Qualitative Causal Models for Autonomous Agents. Position Paper

Jan Lemeire, Jure Žabkar, Domen Šoberl

v1.2 March 2026

Abstract

In this paper, we advocate Qualitative Causal Models (QCMs) as an essential component for cognitive architectures. We outline (1) the improvements that we propose with respect to the state-of-the-art of qualitative models and accompanying algorithms, and (2) the strengths of QCMs in comparison with the current approaches in the field of autonomous agents. By this, we hope to convince the reader of the potential of this line of research.

QCMs describe qualitative relationships between the variables of a dynamical system, more precisely, monotonic relationships between dependent and independent variables. Such relationships are significantly more general than the underlying numerical functions, and divide the state space into subspaces of distinct monotone relationships. We call such such a subspace a *context* of a QCM.

We regard QCMs as a useful link between other components of an effective framework for autonomous AI agents, which consists of methods for learning, reasoning, modeling, state prediction, policy memorization, etc.

We identify the following strengths of QCMs: (1) They allow sample efficient incremental learning from numerical data, which are obtained through experimentation and state space exploration; (2) high level of model generalization, which comes from abstracting away the numerical information; (3) the capability of performing qualitative simulation, planning and control; (4) explainability, which comes from the fact that qualitative representations are closer to human intuition than numerical representations. Moreover, we believe that QCMs are part of the human brain.

We are currently working on the underlying framework and algorithms to deliver convincing results this year.

1 Introduction

Despite the impressive results of current AI technology, several of the long-standing open problems in AI remain unsolved, e.g. explaining agents' decisions, linking of quantitative models with the symbolic world, and learning efficiency (compared to human learning). Deep reinforcement learning remains data-hungry, opaque, and fragile in novel situations. Researchers discuss whether a trained model has acquired a deep understanding of the system under study, or just a set of contextual rules; "a bag of heuristics instead of a world model" (Mitchell, 2025b). Numerous researchers argue that tackling these issues requires world models that are both causal and abstract (LeCun and Courant, 2022; Mitchell, 2025a,b). We propose *Qualitative Causal Models* (QCMs) as an essential component of cognitive architectures. These can form the backbone of a self-learning agent that mimics human-like learning capabilities (Wiley et al., 2014, 2016a; KOŠMERLJ et al., 2011), which means learning efficiently by interacting with the environment and forming structured internal representations.

A QCM can serve as an abstract, "coarse-grained" world model. It describes the qualitative properties of causal relationships among variables in terms of *signs* of cause and effect, such as "Activating the motor causes forward motion" (positive causal relation) or "Turning the steering wheel clockwise results in a leftward turn." (negative causal relation). Variables can also be causally linked with their time derivatives, e.g.: "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. Examples are: "Moving forward decreases your distance to the object when you are oriented toward it; otherwise, the distance increases", or "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". Qualitative relationships are defined by context, specifically by the values of certain state variables that partition the state space into subspaces. These values are called *landmarks*. Some causal relationships only hold in certain subspaces; they are referred to as *context-specific dependencies*. The essential components of a QCM are the qualitative descriptions and the subspaces of monotone behavior.

Section 2 provides a formal definition of the proposed modeling framework. It is compared with the current approach to defining qualitative models in Section 3. We then show how this class of models integrates key principles of qualitative modeling, system modeling, and causal modeling. Next, in Section 4, the strengths of the proposed model are discussed and compared with the state-of-the-art.

Section 5 lists the properties that a cognitive architecture should have. Section 6 shows how these properties could be achieved by combining qualitative causal models with other knowledge components. The relation with the human brain is briefly discussed in Section 7.

2 Definition

Let $\boldsymbol{X}^t = [X_1^t, X_2^t, ..., X_n^t]$ a set of continuous, random variables at time step t[1]. The set of random variables consists of state variables $\boldsymbol{S}$ and action variables $\boldsymbol{A}$. The proposed model class provides a description of $P(\boldsymbol{S}^t j \boldsymbol{S}^{t-1}, \boldsymbol{A}^t)$.

[1]Boldface letters indicate sets.

