"Who's a Good Robot?!" Designing Human-Robot Teaching Interactions Inspired by Dog Training

Patrizia PACI [a,1], Ilaria TIDDI [b], Daniel PRECIADO [b] and Kim BARAKA [b]

[a] *The Open University, UK*
[b] *Vrije Universiteit Amsterdam*

Abstract. Recent work in Human-Robot Interaction (HRI) investigates the role of human users as teachers from which robots can flexibly learn new personalised skills through interaction. However, existing human-robot teaching methods remain largely unintuitive for the end user and require significant effort to adapt to the way the robot learns. This paper envisions the use of dog training methods as a starting point for HRI researchers to develop more intuitive interactions between human teachers and robot learners. We provide a design framework (called FETCH-R) aimed at guiding the conception of interactions between human teachers and robot learners inspired by dog training. This work paves the way towards the use of animal training as an inspiration to create human-robot teaching protocols that promote engagement, ease-of-use, and fosters human-robot relationships.

Keywords. Human-robot interaction, human-robot teaching, social robots, robot learning, hybrid intelligence, ethorobotics, dog training.

1. Introduction

An emerging trend within HRI is to develop teachable robotic agents that can flexibly adapt existing skills or even learn new skills from interaction with humans through natural communication (e.g. verbal and non-verbal cues). This interdisciplinary effort, called Human-Interactive Robot Learning (HIRL) [1], aims at designing algorithms that allow a physical robot to learn interactively from human inputs, which include demonstrations, feedback, advice, correction, and instructions. Unlike a typical user commands/robot responds approach, HIRL techniques allow robot behaviour to be adapted and better aligned with human needs and preferences, and give users more control over their operation without requiring technical knowledge about their inner workings [2].

HIRL existing approaches to robot learning remain largely unintuitive, requiring significant effort from the human to adapt to the way the robot learns [2]. Aspects pertaining to the teacher, such as intuitiveness, willingness to teach, teaching preferences and styles have been largely understudied. Learning from Demonstration (LfD) techniques are applied in social robotics contexts to teach a robot directly from human inputs rather than programming. Yet, a mismatch between teacher and learner's mental models and actions spaces are still reported, e.g. when visual demonstrations are received by too simplistic robotic perception systems than biological ones [2]. These discrepancies contribute to

---

[1]Corresponding Author: Patrizia Paci, patrizia.paci@open.ac.uk

make the teaching process frustrating for the user, heightening their expectations of what a robot can actually perform, and ultimately decreasing their engagement following the robot's failures. It therefore becomes essential to rethink how we should be designing interfaces that allow for seamless and intuitive teaching behaviours on the human side.

To improve users' teaching abilities and experiences towards robots, this paper proposes to follow an animal model of social interactions [3]. Humans have a propensity to keep animals as pets, hence they have developed natural and familiar ways to interact with agents who have different mental models and cognitive capabilities as them. Arguably, this inclination could work when users approach a teaching exercise with artificial agents, too. We therefore propose to exploit the familiar interactions occurring between human trainers and animal learners such as dogs. We build upon the idea that dog training principles, as well as more spontaneous human interactions towards dogs, can be transferred to HRI settings to facilitate human teaching and robot learning processes [3]. In this paper, we present a high-level interaction design framework aimed at guiding the development of human-robot interactions according to human-dog training interactions. The framework – which we called FETCH-R (Framework for Enhanced Teaching Capabilities for Humans with Robots) – encompasses elements typically encountered in canine training (e.g. attention calling, positive/negative reinforcement, etc.) and human interactions with dogs (e.g. verbal commands, body postures, gesturing, etc.). We illustrate how FETCH-R could guide the design of natural human-robot interaction flows based on dog-directed teacher cues.

2. Related Work and Challenges Ahead

In this section, we report about ways to transfer skills interactively from humans to agents. We then introduce the pioneering work in ethorobotics to scope how animal models can be used. Finally, we present a scoped rationale that motivates our approach.

2.1. HIRL tools as building blocks towards dog-inspired human-robot teaching

Natural methods to directly transfer human skills to artificial agents fall under the umbrella of interactive machine learning, with a focus on a reinforcement learning formalism [4]. The most popular methods include imitation learning, whereby the robot learns close-to-optimal policies from human demonstrations [2] or observations [5], learning from different forms of feedback (e.g. evaluative [6], corrective [7], suggestive [8]), and curriculum learning [9]), whereby tasks of increasing difficulty are successively presented to the learner. Some of these teaching signals are related to dog training. For instance, teaching through demonstrations can be related to the concept of modelling in dog training [10]. Teaching through evaluative feedback (i.e. positive and negative reinforcement) can be related to the concept of shaping [10] in dog training, a core concept in the TAMER framework proposed by [11]. Furthermore, teaching through curricula and hierarchical teaching by demonstrations [12] can be related to the concept of scaffolding in dog training (whereby a subset of learned skills help building more complex ones). Note that these teaching signals and strategies are basic building blocks of learning, not only in dogs, but also in humans. They lay the technological foundations needed to start thinking about teaching interfaces, that would allow humans to teach robots in ways that are tailored towards rich interactions inspired by dog training.

2.2. The pioneering work in ethorobotics

Reinforcement Learning has been used by trainers to shape the behaviour of animals well before the AI advent. Based on this idea, [10] proposed the use of animal training techniques based on shaping (e.g. the "clicker training") to teach unusual behaviours and sequences of actions to zoomorphic AIBO robots. The authors demonstrated that such a technique can be used to combine pre-programmed embedded behaviours to create new more complex ones, or teach singular variants of what the robot already performs.

This pioneering work paved the way for ethorobotics, an interdisciplinary area focusing on the design of companion robots through a human-animal relationship lens [13]. Based on the reflection that owning pets is a universal trait in the human species, [3] analyses human-animal partnerships as a source to design familiar interactions in HRI. Since pets have the ability to trigger positive human social experiences [14], [15] also includes the dimension of human-robot attachment (HRA), using companion dogs as a model to shape long-lasting and personalised human-robot partnerships in social contexts.

Ethorobotics is counterposed to the mainstream attitude of shaping robot's social interactions on human-like attributes and psychology. In [13], the authors challenge the assumption that alike forms correspond to social functionality and focus on the "believability" of social interactions instead. As robots often lack analogous human body parts and reaction functions necessary to respond to inputs in an natural way, the authors focus on a complex design of functions, embodiment-behaviour correspondences (e.g. robots' fingers should do what humans' finger do) and technological capabilities (e.g. input perception, motor capacity), that better reflect the functions robots are intended for [3].

These pieces of research sustain that pet companionship can be a functional archetype to design companion robots. They show that human-animal behavioural models (at a more theoretical level) and dog training (on a more applicable extent) can be used in HRI learning contexts to shape believable interactions, congruent user's expectations, and positive teaching experiences.

2.3. Dog training for human-robot teaching: rationale and challenges ahead

Contrary to robots, dogs are sentient beings. They are behaviourally independent from humans, possessing their own individual and species-specific needs, motivations, goals and communication styles. Hence, both dogs and humans need to adapt to each others' way of communication, and somehow formalise their *heterospecific* interactions (i.e. between individuals of different species). For example, formal training enables dogs to learn human-shaped tasks through a combination of cues and structured protocols, though a person would not normally use them with other human individuals.

Arguably, this two-way communicative adaptation does not entirely apply to human-robot interactions (even with state-of-the-art robots). Indeed, artificial agents can be programmed to speak and understand (to some extent) human natural languages, and they can be designed to share the same goals of their users (e.g. in medical applications [16]). Yet, by co-evolving with dogs, humans have developed a familiar form of communication made of "motherese" or "doggerel" speech, simplified gestures, exaggerated facial expressions and postures [17,18,19]. For example, people tend to use high pitch voice tone to catch the dogs' attentions [18], or simple cues such as the pointing gesture to help them finding something [19]. These spontaneous attitudes can form the basis of formal training, where the systematic use of rewards and corrections is employed.

We support the vision that this unique communication process can be successfully transferred to HRI purposes of teachability, enabling non-expert users to interact in an intuitive way with robots. Since the robot's reactions to humans' cues can be shaped through algorithms – and therefore robots are behaviourally dependent from humans – we can focus on what humans do spontaneously when in front of a dog, regardless whether the human interaction is meaningful for real dogs. In other words, we are interested in exploring how humans communicate their intent to dogs through both spontaneous and formal forms of verbal and non-verbal language. Broadly, our goal is to use a human-dog social model to characterise both spontaneous and formal human interactions, to be further exploited in HRI contexts. These forms of communication can then be translated into machine learning models that the robot can use to select relevant actions.

3. A Framework for Enhanced Teaching Capabilities of Humans with Robots

An interdisciplinary approach was used to build our interaction design framework. The process included (i) collaborative discussions among domain experts, to combine notions applicable to dog-inspired robot teaching scenarios; (ii) review of existing theoretical and methodological work in the area of ethorobotics and HRI learning; (iii) the analysis of outcomes of a design sprint session involving a professional dog trainer and the design of a set of tasks using the Choregraphe visual programming tool [20].

3.1. FETCH-R specification

Figure 1 illustrates the FETCH-R framework we designed. This is composed of two core containers describing the teaching interactions, namely the *teacher sphere* and the *learner sphere*. Each sphere includes three types of blocks: decisions/responses from the teacher/learner (hexagons), the different alternatives or choices associated with each decision or response (rectangles), and the sensory capabilities enabling the establishment of an interaction (oblique parallelograms). FETCH-R is meant to be expanded and evolved according to the scenarios and use-cases in which it is applied.

3.1.1. Teacher sphere

The teacher sphere includes decisions, actions and capabilities of the human instructor.

Teachable Tasks are the actual actions that the user would like to teach to the agent. For example, stop walking when a verbal command is given. This is the first level of decision from which other choices can be made.

Teaching Models are related to the mental schemes that the user wants to employ to teach the task. *Inductivity* means having discrete/elementary behaviours, actions or movements in mind and combining them into a sequence to create a serendipitous task. *Deductivity* is about knowing the exact task and subdividing it into discrete/elementary behaviours, actions or movements to create a specific sequence.

Interactive Inputs refer to the type and form of communication the teacher can use according to the sensory capabilities of the learner. For example, communication needs to be established through *attention requests*; elementary actions and behaviours can be taught by giving *teaching signals*; correct or wrong learner's executions can be "marked" through *reinforcement* cues. All these inputs can be delivered through vocal/verbal commands, body postures, facial expressions, touching sensors, etc.
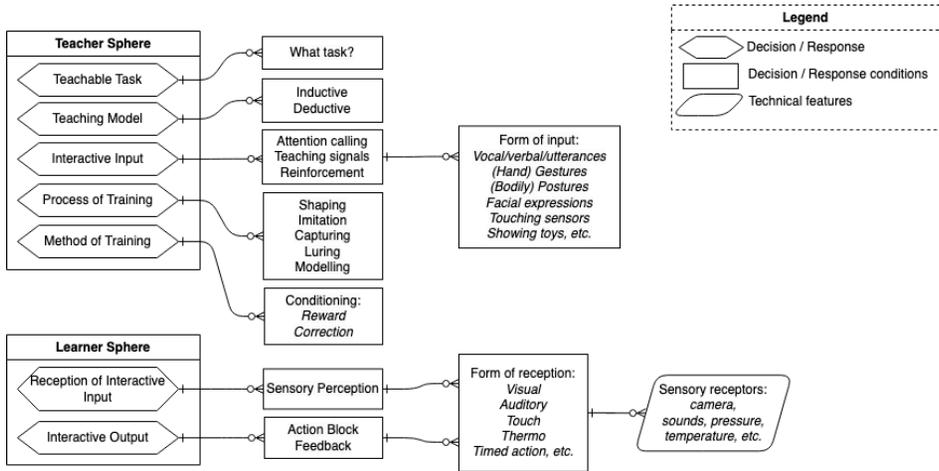
Figure 1. FETCH-R with its elements. In each sphere, hexagons symbolise human's decisions or robot's responses. Rectangles contain lists of decisions or responses. They are affixed to the hexagons through multi-end connectors, showing that they represent a range of possibilities. Oblique boxes include technical items.

Processes of Training map the ways agents can be taught tasks. These are luring, modelling, capturing, imitating, and shaping. In *luring*, the trainer catches the learner's attention by putting a desired item (e.g. a treat) in front of its sensors (e.g. dog's snout), then uses it to guide the agent into the desired position. In *modelling*, the trainer physically manipulates the body of the learner (i.e. kinesthetic manipulation). In *capturing*, the trainer catches a desired learner's behaviour and reinforces it. This can be applied to agents that already know a behaviour, e.g. if a robot randomly sits, the user can reinforce the command *sit*. However, this technique needs spontaneous behaviour to occur by chance, which implies potentially long waiting sessions. In *imitation*, the trainer takes an action and rewards the learner when the action gets imitated. The agent may learn from either copying the form of the action (proper imitation) or directly from the outcome of the action (emulation). A drawback here are the morphological differences between teacher and learner. If the agent learns by copying the trainer's movements, it could be difficult for it to map different body parts from its own. Finally, in *shaping*, the trainer reinforces the learner's small movements (or sounds) close or functional to the desired task. These partial motions will eventually amplify and build the final desired behaviour.

Methods of Training concern the type of conditioning applicable to increase the agent's ability to learn a task. *Rewards* and *corrections* can be given by the teacher through different sensory forms, e.g. utterances, touching sensors, or showing treats.

3.1.2. Learner sphere

The learner sphere concerns the decisions, actions and capabilities of the learner.

The **reception of Interactive Inputs** is linked to the *sensory perception* (e.g. visual, auditory, touch) the learner possesses. This is dependent on the sensory receptors enabling the acquisition of information and making it available for processing (the dog brain, the robot algorithm). Depending on the sensory system of the learner, information can be acquired and processed as streams (videos) and images (including thermo vision) by means of cameras, sounds, haptic inputs (pressure, temperature, motor perception).

Hence, the human's interactive inputs (in the teacher sphere) should be aligned with the receptive ability of the learner.

Interactive Outputs are what the learner communicates back to the teacher. These can be the discrete or complete behaviours, actions and movements taught (i.e. *action blocks*), or feedback showing that the interaction has been established, such as lights/sounds indicating that the learner is in receptive mode. The teacher then knows that the learner is ready to receive teaching signals or reinforcement cues. *Feedback modalities* should be fitting humans' sensory predispositions such as sight and hearing. Hence, they might be visual (e.g. progress bars, LEDs), auditory (e.g. sounds such as beeping) or even timed actions (e.g. when the learner performs a task immediately).

3.2. *Application of FETCH-R: An example of intuitive interaction flow*

We present a use-case describing the full teacher-learner workflow using FETCH-R, illustrating how the framework can be applied on with embodied robots (Figure 2).

Let us imagine that the user of an autonomous vacuum cleaner wants to teach the robot to stop cleaning a precise spot on the floor when a command is given, without operating any button or peripheral device (Figure 3). The *teachable* task to be taught is to stop on a personalised control user command. Given the teaching goal the user has in mind, their mental scheme is deducted from the task itself (a *teaching model*). Here, this is composed by: 1) call the robot to seek attention, 2) use a verbal cue in the form of words to give the command, 3) use visual signal in the form of a gesture to deliver a reward (or correction). The vacuum cleaner robot is equipped with sensors allowing to perceive the surroundings. This sensory capability enables and limits at the same time the acquisition of the personal teaching model of the user. Hence, it is the role of the human to choose *interactive inputs* that are legible by the robot. The whole teaching interaction then results in a two-way communication between the teacher and the learner, including the following back-and-forth signalling sub-interactions: (i) the teacher seeks attention and waits for the robot to signal the achieved attention; (ii) if successful, the teacher can give a command by selecting a *process of training* such as shaping, and wait for the performance of the task; (iii) depending on this, the human chooses a *method of training* to teach whether the action was correctly performed or not. For each sub-interaction, the steps occurring within the teacher and learner communication include the selection of an *interactive input* from the human side (e.g. attention calling for the first sub-interaction) and the *reception of the interactive input* from the robot side. The robot will need to signal back each time whether the input was received and processed, in order for the human to pass to the further steps of the interaction. Depending on whether the learner received the input and conveys this back to the teacher, the interaction can be either interrupted or continued. If interrupted, the teacher will repeat. If continued, the teacher will pass to the subsequent sub-interaction. In the case the robot does not have the receptor for a signal, the teacher needs to change the form of input. Each sub-interaction either restarts or ends when no feedback is given to the user. If the method of training is received, the learning lesson is recorded and an action block is created.

4. Discussion

In the domain of teachable robotics, research in interactive machine learning and ethorobotics aim at building signals (e.g. demonstrations, advice, feedback) easy to use
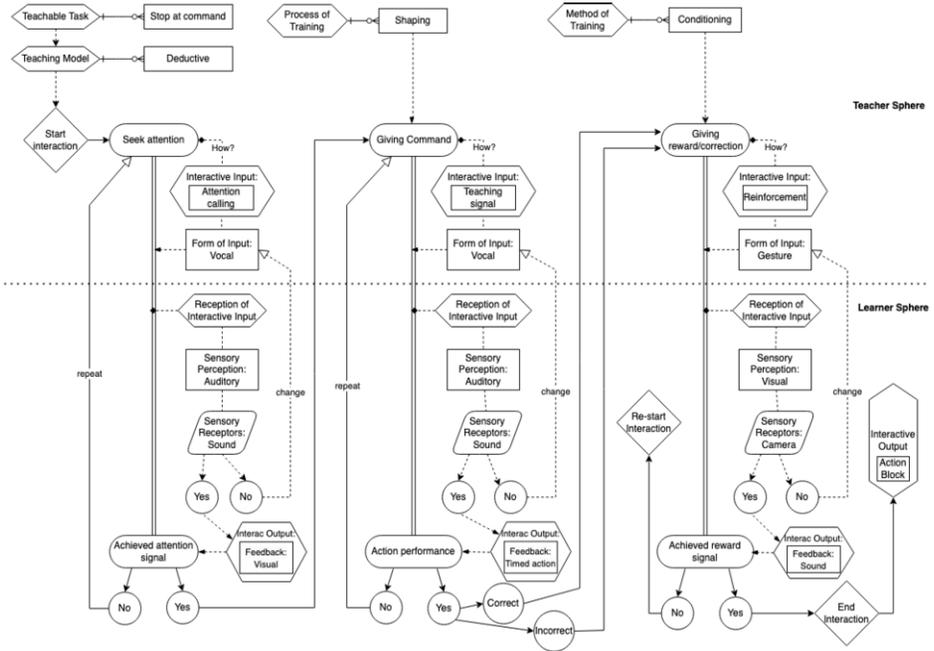
Figure 2. Application of FETCH-R on a use-case (ROOMBA). Hexagons, rectangles and oblique boxes correspond to the framework elements. Multi-end connectors link the hexagons to the rectangles to represent a choice. Rhombuses display the start and end of an interaction. Ovals are the "visible" edges of a two-way communication, illustrated by a double continuous line. Within two ovals, a sequence of "hidden" (mental) choices is represented, where the teacher selects interactive inputs and their forms while the learner receive the input through its sensory system. This hidden layer is depicted through a dotted line. Circles show boolean values.
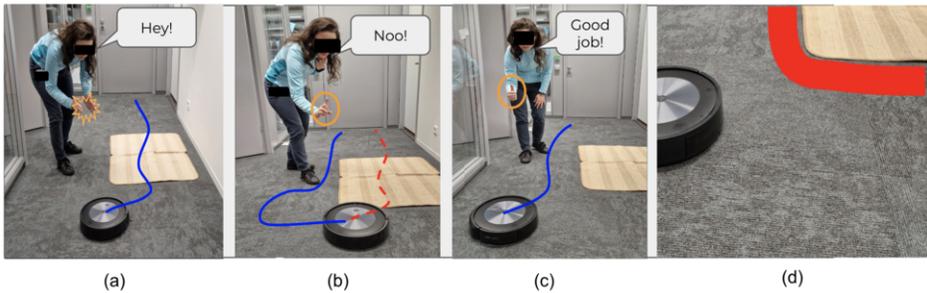

Figure 3. The user teaches a ROOMBA to *avoid* the carpet inspired by the *shaping* method. (a) The user captures the robot's attention using a wakeword and/or sound. The blue line indicates the robot trajectory. (b) The user provides a multimodal negative feedback and the robot acknowledges it, by replanning its trajectory (red dotted line) into a new one (blue line). (c) The user reinforces the adjusted trajectory (blue line) through positive feedback. (d) The robot has learnt to avoid the carpet.

for non-expert users, along with believable and bondlike human-robot interactions. However, these aims would be more usefully approached and achieved by integrating views and methods coming from both domains. While existing control architectures and algorithms are already created and combined to transfer knowledge from humans to robots using the familiar interactivity between humans and dogs [10], a high-level tool to help

shaping such interactivity was still missing. Hence, we provided an interaction design framework that put together dog training concepts and formalisms able to design real-world human-robot teaching interactions. Our framework facilitates the planning of intuitive teaching interactions, potentially applicable to a wide variety of artificial agents, independently of their design, function, or the context in which they are used.

With FETCH-R, we expect a facilitation of the design of robots capable of learning and fine-tuning their behaviour based on natural, spontaneously generated user's input. This in turn would boost the level of engagement and satisfaction with the interaction and the teaching process. This might also enable to ground the planning of machine learning algorithms on natural situations, fulfilling the HRI goal of real-world deployment. The vignette visually describe how FETCH-R can be used to conceive natural and familiar interaction flows inspired by dog-training.

Thinking of robots as if they were dogs might help to bridge the differences in mental models between humans and robots. In other words, the human mind could simplify a different cognitive system than our own more easily. Therefore, it might be more suitable to find a dog-robot correspondence than a human-robot one. However, one limitation of this view is that dog training techniques might result as less natural and intuitive for people not used to interact with dogs. Yet, the principles are easy to understand [10] and we can rely on the evolved, inherent predisposition of the human species towards pets. Indeed FETCH-R, which is today inspired by dog training, could be thought of as inspired to any other animal species with which a person may feel more affiliated with.

Finally, our framework is not intended to be a fit-all solution, but only one way to move forward. Notice that we do not argue that robots should always be treated as dogs, nor that they have strong similarities with them in how they function. We rather stand to the view that from an interaction perspective, our framework has the potential to guide the design of richer human-robot interactions in this context.

5. Conclusions and future work

Transferring human knowledge to robots through natural communication so that the artificial agent can learn interactively from humans has become a prominent goal in HRI. Existing methods still miss to deliver the degree of intuitiveness required by non-programmer end-users, and to close the differences between the teacher and learner's mental models and actions spaces. In this work, we propose to use the hetero-specific communication between humans and dogs to improve users' teaching abilities and experiences towards robots. We presented a design framework called FETCH-R, aimed at guiding the construction of human-robot interactions modelled according to the natural communication existing between pet owners and their animals. Through a scenario, we illustrated how the framework allows setting up demonstrations from human inputs, making them easy to perform by non-experts and learn from by robots, and ultimately promoting users' engagement. Future steps will include the evaluation of the usefulness of the tool through user experiments, and the creation of a database of prototypical interactions that can be implemented using visual programming tools (e.g. Choregraphe).

Acknowledgments

References

[1]   Mirsky R, Baraka K, Faulkner TK, Hart J, Yedidsion H, Xiao X. Human-Interactive Robot Learning (HIRL). In: 2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE; 2022. p. 1278-80.

[2]   Ravichandar H, Polydoros AS, Chernova S, Billard A. Recent Advances in Robot Learning from Demonstration. Annual Review of Control, Robotics, and Autonomous Systems. 2020;3(1):297-330. _eprint: https://doi.org/10.1146/annurev-control-100819-063206. Available from: https://doi.org/10.1146/annurev-control-100819-063206.

[3]   Miklósi A, Gácsi M. On the Utilization of Social Animals as a Model for Social Robotics. Frontiers in Psychology. 2012;3. Available from: http://journal.frontiersin.org/article/10.3389/fpsyg.2012.00075/abstract.

[4]   Sutton RS, Barto AG. Reinforcement learning: An introduction. MIT press; 2018.

[5]   Torabi F, Warnell G, Stone P. Recent Advances in Imitation Learning from Observation. In: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence. Macao, China: International Joint Conferences on Artificial Intelligence Organization; 2019. p. 6325-31. Available from: https://www.ijcai.org/proceedings/2019/882.

[6]   Arzate Cruz C, Igarashi T. A Survey on Interactive Reinforcement Learning: Design Principles and Open Challenges. In: Proceedings of the 2020 ACM Designing Interactive Systems Conference. DIS '20. New York, NY, USA: Association for Computing Machinery; 2020. p. 1195-209. Available from: https://doi.org/10.1145/3357236.3395525.

[7]   Perez Dattari RJ, Celemin C, Ruiz-del Solar J, Kober J. Interactive Learning with Corrective Feedback for Policies Based on Deep Neural Networks. Proceedings of the 2018 International Symposium on Experimental Robotics. 2020;11. Publisher: Springer. Available from: https://repository.tudelft.nl/islandora/object/uuid%3A281fb83a-4f11-4b57-a8af-8f930ced28ea.

[8]   Najar A, Chetouani M. Reinforcement Learning With Human Advice: A Survey. Frontiers in Robotics and AI. 2021;8. Available from: https://www.frontiersin.org/articles/10.3389/frobt.2021.584075.

[9]   Soviany P, Ionescu RT, Rota P, Sebe N. Curriculum Learning: A Survey. International Journal of Computer Vision. 2022 Jun;130(6):1526-65. Available from: https://doi.org/10.1007/s11263-022-01611-x.

[10] Kaplan F, Oudeyer PY, Kubinyi E, Miklósi A. Robotic clicker training. Robotics and Autonomous Systems. 2002 Mar;38(3-4):197-206. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0921889002001689.

[11] Knox WB, Stone P. Interactively shaping agents via human reinforcement: the TAMER framework. In: Proceedings of the fifth international conference on Knowledge capture. K-CAP '09. New York, NY, USA: Association for Computing Machinery; 2009. p. 9-16. Available from: https://doi.org/10.1145/1597735.1597738.

[12] Mohseni-Kabir A, Rich C, Chernova S, Sidner CL, Miller D. Interactive Hierarchical Task Learning from a Single Demonstration. In: Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction. Portland Oregon USA: ACM; 2015. p. 205-12. Available from: https://dl.acm.org/doi/10.1145/2696454.2696474.

[13] Miklósi A, Korondi P, Matellán V, Gácsi M. Ethorobotics: A New Approach to Human-Robot Relationship. Frontiers in Psychology. 2017;8. Available from: https://www.frontiersin.org/articles/10.3389/fpsyg.2017.00958.

[14] McGlynn SA, Kemple S, Mitzner TL, King CHA, Rogers WA. Understanding the Potential of PARO for Healthy Older Adults. International Journal of Human-Computer Studies. 2017 Apr;100:33-47.

[15] Krueger F, Mitchell KC, Deshpande G, Katz JS. Human–dog relationships as a working framework for exploring human–robot attachment: a multidisciplinary review. Animal Cognition. 2021;24:371-85.

[16] van Stijn JJ, Neerincx MA, ten Teije A, Vethman S. Team Design Patterns for Moral Decisions in Hybrid Intelligent Systems: A Case Study of Bias Mitigation. In: AAAI Spring Symposium: Combining Machine Learning with Knowledge Engineering; 2021. .

[17] Paxton DW. Dogma and catma: Coevolution of people, dogs and cats. The Humanistic Psychologist. 2021;49:577-88. Place: US Publisher: Educational Publishing Foundation.

[18] Dotson MJ, Hyatt EM. Understanding dog–human companionship. Journal of Business Research. 2008.

[19] Kaminski J, Nitzschner M. Do dogs get the point? A review of dog–human communication ability. Learning and Motivation. 2013;44:294-302. Place: Netherlands Publisher: Elsevier Science.

[20] Pot E, Monceaux J, Gelin R, Maisonnier B. Choregraphe: a graphical tool for humanoid robot programming. In: RO-MAN 2009 - The 18th IEEE International Symposium on Robot and Human Interactive Communication; 2009. p. 46-51. ISSN: 1944-9437.