

Investigating Haptic Co-Creation with Reinforcement Learning

Juliette Regimbal^{1,2}[0000–0003–4902–046X] and
Jeremy R. Cooperstock^{1,2}[0000–0002–3392–2541]

¹ Department of Electrical and Computer Engineering

² Centre for Interdisciplinary Research in Music Media and Technology

McGill University, Montréal, Canada

juliette.regimbal@mail.mcgill.ca, jer@cim.mcgill.ca

Abstract. Co-creative computer agents can be used to support designers they work alongside. However, this approach to co-creativity has not been studied in haptics. In this paper, we investigate this topic by applying an interactive reinforcement learning technique developed for sound synthesis to a two degree-of-freedom force-feedback design space. We developed a prototype authoring tool to study the interaction between designers and this type of co-creative agent. As these agents are trained on the designer’s feedback in real-time, this method is flexible and could be applied to other hardware and software configurations. Five participants were recruited to use the tool in two artistic tasks and were interviewed about their experiences. An analysis of the participants’ comments and actions indicated that co-creation in haptics can help designers explore new effects if they are sufficiently able to control the extent and direction of exploration.

Keywords: Authoring tool · Co-creativity · Machine learning.

1 Introduction

Quickly iterating on and evaluating new ideas is a crucial aspect of the design process. Creativity Support Tools (CSTs) in visual art [3] and sound synthesis [10] have integrated co-creative agents to suggest new concepts to designers and improve their creative experience through collaboration. Such sharing of control between designer and agent has not been integrated into haptic design tools, which instead primarily focus on collaboration between human designers [9,15]. Co-creative agents could be beneficial to haptic designers in artistic contexts, where there is no optimal solution, and especially those frequently working with unfamiliar hardware who cannot become experts in the design space of every system they use.

To investigate how such co-creation may support designers and to guide future research, we implemented a simple haptic authoring tool which allows users to switch between manipulating haptic parameters and responding to modifications made by a reinforcement learning (RL) computer agent. This tool controls

a third generation Haply 2DIY,³ a 2-DoF force feedback device. We recruited five participants to use this tool in two open-ended design activities and analyzed their use of this system. The participants were then interviewed on their experience and how they believe co-creativity may support their haptic design practices. These initial results were promising and indicated future directions to refine our current approach, the implementation of which is available as a Git repository.⁴

2 Related Work

CSTs are broadly intended to support aspects of exploration, collaboration, and iteration as related to the creative process [11], typically focused on supporting idea generation, evaluation, and implementation [4]. Some CSTs include an agent, often through machine learning methods such as deep learning [8] or RL [3,10], which modify a design alongside the designer. This approach poses additional challenges, such as how to manage control shifting between designer and agent, and agents behaving in ways dissimilar from human collaborators [3,8].

Traditional RL treats a problem as a Markov decision process in which an agent acts upon an environment and learns from changes in the state and rewards produced by that environment on each step. The agent then adjusts its actions in order to maximize its total reward, discounting those received in the future by a factor $\gamma \in [0, 1]$ [12]. The TAMER interactive RL framework modifies this structure by replacing the environment reward after each step with intermittent, positive or negative rewards from a human evaluating the agent’s performance. In other words, a TAMER agent learns a policy to maximize this evaluator’s feedback. Since an evaluator’s criteria may change as performance improves, TAMER typically uses a smaller discount rate than traditional RL [6].

Our tool is primarily based on the Co-Explorer [10], a CST for sound synthesis that uses a deep learning version of TAMER [14]. The Co-Explorer allows users to directly modify sound parameters and use a TAMER agent acting in response to a user’s evolving goals as expressed through two types of feedback. Guiding feedback reflects the user’s assessment of the recent trajectory of changes made by the agent. Zone feedback indicates a like or dislike of sounds similar to the current one, not taking into account the direction of recent changes. Users can also have the agent jump to unexplored zones to immediately begin evaluating new sounds [10].

Generative adversarial networks to design vibrations or textures [5,7,13] have been the primary application of machine learning to haptic design. With the exception of Lu et al.’s system, which allows users to iteratively select their preferred texture [7], these projects have not supported user interaction. Our approach differs from these by focusing on the creative experience of designers rather than realism, allowing direct user manipulation of effects, and not requiring datasets or model pre-training prior to use.

