Riding the Co-Explorers

Hugo Scurto
Ircam - Centre Pompidou
STMS IRCAM–CNRS–SU
Hugo.Scurto@ircam.fr

Bavo Van Kerrebroeck
Ircam - Centre Pompidou
STMS IRCAM–CNRS–SU
Bavo.VanKerrebroeck@ircam.fr

Frédéric Bevilacqua
Ircam - Centre Pompidou
STMS IRCAM–CNRS–SU
Frederic.Bevilacqua@ircam.fr

Abstract
We demonstrate Co-Explorers, a set of interactive agents that collaborate with users in musical space exploration. It implements an interactive reinforcement learning workflow, in which users can guide navigation by giving positive or negative feedback to the agent. It uses a deep reinforcement learning algorithm that optimizes agent exploration toward parts of space that might be of interest for users. Co-Explorers are independent of the type of space explored; the installation that we propose invites attendees to immerse themselves into collaborative sound space exploration.

Author Keywords
Interactive Reinforcement Learning, Exploration, User Experience

Introduction
Exploration is crucial in today's digitalized world [3]. Computer tools need exploration methods to automate information retrieval in ever-increasing parameter spaces. Users also need exploration interfaces to facilitate discovery of subjectively-interesting content. Music is an exemplar application domain for exploration as it holds both such automatic and user-centric features, and requires exploration processes to foster creativity [4].

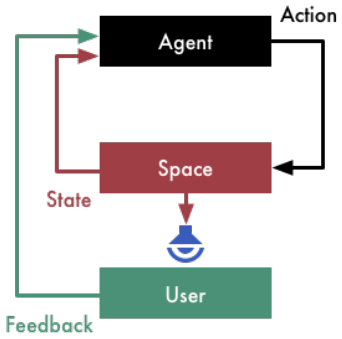

Figure 1: Co-Explorers workflow.

In the field of MIR, computer tools for exploration often consider users as passive. For example, recent recommender systems include notions of novelty and discovery, but often exclude users from customizing the model [5]. We are interested in including users as active in the system's functioning, consciously communicating preferences and guidance to the system throughout exploration. We believe such collaborative paradigm could improve user experience during exploration, while also providing inspiration in more creative and artistic contexts.

Co-Explorers

Co-Explorers are interactive agents that support exploration of large parameter spaces in joint collaboration with their users. They are independent of the type of space explored.

Workflow

Our system workflow is twofold and is shown in Figure 1. On the one hand, the agent explores the musical space at a parametric level. It navigates through the space by observing the space's current parametric configuration (called state) and exploring other parts of the space through direct parameter modification (called actions).

On the other hand, the user explores propositions made by the agent at a perceptual level. It can give positive or negative feedback to the agent regarding its current state (called label feedback), or regarding its current actions (called guidance feedback). It can also give instructions to the agent, having the agent move forward or backward along its exploration path. More details on the workflow design are discussed in [6].

Algorithm

Our system algorithm implements an interactive approach to reinforcement learning, in a way similar to [2].

Reinforcement learning agents intrinsically need to act on their environment to learn optimal behavior [7]. Agents take such exploration actions so as to maximize feedback they will receive from the environment as consequences of their actions. In our interactive approach, feedback is provided by the user. This feedback-driven learning differs from standard supervised learning algorithms, which require training datasets to learn optimal behavior.

In order to tackle approximation to large parameter space, we investigate deep reinforcement learning methods [9]. Deep reinforcement learning support feedback generalization to parts of the space that remain unexplored. We also implemented intrinsic motivation methods to steer the agent toward parts of the space that could offer novelty to users [1]. A complete description of the algorithm is provided in [8].

Interface

The user interface combines tangible and screen-based elements. Visualization of agent trajectories is supported by the screen using different configurations from standard history bars to more advanced dimensionality reduction techniques such as PCA. We are currently implementing coax, a Python software library for human-agent collaborative exploration that uses the OSC protocol to connect agents to musical spaces.

Although Co-Explorers can be connected to any type of musical space, the example we propose focuses on exploring the sonic space defined by a VST. Most mainstream VSTs often present users with complex interfaces requiring technical knowledge of tens of knobs to explore sounds. Co-Explorers thus propose an entirely new paradigm to explore such sonic space, by switching from an analytical, parameter-based view to a more embodied, perception-based interaction.

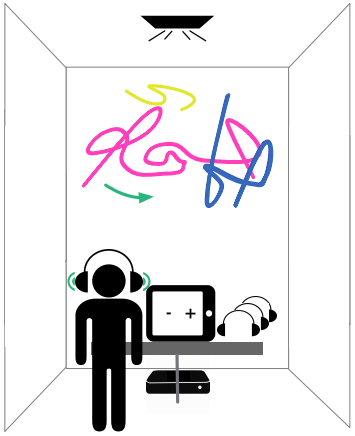

Figure 2: Installation drawing.

Installation

The ISMIR installation invites attendees to immerse themselves in the exploration of a sonic space in collaboration with a Co-Explorer. We envision having the Co-Explorer work continuously across all days of the conference so that all attendees would be allowed to shape its learning.

Description

The installation consists in a minimalist, immersive space, as schematically showed in Figure 2. One tablet allows one attendee to communicate feedback on sound generated by the agent, whose trajectory is projected onto the wall. The more attendees would give feedback on the sounds, the more the agent will learn how to behave sonically—the more trajectories will be projected on the wall. Attendees can spend as much time exploring with the Co-Explorer as they want to. They can use the installation with the help of a supervisor as well as without any supervision. Three pairs of headphones would be available to allow several attendees to share the same experience, even if one attendee actively explores. An example of interaction with the system is demonstrated in the video provided[1].

Technical Specifications

The possible space required for the installation corresponds to approximately 3 meters long, 1 meter wide, and 2 meters high. Table, electrical outlet and extension cord, and area to hang a poster would be required. We would bring our own equipment consisting of a computer, tablet, projector, and headphones. Setup time is about thirty minutes.

[1] https://youtu.be/z7-3ftBMb1I

Acknowledgements

We would like to thank Baptiste Caramiaux and Benjamin Matuszewski for useful discussion and suggestions.

References

[1] Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems* (2016), 1471–1479.

[2] Knox, W. B., and Stone, P. Interactively shaping agents via human reinforcement: The tamer framework. In *Proceedings of the fifth international conference on Knowledge capture*, ACM (2009), 9–16.

[3] Marchionini, G. Exploratory search: From finding to understanding. *Commun. ACM 49*, 4 (Apr. 2006), 41–46.

[4] Resnick, M., Myers, B., Nakakoji, K., Shneiderman, B., Pausch, R., Selker, T., and Eisenberg, M. Design principles for tools to support creative thinking. *Working Paper* (2005).

[5] Schedl, M., Gómez, E., Urbano, J., et al. Music information retrieval: Recent developments and applications. *Foundations and Trends® in Information Retrieval 8*, 2-3 (2014), 127–261.

[6] Scurto, H. e. a. Co-explorers: Designing collaborative ai for exploration in creative design (in preparation).

[7] Sutton, R. S., and Barto, A. G. *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press, 2011.

[8] Van Kerrebroeck, B. Deep reinforcement learning for interactive sonic exploration. Master's thesis, Sorbonne Université, in preparation.

[9] Warnell, G., Waytowich, N., Lawhern, V., and Stone, P. Deep tamer: Interactive agent shaping in high-dimensional state spaces. *arXiv preprint arXiv:1709.10163* (2017).