Journal of Experimental Child Psychology 150 (2016) 165–179

Contents lists available at ScienceDirect

Journal of Experimental Child Psychology

journal homepage: www.elsevier.com/locate/jecp

Automated decoding of facial expressions reveals marked differences in children when telling antisocial versus prosocial lies

Sarah Zanette [a], Xiaoqing Gao [a], Megan Brunet [a], Marian Stewart Bartlett [b], Kang Lee [a,*]

[a] University of Toronto, Toronto, Ontario M5R 2X2, Canada
[b] University of California, San Diego, La Jolla, CA 92093, USA

ARTICLE INFO

Article history:
Received 17 August 2015
Revised 11 May 2016
Available online 16 June 2016

Keywords:
Antisocial lying
Prosocial lying
Nonverbal behavior
Facial expressions
Machine learning
Emotions

ABSTRACT

The current study used computer vision technology to examine the nonverbal facial expressions of children (6–11 years old) telling antisocial and prosocial lies. Children in the antisocial lying group completed a temptation resistance paradigm where they were asked not to peek at a gift being wrapped for them. All children peeked at the gift and subsequently lied about their behavior. Children in the prosocial lying group were given an undesirable gift and asked if they liked it. All children lied about liking the gift. Nonverbal behavior was analyzed using the Computer Expression Recognition Toolbox (CERT), which employs the Facial Action Coding System (FACS), to automatically code children's facial expressions while lying. Using CERT, children's facial expressions during antisocial and prosocial lying were accurately and reliably differentiated significantly above chance-level accuracy. The basic expressions of emotion that distinguished antisocial lies from prosocial lies were joy and contempt. Children expressed joy more in prosocial lying than in antisocial lying. Girls showed more joy and less contempt compared with boys when they told prosocial lies. Boys showed more contempt when they told prosocial lies than when they told antisocial lies. The key action units (AUs) that differentiate children's antisocial and prosocial lies are blink/eye closure, lip pucker, and lip raise on the right side. Together, these findings indicate that children's facial expressions differ while telling antisocial versus prosocial lies. The reliability of CERT in

* Corresponding author.
E-mail addresses: kang.lee@utoronto.ca, kangleetoronto@gmail.com (K. Lee).

http://dx.doi.org/10.1016/j.jecp.2016.05.007

detecting such differences in facial expression suggests the viability of using computer vision technology in deception research.

Introduction

Lying is a typical developmental phenomenon in children (Ahern, Lyon, & Quas, 2011; Darwin, 1877; Hartshorne & May, 1928; Lewis, Stanger, & Sullivan, 1989; Talwar & Crossman, 2011; Talwar & Lee, 2008; see Lee, 2013, for a review). Children tell lies as early as 2-years-old (Evans & Lee, 2013) and their ability to tell convincing lies improves with age. Whereas 2- to 5-year-olds often stumble during lie-telling and reveal their dishonest behavior, they often become sophisticated lie-tellers by 6 years of age (see Lee, 2013, for a review). Just like adults, children also tell different kinds of lies. Antisocial lies are lies told to conceal a misdeed (e.g., denying cheating in a game; Bussey, 1999; Fu, Evans, Xu, & Lee, 2012; Polak & Harris, 1999; Talwar & Lee, 2008). In contrast, prosocial lies, or "white lies," are lies told to spare another's feelings (e.g., claiming to like a disappointing gift; Bussey, 1999; Saarni, 1989; Talwar & Crossman, 2011; Talwar, Murphy, & Lee, 2007). Because antisocial and prosocial lies serve different interpersonal functions, children must learn how to tell these different types of lies according to different social contexts.

Nonverbal behavior during deception

Lying is a multidimensional act. Telling a lie involves expressing a verbal statement that is inconsistent with the truth. To be convincing, the lie-teller must also convey an emotion that matches the social context of the lie being told. This is because different social contexts call for different display rules about how people should express their overt emotions (Bussey, 1999; Davis, 1995; McDowell, O'Neil, & Parke, 2000; Talwar, Crossman, Williams, & Muir, 2011; Tobin & Graziano, 2011). Thus, when telling antisocial or prosocial lies in a social context, children must learn to display the appropriate emotions consistent with the display rules for this social context. Just as adults do, children must adjust their nonverbal behaviors while lying to match their verbal statements told in antisocial or prosocial lying contexts. In particular, the display rules for telling an antisocial lie are different from the display rules for telling a prosocial lie. For example, a child who uses an antisocial lie to cover up his or her own transgression must conceal any facial expressions that could convey guilt for committing the misdeed and must instead simulate expressions of innocence (Lewis et al., 1989; Polak & Harris, 1999; Talwar & Lee, 2002a; Talwar et al., 2011). Conversely, a child who uses a prosocial lie to falsely claim to like a disappointing gift must conceal his or her feelings of disappointment and instead express gratitude toward the gift-giver (Saarni, 1989; Talwar, Murphy, et al., 2007; Talwar et al., 2011).

Considering that each type of lie requires the lie-teller to convey or suppress different emotions, it is likely that the facial expressions of children telling an antisocial lie differ from those of children telling a prosocial lie, but this has not been empirically tested. Prior research has only examined children's nonverbal behaviors associated with antisocial and prosocial lying separately. Thus, it is currently unknown whether children aged 6 to 11 years have developed the ability to differentiate between the display rules of lies told in antisocial contexts from lies told in prosocial contexts. Examining this issue will provide important insights about children's developing ability to manipulate their nonverbal behaviors according to the demands of different social contexts.

Prior researchers have identified that several nonverbal behaviors occur when children tell lies. In an examination of antisocial lies, Lewis and colleagues (1989) recorded 3-year-olds lying about peeking at a toy they were asked not to look at. Researchers then reviewed recorded videos of the session in normal and slow motion to analyze children's facial expressions. They found that children who lied about peeking smiled more often than children who did not peek at the toy and, therefore, did not lie (Lewis et al., 1989). Similarly, Talwar and Lee (2002a) found that children aged 3 to 7 years old were more likely to exhibit a big smile and less likely to display a relaxed mouth while telling an antisocial lie compared with non-liars. In an examination of prosocial lies, Talwar, Murphy, et al. (2007) gave 3- to 11-year-olds an undesirable gift (a bar of soap) and asked whether they liked it. Seventy-seven percent of the children told a prosocial lie and claimed to like the bar of soap. The children who lied about their liking of the undesirable gift were more likely to smile compared to a control group of children who received a desirable gift (a commercial toy) and, thus, were telling the truth when they claimed to like the gift. Older children also smiled more than younger children (Talwar, Murphy, et al., 2007). Interestingly, this study also found gender differences, whereby girls showed more positive expressions and smiled more than boys. Although these existing studies offer important insights to the facial expressions of children during lie-telling, no study has concurrently examined whether the facial expressions associated with children's antisocial lying differ from the facial expressions associated with their prosocial lying.

Automated analysis of facial expressions

Existing studies analyzing children's nonverbal behaviors during lie-telling have relied on trained human coders. This process involves playing recorded videos of children lying at a slow pace several times to reliably and accurately code facial expressions. Unsurprisingly, manually scoring videos is extremely time-consuming, costly, and labor-intensive. To improve efficiency, many researchers focus on coding only a small number of nonverbal behaviors, such as smiling and blinking, potentially overlooking other key emotional expressions as a result. More important, because nonverbal behaviors are fast and fleeting, human coders might not be sensitive enough to detect all of the minute changes in facial expressions. To overcome the limitations of human coders, the current study used a novel computer vision technology, the Computer Expression Recognition Toolbox (CERT), to analyze children's facial expressions while telling antisocial and prosocial lies.

