled emotion data.

While several categories of prompt are supported, the two most germane to emotion recognition and expression are *emoji*, showing exaggerated cartoon representations of emotive faces, and *faces*, which displays real photos of children. Examples of the main game screen when these two prompts are shown can be seen in Figure 1, along with the main deck selection screen that allows users to select any combination of prompts to be shown during the 90-second game session.

When the child acknowledges the correct guess, or when the parent makes the determination that the prompt has been represented correctly based on *a priori* knowledge about the image shown, parents can change the prompt by tilting the phone forward to award a point. By tilting the phone backward, the prompt is skipped without awarding a point. Immediately thereafter, a new prompt is randomly selected until the ninety seconds have elapsed. The parent can determine the correct prompt by rotating the screen laterally to peek at the screen; the prompt only changes if a tilt in the longitudinal direction is detected. To reduce the likelihood of points being awarded accidentally, tilt detection is disabled for the first two seconds following the display of a new prompt.

After the game session, parents can review the footage and elect to share the data by uploading the video to an IRB-approved secure Amazon S3 bucket fully compliant with the Stanford University's High-Risk Application security standards. Meta information is included with the video, which describes the prompts shown, timing data, and the number of points awarded. *Guess What?* is available for download online at: [guesswhat.stanford.edu](http://guesswhat.stanford.edu).

## IV. ALGORITHMS

The development of an emotion classifier that generalizes appropriately to children with ASD requires a substantial database of labeled images from this population. While videos derived from emotion-centric *Guess What?* game sessions are a natural choice to extract these frames, children are rarely able to consistently enact the correct prompts. Therefore, video segments in which images of a particular emotion are displayed cannot be used as a label for the frames contained therein. For example, preliminary results indicated that less than 25% of frames within regions of the video in which *sad* emotion were shown matched the prompt. As manual frame-by-frame analysis is a tedious and error-prone process requiring trained raters, it is desirable to develop an automatic labeling approach that leverages meta information from the game, but that is not solely reliant on it. Therefore, an additional filtering step is necessary to remove invalid frames from video segments associated with a particular emotion.

### A. Ensemble Learning for Emotion Classification

Commercial emotion recognition APIs such as Azure Emotion API [26], Amazon Rekognition [27] and Sighthound [28] are convenient platforms for emotion recognition applications, and can be used to filter out frames within a region that are discordant with the prompt. However, these platforms showed remarkably poor performance on our datasets: no classifier identified over a third of frames in categories manually labeled as *sad*, *disgusted*, or *angry*. By combining the classification scores from multiple classifiers, we are able to effectively average out the nuances associated with each classifier, increasing the robustness of our classification. Two such approaches for combining classification confidence scores to label frames derived from game videos are shown in Figure 2.

Figure 2a shows how three emotion classifiers can be used to obtain three sets of classification confidence scores for each emotion. For example, Sighthound may report a 80% chance that the frame is *happy* with a 20% chance that it is *sad*, while AWS results may indicate 70%/30% probabilities. In the *reduction-based* architecture, classification confidence scores for each emotion are normalized and combined in three different ways: min, max, and average. A final classification layer then predicts the emotion associated with the image based on a reduced feature set consisting of a single probability score per emotion.

Figure 2b shows the *aggregation-based*, approach, in which probability scores from each emotion are not combined. Rather, probabilities are concatenated into a feature vector and all are used for prediction. In this approach, the final classification layer determines the best method to integrate duplicated probabilities from multiple classifiers into a final label assignment. This approach provided generally higher performance compared to the *reduction-based, method*; detailed results are presented in Section VI.

