Signals of Emotion Regulation in a Social Dilemma: Detection from Face and Context

Rens Hoegen*, Jonathan Gratch*, Brian Parkinson† and Danielle Shore†

*University of Southern California, Institute for Creative Technologies, Los Angeles, CA, USA

{rhoegen, gratch}@ict.usc.edu

†University of Oxford, Oxford, UK

{brian.parkinson, danielle.shore}@psy.ox.ac.uk

Abstract—In social decision-making tasks, facial expressions are informative signals that indicate motives and intentions. As people are aware that their expressions influence partner behavior, expressions may be strategically regulated in competitive environments to influence a social partner's decision-making. In this work, we examine facial expressions and their strategic regulation within the context of an iterated prisoner's dilemma. Utilizing video-cued rating procedures, we examine several key questions about the functionality of facial expressions in social decision-making. First, we assess the extent to which emotion and expression regulation are accurately detected from dynamic facial expressions in interpersonal interactions. Second, we explore which facial cues are utilized to evaluate emotion and regulation information. Finally, we investigate the role of context in participants' emotion and regulation judgments. Results show that participants accurately perceive facial emotion and expression regulation, although they are better at recognizing emotions than regulation. Using automated expression analysis and stepwise regression, we constructed models that use action units from participant videos to predict their video-cued emotion and regulation ratings. We show that these models perform similarly and, in some cases, better than participants do. Moreover, these models demonstrate that game state information improves predictive accuracy, thus implying that context information is important in the evaluation of facial expressions.

Index Terms—Facial Expression, Emotion Regulation, Decision Making, Video-Cued Recall

I. INTRODUCTION

Facial expressions play a key role in social decision-making. Extensive research shows that facial expressions function at an interpersonal level to inform social decisions [1], [2]. In everyday interactions, facial expressions are frequently exchanged between social partners. Considering this exchange as a communication between a sender (displaying an expression) and a receiver (observing the expression), a sender's facial expressions signal valuable information about their motives and intentions to the receiver [2], [3]. In accord, research shows receivers use a sender's expression to infer intentions and likely future actions [4]. There have been many studies that use automatic expression analysis to make inferences about user states, such as boredom [5], depression [6] and rapport [7]. However, less is known about the extent to which receivers' perceptions of the sender's expressions are accurate.

Most studies examining interpersonal expression exchanges use prototypical emotion expressions; yet, in everyday interactions people have to evaluate spontaneous expressions, which often do not resemble prototypical expressions. Indeed, recognition of spontaneous expressions is fairly inaccurate (often less than 50%), although relatively reliable [8], [9]. To facilitate beneficial decisions in social interactions, receivers must make valid inferences about sender's intentions. However, given that senders are aware that receivers are making these judgments based on their facial expressions, they may strategically regulate their expressed emotions in order to influence receivers' decision making [10], [11].

Suppressing facial expressions is a commonly deployed emotion-regulation strategy [12]. Previous research demonstrates that when communicated emotions are regulated, the impact of a sender's facial expression on decision-making differs [13], [14]. For example, authentic anger increases concessions in negotiations, whereas inauthentic regulated anger reduces concessions [15]. Research also suggests that facial expressions play a key role in emotional lie detection [16]. While there has been much research on the impact of emotion regulation, little research has examined how people detect emotion regulation in interactions.

As well as facial cues, most interpersonal interactions provide additional contextual cues that likely help a receiver to decode a sender's expressions. Facial expressions do not occur in isolation but arise in response to specific events. Observers attend to these events when judging and extracting meaning from facial displays [8], [17]. Spontaneous expressions in everyday interactions are more ambiguous than prototypical expressions, and context provides additional information that helps disambiguate expressions [18], sometimes leading to changes in evaluation [19].

Building on this, we aimed to address limitations of previous research and investigate the following research questions:

RQ1a: What is the association between senders' ratings of displayed emotion and receivers' perception thereof?

RQ1b: What is the association between senders' ratings of their emotion regulation and receivers' perception thereof?

RQ2a: What facial features are used by senders when they rate their own displayed emotions?

RQ2b: What facial features are used by receivers when they rate the observed emotions of senders?

- To what extent are these ratings shaped by context?

RQ3a: What facial features are used by senders when they rate regulation of their own displayed emotion?

RQ3b: What facial features are used by receivers when they rate the observed emotion regulation of senders?

- To what extent are these ratings shaped by context?