A *Qualitative Causal Model* is a tuple $\langle \boldsymbol{S}, \boldsymbol{A}, \boldsymbol{Q}, G, \boldsymbol{F} \rangle$. $\boldsymbol{Q}$ is the set of Quantity Space Functions $\mathcal{Q}_i$ that convert the continuous variable X_i into a finite set of (qualitative) values. A quantity space (Kuipers, 1994; Forbus, 1984) is defined by a finite and totally ordered set of landmarks (L_i) $[l_1, \ldots, l_{L_i}]$ that partition the domain X_i into a set of $2L_i{+}1$ mutually disjoint open intervals $[(-\infty, l_1), l_1, (l_1, l_2), l_2, ..., l_n, (l_n, +\infty)]$. The function $\mathcal{Q}_i$ returns the index of the interval in which the value of X_i resides. The default landmark is the value 0. The quantity space function then returns the algebraic *sign* of the variable. We will use the term sign as the index returned by $\mathcal{Q}_i$ even when the set of landmarks is extended beyond the value 0. G is a Directed Acyclic Graph (DAG) defined over variables $\boldsymbol{S}^t$, $\boldsymbol{S}^{t-1}$, and $\boldsymbol{A}^t$, with no edges into the variables of $\boldsymbol{S}^{t-1}$ and $\boldsymbol{A}^t$ (only outgoing edges). $Pa(X_i)$ denote the parents of variable X_i in G. $\boldsymbol{F}$ is the parameterization of each variable of $\boldsymbol{S}^t$ given its parents. This parameterization is either a deterministic function $X_i = f_i(Pa(X_i))$ or a qualitative function $\mathcal{Q}(X_i) = qf_i(Pa(X_i)) = f_i(\mathcal{Q}(Pa(X_i))$ over the signs of the parent variables.

The quantity spaces of the variables of $\cup_i Pa(\boldsymbol{S}_i^t)$ is chosen such that the qualitative functions qf_i exist. This is the assumption of *qualitative determinism*.

The causal interpretation of the relations of a QCM follows from the fact that action variables have their effects represented as their descendants in the graph. Those edges are the causal edges of the graph.

Variables defined by a deterministic function in terms of other variables are called *derived variables*. A special case is the time derivative of a state variable, such as velocity, which can be defined using x^t and x^{t-1}.

An edge between a state variable S_j^{t-1} and S_i^t is called a *contextual edge*, as the sign of the variable S_j^{t-1} modulates the qualitative effect of the action variables on S_i^t. Variable S_j^{t-1} determines the context. It partitions the state space into *subspaces*, also called *operating regions* (Šoberl, 2021, Sec. 3.1.3), of specific monotone relations. If a specific context is required for the existence of a causal relationship, it is referred to as a contextual dependency. The edge is inactive when the system is not in the appropriate context.

Figure 1 shows an example QCM of a mobile robot with a gripper. It has 4 possible actuations and can change its orientation, its position, the distance to the object, hold the object, and lift it.

3 Difference with state-of-the art

Qualitative models have been used in several domains, such as qualitative reasoning (Kuipers, 1994; Forbus, 2018), naive physics (De Kleer and Brown, 1984), qualitative simulation (Bredeweg et al., 2009) and planning (Wiley et al., 2016b). There are, however, essential differences between the proposed model definition and the ones used in previous work.

Qualitative model theory started with the pioneering work of (Kuipers, 1986). The model comprises qualitative constraints derived from the Ordinary Differential Equations (ODEs) that describe the system under study. The Q-constraints are defined as

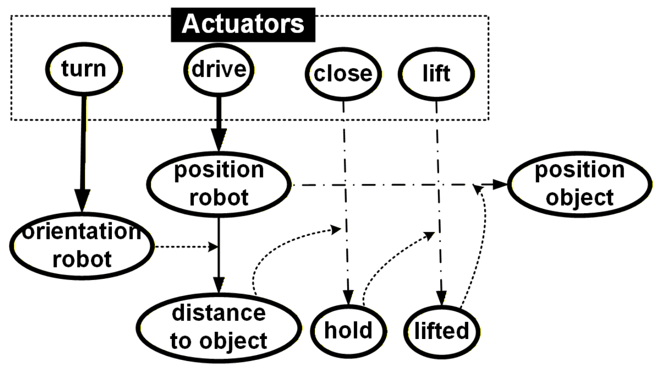

Figure 1: Qualitative Causal Model of a mobile robot with a gripper. Thick arrows represent causal edges, thin arrows represent deterministic relations of a derived variable. Dashed arrows are contextual edges. Alternate point-dash arrows are inactive edges because of a contextual dependency. To simplify the graph, the state variables at $t-1$ are omitted. Contextual edges start in the state variables at $t-1$, while the other edges are connected to the state variables at t.

follows (Žabkar et al., 2007):

$$\begin{array}{lll} y = Q^{+}(x) & \text{means} & \frac{\partial y}{\partial x} > 0, \\ y = Q^{-}(x) & \text{means} & \frac{\partial y}{\partial x} < 0. \end{array} \tag{1}$$

A positive derivative means that the function $y = f(x)$ is monotonically increasing, while a negative derivative corresponds to a monotonically decreasing function.