CERT, developed by Bartlett and colleagues, is a fully automatic and real-time software tool for coding facial expressions (Bartlett et al., 2006; Donato, Bartlett, Hager, Ekman, & Sejnowski, 1999; see Littlewort et al., 2011, for an overview). Using video files, live video, or individual images, CERT automatically and reliably codes 20 facial muscle movements based on the Facial Action Coding System (FACS; Ekman, Friesen, & Hager, 2002) and the following emotions: happiness (joy), sadness, anger, disgust, surprise, fear, and contempt as well as neutral expressions (Littlewort et al., 2011).

FACS is traditionally used as a hand-scoring method of coding facial muscle movements. Ekman and Friesen (1978) referred to the individual facial muscle movements as "action units" (AUs) and numerically coded each AU. For example, blinking is referred to as AU45 and lip raise is referred to as AU10. FACS has been widely used across the behavioral sciences to code facial expressions because it has been proven to be highly useful for the study of the facial movements associated with cognitive and affective states (Ekman & Rosenberg, 2005). However, a major limitation of this system is its reliance on human coders. FACS requires intensive training to learn and approximately 2 hours for a human expert to code each minute of video (Ekman & Friesen, 1978). In contrast, CERT uses a computer vision system to automatically code according to FACS, simultaneously coding 20 AUs frame by frame in real time (20 minutes of video takes 20 minutes to code; see Fig. 1 for sample output). Consequently, CERT can reliably and accurately code large quantities of video data according to FACS in a mere fraction of the time compared with manual coding. CERT has been successfully employed in numerous studies of spontaneous expressions (Bartlett, Littlewort, Frank, & Lee, 2014; Bartlett & Whitehill, 2010; Bartlett et al., 2006; Littlewort, Bartlett, & Lee, 2009). For example, Bartlett and colleagues (2014) used CERT to successfully identify adults who were experiencing genuine pain versus those who were faking experiencing pain with high accuracy (>90%). However, no research has used CERT to examine children's facial expressions while lying in antisocial and prosocial contexts.

The current study

The current study employed CERT for the first time to investigate whether the facial expressions of children (6–11 years old) telling an antisocial lie differ from the facial expressions of children telling a

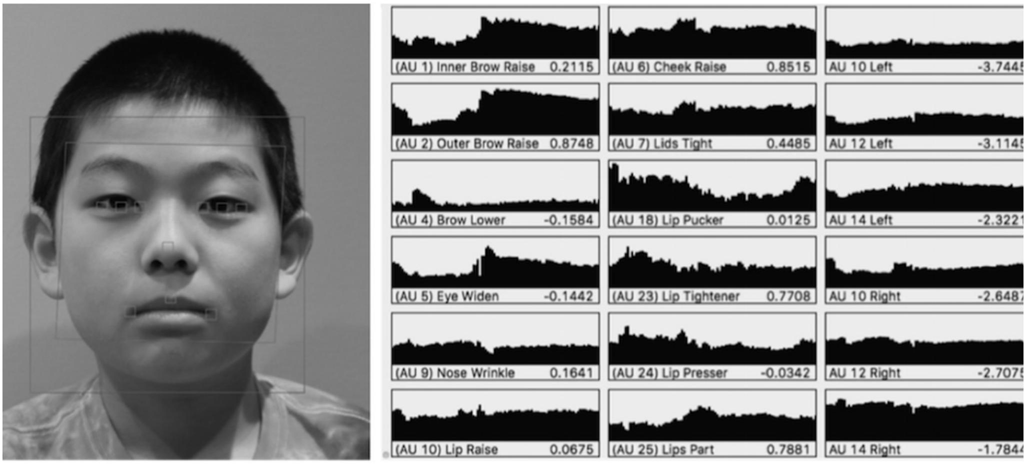

Fig. 1. Sample image of CERT and associated action unit and emotion score output. CERT first locates the face and 10 different points on the eyes, nose, and mouth. AU and emotion scores are then computed for each frame of video. Graphs on the right show a selection of AU intensity scores; the *x*-axis depicts each frame of video in temporal sequence, and the *y*-axis represents the intensity of that AU. Higher scores on the *y*-axis indicate more intense expression of that AU.

prosocial lie. We chose to examine this age group because existing studies have shown that by this time in childhood most children have refined their lie-telling abilities to such a degree that most adults cannot distinguish their lies based on facial expressions (see Lee, 2013, for a review). This created an ideal situation to challenge our computer vision program and to assess its ability to distinguish children's antisocial and prosocial lying.

Half of the sample was randomly assigned to the antisocial lying group, where children were asked not to peek at a gift being wrapped for them (the temptation resistance paradigm). All children in this group peeked at the toy and subsequently lied about their peeking behavior when asked by the experimenter. A second group of children was assigned to the prosocial lying group, where they were given a disappointing gift (a plain white bar of soap) and asked whether they liked it (the disappointing gift paradigm). All children in this group told a prosocial lie by dishonestly claiming that they liked the gift. Thus, all children in the antisocial lying group told an antisocial lie and all children in the prosocial lying group told a prosocial lie (children participated in only one lie-telling paradigm). Although the rate of lying in our sample might seem high, this is not necessarily unusual given that prior research has shown that nearly all children aged 6 to 11 years (93%) will lie when given the opportunity (Talwar, Gordon, & Lee, 2007).

Given that no child told the truth, we were unable to compare lying versus truth telling. The current study focused on comparing the facial expressions of children who spontaneously told an antisocial lie with the facial expressions of children who spontaneously told a prosocial lie. We used hidden video cameras to record children's facial expressions while lying and then used CERT to simultaneously code multiple AUs and emotions frame by frame in the two different lying paradigms. Due to the multidimensional nature of the facial expression data, we used machine learning analyses to accurately differentiate between antisocial lies and prosocial lies as well as to examine which features (AUs or emotions) specifically differentiated between these two lie types.

Based on existing (albeit limited) findings, we expected children's facial expressions to differ between antisocial and prosocial lies. We hypothesized that because the display rules for nonverbal behavior differ in an antisocial lying context from those in a prosocial lying context, the facial expressions of children in the antisocial lying group would significantly differ from the facial expressions of children in the prosocial lying group. Regarding specific AUs and emotions, we expected smiling to be a key feature in differentiating between antisocial and prosocial lies. Previous studies have found that children smile in both antisocial and prosocial lying contexts (Lewis et al., 1989; Talwar & Lee, 2002a; Talwar, Murphy, et al., 2007). However, due to human coders' limited sensitivity to facial intensity differences, it is unclear whether smiling occurs more intensely during antisocial lie-telling or prosocial lie-telling. CERT overcomes this problem due to its high sensitivity to expression intensity. Thus, one possibility would be that children would smile more when telling prosocial lies than when telling antisocial lies because display rules in a gift-giving scenario require one to smile to show appreciation for the gift. Alternatively, children might show greater smiling when telling antisocial lies than when telling prosocial lies because of the so-called "cheater's high," where a lie-teller experiences joy for deceiving another (Ruedy, Moore, Gino, & Schweitzer, 2013).