To address these questions, we recruited a novel combination of a two-player computer-mediated iterated prisoners dilemma task, video-cued recall and automatic facial expression recognition. A prisoner's dilemma task is a two-player task where the payoffs of each player depend on the simultaneous choices of both players. In each round, both players choose to either cooperate or defect. The payoffs are such that the players must resolve the dilemma of cooperation leading to enhanced joint gains but possible exploitation. Players therefore need to be attentive to each other's facial expressions to evaluate intentions and likely actions. In fact, the iterated prisoner's dilemma has long been used for studying how people use facial expressions to resolve this dilemma [20]. In our study we utilize an iterated prisoner's dilemma, where the same players play multiple rounds with each other. Extant research shows expressed emotions shape joint decision-making in an iterated prisoner's dilemma [21]. Moreover, people are likely to engage in strategic emotion regulation in a social-dilemma task where the highest individual payoff requires exploitation of a cooperative partner.

To elucidate the relationship between a sender's actual emotion and regulation and a receiver's decoding of their expression, participants completed video-cued recall immediately after the iterated prisoner's dilemma game [22], [23]. This involved reviewing video-clips from the game of themselves and their partner, and making judgments about the valence of facial expressions and the extent to which expressions were regulated. This enabled assessment of the accuracy of emotion perception and of the extent to which regulation is detected. Combining video-cued recall with automatic analysis of facial expressions also allowed us to identify the facial cues that underpin accurate expression evaluation.

Affective computing methods have demonstrated that an understanding of emotional behavior yields practical benefits such as improving learning outcomes [24], job interview training [25] or achieving personal gains in negotiations [26]. Emotion regulation can change perceptions of behavior and as such there have been several computational models that incorporate emotion regulation and coping behaviors [27]–[29]. However, comparatively little work has explored how emotion regulation is actually perceived by people. By using data from the video-cued recall alongside automatic facial expression recognition, we computationally model the effects of emotion regulation. Using stepwise linear regression, we create models similar to Brunswik's lens model [30], which is commonly used to investigate correlations between displayed and perceived behaviors.

Combining these methods gives novel insight into these research questions by providing a ground truth to build accu-

TABLE I
PAYOFFS FOR EACH COMBINATION OF CHOICES IN THE ITERATED PRISONER'S DILEMMA GAME

		Participant B	
		cooperate	*defect*
Participant A	*cooperate*	A = 5, B = 5	A = 0, B = 10
	defect	A = 10, B = 0	A = 1, B = 1

rate models for decoding facial expressions in interpersonal interactions. Specifically, we examine strategic expression regulation in natural dyadic interactions and explore whether automatic analysis of expressions can provide insight into signals of emotion regulation in face-to-face interactions, and how these signals map onto perceptions of regulation.

II. EXPERIMENTAL SETUP

One-hundred people total (67 female, mean age=26.42 SD=7.44) who came from a community panel (69 participants) or responded to local advertisements (31 participants) participated in the study and were paid £10 for participation. The age, gender and ethnicity of community panel participants did not differ from those who responded to local advertisements. All participants provided written informed consent and the University's Ethics Committee approved the study. One person withdrew consent after participating and both their data and their partner's data were removed from analysis. Due to computer error, the video-cued recall data of one participant were only captured for the first three rounds. Two other dyads were excluded from analysis due to video playback issues during the video-cued recall procedure.

Participants were randomly paired and played 10 rounds of a computer-mediated iterated prisoner's dilemma game, based on the British TV show Golden Balls. We used the design and software developed by Hoegen et al. [31]. In each round, participants played for a set of lottery tickets and selected whether to cooperate or defect (associated payoffs are shown in table I). Participants were informed that each ticket earned would be entered into a lottery draw with one £100 prize and two prizes of £50, and as such, by obtaining more tickets they could increase their chances of winning a prize.

Each participant played in a separate room, seated in front of a computer with a webcam, they could see each other's webcam video stream while playing. In each round, participants selected whether to "Split" (cooperate) or "Steal" (defect) by clicking on the screen. Once both participants had made their choices, the outcome of the round was displayed to both participants. A ticket counter allowed participants to track their current scores. Figure 1 shows the game interface. The webcam video stream did not include audio and participants were instructed not to talk to each other or to use hand gestures to communicate, but that they were free to use facial expressions.

