On the Sensory Commutativity of Action Sequences for Embodied Agents

Extended Abstract

Hugo Caselles-Dupré
U2IS, ENSTA Paris, Institut Polytechnique de Paris & INRIA
casellesdupre.hugo@gmail.com

Michael Garcia-Ortiz
CitAI, SMCSE, City University of London
michael.garcia-ortiz@city.ac.uk

David Filliat
U2IS, ENSTA Paris, Institut Polytechnique de Paris & INRIA
david.filliat@ensta-paris.fr

ABSTRACT

We study perception in the scenario of an embodied agent equipped with first-person sensors and a continuous motor space with multiple degrees of freedom. We consider the commutative properties of action sequences with respect to sensory information perceived by such an embodied agent. We introduce the Sensory Commutativity Probability (SCP) criterion which measures how much an agent's degree of freedom affects the environment in embodied scenarios. We show how to compute this criterion in different environments, including realistic robotic setups and discuss how SCP and the commutative properties of action sequences can be used to learn about objects in the environment and improve sample-efficiency in Reinforcement Learning.

KEYWORDS

Artificial Perception, Unsupervised Learning, Robotics

ACM Reference Format:
Hugo Caselles-Dupré, Michael Garcia-Ortiz, and David Filliat. 2021. On the Sensory Commutativity of Action Sequences for Embodied Agents: Extended Abstract. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 3 pages.

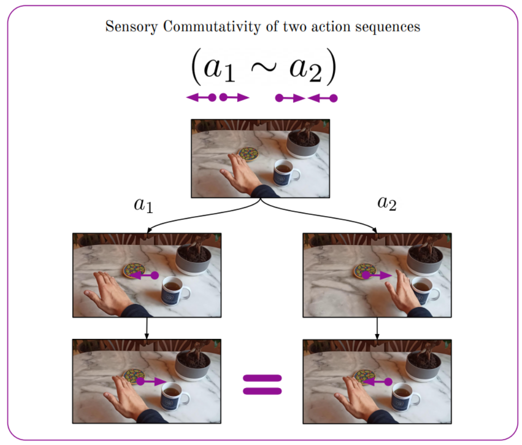

Figure 1: Two action sequences sensory commute if they produce the same sensory state when played in different orders from the same starting position. Here, the action sequences do not commute as an object is in the way of the hand movement.

1 INTRODUCTION

We study the commutativity of action sequences with respect to sensors, which we term sensory commutativity, illustrated in Fig.1. We introduce Sensory Commutativity-experiments (SC-experiments), which consists in having the agent play an action sequence in two different orders from the same starting point, which is our basis for studying sensory commutativity. We define the Sensory Commutativity Probability (SCP) as the probability that a sequence of movements using only one degree of freedom of the agent, an arm joint, for instance, sensory commutes. We show that this value has meaning for the embodied agent: if the SCP is high then the degree-of-freedom has a low impact on the environment (e.g. moving a shoulder is more likely to lead to environment changes than moving a finger, so SCP for a shoulder is lower than for a finger). By computing the SCP for each degree of freedom of the agent, we are able to characterize its motor space and use this information for subsequent tasks. In further experiments [1], we illustrate how SCP, and more generally SC-experiments, can be used to learn about objects in the environment and improve sample-efficiency in a Reinforcement Learning (RL) problem.

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

2 COMMUTATIVE PROPERTIES OF ACTION SEQUENCES

Sensory Commutativity Probability.

We define "degree of freedom" (DOF) as a dimension of the multidimensional continuous action space of the agent. We also define what we term a sensory commutativity experiment: for an action sequence h, the agent plays it in two different orders starting from the same situation.

Definition 1 (Sensory commutativity experiment (SC-experiment)). *Let h be an action sequence of finite length. Let h_p be a random permutation of h (same sequence but different order).*

We define a sensory commutativity experiment (SC-experiment) as playing h and h_p from the same starting point and comparing the two resulting observations in the agent's sensors.

Depending on the nature of the DOF involved in an SC-experiment with an action sequence of length l, the two final observations will either be identical, or different. If the DOFs involved are not associated to displacement, such as opening the eyes/mouth or rotating the head, the action sequence played in two different orders will

Figure 2: Left: External view of the iGibson simulator where the Fetch robot is in a living room. Middle: Fetch's first person view. Right: SCP computed for each of Fetch's degrees of freedom.

commute, and the two resulting observations will be identical. This is what we call Philipona's conjecture, formulated in [3]. The agent is more likely to observe two different outcomes if the action sequence used in the SC-experiment is composed of displacement actions that affect the environment (e.g. movings limbs). There is an even greater chance that they are different if the DOFs are very likely to affect the environment. Consider moving your forearm (elbow joint) compared to moving your whole arm (shoulder joint): the latter is more likely to move things around in the environment and thus induce sensory non-commutativity when played in two different orders (i.e. having two different sensory outcomes). We thus formalize this intuition by defining the Sensory Commutativity Probability (SCP) of a degree of freedom, averaged over all starting situations s. It is defined as the probability that an action sequence composed of actions from one DOF commutes when doing an SC-experiment. In the previous example, an elbow joint should therefore have a higher SCP than a shoulder joint.

