

Continual Learning for Affective Robotics: Why, What and How?

Nikhil Churamani*, Sinan Kalkan[†] and Hatice Gunes*

Abstract—Creating and sustaining closed-loop dynamic and social interactions with humans require robots to *continually* adapt towards their users’ behaviours, their affective states and moods while keeping them engaged in the task they are performing. Analysing, understanding and appropriately responding to human nonverbal behaviour and affective states are the central objectives of affective robotics research. Conventional machine learning approaches do not scale well to the dynamic nature of such real-world interactions as they require samples from stationary data distributions. The real-world is not stationary, it changes continuously. In such contexts, the *training* data and learning objectives may also change rapidly. Continual Learning (CL), by design, is able to address this very problem by learning incrementally. In this paper, we argue that CL is an essential paradigm for creating fully adaptive affective robots (*why*). To support this argument, we first provide an introduction to CL approaches and what they can offer for various dynamic (interactive) situations (*what*). We then formulate guidelines for the affective robotics community on how to utilise CL for perception and behaviour learning with adaptation (*how*). For each case, we reformulate the problem as a CL problem and outline a corresponding CL-based solution. We conclude the paper by highlighting the potential challenges to be faced and by providing specific recommendations on how to utilise CL for affective robotics.

I. INTRODUCTION

With advances in Artificial Intelligence (AI) and Human-Robot Interaction (HRI), intelligent robotic systems are becoming ubiquitous in human life. Moving beyond assisting in industrial tasks that require high precision and accuracy, these robots are now becoming an integral part of our daily lives in the form of assistants, tutors and even companions [1] capable of sensing their users and supporting them through social interactions, with the ultimate goal of fostering their cognitive and socio-emotional well-being. Understanding human socio-emotional signals to enhance HRI, forms the central focus of affective robotics research [2], [3], which is a challenging research topic [4], [5] still in its infancy. These skills are important for robots to provide physical and social support to human users and to engage in and sustain long-term interactions with them in a variety of application domains that require human-robot interaction, including healthcare, education, entertainment, amongst others.

The main challenge for affective robotics is understanding the underlying mechanisms of human behaviour in real

life situations and how to model these mechanisms for the embodiment of naturalistic, human-inspired behaviours in robots. Addressing this challenge successfully requires an understanding of the essential components of social interaction, including nonverbal behavioural cues such as interpersonal distance, body position and posture, arm and hand gestures, head and facial gestures, gaze, silences, vocal outbursts, and their dynamics [6]. To create truly intelligent social robots, these cues need to be interpreted to form an understanding of higher-level phenomena including first-impressions, social roles, interpersonal relationships, focus of attention, synchrony, affective states and emotions [7], personality and engagement, and in turn, manifest optimal behaviours to express these through robotic platforms in an appropriate and timely manner [8], [9]. To add to this challenge, social robots are expected to be sensitive to individual differences (due to culture, gender or personality, among other factors) in how humans manifest socio-emotional behaviours, offering a naturalistic and engaging interaction experience personalised to each user [10], [11].

Although the current (deep) learning-based approaches provide high performance on affect recognition and classification benchmarks (see, e.g., [12]–[14] for an overview), they are not able to translate this performance to real-world situations where robots need to dynamically interact with different users. The development cycle for most learning-based approaches follows a fixed transition from first being trained in isolation on a ‘large enough’ dataset with high variability and then being applied to different real-world applications. With the majority of the existing affect datasets capturing relatively *controlled* expressions recorded in fixed settings, generalisation to real-world scenarios becomes problematic [15]. Even when evaluating affect *in-the-wild* [16], these models follow a similar development cycle, with limited adaptability in their application towards capturing differences in individual expressions [17]–[19].

Affective robotics needs to adopt socio-emotional perception models that not only generalise to real-world application scenarios but also personalise towards individual users and adapt to their context (for example, user and task attributes, and the environment as illustrated in Fig. 1). Additionally, they also require learning mechanisms that can adapt to dynamic interaction contexts, in complex real-world situations. Conventional Machine Learning (ML) focuses on modelling a static data distribution, with all the data for a task available a priori, making these approaches unsuitable or at the least, inefficient in real-world interactions where data distributions shift with each user or task. Continual Learning (CL) research [20], [21] aims to address this very problem of long-term adaptability in agents, enabling them to learn with incrementally acquired data as they interact with their environment. Although commonly applied to learn objects

*N. Churamani and H. Gunes are with the Department of Computer Science and Technology, University of Cambridge, United Kingdom. {nikhil.churamani, hatice.gunes}@cl.cam.ac.uk

[†]S. Kalkan is with the Department of Computer Engineering, Middle East Technical University, Ankara, Turkey, and currently a visiting researcher at the Department of Computer Science and Technology, University of Cambridge, United Kingdom. skalkan@metu.edu.tr

N. Churamani and H. Gunes are supported by EPSRC under grants: EP/R513180/1 (2107412) and EP/R030782/1. S. Kalkan is supported by Scientific and Technological Research Council of Turkey (TÜBİTAK) through BİDEB 2219 International Postdoctoral Research Scholarship Program.

