Towards Emotional Intelligence in Social Robots Designed for Children

De'Aira Bryant
dbryant@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

ABSTRACT

Social robots are robots designed to interact and communicate directly with humans, following traditional social norms. However, many of these current robots operate in discrete settings with predefined expectations for specific social interactions. In order for these machines to operate in the real world, they must be capable of understanding the multiple factors that contribute to human-human interaction. One such factor is emotional intelligence. Emotional intelligence allows one to consider the emotional state of another in order to motivate, plan, and achieve one's desires. One common method of analyzing the emotional state of an individual involves analyzing the emotion displayed on their face. Several artificial intelligence (AI) systems have been developed to conduct this task. These systems are often classifiers trained using a variety of machine learning techniques which require large amounts of training data. As such, they are susceptible to biases that may appear during performance analyses due to disproportions existing in training datasets. Children, in particular, are often less represented in the primary datasets of annotated faces used for training such emotion classification systems. This work seeks to first analyze the extent of these performance differences in commercial systems, then to present new computational techniques that work to mitigate some of the effects of minimal representation in datasets, and to finally present a social robot which utilizes an improved emotional AI to interact with children in various scenarios where emotional intelligence is key to successful human-robot interaction.

ACM Reference Format:

De'Aira Bryant. 2019. Towards Emotional Intelligence in Social Robots Designed for Children. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES '19)*, January 27–28, 2019, Honolulu, HI, USA. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3306618.3314319>

1 ANALYZING EMOTION RECOGNITION SOFTWARE

Emotion classification as discerned from facial configurations is a popular method used to provide emotional intelligence to a variety of smart AI systems. With the intent of utilizing this technology in use cases involving children, my initial desire was to find a pre-existing software that would serve this purpose. I found that several commercial systems were available for integration into other smart

systems. With so many options available, I sought to evaluate the performance of these commercial emotion recognition systems on my target demographic. I utilized a total of five datasets that comprised over 3,000 images of annotated children faces to assess the performance of each of the recognition systems. I stored the ground truth labels of the images, processed each of the images from each of the datasets through each of the emotion recognition systems, normalize the results for an equal comparison, determined the emotion label with the highest confidence value through a maximization function, and finally compared the system-produced predicted label to the ground truth labels to store the results. Each system was assessed in terms of matching score (true positive rate), positive-predictive value, and failure to compute (FTC) rate.

This analysis revealed that these systems performed worse when labeling the images of children faces when compared to human ratings on the same images. Additionally, when compared to similar analyses conducted with adult datasets, these ratings were also lower. This analysis gave evidence that these commercial systems are premature in their use for children.

Yet, there is evidence of such commercial systems being used in schools to target issues surrounding the security and surveillance of students (Vanderklippe 2018). I discuss these results, their implications, and potential solutions in a paper submitted to the AIES 2019 main conference. Additional analysis of the data has also revealed differences in recognition when assessing the children's facial images by gender and ethnicity. These differences are consistent with other results found in similar studies examining the performance of AI systems on groups that are often less represented in datasets (Buolamwini and Gebru 2018).

2 IMPROVING PERFORMANCE FOR CHILDREN

The AIES 2019 submission also introduces a method for assigning a diversity rating to each image dataset based on the variety of images within the dataset. This rating considers participant age, gender, ethnicity, gaze directions, geographic region(s) of collection, clothing, pose, and number of discrete emotion classes. The rating can be extended to apply a diversity rating to each individual image. These ratings can then be used as weights when training a machine learning classifier. The weights will allow the learner to place higher value on images that have higher diversity ratings, as they are rarer to find in large datasets. My current work involves utilizing this technique to assess whether the weighting of training images can better the overall performance of a machine learning classifier. The images used for training will be the same images used to analyze the commercial systems in my initial work. As these images are comprised of entirely children faces, I hypothesize that

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

AIES '19, January 27–28, 2019, Honolulu, HI, USA

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6324-2/19/01.

<https://doi.org/10.1145/3306618.3314319>

a learned classifier using these images will outperform the commercial systems from the initial comparative work on new children's faces.

