



Detecting Gaze Towards Eyes in Natural Social Interactions and Its Use in Child Assessment

EUNJI CHONG and KATHA CHANDA, Georgia Institute of Technology

ZHEFAN YE, University of Michigan

AUDREY SOUTHERLAND and NATANIEL RUIZ, Georgia Institute of Technology

REBECCA M. JONES, Weill Cornell Medicine

AGATA ROZGA* and JAMES M. REHG*, Georgia Institute of Technology

Eye contact is a crucial element of non-verbal communication that signifies interest, attention, and participation in social interactions. As a result, measures of eye contact arise in a variety of applications such as the assessment of the social communication skills of children at risk for developmental disorders such as autism, or the analysis of turn-taking and social roles during group meetings. However, the automated measurement of visual attention during naturalistic social interactions is challenging due to the difficulty of estimating a subject's looking direction from video. This paper proposes a novel approach to eye contact detection during adult-child social interactions in which the adult wears a point-of-view camera which captures an egocentric view of the child's behavior. By analyzing the child's face regions and inferring their head pose we can accurately identify the onset and duration of the child's looks to their social partner's eyes. We introduce the Pose-Implicit CNN, a novel deep learning architecture that predicts eye contact while implicitly estimating the head pose. We present a fully automated system for eye contact detection that solves the sub-problems of end-to-end feature learning and pose estimation using deep neural networks. To train our models, we use a dataset comprising 22 hours of 156 play session videos from over 100 children, half of whom are diagnosed with Autism Spectrum Disorder. We report an overall precision of 0.76, recall of 0.80, and an area under the precision-recall curve of 0.79, all of which are significant improvements over existing methods.

CCS Concepts: •**Human-centered computing** → **Ubiquitous and mobile computing**; •**Computing methodologies** → **Computer vision**; •**Applied computing** → *Psychology*;

Additional Key Words and Phrases: Wearable camera, machine learning, deep learning, computer vision, autism spectrum disorder, eye contact, gaze classification, assessment

ACM Reference format:

Eunji Chong, Katha Chanda, Zhefan Ye, Audrey Southerland, Nataniel Ruiz, Rebecca M. Jones, Agata Rozga*, and James M. Rehg*. 2017. Detecting Gaze Towards Eyes in Natural Social Interactions and Its Use in Child Assessment. *PACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 43 (September 2017), 20 pages.

DOI: <http://doi.org/10.1145/3131902>

* Rozga and Rehg share senior authorship of this work. This study was funded in part by the Simons Foundation under grants 336363 and 383667, as well as the National Science Foundation under grant IIS-1029679.

Authors' addresses: E. Chong, K. Chanda, A. Southerland, Nataniel Ruiz, A. Rozga and J. M. Rehg, Center for Behavioral Imaging and School of Interactive Computing, College of Computing, Georgia Institute of Technology; Z. Ye, Computer Science and Engineering, University of Michigan; R. M. Jones, Weill Cornell Medicine, Center for Autism and the Developing Brain.

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

© 2017 ACM. 2474-9567/2017/9-ART43 \$15.00

DOI: <http://doi.org/10.1145/3131902>




Fig. 1. **Our data collection setup.** Examiner wears a camera embedded in the bridge over the nose, thus the camera is naturally aligned with the adult's eyes and when the child looks to the adult's eyes it will be captured as a look towards the camera.

1 INTRODUCTION

Eye contact is one of the most basic and powerful forms of nonverbal communication that humans use from the first months of life [3]. It plays a crucial role in social interactions where it is used for various purposes: to express interest and attentiveness, to signal a wish to participate, and to regulate interactions [1, 28]. Moreover, eye contact is a key constituting element in joint attention, in which gaze and gestures are used to spontaneously create or indicate a shared point of reference with another person [36]. Individuals with autism spectrum disorder (ASD), a group of developmental disorders characterized by difficulties in engaging with the social world, show atypical patterns of gaze, eye contact, and joint attention; they look less at the eyes of others and respond less to the calling of one's name than their typically developing (TD) peers. These patterns have been identified as among the earliest indicators of autism in the first two years of life [23, 27, 44], and continue to remain distinctive throughout childhood and adolescence [29, 48].

ASD affects a large population of children: 1 in 68 children in the United States have ASD and approximately 1 – 2% of children are affected worldwide [15]. While the cause of ASD is unknown, the consensus among experts is that early diagnosis and intervention can substantially increase positive outcomes [8]. The development of automated measures for the early behavioral signs of autism (e.g., lack of eye contact) would facilitate large-scale screening and decrease the age of diagnosis. In addition, the continuous measurement of social behaviors such as eye contact in children who are receiving treatment would be extremely valuable in tailoring interventions and quantifying their effectiveness.

In spite of its importance as a measure of social communication, there exist very few methods for measuring eye contact during naturalistic social interactions. The only widely-utilized paradigm occurs in the context of psychology research studies, in which social interactions are recorded from one or more cameras and manual annotation by multiple human raters provides a measure of eye contact. This is an inherently subjective determination and it clearly does not scale to broad use in screening and outcome assessment.

Eye tracking technology provides another widely-used approach for automatically assessing gaze behavior in social contexts. However, it imposes substantial constraints on the interaction, requiring a child to either passively view content on a monitor screen or wear head-mounted eye tracking hardware. Neither of these

scenarios are appropriate for our target use case of naturalistic, face-to-face interactions. Since face-to-face interactions between an adult and a child are the basis for all screening, diagnosis, treatment, and assessment, it is a critical use case. While eye tracking technology is likely to soon be available on a large scale through its integration into tablets and laptop screens, it is unclear to what extent a child's gaze behavior in such settings reflects their gaze behavior with real-life social partners [16].

To address the limitations of existing approaches, we propose to utilize a point-of-view (POV) camera worn by an adult social partner to measure a child's eye contact behavior. This egocentric setting is ideally suited for social behavior measurement as the head-worn camera provides high quality visual data, e.g. consistent near-frontal views of faces, high resolution images of the eye regions, and less occlusion. Our approach capitalizes on the increased availability of POV platforms such as Pivothead and Snapchat Spectacles. In these platforms, a high definition outward-facing camera is integrated into a pair of glasses. For example, in the Pivothead, the camera is located at the bridge of the nose. Thus the camera is naturally aligned with the adult's eyes and when the child looks to the adult's eyes it will be captured as a look towards the camera (see Fig 1a).

Detecting looks to the eyes in POV videos is still a challenging problem, however, due to the diverse and varied appearance of the human eye and the role of head pose in determining gaze direction. Specifically, the estimation of head pose is a critical factor in correctly interpreting the gaze direction based on analysis of the eye regions [32]. In our prior work on POV-based eye contact detection [58], we addressed this issue by developing a set of machine learning-based eye contact detectors which were specialized for different configurations of head pose. Thus each detector was tuned for a particular range of child head poses. A strong disadvantage of this approach is the need to discretize the head pose and train independent detection models. In this paper, we remove this limitation through a novel deep learning solution in which head pose is jointly estimated along with eye contact detection.

This paper makes the following contributions:

- (1) We present a novel deep learning architecture for eye contact detection in POV video that outperforms previous approaches [50, 58] by a large margin: 23.8% gain over [58] and 50% gain over [50] in F_1 score; 38.6% gain over [58] and 64.5% gain over [50] in area under the precision-recall curve
- (2) We conducted a systematic evaluation of our method on the largest corpus of naturalistic child social interactions available to-date, consisting of 22 hours of video from 100 children and encompassing four different interaction contexts and subjects with and without ASD
- (3) We will make our model for eye contact detection and our code for training and testing freely-available to the research community¹

2 RELATED WORK

There are three categories of relevant prior work. First, we compare to the small number of previous works that addressed eye contact detection. Second, we describe recent works on large scale eye tracking that also make use of deep models for appearance analysis of eye regions. Third, we briefly review the use of classical eye tracking methods in autism.