Instead of a set of qualitative constraints, we propose a causal model that is based on the approach used in the field of *system modeling*. It is shown that systems described by higher-order ODEs, can be converted into a *state-space representation* by expressing them as a set of coupled first-order ODEs (Franklin et al., 2002)[Sec. 2.2]. For this transformation, additional variables, such as velocity, might be added to link acceleration with position. Such a state-space representation is widely used in control engineering. A mathematical model of a physical system is then specified as a set of input, output, and state variables related by first-order differential equations. These equations can then be discretized. In our model, each equation is specified by a qualitative function. Note the similarity with Pearl's causal model theory which relies on a set of structural equations (Pearl, 2000) resulting in a Bayesian network. The relations among the variables are then given a causal interpretation based on the $do()$-operator. A dynamic system is described by a dynamic Bayesian network. Our model has the same structure, except for the qualitative parameterization of the relations, similar to the work of Mugan et al. (Mugan and Kuipers, 2012).

We provide an objective criterion for partitioning the state space into subspaces and the quantity spaces following from them. Such a clear and unequivocal criterion is lacking in related work. Dekleer (De Kleer and Brown, 1984, Sec. 2.2) on quantity spaces: "make a judicious choice for the qualitative intervals such that as little information as possible is lost". Cartoni (Cartoni et al., 2023) proposes 'a planning approach that dynamically increases abstraction'. The abstraction is based on the distance in the state space between known and novel points, which does not guarantee successful generalization.

These models, as well as those used in reinforcement learning (discussed in more detail in Section 4.3), are probabilistic, whereas QCMs assume qualitative determinism. This assumption is motivated by the observation that most robotic and autonomous agent systems are coarse-grained deterministic—a car, for instance, will not move backward when the motor command is positive. Nevertheless, some stochasticity can still be incorporated into the model, such as the possibility of failing to grasp an object in the example from the previous section. In any case, qualitative simulation supports future bifurcations, since in some situations quantitative state values are required to predict the qualitative evolution of a system (De Kleer and Brown, 1984, Sec. 22): "not possible to determine unambiguously the qualitative behavior".

4 Strengths

The following outlines the potential of our modeling framework and compares it with state-of-the-art methods.

4.1 Ease of Learning and Adaptivity

Learning qualitative models requires only detection of the *sign* of change, not precise numerical values. As such, fewer samples are required than in reinforcement learning, which often requires millions of epochs to reach convergence. Padé is a tool for learning qualitative models, defined by a set of qualitative constraints, from numeric data (Žabkar et al., 2011). We are extending this to the learning of the proposed QCMs. A first version has been implemented and validated (Lemeire et al., 2024). Transitions between monotonic regions, i.e. a change of the sign relationship, signal the emergence of 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 the need for model refinement.

A QCM learning algorithm will be based on the principles of *intrinsic motivation*. It should be driven by exploring novelties, improving predictions, and competencies.

4.2 Generalization and Abstraction

Qualitative models generalize more broadly than parameter-specific models like neural networks (Šoberl and Bratko, 2025), since they do not depend on particular parameterizations of the situation. Qualitative relationships generalize much better than quantitative approaches. For instance, "turning the steering wheel clockwise results in a leftward turn" is valid for almost any vehicle, while the quantitative relation between steering wheel angle and turning angle is vehicle-dependent. Learning a quantitative model