### B. Last-Layer Classification for Label Assignment

In the final classification layer, the feature set generated from the outputs of three emotion classifiers is supplemented by the emotion of the prompt shown to the child at the time the

frame was extracted to provide contextual meta information. This is based on the intuition that a frame is more likely to be *happy* if the prompt shown to the child at the time of the frame is related to a *happy* emotion. An additional classifier leverages this feature vector to assign a final emotion label to the frame. Random Forest [29] is a supervised ensemble learning method for classification that operates by training multiple decision tree classifiers on the dataset and predicting a class based on the number of votes by each decision tree. An advantage of Random Forest over other techniques is the use of bootstrap aggregation to improve predictive accuracy and control over-fitting.

Our classifier is trained for 7-class classification based on the maximum overlapping subset of the Ekman universal emotions [30] shared across the three classifiers: neutral, happy, sad, surprised, disgusted, scared and angry. Two samples were designated as the minimum requirement to split an internal node, with at least 1 sample required for a leaf node.  $N$  features were considered when searching for the best split. For hyperparameters tuning, grid-search Cross Validation was used to determine the number of trees in each forest, as well as the maximum depth of each tree. The regression inputs were binned based on the techniques described in [31]. A total of 250 trees were used per forest, each with a maximum depth of 92. For maximum performance, final classification results are based on Leave One Out cross-validation (LOOCV): K-fold Cross Validation with K equal to the dataset size.

## V. EXPERIMENTAL METHODS

In this section, we describe the methods used to collect videos from *Guess What?*, manually label them to establish our ground truth, and leverage this dataset to validate our automatic labeling algorithms.

### A. Data Collection

To generate a dataset of *Guess What?* videos for algorithm evaluation, eight children with a diagnosis of ASD each played several *Guess What?* games in a single session administered by a member of our research staff on a Google Pixel phone running Android 8.1. The average age of participating children was 4.4 years  $\pm$  .54. Due to the non-uniform incidence of autism between genders [32]-[34] and small sample size, all participants in this study were boys. The session consisted of up to five games with the following decks in: emoji, faces, animals, sports, and jobs. However, we focus this study on the category most strongly correlated with facial affect, *faces*, which yielded a total of 8100 frames subsampled from 30 frames per second (FPS) to 5 FPS.

### B. DATA PROCESSING

To establish a ground truth, two raters manually assigned emotion labels to each frame in the selected videos based on the six Ekman universal emotions [30] with the addition of a *neutral* class. In cases when no face could be located within the frame, or the frame was too blurry to discern, reviewers did not assign a label and the frame was excluded. To reduce the burden of manual annotation, the originally 30 frame-per-second videos were subsampled to five FPS. From the selected videos within the *faces* category, a total of 1350 frames were manually labeled by the two raters. Frames were discarded in cases when the raters

disagreed. This produced a total of 1080 frames from the original 1350. A confusion matrix showing the distribution of the rater's assignments can be seen in Figure 3. The Cohen's Kappa statistic for inter-rater reliability, a metric which accounts for agreements due to chance, was 0.9. This indicates a high level of reliability between the two manual raters.

## VI. RESULTS

In this section, we demonstrate our algorithm's performance in automatic labeled frame extraction based on the ground truth established by two manual raters.

### A. Baseline Labeling Accuracy

To establish a baseline to compare performance, we ran the 1080 derived frames through three Emotion Classification APIs using no meta information and selected the class with the highest probability as the final label assignment. As shown in Table I, results were poor; the highest accuracy was achieved by Azure Emotion [27] at 51.6%. The Sighthound API [27] produced the lowest accuracy at 9.2%, followed by AWS Rekognition [26] at 10.1%. Note that the weighted F1-scores are significantly higher than accuracy based on total percentage of correctly classified instances. This indicates that these platforms are tuned to correctly recognize the most common emotions such as *happy* and *neutral* to the detriment of other less common emotions. These results preclude the integration of these emotion classifiers into *Guess What?* and necessitate novel methods to automatically label frames using ensemble-techniques and game context.