2.1 Eye Contact Detection

Three prior works addressed the direct estimation eye contact from video and constitute the closest related work [47, 50, 58]. The most relevant is our previous paper [58] which introduced the POV camera paradigm for eye contact detection. In comparison, this paper introduces a novel detection architecture based on deep learning which dramatically improves the detection performance. We also present the first experimental results for children with ASD, demonstrating that diagnostic status does not have a significant impact on detection

¹See <http://cbi.gatech.edu/eyecontact/> for details.

performance. In addition, we present the first thorough experimental evaluation of our approach on 100 children across four different interaction contexts. Note that an abbreviated version of our prior POV work also appeared in [57].

Shell et. al. [47] developed an approach to eye contact detection based upon classical gaze tracking methods, using IR diodes on a pair of glasses to create glints on the eyes of the social partner. The reliance on special IR illumination greatly limits the usefulness of the method for naturalistic interaction. In their work on gaze locking [50], Smith et. al. addressed a different application of eye contact, namely the use of gaze to an embedded camera as a user-interface technology in an internet-of-things context. Their approach predated the wide-spread use of deep models and their dataset consisted of subjects in a chinrest, with the consequence that they could not present results for naturalistic interactions.

2.2 Appearance-Based Gaze Estimation and Existing Datasets

Traditional approaches to gaze estimation utilize active IR illumination to both create glints on the surface of the eye and reliably segment the pupil opening using a variety of dark and light pupil methods [21]. Recently, a number of investigators have explored alternative approaches to gaze estimation using appearance-based methods that analyze the eye region in conventional RGB images and avoid the use of structured illumination [30, 51, 60, 61]. We share with these methods the observation that the analysis of the eye region pattern in combination with head pose is a viable alternative to conventional gaze tracking technology. However, there are two main differences between our approach and these others. The first is that these prior works address the traditional eye tracking goal of determining the user's point of gaze on a display surface, as motivated by the widespread availability of user-facing cameras in tablets and laptop screens. These methods cannot be applied directly in our context of naturalistic face-to-face social interactions. Second, previously published methods [30, 51, 60] utilize separately detected keypoints or eye region information as an additional input at testing time, whereas we present a full-face analysis approach which does not require any auxiliary information at testing. The recent unpublished work [61] also describes a full-face analysis approach, but uses a different learning architecture from ours.

Like [30, 60, 61], we develop a deep CNN architecture to learn discriminative facial feature maps that encode gaze information. Unlike these works, we utilize the two stream architecture depicted in Fig 5, which is trained to output a head pose estimate along with the eye contact prediction. In contrast, [61] uses spatial weighting to combine CNN features, while [30] and [60] use a separate estimation pipeline to compute facial landmarks and facial regions which are then combined with the deep CNN feature analysis. Our approach works directly with detected full-face regions and does not require externally-provided head pose estimates or landmark estimation, simplifying our analysis method significantly.

There are several publicly-available gaze datasets available to the research community. However, as these datasets were collected from interactions with screens, they are not suitable for building a model that describes face-to-face social gaze behaviors. Example images for the major public datasets are illustrated in Fig 2. As these images demonstrate, the facial expressions and poses obtained from adult subjects interacting with screens are much less diverse than the variations in children's appearance and pose that we encounter in our setting (see Fig 6 for examples). We will release our model trained with 22 hours of POV video consisting of social interactions with 100 child subjects. A subset of our training data is available as part of the MMDB dataset [42].²

2.3 Gaze Behavior and Autism

A significant amount of prior work has used eye tracking to investigate differences in patterns of looking in individuals with autism, such as reduced looks to social stimuli [7]. For example, toddlers with ASD spend more time looking at geometric shapes than human biological motion [41], children with ASD devote less attention to

²See <http://www.cbi.gatech.edu/mmdb/> for MMDB access instructions.




Fig. 2. **Examples of publicly available gaze datasets.** As these datasets were collected from interactions with screens, they are not suitable for building a model that describes face-to-face social behaviors.

faces while watching videos of social interactions [22], and both adults [29] and children [6] with ASD show preferential fixations to the mouth than to the eyes when viewing social scenes. However, all of these studies have been conducted in a highly-controlled environment in which the subjects were passively viewing a monitor screen for a short period of time. At present, we still lack a comprehensive understanding of similarities and differences in gaze behavior associated with autism [20]. Our work on eye contact can potentially complement this existing literature by providing insight into patterns of looking within a naturalistic social context.

Head-mounted eye tracking systems provide an alternative to monitor-based studies of gaze behavior and have been used in a limited number of studies involving children with ASD [35, 38]. Findings from this work exhibit concordance with monitor-based gaze studies, but have also identified novel gaze patterns [38], suggesting that more research is needed to understand gaze behavior in ecologically valid settings. A basic problem with using any form of eye tracking to analyze face-to-face gaze is the need to identify the gaze target given the estimated gaze direction. In other words, wearable eye tracking gives the location of the point of regard in a POV image, but does not directly answer the question of what gaze target is present at that location. This difficulty substantially increases the complexity of a fully-automated behavior measurement system based on wearable eye tracking. These issues, along with the challenges of compliance in requiring children to wear special hardware [46], have limited the broad-scale applicability of this approach.

We note that a final, classical approach to obtaining social gaze measurements, including measurements of eye contact, is to manually annotate videos recorded in the lab setting or even home videos [62]. While this method is completely non-invasive and naturalistic, it is extremely time consuming and is subject to human error. Our previous work has shown that video from POV cameras is an effective medium for human annotation of children’s social gaze [12]. We utilize such annotations to construct the training and testing sets for our experiments.

3 METHOD

In order to detect eye contact in POV video we must first identify all of the faces that are present and then analyze each face separately to make a determination of eye contact. We describe our novel method for eye contact detection based on full-face analysis in Sec 3.1. This section assumes that all relevant faces have been detected. We describe our approach to face detection in Sec 3.2. In Sec 3.3, we describe an approach to identifying the detected face that corresponds to the child social partner, in the case where more than one face is present in an image. All of these analysis steps make use of modern deep learning architectures, yielding high accuracies relative to the previous state-of-the-art. We test our entire face analysis pipeline and separately evaluate each component in Sec 5.




Fig. 3. **Pose-dependent egocentric eye contact detector (PEEC)**. From an input image, face is detected and head pose estimated. Image is assigned to one of three clusters based on head pose, eye features are extracted, and cluster-specific random forest classifier is used to predict the final output.

3.1 Eye Contact Detection

Here we describe our main approaches to eye contact detection. Our prior work was based on a pose-dependent detection model which is summarized in Sec 3.1.1 for completeness. We then describe our new method based on a deep CNN architecture in Sec 3.1.2.

3.1.1 Pose-Dependent Egocentric Eye Contact (PEEC) Detector. Fig 3 illustrates the pipeline for the PEEC method we introduced in [58]. The input is an egocentric video frame and the output is a prediction of eye contact. Face detection is used to extract a bounding box containing the child's face. Given a bounding box, we identify facial landmark points and estimate the head pose using the IntraFace method from [9]. IntraFace extends cascaded pose regression [10] with a Supervised Descent Method for tracking facial landmarks. The facial landmarks are then used to estimate the head pose with three degrees of freedom- yaw, pitch and roll. These preprocessing stages are illustrated in Fig 4. Using the results of the facial landmark identification, we crop the eye regions and extract Histogram of Oriented Gradients (HOG) features [14] from the cropped eye patches. Using the estimated head pose, the training data is divided into three groups and a binary classifier for eye contact is trained for each group. At testing time, all three classifiers are applied to each frame and the results are aggregated. We now describe each of these steps in more detail.

Head Pose Clustering - This is an important step because the appearance of the human eye during moments of eye contact varies with the head pose. The head pose of the child's face h_t is estimated at frame t by registering the IntraFace landmarks onto an average 3D face. We cluster head poses using a Gaussian mixture model and train a separate classifier for each pose cluster. Since rotation can be resolved by affine transformations, only the pitch and yaw are used for head pose clustering, with the number of clusters set to 3 for all the experiments.

