

MIT Open Access Articles

Toward Personalized Affect-Aware SAR Tutors in Long-Term Interventions for Children with Autism

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Shi, Zhonghao, Groechel, Thomas, Jain, Shomik, Chima, Kourtney, Rudovic, Ognjen (Oggi) et al. 2022. "Toward Personalized Affect-Aware SAR Tutors in Long-Term Interventions for Children with Autism." Transactions on Human-Robot Interaction.

As Published: <https://doi.org/10.1145/3526111>

Publisher: ACM

Persistent URL: <https://hdl.handle.net/1721.1/146429>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Toward Personalized Affect-Aware Socially Assistive Robot Tutors for Long-Term Interventions with Children with Autism

ZHONGHAO SHI, THOMAS R. GROECHEL, SHOMIK JAIN, and
KOURTNEY CHIMA, University of Southern California, Los Angeles, CA
OGNJEN (OGGI) RUDOVIC, Massachusetts Institute of Technology, Cambridge, MA
MAJA J. MATARIĆ, University of Southern California, Los Angeles, CA

Affect-aware socially assistive robotics (SAR) has shown great potential for augmenting interventions for children with autism spectrum disorders (ASD). However, current SAR cannot yet perceive the unique and diverse set of atypical cognitive-affective behaviors from children with ASD in an automatic and personalized fashion in long-term (multi-session) real-world interactions. To bridge this gap, this work designed and validated personalized models of arousal and valence for children with ASD using a multi-session in-home dataset of SAR interventions. By training machine learning (ML) algorithms with supervised domain adaptation (s-DA), the personalized models were able to tradeoff between the limited individual data and the more abundant less personal data pooled from other study participants. We evaluated the effects of personalization on a long-term multimodal dataset consisting of four children with ASD with a total of 19 sessions, and derived inter-rater reliability (IR) scores for binary arousal (IR = 83%) and valence (IR = 81%) labels between human annotators. Our results show that personalized Gradient Boosted Decision Trees (XGBoost) models with s-DA outperformed two non-personalized individualized and generic model baselines not only on the weighted average of all sessions, but also statistically ($p < .05$) across individual sessions. This work paves the way for the development of personalized autonomous SAR systems tailored toward individuals with atypical cognitive-affective and socio-emotional needs.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**; • **Applied computing** → **Computer-assisted instruction**; • **Computer systems organization** → Robotics;

Additional Key Words and Phrases: Human-robot interaction, socially assistive robotics, autism spectrum disorders, personalized machine learning, affective computing

ACM Reference format:

Zhonghao Shi, Thomas R. Groechel, Shomik Jain, Kourtney Chima, Ognjen (Oggi) Rudovic, and Maja J. Matarić. 2022. Toward Personalized Affect-Aware Socially Assistive Robot Tutors for Long-Term Interventions with Children with Autism. *ACM Trans. Hum.-Robot Interact.* 11, 4, Article 39 (September 2022), 28 pages. <https://doi.org/10.1145/3526111>

This research was supported by the National Science Foundation Expedition in Computing Grant NSF IIS-1139148. The work of Ognjen Rudovic was funded by European Union H2020, Marie Curie Action - Individual Fellowship 2016-2019 (EngageMe 701236).

Authors' addresses: Z. Shi, T. R. Groechel, S. Jain, K. Chima, and M. J. Matarić, University of Southern California, Los Angeles, CA, 90007; emails: {zhonghas, groechel}@usc.edu, shomik.jain@alumni.usc.edu, {kchima, mataric}@usc.edu; O. (Oggi) Rudovic, Massachusetts Institute of Technology, Cambridge, MA, 02139; email: orudovic@mit.edu.

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

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-9522/2022/09-ART39 \$15.00

<https://doi.org/10.1145/3526111>

1 INTRODUCTION

Within **human-robot interaction (HRI)**, the field of **socially assistive robotics (SAR)** has emerged at the intersection of assistive robotics and socially interactive robotics [22]. Its central focus is to provide effective assistance and interventions through intelligent interactions with users [22, 37], while also respecting socio-emotional needs of various users. As research in SAR is shifting from constrained laboratory settings toward real-world environments such as in-home interventions, it is critical for SAR systems to be able to perceive the user's emotional (i.e., affective) cues in order to provide naturalistic affect-aware interactions and interventions [6, 41]. Toward this end, recent research has shown that affect-aware SAR tutors have the potential to address the imbalance between cognitive and affective awareness found in contemporary robot-assisted teaching systems [52]. However, existing SAR systems still lack the ability to elicit, perceive, and appropriately respond to user affect in a personalized fashion so that interventions can be tailored towards individual cognitive and affective needs along the learning process [56].

Both the promise and the challenges of SAR for personalized affect awareness are particularly amplified in the context of children with **autism spectrum disorders (ASD)**. ASD affects 1 in 64 children in the United States and has a male-to-female ratio of 4:1 [14]. Traditionally, therapists design interventions by using static toys as tools to induce open-ended and engaging interactions [31]. More recently, research has shown success in adopting SAR tutors as a means of providing more effective interventions for children with ASD [49], due to their interactive and engaging nature as well as the robots' repetitive behaviors [48]. However, every child with ASD has a unique profile of strengths, challenges, and autism-specific characteristics [53]. The lack of personalized affect awareness hinders the ability of existing SAR tutors to perceive and respond to the unique and atypical cognitive-affective behaviors. Therefore, there is a need for affect-aware SAR tutors that can enable personalized SAR interventions for each child with ASD, with the goal of achieving positive cognitive and affective learning outcomes in the long-term.

Affect-aware SAR tutors need to be able to perceive a diverse set of cognitive-affective states in long-term interventions. For instance, engagement is an important metric for evaluating the effectiveness of SAR tutors. For this reason, most prior work has focused on applying supervised **machine learning (ML)** to enable robot perception of engagement directly from users' behavioral data (e.g., a child's vocalizations, facial and body expressions, and autonomic physiology data such as heart rate) [29, 43, 46]. However, ideal SAR tutors still need to perceive a more diverse set of cognitive-affective states, such as confusion and frustration, in order to facilitate cognitive learning [33]. Although the two-dimensional framework for describing arousal and valence has been extensively studied in HRI, computational models for arousal and valence for children with ASD have only been studied in single-session, laboratory settings [43]. Such settings pose many limitations; the feasibility of modeling children's valence and arousal needs further investigation in real-world conditions and over a longer period of time.

Personalized affect-aware SAR for children with ASD needs to be designed and validated in long-term multi-session interventions in order to adequately capture and model children's emotional expressions to target stimuli (e.g., different learning activities as part of an intervention). Although prior work has focused on personalized robot perception for ASD, the methods are only trained and validated on single-session in-lab datasets. For example, Javed et al. [30] collected a dataset of five children with ASD in a single-session study, and evaluated/trained ML models in a child-specific manner (the models were trained/tested on the non-overlapping data of the same child). Such child-specific models suffer from the limited amount of data collected from each individual, as they are unable to leverage data from other participants during training. In another single-session study, Rudovic et al. [43] proposed a **personalized perception of affect deep networks**

(**PPA-nets**), where the model personalization was achieved by the tuning of the network layers to each child's culture, gender, and individual characteristics. However, it is difficult to generalize about the longer-term effects of the personalization since all of the data were collected in a single day-long session. Therefore, to better approximate longer-term use contexts, the impact of model personalization for affect awareness needs to be evaluated on multi-session recordings of children's interactions with SAR.

This article proposes a novel approach to modeling personalization that is specifically tailored to multiple sessions of child-robot interactions of children with ASD in long-term in-home interventions. The sessions included in the dataset are a part of our month-long in-home deployments of SAR interventions with children with ASD. Specifically, the multi-session child data in this study were used to devise personalized ML models for the estimation of children's cognitive-affective states in terms of their valence and arousal levels, as described in Section 2.1. The model personalization is attained by using the notion of supervised domain adaptation—an ML approach that leverages a small portion of supervised data to learn to adapt generic (population-level) models to each individual. The contributions of this work are summarized as follows:

- *A Novel Set of Arousal-Valence Annotations.* This article introduces a novel set of arousal-valence annotations to our previously reported multi-session real-world dataset of SAR interventions [29]. We annotated 19 sessions of SAR interventions from 4 children with ASD, and derived **inter-rater reliability (IR)** scores for binary arousal (IR = 83%) and valence (IR = 81%) labels between human annotators. To the best of our knowledge, this is the first work of its kind that enabled research into the longitudinal modeling of arousal-valence behavioral cues of children with ASD using machine learning. More details can be found in Section 3.1.
- *Modeling Arousal/Valence in Long-Term Intervention.* This article is the first to propose a design and validation of a computational models of arousal and valence from children with ASD on a multi-session real-world dataset of SAR interventions. Although prior work has analyzed arousal and valence in typically developing user populations where individual differences are less pronounced, no prior work to our knowledge has attempted to model valence and arousal from multi-session data of children with ASD. Specifically, we found that the generic XGBoost classifier (our baseline) outperformed the other model candidates (Feed-forward NNs, LogReg, SVM, and KNN) and achieved a 90% AUROC score for arousal and 83% AUROC score for valence across all sessions data. We also found that additional body pose features helped to significantly improve the overall model performance in predicting arousal and valence labels. The results per child participant are found in Section 5.3.
- *Long-Term Model Personalization.* By using the notion of the **supervised domain adaptation (s-DA)** based on loss-reweighting [11], we designed and validated the models for long-term model personalization from the data of multi-session interventions for children with ASD. We demonstrate that personalized models with s-DA significantly outperform their non-personalized counterparts (i.e., the individualized and generic models) in terms of the AUROC curve (with specific improvements over individualized models: arousal: +5%, valence +4%; and generic models: arousal: +2%, valence +3%). We show that these improvements are statistically significant across all four participants ($p < .05$). We also show that the performance boost of personalized models was driven by an improved performance on the challenging negative class (low arousal/valence), while maintaining similar performance on the positive class (high arousal/valence). Results of the effects of model personalization are found in Section 5.

- *Session-Based Model Evaluation.* Different from previous works that report the results in a traditional manner on single-session datasets, due to the nature of our data, we introduce a session-based model evaluation to capture more fine-grained performance of the models on our multi-session HRI dataset. Instead of adapting the percentage-based model evaluation methods used in past work [29], we propose chronological train-test splits based on the recorded sessions; therefore, our method follows the same temporal dependence as real-world deployments of such a system. Consequently, this allows us to better quantify the effects of model personalization and its impact on performance in unseen sessions of the test participant. Results of session-based model evaluation method are found in Section 4.2.

Our experimental evaluations show the importance of model personalization and gains in terms of arousal-valence estimation accuracy on multi-session data. Building upon our prior work on modeling engagement [29], this is the first time that such long-term personalized model design and session-based evaluations have been performed on arousal-valence data of children with ASD across multiple-sessions and as part of in-home interventions. This has important implications for the design of future robot companions and tutors for children with ASD, paving the way for new personalized robot technologies for autism therapy and more effective long-term learning activities for all users.

2 RELATED WORK

This section overviews background work relevant to the main contribution areas of this article: affect-aware SAR tutors (Section 2.1), personalization for assisting ASD therapy (Section 2.2), and model evaluation for HRI (Section 2.3).

2.1 Affect-Aware SAR Tutors

A significant body of SAR research has shown that autonomous SAR tutors enhance cognitive learning gains of children with ASD in a variety of settings [18, 28]. However, those tutoring systems are limited in their ability to autonomously perceive and respond to atypical affective behaviors of children with ASD [29], hindering their ability to personalize toward the specific needs of each child [27, 29, 43]. Previous research has studied perception of basic affective states (fear, anger, happiness, sadness, disgust, and surprise) [10], but child participants rarely experience fear or disgust during the learning process [21, 24]. In contrast, Kort et al. [33] addressed the interplay of affect and learning by modeling cognitive-affective states (engagement, confusion, frustration, etc.) that children with ASD naturally experience in the context of learning, using either a *categorical* or *dimensional* approach.

Categorical approaches map affective states experienced during learning onto a set of basic cognitive-affective states such as engagement, confusion, frustration, boredom, or delight [20]. Since user engagement is considered to be a particularly important metric for evaluating the effectiveness of SAR contexts, robot perception of user engagement has become a crucial capability of autonomous SAR systems [52]. This is especially true for educational SAR tutors that must promote high user engagement levels in order to help users achieve the desired cognitive learning gains [52]. Consequently, the majority of prior work on developing affect-aware SAR tutors has focused on modeling user engagement [29].

Most previous studies have trained and evaluated supervised ML classifiers on multimodal datasets for perception of engagement in HRI interactions. Castellano et al. [9] and Sanghvi et al. [47] trained supervised engagement classifiers with hand-crafted social, physiological, or task-based interaction features. To automate this manual feature engineering process, Rudovic et al. [45] proposed a novel deep reinforcement learning architecture for estimation of engagement

directly from raw video data, where high-dimensional features were automatically extracted using a pre-trained **convolutional neural network (CNN)**, ResNet. However, due to the unique challenges of the atypical cognitive-affective behaviors of children with ASD, it is unclear how these models, trained and validated on a typically developing population, would generalize to children with ASD.

To further enable research on engagement of children with ASD, most previous studies collected multimodal datasets from single-session laboratory HRI studies. For instance, Rudovic et al. [44] collected a multimodal dataset from single-session HRI interventions with children with ASD from different cultural background (Japan and Europe). Based on that dataset, Rudovic et al. [43, 46] proposed a personalized deep learning framework to model engagement intensity of children with ASD that achieved an average agreement of 60% with human experts while estimating user engagement on a continuous scale from -1 to 1. In another single-session study, Javed et al. [30] trained a CNN to model user social engagement. The proposed CNN-based model achieved the accuracy of 78%, but it did not outperform the accuracy of tree-based random forest models (81%), indicating that simpler models may work better in this setting.

Recently, Jain et al. [29] demonstrated the feasibility of modeling user engagement in our more challenging long-term, in-home SAR deployment using the dataset also used in this work (but without model personalization; see Section 2.3 for more details on differences from our work). The authors found that **Gradient Boosted Decision Trees (XGBoost)** outperformed feed-forward neural networks and achieved approximately 90% accuracy for binary classification of user engagement [29]. However, despite this recent progress in modeling user engagement of children with ASD, engagement-aware SAR tutors are still unable to address the diverse cognitive-affective needs of children with ASD during the learning process, such as confusion and frustration that can be directly analyzed from the arousal/valence dimensions of affect that we tackle in this work. Although promoting improved engagement does help child participants to achieve higher cognitive learning gains [52], Lepper and Chabay [35] showed that expert human tutors tend to devote the same amount of attention to child participants' emotional goals as they do to cognitive learning gains. Therefore, it is important for affect-aware SAR tutors to have the ability to perceive and respond to not just engagement but a diverse range of cognitive-affective states of children with ASD.

Dimensional approaches map different cognitive-affective states onto a two-dimensional coordinate system consisting of perceived child participant valence and arousal, as shown in Figure 1 [33]. Furthermore, Kort et al. [33] categorized these cognitive-affective states into four major learning phases: (1) interested and curious about learning a new topic; (2) confused or puzzled by the learning topic; (3) frustrated or disengaged about some misconceptions; and (4) help-seeking for hints or fresh research. They suggested that autonomous tutoring systems should ideally first determine the child participant's cognitive-affective states, and then provide feedback based on that state to help the child participant be positively guided in the learning cycle in order to achieve both cognitive and emotional learning gains, as illustrated in Figure 1(B).

Although a large body of HRI research has studied emotion recognition using this two-dimensional framework [19, 25], very few studies have designed computational models for arousal and valence from children with ASD [43, 56]. Rudovic et al. [43] modeled arousal and valence for children with ASD on a scale from -1 to 1 in a single-session in-lab study. However, due to the limitations of single-session studies in laboratory settings, the feasibility of modeling the children's valence and arousal requires further investigation in the real-world conditions and over a longer period of time. To this end, in this article we address the long-term, in-home SAR interventions for children with ASD. To our knowledge, this is the first time that the modeling of affective states of children with ASD has been attempted in real-world conditions. Toward that goal, we used a novel

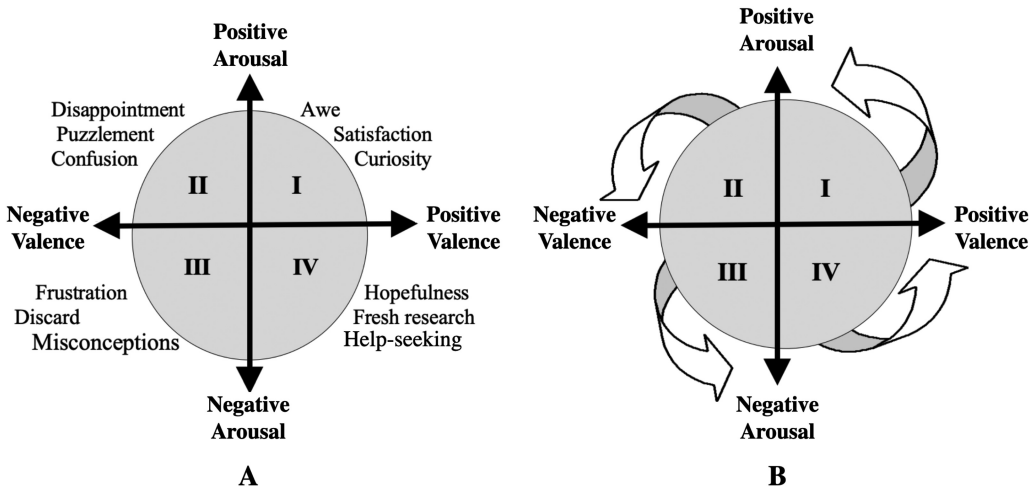


Fig. 1. A: Based on Kort et al. [33], two-dimensional arousal and valence framework for cognitive-affective states. States with positive valence (more pleasurable) are on the right; states with negative valence (more unpleasant) are on the left. Similarly, states with positive arousal (more constructive learning activity) are at the top; states with negative arousal (no learning activity) are at the bottom. B: Ideal circular flow of cognitive-affective states [33].

long-term dataset of arousal and valence collected from month-long in-home deployments of SAR tutors with children with ASD [15]. Using this multi-session dataset, we demonstrate the feasibility of modeling arousal and valence of children with ASD in long-term in-home interventions, while showcasing the potential of using a two-dimensional approach to model a more diverse set of cognitive-affective states.

2.2 Personalization in ASD

Personalization of the learning process is important for delivering effective educational interventions and is therefore a key feature of SAR tutors [16, 23]. This is particularly true in the ASD context since children with autism tend to have atypical and diverse ways of expressing their cognitive-affective states [53]. To address this heterogeneity of the user population, past work has focused on studying the *atypical behavioral patterns of children with ASD*, and *personalized modeling* methods have been developed in the ASD context.

The work on HRI and SAR for autism is informed by the extensive research on ASD in developmental psychology and related fields [7, 26]. It is well established that every child with ASD has a unique profile of strengths, challenges, and specific autism characteristics [53]. This diversity makes it necessary to personalize SAR interventions to meet the unique needs of each child with ASD. Furthermore, past work found that children with ASD had a wide variety of reactions to robots and levels of ability to perform academic tasks [50], also highlighting the need for personalization.

Personalization has been studied extensively in SAR for ASD. While past work explored personalizing the robot action selection [16, 36, 51], this work focuses on studying the model personalization for robot perception of affect. Javed et al. [30] collected individual datasets for each child with ASD based on a single-session study. Separate models were trained and evaluated using each participant's dataset. However, this individual modeling approach did not leverage data from other participants for model training. To improve on this, in another single-session study, Rudovic et al.

[43] proposed a deep learning network, where personalization of the network was achieved using the demographic information (culture and gender), followed by individual network layers for each child. The personalized model outperformed the non-personalized baseline. However, due to the limitation of single-session study, the training and test sets were randomly sampled from the same session for each participant. Therefore, the observed improvement of personalized models is likely the result of the temporal dependence of instances between training and test sets. As suggested in the work, to eliminate the possible correlation between training and test sets, multi-session long-term studies need to be conducted, so the model can be evaluated on the next session to further validate the impact of model personalization.

Toward that end, we trained and evaluated our models on a multi-session dataset collected from month-long in-home deployments with four children with ASD. By applying supervised domain adaptation techniques for personalization, this article shows that long-term personalized models outperform non-personalized baseline models. This article further validates the need for long-term personalization of robot perception for children with ASD.

2.3 Model Evaluation for HRI

Due to the nature of our dataset, the design of the model evaluation scheme is critical for properly assessing the model's performance. Consequently, we first describe traditional evaluation protocols and then ours. Traditional methods for model evaluation differ in how the training and test sets are selected and how evaluation experiments are conducted. Model evaluation aims to estimate the accuracy induced by supervised ML algorithms so as to be able to determine the best-performing model during model selection [32]. Work on affect-aware SAR tutors in ASD has primarily focused on two model evaluation methods: *holdout* and *cross-validation*.

Holdout methods partition the data into training and test sets based on a percentage parameter (e.g., 80% training data to 20% testing data). *Cross-validation methods* split the available data into k mutually exclusive subsets. For each experiment, one of the k subsets is used as the test set with the other $k - 1$ subsets combined to form the training set. Similar to holdout methods, cross-validation methods commonly use random sampling to assign instances from the available data into each subset [32].

Recent work has used randomly sampled holdout or cross-validation for model evaluation. For instance, Rudovic et al. [43] applied the randomly sampled cross-validation method to evaluate supervised ML models trained on a dataset collected from a single-session study involving children with ASD for user engagement and affect [43]. Sanghvi et al. [47] and Lala et al. [34] performed randomly sampled stratified 10-fold cross-validation to evaluate the performance of engagement classifiers trained with different model candidates. More recently, Javed et al. [30] also used randomly sampled cross-validation for every subject's individual dataset to design individualized models for detecting engagement [30].

However, due to the temporal nature of SAR deployments, as shown in Figure 2, randomly sampled model evaluation methods used in past works are impractical for two major reasons. First, since SAR deployments obtain labeled training data chronologically before the testing data, model evaluation should also follow this temporal relationship to obtain accurate estimation of model performance. Random sampling violates this temporal relationship and possibly leads to having data from late in the interaction being assigned to the training set, and data from earlier in the interaction being assigned to the test set. Secondly, participant body pose and facial expressions are highly correlated for data instances that happen close in time. Therefore, with random sampling from the same single-session dataset, the training and test sets also become correlated, leading the model evaluation process to overestimate the model performance [29].

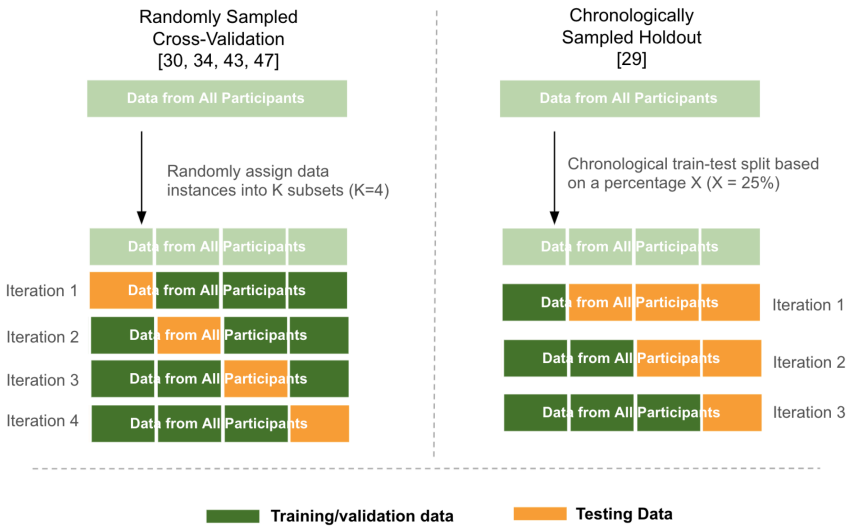


Fig. 2. The weaknesses of model evaluation methods applied in prior work. Left: (1) With random sampling from the same single-session dataset, the training set, and test set can become correlated, leading the model evaluation process to overestimate model performance; (2) Randomly sampled cross-validation method also tends to violate the temporal relationship between data instances. This possibly leads to having data from chronologically later in the interaction being assigned to the training set, and data from earlier in the interaction being assigned to the test set. Right: (1) In the chronologically sampled holdout method, more of the overall data assigned to training data leads to less testing data, so model performance cannot be correctly estimated in all data instances, especially if the chronologically earlier data have substantial variance from later data; (2) Since long-term SAR deployments obtain each session chronologically possibly on different days, the model can only be trained on the earlier sessions and tested on the next deployed session on another day. Since this percentage-based method is likely to split the first half of a session as the training set and the second half as a test set, the model evaluation process violates the temporal dependence between sessions and leads to overestimating model performance. More details are discussed in Section 2.3.

To address the temporal nature of SAR datasets, Jain et al. [29] applied a chronologically sampled holdout method for model evaluation, where an early subset of the data was used for training and the remaining subset was used for testing. As illustrated in Figure 2, the work defined an *early subset* of the data as the first X% of a user’s data sorted chronologically. The evaluation experiments were conducted by varying the percentage of data used for training from the first 10% to the first 90% of a user’s data, in increments of 10%. The study successfully demonstrated the feasibility of recognizing user disengagement in real-world SAR settings. However, despite following the temporal relationship between data instances within each individual session, this method is still not practical for multi-session model evaluation. First, in this chronologically sampled holdout method, more of the overall data assigned to training data leads to less available testing data, so the method cannot accurately capture model performance in all data instances, especially if chronologically earlier data have substantial variance with later data. In addition, since long-term SAR deployments obtain each session of deployment chronologically, possibly on different days, the model can only be trained on the earlier sessions and tested on the next deployed session on another day. Since this percentage-based method is likely to split the first half of a session as the training set and the second half as the test set, the model evaluation process violates the temporal dependence between sessions and leads to overestimating model performance. To address these



Fig. 3. System setup for the month-long in-home deployments. This article introduces a novel set of arousal-valence annotations to our previously reported multi-session real-world dataset of SAR interventions [29]. To the best of our knowledge, this is the first work of its kind that enabled research into the longitudinal modeling of arousal-valence behavioral cues of children with ASD using machine learning.

limitations, we introduce a session-based model evaluation method that follows the temporal dependence both across sessions and across data instances, so that the model evaluation process is able to estimate the real-world results of multi-session SAR deployments for children with ASD. More details on the proposed evaluation protocol can be found in Section 4.2 and Figure 6.

3 DATA COLLECTION AND PREPROCESSING

The insights of this work result from a novel *long-term dataset of arousal and valence data from children with ASD* collected in our month-long in-home deployments of SAR tutors [29]. In those deployments, each child participant with ASD interacted with a SAR tutor over many sessions by playing educational math games on a tablet, as shown in the Figure 3. Based on each child's cognitive performance, the SAR tutor provided both verbal and expressive feedback to promote the child's social and math skill development.

3.1 Participants

The average age of the participants whose data were used for this work was 4.58 ± 0.23 years old; 75% were male and 25% female, reflecting ASD rates in the recruited population. As discussed in our previous work [29], the study was approved by the Institutional Review Board of the University of Southern California (UP-16-0075(v)). Informed consent forms were obtained from the children's caregivers. The child participants had a clinical diagnoses of ASD in mild to moderate ranges according to the *Diagnostic and Statistical Manual of Mental Disorders* [2]. The details about the SAR tutor and intervention design can be found in our previous publications [17, 38].

Due to the challenges of our month-long in-home data collections and the limitations of human annotation, four participants with sufficient and qualified data were selected to train and evaluate the personalized cognitive-affective models proposed in this study. Specifically, Participant 1 had five annotated sessions (total: 2hrs:8min:11sec); Participant 2 had six annotated sessions (total: 2hrs:8min:46sec); Participant 3 had four annotated sessions (total: 2hrs:16min:48sec); and

Participant 4 had four annotated sessions (total: 1hr:42min:35sec). As discussed in our previous work [29], data collected before and after the interventions were truncated for each session to only include the child-robot interaction between the first and last game in the dataset. Moreover, in this work, we also exclude trivial frames of disengagement where participants left the recorded view, so that the models can be trained/evaluated on the relevant portion of the data. The details about the final number of instances for each session and participant before applying the windowing technique can be found in Table 2 in the Appendix.

3.2 Feature Extraction

The feature space of our multimodal dataset consists of 115 visual (body pose and face) features, 6 audio features, and 16 game performance features. As discussed in our previous work [29], the features were extracted using open-source libraries feasible for closed-loop use. The facial features extracted using OpenFace [3] included: (a) eye gaze direction; (b) facial expressions; (c) face detection confidence value; and (d) facial keypoint position. Praat [4] was used to extract pitch, frequency, intensity, and harmonicity features from the recorded audio. Game performance features derived from the SAR system recording included: (a) challenge level and type of game being played; (b) number of games played in the session; (c) elapsed time in a session, game, and since the robot last talked; and (d) number of incorrect responses. Additionally, in this work, we also used OpenPose [8] to extract body pose features including the number of people in the environment and keypoint positions for participant's body, feet, hands, and face.

3.3 Definitions of Arousal and Valence

To enable modeling of a diverse set of cognitive-affective user states, this work demonstrates the feasibility of modeling arousal and valence of children with ASD in long-term in-home interventions. We specifically aim to detect negative cognitive-affective states, so necessary feedback or guidance can be provided to support the child's learning process. For this reason, our work focuses on training ML models to classify negative behaviors from the rest of the behaviors (positive/neutral). Each participant's arousal and valence were annotated as binary labels (positive/neutral or negative) following established practices and standard definitions from the literature (e.g., [25], [39], and [43]). In this work, the *valence label* refers to whether the child's emotion toward the SAR tutor is positive/neutral (1) or negative (0). The *arousal label* refers to whether the active level of the child's interaction with the SAR tutor is positive/neutral (1) or negative (0). Full specifications of positive, neutral, and negative arousal and valence used in this study are found in Table 3. Because the difference between neutral and positive behaviors can be small or ambiguous, this binary definition of arousal/valence helps to ensure high-quality human annotation for negative arousal/valence, and to avoid low inter-rater reliability across annotators. In addition, as shown in our previous work [29], temporal patterns in sequences of predicted binary labels can also provide information about more fine-grained arousal/valence levels. For example, a 4-second sequence of binary labels with 25% positive arousal/valence predictions indicates different arousal/valence levels from a 4-second sequence of binary labels with 75% positive arousal/valence predictions. This more fine-grained temporal pattern can then be used to guide the SARs' feedback action selection process. The lead author of this work annotated the labels of arousal and valence for the four participants. To confirm the absence of annotation bias, three additional annotators independently annotated 20% of the data for both valence and arousal labels. Inter-rater reliability scores were derived using Fleiss' Kappa, and a reliability of 81% and 83% were achieved for valence and arousal, respectively, between the primary and verifying annotators. The detailed information about the class distribution of negative arousal/valence labels can be found in Table 2 in the Appendix. As shown, we observed class imbalance in our multi-session dataset,

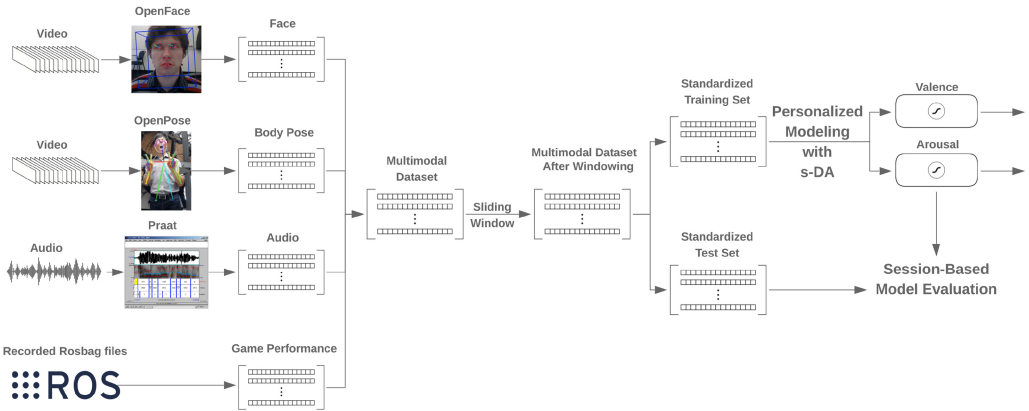


Fig. 4. Multimodal early fusion. Multimodal features were extracted using open-source libraries (OpenFace [3], OpenPose [8], Praat [4], and ROS [42]) that could be used in a closed-loop systems, and then pre-processed using sliding window and standardization methods.

and the severity of the imbalance varies across different participants. This also supports the need for personalizing the models to each participant’s unique cognitive-affective behaviors.

3.4 Data Preprocessing

As shown in Figure 4, we first performed early fusion of features from different modalities and pre-processed the dataset. After we concatenated features from different modalities, we obtained the raw multimodal dataset with a shape of $T \times M$, where T is the total number of frames and M is the total number of multimodal features. To address the uncertainty of extracted features for each individual frame, we applied a 3-second window with a 1-second shift to the raw multimodal dataset (i.e., 0 to 3 seconds, 1 to 4 seconds, 2 to 5 seconds, etc.). This 3-second window size was determined by the empirical observations of window size of participants’ affective behaviors from our annotations. As discussed in our previous work [29], within each window, the means and variances of continuous features were derived and added to the new dataset. For discrete features, a feature indicating whether features changed in the window was derived and added to the new dataset along with the means of the discrete features. This sliding window method also addressed the occasional frame drops in OpenFace and OpenPose that caused missing data. In cases when the participants were partially or completely out of the recorded frames, and OpenFace’s or OpenPose’s features were also shown as missing (NaN), the values of these features were approximated by using the maximum value of the feature after standardization. Moreover, since the multimodal features were on different time scales, all data instances in both the training set and test set method were standardized to have zero mean and unit variance based on the means and variances of the training set [29].

3.5 Model Selection for Arousal/Valence

We tested a collection of supervised ML model candidates during the model selection process. Past work has shown that tree-based XGBoost and **Deep Neural Network (DNN)** perform better than other model candidates for modeling user engagement [29, 30, 43]. Consequently, we implemented XGBoost and DNN as the potential model candidates. Moreover, we also implemented **Support Vector Machine (SVM)** and **K-Nearest Neighbor (KNN)** as baseline models for comparison. Because of large variances across sessions for each participant typical for ASD, hyperparameter

tuning strategies such as grid search tend to overfit to the training data, so we used default hyperparameters for the model candidates. As discussed in our previous work [29], we implemented the ML models using the open-source libraries shown as follows: Scikit-learn version 0.22.1 [40], XGBoost version 0.90 [12], TensorFlow version 1.15.0 [1], and Keras version 2.2.4 [13].

4 METHODS

We introduce a *personalized modeling method with supervised domain adaptation (s-DA)* for training personalized cognitive-affective models for children with ASD in long-term interventions. To validate the effects of long-term personalization, we introduce the *session-based model evaluation method* to evaluate the personalized models by comparing them with two non-personalized baselines: individualized (participant-dependent) and generic (participant-independent) models. In this section, we describe these methods, describe the training and evaluation of the personalized models for arousal and valence of children with ASD, and further demonstrate the feasibility of modeling their cognitive-affective states in long-term in-home interventions.

4.1 Personalization with Supervised Domain Adaptation

As discussed in our previous work [29], two major baseline methods have been used to train cognitive-affective models in long-term interventions: (1) individualized models trained only on available individual data; and (2) generic models trained only on available data from other participants. For individualized modeling, since participant-specific data are likely to be limited and imbalanced, individualized models are prone to overfitting. On the other hand, generic models are participant-independent, thus it is more challenging for such models to account for the heterogeneity in the data from the ASD population. We used both the generic and individualized models as baselines, and compared with our personalized modeling to explore the effects of the model personalization.

Our personalized modeling is achieved by using s-DA to obtain the optimal personalized cognitive-affective models for each child with ASD. By applying domain adaptation via *loss reweighting* [11], as described below, our method is able to trade off between the limited individual data and the more abundant but less personal data pooled from other participants, so that the best-performed personalized model can be obtained for each participant.

Loss reweighting is an instance-based method for supervised domain adaptation that adjusts the loss function of a classifier to weigh the relative importance of target and source datasets [11]. In our personalized modeling method, shown in Figure 5, data from other participants are used as the source domain, and the individual data from the test child participant with ASD are considered the target domain. By performing a hyper-parameter search with 5-fold cross-validation on the target individual data for the optimal weight α between the source and target domains, the loss reweighting s-DA technique searches for the optimal personalized model h that minimizes the α -error $\epsilon_\alpha(h)$. By adopting the optimized implementation of loss reweighting (Chen et al. [11]), this can formally be expressed as:

$$\hat{\epsilon}_\alpha(h) = \alpha \hat{\epsilon}_T(h) + (1 - \alpha) \hat{\epsilon}_S(h), \quad (1)$$

where α -error $\epsilon_\alpha(h)$ is a linear combination of the source domain (other participants' data) error $\hat{\epsilon}_T(h)$ and the target domain (individual data) error $\hat{\epsilon}_S(h)$ for a given $\alpha \in [0, 1]$. In the case of our multi-session data, we generalize the loss re-weighting function from Equation (1) by accounting for its dependence on the sessions that accumulate over time as:

$$\hat{\epsilon}_{\alpha,k}(h) = \alpha \hat{\epsilon}_{T,k}(h) + (1 - \alpha) \hat{\epsilon}_{S,k}(h), \text{ where } k \in \{1, \dots, K\}, \quad (2)$$

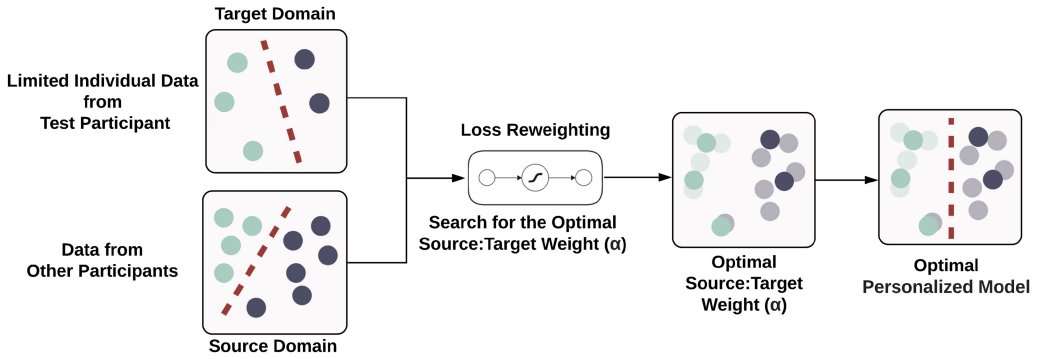


Fig. 5. Loss Reweighting. Loss reweighting is an instance-based method for supervised domain adaptation that adjusts the loss function of a classifier to weigh the relative importance of target and source datasets [11]. It was applied in concert with personalized modeling to obtain the optimal personalized model for each child participant with ASD.

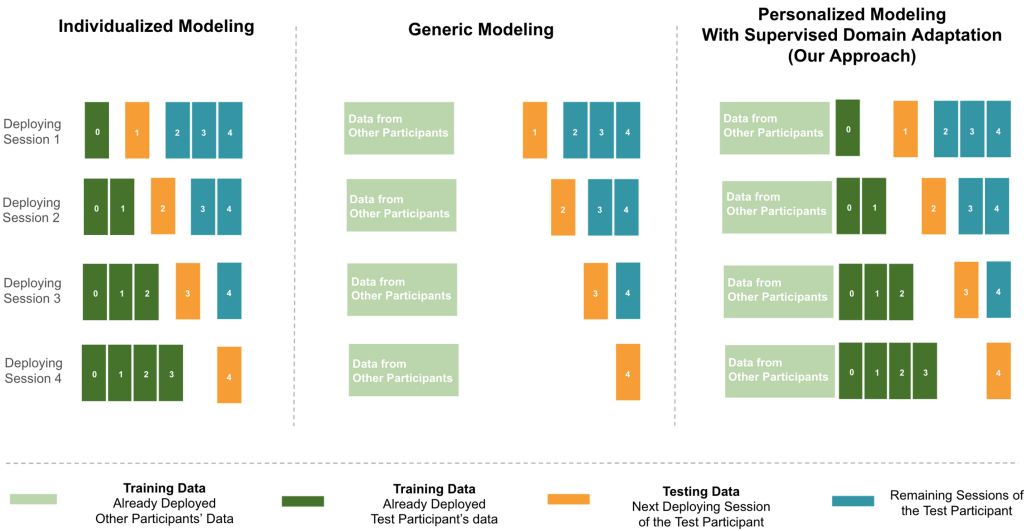


Fig. 6. A Comparison of Modeling Methods. Left: individualized models are trained only on available individual data. Center: generic models are trained only on available data from other participants. Right: personalized models with s-DA are trained on available data from both the test data and other participants' data with the re-weighted loss function.

where the session index k is used to denote the data accumulated from all previous sessions up to session k . In this way, the personalized cognitive-affective models are gradually adapted to the target child's affective behaviors, while more individual data being included in a session-by-session fashion. Figure 6 provides a comparison of our approach to other modeling methods.

4.2 Session-Based Model Evaluation

We introduce a session-based model evaluation to estimate the performance of the newly introduced models for multi-session data from SAR deployments for children with ASD. As discussed in Section 4.2, randomly sampled and percentage-based model evaluation methods used in past

work are not suitable for multi-session SAR deployments. This work performs chronological train-test splits based on sessions; therefore, our method follows the same temporal dependencies both across sessions and across data instances as found in real-world deployments.

One round of session-based model evaluation for different modeling methods is illustrated in Figure 6. For each round, one out of the total N participants is chosen to be the current test participant (P), and the data of the other $N - 1$ participants are available for training. In each round, if there are S sessions collected in the dataset for test participant P , $S - 1$ experiments are then conducted chronologically for that test participant. For example, the current test participant illustrated in Figure 6 has a total of five sessions, four experiments are conducted for that test participant. For each experiment, the train-test split is then based on P 's sessions and follows the same pattern of real-world deployments. For example, in the first experiment of individualized modeling, we assume that we have already obtained and annotated the first session of the test participant. This first session is used as the training set and the second session is used as the test set. In the next experiment of individualized modeling, the second session of the test participant is treated as if the deployment has completed and the data from the second session have been annotated. The model is then retrained on the union of the first and second sessions and tested on the third session. This process continues until every session except the first has been used as the test set. With N different participants, N rounds of experiments are conducted with different participants serving as the test participant for each round.

Similarly, with different training data, generic models and personalized models with s-DA can also be evaluated in the same fashion, as illustrated in Figure 6. Since all modeling methods are tested on the same set of sessions, the model evaluation can be kept consistent across different modeling methods. The final result is derived by averaging the models' performance in a weighted fashion on all tested sessions for each participant and comparing the performance between different modeling methods and model candidates. AUROC was used as the single-value metric to evaluate the overall performance of ML models on the binary classification of arousal and valence scores [5]. Likewise, F1 scores were used to evaluate the models' performance on both the positive and negative classes of arousal and valence.

5 RESULTS

5.1 Effect of Model Personalization with s-DA

We trained and evaluated the personalized models with s-DA with the aim of enabling automatic perception of arousal and valence of children with ASD during multi-session SAR interventions. As discussed in Section 4, to obtain the optimal personalized models during training, the s-DA technique, used for our model personalization, enables the tradeoff between the limited individual data and the more abundant data pooled from other participants. The effects of personalization were further validated by comparison with two baseline models: (1) individualized models trained only on available individual data; and (2) generic models trained only on data from other participants.

As detailed in Section 4.2, for each round of experiments, one of the four participants was chosen to be the current test participant, and data from the other three participants were considered to have been already obtained from deployments and available for training. Within each round, different modeling methods were trained and evaluated on the same sessions of the test participant using the session-based model evaluation method.

5.1.1 Analysis of Participants. Table 1 reports models' AUROC scores of each experiment round with different test participant (P-1 to P-4) for both arousal and valence, so that we can compare the performance of the personalized models and the two baseline (individualized and generic) models. In addition, the weighted average AUROC scores of the four rounds of

Table 1. Comparison of AUROC Scores of Personalized and Baseline XGBoost Models

Task	Test Participant-ID	IND [29]	GEN [29]	Our Approach: PER (s-DA)
Arousal	P-1	0.84	0.91	0.91
	P-2	0.90	0.90	0.91
	P-3	0.90	0.91	0.92
	P-4	0.85	0.90	0.93
	wAVE	0.87	0.90	0.92
Valence	P-1	0.79	0.82	0.84
	P-2	0.85	0.86	0.88
	P-3	0.84	0.83	0.87
	P-4	0.81	0.87	0.87
	wAVE	0.82	0.83	0.86

Personalized (PER) Models with s-DA Outperformed Both the Individualized (IND) and Generic (GEN) Models on the Weighted Average (wAVE), and Always Performed at Least as Well as the Baseline Models Across Individual Participants and Tasks.

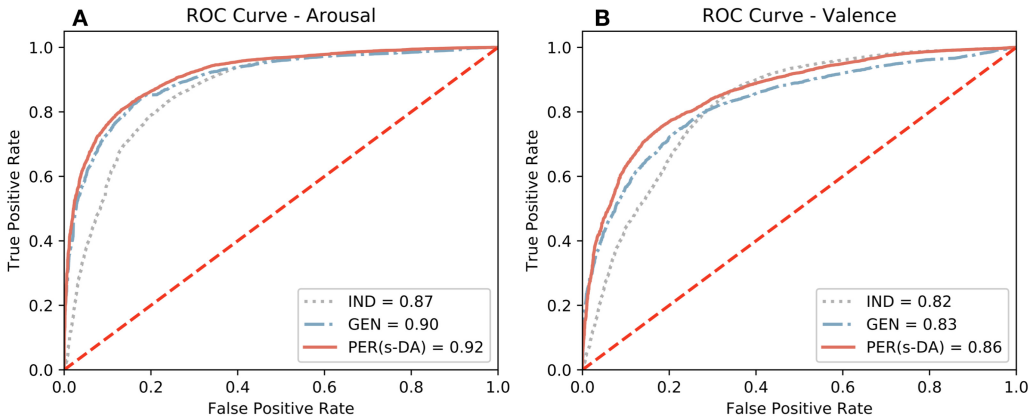


Fig. 7. ROC Curves of Personalized and Baseline XGBoost Models. Personalized (PER) models with s-DA outperformed both the individualized (IND) and generic (GEN) models using the metric of Area Under Receiver Operating Characteristic Curve (AUROC).

experiments for both arousal and valence are also reported to further examine the overall effects of the personalization.

Individually, personalized models outperformed the baseline models in six out of eight rounds of experiments across different test participants and tasks, and also performed at least as well as the baseline models in the other two rounds of experiments. As shown in Figure 7, overall, personalized models outperformed the best baseline models by approximately 2% for arousal and 3% for valence on a wAVE of the four test participants. More specifically, personalized models with s-DA achieved a 92% AUROC score for arousal and an 86% AUROC score for valence. The ROC curves show that our personalized models consistently outperform both model alternatives.

Moreover, as can be seen from Figure 8, personalized models especially outperformed baseline models on the negative class, while maintaining the same performance as the baseline models on the positive class (see also Section 6.1). For the negative class, personalized models achieved 2% and 11% improvements for valence over individualized and generic models, respectively, and also achieved 5% improvements in F1 score for arousal over both baseline models. On the other hand,

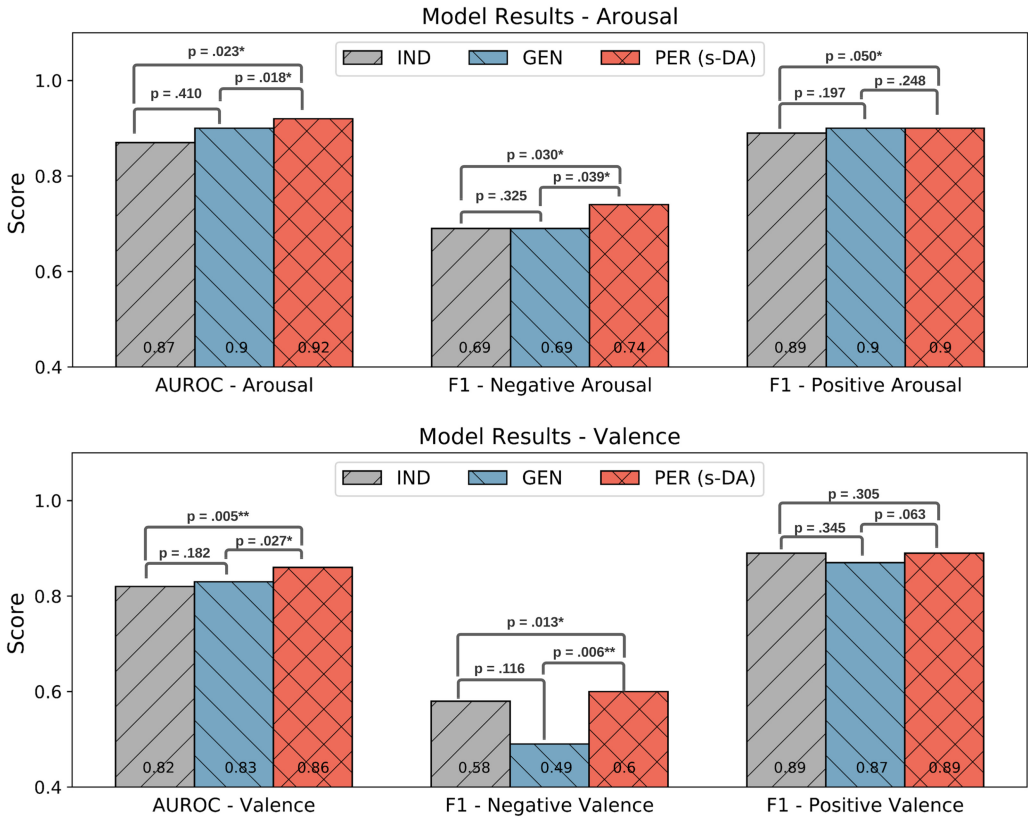


Fig. 8. Effect of Model Personalization with s-DA. (1) Personalized (PER) models with s-DA significantly outperformed the baseline models in AUROC not only on the weighted average but also across individual tested sessions; (2) This performance boost of personalized models was driven by statistically higher performance for the negative class ($p < .05$) across individual sessions, while maintaining performance on the positive class (no statistical difference).

for the positive class, personalized models performed as the best baseline models on each task. Overall, personalized models achieved 90% (arousal) and 89% (valence) in F1 score for the positive instances, and 74% (arousal) and 60% (valence) in F1 score for the negative instances. This gap in performance between the two classes is due to the class imbalance against the negative class of arousal and valence. This is discussed in more details in Section 5.2.

5.1.2 Analysis of Sessions. We also conducted one-sided Wilcoxon signed-rank tests to validate whether the effects of personalization for individual participants on the weighted average also extend to the individual sessions. As detailed in Figure 8, for both arousal and valence in AUROC, the one-sided Wilcoxon signed-rank tests validated that personalized models with s-DA also significantly ($p < .05$) outperformed the two baseline models across sessions in AUROC, while there was no significant difference between individualized and generic models despite the difference in the weighted average.

Furthermore, Wilcoxon signed-rank tests also validated that the performance boost of personalized models was driven by statistically higher performance for the negative class ($p < .05$) in F1 scores across individual sessions, while maintaining performance on the positive class

Table 2. Session-by-Session Effects of Personalization with s-DA in AUROC

Task	Model	Test Participant																		
		P-1				P-2					P-3				P-4					
		S-1	S-2	S-3	S-4	wAVE	S-1	S-2	S-3	S-4	S-5	wAVE	S-1	S-2	S-3	wAVE	S-1	S-2	S-3	wAVE
Arousal	IND	0.72	0.94	0.95	0.92	0.84	0.83	0.93	0.91	0.93	0.95	0.90	0.81	0.95	0.97	0.90	0.74	0.93	0.93	0.85
	GEN	0.90	0.94	0.94	0.94	0.91	0.88	0.91	0.91	0.89	0.94	0.90	0.85	0.95	0.94	0.91	0.89	0.88	0.93	0.90
	PER (s-DA)	0.90	0.95	0.94	0.93	0.91	0.85	0.93	0.93	0.93	0.96	0.91	0.86	0.96	0.94	0.92	0.91	0.93	0.93	0.93
Valence	IND	0.60	0.84	0.92	0.81	0.79	0.70	0.92	0.89	0.95	0.90	0.85	0.74	0.92	0.96	0.84	0.73	0.88	0.88	0.81
	GEN	0.67	0.94	0.77	0.88	0.82	0.76	0.96	0.87	0.90	0.90	0.86	0.73	0.89	0.93	0.83	0.86	0.85	0.93	0.87
	PER (s-DA)	0.70	0.91	0.90	0.86	0.84	0.80	0.97	0.89	0.95	0.91	0.88	0.79	0.91	0.96	0.87	0.80	0.90	0.92	0.87

Personalized (PER) models with s-DA were able to trade off between the alternating performance of individualized (IND) and generic (GEN) models, thus helping PER models with s-DA to outperform the baseline models not only on the weighted average but also statistically across sessions.

(no statistical difference). More specifically, in AUROC, the Wilcoxon signed-rank tests found significantly higher AUROC for personalized over generic models (arousal: $Z = 2.101$, $p = .018$, $r = .384$; valence: $Z = 1.931$, $p = .027$, $r = .353$), and for personalized over individualized models (arousal: $Z = 1.988$, $p = .023$, $r = .363$; valence: $Z = 2.556$, $p = .005$, $r = .467$). However, there was no significant difference in AUROC between generic and individualized models (arousal: $Z = 0.227$, $p = .41$, $r = .041$; valence: $Z = 0.909$, $p = .166$, $r = .182$). For F1 scores of both negative and positive classes, the detailed results of Wilcoxon signed-rank tests can be found in the Appendix A.1.

In addition, for each individual session, personalized models with s-DA were able to trade off between the alternating performance of individualized and generic models, so it helped personalized models with s-DA to outperform the baseline models not only on the weighted average but also statistically across sessions. As can be noted from Table 2, the better-performed baseline model alternated between individualized and generic models, explaining the insignificant difference between individualized and generic models. As discussed in Section 4, by taking advantage of individual data and data from other participants, personalized models achieved the best performance among the model candidates in the majority of test sessions. In the rest of the sessions, personalized models still consistently performed close to the better-performed baseline model. This further explains the significant increases from individualized/generic to personalized models across both participants and sessions. Finally, for the majority of the sessions/participants, the main increase in performance by the personalized models happens after adding the data of the first session of the target participant to the training set. After this, the performance trend remains relatively constant. This is the case for all but P-2, where we see a drop in the performance after the data from session 2 are added. This indicates that not all data from every session/participant are relevant for improving the model's performance in future sessions. One way to handle this is to learn how to select the most informative data and include only those in the adaptation set for model personalization. Various approaches have been proposed to tackle this issue by using the notion of Active/Reinforcement Learning (e.g., see [45]); this is outside of the scope of this work and will be explored in future work.

5.2 Class and Data Distribution

As detailed in Figures 9(A) and 9(B), for both arousal and valence, we observed class imbalance in our multi-session dataset, helping to explain the less accurate performance of models on negative classes. Across different test participants with ASD, arousal and valence varied considerably from each other. From all instances of valence and arousal, average percentage of negative arousal for each child participant ranged from 22.7% to 35.4% with the standard deviation of 6%, and average percentage of negative valence ranged from 14% to 27% with the standard deviation of 10%.

This variance between participants further supports the need for personalizing the models to each participant's unique cognitive-affective behaviors. Within each participant, different sessions

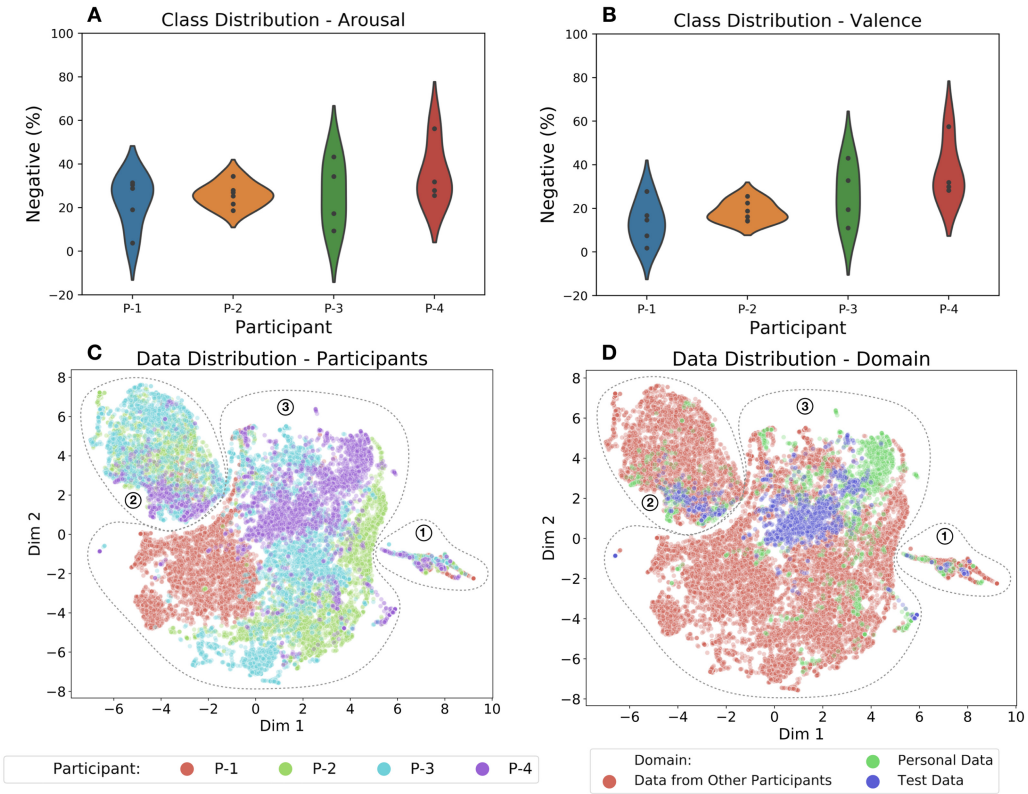


Fig. 9. Class and Data Distribution. A: Percentage of negative arousal out of all arousal instances for each session of all participants. B: Percentage of negative valence out of all arousal instances for each session of all participants. C: Visualization of the data distribution for each participant using the T-distributed t-SNE. D: Visualization of the data distribution for each domain of data within Participant 4 (Cluster 1: Instances without the majority of the OpenFace and OpenPose features; Cluster 2: Instances without the majority of the OpenFace features; Cluster 3: Instances with the majority of the features).

also had substantial variance, highlighting the challenges of personalization in long-term in-home settings. The standard deviation between sessions of each participant ranged from 5% to 14% for arousal, and from 5% to 15% for valence.

To better interpret the data distribution, this work also visualized the feature space of different participants using the **T-distributed Stochastic Neighbor Embedding (t-SNE)**, a commonly used unsupervised dimensionality reduction technique. As shown in Figure 9(C), across different participants, high variance was observed as expected between the feature space of each participant. Within each participant, substantial variances were also observed in the feature space. For P-4, instead of clustering together, the data are located in smaller clusters separated from one another. This is consistent with our observations from annotation indicating that P-4 was more likely to move around during the interaction in the in-home environment compared to other participants.

Figures 9(C) and 9(D) shows that the extreme clusters (clusters 1 and 2) consist of instances where the participants were partially out of the recorded frames. Specifically, cluster 2 consists of instances where OpenFace failed to detect the majority of the facial features and only part of the body pose features were extracted. For cluster 1 (Figure 9(C)), both OpenFace and OpenPose

Table 3. Results of Different Model Candidates

Test Participant	Task	KNN	SVM	LogReg	NN	XGBoost
wAVE	Arousal	0.76	0.86	0.86	0.84	0.90
	Valence	0.70	0.74	0.79	0.78	0.83

Tree-based XGBoost models were the most successful among the supervised ML model candidates we evaluated.

failed to detect the majority of the features from the recorded frames. The data instances in these two extreme clusters were mostly classified as negative arousal or valence. This is intuitive, since the models treated the missing features as negative behavioral cues of participants' arousal and valence. We could have removed these negative instances to make modeling less challenging, as we did when the child was completely missing from the camera-view. Instead, we left those examples as they correspond to the partially observed data that still provide a proxy for the negative valence/arousal, since the participants are not showing positive valence/arousal toward the interaction with the robot.

Moreover, we also visualized the feature space of different data domains (data from other participants, individual data, and test data). As shown in Figure 9(D), in the main cluster (cluster 3), the distribution of the available individual data closely resembles the distribution of the test data. This further supports the importance of leveraging individual data and the value of personalization with s-DA.

5.3 Alternative Modeling Approaches

From our initial model selection, we found that the tree-based XGBoost outperformed all the other model candidates when trained with the generic modeling method. In addition to its better performance, XGBoost also offers superior interpretability and requires less hyper-parameter tuning compared to the second best-performed NN. As shown in Table 3, we also trained and tested the following model candidates with the generic modeling method: K-Nearest Neighbors (KNN), Support Vector Machines (SVM), Logistic Regression (LogReg), Neural Networks (NN), and Gradient Boosted Decision Trees (XGBoost). Due to the challenges of the noisy and high-dimensional nature of the data we used, KNN performed poorly on both tasks (arousal: drop of 14%; valence: drop of 13%). LogReg was also not able to account for the highly nonlinear dependencies in the data (arousal: drop of 4%; valence: drop of 4%). On the other hand, the non-linear kernel method (SVM) also failed to achieve similar performance as XGBoost due to the absence of the hierarchical structure (arousal: drop of 4%; valence: drop of 9%). Finally, despite being hierarchical, NN had limited success because of the amount of data collected and the class imbalance (arousal: drop of 6%; valence: drop of 5%).

Our modality selection found that additional body pose features helped to improve the overall model performance in predicting arousal and valence labels. In our prior work [29], body pose features were dropped from the final set of multimodal features, because they did not significantly contribute to the task of engagement modeling. However, in this work, as shown in Table 4, the multimodal models trained with body pose features extracted by OpenPose significantly outperformed the models proposed in our prior work [29]. This also indicates that the focus of engagement modeling may be mostly on the visual attention of the participants, which correlates mainly with participants' head pose and gaze. On the other hand, the task of modeling arousal and valence requires an understanding of a more holistic view of the participants' cognitive-affective behaviors, so body pose features including keypoint positions for participant's body, feet, and hands would also be needed for better model performance.

Table 4. Performance Comparison Between Multimodal Models Trained without and with Body Pose Features

Test Participant	Task	All Modalities without Body Pose [29]	All Modalities (Our Approach)	Statistical Analysis of Sessions: p value
wAVE	Arousal	0.89	0.92	$p = .001^{***}$
	Valence	0.83	0.86	$p = .017^*$

Multimodal models trained with additional body pose features significantly outperformed the baseline multimodal models from [29] in AUROC not only on the weighted average but also across individual tested sessions ($p < .05$).

We also explored the personalized modeling method with a **unsupervised domain adaptation (u-DA)** technique called **CORrelation ALignment (CORAL)** [54]. CORAL minimizes the domain shift (in our case, each child participant is seen as a different domain) by aligning the second-order statistics of source and target distributions, without requiring any target labels [54]. As detailed in Table 1 in Appendix, the personalized models with u-DA performed poorly compared to the personalized models with s-DA (arousal: drop of 15%; valence: drop of 15%). This suggests that unsupervised domain adaptation methods such as CORAL may not be suitable for the noisy and nonstationary data typical for long-term interventions. Our future work will further explore this insight.

6 DISCUSSION

Personalized robot perception of arousal and valence from children with ASD is an important HRI research goal previously unexplored in the context of long-term in-home SAR interventions. By introducing a novel multi-session dataset, this work is the first to train and evaluate personalized ML models of arousal and valence in this challenging context. This section focuses on discussing the major contributions and findings of this work in the following two areas: effect of personalization in long-term interventions (Section 6.1) and limitations and future work (Section 6.2). These contributions aim to inform the design of more affect-aware and personalized HRI, so that the effectiveness of SAR can be further improved to augment both the emotional and cognitive learning gains of children with ASD.

6.1 Personalized Affect-Aware SAR in Long-Term Interventions

One of the main goals of this work was to demonstrate the feasibility of personalized cognitive-affective models in long-term HRI interventions for children with ASD. The results of this work suggest that supervised domain adaptation can effectively personalize ML models to better perceive the negative cognitive-affective behaviors of children with ASD. As shown in Section 5, personalized models with s-DA outperformed the best baseline models in AUROC on weighted average across each participant for both tasks. In particular, personalized models improved the predictions of negative arousal and valence when compared to the baseline models. Personalized models achieved 2% and 11% improvements in F1 score for negative valence over individualized and generic models, and also achieved 5% improvements in F1 score for negative arousal over both the baseline models. This is particularly important because the purpose of our cognitive-affective models is to detect negative behaviors, so that corresponding feedback can be provided to help the child user be positively guided. Since the data imbalance and unique atypical negative affective behaviors from children with ASD further increased the challenges for perceiving negative affect, both baseline models performed worse on the negative class than the positive class. Therefore, personalized models with s-DA effectively improved the models' capability to perceive the unique negative arousal and valence of each child participant.

One-sided Wilcoxon signed-rank tests further validated that the better performance of personalization for individual participants on the weighted average also extended to the individual sessions. Statistically, personalized models with *s*-DA also significantly ($p < .05$) outperformed the two baseline models across sessions in AUROC, while there was no significant difference between individualized and generic models despite the difference in the weighted average. This indicates that, with our dataset, it did not matter which method we chose between these two baselines; this may not hold for other datasets. However, for the dataset reported here, personalized models with *s*-DA were able to trade off between the strengths of both individualized and generic models.

The other main finding of this work is demonstrating the challenges of personalized affect awareness in a long-term real-world HRI setting due to the high variances and class imbalance. As detailed in Figure 9, we observed class imbalance in our long-term dataset, which caused the models to perform worse on the negative instances than the positive ones. This is expected since negative affective behaviors happened less frequently in our dataset than positive or neutral behaviors. Second, due to the unconstrained nature of the in-home setting, the more varied body movements and behaviors of the child participants resulted in more frequent failures of OpenFace or OpenPose capture of features. Therefore, the feature space was complex and challenging due to the substantial variances within each test participant. Due to the variances across sessions for each test participant, the variance between validation and test data further increase the challenges of long-term personalization.

Compared to the other alternative model candidates, XGBoost outperformed other supervised ML model candidates. This correlates with past findings for modeling engagement [29, 30] and expectations given the advantages of a hierarchical gradient boosted structure of XGBoost to perform on constrained datasets. Despite the recent success of deep learning, the challenging nature of the HRI study led to smaller amounts of noisy data with high class imbalance resulting in poorer NN performance. This work suggests that tree-based XGBoost is more suitable for real-world HRI dataset using constructed features extracted from libraries such as OpenFace and OpenPose. Furthermore, we also explored a naive instance-based unsupervised domain adaptation method called CORAL [54], and our results show that supervised domain adaptation had great advantages over CORAL. For this reason, we excluded those comparisons from the analysis reported here.

6.2 Limitations and Future Work

The major limitations of this work come from the challenging nature of real-world HRI deployments for children with ASD in long-term interventions. The highly resource-intensive nature of conducting in-home deployments and collecting human annotations makes it difficult to obtain large-scale high-quality datasets of cognitive-affective behaviors from children with ASD. At the same time, the quality and scale of an annotated dataset is crucial for effectively training ML models such as NN-based approaches. In addition, the class imbalance of our dataset led to worse performance on negative arousal and valence. Despite the improvements from model personalization, the performance on negative affect still needs to be improved to enable effective interventions with robot feedback to help children users to be positively guided. Lastly, the high variance and noise in our autism dataset further challenges personalization. As discussed in Figure 9, the unconstrained in-home setting contributes to the high variance and noise in the dataset within each participant. Moreover, such data properties are expected from ASD datasets, since it is well known that every child with ASD has a unique profile of autism characteristics [53]. Because both the source and

target domain data in this work are collected from different children with ASD, this also potentially led to higher variance both within and across domains (and children participants).

One direction for future work is to explore annotation-efficient methods for model personalization so that the same modeling performance can be achieved more efficiently in terms of time and expense with fewer manual annotations. As we found in this work, not all the data from every session/participant are relevant for improving model performance in the future deployed sessions. This suggests that the amount of human annotations can be potentially reduced by selectively annotating the most informative personal data instances. Prior work has proposed to tackle this problem by applying Active/Reinforcement Learning [45]. This is outside of the scope of this work, but we plan to explore it in the future.

Another promising approach is to continue using supervised domain adaptation with larger, more balanced and constrained datasets as source domains. For example, instead of using data from other children with ASD as the source domain, a more general dataset collected from typically developing children in a laboratory setting could potentially be used as the source domain. By keeping the target domain as the data from the target child with ASD, this method may allow personalized ML models to be trained more efficiently and effectively.

Finally, u-DA remains a promising approach for personalized modeling with limited human annotation of the target domain data, despite our findings for CORAL. Preliminary work has shown that u-DA has great potential to successfully train personalized affective labels with no human annotation in the target domain [57]. However, most of the u-DA techniques have only been validated on unimodal (e.g., vision-only) datasets in constrained settings [55]. Consequently, open research questions remain as to how to effectively train personalized affective models with multimodal datasets involving children in unstructured settings like homes.

6.3 Conclusions

This work developed and validated personalized models for robot perception of arousal and valence from children with ASD in the previously unexplored context of a long-term in-home setting. Our motivation was to design personalized cognitive-affective models to perceive the unique cognitive-affective behaviors of each child with ASD. When deployed, these models should enable personalized interventions to help each child to be positively guided in the learning cycle and achieve improved cognitive and emotional learning gains.

We introduced a novel long-term multimodal dataset of arousal and valence collected from month-long in-home deployments of SAR tutors with children with ASD. Based on this dataset, we trained models with supervised domain adaptation (loss reweighting) to trade off between the limited individual data (target domain) and the more abundant generic data pooled from other participants (source domain). The models were also validated using a session-based method that follows the temporal dependence across sessions and data instances, so that the model evaluation process is able to estimate the real-world results of multi-session SAR deployments for children with ASD. The results showed that personalized models outperformed non-personalized (individualized and generic) models in both weighted average and statistically ($p < .05$) across individual sessions. Personalized models especially improved the predictions of negative arousal and valence when compared to the baseline models. Our research demonstrates the potential of using supervised domain adaptation to improve the personalized affect awareness of SAR tutors with children with ASD in long-term interventions. We hope to inform the future development of personalized affect-aware robot tutors tailored towards individuals with atypical cognitive-affective needs such as children with ASD.

A APPENDIX

A.1 Wilcoxon Signed-Rank Tests for F1 Scores of Negative and Positive Classes

For the predictions of negative classes, Wilcoxon signed-rank tests also indicated significant increases in F1 score for negative arousal/valence between personalized and generic models (arousal: $Z = 1.761$, $p = .039$, $r = .321$; valence: $Z = 2.499$, $p = .006$, $r = .456$), and in F1 score for negative arousal/valence between personalized and individualized models (arousal: $Z = 1.874$, $p = .03$, $r = .342$; valence: $Z = 2.215$, $p = .013$, $r = .404$). Similar to the prior AUROC results, there was also no significant increases in F1 scores for negative arousal/valence indicated between generic and individualized models (arousal: $Z = 0.454$, $p = .325$, $r = .083$; valence: $Z = 1.193$, $p = .116$, $r = .218$).

For the predictions of positive classes, except that personalized models significantly outperformed individual models for arousal, there were no significant increases between different modeling methods. Wilcoxon signed-rank tests indicated significant increases in F1 scores for positive arousal between personalized and individualized models ($Z = 1.647$, $p = .05$, $r = .301$). However, no significant increases were indicated in F1 scores for positive valence between personalized and individualized models ($Z = 0.511$, $p = .305$, $r = .093$), in F1 scores for positive arousal/valence between personalized and generic models (arousal: $Z = 0.682$, $p = .248$, $r = .124$; valence: $Z = 1.533$, $p = .063$, $r = .280$), and in F1 scores for positive arousal/valence between individualized and generic models (arousal: $Z = 0.852$, $p = .197$, $r = .156$; valence: $Z = 0.398$, $p = .345$, $r = .073$).

A.2 Table of Comparison between Supervised and Unsupervised Domain Adaptation

Table 1. Comparison between Supervised and Unsupervised Domain Adaptation

Task	Test Participant-ID	PER (s-DA)	PER (u-DA)
Arousal	P-5	0.91	0.81
	P-7	0.92	0.76
	P-9	0.92	0.77
	P-17	0.92	0.73
	wAVE	0.92	0.77
Valence	P-5	0.84	0.72
	P-7	0.88	0.77
	P-9	0.87	0.66
	P-17	0.87	0.73
	wAVE	0.86	0.71

Personalized models with u-DA performed poorly comparing to the personalized models with s-DA (Arousal: drop of 15%; Valence: drop of 15%).

A.3 Table of Number of Instances and Class Distribution for Each Session and Participant

Table 2. Number of Instances and Class Distribution for Each Session and Participant

Participants	Sessions	Number of Instances	Negative Arousal	Negative Valence
P-0	S-0	35,941	4%	2%
	S-1	38,851	31%	7%
	S-2	47,510	29%	17%
	S-3	30,143	19%	15%
	S-4	38,011	30%	28%
	Total	190,456	23%	14%
P-2	S-0	60,437	25%	16%
	S-1	25,387	27%	14%
	S-2	13,223	34%	26%
	S-3	28,417	28%	22%
	S-4	25,538	22%	14%
	S-5	17,653	19%	19%
Total	170,655	25%	18%	
P-3	S-0	22,547	34%	33%
	S-1	94,092	17%	19%
	S-2	59,177	43%	43%
	S-3	3,603	9%	11%
	Total	179,419	28%	29%
P-4	S-0	27,823	26%	30%
	S-1	48,298	32%	32%
	S-2	50,032	28%	28%
	S-3	12,550	56%	57%
	Total	138,703	31%	32%

Number of Instances: The number of instances for each session of every participant before applying the windowing technique; Negative Arousal/Valence: The percentage of negative arousal/valence out of all arousal/valence instances for each session of every participant.

A.4 Table of the Annotation Criteria

Table 3. Arousal/Valence Annotation Criteria

Level	Label	Definition
Positive	Arousal	The child performs a high level of interactive activities with the socially assistive robot (SAR) system. The child is engaged with the robot and showing clear, visible signs of physical activities (e.g., hand clapping, talking to the robot, touching the tablet/robot, and actively playing with the robot).
Neutral	Arousal	The child performs a low level of interactive activities with the SAR system. The child shows no clear, visible signs of physical activity with the robot, but seems calm, thinking, and engaged with the SAR system.
Negative	Arousal	The child performs no interactive activity with the SAR system (e.g., looking away, not showing interest in the interaction with the robot, seeking help from the parents, discussing the game or puzzle with parents, playing with toys, bored, sleepy).
Positive	Valence	The child shows clear positive emotions toward the SAR system (e.g., laughing or smiling with the robot, showing signs of happiness or joy such as clapping or celebrating).
Neutral	Valence	The child shows no clear, visible signs of pleasant or unpleasant emotions (e.g., calm and engaged with the robot, seeking help from parents, thinking and working with the robot).
Negative	Valence	The child shows clear and visible negative emotions and disinterest toward the SAR system (e.g., being unhappy or frustrated, showing confusion or dissatisfaction, visibly upset, rubbing eyes, putting hands in front of the face or shaking head to show dissatisfaction and embarrassment, looking away and showing disinterest (not seeking help), asking to play with toys instead of the SAR system).

Arousal and valence were annotated as binary labels (positive/neutral or negative) for each participant following standard definitions grounded in established practices from emotion science, automated affect recognition, and human-robot interaction [25, 39, 43].

ACKNOWLEDGMENTS

The authors want to thank Balasubramanian Thiagarajan for his help with data analysis, and Sophia Pei, Majlis Whetsel, and Haocheng Yuan for their assistance with annotations. The authors are also very grateful to Caitlyn Clabaugh, Gisele Ragusa, Roxanna Pakkar, David Becerra, Eric C. Deng for their key roles in the original study design, recruitment, assessments, and more. The entire research team thanks the children and families who participated in the study that generated the dataset.

REFERENCES

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mane, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viegas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467* (2016).

- [2] American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders (DSM-5®)*. American Psychiatric Pub.
- [3] Tadas Baltrušaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. 2018. OpenFace 2.0: Facial behavior analysis toolkit. In *Proceedings of the 2018 13th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2018)* (2018), 59–66.
- [4] Paul Boersma and David Weenink. 2001. PRAAT, a system for doing phonetics by computer. *Glott International* 5 (2001), 341–345.
- [5] Andrew P. Bradley. 1997. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition* 30, 7 (1997), 1145–1159.
- [6] Cynthia Breazeal. 2003. Emotion and sociable humanoid robots. *International Journal of Human-Computer Studies* 59, 1–2 (2003), 119–155.
- [7] John-John Cabibihan, Hifza Javed, Marcelo Ang, and Sharifah Mariam Aljunied. 2013. Why robots? A survey on the roles and benefits of social robots in the therapy of children with autism. *International Journal of Social Robotics* 5, 4 (2013), 593–618.
- [8] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. 2019. OpenPose: Realtime multi-person 2D pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019).
- [9] Ginevra Castellano, André Pereira, Iolanda Leite, Ana Paiva, and Peter W. McOwan. 2009. Detecting user engagement with a robot companion using task and social interaction-based features. In *Proceedings of the 2009 International Conference on Multimodal Interfaces*. 119–126.
- [10] Oya Celiktutan, Evangelos Sariyanidi, and Hatice Gunes. 2018. Computational analysis of affect, personality, and engagement in human–robot interactions. In *Computer Vision for Assistive Healthcare*. Elsevier, 283–318.
- [11] Justin Chen, Edward Gan, Kexin Rong, Sahaana Suri, and Peter Bailis. 2019. CrossTrainer: Practical domain adaptation with loss reweighting. In *Proceedings of the 3rd International Workshop on Data Management for End-to-End Machine Learning*. 1–10.
- [12] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 785–794.
- [13] Francois Chollet. 2015. Keras. <https://keras.io>.
- [14] Deborah L. Christensen, Kim Van Naarden Braun, Jon Baio, Deborah Bilder, Jane Charles, John N. Constantino, Julie Daniels, Maureen S. Durkin, Robert T. Fitzgerald, Margaret Kurzius-Spencer, Li-Ching Lee, Sydney Pettygrove, Cordelia Robinson, Eldon Schulz, Chris Wells, Martha S. Wingate, Walter Zahorodny, and Marshalyn Yeargin-Allsopp. 2018. Prevalence and characteristics of autism spectrum disorder among children aged 8 years—autism and developmental disabilities monitoring network, 11 sites, United States, 2012. *MMWR Surveillance Summaries* 65, 13 (2018), 1.
- [15] Caitlyn Clabaugh, Shomik Jain, Balasubramanian Thiagarajan, Zhonghao Shi, Leena Mathur, Kartik Mahajan, Gisele Ragusa, and Maja Mataric. 2018. Attentiveness of children with diverse needs to a socially assistive robot in the home. In *Proceedings of the 2018 International Symposium on Experimental Robotics*.
- [16] Caitlyn Clabaugh, Kartik Mahajan, Shomik Jain, Roxanna Pakkar, David Becerra, Zhonghao Shi, Eric Deng, Rhianna Lee, Gisele Ragusa, and Maja Mataric. 2019. Long-term personalization of an in-home socially assistive robot for children with autism spectrum disorders. *Frontiers in Robotics and AI* 6 (2019), 110. <https://doi.org/10.3389/frobt.2019.00110>
- [17] Caitlyn Clabaugh, Kartik Mahajan, Shomik Jain, Roxanna Pakkar, David Becerra, Zhonghao Shi, Eric Deng, Rhianna Lee, Gisele Ragusa, and Maja Mataric. 2019. Long-term personalization of an in-home socially assistive robot for children with autism spectrum disorders. *Frontiers in Robotics and AI* (2019).
- [18] Caitlyn Clabaugh and Maja Mataric. 2019. Escaping Oz: Autonomy in socially assistive robotics. *Annual Review of Control, Robotics, and Autonomous Systems* 2 (2019), 33–61.
- [19] Adam K. Coyne, Andrew Murtagh, and Conor McGinn. 2020. Using the Geneva emotion wheel to measure perceived affect in human-robot interaction. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (Cambridge, United Kingdom) (HRI'20)*. ACM, New York, 491–498. <https://doi.org/10.1145/3319502.3374834>
- [20] Sidney D’Mello and Art Graesser. 2011. The half-life of cognitive-affective states during complex learning. *Cognition & Emotion* 25, 7 (2011), 1299–1308.
- [21] Sidney D’Mello, Rosalind W. Picard, and Arthur Graesser. 2007. Toward an affect-sensitive AutoTutor. *IEEE Intelligent Systems* 22, 4 (2007), 53–61.
- [22] David Feil-Seifer and Maja J. Mataric. 2005. Defining socially assistive robotics. In *Proceedings of the 9th International Conference on Rehabilitation Robotics, 2005 (ICORR 2005)*. IEEE, 465–468.
- [23] Jillian Greczek, Elaine Short, Caitlyn E. Clabaugh, Katelyn Swift-Spong, and Maja Mataric. 2014. Socially assistive robotics for personalized education for children. In *Proceedings of the 2014 AAAI Fall Symposium Series*.

- [24] Hatice Gunes, Oya Celiktutan, and Evangelos Sariyanidi. 2019. Live human–robot interactive public demonstrations with automatic emotion and personality prediction. *Philosophical Transactions of the Royal Society B* 374, 1771 (2019), 20180026.
- [25] Hatice Gunes and Maja Pantic. 2010. Automatic, dimensional and continuous emotion recognition. *International Journal of Synthetic Emotions (IJSE)* 1, 1 (2010), 68–99.
- [26] Amanda D. Heidgerken, Gary Geffken, Avani Modi, and Laura Frakey. 2005. A survey of autism knowledge in a health care setting. *Journal of Autism and Developmental Disorders* 35, 3 (2005), 323–330.
- [27] Chien-Ming Huang and Bilge Mutlu. 2014. Learning-based modeling of multimodal behaviors for humanlike robots. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction (Bielefeld, Germany) (HRI'14)*. ACM, New York, 57–64. <https://doi.org/10.1145/2559636.2559668>
- [28] Luthffi Idzhar Ismail, Thibault Verhoeven, Joni Dambre, and Francis Wyffels. 2019. Leveraging robotics research for children with autism: A review. *International Journal of Social Robotics* 11, 3 (2019), 389–410.
- [29] Shomik Jain, Balasubramanian Thiagarajan, Zhonghao Shi, Caitlyn Clabaugh, and Maja J. Mataric. 2020. Modeling engagement in long-term, in-home socially assistive robot interventions for children with autism spectrum disorders. *Science Robotics* 5, 39 (2020).
- [30] Hifza Javed, WonHyong Lee, and Chung Hyuk Park. 2020. Toward an automated measure of social engagement for children with autism spectrum disorder—a personalized computational modeling approach. *Frontiers in Robotics and AI* 7 (2020), 43. <https://doi.org/10.3389/frobt.2020.00043>
- [31] Connie Kasari, Alexandra Sturm, and Wendy Shih. 2018. SMARTer approach to personalizing intervention for children with autism spectrum disorder. *Journal of Speech, Language, and Hearing Research* 61, 11 (2018), 2629–2640.
- [32] Ron Kohavi. 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 2 (IJCAI'95)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, 1137–1143.
- [33] Barry Kort, Rob Reilly, and Rosalind W. Picard. 2001. An affective model of interplay between emotions and learning: Reengineering educational pedagogy—building a learning companion. In *Proceedings of the IEEE International Conference on Advanced Learning Technologies*. IEEE, 43–46.
- [34] Divesh Lala, Koji Inoue, Pierrick Milhorat, and Tatsuya Kawahara. 2017. Detection of social signals for recognizing engagement in human-robot interaction. *arXiv preprint arXiv:1709.10257* (2017).
- [35] Mark R. Lepper and Ruth W. Chabay. 1988. Socializing the intelligent tutor: Bringing empathy to computer tutors. In *Learning Issues for Intelligent Tutoring Systems*. Springer, 242–257.
- [36] Daniel Leyzberg, Aditi Ramachandran, and Brian Scassellati. 2018. The effect of personalization in longer-term robot tutoring. *ACM Transactions on Human-Robot Interaction (THRI)* 7, 3 (2018), 1–19.
- [37] Maja J. Mataric and Brian Scassellati. 2016. Socially assistive robotics. In *Springer Handbook of Robotics*. Springer, 1973–1994.
- [38] Roxanna Pakkar, Caitlyn Clabaugh, Rhianna Lee, Eric Deng, and Maja J. Mataric. 2019. Designing a socially assistive robot for long-term in-home use for children with autism spectrum disorders. In *Proceedings of the 2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 1–7.
- [39] Silke Paulmann, Martin Bleichner, and Sonja Kotz. 2013. Valence, arousal, and task effects in emotional prosody processing. *Frontiers in Psychology* 4 (2013), 345. <https://doi.org/10.3389/fpsyg.2013.00345>
- [40] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. Scikit-learn: Machine learning in Python. *The Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [41] Rosalind W. Picard. 1997. *Affective Computing*. MIT Press, Cambridge, MA.
- [42] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y. Ng. 2009. ROS: An open-source robot operating system. In *Proceedings of the ICRA Workshop on Open Source Software*, Vol. 3. Kobe, Japan, 5.
- [43] Ognjen Rudovic, Jaeryoung Lee, Miles Dai, Björn Schuller, and Rosalind W. Picard. 2018. Personalized machine learning for robot perception of affect and engagement in autism therapy. *Science Robotics* 3, 19 (2018).
- [44] Ognjen Rudovic, Jaeryoung Lee, Lea Mascarell-Maricic, Björn W. Schuller, and Rosalind W. Picard. 2017. Measuring engagement in robot-assisted autism therapy: A cross-cultural study. *Frontiers in Robotics and AI* 4 (2017), 36. <https://doi.org/10.3389/frobt.2017.00036>
- [45] Ognjen Rudovic, Hae Won Park, John Busche, Björn Schuller, Cynthia Breazeal, and Rosalind W. Picard. 2019. Personalized estimation of engagement from videos using active learning with deep reinforcement learning. In *Proceedings of the 2019 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR'W)*. IEEE, 217–226.
- [46] Ognjen Rudovic, Yuria Utsumi, Jaeryoung Lee, Javier Hernandez, Eduardo Castelló Ferrer, Björn Schuller, and Rosalind W. Picard. 2018. CultureNet: A deep learning approach for engagement intensity estimation from face images

- of children with autism. In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 339–346.
- [47] Jyotirmay Sanghvi, Ginevra Castellano, Iolanda Leite, André Pereira, Peter W. McOwan, and Ana Paiva. 2011. Automatic analysis of affective postures and body motion to detect engagement with a game companion. In *Proceedings of the 6th International Conference on Human-Robot Interaction*. 305–312.
- [48] Brian Scassellati, Henny Admoni, and Maja Matarić. 2012. Robots for use in autism research. *Annual Review of Biomedical Engineering* 14 (2012).
- [49] Brian Scassellati, Laura Boccanfuso, Chien-Ming Huang, Marilena Mademtzi, Meiying Qin, Nicole Salomons, Pamela Ventola, and Frederick Shic. 2018. Improving social skills in children with ASD using a long-term, in-home social robot. *Science Robotics* 3, 21 (2018).
- [50] Elaine Short, Katelyn Swift-Spong, Jillian Greczek, Aditi Ramachandran, Alexandru Litoiu, Elena Grigore, David Feil-Seifer, Samuel Shuster, Jin Lee, Shaobo Huang, Svetlana Levonisova, Sarah Litz, Jamy Li, Gisele Ragusa, Donna Spruijt-Metz, Maja Matarić, and Brian Scassellati. 2014. How to train your dragonbot: Socially assistive robots for teaching children about nutrition through play. In *Proceedings of the IEEE International Workshop on Robot and Human Interactive Communication* 2014. <https://doi.org/10.1109/ROMAN.2014.6926371>
- [51] Samuel Spaulding and Cynthia Breazeal. 2019. Frustratingly easy personalization for real-time affect interpretation of facial expression. In *Proceedings of the 2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 531–537.
- [52] Samuel Spaulding, Goren Gordon, and Cynthia Breazeal. 2016. Affect-aware student models for robot tutors. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. 864–872.
- [53] Mary E. Stewart, Natalie Russo, Jennifer Banks, Louisa Miller, and Jacob A. Burack. 2009. Sensory characteristics in ASD. *McGill Journal of Medicine: MJM* 12, 2 (2009).
- [54] Baochen Sun, Jiashi Feng, and Kate Saenko. 2015. Return of frustratingly easy domain adaptation. *arXiv preprint arXiv:1511.05547* (2015).
- [55] Garrett Wilson and Diane J. Cook. 2020. A survey of unsupervised deep domain adaptation. *ACM Trans. Intell. Syst. Technol.* 11, 5, Article 51 (July 2020), 46 pages. <https://doi.org/10.1145/3400066>
- [56] Beverly Woolf, Winslow Burleson, Ivon Arroyo, Toby Dragon, David Cooper, and Rosalind Picard. 2009. Affect-aware tutors: Recognizing and responding to student affect. *IJLT* 4 (01 2009), 129–164. <https://doi.org/10.1504/IJLT.2009.028804>
- [57] Gloria Zen, Enver Sangineto, Elisa Ricci, and Nicu Sebe. 2014. Unsupervised domain adaptation for personalized facial emotion recognition. In *Proceedings of the 16th International Conference on Multimodal Interaction*. 128–135.

Received December 2020; revised October 2021; accepted January 2022