### B. Ensemble-Classifer Results: No Meta Information

The performance of the ensemble-labeling technique with no meta information is shown in Table II. Specifically, this approach is based on Random-Forest classification of a feature set of confidence scores derived from three emotion classifiers, but does not include the prompt shown at the time the frame was extracted as a feature. These results indicate that the *aggregation*-based technique shown in Figure 2b is a better labeling approach than the *reduction*-based techniques of 2a for overall classification accuracy. Furthermore, we can conclude that both *aggregation* and *reduction*-based techniques outperform all three commercial emotion recognition APIs even without the inclusion of game context.

### C. Ensemble-Classifer Results: With Meta Information

Results for ensemble-based approaches that use probability scores from all three classifiers in addition to game meta information are shown in Table III. The *aggregation* approach, which treats scores from each classifier as separate features, was associated with the highest accuracy, at 83.4%: a significant improvement over the best emotion classification API's 62.6% accuracy. While all *reduction*-based approaches had lower accuracy, performance of the max-probability and average-probability approaches were similar at 78.3% and 77.6% respectively. As before, the minimum-probability approach had the lowest accuracy at 61.0%.



## D. Comparison of Methods

Table IV shows the best possible accuracy achieved by the baseline classifier with no meta information (Azure), the ensemble technique with no meta information, and the technique with the highest accuracy: the ensemble classifier that includes game-meta information, which correctly identified 83.4% of frames.

## E. Emotion-Specific Results

Figure 4 shows the percentage of frames within each region that were correctly identified by the best commercial API evaluated (Azure), the best reduction technique (max), and the aggregation-based technique. Specifically, the reported accuracy for each emotion represents the percentage of identified frames of that class within game periods in which the image associated with that emotion were shown. It should be noted that the Azure classifier operates solely on the data frame and does not take into consideration any meta information from the game session.

Several conclusions can be drawn from this data. First, per-class results indicate that the Azure API performance was highly inconsistent between different emotions: the system performed well on *neutral* and *happy* frames but performed poorly on others. Secondly, results show that the aggregation-based technique does not outperform the max-reduction algorithm for every emotion: max-reduction identified a greater percentage of frames in *disgust* and *surprise* categories. Lastly, even the best of the two labeling techniques still could not correctly identify a majority of frames associated with *angry* and *scared* classes; this may be a consequence of a limited training set and will be addressed in future work.

## VII. CONCLUSION

In this paper, we have presented a mobile charades-style game in active development, designed to deliver emotion- recognition training to children with Autism Spectrum Disorder. We describe how this platform can be used to derive emotion-rich egocentric video for processing with automatic labeling algorithms that fuse game-meta information with probability scores to predict emotion with a higher accuracy than commercial emotion recognition APIs. In future work, the labeled frames extracted from these videos will be used to train an emotion classifier that generalizes to children with ASD, which will be integrated into the game to provide reinforcement as social deficits are addressed through gameplay via at-home discrete trial training.

## ACKNOWLEDGMENT

This study was supported by awards to D.P.W. by the National Institutes of Health (1R21HD091500-01 and 1R01EB025025-01). Additionally, we acknowledge support to D.P.W. from the Hartwell Foundation, the David and Lucile Packard Foundation Special Projects Grant, Beckman Center for Molecular and Genetic Medicine, Coulter Endowment Translational Research Grant, Berry Fellowship, Child Health Research Institute, Spectrum Pilot Program, Stanfords Precision Health and Integrated Diagnostics Center (PHIND) and Stanford's Human Centered Artificial Intelligence Program. The Dekeyser and Friends Foundation, the Mosbacher Family Fund for Autism Research and Peter Sullivan provided additional funding. Haik Kalantarian would like to acknowledge support from the Thrasher Research Fund and Stanford NLM Clinical Data Science program (T-15LM007033-35).



## Biography



**Haik Kalantarian** is a Postdoctoral Scholar in the Department of Pediatrics and Biomedical Data Science at Stanford University, where he is investigating mobile and wearable therapies for autism. He completed his PhD in Computer Science at the University of California, Los Angeles in 2016.