In addition to the above primary hypotheses, we also explored two secondary hypotheses. Specifically, because previous research has established that children become better lie-tellers as they get older (e.g., Lee, 2013), we predicted that children's facial expressions would become more differentiated between antisocial and prosocial lies with increased age as they become more experienced with display rules for different social contexts. Lastly, we hypothesized that, consistent with Talwar, Murphy, et al. (2007), gender differences would occur such that girls smile more than boys during prosocial lies.

Method

Participants

In total, the videos of 57 children aged 6 to 11 years (M = 108.33 months [9 years 4 months], SD = 19.09) were analyzed. Of this sample, 53% of participants (n = 30) were male. Participants were recruited from Toronto, Ontario, Canada, and the composition of ethnicity closely represented the demographics of the city of Toronto. According to parental reports, 58% of the children were White, 9% were Asian, 9% were Black, 2% were Latino, 8% were another ethnicity, and 8 parents did not indicate their children's ethnicity (14%). Additional children participated in the study; however, due to the natural settings in which the data were collected, children often turned away from the hidden cameras. CERT specifications require faces to be pointed toward the camera (within 15 degrees), with no objects blocking the view of the face, to obtain an accurate facial expression analysis. As a result, 26 participants who did not meet this criterion (12 from the antisocial group and 14 from the prosocial group) were eliminated from the current study. All children in the excluded sample lied and did not systematically vary on age, gender, or ethnicity variables. In the end, 28 children were in the antisocial lying group and 29 were in the prosocial lying group.

Design and procedure

All participants were tested individually. Testing began after obtaining consent from parents and oral assent from children. The testing room contained a hidden camera disguised as a pencil holder positioned in front of children to capture their facial expressions throughout the session. Children were not made aware of the hidden cameras until after the session was completed to ensure that their behavior during the testing session remained as natural as possible.

Antisocial lying paradigm

Children in the antisocial lying group first completed a series of filler tasks with the experimenter. Next, a temptation resistance paradigm was used to prompt children to tell an antisocial lie. The experimenter told children that they would receive a gift for their hard work but that it first needed to be wrapped. Children were asked to close their eyes and were told not to peek at the gift being wrapped for them. The experimenter set a mug on the table in front of the children atop a piece of plain brown wrapping paper and then pretended to look for tape to wrap the gift. The experimenter, after saying "Where's my tape?" to herself, but in a voice audible to the children, pretended to look for tape in a bin behind children's back for 20 seconds (M = 19.41 seconds, SD = 5.41), giving children an opportunity to peek at the gift without getting caught committing the transgression. Because this is a highly tempting situation for children, all participants peeked at the gift. To signal that they were returning to where children were seated, the experimenter stated, "Here's my tape!" and walked to the table to finish wrapping the item. The experimenter placed a large box over the mug to facilitate wrapping the item and help conceal its identity. Once the gift was wrapped, children were told to open their eyes and the experimenter asked the target question, "While I was wrapping the gift, did you peek?" All children peeked at the gift and subsequently lied about their actions. After children responded to the question, they were permitted to open the gift.

Prosocial lying paradigm

Similar to previous research examining prosocial lying, we used a disappointing gift paradigm to prompt children in the prosocial lying group to tell a prosocial lie (Talwar, Murphy, et al., 2007). Children in the prosocial lying group first completed a manipulation check with a research assistant. The research assistant asked children to rank four items in order from the one they liked the best to the one they liked the least; three of the items were ranked by children to be desirable items (e.g., dolls, toy trucks), whereas one gift, a package of plain soap, was ranked as the undesirable fourth gift by all participants. After ranking the gifts, the main experimenter completed the same filler tasks with children as were completed with children in the antisocial group. Afterward, children were told that for successfully completing the study they would receive a gift. The experimenter told children, "I know that you selected a gift with [research assistant], but I picked out a gift for you before you got here." The experimenter then gave children a closed paper bag containing a bar of white soap (the disappointing gift), and children were instructed to open their gift. Meanwhile, the experimenter was turned away from the children while cleaning up the room for 30 seconds. This was done to closely replicate previous studies where children were given disappointing gifts and briefly left alone in the room to allow them to naturally react to the gift (Talwar, Murphy, et al., 2007). Following the 30-second delay, the experimenter turned around and returned to the seat across from the children before asking them the target question, "Do you like your gift?" All children lied to the experimenter by falsely claiming to like the gift. As a final manipulation check, the experimenter asked children whether they wished to switch their awarded gift for one that they had previously selected. All children switched their gift, confirming that the bar of soap was indeed something that they disliked.

Data analysis

We used CERT software to analyze children's facial expressions frame by frame. CERT uses a method called support vector machine (SVM) to analyze the actions of different features on the face. SVM finds what is referred to as a "hyperplane" in a multidimensional space, which maximally separates the data into the different features. Using this hyperplane, CERT then calculates how intense each feature is being displayed (i.e., a subtle smile compared with a more pronounced smile). CERT provides two types of feature scores: action unit scores and emotion scores.

Action unit scores code the intensity of 20 different action unit movements (see Littlewort et al., 2011, for details) plus an additional 6 unilateral AUs. The CERT output is the distance to the separating hyperplane. There is one SVM per AU, detecting the presence or absence of the AU. These scores indicate the degree to which each muscle group is activated. CERT outputs correlate with expression intensity, as measured through multiple methods: CERT estimates of expression intensity correlated with FACS expert intensity codes (Bartlett et al., 2006). CERT also correlates with zygomatic facial EMG (electromyography) in simultaneous recordings of 15 participants and 300 trials ($r = .61$, $p < .0001$; Littlewort et al., 2011). In addition, smile intensity is correlated with human intensity labels at $r = .89$ (Whitehill, Littlewort, Fasel, Bartlett, & Movellan, 2009). This precise information on the dynamics of each facial action that CERT provides would normally be impractical to obtain through manual FACS coding.

In contrast to action unit scores, CERT uses a multinomial logistic regression (MLR) to transform the output of the 20 AUs to identify basic emotions. These emotion detectors were trained using the Cohn–Kanade dataset, a large high-quality database containing thousands of images of each type of emotional expression (Kanade, Cohn, & Tian, 2000; Littlewort et al., 2011). The output for each of the basic emotions and neutral state is passed through a softmax competition. The final output codes the likelihood in logarithm that eight different emotional expressions (joy, sadness, anger, disgust, surprise, fear, contempt, and a neutral expression) are occurring on the face in each frame (e.g., the log [p(joy)/p(not_joy)]. These scores range from 0 to 100%, with higher scores indicating a greater probability that the emotion is shown. Within each frame, the scores of each of the eight emotional expressions add up to 100%. Thus, a child's scores in a single frame might indicate he or she is 80% likely to be showing anger, 20% likely to be showing contempt, and 0% likely to be showing any of the other six emotions.