By training multiple classifiers with different weights on diversity attributes, several analyses can be made to determine the reliance of overall performance on dataset compilation. This work will shed further light on the impact that low or nonexistent representation of a particular group within a training set may have on the performance of a classifier. I hypothesize that a more balanced dataset in terms of diversity will yield more consistent performance across the board for subgroups of the overall population. The next phase of this overall work will seek to implement this technique on several embodied platforms that will interact with a diverse set of children in terms of gender, ethnicity, and ability.

3 PREVIOUS HRI STUDIES USING SOCIAL ROBOTS

In prior work, a social robot is used to interact with adults and children in a rehabilitative game by providing encouraging and corrective feedback. The system targets upper extremity function by having users pop virtual bubbles (Xu, Bryant and Howard 2018). A 3D model for movement time is used to calculate the target movement time for each respective user in a phase used for calibration. This system also utilizes a script for the robot to interact with the user. The feedback is generated within the gaming system and then sent to the robot via TCP connection. As such, in these prior studies, the robot did not take in any input from the direct interaction with the user.

We also evaluate this system in terms of user trust, another key component in social interaction (Xu, Bryant and Howard 2018). Survey results indicate trust in a therapy intervention coupled with a robotic agent is comparable to that of trust in an intervention coupled with a human agent.

The AI developed throughout this work has the goal of being implemented onto an embodied robotic platform that will interact directly with humans. As such, the chosen robotic embodiment may also impact the overall experience for a user of the technology. One potential factor that I also investigated is the impact that robot "skin" color could have on the human perception of robotic emotional expression. In this HRI study presented at the 2018 FAT-ML conference, the Darwin mini robot expresses several emotions as participants are instructed to classify the emotion and its perceived intensity (Ogunyale, Bryant and Howard 2018). By eliminating stereotype priming from the methodology, we found no substantial differences in the perception of the robot's expressions of happy and angry on black and white robots. Current work consists of using a variety of robot "skin" colors on several robotic embodiments to determine if this result remains consistent.

4 EMOTIONALLY INTELLIGENT SOCIAL ROBOTS

The additional user responses from these collective HRI works suggested a need for a more interactive experience for the users. In future work, the emotionally intelligent AI system described to improve performance for children will be integrated onto a variety of robotic embodiments (NAO, Cozmo, and Pepper). User studies

will be conducted to better analyze the performance of the integrated emotional intelligence when compared to a scripted robot interaction, a human interaction, and a control no-robot interaction. These various conditions will validate the use of diversity weighting as a method to improve performance for applications designed specifically for special subsets of the population. Overall, this work will contribute novel approaches to child-robot interaction through emotionally intelligent AI integrations.

REFERENCES

- [1] Buolamwini, J., & Geburu, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In Conference on Fairness, Accountability and Transparency.
- [2] Ogunyale, T., Bryant, D., & Howard, A. 2018. Does Removing Stereotype Priming Remove Bias? A Pilot Human-Robot Interaction Study. In Proceedings of the 2018 International Workshop on Fairness, Accountability, and Transparency in Machine Learning (FAT-ML).
- [3] Vanderklippe, N. 2018. In China, Classroom Cameras Scan Students' Faces for Emotion, Stroking Fears of New Form of State Monitoring. <https://www.theglobeandmail.com/world/article-in-china-classroom-cameras-scan-student-faces-for-emotion-stoking/>.
- [4] Xu, J., Bryant, D., & Howard, A. 2018. Would You Trust a Robot Therapist? Validating the Equivalency of Trust in Human-Robot Healthcare Scenarios. In The 27th IEEE International Symposium On Robot And Human Interactive Communication (Ro-Man).
- [5] Xu, J., Bryant, D., Chen, Y. P., & Howard, A. 2018. Robot therapist versus human therapist: Evaluating the effect of corrective feedback on human motor performance. In Proceedings of the 2018 IEEE International Symposium on Medical Robotics (ISMR). pp. 1-6.