After the game, participants completed a video-cued recall procedure, in which they reviewed video clips and evaluated the valence and regulation of both their own and their partner's
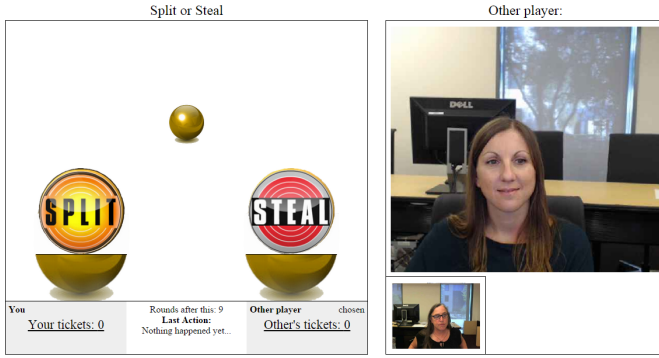
Fig. 1. The iterated prisoner's dilemma game interface. Participants make choices by selecting either the "split" or "steal" ball on each round. They can also see the number of tickets won by themselves and by the other player as well as the other player's webcam video (big window) and their own webcam video (small video). (Image taken from Hoegen et al. [31])

expressions [22], [23]. By using the specific timing of the participants' decisions as logged in the database, we could automatically generate video clips of specific events. For every participant we created sets of two five-second video clips for all 10 rounds played. The first clip ('decision' clip) in this set captured the decision of the participant while opting whether to "Split" or "Steal". The second clip ('outcome' clip) captured both participants for the five seconds following the reveal of the round outcome. For 'decision' clips, participants were instructed to rate expressions that occurred immediately before the "Split" or "Steal" choice was made. For 'outcome' clips, participants were instructed to rate expressions that occurred immediately after both players' decisions and the outcome of the round were revealed.

The video clips were presented in chronological order; for each 'decision' and 'outcome' event participants first rated the video of themselves, and then the video of their partner. For the 'outcome' clip, the outcome of the round was displayed to participants when they made their ratings to provide information about the context in which the expression was displayed. In this paper, we examine ratings of the 'outcome' clips because we are interested in how receivers decode senders' facial responses to specific game events.

III. MEASURES

We collected behavioral data about participants' decisions during the games and two sets of facial expression annotations; one created by the participants themselves during the video-cued recall, and one automatically extracted from the video using facial expression detection software.

A. Behavioral Data

The actions of participants throughout the game were logged in a database as timestamped data. The database included information about the decision to cooperate or defect, and other important game events, such as the moment when the results of the joint decision was revealed.

From the joint decision made by each pair of participants in each round, we defined a set of 'game states'. When both participants chose to cooperate, this results in a "joint cooperation" or CC-state, while both participants defecting results in a "joint defection" or DD-state. When one participant defects while the other participant cooperates this results in an unequal outcome; the participant who defected and as such exploits their partner is in the DC-state, while the participant who was exploited is in the CD-state. By defining these specific states, we are able to investigate whether contextual information such as these game states affect the displays and perceivers' ratings of them.

B. Video-Cued Annotations

Using a video-cued recall procedure, we obtained ratings for each of the video clips of the participant, with ratings by both the participant themselves (sender) and their partner (receiver. For each video clip participants rated the valence of the displayed expression on a scale from -50 very negative to +50 very positive. Following this, participants rated the videos on both positive and negative expression regulation using a -50 (suppressed) to +50 (exaggerated) scale.

C. Automatic Facial Expression Annotations

After the study was completed, we extracted facial expression ratings automatically from the webcam videos using a commercial software system based on CERT [32]. Using this system, we extracted Action Units (AUs) as defined by the Facial Action Coding System (FACS) from each participant's video on a frame-by-frame basis. For each frame the system assigns each AU an 'evidence' value. This value represents the odds of the AU being activated, with positive values indicating that a particular AU is activated and negative values indicate that the AU is not active. The system uses a logarithmic scale base 10 scale for scoring AUs (i.e. a score of 1 indicates an AU is 10 times more likely to be rated as active by a human coder than inactive, a score of 2 indicates the AU is 100 times more likely to be rated as active, etc.). We then computed overall AU evidence scores for each of the video clips that were rated by the participants during the video-cued recall. This was done by finding data for the set of frames that were shown during the clip, and then calculating the mean value for each AU using this data.

IV. RESULTS

A. RQ1: Correlations between sender and observer ratings

To explore whether there is an association between senders' and observers' ratings of displayed expressions (RQ1), we computed a series of correlation coefficients.

We first addressed RQ1 for expressed emotion ratings (RQ1a). We found a significant positive correlation between senders' and observers' ratings of displayed emotions (r = .40, n = 909, p < .001), but the effect size was relatively small.

Next, we investigated senders' and observers' positive emotion regulation ratings (RQ1b). Again we found a significant

positive correlation (r = .13, n = 909, p <.001), but the effect size was even smaller than for RQ1a.

