Is Machine Learning Enough to Train Robotic Pets?*

Minqiu Zhou[1], Isobel Voysey[1,2] and J. Michael Herrmann[1,2]

***Abstract*— We discuss the problem of learning in robotic pets asking whether the core machine learning paradigm, namely the optimisation of a bounded error function, is sufficient in this context. In pet robots, it seems that the learning process itself rather than the result of this process is the main criterion for the quality of the interaction. Potential extensions of the optimisation paradigm include emotional, self-organising, and exploratory mechanisms to support desirable learning capabilities of a robotic pet. We also propose a co-design process that develops a personalised interaction experience and mutual learning with active contributions from both robotic pets and their owners.**

I. ROBOTIC PETS

Animal companions provide various benefits to their human owners, from longevity and the prevention of coronary heart disease to improved bonding within the family [1]. Pet ownership can be particularly beneficial for children, as the emotional attachments they form with pets can have positive impacts on their socio-emotional development [2]. Owning and caring for a pet can also help to develop children's understanding of biological concepts, such as inheritance [3].

However, there are many reasons that people do not keep pets, including cost, the responsibility involved, unsuitable housing, and allergies [4], [5]. Robotic pets appear a promising alternative but, despite the early success of Tamagotchi toys [6], most pet-like robots do not reach the level of a companion and are thereby limited to educational support for children [7] or short-term interventions for older adults in care settings [8], [9]. This is despite the potential that older adults living independently see for a robotic pet to enhance their social relationships [10], and it disregards the levels of emotional attachment children show after brief interactions with the robotic dog AIBO [11]. Robotic pets could also be used to teach children (and adults) about various aspects of pet care and appropriate behaviour towards animals. This education prior to the ownership of a live pet, whether in the home or through animal welfare education programmes, would contribute to a reduction of animal suffering [12].

One of the most basic requirements for a robotic or live animal pet is physical companionship. Companion animals are animals kept for the purpose of companionship and comfort [13]. They include a broad range of species, from cats and dogs to reptiles, birds, and fish [14]. Attachments can be formed even with species that have limited capacity for social interaction [15], and a study comparing users' feelings of companionship between living, robotic, and virtual pets found that living and robotic pets provided similar levels of companionship, while virtual pets provided markedly less companionship [16], suggesting the mere physical existence of a pet can provide a framework onto which people may project social connections.

For a strong sense of companionship to form with non-human entities, requirements may include an appealing or "cute" appearance [17], a high degree of animacy [18], responsiveness [19], and emotional expression [20], [21], [22]. On a higher behavioural level, it may also be desirable for a robotic pet to provide some level of emotional support, through attunement to the owner's emotional state [23]. Additionally, the owner may want the pet to display reciprocal attachment, which can be realised by the pet's behaviour signifying interest and care, such as greeting the owner when they return home [23]. However, each individual will have different motivations and circumstances behind acquiring a pet [24], so the detailed specification of the robotic pet will be highly personal, and individual preferences and level of understanding about animals will need to be taken into account in the design. However, this does not refer a fixed specification of an objective for a pet's adaptation, as these preferences are developing with the interaction or pet and human, and will not be observable outside the interaction that is to be designed to begin with.

One of the early studies on artificial pets proposes the *uselessness principle* for their design [25]. It suggests that a robotic pet's primary function is not to provide any service to its owner, thereby necessitating a radically different approach to design compared to the majority of robots. The robotic pet should be driven largely by its own goals and may ignore orders that do not align with its goals. The author argues that this autonomy is a necessary, although not sufficient, feature for the development of an interesting and engaging relationship [25]. In the years since, this idea has been refined in the field of autonomous learning, and in particular in self-motivated [26] and reflexive reinforcement learning [27], skill discovery [28], and others, which we study in the present project.

Recent developments in machine learning enable advanced sensing, planning and action in robotics which we start to appreciate in service robots, but we will argue that an improvement of the information processing capabilities together with increased performance and appearance will not be sufficient to improve the acceptance of robots as pets.

We do not consider the implications of natural language processing here which may be perceived as unnatural in

*This work received partial support from UKRI via CDT RAS Edinburgh.