Evaluation of CERT on benchmark datasets shows exceptional performance for both the recognition of basic emotions and facial action units. That is, CERT was established as the most accurate and efficient software program to date for the analysis of facial expressions. Detection accuracy for facial actions has a mean area under the curve (AROC) of .93 for posed expressions and a mean AROC of .84 for spontaneous expressions that include head movements and speech (Littlewort et al., 2011). Emotion scores have been shown to be extremely accurate, with .98 AROC when tested on a dataset of posed expressions (Littlewort et al., 2011). Thus, CERT provides information on facial expression intensity and dynamics at temporal resolutions that were previously impractical via FACS manual coding and without requiring application of electrodes to the face.

Before analyzing videos through CERT, we cropped the video footage of participants' faces to isolate the key areas of the lying scenarios. Videos from children in the antisocial lying group were edited to begin immediately after the experimenter asked children, "While I was wrapping the gift, did you peek?" Videos from children in the prosocial lying group began immediately after the experimenter asked children, "Do you like your gift?" In both cases, the video clip ended right *before* children gave their response so that no dialogue was present in the video. This was done so that the facial movements required for speaking did not influence CERT analysis, as FACS coding with speaking present is unreliable. Furthermore, children gave different responses in the two lying groups—"no" in the antisocial lying group and "yes" in the prosocial lying group—in order to lie. Thus, using facial expression data extracted during speaking would inaccurately distort the data. Consequently, to preserve the integrity of the CERT data, each video began immediately after the experimenter finished asking children the target question and stopped immediately before children began to give a response (either verbal or nonverbal). After editing the videos, we used CERT to obtain a frame-by-frame (24 frames per second) coding of action unit scores and emotion scores for each participant.

Due to the nature of the experimental paradigms, the number of available frames that CERT could analyze varied per participant. Because children's lies were spontaneous, response times for children to answer the target question varied. The number of available frames ranged from 7 to 16 frames per participant ($M_{\text{frame count}}$ = 9.13, SD = 2.05). Recall that video was recorded at 24 frames per second. This meant that all children lied roughly half a second after the question was asked. We took the scores for each emotion and action unit and averaged them across all frames within each participant to create a single score for each of the action units and a single score for each of eight emotions per participant.

Machine learning

We examined whether we could use the multidimensional information extracted from the facial videos by CERT to differentiate children who told an antisocial lie from those who told a prosocial lie. To answer this question, we applied a machine learning technique to the facial videos encoded by CERT. The core of machine learning is what is referred to as a "classifier". We used a specific type of classifier called a linear support vector machine. The linear SVM finds what is referred to as a hyperplane in a multidimensional space defined by the features in the training set, which maximally separates between two groups. This hyperplane is then applied to the corresponding multidimensional space defined by the features in the testing set to correctly classify each case into one of the two groups (Chang & Lin, 2011; Hsu, Chang, & Lin, 2003; Vapnik, 2000). In the current study, the goal of the classifier was to use the CERT data (action unit scores and emotion scores) to correctly categorize each participant based on whether the child told an antisocial lie or a prosocial lie. We conducted two separate sets of linear SVM analyses to perform pattern classification based on the facial movements measured by CERT using emotion scores and action unit scores separately.

SVM analyses require the dataset to be split into two groups: one group of data for training the SVM model and a second group of data for testing the model's accuracy (Chang & Lin, 2011). This ensures that we are not creating a biased model by training and testing model accuracy on the same data, as would be the case had we used the entire dataset for training and testing. Using the training set, the linear SVM first ''learns'' the intrinsic association between the features (i.e., emotions/action units and their corresponding lie types), identifying the combination of features that maximally differentiates the two types of lies. In other words, it identifies which emotions/action units show distinctly different patterns between antisocial lies and prosocial lies. For example, mean scores in one emotion/action unit may be consistently higher in one lie type compared with the other. Using the knowledge learned from the training set, the classifier (a linear equation or model) is then developed with the goal of accurately predicting lie type based on emotion/action unit scores. For example, if high scores for anger are always in the videos of prosocial lies but not in those of antisocial lies in the learning set, then the classifier will use eyebrows (along with other learned cues) to classify videos as belonging to the prosocial lying group or antisocial lying group using the videos in the testing set. However, if scores for anger do not systematically differ between the lying groups, then the classifier will not use that feature to classify the testing set and will instead choose more distinct features to inform the classifier. After the classifier makes predictions on the testing set, these predictions are then compared with the actual lying type to which the test data belong (which was unknown to the classifier during the testing session).

As mentioned above, SVM analyses require the dataset to be split into two groups: a training set and a testing set. That way, we are able to evaluate the performance of SVM using data that were *not* involved in calculating the SVM parameters (training the classifier). Given our relatively small sample size, we used a procedure commonly referred to as a leave-one-out cross-validation method to maximize the use of our data set. This means that of our total sample of 57 participants, we set aside one video to use as the testing set and the remaining 56 videos to use as the training set. We then performed the learning and testing procedure described in the previous paragraph. This procedure was repeated independently a total of 57 times, each time using a different video as the testing set. Thus, each participant's video was used as the test set once, and the training of the classifier in each iteration was independent of the training that occurred in prior iterations.

Each time the classifier was tested on the testing data, a score of either correct (1) or incorrect (0) classification was then given. If the emotion/action unit scores do carry useful information that can be used by the classifier to differentiate the two types of lies, then we would expect the average prediction accuracy to be above chance, which in this case is 50%. Successful prediction during the testing stage would suggest that the information learned by the classifier can be applied to new facial videos to make accurate predictions of lie type in new participants.

Feature selection

The SVM classifier takes into account each of the different features (i.e., action units and emotions) in the classification process; however, different features may have greater importance for the classification process than others. It is possible that by selecting a subset of features that are the most useful to classification performance, we will optimize the accuracy of the classifier. We employed a leave-one-out cross-validation method to detect which features were the most important, or most influential, in the model. Using a recursive procedure, we are able to determine which set of features would obtain the greatest classification accuracy.

After obtaining the mean accuracy score using all the features together, we identified the feature that was least important to the classifier. To identify the least important feature, we examined how strongly each feature was weighted in the classifier. Strongly weighted features indicate that scores on these features systematically differ between antisocial and prosocial lies. In contrast, weakly weighted features indicate that the scores on these features do not differentiate the two lie types well and, therefore, are not as useful to the model. We eliminated the least important feature and repeated the same training and testing procedure previously described in a second cycle, obtaining a new mean prediction accuracy.

This process of feature elimination and retesting was repeated with one less feature present until there was only one feature left for the classifier to use. Each time, the classifier was trained and tested independently from prior iterations. Finally, we examined the accuracy scores for all observed feature combinations. We looked for an inverted U shape in the performance curve as a function of the number of features and chose the feature combination that provided the highest classification performance accuracy (at the peak of the inverted U) as the final model.