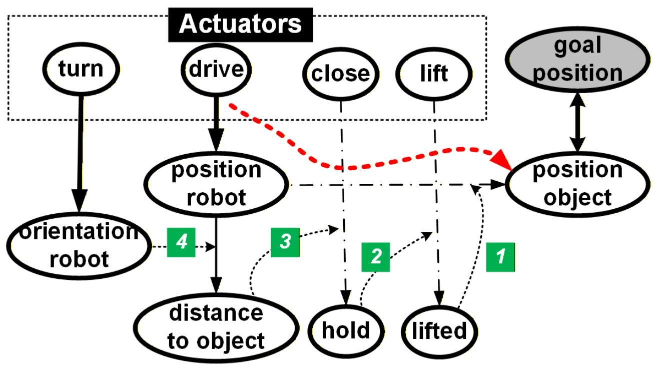

Figure 2: Planning with a Qualitative Causal Model: to activate the path between an action and the target variable (in red), inactive edges have to be activated by reaching the right context. This results in a sequence of subgoals. First turn towards the object (4), then drive towards it (3), grab the object (2), and lift it (1). Then the object can be moved to the goal location.

is easier when it is based on a generic qualitative model, as discussed in Section 4.4. Next, the validity and boundary of each qualitative relationship is clearly defined by an easily detectable change in dependency or tone. Incremental online learning - see the previous subsection - assists in identifying these boundaries as the agent explores.

4.3 Planning and Control

Autonomous agents that must carry out tasks require both high-level planning and low-level motion control. Qualitative models have been used for control (Šoberl and Bratko, 2017). The proposed QCMs can also be used for high-level planning. 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. Figure 2 shows the construction of a plan for the task of picking up an object and bringing it to the goal location: the contexts define the relevant subgoals (Lemeire et al., 2024).

Next, to accomplish each subgoal, the agent's action parameters must be carefully controlled. The signs of the relationships indicate the direction of change resulting from actions. This enables the agent to reach the subgoal (a subspace defined by the context) via a form of gradient descent, as shown by Figure 3. However, it does not always lead to the most efficient or optimal trajectory. To address this limitation, a quantitative model may be required, which is discussed in the next subsection.

Reinforcement learning (RL) is the major approach for self-learning autonomous agents. It assumes that the underlying system is a Markov Decision Problem (Sutton et al., 1998) that can be stated as $\boldsymbol{S}^{t+1} = P(\boldsymbol{S}^t, \boldsymbol{A}^t)$. Model-free RL learns a policy π that maximizes the reward signal. The policy is a function from state $\boldsymbol{S}$ to action $\boldsymbol{A}$, while in a model-based approach, the optimal action has to be derived by solving an inverse problem with the model of $\boldsymbol{S}^{t+1} = P(\boldsymbol{S}^t, \boldsymbol{A}^t)$. We outlined here how this can be done with a QCM.

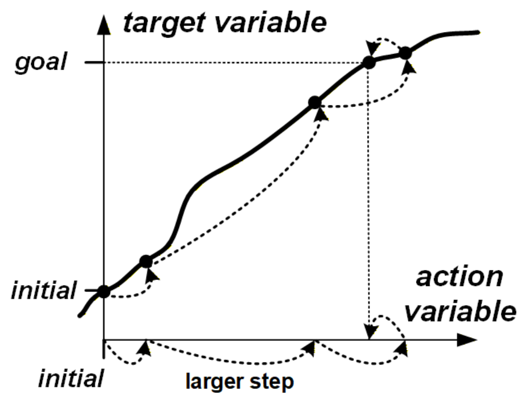

Figure 3: Due to the monotonicity of relations within a subspace, a basic feedback controller can achieve the desired value of a target variable by regulating the corresponding action variable, taking into account the sign relationship.

Pure symbolic planning—such as approaches based on the Planning and Domain Definition Language (PDDL)—fails to capture key qualitative structures, such as monotonic trends, landmarks, and orders of magnitude, that humans and agents rely on for reasoning.

4.4 Guide 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 (Šoberl et al., 2024), 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.

4.5 The Link with the Symbolic World

Qualitative models provide an objective bridge between the subsymbolic world and symbolic concepts. Subspaces, each with its distinct behavior, can be mapped 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.

In the field of self-learning embodied agents, bottom-up learning of the symbolic language for planning, see also Section 4.3, PDDL has been studied by various researchers (Konidaris et al., 2014; Ahmetoglu et al., 2022, 2025). They, however, lack an objective criterion for discretizing the continuous world. Ahmetoglu et al. (Ahmetoglu et al., 2022) base the symbol formation and rule extraction on the latent vector trained in an encoder-decoder setup. Our partitioning of the continuous state space is defined by changes in tone in the cause-and-effect relations, providing an objective method for discretization. This grounding of semantics offers a novel, objective route toward symbol emergence, enabling reasoning, abstraction, and verification. This may provide a partial, principled basis for language and representation, one grounded in the *interaction* of agents with their environment.