Feature Extraction - Left and the right eye patches from each face are cropped using IntraFace facial landmarks and an affine transformation is used to align tilted eyes. We resize all of the images to a fixed size of 73×37 . We extract HOG appearance features A_t from both the left and right eyes at frame t and concatenate them to create a feature vector.

Eye-Contact Detection - We assign a label y_t to each frame t , where y_t is 1 if there is eye contact and 0 otherwise. Let the appearance features be denoted by A_t and the head pose by H_t for the frame t . Then the conditional probability $P(y_t|H_t, A_t)$ is given by:

$$P(y_t|H_t, A_t) = \sum_{h_c=1}^k P(h_c|H_t)P(y_t|h_c, A_t),$$




Fig. 4. **Preprocessing steps in PEEC prior to eye contact detection.** Left to right: 1) Input image; 2) All faces are detected; 3) Child's face is selected; 4) Facial landmarks are localized; and 5) Head pose is estimated.

where h_c enumerates the head pose clusters in the Gaussian mixture model. Each cluster is considered as a sub-problem and the classifier outputs from each cluster are averaged together. We utilize a Random Forest classifier [4] with 100 trees and 10 splits per node.

3.1.2 Pose-implicit Convolutional Neural Networks (PiCNN) Detector. There are two major problems with the PEEC approach from Sec 3.1.1. First, due to the fact that the method requires head pose estimates (to assign each sample to one of the pose clusters) and eye localizations (to extract features), gaze estimation cannot be performed when landmark detection fails. Thus any failure in steps 4 and 5 of the preprocessing pipeline of Fig 4 will result in a missed detection. Facial landmark detection is a more difficult problem than face detection. It is challenging to detect the landmarks reliably when the face is occluded or foreshortened due to out-of-plane rotation, which happens quite frequently during dynamic social interactions with children (see Fig 7). Our experiments reveal that, in our dataset, landmarks are successfully found in only 75.56% of the total eye contact frames, whereas faces are detected 97.94% of the time (Table 3). The second problem with PEEC is its reliance on hand-designed HOG features and a random forest classifier, which are known to be inferior to modern deep neural network classifiers which support end-to-end feature learning.

We therefore propose a new classification architecture based on Convolutional Neural Networks (CNN) [31] that addresses these two issues. Our new approach learns an image representation that *jointly* predicts head pose and eye contact, instead of being dependent on precomputed head pose and facial landmarks. Since the representation is learned end-to-end (i.e., the input is raw pixels and the output is a binary prediction) there is the opportunity to learn features that are specific to the task. This approach is particularly promising in light of the large-scale social interaction dataset that we have assembled. Our proposed model is called “*Pose-implicit CNN Detector*” or *PiCNN*.

Fig 5 illustrates the PiCNN architecture. The input to the network is a rectangle of pixels corresponding to a face bounding box produced by a face detector (see Sec 3.2). The detected face patch is resized to 227×227 . The network has five convolutional and pooling layers, similar to AlexNet [31], but with a smaller filter size (7 by 7) with strides of 2 in the first layer to capture the finer details in the face. Layers 6 to 8 are fully connected. The primary difference between our approach and AlexNet is the presence of two branches in the 7th and 8th fully-connected layers. The upper branch outputs a prediction of the three axes of head rotation (i.e., regression in yaw, pitch, roll) and the lower branch is tasked with the binary classification of eye contact. The weights and parameters are the same and are linked during training in the first 6 layers (5 convolutional and 1 fully connected) and branched in the last 2 layers to facilitate multi-task learning. Thus the prediction of eye contact can benefit




Fig. 5. **Pose-implicit Convolutional Neural Networks (PiCNN) architecture for eye contact detection.** We utilize a two stream network structure with 8 layers modeled after AlexNet. One output branch predicts head pose and the other predicts eye contact. The inputs to the network are face bounding boxes. A representative face bounding box is illustrated with a red outline in the image at the left.

from implicitly learning the variability in eye appearance resulting from head pose change. This model achieves the best performance, by a large margin, among all existing methods on our dataset, as presented in Sec 5.3.

In training the PiCNN model, we have ground truth eye contact labels for every frame in the training dataset. This allows us to backpropagate training error through the eye contact detection branch in every training batch.³ However, we do not have frame-level human annotations of head pose for any frames. We solve this problem by using the IntraFace system as a source of head pose training data. In frames for which head pose estimates from IntraFace are available, we additionally backpropagate training error through the head pose branch. Note that this strategy has the benefit that the network is forced to learn representations for predicting eye contact even when a head pose reference is not available, making it potentially more robust than a sequential approach in which head pose must be computed before detection can be performed. At the same time, our method can take advantage of sparse head pose annotations where they are available and improve the detection performance. During testing time, we apply the PiCNN model in feedforward mode and do not utilize any separate head pose estimates. Now that we have presented our two approaches to eye contact detection, PEEC and PiCNN, we briefly describe the related methods which are used as baselines in our experiments in Sec 5.

3.1.3 Modified AlexNet. One important performance baseline for our method is the architecture in Fig 5 with the head pose prediction branch removed. Modulo some small differences in the convolutional layers, this is equal to the standard AlexNet implementation [31]. This approach requires the model to learn features which encode head pose cues without having access to an explicit training signal. Our performance gains over this model demonstrate the utility of our two branch approach and the benefit of providing an explicit reference for head pose during training.

3.1.4 Gaze Locking. We have implemented the Gaze Locking method from [50] to compare its performance against our method. Using the facial landmark detection results, we detect the location of the eyes in each frame. Eye tilt is corrected using affine transformations and each eye is cropped to a size of 37 x 73. We then concatenate the intensity pixels from the left and the right eye to form a high dimensional feature vector. This is then projected onto a low dimensional space by performing Principal Component Analysis, (PCA) [52], followed by Multiple Discriminant Analysis, (MDA) [11]. The resultant feature vector is then fed into a support vector machine [5], which is trained offline on the training set. We choose to reduce our high dimensional feature vector into 200 dimensions using PCA. We find that compressing this further with MDA into a 6 dimensional subspace gives good results. Our support vector machine is used with default parameters as described in the original paper. The SVM performs a binary classification on each frame.

³See [17] for details on neural network training in general, and Sec 5.3 for the details in our approach.

3.2 Face Detection

All of the eye contact detection methods described above require a detected face bounding box as input. Here we describe the approach to face detection used in our experiments. There exists a substantial literature on face detection, with approaches such as [53] seeing widespread use prior to the advent of deep learning. In the same way that deep CNNs have come to dominate the fields of object recognition and object detection, face detection has similarly turned to CNNs, yielding impressive results in benchmark datasets such as [24] and [56]. Recently, complex pipelines have been built for this task to deal with issues such as distant faces and low quality images [55, 56, 59]. Since we do not encounter these problems in our dataset, we adopt a simpler approach.

The general purpose object detection method Faster R-CNN [43] achieved impressive performance in object detection challenges such as Microsoft COCO [33] and PASCAL VOC [13]. Faster R-CNN generates object proposals from convolutional feature maps of the image using a trainable Region Proposal Network. These proposals are then evaluated by the Fast R-CNN detection network which outputs high confidence object detections. Jiang et al. [26] demonstrated that Faster R-CNN can be trained to obtain a highly-accurate face detector. Similarly to [26], we train Faster R-CNN on the WIDER FACE training dataset. We analyze our entire dataset using this detector and output a list of face bounding boxes for every frame.

3.3 Child Face Selection

In general, face detectors will identify all faces in the environment, including a parent's face (see Fig 4), a sibling's face, and occasional face-like spurious patches, along with the child's face. Since we are only interested in the child's social behavior, we need to select the child's face and reject the other face patches. To classify individual faces we adopt an online appearance learning approach, where in the beginning a pre-defined number of face classes are initialized (ranging from one to four in our experiments), a simple classifier is learned, in the next frame each detected face is added to one of the classes, and then the classifier is updated and the process is repeated. For facial appearance feature extraction, we use the VGG-face model [40], a pre-trained deep network for face recognition which is well-matched to our application. For online classification, we use logistic regression with sequential updates. The combination of deep network features with on-line classifier learning successfully identifies the child's face at very high precision and recall (both greater than 90%, details in Table 2 and 3).