These findings suggest that it is quite difficult for people to accurately decode each other's expressions and even harder for them to determine the extent to which regulation is taking place. It is also worth noting that participants did not appear to improve in recognizing expressions over time, as there was no significant correlation between accuracy of participants' expression recognition (a measure based on the absolute distance between the sender's and observer's ratings) and the round number (r = -0.02, n = 909, p = .518). Participants that were more accurate at recognizing expressions also were not capable of using this information to improve their scores, as shown by a second correlation between the average accuracy of a participant's expression recognition and their score (r = .10, n =82, p = .363).

B. Regression model approach

Following RQ1, we investigated the research questions about the facial features used by people when rating expressed emotion and emotion regulation (RQ2 and 3). We created several forward stepwise multiple linear regression models using the AUs that were automatically extracted from the videos as predictors of rated emotion and emotion regulation. For all of these stepwise regression models, the enter criterion was a p-value below 0.05 and the removal criterion was a p-value higher than 0.1. By tracking which AUs are selected by the stepwise selection process, we can infer which AUs account for more variance in our data.

Using these models, we also investigated the secondary part of RQ2 and 3 (To what extent are participants' ratings shaped by context?). This involved adding a predictor variable representing context to the regression equation. In this case, we define the context of an iterated prisoner's dilemma as the 'game state' (CC, DD, CD and DC) defined in the measures section of this paper. Game state was coded as a categorical variable, using dummy coding for each possible game state. By either including or excluding the dummy code for a game state in the regression equation, we were able to add or remove the context of the situation to our models. As such, we created two linear regression models for each RQ, one that included the game state and one that did not. This allowed us to directly compare the performance of these two models and gain understanding of the importance of context, by comparing the variance in the data explained by each model.

C. RQ2: Expressed emotion

Using the above approach, we first investigated which features are used by senders when they rate their own expressed emotions (RQ2a). Using forward stepwise regression with the aforementioned criteria, we obtained a significant model (F(6,902) = 14.50, p <.001 with an R^2 of .088), that used AUs 9 (Nose Wrinkler), 12 (Lip Corner Puller), 14 (Dimpler), 25 (Lips Part), 26 (Jaw Drop) and 43 (Eyes Closed). We can compare this performance to that of receivers' ratings based on observation of the same expressions. The correlation between
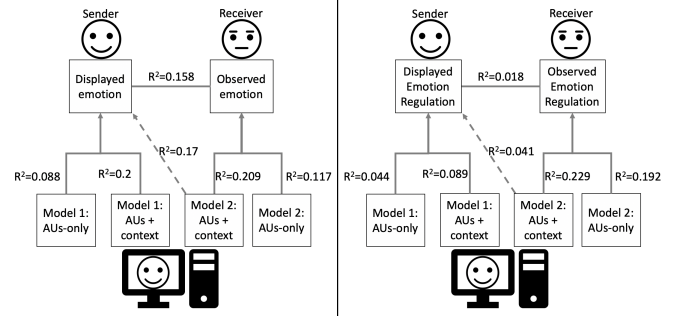

Fig. 2. Performance in R^2 of the models in predicting displayed and observed emotion (RQ2) on the left, and displayed and observed emotion regulation on the right (RQ3). The top line in the diagrams represents the R^2 of the correlation between participant ratings (RQ1)

the sender's and receiver's score had an r of around .4, and the R^2 is .158. In other words, the model based on automatically coded expressions explains less of the variance in senders' ratings than observers' judgements are able to do.

Next, we additionally included game state as a predictor of senders' emotion ratings, in order to investigate the impact of context. We found a significant multiple linear regression model (F(8,900) = 28.10, p <.001 with an R^2 of .200) that used as predictors AUs 4 (Brow Lowerer), 10 (Upper Lip Raiser), 12, 26 and 43 as well as the game state, encoded as three dummy variables (see Figure 3 left). Adding game state to the model meant that it explained slightly more of the variance (R^2 of .200) than observers' ratings were able to do (R^2 of .158).

Following this, we investigated which features are used by receivers when they rate the observed expressed emotions of senders (RQ2b). Using forward stepwise regression again, we obtained a significant multiple linear regression model (F(7,901) = 17.00, p <.001 and an R^2 of .117), that included AUs 2 (Outer Brow Raiser), 9, 12, 15 (Lip Corner Depressor), 17 (Chin Raiser), 26 and 28 (Lip Suck) as predictors.