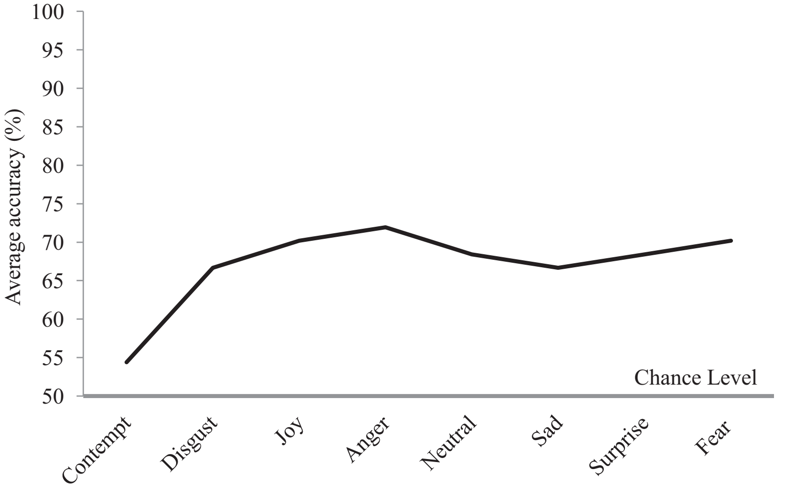

Fig. 2. Classification accuracy using emotion scores. Each data point represents the average accuracy score when all emotions listed to the left of the *x*-axis label are included in the model. Thus, the first accuracy level is achieved when contempt is added to the model alone. The second accuracy level is achieved when disgust is added to contempt. The third accuracy level is achieved when joy is added to contempt and disgust, and so on.

This entire procedure was conducted first with emotion scores as the features and then with action unit scores as the features.

Results

Results based on emotions

To examine whether emotion scores significantly differed between children who told an antisocial lie and children who told a prosocial lie, we started with all eight emotions and created and tested eight different models, each time removing the least important emotion. The accuracy scores of each model are displayed in Fig. 2. The model with the highest classification success obtained an average accuracy score of 72% and contained four emotions; anger, contempt, disgust, and joy. To test whether this was above chance-level accuracy, we randomized the dataset both between participants and between emotions and then tested the accuracy of our final model using the leave-one-out method described above. This randomization procedure was repeated over 1000 iterations in order to obtain a chance distribution. We found the average chance-level accuracy to be 50% over 1000 iterations, 95% *CI* [7.02, 64.91]. An examination of the confidence intervals reveals that our model is significant above chance level. Thus, the emotion probability scores of children when telling either an antisocial lie or a prosocial lie can be significantly differentiated between the two types of lies. Age was not significantly correlated with successfully predicting lie type, $b = -0.01$, $p = .46$. This means that the classifier was able to differentiate between the emotional expressions of younger children just as well as it was for the older children, suggesting that the emotional expressions of younger children may be just as differentiated between antisocial and prosocial lying as the expressions of older children in our sample.

To determine which emotions were more likely to occur in each lying group, we conducted a series of generalized linear model (GLM) analyses with lie type (antisocial vs. prosocial), gender (male vs. female), and age (as a continuous variable) as the independent variables. We examined all emotions used in the final emotion SVM model (anger, contempt, disgust, and joy) as the dependent variables. Results showed no significant effects of lie type for both anger and disgust (all *p*s > .05). However, we found a marginal effect of lie type for contempt, whereby children in the prosocial group were more likely to show contempt ($M = 36\%$, $SD = 24$) than children in the antisocial group ($M = 25\%$, $SD = 17$), $F(1, 53) = 3.38$, $p = .071$, $\eta_p^2 = .060$. Furthermore, whereas we found no significant main effect of gender, a significant interaction between lie type and gender was revealed for contempt, $F(1, 53) = 6.49$, $p = .014$, $\eta_p^2 = .109$ (Fig. 3). Follow-up analyses revealed that boys showed greater contempt during the prosocial lie ($M = 44\%$, $SD = 24$) than during the antisocial lie ($M = 21\%$, $SD = 15$), $t(28) = -3.02$, $p = .005$. However, girls did not significantly differ in the expression of contempt between antisocial ($M = 28\%$, $SD = 18$) and prosocial ($M = 25\%$, $SD = 19$) lying groups, $t(25) = 0.52$, $p = .608$ (Fig. 3).

Overall, children in the prosocial lying group were more likely to display joy ($M = 11\%$, $SD = 20$) than children in the antisocial lying group ($M = 4\%$, $SD = 7$), $F(1, 53) = 4.97$, $p = .030$, $\eta_p^2 = .086$, and girls were more likely to display joy ($M = 12\%$, $SD = 21$) than boys ($M = 4\%$, $SD = 6$), $F(1, 53) = 5.03$, $p = .029$, $\eta_p^2 = .087$. However, we also found a significant gender by lie type interaction, $F(1, 53) = 6.60$, $p = .013$, $\eta_p^2 = .111$. Girls were more likely to show joy ($M = 22\%$, $SD = 23$) than boys ($M = 4\%$, $SD = 4$) in the prosocial group, but no gender differences were found in the antisocial group (Fig. 3). Finally, we found no significant age effects, indicating that the expression of emotions during lie-telling does not change with age.

Results based on AUs

To examine whether AU scores significantly differed between children who told an antisocial lie and children who told a prosocial lie, we conducted the same SVM procedure as we did with the emotion scores. This time, we started with all 26 AUs and created and tested 26 different models, each time removing the least important emotion. The accuracy scores of each model are displayed in Fig. 4. Based on our SVM analysis, the model with the highest classification success achieved an average accuracy rate of 79% classification and contained three action units: lip pucker (AU18), blink/eye closure (AU 45), and unilateral lip raise on the right side (AU 10). The randomized chance-level distribution obtained an average accuracy of 49%, 95% *CI* [14.91, 64.91]. Therefore, our final model using emotion scores classified whether children told an antisocial lie or a prosocial lie significantly above chance. Age was not significantly correlated with the successful lie-type classification, $b = -0.01$, $p = .51$. This means that, similar to the results using emotion scores, the classifier was able to differentiate between the facial expressions of younger children just as well as it was for the older children, suggesting that the facial expressions of younger children may be just as differentiated between antisocial lying and prosocial lying as the expressions of older children.

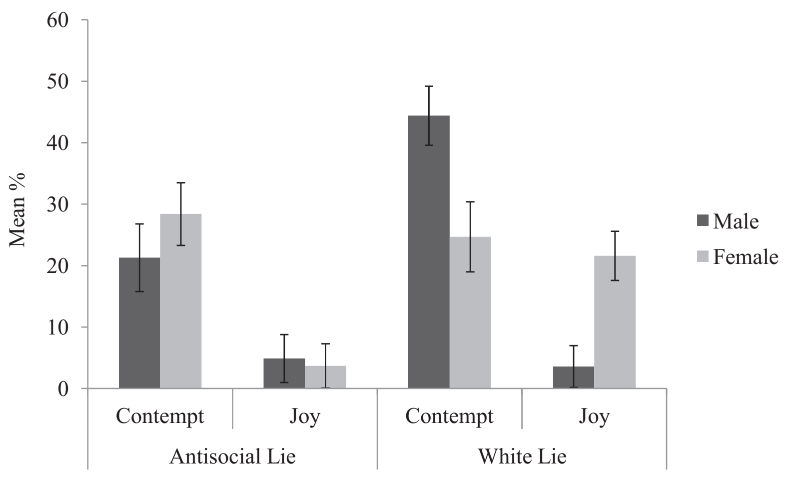

Fig. 3. Average probability of a contempt and joyful expression according to gender. Error bars represent standard error values.