4.6 Explainability and intuitiveness

The models address the need for explainability in current AI technology. Qualitative causal models appear to be very understandable and highly intuitive. Consider the plan and control loop inferred to reach a goal, as discussed in Section 4.3. Each decision can be explained by referring to the causal model: why a subgoal should be attained and why a certain action will bring you closer to the (sub)goal.

5 Necessary properties

An embodied agent that can effectively and efficiently learn, adapt, and deal with novel situations, it should be able to quickly acquire knowledge and 'juggle' with its knowledge in a flexible way. The essential properties of a knowledge system are generalization, adaptability, reuse across similar settings, composition, explainability, and links between different knowledge representations (as discussed in section 6).

An hypothesis is that *explicitness* and the *disentanglement* of knowledge (as opposed to its dispersion or obscurity in a network) are necessary (or at least helpful) for several of the above properties. Although RL can achieve convincing results (after extensive training), when based on monolithic neural networks, it does not exhibit most of these properties.

6 Component of a cognitive architecture

QCMs are an essential component of a successful cognitive architecture if combined with other modules, as shown in Figure 4. QCMs are rooted in the subsymbolic world, but need a correct state estimation from observations, and, where necessary, a map of the environment and/or 3D models of the objects to enable spatial reasoning. As discussed earlier, QCMs can provide a link between the subsymbolic and symbolic worlds and, as such, open the gate to the power of symbolic approaches. It can also help tune the parameters of quantitative models. Next, an agent stores episodes, plans, decisions, and outcomes in memory to guide future decisions, steered by a form of reinforcement learning.

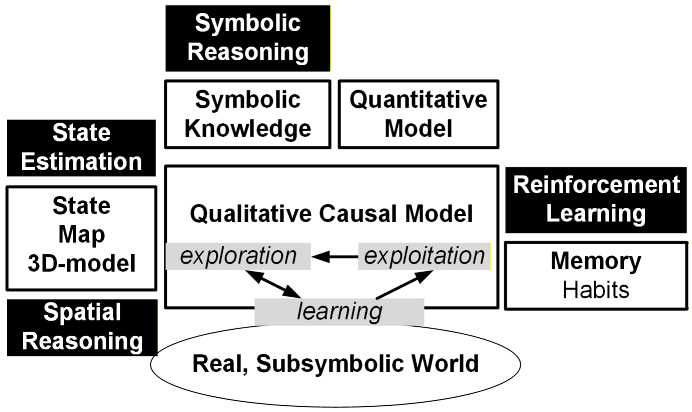

Figure 4: The Qualitative Causal Model and the other modules for constituting a cognitive architecture for an autonomous agent.

It is known that the human brain consists of several layers with different functions that, together, make decision-making.

7 The human brain

QCMs provide a hypothesis for part of the human cognitive machinery. This can be motivated by the following example. Learning to drive a car starts with understanding the qualitative relationship between turning the steering wheel in a certain direction and the car turning. Then one learns a quantitative model of how much steering is needed for a certain turn. For driving another car, the quantitative model's parameters need to be tuned while the qualitative relationships remain the same. It shows how qualitative relationships allow for generalization and reuse and can help in the learning of quantitative models, since the functions' monotonicity is known.

8 Conclusions

The AI community has been actively seeking new modeling paradigms to bridge the gap created by the inability of modern quantitative models to perform genuine reasoning. We believe that qualitative causal models represent a crucial yet underutilized component in the pursuit of intelligent embodied agents. The initial results are promising, but further research is necessary to fully evaluate the potential of this paradigm. Can qualitative causal models form the cognitive architecture's backbone on top of which other cognitive components can be constructed, such as qualitative reasoning, quantitative descriptions, habits, and symbolic formalisms? Does they reveal anything about human cognition?

References

Ahmetoglu, A., Öztop, E., and Ugur, E. (2025). Symbolic manipulation planning with discovered object and relational predicates. *IEEE Robotics Autom. Lett.*, 10(2).

Ahmetoglu, A., Seker, M. Y., Piater, J., Oztop, E., and Ugur, E. (2022). Deepsym: Deep symbol generation and rule learning for planning from unsupervised robot interaction. *J. Artif. Int. Res.*, 75.