In a further analysis, we included game state in the regression equation. We found a significant multiple linear regression model (F(8,900) = 29.60, p <.001, R^2 = .209) using AUs 9, 12, 14, 17, 26 and game state as predictors. Once again, game state contributed significantly to model's performance, thus suggesting that participants use the context to inform their judgments of observed expressions. As an additional test, we compared this model to ratings made by senders themselves. We did this by using the predictors included in the final multiple regression model to predict senders' ratings of the displayed emotion. This resulted in an R^2 of .170. This R^2 is higher than the association that we found between receivers' judgements of senders' expressions and senders' own ratings, indicating that the model captured important variance in how people rate their partner.

Summarizing our findings for RQ2 (see Figure 2 left), we created several regression models that were able to reliably predict expressed emotion. However, the absolute accuracy of prediction remained relatively low. Regression models were

able to predict observed emotion ratings at a slightly higher level, with the model that included both AUs and context as predictors achieving an R^2 of .209.

D. RQ3: Emotion regulation

Lastly, we investigated which features are used by senders when they rate their own displayed regulation of emotion (RQ3a). Taking the same approach as for RQ2 we performed a stepwise regression without game state, then attempted to improve upon this model by including game state. Our stepwise linear regression without game state resulted in a significant model ($F_{(3,905)} = 13.90$, $p < .001$ with R^2 of .044) using AUs 4, 12 and 20 (Lip Stretcher) as predictors. Although this model performed slightly better than observers (R^2 of .018), the variance explained by the model remains low.

Next, we included context information in the stepwise regression, which resulted in another significant linear regression model ($F_{(6,902)} = 14,6$, $p < .001$ with an R^2 of .089), using AU4, 12 and 20, and game state as predictors. Once again, we were able to increase the amount of variance explained by the model by including game state information, but even here it remains low. This shows that predicting senders' video-cued ratings of their own emotion regulation is a hard task.

Finally, we investigated which features are used by receivers when they rate observed emotion regulation (RQ3b). The first stepwise regression without game state resulted in a significant model ($F_{(4,904)} = 53.7$, $p < .001$ and R^2 of .192) using AU10, 12, 17 and 43 as predictors.

In the second stepwise regression (for RQ3b), we included game states again. This yielded a significant regression model ($F_{(8,900)} = 33.40$, $p < .001$ and an R^2 of .229) that used AUs 5 (Upper Lid Raiser), 10, 12, 14 and 43, and the game states as predictors. This improved on the R^2 of the previous model. Similarly to the model predicting observed emotions (RQ2b), we used the model to predict participants' emotion regulation ratings of their own videos. This resulted in an R^2 of .041, similar to the R^2 based on the correlation between senders' and receivers' ratings of emotion regulation. This provides confidence that we are capturing some of the variance in participants' ratings of emotion regulation.

Similarly to RQ2, our findings for RQ3 (see Figure 2 right), showed that, although we found significant regression models, the effect sizes were very small for models of *displayed* emotion regulation. The effect sizes for *observed* emotion regulation were quite higher: without including context we achieved an R^2 of .19 and we further improved upon this by including game state as additional information.

V. Discussion

This research explored accuracy of expression and regulation evaluations in face-to-face interactions, and modelled face and contextual cues associated with these evaluations. By utilizing a novel combination of participants' video-cued ratings of both displayed and observed expressions and auto-mated expression analysis in an iterated prisoner's dilemma, we obtained a clearer picture of how people encode and decode
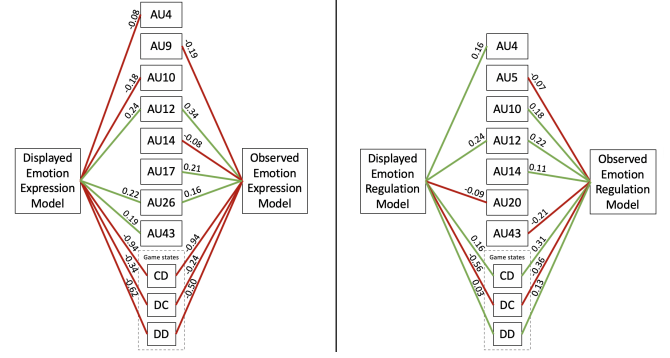

Fig. 3. Emotion expression lens diagram of features used by multiple linear regression models (game state was encoded in these models using dummy coding). The left side shows the models for expressed emotion, the right shows models for emotion regulation. The coefficients of the variables used by the models are given next to the lines.

facial expressions in context. Having separate ratings for both displayed and observed expressions allowed exploration of how accurate people were in their judgments of observed expressions and regulation. For expressed emotion, there was a significant association between what senders displayed and receivers observed. There was also a significant association between senders' and receivers' ratings for emotion regulation, but the effect was smaller. This relatively small association suggests that detecting emotion regulation is a more difficult task than detecting expressions themselves. The accuracy of expression detection was fairly low, however this is in line with previous research on spontaneous facial expressions [8].