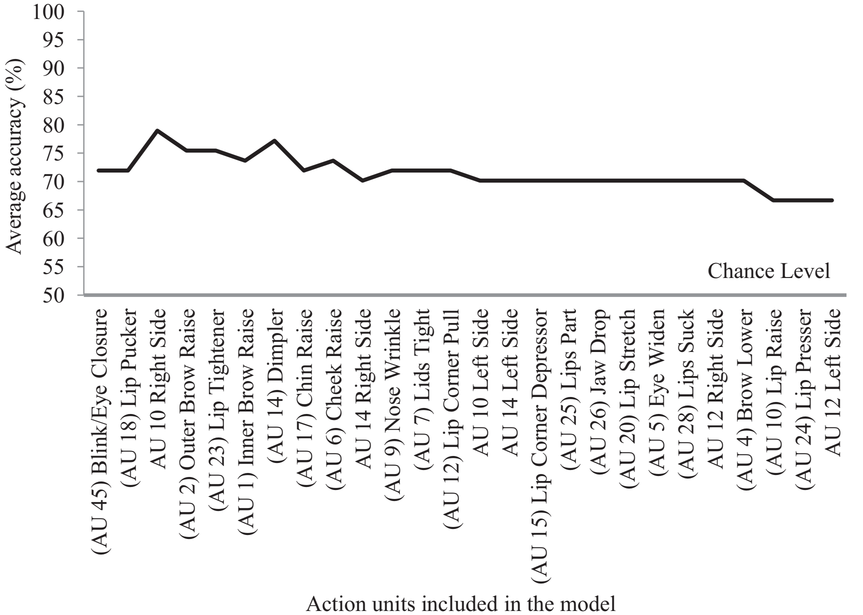

Fig. 4. Classification accuracy using action unit scores. Each data point represents the average accuracy score when all action units listed to the left of the *x*-axis label are included in the model. Thus, the first accuracy level is achieved when blink/eye closure is used in the model alone. The second accuracy level is achieved when lip pucker is added to blink/eye closure. The third accuracy level is achieved when lip raise on the right side is added to the blink/eye closure and lip pucker, and so on.

We then conducted further analyses to determine which specific action units differed in intensity between the antisocial lying group and the prosocial lying group and how they differed. To do so, we performed a series of GLM analyses with lie type (antisocial vs. prosocial), gender (male vs. female), and age (as a continuous variable) as the independent variables. We examined all AUs used in the final AU SVM model (lip pucker [AU 18], blink/eye closure [AU 45], and unilateral lip raise on the right side [AU 10]) as the dependent variables. We discovered that children who told an antisocial lie displayed lip pucker ($M = -0.02$, $SD = 0.02$) with greater intensity than children who told a prosocial lie ($M = -0.03$, $SD = 0.01$), $F(1, 53) = 5.85$, $p = .019$, $\eta_p^2 = .099$. In contrast, children in the prosocial lying group displayed both blink/eye closure ($M = -0.23$, $SD = 0.85$) and unilateral lip raise on the right side ($M = -3.80$, $SD = 0.47$) with greater intensity than children in the antisocial lying group (blink/eye closure $M = -0.83$, $SD = 0.40$, unilateral lip raise on the right side $M = -4.11$, $SD = 0.61$), $F(1, 53) = 10.26$, $p = .002$, $\eta_p^2 = .162$, $F(1, 53) = 4.73$, $p = .034$, $\eta_p^2 = .082$. No significant gender differences or interaction effects were found with any of the action unit scores. Finally, we found no significant age effects (all ps > .05), indicating that the intensity of AUs during lie-telling does not change with age.

Discussion

The current study used a computer vision program to decode facial expressions associated with children's antisocial and prosocial lies. We obtained four major findings. First, children's facial expressions during antisocial and prosocial lying can be accurately and reliably differentiated in terms of either emotions or AUs. Second, joy is the key emotion that differentiates children's antisocial and prosocial lies. Third, gender differences occurred, whereby boys showed greater contempt during the prosocial lie than during the antisocial lie, whereas girls showed more joy and less contempt compared with boys when they told the prosocial lie. Finally, the key AUs that differentiate children's antisocial and prosocial lies are blink/eye closure, lip pucker, and lip raise on the right side. We discuss these major findings in sequence.

Our success in using CERT to accurately and reliably differentiate children's antisocial and prosocial lies demonstrates for the first time that computer vision technology can be used to study children's facial expressions during deception. Its advantages over using human coders are numerous. First, it is reliable. CERT will produce identical results when the same video is coded multiple times, unlike trained human coders who inevitably will produce inconsistencies due to attention lapses, and fatigue (Littlewort et al., 2011). Second, CERT can look at multiple AUs simultaneously, whereas human coders are limited to coding one or two AUs at a time. Thus, CERT not only significantly improves coding efficiency but also enables examining a much larger range of AUs at once. The same is also true for emotions. Third, CERT can obtain measurements for each AU or emotion coded at a frame-by-frame temporal resolution, which is impossible for human coders to achieve. Fourth, because CERT runs on a computer, it can analyze facial expressions with much greater speed and at a far lower cost than human coders.

Children in the prosocial lying group told the experimenter that they liked the gift given to them even though they did not. In a typical gift-giving situation, social expectations require that the recipient express gratitude toward the gift-giver by smiling and giving thanks. In contrast, children in the antisocial lying group told a lie to cover up a misdeed. Display rules in this scenario do not require children to smile. In fact, smiling in this context may arouse the lie recipient's suspicion that the child is lying given that many adults cite smiling as a sign of deceit (Vrij & Semin, 1996). Thus, the fact that our participants showed more joy in the prosocial lie than in the antisocial lie suggests that they might be aware of these differing expectations and were able to manipulate their facial behaviors accordingly so as to tell a convincing lie.

We had predicted that, based on previous research on prosocial lie-telling (Talwar, Murphy, et al., 2007), girls would smile more compared with boys during the prosocial lie. Our results supported this hypothesis. Girls showed more joy than boys during the prosocial lie, whereas no gender differences were found in any emotions during the antisocial lie. Specifically, girls in the prosocial lying group were 22% likely to show joy, whereas boys were only 4% likely to show joy. In fact, the likelihood of showing joy was the same for boys who told an antisocial lie as it was for boys who told a prosocial lie. The gender difference in the expression of joy during prosocial lying may be because girls are more sensitive to the social display rules for receiving a gift. Existing literature on display rules during a disappointing situation suggests that gender differences emerge due to different cultural expectations for boys and girls (Davis, 1995; McDowell et al., 2000; Tobin & Graziano, 2011). Politeness is generally viewed as a more salient cultural value for girls, whereas expressions of negative emotions and behaviors (e.g., aggression) are more readily tolerated in boys. Given this increased pressure on girls to be polite, the gender differences in facial expression while telling a prosocial lie could be the result of girls' greater practice at suppressing negative feelings compared with boys. Alternatively, a biological explanation could suggest that boys and girls experience different developmental timetables in obtaining the ability to suppress negative emotions and adhere to social display rules (Davis, 1995; McDowell et al., 2000; Tobin & Graziano, 2011). As discussed below, the current study did not have a large enough sample size to adequately measure age differences to test this developmental hypothesis.