³ <https://2diy.haply.co/>

⁴ <https://github.com/JRegimbal/haptic-swatch>

3 System Description

The tool was developed to test computer co-creation in force-feedback haptic design following the approach to user interaction taken in the Co-Explorer [10]. A simple interface (Figure 1) was developed to support the creation, repositioning and manipulation of circular haptic elements in a two-dimensional workspace. Virtual knobs for each element allowed the user to control 1) a spring force originating at the centre of the element, 2) a simple model of friction based on the technique of Culbertson et al. [2], and 3) two vibrational texture components for which the frequency and maximum amplitude of vibration can be adjusted separately. These effects are felt when the user moves the avatar for the haptic device’s end effector into an element.

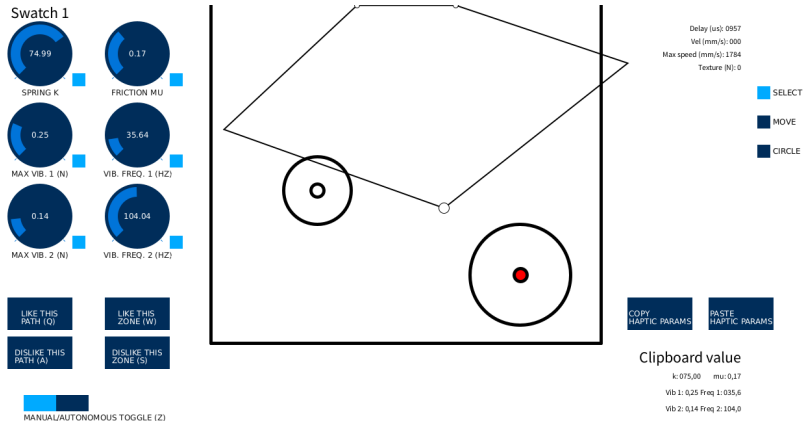


Fig. 1: The tool UI showing the haptic parameter knobs (top left), agent feedback and toggle (bottom left), workspace (centre), and element controls (right).

By default, the designer is able to control these elements without an agent’s input. A toggle permits the user to switch between this mode and an autonomous, agent-driven one. In autonomous mode, the system continuously makes slight adjustments, at a rate of 10 Hz, to the parameters of the element over which the designer has positioned the end effector. Checkboxes next to each knob allow the designer to “lock” the value of that parameter, preventing the agent from modifying it for that element.

A separate TAMER agent is responsible for each of the elements in the workspace since each may have different desired effects. A shallow neural net with 16 hidden nodes was used to estimate how a user would reward state-action pairs. It was updated in response to guiding and zone feedback using stochastic gradient descent (learning rate $\alpha = 0.02$). Agent actions were determined via an ϵ -greedy policy ($\epsilon = 0.1$), in which the highest valued action is selected with probability $1 - \epsilon$ and an action is chosen randomly to promote exploration with

probability ϵ . Future rewards were discounted by $\gamma = 0.5$, each action adjusted one parameter by 5% of its total range, and actions on locked parameters could not be selected. The weight of the reward assigned to past state-action pairs in guiding feedback uses the same gamma distribution as in TAMER [6].

These features were added so that, in autonomous mode, the user would feel the properties of the element change quickly, but not too quickly as to be overwhelming. While initially the agent would essentially select actions at random, it should increasingly make decisions that align with user preference expressed through feedback. As the user finds that parameter values become satisfactory, whether by manual or automatic changes, the locking feature can prevent these from being overwritten by the agent if the reward model has not yet identified these values as being optimal for the particular parameters.

4 Methodology

We recruited individuals with prior experience in haptic design in order to observe how they interacted with our system and to identify areas of this interaction to develop and refine. To begin, each participant completed a demographic questionnaire and was instructed on the use of the 2DIY and of the design tool, walking through a brief scenario with the experimenter.

Following this, the participant received two design prompts inspired by the “design to create” exercise conducted by Scurto et al. [10]. For each prompt, participants received a photograph or abstract art and were instructed to make a haptic representation of its “atmosphere”. These instructions and graphics were selected to encourage creative exploration and avoid unintentionally communicating that a particular sensation should be recreated. Participants were told that their representations would not be assessed for quality, and that the purpose of the exercise was to evaluate the tool and co-creative functionality.