We also investigated how facial AUs contributed to encoding and decoding processes. By using automatic facial expression detection, we assessed how individual AUs relate to senders' and receivers' judgments of emotion and emotion regulation. We performed several stepwise regression analyses, using AU measurements that were automatically extracted from the videos. We also compared regression models that only use AUs as predictors with those that also included contextual information in the form of dummy-coded game state. Overall, these models highlight the importance of contextual information in the evaluation of facial expressions. Results indicate that game-state information improves predictive accuracy and support the view that facial expressions communicate a sender's response to a specific event [33], meaning that contextual cues disambiguate senders' expressions and help receivers to interpret senders' emotion and regulation [18].

Across the regression models, several key AUs were identified that predict senders' and receivers' expression judgements. The AUs identified in the regression models respond to mouth and eye/brow areas, which are generally recognized as key regions for emotion expression identification. For example, AU12 is typically associated with happy expressions, while AU10 is associated with disgust [34]. Thus, the models indicate that participants appear to utilize AUs which differentiate between expressions of emotion to decode valence and regulation information from facial expressions.

For the regression model predicting emotion judgments, AU12 (an important component of smiles) and AU26 (associated with surprise) were positively associated with both senders' (displayed) and receivers' (observed) ratings (Figure 3 left). In the observed model, the negative coefficients of AU9 indicate that it has an inverse relationship with positive valence (expressions are rated on a negative to positive scale). The same goes for AU14. The displayed expression models used AU10 (which commonly co-occurs with AU9) and AU4, which both have an inverse relationship with positive expressions. The roles of AU43 for the displayed model and AU17 for the observed model are less clear.

Game state appears quite important for both displayed and observed expression, as adding this predictor to the model improved its predictive power (the game states were dummy-coded so CC is absent in the figure and the coefficients of the other game states are thus related to this state: i.e. a negative coefficient indicates the state has an inverse relationship with positive expressions). The CD-state specifically (participant exploited by other) has a high negative coefficient value, indicating that participants tend to rate both their displayed expressions and observe others' expressions more negatively in states associated with negative outcomes.

The regression models predicting both displayed and observed emotion regulation ratings (Figure 3 right) again selects AU12, highlighting its importance for both encoding and decoding expression regulation. However, the interpretation of other AUs is less clear. The observed model selects AU10 and 14, and the displayed model selects AU4, which tend to be associated with negative expressions. This suggests that when evaluating upregulation of positive expressions, participants rely on action unis that are typically associated with negative displays. The negative coefficients in the observed model are AU5 and 43, while the displayed model includes AU20. Similar to the expressed emotion models, adding game state information to these models improved their predictive power. In the case of regulation, it appears that the DC-state (exploits other) is related to suppression, while the CD and DD states are not. Looking at the mean scores of the senders' regulation ratings, we can see that participants overall tend to downregulate their positive expressions (M = -3.51). However participants downregulate more in the DC-state (M = -10.35) than they do in CC (M = -3.28) or DD (M = -1.63), or the CD-state (M = 0.81). As such, the model appears to be sensitive to the fact that participants tend to report more downregulation after defecting against a cooperative partner. These patterns of regulation make sense in a social context; when participants exploit someone, they downregulate their positive expressions to reduce negative reactions to their defection, whereas when participants are defected upon, they perform less downregulation on their positive expressions to encourage future cooperation. Indeed, previous research indicates that negative expressions reduce joint gains over time in a prisoner's dilemma task [21].

Overall, the variance explained by the displayed models is low and we were relatively more successful in building models for observed ratings. However, compared to other work that attempts to find associations between expressions and judgments, our models explains similar amounts of variance. By adding context to these models, we were able to explain further variance. As such it appears that people do in fact use contextual information when they rate emotions and emotion regulation [18], [33].

There are some limitations to the present study that we hope to address in future research. In this study, the multiple linear regression models use individual AUs as separate predictors. However, because these AUs occur together they also probably interact with each other. These interactions are currently not captured in our models.

While the video-cued recall allowed us to gain unique and interesting insights into facial expression decoding processes, it may also have introduced some degree of self-serving bias to ratings, especially since all our data came from people who were directly involved in the interaction under investigation. In order to gain a better understanding of the influence of context on these ratings, future research should also obtain judgments from external raters who were not involved in the game.