SCP computation. We propose a simple procedure to estimate the SCP of each degree of freedom of the agent. We initialize the SCP value to 0 (`SCP←0`). We then repeat the following process n times for each DOF:

- Sample an action sequence of length l using the selected degree of freedom (a sequence of action where each action is a value between -1 and 1).

- Play it in 2 different orders starting from the same randomly chosen state and save the two final sensor images s_1 and s_2. Compute the distance between the two images $d(s_1, s_2)$.

- Count one (`SCP+=1`) if $d(s_1, s_2) \leq t$, zero otherwise.

Finally, the estimator of the SCP is the average over the number of trials (`SCP/`n).

3 EXPERIMENTS IN REALISTIC SIMULATORS

In this experimental section, we compute and interpret the SCP for a realistic embodied agent scenario using the interactive Gibson environment (iGibson) [6].

Simulation description. iGibson is a simulation environment for robotics providing fast visual rendering and physics simulation. It is packed with a dataset with hundreds of large 3D environments reconstructed from real homes and offices, and interactive objects that can be pushed and actuated. In our experiments, we use the Rs environment, which is basically a regular apartment. We place the Fetch robot in this environment, see Fig.2. Fetch is originally a 10-DOF real robot [5] equipped with a 7-DOF articulated arm, a base with two wheels, and a liftable torso. Fetch perceives the environment through a camera placed in his head, see Fig.2.

SCP computation. We use a perceptual distance using the VGG16 [4] features of each observation. We thus have $d(s_1, s_2) = ||VGG16(s_1) - VGG16(s_2)||_2^2$. The choice of the threshold t is arbitrary, we verify in our experiments that a large choice of t leads to equivalent results.

Results. Presented on Fig.2, the results are consistent with Philipona's conjecture. The torso lift DOF is not associated with displacement in the environment, so it has an SCP of 1, i.e. it always sensory commutes.

Moreover, SCP is inversely proportional to how each degree of freedom affects the environment, which is the main claim of the paper. The wheels have the lowest SCP since they provide longitudinal movement and rotations for the robot. Then comes the first DOF of the articulated arm, i.e. the ones that are closer to its base (as expected in the shoulders vs. elbows example). Finally, the highest SCP values correspond to the arm DOF that are further on its arm and the torso lift. We obtain a hierarchical organization of the action space in which the less important dimensions for manipulation and navigation are separated from the dimensions that are not crucial for such tasks.

About the choice of the threshold to compute the SCP, we tried a range of values for t, from 20 to 100, and in each case, we obtain the same hierarchy and thus the same conclusion, only the nominal values change, which is irrelevant for the use of SCP.

In additional experiments, we verified the robustness of these results by computing the SCP for a different type of robot called JackRabbot [2] for which We reach the same conclusions.

Conclusion. We studied the sensory commutativity of action sequences for embodied agent scenarios. We introduced SC-experiments and the SCP criterion. We showed that SCP is a good proxy for estimating the effect of each action on the environment, for 3D realistic embodied scenarios. We also illustrate, in a longer report [1], the potential usefulness of such criterion and SC-experiments in general by performing movable and immovable object detection and improving sample-efficiency in a RL problem.

REFERENCES

[1] Hugo Caselles-Dupré, Michael Garcia-Ortiz, and David Filliat. 2020. On the Sensory Commutativity of Action Sequences for Embodied Agents. arXiv preprint arXiv:2002.05630 (2020).

[2] Roberto Martín-Martín, Hamid Rezatofighi, Abhijeet Shenoi, Mihir Patel, JunYoung Gwak, Nathan Dass, Alan Federman, Patrick Goebel, and Silvio Savarese. 2019. JRDB: A Dataset and Benchmark for Visual Perception for Navigation in Human Environments. arXiv preprint arXiv:1910.11792 (2019).

[3] David Philipona. 2008. Développement d'un cadre mathématique pour une théorie sensorimotrice de l'expérience sensorielle. (2008).

[4] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014).

[5] Melonee Wise, Michael Ferguson, Derek King, Eric Diehr, and David Dymesich. 2016. Fetch and freight: Standard platforms for service robot applications.

[6] Fei Xia, William B Shen, Chengshu Li, Priya Kasimbeg, Micael Edmond Tchapmi, Alexander Toshev, Roberto Martín-Martín, and Silvio Savarese. 2020. Interactive Gibson Benchmark: A Benchmark for Interactive Navigation in Cluttered Environments. IEEE Robotics and Automation Letters 5, 2 (2020), 713–720.