Qualitative Models – The Missing Link for Self-Learning Embodied Agents

Jan Lemeire[1], Jure Žabkar[2], Domen Šoberl[3], Marco Van Cleemput[1] and Ruben Spolmink[1]

***Abstract*— Self-learning embodied agents offer a promising path toward achieving general, explainable, and adaptive intelligence. Inspired by the way humans interact with and learn from their environment, we propose qualitative models as a missing component in current AI architectures. These models offer a robust, interpretable, and sample-efficient alternative to dominant data-driven approaches like reinforcement learning. We argue that qualitative models — based on causality and sign relationships — provide foundational capabilities for learning, generalization, planning, quantitative skill acquisition, and symbol grounding. This extended abstract outlines the conceptual framework and empirical potential of qualitative models, advocating their role in advancing autonomous intelligent systems.**

I. Introduction

While deep reinforcement learning has achieved impressive results, it remains data-hungry, opaque, and fragile in novel situations. In contrast, humans learn efficiently by interacting with their environment and forming structured internal representations. *Qualitative models* can form the backbone of a self-learning agent architecture [1] that mimics these human-like learning capabilities.

Qualitative models describe causal relationships in terms of the signs of cause and effect, such as "Activating the motor causes forward motion" or "Turning the steering wheel clockwise results in a leftward turn." Derivatives can also be described by their sign relationship: "An increased motor activation results in a faster acceleration". A sign relationship remains the same for monotone functions. A qualitative relationship holds for the part of the state space in which monotonicity holds. The state space can thus be divided into subspaces in which relationships are present and monotonicity holds. The former is a context-specific dependency. Subspaces are defined by the context: specific values for certain state variables. Examples are: "An object can be moved when you hold it or be next to it and push it" or "An object can be lifted only when it is being held".

II. Five Key Advantages of Qualitative Models

The following 5 properties show the strength of qualitative models.

[1]Jan Lemeire, Ruben Spolmink and Marco Van Cleemput are with the Dept. of Industrial Sciences (INDI) and the Dept. of Electronics and Informatics (ETRO) of the Vrije Universiteit Brussel (VUB) Pleinlaan 2, B-1050 Brussels. jan.lemeire@vub.be

[2]Jure Žabkar is with the Faculty of Computer and Information Scienceof the University of Ljubljana, Ljubljana, Slovenia. jure.zabkar@fri.uni-lj.si

[3]Domen Šoberl is with the the Faculty of Mathematics, Natural Sciences and Information Technologies, University of Primorska, Primorska, Slovenia. domen.soberl@famnit.upr.si

A. Ease of Learning and Adaptivity

Learning qualitative models requires only detection of the *sign* of change [2], not precise numerical values. Transitions between monotonic regions signal new subspaces; boundaries can then be discovered through active exploration. Agents use intrinsic motivation to explore: seeking to maximize influence over variables or probing unexplored regions. When a new actionable variable is discovered, its effects are isolated and modeled. Learning is incremental: newly discovered actionable variables and subspaces are added gradually. The learning is also adaptive: unexpected results - *surprise* - trigger refinement. A first version of this method has been implemented and validated [3].

B. Generalization and Abstraction

Qualitative models generalize more broadly than parameter-specific models like neural networks [4], since they do not depend on particular parameterizations of the situation. Qualitative relationships are valid throughout their subspaces, not limited to the part of the state space covered by the training set. The boundary of each relationship is clearly defined by a change in dependency or tone. Incremental online, see the previous subsection, learning assists in identifying these boundaries as the agent explores.

C. Planning and Control

Qualitative models can be used for high-level planning and, to a certain extent, also for low-level control. To achieve a goal, an agent must first construct a high-level plan consisting of a sequence of subgoals, each representing a necessary condition for reaching the final goal. Next, to accomplish each subgoal, the agent's action parameters must be carefully controlled. We demonstrated that both planning and control can be supported by a contextual qualitative causal model [3]: the contexts define the relevant subgoals, while the qualitative relationships indicate the direction of change, enabling the agent to reach a subgoal via a form of gradient descent. This approach requires the possibility of iterative control. However, in such cases, it does not always lead to the most efficient or optimal trajectory. To address this limitation, a quantitative model may be required. Fortunately, as we will discuss in the next section, qualitative models can support and guide the learning of such quantitative models.

D. Help the Learning of Quantitative Models

When precision is necessary, qualitative models provide directional cues for parameter tuning. This hybrid strategy was used in learning skills such as ball throwing, reverse parking, and reaching. Initially, qualitative causal graphs are

learned from sensorimotor data; then, their structure guides quantitative optimization (e.g., gradient descent). Experiments show skills are learned in under 15 trials, outperforming standard methods by constraining the search space. The proposed approach provides a hypothesis of how humans learn skills: qualitative reasoning happens consciously, while quantitative tuning happens subconsciously.

E. The Link with the Symbolic World

Qualitative models provide an objective bridge between the subsymbolic world and symbolic concepts. Subspaces, each having its distinct behavior, can map to discrete symbols. Consider the following examples. A *wall* blocks the agent from moving from location A to location B. An *obstacle* does as well, but unlike a wall, the agent can navigate around it. A *door*, in contrast, can be opened—meaning it has two states, closed and open, which the agent can actively change—allowing access to location B. An *object* is defined as something the agent can move. A *button*, when pressed, causes a state change in another variable. This grounding of semantics offers a novel, objective route toward symbol emergence, enabling reasoning, abstraction, and verification. We hypothesize this may offer a principled basis for language and representation, though this aspect remains to be explored.

III. Conclusions

We believe that qualitative models represent a crucial yet underutilized component in the pursuit of intelligent embodied agents. Current results are promising, but more research is needed to test the full capacity of this paradigm.

References

[1] T. Wiley, C. Sammut, B. Hengst, and I. Bratko, "A Planning and Learning Hierarchy using Qualitative Reasoning for the Online Acquisition of Robotic Behaviors," *Advances in Cognitive Systems*, 2016.

[2] J. Žabkar, M. Možina, I. Bratko, and J. Demšar, "Learning qualitative models from numerical data," *Artificial Intelligence*, 2011.

[3] J. Lemeire, N. Wouters, M. Van Cleemput, and A. Heirman, "Contextual qualitative deterministic models for self-learning embodied agents," in *Active Inference*, C. L. Buckley, D. Cialfi, P. Lanillos, M. Ramstead, N. Sajid, H. Shimazaki, T. Verbelen, and M. Wisse, Eds. Cham: Springer Nature Switzerland, 2024, pp. 3–13.

[4] D. Šoberl and I. Bratko, "Qualitative control learning can be much faster than reinforcement learning," *Machine Learning*, vol. 114, 01 2025.