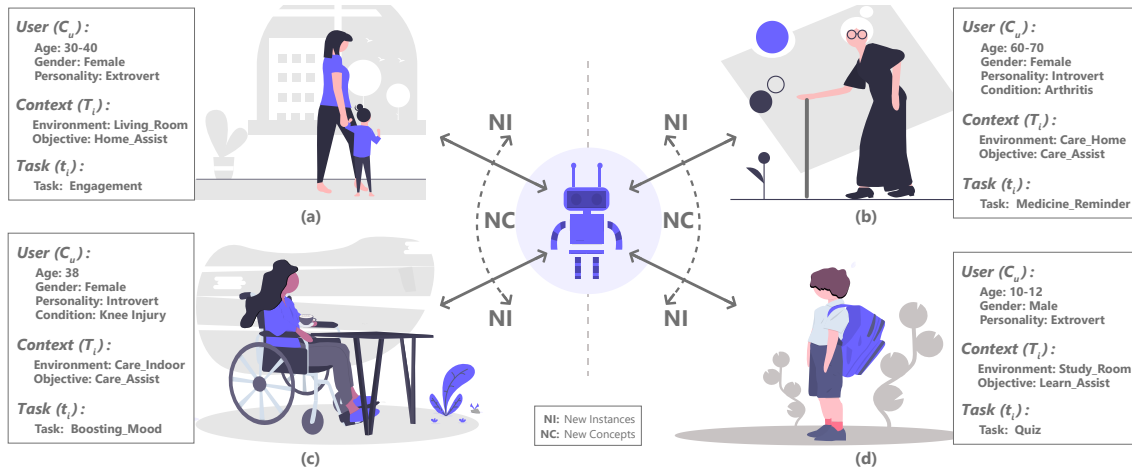


Fig. 1. Robot adapting its **Perception** (left; (a), (c)) and **Behaviour** (right; (b), (d)) interacting with users under varying contexts. User (C_u) and context-based (T_i) attributes personalise robot learning in tasks (t_i). NI updates enable user-specific adaptation; NC updates enable generalisation.

or task-based learning [21], [22], the learning principles of CL can be applied to affective robotics, learning to perceive users' socio-emotional behaviours and states over repeated interactions [18], [19], [23] as well as to generate appropriate behaviours [24]. This can be helpful for affective robots not only to understand present responses but also to predict the future socio-emotional behaviours of their users.

In light of the above, we argue that CL is an essential paradigm for creating fully adaptive affective robots (*why*). To support this argument, we first introduce CL as a learning paradigm, providing a general outline of CL approaches and their learning settings (*what*). We then re-formulate *personalisation* and learning *context-driven* interaction behaviour in affective robotics as a CL problem. Finally, we present guidelines for developing CL solutions for affective robotics, discussing the potential challenges as well as opportunities that lie ahead and how CL can offer solutions to these (*how*).

This paper complements the discussion put forth in other survey articles that effectively summarise CL literature for neural networks [21] and robotics [22] research, respectively. We add to this discussion by providing specific recommendations for adopting CL for affective robotics research, cross-fertilising insights from affective computing, robotics and human-robot interaction fields.

II. THE CONTINUAL LEARNING PARADIGM

A. Definition

The ability of agents to *continually* learn and adapt throughout their lifetime, acquiring new information while retaining the previously learnt knowledge, is termed as *Lifelong* or *Continual Learning (CL)* [20], [21]. This is particularly beneficial for agents that interact with uncertain and changing environments, for example, environments that include interactions with humans.

In CL problems, observations ($X \times Y$) are assumed to follow an infinite sequence of unknown distributions, $\mathcal{D} = \{D_1, \dots, D_N\}$. At timestep i , the agent obtains a training set Tr_i following distribution D_i to learn a task with label t in the form of a prediction function: $h^*(x, t)$. Lesort et al. [22] formulate a CL algorithm A_i^{CL} that learns a general (target) prediction model h^* as follows:

$$\forall D_i \in \mathcal{D}, A_i^{CL} : \langle h_{i-1}, Tr_i, M_{i-1}, t_i \rangle \rightarrow \langle h_i, M_i \rangle, \quad (1)$$

where h_i is the current model hypothesis, M_i is the memory storing previous training samples up to time-step i ; t_i is the task label; Tr_i is the training set of examples $e_j^i = \langle x_j^i, y_j^i \rangle$ with $j \in [1 \dots, m]$ drawn from the current data distribution D_i . For each Tr_i , A_i^{CL} adapts its model hypothesis acquiring this new information ($h_{i-1} \rightarrow h_i$), at the same time, updating its memory to represent past learning ($M_{i-1} \rightarrow M_i$). We adopt and update these notations to reformulate affective robotics challenges as CL problems (see Section III).

B. Why Continual Learning?

A major challenge faced by conventional ML models, whether shallow or deep, is their *applicability to real-world* interactions. Given the unpredictability of the real-world, these models constantly encounter novel information and tasks, requiring them to adapt their learning. However, they are not able to integrate this new information *on-the-fly* without retraining (partially or from scratch). As the agent acquires data incrementally (and sequentially as in the case of *online CL* settings [25]) through interactions with the environment, adapting to such dynamic conditions becomes computationally intractable for conventional ML models.

Any adaptation or learning is achieved at the cost of previous knowledge being forgotten or 'overwritten' [26], leading to *catastrophic forgetting* [27]. Gradient-based ML paradigms, in particular, rely on the assumption that training samples are *independently and identically drawn (i.i.d.)* from a training set. This assumption is violated in real-world conditions [28] where data is available only incrementally. As the model learns new tasks, its performance on previously learnt tasks progressively deteriorates [29]. Conventional ML models may also experience capacity saturation, where, as the agent acquires more information, and adapts to this new knowledge, its overall capacity to represent and preserve knowledge saturates [30]. This can result from the complexity of the model not being enough to retain information or the learnt feature representations not sufficient to distinguish between the learnt tasks [31].

C. An Overview of Existing Approaches

Although CL approaches may also employ deep neural architectures, they are designed to equip agents with learn-

ing capabilities that acquire and integrate new information without interfering with previously learnt knowledge (see e.g. [20], [21] for a review). This is achieved by regulating model updates, storing and replaying already seen information to simulate i.i.d. conditions or dynamically expanding the models to compensate for new information. CL approaches can be summarised under four categories based on the strategy employed for balancing novel vs. past learning.

1) *Regularisation-based Approaches*: Regularisation is a family of techniques for learning models to guard against over-fitting. For CL, regularisation can reduce *destructive interference*, preventing newly learnt tasks from ‘*overwriting*’ previous information. This is achieved by either freezing parts of the model that correspond to learnt information [32] and updating only newly added parameters [33], penalising weight-updates that deteriorate performance on previously learnt tasks [34], or prioritising weight-updates for different parameters given their relevance to different tasks [30], [35]. Constraining weight-updates of the model (or parts of it), preserves prior knowledge, avoiding catastrophic changes to learnt parameters.

2) *Rehearsal-based Approaches*: To mitigate forgetting in incremental learning, a straightforward approach can be to physically store the encountered data in memory and regularly replay it (known as *rehearsal*), interleaved with new samples [36]. This replicates offline i.i.d. settings as the model is trained on mixed batches of data consisting of samples from all the classes (or tasks). Although this works when the number of tasks is small [37], it does not scale well as the number of tasks increases. In case of high-dimensional data (e.g., images) with a large number of classes, physically storing and replaying training samples becomes computationally intractable. A generative or probabilistic model may be used to learn data statistics to draw pseudo-samples [38] from the memory (known as *pseudo-rehearsal*), reducing the cost of these models significantly [28], [39], [40]. Yet, as the number of tasks increases, it becomes harder to train the generative/probabilistic models to represent all the tasks.

3) *Dynamic Architectures*: As the complexity of the data and tasks increases, models trained with the previously described approaches are not able to scale up. This is due to capacity saturation - i.e., due to weights frozen from previously learnt tasks or memory-exhaustion from storing samples for rehearsal [30]. To alleviate this problem, additional neural resources can be allocated to extend the capacity, either by expanding trainable parameters [41] or allowing the architecture itself to grow [42] to account for the increased complexity. Starting with a relatively simple architecture, the model is extended by allocating additional neurons [42]–[45] or network layers [32], [46], [47] as and when required. This growth can be regulated using the model’s performance on previously learnt tasks, its neural activation in response to data samples or the contribution of existing parameters towards solving new tasks. Despite the additional overhead of adding new neural resources, these models are shown to work well in mitigating catastrophic forgetting, enabling continual learning of information [21].

4) *Neuro-inspired Approaches*: An enhanced understanding of Complementary Learning Systems (CLS) [27] in the

human brain has inspired a new approach for CL [21], [48]. This approach implements learning over multiple memory models, each of which adapts to learning at different stages, alleviating catastrophic forgetting. While an *episodic memory* is employed to realise active learning of novel experiences for the agent, a *semantic memory* responds to long-term retention of information by slowly replaying episodic experiences. This replay of experiences is facilitated by generative or probabilistic models [19], [40], [49], [50] that transfer experiences between the different memories using *pseudo-rehearsal*. Other CLS-based approaches employ self-organising neural models for encoding sensory experiences in the memory [31], [37], [42]. These models regulate levels of neural-growth based on the capability of the model to integrate new information and retain previous knowledge.

D. Learning Types

As CL models aim to incrementally integrate continuous sequences of new data samples while preserving previous knowledge, this can result in three main learning types based on the nature and availability of sequential data [51].

(1) *New Instances (NI)*: The model receives samples from all the tasks in the very first instance, and all incoming data samples adhere to these seen tasks or *concepts*. The model does not learn a new task but instead learns variation in the data distribution for already learnt tasks.

(2) *New Concepts (NC)*: For each sequential batch, the model receives samples only from a new task or concept and is evaluated on its ability to learn this new task while still maintaining its performance on the previously learnt tasks.

(3) *New Instances and Concepts (NIC)*: The model not only receives more samples for already learnt tasks but also needs to learn new tasks, with each sequential batch of input (that is, a combination of NI and NC).

E. Model Evaluation

The dynamic nature of the CL paradigm requires different evaluation strategies from those used for conventional ML models [22] measuring how well the model adapts to changes and can cope with the challenges outlined in Section II-B. These evaluations focus on answering questions such as: How much each task contributes to learning a new task (as opposed to learning that new task from scratch)? Or how much the performance of a previously learnt task [26] worsens? This relates to assessing the model’s ability to retain the previously learnt information and transfer experience to new learning as much as possible. And what is the accuracy of the model on all the data observed so far? This is interpreted *relative* to the accuracy of the corresponding conventional model trained with all the observed data.

III. CONTINUAL LEARNING FOR AFFECTIVE ROBOTS

For robots to effectively interact with humans, it is important that they proactively participate in the human *affective loop* [2]. This requires them to not only perceive and analyse human socio-emotional behaviours across sensory modalities but also learn to respond in a manner conforming to the context and the evolution of the interaction [9]. This is particularly beneficial when using robots in interventions

with sensitive user groups such as providing care for the elderly and assisting children in learning [52], [53].

Consequently, the desiderata from affective robots (as exemplified in Fig. 1) include (i) perception models that are robust to real-world interaction settings [16] while being sensitive to each individual’s socio-emotional behaviours [18], [19], [54], and (ii) generation of context-specific behaviour attributing both the users’ behaviour as well as the interaction settings [8], [55]. As robots acquire data about their environment incrementally and sequentially by interacting with different users, they need to be able to learn and integrate this information *on-the-fly*. Hence, we argue that adopting CL as a learning paradigm, in particular using *online CL methodologies* [25], is crucial for affective robotics and HRI research. The ability to balance novel vs. past learning gives CL models an advantage over conventional ML solutions. With this in mind, in this section, we adapt the theoretical definition of CL algorithms (see Section II-A) to formulate *personalised affect perception* and *context-specific behaviour generation* in affective robots as CL problems.

A. Personalised Affect Perception

Personalisation in this context is the ability of an agent to adapt to the *socio-emotional behaviour* of a user during interactions (see Fig. 1). This requires the agent to adapt its perception model with each user, accounting for individual differences in nonverbal behaviour and expression [17], [19]. This adaptation needs to adhere to both at *individual level* for learning to be sensitive towards the individual behaviour of a user and *across individuals* for generalising its learning to interact with different users.

Continual Learning Formulation: We formulate personalised affect perception (P^{CL}) as a CL problem, adapting Eq. 1 to depict the requirements from affective robotics. Following such a formulation not only allows for perception models to adapt to individual users but also enable generalisation to novel experiences in changing interaction settings. P^{CL} can be formalised as follows:

$$\forall u \in \mathcal{U}, \forall i \in \mathcal{I}, P_{u,i}^{CL} : \langle h_{i-1}, Tr_i, M_{i-1}, C_u, t_i \rangle \rightarrow \langle h_i, M_i \rangle, \quad (2)$$

where h_i is the current affect perception model and M_i is the memory storing previously seen training samples up to interaction state i . t_i represents the current expression recognition task label (for example, expression category) summarising the affective state of user u , C_u is the set of user-specific attributes (for example, user preferences, contextual attributions or personality-specific traits) that may be known, and Tr_i is the current training data (for example, face images or speech signals) obtained during the interaction.

Continual Learning Scenarios: Throughout an interaction, a robot may either observe multiple samples for the same expression/social signal (type *NI*) or observe a user under different socio-emotional contexts (type *NC*) requiring its perception model to not only be robust to the variation in expressiveness for a learnt expression but also learn different expressions of the user. As a robot interacts with multiple users, this individual-level learning ($P_{u,i}^{CL}$) is aggregated across multiple users resulting in the overall perception

model (P^{CL}) for the robot. This aggregation can be achieved by maintaining several individual models that can be *loaded* upon identifying the user [18] or by learning semantic representations that aggregate robot’s knowledge to generalise learning [19], [56].

Existing approaches: Most personalised affect perception approaches focus on contextual attributions for each user [53], perform selective weighting of subject-specific data [17], or apply unsupervised clustering of person-specific feature representations [56] to adapt to individual users. Despite their success on benchmark evaluations, they suffer from the same problems as conventional ML algorithms (see Section II-B), as data is only acquired during interactions. CL principles of learning with incrementally acquired data have been applied in some studies for personalised affect perception either focusing on learning individual affective memories [18] or applying CLS-based learning [19], adapting with each user. Building memory representations (M_i) as they acquire more data, these can personalise towards each user by remembering past interactions (by *rehearsal*) while using this memory as an influence on the learning of novel expressions (t_i). Yet, they do not take into account user-specific contextual attributes (C_u) that can improve learning.

B. Context-specific Robot Behaviour Generation

Recent works on learning robot behaviour generation investigate the role of affect for modulation in Reinforcement Learning (RL) algorithms, either as an intrinsic motivation [24] to drive robot learning or as an evaluation of human affective behaviour [57] to learn optimal interaction policies. As learning dynamics are dependent on the environment and how the agent dynamically interacts with it, most RL formulations can be directly compared to CL settings [22]. This can be seen in most of the popular RL algorithms that either implement the use of external memory (*replay buffer*) to store and *rehearse* previously seen examples [58], [59], consolidate knowledge using multiple agents in parallel [60] or constrain shifts in learning [61] to improve learning across different tasks. Yet, these require a lot of training data to yield good results which may not be possible while interacting with humans. Interactive RL (IRL) techniques, with the *human in the loop* [62]–[64], offer potent solutions for embedding such dynamic real-time adaptation in robots. Receiving feedback directly from the user speeds up convergence, boosting learning in the model.

Thus, we propose that complementing aspects of IRL and CL should be combined. That is, learning with human feedback combined with knowledge rehearsal and distillation to control shifts in learning can prove useful for affective robots. Approaching learning interaction-driven behaviour in HRI from a CL perspective can enable robots to learn context or task-specific behaviours from their experiences with different users (see Fig. 1).

Continual Learning Formulation: Following Eq. 1, we formalise behavioural learning ($B_{u,i}^{CL}$) as a CL problem where the agent learns optimal interaction behaviour in an interaction state i while interacting with a particular user u :

$$\forall u \in \mathcal{U}, \forall i \in \mathcal{I}, B_{u,i}^{CL} : \langle h_{i-1}, Tr_i, M_{i-1}, C_u, T_i, t_i \rangle \rightarrow \langle h_i, M_i \rangle, \quad (3)$$

where h_i is the current behaviour learning model for the robot (for example, an RL model) and M_i is the memory or *experience buffer* storing past interaction samples up to the current state i in the overall interaction \mathcal{I} , t_i is the current task of the robot in interaction state i for which it needs to learn the optimal response, C_u is the set of user-specific contextual attributes that can influence robot behaviour, T_i are the task-specific contextual attributes derived from the environment or the rules that govern the entire interaction. Tr_i is the current training data acquired by the robot, for example, affective feedback from the user, robot’s state in its environment or the sensory evaluations resulting from its perception model P^{CL} , during its interaction with user u .

Continual Learning Scenarios: Learning to interact with different users under a similar context may constitute instance-level (type NI) adaptation where the robot becomes robust to the variation in human behaviour under similar contextual settings, learning how to respond to them. The same robot, however, might need to interact with users under different contexts or learn to perform different tasks, requiring concept-level (type NC) adaptation. The proposed formulation, using C_u and T_i allows for adaptation to these scenarios by providing the relevant contextual information.

Existing approaches: Most behaviour generation approaches focus on generating robot emotional expressions [55] as back-channels to support conversational HRI [65]. Only a few focus on learning task-oriented behaviours [66] during interactions and these are limited to generating relatively low-level atomic behaviours such as verbal utterances or atomic body gestures. High-level behaviours, such as context-dependent interaction switching, are mostly handled by expert planners. CL has the potential to enable robots to learn dynamic interaction behaviours both at an atomic level [22] and at context-level. This can be achieved by first learning to extract state representations from robot perception (using P_u^{CL}) and then, learning behaviour policies forming contextualised task-representations (using T_i) that enable robots to handle complex interaction scenarios.

IV. OPPORTUNITIES AND CHALLENGES

A. Opportunities

To enable social robots to become human companions, they need to be equipped with *continually* adapting perception and behaviour models that can cater to the changing dynamics of real-world interactions. Placing robots in household settings would mean that there cannot be broad assumptions made about user demographics or HRI contexts governing interactions. For example, the same robot may be required to care for the elderly, assist adults in day-to-day chores as well as be a learning companion for the young (see Fig. 1). CL offers a learning paradigm that is very suitable for such affective robotics applications. Equipped with state-of-the-art CL-based learning models, right out of the box, such robots will be able to adapt and learn with each user, while continuously improving their socio-emotional intelligence.

B. Challenges

There are fundamental challenges that need to be addressed for a successful application of CL for affective

robotics. These might arise from how the robot gathers and manages data, obtains ground truth evaluations for user-specific socio-emotional behaviour or learns context-specific task representations. Below we discuss some of these challenges in detail:

1) *Gathering Person-specific Data:* As the only source of data for the robot is interactions with a user, it might require a lot of interactions before the model can successfully adapt, negatively impacting the initial user experience. *Adversarial training* [67] can be used as a mitigation strategy as it enables simulation of person-specific data [19], [68], [69] allowing the robot to *imagine* interactions with a user [19] and learn from such *imagined contact* [70]. However, even such models need large amounts of training data before reasonable person-specific samples can be generated.

2) *Obtaining Ground Truth Data:* Affective interactions can be highly subjective. Obtaining ground truth for the data sequentially through interactions is challenging and varies from user-to-user. Many CL approaches have looked at this problem from an object recognition point-of-view (see [22], [25] for a review) and tackled it by using self-supervision mechanisms driven by curiosity or novelty detection to aid learning. Alternatively, unsupervised clustering of information and applying Hebbian-like learning [18], [19] can help improve the robustness of the model. However, there is a need for more established approaches taking inspiration from findings from human interaction studies.

3) *Multi-Task Learning without Task Boundaries:* Real-world human interactions are fluid and may toggle between different contexts. CL approaches deal with such sudden context shifts rather robustly by sensing and adapting to changing data distributions arising from different tasks [22], [30]. Yet, in affective HRI, this may not be as straightforward as the change may be too subtle or the contextual attributions of different tasks may overlap, without clear and distinct task boundaries. Hence, affective robots need stronger context-awareness to learn different context-dependent task representations [71], at the same time.

4) *Robot Hardware and Memory:* Integrating dynamic adaptation in robots requires the robot to not only store the gathered data in memory but also run comprehensive computations to update its learning. Despite technological advancements that make computation cheaper, robots are still configured with relatively ‘light-weight’ hardware capabilities. Most CL approaches that focus on realising online learning capabilities in agents [25] reduce the memory foot-print of the models by computationally modelling inherent data statistics using a generative or probabilistic model (known as *pseudo-rehearsal*), making a trade-off between the on-board storage and computational resources. More recently, with several cloud-based services (for example, Robotics as a Service (RaaS) platforms such as Amazon AWS RoboMaker) providing a host of solutions, some of the computation and memory load can be offset over the cloud, facilitating real-time adaptation in the models. Yet, latency-ridden cloud-based computations in complex interaction scenarios can negatively impact the HRI experience.

TABLE I
RECOMMENDATIONS FOR AFFECTIVE ROBOTICS

Recommendation	Why is this important and needed?	How can this be achieved?
Acquire person-specific data	Adapting learning models to individual preferences requires large amounts of data that can only be sourced through interactions with users.	(1) Conduct introductory HRI rounds to enable the robot to collect additional data about the user. (2) Leverage adversarial learning to train a generative model to simulate additional person-specific data.
Obtain normative baselines	The robot needs to know the behavioural <i>norm</i> for each user against which deviations can be observed. Deviations help identify shifts in user socio-emotional behaviours and infer changes in interaction context.	(1) Conduct interactions under contextually inert (neutral) situations during introduction rounds. (2) Use the (subtle) deviations from this baseline, given the interaction context, to analyse shifts.
Extract semantic associations	Adapting the learning for a large number of users is computationally intractable. Learning models will get saturated, not able to remember previous information or learn with new individuals.	(1) Form user groupings, using person-specific attributes (C_u in Eq. 2-3) to learn group-based adaptations. (2) Use unsupervised data clustering to facilitate learning semantic groupings of users.
Learn contextual affordances	Interactions are driven by context and humans switch between contexts without clear boundaries. Contextual attributions may not always be implicit and need to be learnt separately	(1) Learn context-aware embeddings to distinguish between task boundaries. (2) Use contextual affordances (e.g. T_i in Eq. 3) to facilitate smooth switching between affective HRI contexts.
Balance memory with computation	The memory-computation trade-off needs to be considered w.r.t the application domain. Adding more memory facilitates rehearsal of past knowledge, while additional computation power improves adaptation to novel experiences.	(1) Use generative models for pseudo-rehearsal to reduce model's memory foot-print. (2) Offload part of the computation/memory load to Robotics as a Service (RaaS)-based solutions to balance old vs. novel learning.
Allow controlled forgetting	When learning is continuous, redundant information in the memory/model, is not released, hindering learning capacity of the model.	(1) Utilise forgetting mechanisms (inspired by biological organisms) on unused memory locations or parts of the model, to learn new knowledge.
Use multiple performance metrics	Benchmark evaluations from conventional ML and CL perspectives are needed for reproducibility and fairness guarantees, and to evaluate model's robustness to dynamic shifts in data distributions.	(1) Report CL performance metrics (Section II-E), along with the classification metrics of F-measure and AUC-ROC scores or reward-function dynamics for behaviour learning.

V. RECOMMENDATIONS AND CONCLUDING REMARKS

In this paper, we have argued that CL is an essential paradigm for affective robotics. We discussed the *why*, *what* and *how* of this argument and provided a CL formulation for personalised affect perception and context-specific robot behaviour generation. From these, we distilled a set of concrete recommendations in Table. I across several dimensions that are crucial to consider when integrating CL into affective robotics - ranging from data to scalability to a large number of users and a long lifespan, as well as memory with respect to computing limitations, and benchmarking and evaluation.

It is our genuine hope that these discussions, formulations and recommendations will create a stepping stone for social robotics and HRI studies that consider taking a CL approach to building fully autonomous and adaptive robots that are purposeful and engaging in their interactions with their human users.

REFERENCES

- [1] E. van Oost *et al.*, "Towards a Sociological Understanding of Robots as Companions," in *Human-Robot Personal Relationships*, M. H. Lamers *et al.*, Eds. Springer Berlin Heidelberg, 2011, pp. 11–18.
- [2] A. Paiva *et al.*, "Emotion Modelling for Social Robots," *The Oxford Handbook of Affective Computing*, pp. 296–308, 2014.
- [3] T. Fong *et al.*, "A survey of socially interactive robots," *Robotics and Autonomous Systems*, vol. 42, no. 3-4, pp. 143–166, Mar. 2003.
- [4] C. Breazeal, "Emotion and Sociable Humanoid Robots," *International Journal of Human-Computer Studies*, pp. 119–155, 2003.
- [5] L. Bishop *et al.*, "Social robots: The influence of human and robot characteristics on acceptance," *Paladyn, Journal of Behavioral Robotics*, vol. 10, no. 1, pp. 346 – 358, 2019.
- [6] H. Gunes *et al.*, "Emotion representation, analysis and synthesis in continuous space: A survey," in *Face and Gesture 2011*, pp. 827–834.
- [7] K. Scherer, "Emotion," in *Introduction to social psychology: A European perspective*. Blackwell Oxford, 2000, pp. 151–191.
- [8] R. Kirby *et al.*, "Affective social robots," *Robotics and Autonomous Systems*, vol. 58, no. 3, pp. 322–332, 2010.
- [9] K. Dautenhahn, "Robots we like to live with?! - a developmental perspective on a personalized, life-long robot companion," in *IEEE International Workshop on Robot and Human Interactive Communication (Ro-Man)*, Sept 2004, pp. 17–22.
- [10] M. K. Lee *et al.*, "Personalization in HRI: A longitudinal field experiment," in *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2012, pp. 319–326.
- [11] M. Dziergwa *et al.*, "Long-Term Cohabitation with a Social Robot: A Case Study of the Influence of Human Attachment Patterns," *IJRS*, vol. 10, no. 1, pp. 163–176, Jan 2018.
- [12] S. Poria *et al.*, "A review of affective computing: From unimodal analysis to multimodal fusion," *Information Fusion*, vol. 37, pp. 98–125, 2017.
- [13] C. A. Corneanu *et al.*, "Survey on rgb, 3d, thermal, and multimodal approaches for facial expression recognition: History, trends, and affect-related applications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 8, pp. 1548–1568, Aug 2016.
- [14] S. Li *et al.*, "Deep facial expression recognition: A survey," *IEEE Transactions on Affective Computing*, March 2020.
- [15] M. F. Valstar *et al.*, "Spontaneous vs. posed facial behavior: Automatic analysis of brow actions," in *Proceedings of the International Conference on Multimodal Interfaces (ICMI)*. ACM, 2006, pp. 162–170.
- [16] S. Zafeiriou *et al.*, "Aff-wild: Valence and arousal 'in-the-wild' challenge," in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, July 2017, pp. 1980–1987.

- [17] W. Chu *et al.*, “Selective transfer machine for personalized facial expression analysis,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 3, pp. 529–545, March 2017.
- [18] P. Barros *et al.*, “A Personalized Affective Memory Model for Improving Emotion Recognition,” in *Proceedings of the International Conference on Machine Learning*. PMLR, Jun 2019, pp. 485–494.
- [19] N. Churamani *et al.*, “CLIFER: Continual Learning with Imagination for Facial Expression Recognition,” in *Proceedings of the 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, 2020, pp. 693–699.
- [20] S. Thrun *et al.*, “Lifelong Robot Learning,” *Robotics and Autonomous Systems*, vol. 15, no. 1, pp. 25 – 46, 1995, the Biology and Technology of Intelligent Autonomous Agents.
- [21] G. Parisi *et al.*, “Continual Lifelong Learning with Neural Networks: A review,” *Neural Networks*, vol. 113, pp. 54–71, 2019.
- [22] T. Lesort *et al.*, “Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges,” *Information Fusion*, vol. 58, pp. 52–68, 2020.
- [23] J. Gideon *et al.*, “Progressive Neural Networks for Transfer Learning in Emotion Recognition,” in *Proc. Interspeech*, 2017, pp. 1098–1102.
- [24] T. M. Moerland *et al.*, “Emotion in reinforcement learning agents and robots: a survey,” *Machine Learning*, vol. 107, no. 2, pp. 443–480, Feb 2018.
- [25] G. I. Parisi *et al.*, “Online continual learning on sequences,” in *Recent Trends in Learning From Data*. Springer International Publishing, 2020, pp. 197–221.
- [26] D. Lopez-Paz *et al.*, “Gradient episodic memory for continual learning,” in *Proceedings of the International Conference on Neural Information Processing Systems*, ser. NIPS’17. USA: Curran Associates Inc., 2017, pp. 6470–6479.
- [27] M. McCloskey *et al.*, “Catastrophic interference in connectionist networks: The sequential learning problem,” ser. *Psychology of Learning and Motivation*. Academic Press, 1989, vol. 24, pp. 109–165.
- [28] A. Seff *et al.*, “Continual learning in generative adversarial nets,” *CoRR*, vol. abs/1705.08395, 2017.
- [29] R. Kemker *et al.*, “Measuring Catastrophic Forgetting in Neural Networks,” in *32nd AAAI Conference on Artificial Intelligence*, 2018.
- [30] J. Kirkpatrick *et al.*, “Overcoming catastrophic forgetting in neural networks,” *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [31] G. I. Parisi *et al.*, “Lifelong learning of spatiotemporal representations with dual-memory recurrent self-organization,” *Frontiers in Neurobotics*, vol. 12, p. 78, 2018.
- [32] C. Fernando *et al.*, “Pathnet: Evolution channels gradient descent in super neural networks,” *CoRR*, vol. abs/1701.08734, 2017.
- [33] A. S. Razavian *et al.*, “CNN Features Off-the-Shelf: An Astounding Baseline for Recognition,” in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2014, pp. 512–519.
- [34] Z. Li *et al.*, “Learning without forgetting,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 12, pp. 2935–2947, Dec 2018.
- [35] F. Zenke *et al.*, “Continual learning through synaptic intelligence,” in *Proceedings of the International Conference on Machine Learning*, ser. ICML’17, vol. 70. JMLR, 2017, pp. 3987–3995.
- [36] A. Robins, “Catastrophic forgetting in neural networks: the role of rehearsal mechanisms,” in *Proc. The 1st New Zealand International Two-Stream Conference on Artificial Neural Networks and Expert Systems*, Nov 1993, pp. 65–68.
- [37] A. Gepperth *et al.*, “A bio-inspired incremental learning architecture for applied perceptual problems,” *Cognitive Computation*, vol. 8, no. 5, pp. 924–934, Oct 2016.
- [38] A. Robins, “Catastrophic forgetting, rehearsal and pseudorehearsal,” *Connection Science*, vol. 7, no. 2, pp. 123–146, 1995.
- [39] C. Wu *et al.*, “Memory Replay GANs: Learning to Generate Images from New Categories Without Forgetting,” in *Proceedings of the International Conference on Neural Information Processing Systems*, ser. NIPS’18. USA: Curran Associates Inc., 2018, pp. 5966–5976.
- [40] H. Shin *et al.*, “Continual Learning with Deep Generative Replay,” in *Advances in Neural Information Processing Systems 30*. Curran Associates, Inc., 2017, pp. 2990–2999.
- [41] T. J. Draelos *et al.*, “Neurogenesis deep learning: Extending deep networks to accommodate new classes,” in *International Joint Conference on Neural Networks (IJCNN)*. IEEE, May 2017, pp. 526–533.
- [42] G. I. Parisi *et al.*, “Lifelong learning of human actions with deep neural network self-organization,” *Neural Networks*, vol. 96, no. Supplement C, pp. 137–149, 2017.
- [43] G. Zhou *et al.*, “Online incremental feature learning with denoising autoencoders,” in *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, vol. 22. PMLR, 21–23 Apr 2012, pp. 1453–1461.
- [44] J. L. Part *et al.*, “Incremental online learning of objects for robots operating in real environments,” in *2017 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, Sep. 2017, pp. 304–310.
- [45] F. I. Doğan *et al.*, “Cinet: A learning based approach to incremental context modeling in robots,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 4641–4646.
- [46] A. A. Rusu *et al.*, “Progressive neural networks,” *CoRR*, vol. abs/1606.04671, 2016.
- [47] F. I. Doğan *et al.*, “A deep incremental boltzmann machine for modeling context in robots,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2411–2416.
- [48] D. Hassabis *et al.*, “Neuroscience-Inspired Artificial Intelligence,” *Neuron*, vol. 95, no. 2, pp. 245 – 258, 2017.
- [49] N. Kamra *et al.*, “Deep generative dual memory network for continual learning,” *CoRR*, vol. abs/1710.10368, 2017.
- [50] R. Kemker *et al.*, “Fearnert: Brain-inspired model for incremental learning,” *CoRR*, vol. abs/1711.10563, 2018.
- [51] V. Lomonaco *et al.*, “CORE50: A New Dataset and Benchmark for Continuous Object Recognition,” in *Proceedings of the 1st Annual Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, vol. 78. PMLR, 13–15 Nov 2017, pp. 17–26.
- [52] K. Dautenhahn, “Socially intelligent robots: dimensions of human-robot interaction,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 362, no. 1480, pp. 679–704, Apr. 2007.
- [53] O. Rudovic *et al.*, “Personalized machine learning for robot perception of affect and engagement in autism therapy,” *Science Robotics*, vol. 3, no. 19, 2018.
- [54] G. Zen *et al.*, “Learning personalized models for facial expression analysis and gesture recognition,” *IEEE Transactions on Multimedia*, vol. 18, no. 4, pp. 775–788, April 2016.
- [55] M. FicoCELLI *et al.*, “Promoting interactions between humans and robots using robotic emotional behavior,” *IEEE Transactions on Cybernetics*, vol. 46, no. 12, pp. 2911–2923, 2016.
- [56] G. Zen *et al.*, “Unsupervised domain adaptation for personalized facial emotion recognition,” in *Proceedings of the International Conference on Multimodal Interaction (ICMI)*. ACM, 2014, pp. 128–135.
- [57] J. Broekens, “Emotion and reinforcement: Affective facial expressions facilitate robot learning,” in *Artificial Intelligence for Human Computing*. Springer Berlin Heidelberg, 2007, pp. 113–132.
- [58] T. P. Lillicrap *et al.*, “Continuous control with deep reinforcement learning,” *CoRR*, vol. abs/1509.02971, 2015.
- [59] T. Schaul *et al.*, “Prioritized experience replay,” in *International Conference on Learning Representations, ICLR*, 2016.
- [60] V. Mnih *et al.*, “Asynchronous methods for deep reinforcement learning,” in *Proc. ICML*, vol. 48. PMLR, 2016, pp. 1928–1937.
- [61] J. Schulman *et al.*, “Proximal policy optimization algorithms,” *CoRR*, vol. abs/1707.06347, 2017.
- [62] A. L. Thomaz *et al.*, “Reinforcement learning with human teachers: Understanding how people want to teach robots,” in *IEEE Ro-Man*, 2006, pp. 352–357.
- [63] F. Cruz *et al.*, “Training agents with interactive reinforcement learning and contextual affordances,” *IEEE Transactions on Cognitive and Developmental Systems*, vol. 8, no. 4, pp. 271–284, Dec 2016.
- [64] W. B. Knox *et al.*, “Training a robot via human feedback: A case study,” in *Social Robotics*. Springer, 2013, pp. 460–470.
- [65] Y. Li *et al.*, “Expressing reactive emotion based on multimodal emotion recognition for natural conversation in human–robot interaction,” *Advanced Robotics*, vol. 33, no. 20, pp. 1030–1041, 2019.
- [66] J. Hemminghaus *et al.*, “Towards adaptive social behavior generation for assistive robots using reinforcement learning,” in *Proceedings of International Conference on Human-Robot Interaction, ACM/IEEE*, 2017, p. 332–340.
- [67] J. Han *et al.*, “Adversarial Training in Affective Computing and Sentiment Analysis: Recent Advances and Perspectives,” *IEEE Computational Intelligence Magazine*, vol. 14 (2), pp. 68–81, 2019.
- [68] A. Lindt *et al.*, “Facial Expression Editing with Continuous Emotion Labels,” in *Proceedings of the International Conference on Automatic Face and Gesture Recognition (FG)*. IEEE, 2019, pp. 1–8.
- [69] D. Saez-Trigueros *et al.*, “Generating photo-realistic training data to improve face recognition accuracy,” *CoRR*, vol. abs/1811.00112, 2018.
- [70] R. Wullenkord *et al.*, “Imagine how to behave: the influence of imagined contact on human–robot interaction,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 374, no. 1771, p. 20180038, Apr. 2019.
- [71] I. Bozcan *et al.*, “Cosmo: Contextualized scene modeling with boltzmann machines,” *Robotics and Autonomous Systems*, vol. 113, pp. 132–148, 2019.