Bredeweg, B., Linnebank, F., Bouwer, A., and Liem, J. (2009). Garp3 — workbench for qualitative modelling and simulation. *Ecological Informatics*, 4(5):263–281. Special Issue: Qualitative models of ecological systems.

Cartoni, E., Montella, D., Triesch, J., and Baldassarre, G. (2023). Real-x—robot open-ended autonomous learning architecture: Building truly end-to-end sensorimotor autonomous learning systems. *IEEE Transactions on Cognitive and Developmental Systems*, 15(4).

De Kleer, J. and Brown, J. S. (1984). A qualitative physics based on confluences. *Artificial Intelligence*, 24(1):7–83.

Forbus, K. D. (1984). Qualitative process theory. *Artificial intelligence*, 24:85–168.

Forbus, K. D. (2018). *Qualitative representations: How People Reason and Learn about the Continuous World*. MIT Press.

Franklin, G., Powell, J., and Emami-Naeini, A. (2002). *Feedback Control Of Dynamic Systems, 4th edition*. Prentice Hall.

Konidaris, G., Kaelbling, L., and Lozano-Perez, T. (2014). Constructing Symbolic Representations for High-Level Planning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 28(1). Number: 1.

KOŠMERLJ, A., BRATKO, I., and ŽABKAR, J. (2011). Embodied concept discovery through qualitative action models. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 19(03):453–475.

Kuipers, B. (1986). Qualitative simulation. *Artificial Intelligence*, 29(3):289–338.

Kuipers, B. (1994). *Qualitative reasoning: modeling and simulation with incomplete knowledge*. MIT Press, Cambridge, MA, USA.

LeCun, Y. and Courant (2022). A path towards autonomous machine intelligence version 0.9.2, 2022-06-27.

Lemeire, J., Wouters, N., Van Cleemput, M., and Heirman, A. (2024). Contextual qualitative deterministic models for self-learning embodied agents. In Buckley, C. L., Cialfi, D., Lanillos, P., Ramstead, M., Sajid, N., Shimazaki, H., Verbelen, T., and Wisse, M., editors, *Active Inference*, pages 3–13, Cham. Springer Nature Switzerland.

Mitchell, M. (2025a). Llms and world models, part 1.

Mitchell, M. (2025b). Llms and world models, part 2.

Mugan, J. and Kuipers, B. (2012). Autonomous Learning of High-Level States and Actions in Continuous Environments. *IEEE Transactions on Autonomous Mental Development*, 4(1):70–86.

Pearl, J. (2000). *Causality. Models, Reasoning, and Inference*. Cambridge University Press.

Sutton, R. S., Barto, A. G., et al. (1998). *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.

Šoberl, D., Lemeire, J., Van Cleemput, M., and Žabkar (2024). Safe parking of a nonholonomic autonomous vehicle by qualitative reasoning. In *Proceedings of the 37th International Workshop on Qualitative Reasoning*.

Žabkar, J., Bratko, I., and Demšar, J. (2007). Learning qualitative models through partial derivatives by Padé. In *Proceedings of the 21st Annual Workshop on Qualitative Reasoning*, pages 193–202.

Wiley, T., Sammut, C., and Bratko, I. (2014). Qualitative planning with quantitative constraints for online learning of robotic behaviours. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 2578–2584, Palo Alto, California USA. AAAI Press.

Wiley, T., Sammut, C., and Bratko, I. (2016a). A planning and learning hierarchy using qualitative reasoning for the on-line acquisition of robotic behaviors. *Advances in Cognitive Systems*, 4:93–112.

Wiley, T., Sammut, C., Hengst, B., and Bratko, I. (2016b). A Planning and Learning Hierarchy using Qualitative Reasoning for the Online Acquisition of Robotic Behaviors. *Advances in Cognitive Systems*, page 19.

Šoberl, D. (2021). *Automated planning with induced qualitative models in dynamic robotic domains*. PhD thesis, University of Ljubljana.

Šoberl, D. and Bratko, I. (2017). Reactive motion planning with qualitative constraints. In *Advances in Artificial Intelligence: From Theory to Practice*, pages 41–50. Springer International Publishing.

Šoberl, D. and Bratko, I. (2025). Qualitative control learning can be much faster than reinforcement learning. *Machine Learning*, 114.

Žabkar, J., Možina, M., Bratko, I., and Demšar, J. (2011). Learning qualitative models from numerical data. *Artificial Intelligence*, 175(9):1604–1619.