**Khaled Jedoui** is currently an undergraduate student at Stanford University, where he is studying Mathematics and Computer Science. He is investigating mobile and ubiquitous computing technologies in the Wall Lab.



**Peter Washington** is currently a Graduate student in the Department of Bioengineering at Stanford University. His research interests include ubiquitous technology and human-computer interaction.



**Dennis P. Wall** is an Associate Professor of Pediatrics and Biomedical Data Science at Stanford Medical School. He is the director of the Wall Lab, focused on developing methods in biomedical informatics to disentangle complex conditions that originate in childhood and perpetuate through the life course, including autism and related developmental delays.

## References

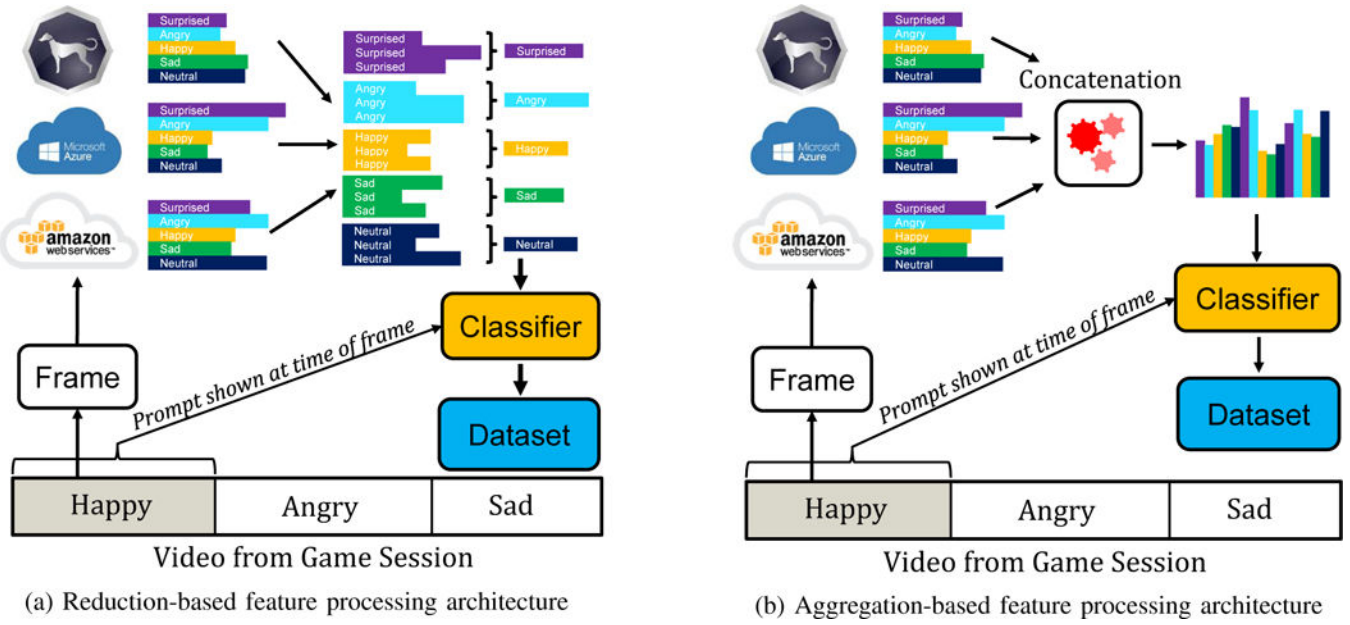
- [1]. Baio J, Wiggins L, Christensen DL, Maenner MJ, Daniels J, Warren Z, Kurzius-Spencer M, Zahorodny W, Rosenberg CR, White T et al., “Prevalence of autism spectrum disorder among children aged 8 years,” *MMWR Surveillance Summaries*, vol. 67, no. 6, p. 1, 2018.
- [2]. Foxx RM, “Applied behavior analysis treatment of autism: The state of the art,” *Child and Adolescent Psychiatric Clinics*, vol. 17, no. 4, pp. 821–834, 2008.