4 DATASETS

The data utilized in our experiments comes from four separate studies, conducted at one or more of the following three sites: the Georgia Tech Child Study Lab in Atlanta, GA (GT); the Center for Autism and the Developing Brain in White Plains, NY (CADB); and the Marcus Autism Center in Atlanta, GA (MAC). Descriptive information for the subset of participants whose data was included in the current analysis and details of the data collection protocol for each study is detailed below, separately for each dataset, and in aggregate in Table 1. Sample images from our dataset are illustrated in Fig 6.

Generally, all four data collection protocols involved a semi-structured play interaction between an adult and a child, who sat across a small table from each other. The specific protocols chosen were selected based on their prior use in research on social attention and communication in typically developing children and children with autism, and because they have been shown to reliably elicit eye contact from these groups. In all four studies, the examiner interacting with the child wore a pair of commercially-available glasses - Pivthead Kudu - which have an outward-facing camera embedded in the bridge over the nose. By virtue of its placement, this camera reliably captures a close-up image of the child's face as the examiner interacts with the child. The lenses were removed from the glasses to provide the child an unobstructed view of the examiner's eyes. Our prior research indicates that the presence of the glasses does not affect the gaze behavior of children [12].



Fig. 6. **Sample child faces in our dataset.** Green box indicates ground truth eye contact, red indicates no eye contact. Notice the diversity and richness in the children’s facial expressions and head pose during natural social interactions.

Research assistants annotated the pivothead videos from each dataset using one of two video-annotation software tools: ELAN⁴ and INTERACT Mangold, 2017.⁵ Ground truth coding involved flagging the frame-level onset and offset of each instance of the child making eye contact with the examiner, as captured by looks into the camera. Kappas for frame-level agreement between pairwise comparisons of the 6 coders ranged from .89 to .94.

All studies were approved by the Institutional Review Boards of the respective institutions. Caregivers of the children participating in the studies provided written consent for their child’s participation, video recording, and data sharing.

4.1 Dataset 1

Twenty-eight children with a diagnosis of Autism Spectrum Disorder (ASD) (3 females) between the ages of 5 years and 13.7 years (mean age = 7.2 years) were recruited through the Center for Autism and the Developing Brain (CADB) in White Plains, NY. The sample was 67.9% Caucasian, 17.9% mixed ethnicity, 3.6% Hispanic, and 3.6% Asian (2 missing). All participants completed the play-based assessment during a single visit; a subset (n = 15) participated in a second assessment, eight weeks from their initial visit. Diagnosis of ASD was confirmed prior to participation by a licensed clinical psychologist at CADB. A best estimate diagnosis was based upon information collected from the Autism Diagnostic Observation Schedule (ADOS) [34] and the Autism Diagnostic Interview-Revised (ADI-R) [45]. All participants were recorded in a modified version of the Brief Observation of Social Communication Change (**v-BOSCC**) [19], a 12-minute examiner-child interaction that consisted of two 5-minute play segments with standardized sets of toys, separated by a 2-minute conversation segment. The examiner and participant were seated across from each other (face-to-face) at a small table. During the play, the child was given the option to choose a single toy from a box that contained a standardized set of toys. The set of toys available in the box differed between the first play segment and the second play segment. Within each play segment, the child was free to chose a new toy if they no longer wanted to play with the toy originally chosen. The child was free to play with each toy as he or she saw fit; the examiner joined the child’s play but did not guide it, and maintained an amount and level of language commensurate with the child’s. The transition to the conversation segment was signaled by the examiner stating that it was time to clean up and then introducing an open-ended conversation topic (e.g., “I went to the park this weekend”). During the conversation, no toys were present on the table.

⁴<http://tla.mpi.nl/tools/tla-tools/elan/>

⁵<https://www.mangold-international.com/en/products/software/behavior-research-with-mangold-interact>

age (in months)		gender		diagnosis		ethnicity		protocol		annotation	
	# subjects		# subjects		# subjects		# subjects		# sessions	# frames	minutes
less than 20	8	male	74	TD	50	Caucasian	60	ESCS	30	2,364,773	1,314
20 ~ 29	33	female	26	ASD	50	mixed	16	R-ABC	34		
30 ~ 39	15					African American	10	v-BOSCC	43		
40 ~ 49	7					Asian	6	nv-BOSCC	26		
50 ~ 59	4					Hispanic	4	Marcus	23		
60 ~ 69	14					unknown	4				
70 ~ 79	5										
80 ~ 89	5										
greater than 90	9										

Table 1. Dataset summary.

4.2 Dataset 2

Two groups of participants were recruited for this study: a sample of typically developing children between the ages of 18 and 36 months recruited at Georgia Tech, and a sample of 3-6 year old children with a diagnosis of ASD recruited at CADB. The sample included in the current analysis consisted of 19 children with ASD (6 females) and 16 typically developing children (7 females). The ASD sample ranged in age from 23 to 60 months (mean = 45 months), and was 63% Caucasian, 15.7% Asian, 10.5% Hispanic, and 5.2% mixed ethnicity (1 missing). The TD sample ranged in age from 20 to 36 months (mean = 28.3 months), and was 81.3% Caucasian, 12.5% mixed ethnicity, and 6.3% African American.

Children participated in two assessments, described in more detail below: the nonverbal version of the Brief Observation of Social Communication Change (**nv-BOSCC**; 13 ASD, 13 TD) and the Early Social Communication Scales (**ESCS**; 13 ASD, 14 TD). 10 ASD and 11 TD children contributed data to both assessments.

The ESCS [37] is a 15-25 minute structured assessment that uses standardized toys, examiner-initiated prompts (e.g., points to pictures in a book and posters on the wall) and contextual presses (e.g., wind-up toys activated out of reach of the child) to elicit gaze shifts relevant to different communication functions (sharing attention, requesting, maintaining social interaction). The examiner presents the toys to the child, one at a time, first activating the toy (e.g., blowing up balloon and then slowly letting the air out to make a squeaking noise; pressing on a trapeze toy to make a monkey swing around; winding up small toys that move across the table) and then handing the toy to the child before retrieving the toy and activating it two more times. This assessment has been shown to reliably elicit shifts of attention from objects to the examiner’s eyes in typically developing children and children with autism [44, 49].

As described above for dataset 1, the BOSCC is a naturalistic examiner-child play interaction that involves two play segments involving toys, separated by a segment of non-object-mediated social interaction. Due to the younger age of the children participating in this study, the nonverbal version of the BOSCC was used [18]. This version follows the same procedures as described above in dataset 1, but utilizes more age-appropriate materials and activities. Hence, the boxes include toys selected based on their appropriateness for younger children, and the length of play per box is reduced to 4 minutes. The conversation segment is replaced with a 2-minute snack in which the examiner offers the child a choice of crackers and cookies and then engages the child socially by commenting on the activity (e.g., “Goldfish crackers, yum yum!”). During the snack, no toys are present on the table. The transition from the first toy play segment to the snack segment is signaled by the examiner stating it was time to clean up for snack, and the transition from the snack to the second toy play segment is signaled by the examiner stating “Let’s play with some new toys.”

4.3 Dataset 3

Typically developing children between 15 and 36 months of age were recruited from the community at large in metropolitan Atlanta via fliers posted at daycare centers and a mailing to parents of children in the appropriate age range using a mailing list purchased from Experian. The sample included in the current analysis consists of 34 children (9 females) who ranged from 16 to 34 months (mean = 23.9 months). The sample was 53% Caucasian, 23.5% African American, 17.6% mixed ethnicity, and 5.9% Asian. Twelve children in the current analysis were also included in the analyses we reported in [58].

Children participated in the Rapid-ABC assessment (**R-ABC**) [39], a 2-4 minute structured interaction led by the examiner that consists of five activities: (1) examiner greets the child; (2) examiner presents a ball to the child and then initiates a turn-taking game by rolling the ball back and forth across the table; (3) examiner presents a book to the child and then reads it out loud while encouraging the child to turn the pages and to point to pictures; (4) examiner surprises the child by placing the book on her head as if it were a hat and waits for the child's reaction; and (5) the examiner initiates a gentle tickle game with the child. More details about this assessment can be found in Rehg et al [42].

4.4 Dataset 4

As part of a pilot study to trial the use of the wearable camera in a clinical setting, data was collected from three boys with a diagnosis of autism between the ages of 3 and 6 years. One boy was Caucasian, one was African-American, and the ethnicity of the third is unknown. Participants were recruited from the Language & Learning Clinic (LLC) at the Marcus Autism Center, where all were receiving intervention services at the time. Data collection involved sessions intended to resemble treatment sessions at the clinic, including sessions where the therapist engaged the child in toy play while placing no demands on the child to respond (Pairing, 12 sessions), sessions in which the therapist placed demands on the child, (Demands, 5 sessions), and sessions where the therapist periodically withheld toys from the child in an effort to get the child to request the toy (Mands; 6 sessions). We denote these plays as **Marcus** in this paper. All but two sessions were 10 minutes in length; the remaining two were 3 and 9 minutes in length, respectively.

4.5 Dataset Summary

In summary, when all of the four datasets are considered together, there are 100 unique children in 156 different play sessions (some participants completed two sessions). Half of the children were diagnosed with ASD and 74% of the participants were boys. The total number annotated frames exceeds 2 million, corresponding to approximately 22 hours of video. The complete statistics for this dataset are given in Table 1.

5 EXPERIMENTS AND RESULTS

In this section, we provide experimental evaluations for each of the components of our method.

5.1 Child Face Detection Results

In order to evaluate the accuracy of face detection, we performed additional annotations in a sample of frames. We annotated child faces with bounding boxes for a total of 5 minutes of video containing a single child participant, and 5 minutes of video containing a child and a caregiver together, corresponding to approximately 18,000 frames in total. Using these annotations as ground truth, we compared the accuracy of the Faster R-CNN face detector from Sec 3.2 to the OMRON OKAO face detector [54] which was used in our previous experiments in [58]. We obtained detections for every frame using both methods, and filtered spurious false positives and adult faces for both methods using the online selection method from Sec 3.3. The precision-recall results at Intersection over Union (IoU) thresholds of 0.5 and 0.75 are given in Table 2. IoU score is defined as overlapping ratio between

IoU threshold	0.5		0.75	
	Recall	Precision	Recall	Precision
OMRON	0.437	0.941	0.015	0.212
Faster R-CNN	0.947	0.963	0.828	0.870

Table 2. Precision and recall comparison at two IoU thresholds between the OMRON OKAO face detector and Faster-RCNN trained on the WIDER FACE dataset. Testing set consists of approximately 18k child face examples.




Fig. 7. Landmark detection failures due to occlusions or extreme head pose.

the predicted and ground truth bounding boxes. See Sec. 4.2 in [13] for a description of IoU score and its use in evaluating object detectors. The results show a high recall discrepancy between the two methods, with Faster R-CNN delivering substantially greater recall at a comparable precision for IoU 0.5. High recall is a critical property for our application, in order to not miss brief moments of eye contact.

5.2 Landmark Detection and Pose Estimation Results

To evaluate the frequency of successful landmark detection (and by extension, head pose estimation) in our dataset, we calculated the percentage of frames labeled as eye contact for which landmarks were detected, shown in Table 3. Note that this is not a statement about the accuracy of the landmarks, only about their availability. This is meaningful because state-of-the-art methods such as IntraFace will only output landmarks if they are of sufficiently high quality. The percentage of frames with missing landmarks is the portion of the error rate for eye contact detection which is attributable to landmark detection failure. As can be seen in Table 3, landmark detection is successful in only 75.56% of frames, in comparison to 97.94% for face detection. This suggests that around a quarter of the errors are due to landmark detection failure, and a negligible amount are due to face detection failure. These findings explain the low recall rate of the two eye contact methods (PEEC and GazeLocking) that rely on landmark detection. Fig 7 shows some examples of situations of occlusion and extreme rotation that result in landmark detection failures.

5.3 Eye Contact Detection Results

We now describe our experimental evaluation of four eye contact detection methods, PiCNN, PEEC, Modified AlexNet, and GazeLocking, using the entire dataset from Sec 4.

5.3.1 Experiment Design. We divide the total dataset summarized in Table 1 into 5 disjoint train/test splits for 5-fold cross validation, where each split has 80% of training and 20% of testing sessions and subjects included in the training set are not present in the testing set. This ensures that testing is always done on unseen participants.

	ASD	TD	ESCS	v-BOSCC	nv-BOSCC	R-ABC	Marcus	All
total frames	118,464	62,220	29,305	74,867	38,891	19,718	17,903	180,684
face detected (%), mean	98.21	97.52	97.45	98.90	95.34	98.30	99.32	97.94
face detected (%), std	0.03	0.03	0.03	0.02	0.05	0.02	0.01	0.03
landmark detected (%), mean	73.65	70.93	81.23	80.06	80.23	59.99	55.77	75.56
landmark detected (%), std	0.19	0.22	0.14	0.13	0.16	0.24	0.21	0.21

Table 3. Total number of frames within the moments that are marked as eye contact by human annotators, and the success rate (mean and standard deviation) of face detection and landmark detection, grouped by different conditions.

Each training set split is sampled uniformly across combinations of diagnosis (TD/ASD) and play protocols. For example, we sample 80% for a training set from the TD ESCS group, from the ASD ESCS group, from the TD nv-BOSCC group, from the ASD nv-BOSCC group, and so on, such that different conditions are fairly represented across the five splits. We use the same five folds to train, test and compare the four eye contact detection models described in Sec 3.1 – our PiCNN model, PEEC [58], Modified AlexNet [31], and Gaze Locking [50]. With this cross validation approach, we are able to obtain testing results on all of our datasets, and we report the overall performance averaged over all sessions as well as within different populations. Table 4 summarizes the results.

On average, the training set in each split initially had 145k positive (eye contact) frames and 1,746k negative (no eye contact) frames, which is highly imbalanced. To overcome the data skewness issue, we resampled the training sets with horizontal flip, slight rotation and color jittering with positive set oversampling and negative set subsampling to make it more balanced at positive:negative = 561k:842k = 4:6 ratio.

For the two deep models (PiCNN and AlexNet), we used the Caffe deep learning framework [25] on an Nvidia Titan X 12GB GPU, with batch size 128, learning rate 0.005 until the 100,000th iteration and 0.0005 until the final 200,000th iteration. For the other two models (PEEC and Gaze Locking), we used MATLAB on an Intel Core i7-5930K CPU.

Note that our analysis does not include direct performance comparisons to other appearance-based gaze tracking methods such as [30, 60]. The primary reason is that these methods are designed for a different task, accurately mapping the participant’s point of gaze on a mobile screen. These methods can identify *where* a participant is looking, but they only know *what* the user is looking at if they have access to ROI masks for the gaze stimulus. In contrast, our method automatically identifies *what* the participant is looking at, but only for the specific gaze event of eye contact. Moreover, the domain of screen viewing is quite different from naturalistic social interactions. This can be seen by comparing Figs 2 and 6. In particular, [30, 60] rely on landmark localization [2] to extract the eye region from the face. Based on our results from Sec 5.2, these methods will take an immediate 22.4% miss in recall (see Table 3) due to the difficulty of detecting landmarks under challenging conditions. These methods are therefore unlikely to be competitive with PiCNN for our task.

5.3.2 Experimental Results. All four eye contact detection methods output a confidence score between 0 and 1 at each frame, with a higher score indicating increased likelihood of eye contact. We use these per-frame scores to evaluate how well the model is detecting eye contact. Since our dataset has more than 92% negative samples, accuracy ($\frac{tp+tn}{p+n}$) is not a good measure of performance. For example, simply predicting everything as negative will give 92% accuracy but that is not a useful detector. Instead, we report the detector’s performance with respect to three metrics, namely F_1 , Matthews correlation coefficient (MCC), and the Area Under Curve of Precision Recall curve (AUC-PR).

The F_1 score is defined as

$$F_1 = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} \quad (1)$$




Fig. 8. Average Precision-Recall curve.




Fig. 9. Effect of training set size.

It can be interpreted as a weighted average of the precision and recall, where 1 is the best value and 0 is the worst. In our analysis, we report our overall precision and recall at the threshold value that maximizes the F_1 score.

MCC is defined as

$$MCC = \frac{tp \cdot tn - fp \cdot fn}{\sqrt{(tp + fp)(tp + fn)(tn + fp)(tn + fn)}} \quad (2)$$

An MCC of 1 represents a perfect prediction, 0 is no better than random prediction and -1 indicates total disagreement between prediction and ground truth.

In the equations above, $precision = \frac{tp}{tp+fp}$, $recall = \frac{tp}{tp+fn}$, tp is the number true positive samples, fp is the number of false positive samples, tn is the number true negative samples, and fn is the number of false negative samples.

Both F_1 and MCC are widely used in machine learning as a measure of the quality of binary classifiers. Because the prediction output of the models we used in our analysis is a real-numbered value instead of a hard binary output, we calculate the F_1 and MCC scores at every cutoff points and report the maximum. Additionally, we also compute the Area Under Curve of Precision Recall curve (AUC-PR). Typically, AUC is computed using Receiver Operating Characteristic (ROC) instead of Precision and Recall curve (PR curve), but AUC of ROC is not sensitive to the uneven class sizes, thus we use the PR curve. Like F_1 score, AUC-PR is maximum at 1 and minimum at 0, but it is an aggregated measure across all prediction cutoffs. Our final results by these criteria are summarized in Table 4. Clearly, our PiCNN method outperforms all other methods when evaluated on all datasets as well as on individual groups with different conditions. When the results are broken down into different diagnostic groups, play protocols, gender and ethnicity, in all cases and in all metrics, our PiCNN method achieves the best performance, followed by AlexNet, PEEC and GazeLocking in this order. The standard deviation of the scores is greatest under different play protocol settings ($F_1=0.0465$, $MCC=0.0492$, $AUC-PR=0.065$), and second greatest under different ethnic groups ($F_1=0.034$, $MCC=0.035$, $AUC-PR=0.0336$).

If smaller number of training sessions were used in a train-test split, precision tends to increase proportionally to training set size both in PEEC and PiCNN with PiCNN always performing better. However, PEEC recall is higher with smaller training set. This difference in recall is reversed after around 30 sessions and becomes more evident as more training samples are used (Fig 9). If one wants to train PiCNN model from scratch, they would need to use at least 30 sessions (100K annotated frames) to take advantage of it. If they want to include their own

dataset in our framework, it is more advisable to fine-tune our pre-trained model that already has learned from more than 128 sessions.

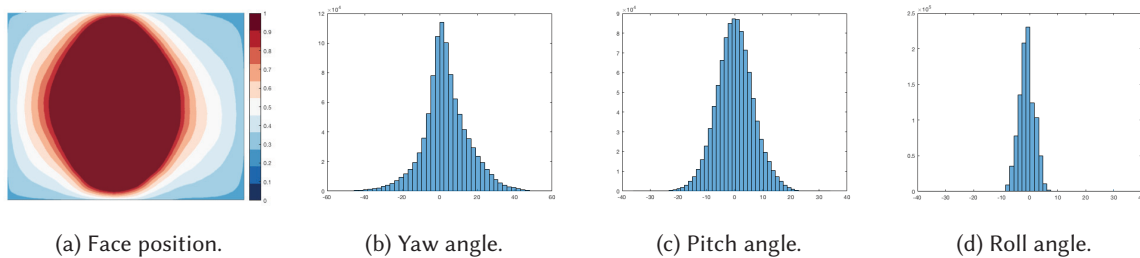


Fig. 10. **Our dataset statistics.** (a) distribution of child's face location in the image shows that captured faces making eye contact often appear around image center with bigger mass vertically. (b)–(d) estimated head pose distribution in degrees.

Fig 10 illustrates the ranges of face that our dataset covers. Although these are not a perfect representation of the data distribution because position is from the face detector's measure and pose is from the landmark detector's output, it gives an idea about the rough scope of what the current training set deals with. In case others want to apply the model to a very different setup (e.g., subject and examiner are physically configured substantially different), it might be better to do fine-tune with their own data rather than using the model as-is. Such process could make the model additionally learn unfamiliar facial features associated with eye contact, but slight variation in optical parameters (e.g., noise, exposure, motion blur, etc) would not likely need the same process because of the data augmentation step during training.

5.3.3 Deep Visualization. Lastly, we visualize some of the weights of our PiCNN model to qualitatively observe what the model has learned. In Fig 11, the child face image (a) is an example input, (b) shows the 96 filters learned in the first convolutional layer, and (c)–(e) are the activation maps from the first three convolutional layers when the input image is convolved with their corresponding filters. Although is not entirely clear how we can assess the usefulness of the learned weights, it is recognized that the networks have learned low-level color and round edge filters that seem useful for facial feature extraction. As shown in the activation maps, high response is produced on the regions that are valuable for eye contact detection such as the eyes and face contour.

		F_1	MCC	AUC-PR	Precision	Recall
All	PiCNN (Ours)	0.78	0.77	0.79	0.75	0.80
	AlexNet [31]	0.73	0.72	0.75	0.71	0.77
	PEEC [58]	0.63	0.62	0.57	0.67	0.59
	GazeLocking [50]	0.52	0.50	0.48	0.52	0.52
ASD	PiCNN (Ours)	0.76	0.75	0.78	0.75	0.78
	AlexNet [31]	0.72	0.71	0.74	0.69	0.74
	PEEC [58]	0.64	0.64	0.60	0.68	0.61
	GazeLocking [50]	0.51	0.49	0.49	0.50	0.52
TD	PiCNN (Ours)	0.79	0.78	0.78	0.74	0.83
	AlexNet [31]	0.75	0.74	0.74	0.69	0.81
	PEEC [58]	0.63	0.62	0.55	0.65	0.62
	GazeLocking [50]	0.54	0.53	0.49	0.54	0.54
ESCS	PiCNN (Ours)	0.75	0.75	0.76	0.69	0.82
	AlexNet [31]	0.69	0.69	0.72	0.62	0.78
	PEEC [58]	0.57	0.56	0.47	0.55	0.60
	GazeLocking [50]	0.48	0.46	0.41	0.46	0.50
v-BOSCC	PiCNN (Ours)	0.76	0.74	0.79	0.75	0.77
	AlexNet [31]	0.73	0.71	0.75	0.72	0.74
	PEEC [58]	0.69	0.67	0.67	0.72	0.65
	GazeLocking [50]	0.56	0.53	0.56	0.53	0.60
nv-BOSCC	PiCNN (Ours)	0.82	0.81	0.83	0.82	0.81
	AlexNet [31]	0.77	0.76	0.78	0.76	0.78
	PEEC [58]	0.71	0.71	0.68	0.76	0.67
	GazeLocking [50]	0.58	0.56	0.55	0.57	0.58
R-ABC	PiCNN (Ours)	0.77	0.77	0.71	0.72	0.84
	AlexNet [31]	0.72	0.72	0.68	0.66	0.80
	PEEC [58]	0.59	0.59	0.49	0.66	0.54
	GazeLocking [50]	0.52	0.51	0.43	0.53	0.51
Marcus	PiCNN (Ours)	0.86	0.86	0.88	0.86	0.87
	AlexNet [31]	0.80	0.79	0.83	0.77	0.83
	PEEC [58]	0.54	0.54	0.49	0.67	0.45
	GazeLocking [50]	0.46	0.44	0.41	0.49	0.44
male	PiCNN (Ours)	0.77	0.75	0.77	0.74	0.79
	AlexNet [31]	0.72	0.71	0.73	0.69	0.76
	PEEC [58]	0.64	0.63	0.57	0.66	0.61
	GazeLocking [50]	0.52	0.50	0.48	0.51	0.53
female	PiCNN (Ours)	0.78	0.78	0.80	0.76	0.81
	AlexNet [31]	0.74	0.74	0.76	0.70	0.80
	PEEC [58]	0.64	0.63	0.59	0.68	0.60
	GazeLocking [50]	0.53	0.52	0.49	0.54	0.52
Caucasian	PiCNN (Ours)	0.76	0.75	0.76	0.73	0.79
	AlexNet [31]	0.72	0.71	0.72	0.68	0.76
	PEEC [58]	0.64	0.63	0.57	0.66	0.62
	GazeLocking [50]	0.51	0.50	0.47	0.51	0.52
mixed	PiCNN (Ours)	0.76	0.75	0.81	0.74	0.79
	AlexNet [31]	0.75	0.74	0.76	0.72	0.79
	PEEC [58]	0.69	0.67	0.66	0.70	0.67
	GazeLocking [50]	0.60	0.58	0.59	0.57	0.63
African American	PiCNN (Ours)	0.82	0.82	0.83	0.81	0.84
	AlexNet [31]	0.79	0.79	0.79	0.76	0.83
	PEEC [58]	0.62	0.61	0.57	0.68	0.56
	GazeLocking [50]	0.51	0.50	0.48	0.51	0.52
Asian	PiCNN (Ours)	0.79	0.78	0.75	0.72	0.86
	AlexNet [31]	0.67	0.66	0.63	0.61	0.74
	PEEC [58]	0.65	0.65	0.58	0.66	0.65
	GazeLocking [50]	0.49	0.49	0.41	0.43	0.57
Hispanic	PiCNN (Ours)	0.73	0.73	0.78	0.68	0.79
	AlexNet [31]	0.70	0.70	0.72	0.66	0.75
	PEEC [58]	0.70	0.69	0.67	0.69	0.71
	GazeLocking [50]	0.58	0.57	0.58	0.55	0.61