Although the models including context improved prediction accuracy, a fair amount of variance remained unexplained by our regression analyses. In addition to context, future research could examine other factors that may account for some of variance our models do not yet capture.

The models presented in this paper use mean values of the AUs displayed in the video, but encoding data in this fashion excludes information about features such as timing and dynamics. For example, previous research shows that timing information is important in expression evaluation [35], and there might be a difference in onset of specific expressions, especially when they are regulated.

For this analysis, we examined the valence of expressions rather than specific discrete expressions. In future work, more fine-grained analyses of specific displayed emotions may further explain encoding and decoding of facial expressions in social-decision making.

Finally, in this paper we only analyzed participants' ratings of the 'outcome' clips, since we were interested in responses to specific contexts. However, it is likely that participants also engaged in some strategic regulation in anticipation of their planned behavior in order to influence the other player's decision. Future work will examine expressions and regulation from the 'decision' clips as well as examine participants' reported motives for regulation to further elucidate the extent to which expression regulation is strategic.

We performed regression models on video-cued ratings of facial expressions and expression regulation from an iterated prisoner's dilemma game. We showed that people are able to accurately evaluate expressions and detect regulation, although they are not particularly good at decoding spontaneous expressions and are worse at detecting regulation. Furthermore, we revealed specific actions units that are utilized in these processes and demonstrate the importance of context information in the evaluation of expressions.

REFERENCES

[1] B. Parkinson, "Emotions are social," *British journal of psychology*, vol. 87, no. 4, pp. 663–683, 1996.

[2] G. A. Van Kleef, C. K. De Dreu, and A. S. Manstead, "An interpersonal approach to emotion in social decision making: The emotions as social information model," in *Advances in experimental social psychology*. Elsevier, 2010, vol. 42, pp. 45–96.

[3] B. Parkinson, "Do facial movements express emotions or communicate motives?" *Personality and Social Psychology Review*, vol. 9, no. 4, pp. 278–311, 2005.

[4] C. de Melo, J. Gratch, P. Carnevale, and S. Read, "Reverse appraisal: The importance of appraisals for the effect of emotion displays on peoples decision making in a social dilemma," in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 34, no. 34, 2012.

[5] R. S. Baker, S. K. D'Mello, M. M. T. Rodrigo, and A. C. Graesser, "Better to be frustrated than bored: The incidence, persistence, and impact of learners cognitive–affective states during interactions with three different computer-based learning environments," *International Journal of Human-Computer Studies*, vol. 68, no. 4, pp. 223–241, 2010.

[6] S. Scherer, G. Stratou, G. Lucas, M. Mahmoud, J. Boberg, J. Gratch, L.-P. Morency *et al.*, "Automatic audiovisual behavior descriptors for psychological disorder analysis," *Image and Vision Computing*, vol. 32, no. 10, pp. 648–658, 2014.

[7] N. Wang and J. Gratch, "Rapport and facial expression," in *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*. IEEE, 2009, pp. 1–6.

[8] U. Hess and S. Hareli, "The influence of context on emotion recognition in humans," in *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, vol. 3. IEEE, 2015, pp. 1–6.

[9] I. Hühnel, M. Fölster, K. Werheid, and U. Hess, "Empathic reactions of younger and older adults: no age related decline in affective responding," *Journal of Experimental Social Psychology*, vol. 50, pp. 136–143, 2014.

[10] B. Barry, "The tactical use of emotion in negotiation," *Research on negotiation in organizations*, vol. 7, pp. 93–124, 1999.

[11] M. S. Clark, S. P. Pataki, and V. H. Carver, "Some thoughts and findings on self-presentation of emotions in relationships," *Knowledge structures in close relationships: A social psychological approach*, vol. 247274, 1996.

[12] J. J. Gross and O. P. John, "Revealing feelings: facets of emotional expressivity in self-reports, peer ratings, and behavior." *Journal of personality and social psychology*, vol. 72, no. 2, p. 435, 1997.

[13] B. Parkinson, N. Phiri, and G. Simons, "Bursting with anxiety: Adult social referencing in an interpersonal balloon analogue risk task (bart)." *Emotion*, vol. 12, no. 4, p. 817, 2012.

[14] D. M. Shore and B. Parkinson, "Interpersonal effects of strategic and spontaneous guilt communication in trust games," *Cognition and Emotion*, vol. 32, no. 6, pp. 1382–1390, 2018.

[15] S. Côté, I. Hideg, and G. A. Van Kleef, "The consequences of faking anger in negotiations," *Journal of Experimental Social Psychology*, vol. 49, no. 3, pp. 453–463, 2013.

