# Low-level causality in a robotic sensorimotor behavior*

**Miroslav Cibula, Igor Farkaš**

Faculty of Mathematics, Physics and Informatics
Comenius University Bratislava
cibula25@uniba.sk, igor.farkas@fmph.uniba.sk

**Abstract**

Humans and other higher animals are able to observe and learn the causal relationships between the actions they take and their perceptual consequences in the environment. This concept has inspired the field of robotics with an argument that understanding causal relations as a fundamental ability is a prerequisite for building advanced AI systems with common sense. In this work, we study and evaluate low-level causal mechanisms related to a robotic arm that learns the sensorimotor dependencies as well as the effects of its motor actions on the environment.

## 1 Introduction

Observing and learning causal relations in an environment is essential to human cognition (Gerstenberg and Tenenbaum, 2017). Thanks to this ability, humans can form intuitive theories from multiple observations and use them to predict the environment behaviour in response to their actions (Gerstenberg and Tenenbaum, 2017). This common sense understanding includes the knowledge of intuitive physics, a key ingredient of early cognitive development (Lake et al., 2016).

In this paper, we take inspiration from causal learning in humans (Lombard and Gärdenfors, 2017) and apply it in the field of robotics (Hellström, 2021). We focus on a low-level approach using a robotic arm in a simulated environment, where the arm (an artificial agent) performs random actions and learns by observing their subsequent effects. This process is implemented by training two complementary models that implement learned knowledge. Then, we use methods to interpret that knowledge in terms of causal relations.

## 2 Problem formulation

We understand a low-level causality as a (transition) function $\mathcal{T}\colon\ [s(t), a(t)] \mapsto s(t+1)$, where $s(t), s(t+1) \in \mathcal{S}$ are the current (pre-action) and the next (post-action) state of the environment, respectively, from a state space $\mathcal{S}$, and $a(t) \in \mathcal{A}$ is an action from an action space $\mathcal{A}$ performed at time $t$. We refer to $\Delta s_a(t+1) = s(t+1) - s(t)$ as an effect of action $a$, reflected in changes of some features of the environment state. Additionally, we understand $\mathcal{T}$ as a low-level intuitive theory encapsulating the accumulated knowledge of causal relations in a given environment.

## 3 Methods

The methods we use involve two components: causal learning and subsequent knowledge extraction. Both are described below.

### 3.1 Causal learning

Causal cognition in robots has been proposed to include a range of categories, varying in terms of complexity (Hellström, 2021). Here, we focus on the lower end of this spectrum and discuss low-level causality regarding two categories: sensorimotor self-learning (C1) and learning the consequences of one's own actions on objects in the environment (C2). Both categories naturally imply embodied knowledge, reflected in arm geometry and kinematics.

Learning of both categories of causality required first offline collection of observation data using motor babbling in a simulated environment for which we used myGym toolkit (Vavrečka et al., 2021). In each step, the agent (robotic arm) executes a randomly selected action $a$ and observes a new state $s(t+1)$. Motor babbling is a natural process observed in infants during their first months. In the case of interaction with objects, the concept of intuitive physics becomes relevant.

In the case of C1, the arm performs motor babbling and records its joint configuration and Cartesian effector position before and after the execution of action $a$. In the case of C2, an object is added on the table in the simulated environment, and the arm has the possibility to interact with it using constrained motor babbling. During an episode, the agent observes potential changes in position, rotation and other defined features of the object, arm and environment in response to the arm actions.

Observations before and after executed action along with the action vector (as per definition of $\mathcal{T}$) collected from the data generation stage are subsequently used for

*supported by project KEGA 022UK-4/2023

the training of two standard models in robotics:

- *forward model* FM: $[s(t), a(t)] \mapsto s(t+1)$ and
- *inverse model* IM: $[s(t), s(t+1)] \mapsto a(t)$.

We implemented both models using supervised feed-forward neural networks. FM is a well-defined, causal model trying to approximate $\mathcal{T}$. In contrast, IM is non-causal since it reverses effects and causes in time. The FM contributes to causal knowledge and provides the agent with predictions of the next state caused by a current action. A trained FM can be utilized for chained inference of several steps ahead (mental simulation) if provided with a sequence of actions.