Table 4. Result scores in five metrics, grouped by different conditions.




Fig. 11. **Visualization of PiCNN.** (a) example input, (b) conv1 filters, (c)–(e) response map of the first 3 layers.

6 CONCLUSION

In this work, we have proposed a novel approach that measures eye contact in naturalistic social interactions. We have presented a fully automated deep learning system that detects moments of eye contact from egocentric videos of adult-child interactions while also implicitly estimating the head pose. We performed a thorough and systematic evaluation of our methods on the largest ever dataset of naturalistic social interactions comprising of 22 hours of play session videos of 100 individual children. We found that our results are a significant improvement over other existing methods.

Our approach will be instrumental to understanding atypical gaze behavior in natural social settings, as well as improving the diagnosis, screening and treatment of Autism Spectrum Disorder. We also expect our work to have applications in other domains such as in the automatic analysis of turn-taking and social roles in group interactions or in the development of models of social intelligence for robots, which could allow them to interact naturally with humans based on eye contact.

ACKNOWLEDGMENTS

The authors would like to thank Yin Li for helpful discussions, and Marcus Autism Center for sharing their dataset with the authors. This study was funded in part by the Simons Foundation under grants 336363 and 383667, as well as the National Science Foundation under grant IIS-1029679.

REFERENCES

- [1] Michael Argyle and Janet Dean. 1965. Eye-contact, distance and affiliation. *Sociometry* (1965), 289–304.
- [2] Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. 2014. Continuous conditional neural fields for structured regression. In *European Conference on Computer Vision*. Springer, 593–608.
- [3] T Berry Brazelton, Edward Tronick, Lauren Adamson, Heidelise Als, and Susan Wise. 1975. Early mother-infant reciprocity. *Parent-Infant Interaction* 3 (1975), 137.
- [4] Leo Breiman. 2001. Random forests. *Machine Learning* 45, 1 (2001), 5–32.
- [5] Chih-Chung Chang and Chih-Jen Lin. 2011. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)* 2, 3 (2011), 27.
- [6] Katarzyna Chawarska and Frederick Shic. 2009. Looking but not seeing: Atypical visual scanning and recognition of faces in 2 and 4-year-old children with autism spectrum disorder. *Journal of Autism and Developmental Disorders* 39, 12 (2009), 1663.
- [7] Meia Chita-Tegmark. 2016. Social attention in ASD: a review and meta-analysis of eye-tracking studies. *Research in Developmental Disabilities* 48 (2016), 79–93.
- [8] Amy M Daniels and David S Mandell. 2014. Explaining differences in age at autism spectrum disorder diagnosis: A critical review. *Autism* 18, 5 (2014), 583–597.
- [9] Fernando De la Torre, Wen-Sheng Chu, Xuehan Xiong, Francisco Vicente, Xiaoyu Ding, and Jeffrey Cohn. 2015. Intraface. In *Proceedings of the IEEE Conference on Automatic Face and Gesture Recognition (FG)*, Vol. 1. IEEE, 1–8.
- [10] Piotr Dollár, Peter Welinder, and Pietro Perona. 2010. Cascaded pose regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1078–1085.

- [11] Richard O Duda, Peter E Hart, and David G Stork. 2012. *Pattern classification*. John Wiley & Sons.
- [12] Sarah R Edmunds, Agata Rozga, Yin Li, Elizabeth A Karp, Lisa V Ibanez, James M Rehg, and Wendy L Stone. 2017. Brief Report: Using a Point-of-View Camera to Measure Eye Gaze in Young Children with Autism Spectrum Disorder During Naturalistic Social Interactions: A Pilot Study. *Journal of Autism and Developmental Disorders* (2017), 1–7.
- [13] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. 2010. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* 88, 2 (2010), 303–338.
- [14] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. 2010. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 9 (2010), 1627–1645.
- [15] Centers for Disease Control and Prevention. 2016. Summary of Autism Spectrum Disorder Prevalence Studies. <https://www.cdc.gov/ncbddd/autism/documents/ASDPrevalenceDataTable2016.pdf>. (2016). Accessed: 2017-05-03.
- [16] Tom Foulsham, Esther Walker, and Alan Kingstone. 2011. The where, what and when of gaze allocation in the lab and the natural environment. *Vision Research* 51, 17 (2011), 1920–1931.
- [17] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- [18] Rebecca Grzadzinski, Themba Carr, Costanza Colombi, Kelly McGuire, Sarah Dufek, Andrew Pickles, and Catherine Lord. 2016. Measuring changes in social communication behaviors: preliminary development of the Brief Observation of Social Communication Change (BOSCC). *Journal of Autism and Developmental Disorders* 46, 7 (2016), 2464–2479.
- [19] R Grzadzinski, K Martinez, G Gunin, E Ajodan, S Kim, and C Lord. 2017. Development of the Brief Observation of Social Communication Change (BOSCC) for Verbally Able Children with ASD. *Biennial Meeting of the Society for Research on Child Development (SRCD)* (2017).
- [20] Quentin Guillon, Nouchine Hadjikhani, Sophie Baduel, and Bernadette Rogé. 2014. Visual social attention in autism spectrum disorder: Insights from eye tracking studies. *Neuroscience & Biobehavioral Reviews* 42 (2014), 279–297.
- [21] Dan Witzner Hansen and Qiang Ji. 2010. In the eye of the beholder: A survey of models for eyes and gaze. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 3 (2010), 478–500.
- [22] Mariko Hosozawa, Kyoko Tanaka, Toshiaki Shimizu, Tamami Nakano, and Shigeru Kitazawa. 2012. How children with specific language impairment view social situations: an eye tracking study. *Pediatrics* 129, 6 (2012), e1453–e1460.
- [23] Ted Hutman, Mandeep K Chela, Kristen Gillespie-Lynch, and Marian Sigman. 2012. Selective visual attention at twelve months: Signs of autism in early social interactions. *Journal of Autism and Developmental Disorders* 42, 4 (2012), 487–498.
- [24] Vidit Jain and Erik G Learned-Miller. 2010. Fddb: A benchmark for face detection in unconstrained settings. *UMass Amherst Technical Report* (2010).
- [25] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*. ACM, 675–678.
- [26] Huaizu Jiang and Erik Learned-Miller. 2016. Face detection with the faster R-CNN. *arXiv preprint arXiv:1606.03473* (2016).
- [27] Warren Jones, Katelin Carr, and Ami Klin. 2008. Absence of preferential looking to the eyes of approaching adults predicts level of social disability in 2-year-old toddlers with autism spectrum disorder. *Archives of General Psychiatry* 65, 8 (2008), 946–954.
- [28] Chris L Kleinke. 1986. Gaze and eye contact: a research review. *Psychological Bulletin* 100, 1 (1986), 78.
- [29] Ami Klin, Warren Jones, Robert Schultz, Fred Volkmar, and Donald Cohen. 2002. Visual fixation patterns during viewing of naturalistic social situations as predictors of social competence in individuals with autism. *Archives of General Psychiatry* 59, 9 (2002), 809–816.
- [30] Kyle Krafka, Aditya Khosla, Petr Kellnhofer, Harini Kannan, Suchendra Bhandarkar, Wojciech Matusik, and Antonio Torralba. 2016. Eye tracking for everyone. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2176–2184.
- [31] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*. 1097–1105.
- [32] Michael Land and Benjamin Tatler. 2009. *Looking and acting: vision and eye movements in natural behaviour*. Oxford University Press.
- [33] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*. Springer, 740–755.
- [34] Catherine Lord, Pamela C DiLavore, and Katherine Gotham. 2012. *Autism diagnostic observation schedule*. Western Psychological Services Torrance, CA.
- [35] Silvia Magrelli, Patrick Jermann, Noris Basilio, François Ansermet, François Hentsch, Jaqueline Nadel, and Aude Billard. 2013. Social orienting of children with autism to facial expressions and speech: a study with a wearable eye-tracker in naturalistic settings. *Frontiers in Psychology* 4 (2013), 840.
- [36] P Mundy and C Françoise Acra. 2006. Joint attention, social engagement, and the development of social competence. *The Development of Social Engagement: Neurobiological Perspectives* (2006), 81–117.
- [37] Peter Mundy, Christine Delgado, Jessica Block, Meg Venezia, Anne Hogan, and Jeffrey Seibert. 2003. Early social communication scales (ESCS). *Coral Gables, FL: University of Miami* (2003).
- [38] Basilio Noris, Jacqueline Nadel, Mandy Barker, Nouchine Hadjikhani, and Aude Billard. 2012. Investigating gaze of children with ASD in naturalistic settings. *PLoS One* 7, 9 (2012), e44144.

- [39] Opal Ousley, Rosa Arriaga, Michael Morrier, Jennifer Mathys, Monica Allen, and Gregory Abowd. 2013. *Beyond parental report: findings from the rapid-abc, a new 4-minute interactive autism*. Technical Report. Technical report, Georgia Institute of Technology.
- [40] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. 2015. Deep Face Recognition.. In *British Machine Vision Conference*, Vol. 1. 6.
- [41] Karen Pierce, David Conant, Roxana Hazin, Richard Stoner, and Jamie Desmond. 2011. Preference for geometric patterns early in life as a risk factor for autism. *Archives of General Psychiatry* 68, 1 (2011), 101–109.
- [42] James Rehg, Gregory Abowd, Agata Rozga, Mario Romero, Mark Clements, Stan Sclaroff, Irfan Essa, O Ousley, Yin Li, Chanho Kim, and others. 2013. Decoding children’s social behavior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3414–3421.
- [43] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*. 91–99.
- [44] Agata Rozga, Ted Hutman, Gregory S Young, Sally J Rogers, Sally Ozonoff, Mirella Dapretto, and Marian Sigman. 2011. Behavioral profiles of affected and unaffected siblings of children with autism: Contribution of measures of mother–infant interaction and nonverbal communication. *Journal of Autism and Developmental Disorders* 41, 3 (2011), 287–301.
- [45] Michael Rutter, A Le Couteur, and C Lord. 2003. Autism diagnostic interview-revised. *Los Angeles, CA: Western Psychological Services* 29 (2003), 30.
- [46] Noah J Sasson and Jed T Elison. 2012. Eye tracking young children with autism. *Journal of Visualized Experiments* 61 (2012), e3675–e3675.
- [47] Jeffrey S Shell, Roel Vertegaal, Daniel Cheng, Alexander W Skaburskis, Changuk Sohn, A James Stewart, Omar Aoudeh, and Connor Dickie. 2004. ECSSGlasses and EyePliances: using attention to open sociable windows of interaction. In *Proceedings of the 2004 Symposium on Eye Tracking Research & Applications (ETRA)*. ACM, 93–100.
- [48] Marian Sigman. 1998. The Emanuel Miller Memorial Lecture 1997: Change and continuity in the development of children with autism. *Journal of Child Psychology and Psychiatry* 39, 6 (1998), 817–827.
- [49] Marian Sigman, Peter Mundy, Tracy Sherman, and Judy Ungerer. 1986. Social interactions of autistic, mentally retarded and normal children and their caregivers. *Journal of Child Psychology and Psychiatry* 27, 5 (1986), 647–656.
- [50] Brian A Smith, Qi Yin, Steven K Feiner, and Shree K Nayar. 2013. Gaze locking: passive eye contact detection for human-object interaction. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST)*. ACM, 271–280.
- [51] Yusuke Sugano, Yasuyuki Matsushita, and Yoichi Sato. 2014. Learning-by-synthesis for appearance-based 3d gaze estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1821–1828.
- [52] Matthew A Turk and Alex P Pentland. 1991. Face recognition using eigenfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 586–591.
- [53] Paul Viola and Michael J Jones. 2004. Robust real-time face detection. *International Journal of Computer Vision* 57, 2 (2004), 137–154.
- [54] OMRON OKAO vision. 2017. <https://www.omron.com/ecb/products/mobile/okao01.html>. (2017). Accessed: 2017-05-03.
- [55] Shuo Yang, Ping Luo, Chen-Change Loy, and Xiaoou Tang. 2015. From facial parts responses to face detection: A deep learning approach. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 3676–3684.
- [56] Shuo Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. 2016. WIDER FACE: A Face Detection Benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [57] Zhefan Ye, Yin Li, Alireza Fathi, Yi Han, Agata Rozga, Gregory D Abowd, and James M Rehg. 2012. Detecting eye contact using wearable eye-tracking glasses. In *Proceedings of the ACM Conference on Ubiquitous Computing*. ACM, 699–704.
- [58] Zhefan Ye, Yin Li, Yun Liu, Chanel Bridges, Agata Rozga, and James M Rehg. 2015. Detecting bids for eye contact using a wearable camera. In *Proceedings of the IEEE Conference on Automatic Face and Gesture Recognition (FG)*, Vol. 1. IEEE, 1–8.
- [59] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. 2016. Joint Face Detection and Alignment Using Multitask Cascaded Convolutional Networks. *IEEE Signal Processing Letters* 23, 10 (2016), 1499–1503.
- [60] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. 2015. Appearance-based gaze estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4511–4520.
- [61] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. 2016. It’s Written All Over Your Face: Full-Face Appearance-Based Gaze Estimation. *arXiv preprint arXiv:1611.08860* (2016).
- [62] Lonnie Zwaigenbaum, Susan Bryson, and Nancy Garon. 2013. Early identification of autism spectrum disorders. *Behavioral Brain Research* 251 (2013), 133–146.

Received May 2017; accepted June 2017