After the design activities were completed, each participant was interviewed on aspects of their experience using the tool, and especially their use of the co-creative agents. Participants were also asked to discuss the role of co-creative agents in the haptic design process. Audio recordings were made of the interviews and the design tasks, during which the participants were instructed to think aloud as they worked. User actions in the tool were logged for analysis.

5 Results

Five participants (ages 21–37, mean 27, all male) who designed haptic effects as part of their undergraduate or graduate studies were recruited for the experiment. Each was compensated CA\$20. The first author performed a thematic analysis on visualizations of log data and transcriptions of the recordings, following the method outlined by Braun and Clarke [1], developing two themes for interactions with the agents: increasing discovery and the importance of control.

Increasing discovery: Participants preferred to use the co-creative agents to help discover interesting effects possible with the tool and device. We identified two primary uses of the tool for this purpose: initial ideation and design refinement. First, given the nature of the task, participants often started without a particular goal in mind, electing to begin working with the co-design tool in autonomous mode. As they needed to evaluate a particular effect or make adjustments, they would then switch to manual mode. Second, in cases when the participant was able to create a rough version of a desired effect, they would sometimes use the tool to refine their design by evaluating different parameter combinations. This was often more efficient than could be done by hand. Participants would frequently alternate between autonomous and manual controls regardless of which was used first. They also used both kinds of controls to explore multi-element effects.

Importance of control: Participant feedback generally indicated a desire for control, i.e., a sense of agency, which was often determinative of whether they enjoyed an interaction with the system. This was the case when they were able to understand an agent’s operations and meaningfully shape its future actions. This aligns with results studying co-creativity in other contexts [8,10], suggesting that broad guidelines for co-creative interfaces are applicable to haptics. When “chaotic” (P1) or otherwise broad effects were desired, seeing the effects of an agent’s actions on the knobs of a selected element provided enough visibility into its state. However, when subtle effects were targeted and could not quickly be reached, additional insights were desired, such as explanations for why certain actions were taken, or how guiding and zone feedback impacted the agent differently. When working with multiple elements, participants struggled with the lack of visibility into the states of the elements and their respective agents.

Future improvements: Participants further desired flexibility in collaborating with the tool rather than strict turn taking. To ensure that the tool supports the variety of strategies along the “user-as-leader/agent-as-leader” turn-taking continuum [10] adopted by our participants, future research into more flexible means of providing feedback to the agents is warranted. Parameter locking was seen as a step towards such flexibility since it provides a more immediate way to direct the agent’s future actions than guiding and zone feedback. Participants also suggested possible improvements to the tool of adjusting the rate at which, or range in which, actions are taken, and having the agent learn from manual adjustments.

Difficulties forming accurate conceptual models of autonomous and, in one instance, of the haptic elements themselves, motivates work to visualize the state of these co-creative agents and the haptic elements being controlled. Future designers may benefit from additional means of interacting with RL agents and greater insight into the changes these interactions produce. The low computational requirements of the agents and on-the-fly training also encourage applying them in other contexts, such as with other haptic modalities than force feedback or in audio-haptic contexts.

6 Conclusion

Our early results studying haptic designer interactions with low-complexity, co-creative agents built on an interactive reinforcement learning framework are encouraging. Considerable development will be necessary before this approach is usable in everyday work, but our work has identified several directions that could lead to this point. It is our hope that further advances in such co-creative tools, enabling continued improvements in the experience of those performing creative haptic work, will translate into a greater adoption of haptic devices in design practices, and more compelling haptic experiences for end-users.

Acknowledgments. The first author is supported by the Natural Sciences and Engineering Research Council of Canada (scholarship no. 569236), the Fonds de recherche du Québec - Nature et technologies (scholarship no. 315050), and a travel award from the Centre for Interdisciplinary Research in Music Media and Technology. Our study was approved by the McGill University Research Ethics Board (file 22-08-090).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Braun, V., Clarke, V.: Using thematic analysis in psychology. *Qualitative Research in Psychology* **3**(2), 77–101 (Jan 2006). <https://doi.org/10.1191/1478088706qp063oa>
2. Culbertson, H., Delgado, J.J.L., Kuchenbecker, K.J.: One hundred data-driven haptic texture models and open-source methods for rendering on 3D objects. In: 2014 IEEE Haptics Symposium. pp. 319–325 (Feb 2014). <https://doi.org/10.1109/HAPTICS.2014.6775475>
3. Davis, N., Hsiao, C.P., Yashraj Singh, K., Li, L., Magerko, B.: Empirically Studying Participatory Sense-Making in Abstract Drawing with a Co-Creative Cognitive Agent. In: 21st International Conference on Intelligent User Interfaces. pp. 196–207 (Mar 2016). <https://doi.org/10.1145/2856767.2856795>
4. Frich, J., MacDonald Vermeulen, L., Remy, C., Biskjaer, M.M., Dalsgaard, P.: Mapping the Landscape of Creativity Support Tools in HCI. In: 2019 CHI Conference on Human Factors in Computing Systems. pp. 1–18 (May 2019). <https://doi.org/10.1145/3290605.3300619>
5. Hernandez-Mejia, C., Ren, X., Thabuis, A., Chavanne, J., Germano, P., Perriard, Y.: Generative Adversarial Networks for Localized Vibrotactile Feedback in Haptic Surfaces. In: 24th International Conference on Electrical Machines and Systems. pp. 105–110 (Oct 2021). <https://doi.org/10.23919/ICEMS52562.2021.9634513>
6. Knox, W.B., Stone, P.: Interactively shaping agents via human reinforcement: the TAMER framework. In: Proceedings of the fifth international conference on Knowledge capture. pp. 9–16 (Sep 2009). <https://doi.org/10.1145/1597735.1597738>
7. Lu, S., Zheng, M., Fontaine, M.C., Nikolaidis, S., Culbertson, H.: Preference-Driven Texture Modeling Through Interactive Generation and Search. *IEEE Transactions on Haptics* **15**(3), 508–520 (Jul 2022). <https://doi.org/10.1109/TOH.2022.3173935>

8. Oh, C., Song, J., Choi, J., Kim, S., Lee, S., Suh, B.: I Lead, You Help but Only with Enough Details: Understanding User Experience of Co-Creation with Artificial Intelligence. In: 2018 CHI Conference on Human Factors in Computing Systems. pp. 1–13 (Apr 2018). <https://doi.org/10.1145/3173574.3174223>
9. Schneider, O., MacLean, K., Swindells, C., Booth, K.: Haptic experience design: What hapticians do and where they need help. *International Journal of Human-Computer Studies* **107**, 5–21 (Nov 2017). <https://doi.org/10.1016/j.ijhcs.2017.04.004>
10. Scurto, H., Van Kerrebroeck, B., Caramiaux, B., Bevilacqua, F.: Designing Deep Reinforcement Learning for Human Parameter Exploration. *ACM Transactions on Computer-Human Interaction* **28**(1), 1–35 (Feb 2021). <https://doi.org/10.1145/3414472>
11. Shneiderman, B.: Creativity support tools: accelerating discovery and innovation. *Communications of the ACM* **50**(12), 20–32 (Dec 2007). <https://doi.org/10.1145/1323688.1323689>
12. Sutton, R.S., Barto, A.G.: Finite Markov Decision Processes. In: Reinforcement learning: an introduction, pp. 79–106. The MIT Press, Cambridge, Massachusetts, second edn. (2018)
13. Ujitoko, Y., Ban, Y.: Vibrotactile Signal Generation from Texture Images or Attributes Using Generative Adversarial Network. In: Haptics: Science, Technology, and Applications. pp. 25–36 (2018). https://doi.org/10.1007/978-3-319-93399-3_3
14. Warnell, G., Waytowich, N., Lawhern, V., Stone, P.: Deep TAMER: Interactive Agent Shaping in High-Dimensional State Spaces. In: AAAI Conference on Artificial Intelligence (Apr 2018). <https://doi.org/10.1609/aaai.v32i1.11485>
15. Wittchen, D., Spiel, K., Fruchard, B., Degraen, D., Schneider, O., Freitag, G., Strohmeier, P.: TactJam: An End-to-End Prototyping Suite for Collaborative Design of On-Body Vibrotactile Feedback. In: 16th International Conference on Tangible, Embedded, and Embodied Interaction. pp. 1–13 (Feb 2022). <https://doi.org/10.1145/3490149.3501307>