We also found that boys were more likely to show contempt while telling a prosocial lie compared with girls. A contempt facial expression is shown by tightening and raising the corner of the lip on one side and is considered a universal expression for contempt (see Ekman & Friesen, 1986 for images; Ekman & Heider, 1988; Matsumoto, 1992). When children are told that they will be receiving a gift, they are likely to get excited in the anticipation of receiving something desirable for them to enjoy. However, when children open up the wrapping to reveal a disappointing gift, it is possible that they experience feelings of contempt toward the gift-giver for falsely getting their hopes up. We found that, within the prosocial lying group, boys were 44% likely to display contempt, whereas girls were only 25% likely to show contempt. In fact, boys in the prosocial lying group expressed greater contempt compared with boys in the antisocial lying group, whereas girls did not differ between the two groups. We propose two possible explanations for this gender difference. It could be that boys actually experience more contempt than girls in this scenario or that female children are simply more skilled at concealing their feelings of contempt during a prosocial lie. The latter explanation is consistent with the previously described interpretation that girls are better able to control negative emotions and adhere to social display rules compared with boys (Davis, 1995; McDowell et al., 2000; Tobin & Graziano, 2011). Further research is needed to compare facial expressions when a child receives a desired gift in order to tease apart these possibilities.

CERT analyses revealed that three key AUs are sufficient to differentiate antisocial lies from prosocial lies in children above chance level: lip pucker, blink/eye closure, and lip raise on the right side. Lip pucker occurred with greater intensity during the antisocial lie, whereas blink/eye closure and lip raise on the right side occurred with greater intensity during the prosocial lie. Furthermore, four emotions together were sufficient to differentiate antisocial and prosocial lies in children significantly above chance level: joy, anger, disgust, and contempt. When we examined the unique contributions of each emotion, joy was significantly different between the two lying types. Children were more likely to show joy during prosocial lying than during antisocial lying. This finding was consistent with our prediction that children would smile more during prosocial lying than during antisocial lying. Also consistent with this interpretation, the intensity of lip pucker was larger during the antisocial lie. Because lip puckering is an antagonistic muscle movement to smiling, it could be that children are trying to prevent smiling during the antisocial lie because they believe that they will get caught lying if they smile. We found a marginal effect where contempt occurred more during prosocial lying than during antisocial lying overall, an effect that was significant for boys in the sample. The right lip raise result from the analysis based on AUs complements the contempt finding because unilateral lip raise is part of the muscle movements associated with contempt. One unexpected finding is that children blinked less during antisocial lying than during prosocial lying. The exact reason for this difference is unclear. Although the existing research has shown that increased blinking in adults while lying is reflective of experiencing an increase in cognitive load (Leal & Vrij, 2008), it is unclear why children might experience a greater cognitive load while telling a prosocial lie than while telling an antisocial lie. This issue needs to be addressed specifically in future studies.

We found no evidence to support the prediction that children's facial expressions would become more differentiated between the two lie types with increased age. Furthermore, children's facial expressions did not differ with age when both AU and emotion scores were examined independently. This null finding is inconsistent with findings by Talwar, Murphy, et al. (2007), who found that children's smiling during a gift-giving scenario increased with age regardless of whether they liked the gift or not. However, a major limitation of the current study is our sample size. It is possible that with an overall sample size of 57 participants, we did not have enough children to represent each of the age groups. In addition, we may need to examine children under age 6 years to accurately capture developmental trends in facial expressions because the preschool years have been found to be the age period when children rapidly develop their lying skills (Lee, 2013). Future research with a larger sample size and broader age ranges may be better suited to examining whether or not children's facial expressions change as they develop from being poor lie-tellers to skilled ones.

The current study has several additional limitations that warrant attention. First, due to the lack of truth-tellers in both antisocial and prosocial lying paradigms, we were unable to examine whether child liars would display different facial expressions from truth-tellers in either the antisocial or prosocial paradigm. Such research will provide further important insights into children's ability to manipulate their facial expressions to tell a convincing lie. Second, using linear SVM, we obtained a prediction accuracy of 79% based on AUs and 72% based on emotions. Although these accuracies are higher than those of naive human observers (Talwar & Lee, 2002a, 2002b), there is much room for improvement. This could be achieved by using more sophisticated machine learning models than linear SVM. Due to the limitations of linear SVM, we used only means and standard deviations of AU and emotion scores. However, CERT provides much richer temporal information at a frame-by-frame level. Future studies using machine learning algorithms that incorporate such temporal information should improve the detection accuracy. Third, the current study used a between-participants design, whereby half of the children were in the antisocial lying group and the other half were in the prosocial lying group. Future studies could concurrently test the same children in both the antisocial and prosocial paradigms to examine whether the same children's facial expressions differ when telling antisocial and prosocial lies.

Fourth, we were able to examine only one type of antisocial lying (to conceal a transgression) and one type of prosocial lying (when receiving a disappointing gift) when in fact there are many other circumstances where an antisocial lie or a prosocial lie may take place. Thus, it is not clear whether our findings would replicate across similar antisocial and prosocial lie-telling situations, where perhaps both the social display rules and motivations for telling a lie may differ. Furthermore, it should be noted that the current study focused only on children spontaneously lying, which may be different from coached lying (a form of lying that has important legal implications; Ahern et al., 2011). Whether children tell lies out of their own volition or due to coaching may influence children's facial expressions. In spontaneous lie-telling, children lie because they themselves choose to do so in the moment. In contrast, in coached lie-telling situations, children lie because the experimenter asks them to lie well before the actual opportunity to lie presents itself (Lyon, Malloy, Quas, & Talwar, 2008). This means that children have a longer period of time to prepare themselves to tell a lie compared with the in-the-moment decision that occurs in spontaneous lie-telling. The facial muscles associated with voluntary and involuntary facial movements are controlled by different cortical structures in the brain. The location where these movements originate in the brain influences both the type of facial movements portrayed and their dynamics (Bartlett et al., 2014; Ekman, 2001; Ekman & Rosenberg, 2005; Littlewort et al., 2009; Rinn, 1984). Because the motivation behind why a child tells a lie differs in spontaneous deception compared with coached deception, the facial expressions that occur during spontaneous lie-telling may differ from the facial expressions that occur during coached lie-telling. Thus, future studies need to explore these potential differences.

Conclusions

The current study used computer vision technology to reveal for the first time that the facial expressions of children telling an antisocial lie are not the same as those when telling a prosocial lie. Using CERT, we were able to correctly classify children's lies based on facial expressions with high accuracy. Overall, four key features were shown to differ between children's antisocial and prosocial lies: the expression of joy, blink/eye closure, lip pucker, and lip raise on the right side. We also discovered gender differences during prosocial lying, whereby girls showed more joy and less contempt compared with boys when they told a prosocial lie, and boys showed more contempt during the prosocial lie than during the antisocial lie. Together, these findings offer new insights into children's developing abilities to manipulate their nonverbal behaviors according to the demands of different social contexts of lie-telling during childhood.

Acknowledgments

This study was supported in part by the Canadian Institutes of Health Research (130259), the Social Sciences and Humanities Research Council of Canada, and the Canadian Foundation for Innovation.