[16] G. Warren, E. Schertler, and P. Bull, "Detecting deception from emotional and unemotional cues," *Journal of Nonverbal Behavior*, vol. 33, no. 1, pp. 59–69, 2009.

[17] Y. Wu, C. L. Baker, J. B. Tenenbaum, and L. E. Schulz, "Rational inference of beliefs and desires from emotional expressions," *Cognitive science*, vol. 42, no. 3, pp. 850–884, 2018.

[18] H. Aviezer and R. R. Hassin, "Inherently ambiguos: an argument for contextualized emotion perception," in *The science of facial expression*, 2017, pp. 333–352.

[19] H. Aviezer, R. R. Hassin, J. Ryan, C. Grady, J. Susskind, A. Anderson, M. Moscovitch, and S. Bentin, "Angry, disgusted, or afraid? studies on the malleability of emotion perception," *Psychological science*, vol. 19, no. 7, pp. 724–732, 2008.

[20] R. H. Frank, *Passions within reason: The strategic role of the emotions*. WW Norton & Co, 1988.

[21] G. Stratou, R. Hoegen, G. Lucas, and J. Gratch, "Emotional signaling in a social dilemma: An automatic analysis," in *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2015, pp. 180–186.

[22] J. M. Gottman and R. W. Levenson, "A valid procedure for obtaining self-report of affect in marital interaction." *Journal of consulting and clinical psychology*, vol. 53, no. 2, p. 151, 1985.

[23] B. Parkinson, G. Simons, and K. Niven, "Sharing concerns: Interpersonal worry regulation in romantic couples." *Emotion*, vol. 16, no. 4, p. 449, 2016.

[24] N. Wang, W. L. Johnson, R. E. Mayer, P. Rizzo, E. Shaw, and H. Collins, "The politeness effect: Pedagogical agents and learning outcomes," *International journal of human-computer studies*, vol. 66, no. 2, pp. 98–112, 2008.

[25] K. Anderson, E. André, T. Baur, S. Bernardini, M. Chollet, E. Chryssafidou, I. Damian, C. Ennis, A. Egges, P. Gebhard *et al.*, "The tardis framework: intelligent virtual agents for social coaching in job interviews," in *International Conference on Advances in Computer Entertainment Technology*. Springer, 2013, pp. 476–491.

[26] J. Gratch, Z. Nazari, and E. Johnson, "The misrepresentation game: How to win at negotiation while seeming like a nice guy," in *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 728–737.

[27] S. W. McQuiggan, J. L. Robison, R. Phillips, and J. C. Lester, "Modeling parallel and reactive empathy in virtual agents: An inductive approach," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 167–174.

[28] T. Bosse, M. Pontier, G. F. Siddiqui, and J. Treur, "Incorporating emotion regulation into virtual stories," in *International Workshop on Intelligent Virtual Agents*. Springer, 2007, pp. 339–347.

[29] S. Marsella, J. Gratch, N. Wang, and B. Stankovic, "Assessing the validity of a computational model of emotional coping," in *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*. IEEE, 2009, pp. 1–8.

[30] R. Gifford, "A lens-mapping framework for understanding the encoding and decoding of interpersonal dispositions in nonverbal behavior." *Journal of personality and social psychology*, vol. 66, no. 2, p. 398, 1994.

[31] R. Hoegen, G. Stratou, G. M. Lucas, and J. Gratch, "Comparing behavior towards humans and virtual humans in a social dilemma," in *International Conference on Intelligent Virtual Agents*. Springer, 2015, pp. 452–460.

[32] G. Littlewort, J. Whitehill, T. Wu, I. Fasel, M. Frank, J. Movellan, and M. Bartlett, "The computer expression recognition toolbox (cert)," in *Face and gesture 2011*. IEEE, 2011, pp. 298–305.

[33] C. M. de Melo, P. J. Carnevale, S. J. Read, and J. Gratch, "Reading peoples minds from emotion expressions in interdependent decision making." *Journal of personality and social psychology*, vol. 106, no. 1, p. 73, 2014.

[34] P. Ekman, W. V. Friesen, and J. C. Hager, "The facial action coding system, research nexus ebook," *Salt Lake City, UT*, 2002.

[35] E. Krumhuber, A. S. Manstead, D. Cosker, D. Marshall, P. L. Rosin, and A. Kappas, "Facial dynamics as indicators of trustworthiness and cooperative behavior." *Emotion*, vol. 7, no. 4, p. 730, 2007.