- [3]. Council NR et al., *Educating children with autism*. National Academies Press, 2001.
- [4]. Smith T, “Discrete trial training in the treatment of autism,” *Focus on autism and other developmental disabilities*, vol. 16, no. 2, pp. 86–92, 2001.
- [5]. Harms MB, Martin A, and Wallace GL, “Facial emotion recognition in autism spectrum disorders: a review of behavioral and neuroimaging studies,” *Neuropsychology review*, vol. 20, no. 3, pp. 290–322, 2010. [PubMed: 20809200]
- [6]. Macdonald H, Rutter M, Howlin P, Rios P, Conteur AL, Evered C, and Folstein S, “Recognition and expression of emotional cues by autistic and normal adults,” *Journal of Child Psychology and Psychiatry*, vol. 30, no. 6, pp. 865–877, 1989. [PubMed: 2592470]
- [7]. Horlin C, Falkmer M, Parsons R, Albrecht MA, and Falkmer T, “The cost of autism spectrum disorders,” *PloS one*, vol. 9, no. 9, p. e106552, 2014.
- [8]. He K, Zhang X, Ren S, and Sun J, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [9]. Cohn J et al., “Cohn-kanade au-coded facial expression database,” Pittsburgh University, 1999.
- [10]. Douglas-Cowie E, Cowie R, and Schröder M, “A new emotion database: considerations, sources and scope,” in *ISCA tutorial and research workshop (ITRW) on speech and emotion*, 2000.
- [11]. [Online]. Available: [guesswhat.stanford.edu](http://guesswhat.stanford.edu)
- [12]. Lin W-H and Hauptmann A, “News video classification using svm- based multimodal classifiers and combination strategies,” in *Proceedings of the tenth ACM international conference on Multimedia ACM*, 2002, pp. 323–326.
- [13]. Gamon M, “Using mostly native data to correct errors in learners’ writing: a meta-classifier approach,” in *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics Association for Computational Linguistics*, 2010, pp. 163–171.
- [14]. Papageorgiou EI and Kannappan A, “Fuzzy cognitive map ensemble learning paradigm to solve classification problems: Application to autism identification,” *Applied Soft Computing*, vol. 12, no. 12, pp. 3798–3809, 2012.
- [15]. Chaibi MW, “An ensemble classifiers approach for emotion classification,” in *International Conference on Intelligent Interactive Multimedia Systems and Services Springer*, 2017, pp. 99–108.
- [16]. Perikos I and Hatzilygeroudis I, “A classifier ensemble approach to detect emotions polarity in social media.” in *WEBIST (1)*, 2016, pp. 363–370.
- [17]. Burmania A, Parthasarathy S, and Busso C, “Increasing the reliability of crowdsourcing evaluations using online quality assessment,” *IEEE Transactions on Affective Computing*, vol. 7, no. 4, pp. 374–388, 2016.
- [18]. Le J, Edmonds A, Hester V, and Biewald L, “Ensuring quality in crowdsourced search relevance evaluation: The effects of training question distribution,” in *SIGIR 2010 workshop on crowdsourcing for search evaluation*, vol. 2126, 2010.
- [19]. Barsoum E, Zhang C, Ferrer CC, and Zhang Z, “Training deep networks for facial expression recognition with crowd-sourced label distribution,” in *Proceedings of the 18th ACM International Conference on Multimodal Interaction ACM*, 2016, pp. 279–283.
- [20]. Yu Z and Zhang C, “Image based static facial expression recognition with multiple deep network learning,” in *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction ACM*, 2015, pp. 435–442.
- [21]. Tan CT, Sapkota H, and Rosser D, “Befaced: a casual game to crowdsource facial expressions in the wild,” in *CHI’14 Extended Abstracts on Human Factors in Computing Systems ACM*, 2014, pp. 491–494.
- [22]. Tuite K and Kemelmacher I, “The meme quiz: A facial expression game combining human agency and machine involvement.” in *FDG*, 2015.
- [23]. Maung D, Crawfis R, Gauthier LV, Worthen-Chaudhari L, Lowes LP, Borstad A, McPherson RJ, Grealy J, and Adams J, “Development of recovery rapids-a game for cost effective stroke therapy.” in *FDG*, 2014.