### 3.2 Knowledge extraction

Trained causal models can be analyzed by extracting information about the original environment and a learning session. Our primary focus is on analyzing feature importance, which allows us to highlight state features that cannot be manipulated by the agent actions and thus can be removed, hence reducing the dimensionality of the state space.

Recent related work by Lee et al. (2021), which served as an inspiration, focused on determining relevant state features by conducting intervention on one feature at a time and testing whether the same policy execution led to successful task completion or not. This way, causal dependencies were found.

On the contrary, we do not study causality by direct interaction with an environment but by using trained causal models as proxies containing this information. Using the learned FM we can determine the relevance of state features in relation to action features by analyzing the feature importance within the FM.

To do this, we take a sample of generated causal data and explain the prediction made by the FM for each instance using Deep SHAP method (Lundberg and Lee, 2017). This method uses attribution rules of DeepLIFT (Shrikumar et al., 2017) analysis technique to approximate Shapley value of each input feature in relation to an output feature. Shapley value represents an input feature contribution to the output feature prediction. This way, we can determine a contribution of each action variable to each state variable. Local explanations are then aggregated across the data sample.

The resulting global contribution heatmap generated from the FM trained on the C2 task is shown in Figure 1. The task consisted of a robotic arm randomly switching its magnetic endpoint. Upon turning the magnet on, the arm navigated to the cube on the table, picking it up and randomly manoeuvring with it in the space. Here, the y-axis denotes the action of each joint and a magnetic endpoint of the arm. The x-axis then contains defined environment state features. The colour of each square corresponds to the magnitude of contributions of action features to the state features averaged across 200

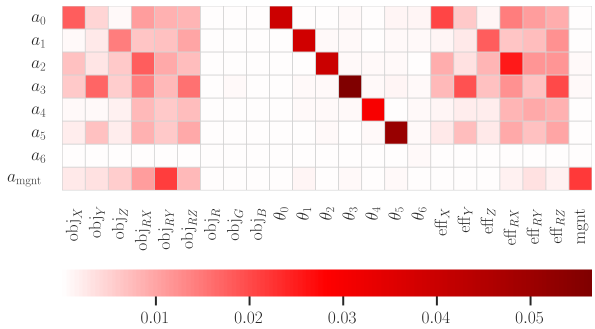


**Fig. 1:** A heatmap generated by Deep SHAP method on the FM showing the contributions of action features' to the state features (object position and rotation, joint angles, effector position and rotation, and magnet state).

samples. The figure shows, for instance, that joint 6 is not used in the sampled observation data. In addition, the colour of an object ($\mathrm{obj}_{\{R,G,B\}}$) is irrelevant in this case as no action can affect it, and thus could be removed (or ignored). On the other hand, all action features affect most object features. This low-level knowledge can be useful for causal analysis at higher levels.

## References


Gerstenberg, T. and Tenenbaum, J. B. (2017). Intuitive theories. In *The Oxford Handbook of Causal Reasoning*. Oxford University Press.

Hellström, T. (2021). The relevance of causation in robotics: A review, categorization, and analysis. *Paladyn, Journal of Behavioral Robotics*.

Lake, B., Ullman, T., Tenenbaum, J., and Gershman, S. (2016). Building machines that learn and think like people. *Behavioral and Brain Sciences*.

Lee, T., Zhao, J., Sawhney, A., Girdhar, S., and Kroemer, O. (2021). Causal reasoning in simulation for structure and transfer learning of robot manipulation policies. In *Int. Conf. on Robotics and Automation*.

Lombard, M. and Gärdenfors, P. (2017). Tracking the evolution of causal cognition in humans. *Journal of Anthropological Sciences*.

Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Int. Conference on Neural Information Processing Systems*.

Shrikumar, A., Greenside, P., and Kundaje, A. (2017). Learning important features through propagating activation differences. In *Int. Conf. on Mach. Learning*.

Vavrečka, M., Sokovnin, N., Mejdrechová, M., and Šejnová, G. (2021). MyGym: Modular toolkit for visuomotor robotic tasks. In *Int. Conf. on Tools with AI*.