References

Ahern, E. C., Lyon, T. D., & Quas, J. A. (2011). Young children's emerging ability to make false statements. *Developmental Psychology, 47*, 61–66.
Bartlett, M. S., Littlewort, G. C., Frank, M. G., Lainscsek, C., Fasel, I., & Movellan, J. R. (2006). Automatic recognition of facial actions in spontaneous expressions. *Journal of Multimedia, 1*(6), 22–35.
Bartlett, M. S., Littlewort, G. C., Frank, M. G., & Lee, K. (2014). Automatic decoding of facial movements reveals deceptive pain expressions. *Current Biology, 24*, 738–743.
Bartlett, M. S., & Whitehill, J. (2010). Automated facial expression measurement: Recent applications to basic research in human behavior, learning, and education. In A. Calder, G. Rhodes, J. V. Haxby, & M. H. Johnson (Eds.), *Handbook of face perception* (pp. 489–513). Oxford, UK: Oxford University Press.
Bussey, K. (1999). Children's categorization and evaluation of different types of truths and lies. *Child Development, 70*, 1338–1347.
Chang, C., & Lin, C. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology, 2*(3), 1–27. Software available at ~cjlin/libsvm" xlink:type="simple" id="ir005">http://www.csie.ntu.edu.tw/~cjlin/libsvm.
Darwin, C. (1877). A biographical sketch of an infant. *Mind, 2*, 285–294.
Davis, T. L. (1995). Gender differences in masking negative emotions: Ability or motivation? *Developmental Psychology, 31*, 660–667.
Donato, G. L., Bartlett, M. S., Hager, J. C., Ekman, P., & Sejnowski, T. J. (1999). Classifying facial actions. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 21*, 974–989.

Ekman, P. (2001). Facial expressions. In C. Blakemore & S. Jennett (Eds.), *Oxford companion to the body* (pp. 113–134). Oxford, UK: Oxford University Press.
Ekman, P., & Friesen, W. V. (1978). *Facial action coding system*. Palo Alto, CA: Consulting Psychologists Press.
Ekman, P., & Friesen, W. V. (1986). A new pan-cultural facial expression of emotion. *Motivation and Emotion, 10*, 159–168.
Ekman, P., Friesen, W. V., & Hager, J. C. (2002). *Facial action coding system*. Salt Lake City, UT: Network Information Research Corporation.
Ekman, P., & Heider, K. G. (1988). The universality of a contempt expression: A replication. *Motivation and Emotion, 12*, 303–308.
Ekman, P., & Rosenberg, E. L. (Eds.). (2005). *What the face reveals: Basic and applied studies of spontaneous expression using the FACS*. Oxford, UK: Oxford University Press.
Evans, A. D., & Lee, K. (2013). Emergence of lying in very young children. *Developmental Psychology, 49*, 1958–1963.
Fu, G., Evans, A. D., Xu, F., & Lee, K. (2012). Young children can tell strategic lies after committing a transgression. *Journal of Experimental Child Psychology, 113*, 147–158.
Hartshorne, H., & May, M. A. (1928). *Studies in the nature of character. I. Studies in deceit*. New York: Macmillan.
Hsu, C. W., Chang, C. C., & Lin, C. J. (2003). *A practical guide to support vector classification* Unpublished manuscript : . Department of Computer Science, National Taiwan University. Retrieved from https://www.cs.sfu.ca/people/Faculty/teaching/726/spring11/svmguide.pdf.
Kanade, T., Cohn, J. F., & Tian, Y. L. (2000). Comprehensive database for facial expression analysis. In *Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition (FG'00)* (pp. 46–53). Grenoble, France: IEEE.
Leal, S., & Vrij, A. (2008). Blinking during and after lying. *Journal of Nonverbal Behavior, 32*, 187–194.
Lee, K. (2013). Little liars: Development of verbal deception in children. *Child Development Perspectives, 7*(2), 91–96.
Lewis, M., Stanger, C., & Sullivan, M. W. (1989). Deception in 3-year-olds. *Developmental Psychology, 25*, 439–443.
Littlewort, G. C., Bartlett, M. S., & Lee, K. (2009). Automatic coding of facial expressions displayed during posed and genuine pain. *Image and Vision Computing, 27*, 1797–1803.
Littlewort, G., Whitehill, J., Wu, T., Fasel, I., Frank, M., Movellan, J., et al (2011). The computer expression recognition toolbox (CERT). *2011 IEEE International Conference on Automatic Face & Gesture Recognition and Workshops (FG 2011)* (pp. 298–305). Santa Barbara, CA: IEEE.
Lyon, T. D., Malloy, L. C., Quas, J. A., & Talwar, V. A. (2008). Coaching, truth induction, and young maltreated children's false allegations and false denials. *Child Development, 79*, 914–929.
Matsumoto, D. (1992). More evidence for the universality of a contempt expression. *Motivation and Emotion, 16*, 363–368.
McDowell, D. J., O'Neil, R., & Parke, R. D. (2000). Display rule application in a disappointing situation and children's emotional reactivity: Relations with social competence. *Merrill-Palmer Quarterly, 46*, 306–324.
Polak, A., & Harris, P. L. (1999). Deception by young children following noncompliance. *Developmental Psychology, 35*, 561–568.
Rinn, W. E. (1984). The neuropsychology of facial expression: A review of the neurological and psychological mechanisms for producing facial expressions. *Psychological Bulletin, 95*(1), 52–77.
Ruedy, N. E., Moore, C., Gino, F., & Schweitzer, M. E. (2013). The cheater's high: The unexpected affective benefits of unethical behavior. *Attitudes and Social Cognition, 105*, 531–548.
Saarni, C. (1989). Children's understanding of the interpersonal consequences of dissemblance of nonverbal emotional–expressive behavior. *Journal of Nonverbal Behavior, 12*, 275–294.
Talwar, V., & Crossman, A. M. (2011). Predictors of children's prosocial lie-telling: Motivation, socialization variables, and moral understanding. *Journal of Experimental Child Psychology, 110*, 373–392.
Talwar, V., Crossman, A., Williams, S., & Muir, S. (2011). Adults' detection of children's selfish and polite lies: Experience matters. *Journal of Applied Social Psychology, 41*, 337–359.
Talwar, V., Gordon, H. M., & Lee, K. (2007a). Lying in the elementary school years: Verbal deception and its relation to second-order belief understanding. *Developmental Psychology, 43*, 804–810.
Talwar, V., & Lee, K. (2002a). Development of lying to conceal a transgression: Children's control of expressive behavior during verbal deception. *International Journal of Behavioral Development, 26*, 436–444.
Talwar, V., & Lee, K. (2002b). Emergence of white lie-telling in children between 3 and 7 years of age. *Merrill-Palmer Quarterly, 48*, 160–181.
Talwar, V., & Lee, K. (2008). Social and cognitive correlates of children's lying behavior. *Child Development, 79*, 866–881.
Talwar, V., Murphy, S. M., & Lee, K. (2007b). White lie-telling in children for politeness purposes. *International Journal of Behavioral Development, 31*, 1–11.
Tobin, R. M., & Graziano, W. G. (2011). The disappointing gift: Dispositional and situational moderators of emotional expressions. *Journal of Experimental Child Psychology, 110*, 227–240.
Vapnik, V. (2000). *The nature of statistical learning theory*. New York: Springer.
Vrij, A., & Semin, G. R. (1996). Lie experts' beliefs about nonverbal indicators of deception. *Journal of Nonverbal Behavior, 20*, 65–80.
Whitehill, J., Littlewort, G., Fasel, I., Bartlett, M., & Movellan, J. (2009). Towards practical smile detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 99*, 2106–2111.