- [24]. Harrold N, Tan CT, and Rosser D, "Towards an expression recognition game to assist the emotional development of children with autism spectrum disorders," in Proceedings of the Workshop at SIGGRAPH Asia ACM, 2012, pp. 33–37.
- [25]. Von Ahn L and Dabbish L, "Designing games with a purpose," Communications of the ACM, vol. 51, no. 8, pp. 58–67, 2008.
- [26]. [Online]. Available: <https://azure.microsoft.com/en-us/services/cognitive-services/emotion/>
- [27]. [Online]. Available: <https://aws.amazon.com/rekognition/>
- [28]. [Online]. Available: <https://www.sighthound.com/products/cloud>
- [29]. Breiman L, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 10 2001 [Online]. Available: 10.1023/A:1010933404324
- [30]. Ekman P, Friesen WV, O'sullivan M, Chan A, Diacoyanni- Tarlatzis I, Heider K, Krause R, LeCompte WA, Pitcairn T, Ricci-Bitti PE et al., "Universals and cultural differences in the judgments of facial expressions of emotion." Journal of personality and social psychology, vol. 53, no. 4, p. 712, 1987. [PubMed: 3681648]
- [31]. Friedman J, Hastie T, and Tibshirani R, The elements of statistical learning. Springer series in statistics New York, 2001, vol. 1.
- [32]. "Autism society: What is autism?" <http://www.autism-society.org/what-is/>, accessed: 2017-010-30.
- [33]. Dawson G, "Early behavioral intervention, brain plasticity, and the prevention of autism spectrum disorder," Development and psychopathology, vol. 20, no. 3, pp. 775–803, 2008. [PubMed: 18606031]
- [34]. Dawson G and Bernier R, "A quarter century of progress on the early detection and treatment of autism spectrum disorder," Development and psychopathology, vol. 25, no. 4pt2, pp. 1455–1472, 2013. [PubMed: 24342850]