[1]Minqiu Zhou, Isobel Voysey, and J. Michael Herrmann are with the School of Informatics, University of Edinburgh, Edinburgh, EH8 9BT, UK, m.zhou-39@sms.ed.ac.uk, i.a.voysey@sms.ed.ac.uk, michael.herrmann@ed.ac.uk

[2]Isobel Voysey and J. Michael Herrmann are with the Edinburgh Centre for Robotics, Edinburgh, UK

animal-like robotic pets and has been shown [29] to evoke ambivalent responses in elderly participants, while non-talking robots have a good chance to be perceived as "beautiful" or "compassionate". We also do not focus on the question of quantifying the performance of robotic pets. There exist subjective measures for people's feeling of companionship [30] or attachment to pets [31], but in future it will be important to select or develop objective measures to evaluate the performance of machine learning techniques, for example, time spent with the pet.

II. Machine learning for robotic pets

Machine learning methods are yielding increasingly reliable results in tasks with a clearly specified goal, such as path planning or face recognition. This is achieved by the minimisation of a loss function over a data set, or, in the case of reinforcement learning, the maximisation of a reward average. Thus, the optimisation of an objective functions can be seen as the main feature of machine learning algorithms. It has even been claimed that it is possible to design algorithms that realise in this way any form of intelligence [32], if large data sets, complex computational architectures, and sufficiently long learning times are realisable. Yet, we argue that it may not be enough to control a modest pet-like robot, where the eventual result is unimportant.

Although efficient algorithms can achieve few-shot learning or use transfer learning to generalise learned behaviour to new domains, and may produce results comparable to the expectations one might have in an animal pet, it seems that learning success is less important than familiarity, reliability, and some level of creativity. The owner tends to value the time spent with the pet, which is at least partially due to the process of learning itself being more joyful and more important than the final performance of the system.

Error minimisation is, nevertheless, an important component in some of the vital functions of a pet robot as a product, for example, in order to provide basic behaviours (similar to natural traits in a pet animal) and to implement safety regulations. The application of reinforcement learning on a high level, for example to improve user satisfaction, increase engagement duration, and uninterrupted function, is non-trivial, because the exploration of the vast space of potentially useful behaviours requires strategies that need to be designed as well. Therefore, there is a need for active learning mechanisms that can enable a pet robot to find intrinsic motivation to guide reinforcement learning, which can improve smoothness and predictability of movements and supports versatility by skill discovery.

It should also be noted that the unsuccessful execution of a new skill by a pet robot can be appreciated by the owner as it can appear as an intention to learn, to cooperate in this learning process that may be considered as rewarding also to the participating human. This support of the robot's progress by the human can be considered as a goal of the robot's learning, such that the cooperative active learning process is characterised by mutual guidance and support. Likewise, exploratory behaviour will be seen as an attempt to gain information, which can be expected to be fancied in particular if the pet shows a tendency to explore the repertoire of the human. However, to keep user interaction at an acceptable level, it may be suitable to include also emotional dynamics. This has the additional benefit that it can help to supervise the system that is driven by various sources of information in various subsystems. So, a supersystem that monitors the learning progress and the state of the interaction would be useful. This emotional system complements the body of work that has been spent on the recognition and emulation of emotions by the robot. It also provides the drives that control flexibility and adaptability in the learning system that in this way realises a form of self-organisation which, however, deserves further study. It will also support a personalised experience by including sensitivity to a suitable amount of co-operative engagement which needs to be included as a design feature in the pet robot's control strategy.

III. Learning in pets as a co-design process

The development of abilities in a robotic pet should reflect the preferences of the human, but it is not a task to be imposed on them. Instead, the opportunity to continuously adapt the behavioural design of the robot should be given to the human. This and the complementary task of the robot to engage with human, leads to the idea of a co-design process. Co-design is known as a design process where all members contribute in the design process as equal collaborators in a way that fuses optimally the expertise of specialists with the problem-awareness of the users, care-givers, technicians, and other relevant stakeholders. Co-design is common practice in many fields, see e.g. [33], but is considered here a mode of interaction to be realised not between groups of people, but as a dominant mode of human-robot interaction between a pet robot and its owner. This co-design process can be seen as a maturation process that includes various learning processes rather than being simply a learning process itself. Its realisation would include the following points.

Robotic pets are sold with pre-trained sensing capabilities and a few basic behaviours, in other words, the robotic pets arrive in an immature state. After a period of interaction with the environment and the human, the robotic pet gradually grows and develops a unique behavioural organisation, based on the learning experiences they share with the human owners. Research in robotic pets will thus focus on design of co-design which will include the following features.

a) Error tolerance: In contrast to general machine learning that suppresses errors, the maturation process embraces errors which are expected in two ways: In the exploration and self-motivation scenario in reinforcement learning, errors are expected and drive robotic pets. In regard to the use of robotic pets as an entertainment companion, they are allowed to produce slips now and then. Thus, referring to the 'uselessness' or 'cuteness', some errors, for example in locating owners, could make them appear more alive.

Prior research has shown that forgetfulness might enable a more natural and believable attachment bond between human and robot companion [34] as opposed to choosing selected "error" behaviours to incorporate, taking thus more holistic view of error tolerance.

b) Personalisation: Individual difference are encouraged in a relationship. With the underlying subjective view, personalisation enriches the human-robot attachment. During the process of maturation, the human-robot interaction becomes more and more specific. It represents a challenge for the behavioural organisation in the robot. In addition to the activation of behaviours by trigger stimuli, any autonomous behaviours need to be grouped according to sequentiality, intended state changes, and context. Although this meta-organisation of behaviour can in principle be learned as well, it may be advisable to provide the robot control architecture with an expressive structure that will become partially inhabited during the maturation process.

c) Active learning: During the maturing period in an individual setting, each robotic pet experiences different interactions and uses any learning successes in the search for new learning data. Thus, in contrast to the assumptions in machine learning, the data sets do not only vary case by case, but are also essentially nonstationary. This leads to complexity and difficulty in the implementation of capturing and defining states or events. Moreover, although robotic pets have access to an enormous number of data during their life, data sets for a specific task are comparably small, which further increases the learning difficulty. However, as errors are not to be avoided, this technical limitation can be experienced as part of the character of the pet.

d) Mutuality: It is not only the robotic pet learning from its owner. Also, owners learn from their robotic pets. This happens while the owners spend time and effort to understand, interpret and control the robotic behaviours. In this way a mutual interaction is formed which is more engaging for the human owners than the mere operation of a machine. Learning in robots and humans works differently in many respects, and it is critical for the robot to be able to access the superior capabilities of humans to adapt. This is possible as demonstrated by the success of computer games.

The implementation of these principles is obviously not straightforward, but can include advances in various fields of machine learning incorporating active learning, imitation learning, transfer learning, reflexive reinforcement learning, and other methods, however, with standards implied by the enjoyable interaction during the learning process, the acceptance of characteristic errors and inconsistencies on both sides, and the development of mutually agreeable repetitive behaviours that would appear as joint rituals.

An evaluation of the quality of a robotic pet as characterised here, would primarily be based on the statistics of the rating of the human user experience. For a more objective account, the duration of daily interactions, and the development of interactions over time can be measured. Beyond this, the complexity of the robotic behaviours can be monitored in pilot cases and information-theoretically analysed. An increase in behavioural complexity with a simultaneous increase of predictability can be seen as indicative of a rich and reliable companionship. It is expressed by the concept of *predictive information* [35] which, in addition to other applications, has been proposed to enable autonomous learning.

IV. Discussion

The project that is described here is still in the making, but it is important to reconsider the principles of the design of a pet robot and to contrast it to the design goals of control architectures of other types of robots. In this way, we can create a niche for a new species of robots that benefits from a symbiosis with its human companions just like the human owners benefit from the robotic companions. Although full functionality of the envisioned pet robot seems a long way ahead, it is still necessary to consider the risks and limitations of this research.

Robotic pets, in particular in applications with elderly persons [8], [9], have been promoted for the purpose of monitoring health, learning progress, and safety, although there are reports of negative side effects of the purposeful sneaky usage of the robot in place of a pet [10]. Similarly, one may object that a fully-functional robotic pet might create a dependency in the interacting human which would not be excusable, *unless* there is a necessary purpose for the presence of the pet in the private domain of the human.

If the robotic pet is used for preparing a (young) person to the responsibilities of owning an animal pet, then any risks of the interaction are contained and will typically be monitored towards the decision of acquiring an animal pet. Likewise, as a temporary companion such as under the conditions of a lockdown or hospital stay the benefits of a robotic pet would outweigh the risks in most cases.

Limitations of the proposed approach consist in the conceptual problem that the robot has no intrinsic need to interact with humans, i.e. the interactive behaviour needs to be explicitly rewarded within the robot's behavioural module. It may be possible for the pet robot to find out that the interaction does support its intrinsic motivation towards behavioural learning, but then the problem is shifted towards a will to learn. Also, at some point learning progress will saturate due to the limited capabilities of the robot, so that this motivation will diminish.

A related point is the limited function of current hardware. In order for a robot to show impressive behaviour, appropriate hardware needs to be designed, which is beyond the control- and learning-related approach taken here. Needless to mention that here lies much potential for trainable hardware, analogous to muscles, bones, and energy consumption that are known to respond to training in animals.

Although beyond the scope of the present paper, we also need to consider the use of bio-degradable materials, sustainable power sources, as well as minimal-impact behaviours in natural environments.

More generally, the research circumscribed here may be a important component in autonomous robots although probably in combination with trustworthy mechanisms that guarantee a minimal function and limit any risks. It may as well be useful in setting a scope for the modelling of animal behaviour.

V. Conclusion

We discussed the appropriateness of the machine learning techniques for different aspects of robotic pets and concluded that although the machine learning paradigm performs well at sense-and-act levels, it may be too narrow at intentional and cognitive levels. We arrived at the proposal of a co-design learning process for robotic pets, resembling the process of maturation, that cannot simply be learned from rewards, but needs to be accommodated by the design of interaction modes that is an essential part of the design of robotic pets.

References

[1] A. M. Beck and A. H. Katcher, *Between pets and people: The importance of animal companionship*. Purdue University Press, 1996.

[2] G. F. Melson, *Why the wild things are: Animals in the lives of children*. Harvard University Press, 2005.

[3] J. M. Williams and L. A. Smith, "Social and experiential influences on the development of inheritance concepts," *International Journal of Behavioral Development*, vol. 30, no. 2, pp. 148–157, 2006.

[4] K. A. Anderson, L. K. Lord, L. N. Hill, and S. McCune, "Fostering the human-animal bond for older adults: Challenges and opportunities," *Activities, Adaptation & Aging*, vol. 39, no. 1, pp. 32–42, 2015.

[5] M. C. Gates, J. Walker, S. Zito, and A. Dale, "Cross-sectional survey of pet ownership, veterinary service utilisation, and pet-related expenditures in New Zealand," *New Zealand Veterinary Journal*, vol. 67, no. 6, pp. 306–314, 2019.

[6] L. Lawton, "Taken by the Tamagotchi: How a toy changed the perspective on mobile technology," *The iJournal: Student Journal of the Faculty of Information*, vol. 2, no. 2, 2017.

[7] J. D. Decuir, T. Kozuki, V. Matsuda, and J. Piazza, "A friendly face in robotics: Sony's aibo entertainment robot as an educational tool," *Computers in Entertainment (CIE)*, vol. 2, no. 2, pp. 14–14, 2004.

[8] L. Hung, C. Liu, E. Woldum, A. Au-Yeung, A. Berndt, C. Wallsworth, N. Horne, M. Gregorio, J. Mann, and H. Chaudhury, "The benefits of and barriers to using a social robot PARO in care settings: A scoping review," *BMC Geriatrics*, vol. 19, no. 1, pp. 1–10, 2019.

[9] S. Petersen, S. Houston, H. Qin, C. Tague, and J. Studley, "The utilization of robotic pets in dementia care," *Journal of Alzheimer's Disease*, vol. 55, no. 2, pp. 569–574, 2017.

[10] A. Lazar, H. J. Thompson, A. M. Piper, and G. Demiris, "Rethinking the design of robotic pets for older adults," in *Proc. ACM Conf. Designing Interactive Systems*, pp. 1034–1046, 2016.

[11] A. Weiss, D. Wurhofer, and M. Tscheligi, ""I love this dog" – Children's emotional attachment to the robotic dog aibo," *International Journal of Social Robotics*, vol. 1, no. 3, pp. 243–248, 2009.

[12] J. M. Williams, M. P. Cardoso, S. Zumaglini, A. L. Finney, S. SPCA, and M. A. Knoll, ""Rabbit rescuers": A school-based animal welfare education intervention for young children," *Anthrozoös*, vol. 35, no. 1, pp. 55–73, 2022.

[13] A. Steiger, "Pet animals: housing, breeding and welfare," *Animal Welfare*, pp. 111–133, 2006.

[14] Royal Society for the Prevention of Cruelty to Animals, "RSPCA Policy A01 Responsible companion animal ownership." https://kb.rspca.org.au/knowledge-base/rspca-policy-a01-responsible-companion-animal-ownership, 2022. [Online; accessed 11/08/2022].

[15] J. Langfield and C. James, "Fishy tales: experiences of the occupation of keeping fish as pets," *British Journal of Occupational Therapy*, vol. 72, no. 8, pp. 349–356, 2009.

[16] J.-L. Rault, "Pets in the digital age: Live, robot, or virtual?," *Frontiers in Veterinary Science*, p. 11, 2015.

[17] C. Caudwell and C. Lacey, "What do home robots want? The ambivalent power of cuteness in robotic relationships," *Convergence*, vol. 26, no. 4, pp. 956–968, 2020.

[18] L. Voss, *More than machines?: The attribution of (in) animacy to robot technology*. transcript Verlag, 2021.

[19] G. E. Birnbaum, M. Mizrahi, G. Hoffman, H. T. Reis, E. J. Finkel, and O. Sass, "Machines as a source of consolation: Robot responsiveness increases human approach behavior and desire for companionship," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 165–172, IEEE, 2016.

[20] S. M. F. Hosseini, D. Lettinga, E. Vasey, Z. Zheng, M. Jeon, C. H. Park, and A. M. Howard, "Both "look and feel" matter: Essential factors for robotic companionship," in *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 150–155, IEEE, 2017.

[21] N. Martelaro, V. C. Nneji, W. Ju, and P. Hinds, "Tell me more designing hri to encourage more trust, disclosure, and companionship," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 181–188, IEEE, 2016.

[22] I. Olaronke and I. Rhoda, "A systematic review of emotional intelligence in social robots," 2017. Subm. to International Conference of Applied ICT (ICAICT).

[23] K. Riddoch, R. Hawkins, and E. S. Cross, "Exploring behaviours perceived as important for human—dog bonding and their translation to a robotic platform," 2021.

[24] K. E. Holland, "Acquiring a pet dog: A review of factors affecting the decision-making of prospective dog owners," *Animals*, vol. 9, no. 4, p. 124, 2019.

[25] F. Kaplan, "Free creatures: The role of uselessness in the design of artificial pets," in *Proceedings of the 1st Edutainment Robotics Workshop*, pp. 45–47, Sankt Augustin, 2000.

[26] N. Chentanez, A. Barto, and S. Singh, "Intrinsically motivated reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 17, p. 1281–1288, 2004.

[27] B. I. Lyons and J. M. Herrmann, "Reflexive reinforcement learning: Methods for self-referential autonomous learning.," in *IJCCI*, pp. 381–388, 2020.

[28] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine, "Diversity is all you need: Learning skills without a reward function," *arXiv preprint arXiv:1802.06070*, 2018.

[29] S. Coghlan, J. Waycott, A. Lazar, and B. Barbosa Neves, "Dignity, autonomy, and style of company: Dimensions older adults consider for robot companions," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5:CSCW1, no. 104, pp. 1–25, 2021.

[30] D.-B. Luh, E. C. Li, and Y.-J. Kao, "The development of a companionship scale for artificial pets," *Interacting with Computers*, vol. 27, no. 2, pp. 189–201, 2015.

[31] T. P. Johnson, T. F. Garrity, and L. Stallones, "Psychometric evaluation of the Lexington attachment to pets scale (LAPS)," *Anthrozoös*, vol. 5, no. 3, pp. 160–175, 1992.

[32] D. Silver, S. Singh, D. Precup, and R. S. Sutton, "Reward is enough," *Artificial Intelligence*, vol. 299, p. 103535, 2021.

[33] L. Huybrechts, H. Benesch, and J. Geib, "Institutioning: Participatory design, co-design and the public realm," *CoDesign*, vol. 13, no. 3, pp. 148–159, 2017.

[34] M. Biswas and J. Murray, "Robotic companionship: how forgetfulness affects long-term human-robot interaction," in *International Conference on Intelligent Robotics and Applications*, pp. 37–48, Springer, 2015.

[35] W. Bialek, I. Nemenman, and N. Tishby, "Predictability, complexity, and learning," *Neural computation*, vol. 13, no. 11, pp. 2409–2463, 2001.