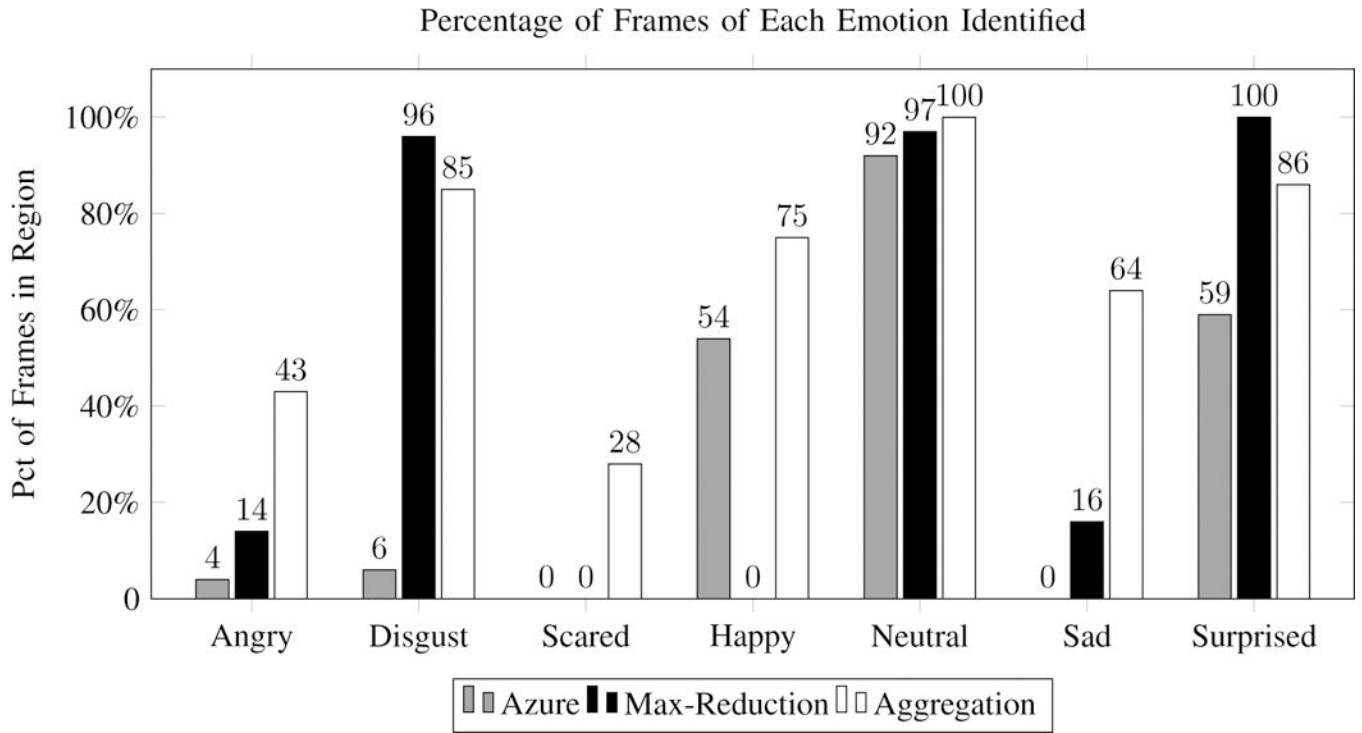
**Fig. 1:** In this mobile charades-style game, various prompts to the child during a 90-second game session. The parent attempts to guess the prompt as the child acts it out.



**Fig. 2:**  
 (a) In the reduction-based feature processing algorithm, probability scores from each classifier are combined on a per-class basis before the final classification layer using *min*, *max*, or *average* functions. (b) In the aggregation-based feature processing algorithm, redundant probability scores from different classifiers are tagged as separate features.



**Fig. 3:** The confusion matrix of the two raters assignments of frames into emotion categories. The abbreviations are: happy, sad, angry, disgust, neutral, scared, and surprised.



**Fig. 4:** A comparison of Azure Emotion API’s performance with our max-reduction and aggregation techniques on a per-class basis shows that our algorithm can correctly label the majority of frames in every category except *scared* and *angry*, and our aggregation-based method outperforms the best of the three commercial emotion recognition APIs in every category.



**TABLE I:**

Baseline accuracy with existing platforms

<b>Technique</b>	<b>Accuracy</b>	<b>Weighted F1-Score</b>
Azure Emotion	51.6%	0.62
AWS Rekognition	10.1%	0.12
Sighthound	9.2%	0.11

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE II:**

Ensemble classification without meta information

Technique	Features	Accuracy	Weighted F1-Score
Minimum probability	7	47.5%	0.36
Maximum probability	7	66.8%	0.64
Average probability	7	66.4%	0.64
All probabilities	21	76.6%	0.75

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE III:**

Ensemble classification with meta information

Technique	Features	Accuracy	Weighted F1-Score
Minimum probability	8	61.0%	0.54
Maximum probability	8	78.3%	0.77
Average probability	8	77.5%	0.77
All probabilities	22	83.4%	0.84

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE IV:**

Overall accuracy comparison

<b>Technique</b>	<b>Accuracy</b>	<b>Weighted F1-Score</b>
Baseline (best)	51.6%	0.62
Ensemble without meta (best)	76.6%	0.75
Ensemble with meta (best)	83.4%	0.84

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript