Complex Model Transformations by Reinforcement Learning with Uncertain Human Guidance

Kyanna Dagenais*[iD], Istvan David*†[iD]
*McMaster University, Hamilton, Canada
†McMaster Centre for Software Certification, Hamilton, Canada

arXiv:2506.20883v2 [cs.SE] 6 Aug 2025

Abstract—**Model-driven engineering problems often require complex model transformations (MTs), i.e., MTs that are chained in extensive sequences. Pertinent examples of such problems include model synchronization, automated model repair, and design space exploration. Manually developing complex MTs is an error-prone and often infeasible process. Reinforcement learning (RL) is an apt way to alleviate these issues. In RL, an autonomous agent explores the state space through trial and error to identify beneficial sequences of actions, such as MTs. However, RL methods exhibit performance issues in complex problems. In these situations, human guidance can be of high utility. In this paper, we present an approach and technical framework for developing complex MT sequences through RL, guided by potentially uncertain human advice. Our framework allows user-defined MTs to be mapped onto RL primitives, and executes them as RL programs to find optimal MT sequences. Our evaluation shows that human guidance, even if uncertain, substantially improves RL performance, and results in more efficient development of complex MTs. Through a trade-off between the certainty and timeliness of human advice, our method takes a step towards RL-driven human-in-the-loop engineering methods.**

Index Terms—**human guidance, machine learning, uncertainty**

I. INTRODUCTION

Modeling activities are often more complex than an atomic model transformation (MT) and rely on *sequences of MTs*. Pertinent examples can be found in model synchronization [1], model refactoring [2], and rule-based design-space exploration [3]. Typically, there might be more than one MT sequence that can successfully transform the source model into the target state, and choosing the most appropriate (cost-effective, efficient, safe) one manually is not tractable. This raises the **need for automated methods for developing complex MTs**, in which MTs are chained in sequences.

Reinforcement learning (RL) is a machine learning (ML) paradigm, in which an agent explores the state space, and through trial and error, learns beneficial action sequences [4]. This ability positions RL as a potential automation method for developing complex MTs. The benefits of RL have been recognized in MDE before, e.g., in model repair [5], in-place MTs [6], and model-driven service migration [7]. However, RL methods perform poorly in the early phases of learning due to suboptimal decisions before converging to better policies [4]. That is, RL agents may require substantial training time and resources before learning an optimal sequence of MTs.

This shallow learning curve limits the applicability of RL, particularly in cases where early mistakes can be costly, e.g., synchronizing models in model-based design of safety-critical systems. While the performance of RL agents improves over time, this improvement often comes too late to be useful.

Human guidance is an apt technique to mitigate problems, such as the shallow learning curve of RL agents [8]. In such a setup, human advice acts as a heuristic to hint the otherwise autonomous agent towards desirable goals, thereby improving its performance. Involving human guidance in automated MDE methods has been explored before, e.g., driving RL-based MTs for model repair by personalized preferences [9], and driving search-based methods for MT-based design space exploration [10]. A common theme of existing techniques is the assumption that firm advice can be produced by the human in a timely and cost-effective fashion. This is a strong assumption that cannot always be guaranteed, especially when the human is asked to provide guidance about complex MTs (e.g., identifying consistent intermediate states in a long refactoring, or known locally optimal configurations in design-space exploration). Human advice, though often uncertain, can still be of high value in driving RL agents [11]. This raises the **need for incorporating human advice into RL-driven MTs through sound modeling and quantification of uncertainty**.

In this paper, **we develop an approach and technical framework to infer complex MTs through RL, guided by potentially uncertain human advice**. The sound formalization and assessment of uncertainty allows for producing useful human advice early in the MT learning process. Depending on the specific situation and engineering problem, early uncertain advice might be of higher value than late certain advice. We evaluate our approach in synthetic and real human guidance scenarios. Our results indicate that human guidance, even if uncertain, substantially improves RL performance, and allows for the efficient inference of complex MTs.[1]

Through a trade-off between the certainty and timeliness of guidance, our method takes a firm step towards ML-driven human-in-the-loop engineering methods. Our approach alleviates the manual labor required for developing complex MTs while enabling domain experts to steer the inference process. This convergence of human and artificial intelligence is much sought-after in complex engineering domains.

---

[1]A replication package with detailed replication instructions, results, and analyses is publicly available at https://doi.org/10.5281/zenodo.16321499.

II. Background and Related work

Here, we introduce the key concepts of our approach: reinforcement learning (Sec. II-A) and its guidance by humans (Sec. II-B); and we discuss the related work (Sec. II-D).

A. Reinforcement learning

In RL, an *agent* learns optimal control of an *environment* by taking sequential actions to explore the environment and update its strategy based on feedback in the form of *rewards* [4]. This can be formalized as a Markov decision process [12] $\langle S, A, \pi(a|s), r_t(s, a) \rangle$. S denotes the set of observable states the environment can produce, and A is the set of actions the agent can take. $r_t(s, a) \in \mathbb{R}$ is the reward the environment produces when the agent chooses action $a \in A$ while observing state $s \in S$, at time step t. At time step t, the agent observes a state s_t and chooses an action a_t. In response, at time step $t + 1$, the environment produces a reward r_{t+1} and a new state s_{t+1}. The agent maintains a policy $\pi(a|s)$, which gives the conditional probability of choosing action a in state s. The agent's goal is to learn and act according to an optimal policy π_* that maximizes the sum of future rewards.

RL has seen growing use in support of MDE, e.g., in MT engineering [6], model repair [9], fault injection testing [13], and inference of simulators for digital twins [14]. **In this work**, we use RL to learn complex MTs, i.e., long sequences of MTs that are infeasible to reason about for human modelers.

B. Human guidance in RL

While RL algorithms can operate autonomously, evidence suggests that human guidance is beneficial to learning dynamics. Guidance is a method in which human input is used to bias the agent's exploration strategy. Techniques differ by when and how advice is provided, as described in the taxonomy of Najar and Chetouani [8]. For example, aligned with our approach, Cruz *et al.* [15] emphasize the role of *early advice*.

Strategies to incorporate guidance into RL are called *policy shaping* methods, which are defined by which part of the RL process they are incorporated in. In value shaping, human advice acts as a function of action preferences [16]. In reward shaping, advice is translated into numerical rewards given to the agent [17]. In decision biasing, human advice is used to directly affect the policy output [18]. Finally, in policy shaping, human advice is directly incorporated into the agent's policy to bias the exploration strategy. For example, Griffith *et al.* [19] estimate the human's feedback policy and incorporate this information into the agent's policy.

By this taxonomy, **our work** falls under the category of guidance, and specifically, **policy shaping**, since we aim to bias the agent's exploration by human guidance.

C. Uncertainty and Subjective Logic

Using human guidance to bias the RL agent's policy requires a formal approach to appropriately capture human opinions. In particular, the epistemic uncertainty of human opinions must be modeled appropriately to make guidance meaningful. Subjective Logic (SL) [20] is a formalism that defines an opinion, a construct composed of belief and an orthogonal component of uncertainty. A binomial opinion is defined as $\omega_x = (b_x, d_x, u_x, a_x)$. This opinion is about the truth of a Boolean predicate x. Here b_x is belief in x, d_x disbelief in x, u_x vacuity of evidence of x, and a_x prior probability of x. The parameters must adhere to the constraints $b_x, d_x, u_x, a_x \in [0, 1]$ and $b_x + d_x + u_x = 1$. A binomial opinion is transformed into a probability $P(x) = b_x + a_x u_x$. Likewise, probability $p(x)$ translates to binomial opinion $\omega = (p, 1 - p, 0, p)$. SL defines fusion operators that allow for multiple opinions to be fused into one opinion.

The utility of SL in MDE has been demonstrated, e.g., in modeling with uncertain types [21–23] and prioritizing inconsistency resolution steps [24]. **In this work**, we use SL to model the uncertainty of guidance for RL in MT development.

D. Related Work

Closest to our work are approaches that use RL for MTs, particularly with human in the loop for guidance or intervention. The PARMOREL framework [9, 25] uses RL to learn user-preferred sequences of model transformations to fix invalid models. While this approach features human input for choosing an appropriate reward function, it does not accommodate the uncertainty of human guidance.

Eisenberg *et al.* [6, 26] use RL for learning in-place MTs. In this approach, the RL agent learns a sequence of MTs that optimizes objectives in the final model. The authors demonstrate that RL is a feasible alternative to traditional search heuristics, e.g., genetic algorithms. Our work improves over these approaches by including human guidance in RL.

Basciani *et al.* [27] introduce the CITRIC tool to recommend developers optimal MT sequences between a source and target metamodel. The approach ranks MT sequences through shortest path algorithms by two relevant criteria: metamodel coverage and information loss. Our approach is similar in aims but differs in technological solution, including the role of the human, who is an active advisor in our approach.

Finally, some adjacent techniques that automate MT development, but not by RL, include the following. Eisenberg *et al.* [28] use multi-objective optimization to explore MT sequences using meta-heuristic search algorithms. Abdeen *et al.* [3] use a genetic algorithm for multi-objective optimization in rule-based design-space exploration. Mokaddem *et al.* [2] learn refactoring rules from examples using genetic programming.

III. Illustrative example

To illustrate our approach, we rely on the following running example of a self-driving vehicle through a frozen lake (Fig. 1). The example is analogous to Open AI's Gym's Frozen Lake environment[2], but it is fully modeled in EMF and its operational semantics are defined by MTs, akin to the Pacman case, a frequently used example in MDE [6, 29, 30].

The vehicle explores the lake by choosing to move up, down, left, or right one tile at a time. It begins its exploration

---

[2]https://gymnasium.farama.org/environments/toy_text/frozen_lake/

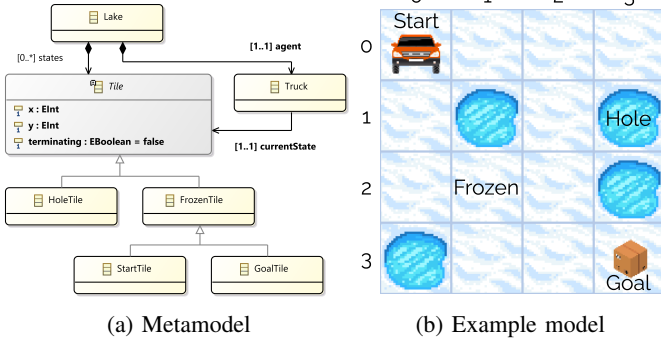(a) Metamodel   (b) Example model

Fig. 1: The Frozen Lake running example

at the top leftmost tile (*Start*). Its aim is to reach the bottom rightmost tile (*Goal*) while avoiding terminating states (*Hole*s).

Model transformations

Operational semantics are given by MTs. As shown in Fig. 2, moving to the tile to the right is achieved by executing the `moveRightRule` MT rule, detailed in Listing 1. To illustrate our approach, we use the VIATRA framework [31].
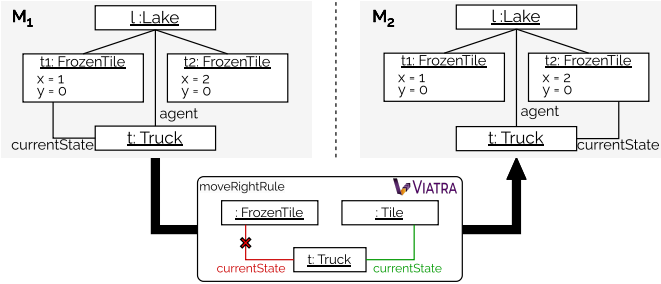

Fig. 2: Example MT rule and its execution

```
1  val moveRightRule =
2    createRule(agentOnFrozenTile)
3  .name("MoveRightRule")
4  .action[
5    val newTile = tiles.findFirst[
6      it.x==tile.x+1 && it.y==tile.y].head
7    truck.tile = newTile].build
```

Listing 1: Example VIATRA MT rule moving the agent right

To successfully navigate through the lake, the vehicle must find a MT sequence that brings it to the goal state while avoiding terminating states (holes).

Reinforcement learning

In RL terms, the MT sequence will generate *reward* upon encountering different states as follows.

Frozen tile Nothing happens. Reward = 0.
Hole The episode ends unsuccessfully. Reward = 0.
Goal The episode ends successfully. Reward = 1.

Through trial-and-error, the agent gradually learns which actions are beneficial in a specific situation. Eventually, it will

develop a policy that allows for navigating from *Start* to *Goal* without falling into *Hole*s. By that, the policy encodes the MT sequence of interest, which can be materialized by selecting the actions (MTs) with the highest probability in each state.

The complexity of the problem prevents the agent from learning at a high pace. Typically, the agent struggles at the beginning of the learning process [4], especially when obstacles hinder exploration (e.g., the *Hole* in [1, 1], and the *Hole*s around [2, 1], forming a narrow passage).

Human guidance

Human guidance can be of significant value to accelerate the agent's learning rate. For example, a human advisor might suggest the agent to avoid [1, 1] as there is a hole (negative advice); or to visit [2, 1] or [1, 2], because those tiles are part of every viable path between *Start* and *Goal* (positive advice).

The advisor might **not be completely certain** in the advice. For example, advisors might be more certain about the vicinity of their location (situational expertise), or might be experts in recognizing *Hole*s but not *Goal*s (domain-specific expertise).

Result of the learning process

The result of the RL process is a policy that is able to assist the execution of MTs by providing priority information about which MT to execute in specific states, i.e., which direction to take while standing on specific frozen tiles. For example, the `moveRightRule` MT rule should be severely deprioritized when the vehicle is on tile [0, 1] as it would lead to a termination by falling into a hole.

IV. OVERVIEW OF THE APPROACH

A. Framework prototype

To implement and evaluate our approach, we prototyped a framework. The goal of the framework is to support the registration of MTs or model manipulation primitives from which complex MT sequences are inferred via RL. The framework is available as an open-source software.[3]

Fig. 3 presents the architectural view of the framework (in blue), in the context of the typical domain concepts, activities, and users. The core element of the framework is the *RL Engine*. The engine is fully modeled: the *RL metamodel* captures the domain-agnostic RL concepts (e.g., agent, policy, state) in the Eclipse Modeling Framework (EMF), and the operational semantics of RL (i.e., the underlying Markov Decision Process) is captured in MTs implemented on the VIATRA platform. These MTs span the *RL lifecycle activations* that drive the engine. (Described in detail in Sec. IV-D.) The *Policy* is maintained by specific activations. To resolve ambiguous MT activations (i.e., when multiple MTs can fire), we implemented a *ConflictResolver* that chooses the activation of the MT that is preferred in the Policy (i.e., has the highest probability in the state the agent exhibits).
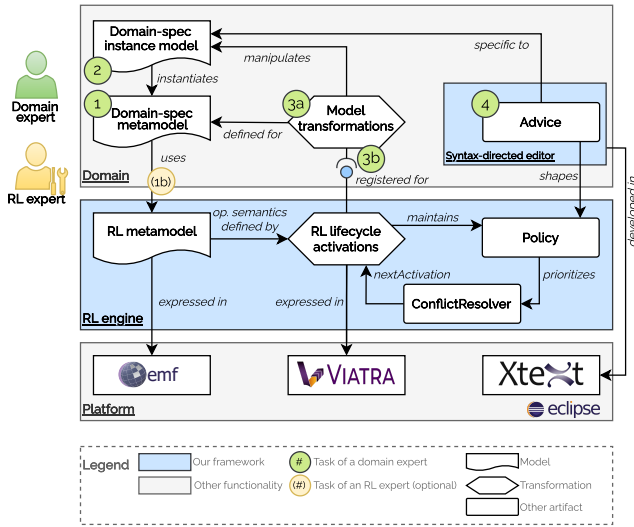
[3]https://github.com/ssm-lab/rl4mt

Fig. 3: Architectural overview

B. Domain-specific (meta)modeling

The framework allows for executing MTs as RL programs in an MDE process, with minimal development effort.

In step ①, a *Domain-specific metamodel* is developed by the *Domain expert* that reuses some key concepts of the *RL metamodel*. Specifically, this step might be aided by an *RL expert* in case the domain expert is not familiar with RL concepts. These cases are typical when a domain expert works with RL for the first time; the concepts to work with (e.g., terminating states and rewards) are trivial and could be circumvented by generative techniques.

Mapping MTs to RL: Related works have provided a clear mapping between MTs and RL [5, 6]. Here, we only recall that graph morphisms or MT left-hand sides (LHS) correspond to states in RL programs and model manipulations correspond to actions in RL. Reward is automatically calculated (e.g., number of violations in model repair [5]), or assumes a reward structure (e.g., quality of solutions in model exploration [3]).

In the running example, the domain expert may choose to refine the concept of the agent into the concept of the *Truck*, and define the tiles on the lake as RL states, with frozen tiles being non-terminating states (the agent may continue the exploration) and holes and the goal being terminating states (the agent will stop exploring).

In Step ②, the domain expert instantiates the domain metamodel, e.g., as illustrated in Fig. 1.

C. Defining MTs

In Step ③a, the domain expert defines the MTs for the domain model. In the running example, this means defining the four movements (left, down, right, up) for the autonomous vehicle. Once the MTs have been defined, they are registered in the RL engine in Step ③b to be executed as an RL program.

Typically, these MTs react to matched graph patterns expressed in terms of the domain-specific metamodel. This is why in the previous step, the domain expert (and in some

cases, the RL expert) needed to ensure that the domain-specific metamodel uses some key concepts of the RL metamodel. Alternatively, users of the framework may opt for defining plain model manipulations without an LHS, instead of complete MTs. Upon registering model manipulations, the framework packages them into MTs with a pattern in the RHS that describes non-terminating states in the RL agent's state space. This allows for the execution of model manipulations as long as the RL program does not terminate.

D. Operational semantics

We developed our RL engine in a fully modeled fashion, in which the RL metamodel is operationalized by MTs. These MTs are internal to the RL engine and different from the registered domain-specific MTs. For example, updating the policy upon encountering a terminating state is formulated as an internal MT. The activations of domain-specific and internal MTs span a lifecycle, shown in Fig. 4.
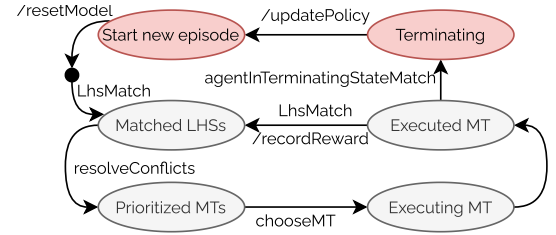

Fig. 4: RL lifecycle for learning MT sequences

The execution of the externally registered MTs starts by matching their LHS rules, resulting in some number of *Matched LHSs*. Typically, multiple MTs' LHSs are matched, which requires resolving conflicts. In our implementation platform, VIATRA, conflict resolution is achieved through a conflict resolver over a conflict set.[4] We implement our conflict resolution logic in a way that MTs are dynamically prioritized by their probability in the policy. Listing 2 shows the relevant excerpt of the `PolicyBasedConflictSet` class in our implementation, which returns the next LHS activation.

```
1 override getNextActivation() {
2    val nextActivation =
3        policyService.chooseAction(container)
4    container.clear
5    return nextActivation
6 }
```

Listing 2: Excerpt of the *PolicyBasedConflictSet* class through which the policy-based conflict resolver prioritizes MTs

After resolving conflicts, the engine has a list of *Prioritized MTs*. After choosing one of the prioritized MTs, the engine starts *Executing the MT*. After execution, the result of the *Executed MT* determines how the lifecycle proceeds.

If the agent is situated in a terminating state at this point, the lifecycle continues by *Terminating* the episode, which includes

---

[4]See `ConflictResolver.java` and `ConflictSet.java` on the API of VIATRA in https://github.com/eclipse-viatra/org.eclipse.viatra.

updating the policy and *Starting a new episode* by resetting the model (e.g., placing the agent in the initial state). This internal MT is shown in Listing 3 and Listing 4.

```
1  pattern agentInTerminatingState(
2    agent: Agent,
3    state: TerminatingState,
4    environment: Environment){
5      find agentInState(agent, state);
6      Environment.agent(environment, agent);
7  }
```

Listing 3: Domain-agnostic graph query matching morphisms in which the RL agent is in a terminating state

```
1  val agentIsTerminatedRule =
2    createRule(agentInTerminatingState)
3    .name("AgentIsTerminatedRule")
4    .action[agent.terminate].build
```

Listing 4: Domain-agnostic MT rule terminating the RL episode upon the agent encountering a terminating state

If the agent is situated in a non-terminating state, the execution continues by recording the reward for the successful step and matching LHSs in the new state.

Through this lifecycle, the RL agent learns beneficial MT sequences based on the rewards for successful (non-terminating and rewarded) MT executions. The result of the RL process is a policy that can guide subsequent executions of the MTs to form the most efficient complex MT.

V. GUIDING RL AGENTS BY OPINIONS

In this section, we elaborate on Step ④. Our previous work [11, 32] presents the general method in great detail. Here, we focus on adopting its essential elements required for this paper. We refer to Fig. 5 in the remainder of this section.

A. Advisor input

First, the advisor provides advice to the agent. This step (shown in ④.1 in Fig. 5) is the only manual step in the method; the rest of the steps are automated in our technical framework.

The advisor expresses their subjective belief about a state being beneficial to the agent. We define advice as $\alpha : s \in S' \mapsto v$, i.e., a mapping from a state $s \in S'$ to a value v. Here, $S' \subseteq S$ is the subset of the problem space that the advisor can feasibly provide advice about, and v is a value that describes the benefit of occupying or exhibiting that state. In the running example, a domain expert could formulate the advice "*Cell (1,1)* [a hole] *is likely not beneficial*". Here, s refers to cell [1,1] in the grid world, and "not beneficial" is the value of occupying cell [1,1]. This advice is used to update the probability of choosing the action in adjacent states that leads to the advised cell, thereby shaping the policy.

Advice is provided through an appropriate domain-specific language (DSL) [33] and ideally, in a *Syntax-directed editor*. To demonstrate principles and conduct our experiments, we developed a DSL and an Xtext-based editor for the running example. Such editors can be developed or generated for the specific case at hand. The DSL to provide advice about the
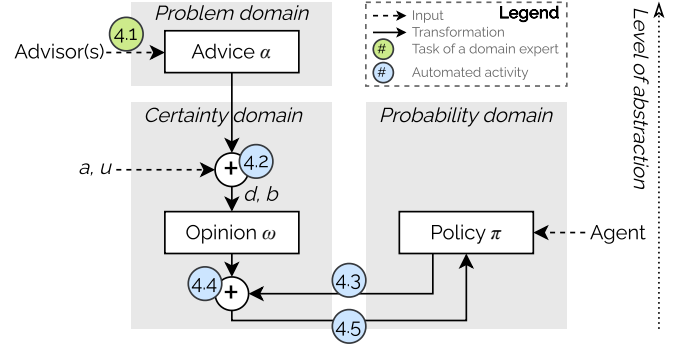

Fig. 5: Method for incorporating advice into RL

grid world of the running example is shown in Listing 5. An example set of advice expressed in the DSL about the running example (Fig. 1) is shown in Listing 6.

```
1  <AdviceList> ::= <Advice>+
2  <Advice> ::= <AdviceLocation> ',' <AdviceValue>
3  <AdviceLocation> ::= '['[0-9]+, [0-9]+']'
4  <AdviceValue> ::= ^-?[0-2]
```

Listing 5: GridWorld grammar in EBNF

```
1  [1,1], -2
2  [3,1], -2
3  [3,0], -1
4  [3,3], +2
```

Listing 6: Example set of opinions for the problem in Fig. 1

The value of the advice represents the advisor's idea about how beneficial the given cell can be for the agent. In the example in Listing 6, the list of advice corresponds to the following cells of the running example (Fig. 1): two holes on the grid, which are not beneficial ([1,1], -2) and ([3,1], -2); the goal, which is beneficial ([3,3], +2); and a frozen tile that is not hazardous but still rather disadvantageous ([3,0], -1) due to the holes in its vicinity.

Uncertainty in advice may arise due to various reasons. Typical to multi-view modeling settings, a subject-matter expert may be required to give advice about adjacent domains' concepts. For example, a mechanical engineer may need to express the quality of the electro-mechanical design.

B. Compiling advice into opinion

Next, in step ④.2, advice is transformed into opinions by mapping to the parameters of **subjective logic** [34]—briefly introduced in Sec. II-C. This step is completely automated and is achieved in three sub-steps as follows.

1) Calibrating base rate: First, we set base rate a_i, which represents the prior probability of the agent taking an action in a state. For a set of actions A, the base rate is derived as

$$a_i = 1/|A|. \tag{1}$$

In the running example, there are four actions the agent can choose in any state, thus $|A| = 4$ and the base rate $a_i = 0.25$. (Note that if the agent chooses an action that would lead it out of the grid world, it stays in place.)

2) Calibrating uncertainty: Next, we calibrate uncertainty u_i. While uncertainty can be set manually by the advisor, it can prove difficult to quantify, and thus, we recommend deriving uncertainty using an automated method. For example, uncertainty may be derived from an appropriate distance metric, such that uncertainty increases with distance. To calibrate uncertainty using a distance measure, it is essential to choose (i) an appropriate distance metric and (ii) a discount function that successively discounts uncertainty as distance increases.

In general the discount function is given as $\gamma : (u_{max}, \delta) \mapsto u_i \in (0,1)$, where δ is a distance measured by a distance measure Δ, and u_{max} is the upper bound of uncertainty. From the identity of $u_i = 1 - (b_i + d_i)$ (Sec. II), it follows that $u_{max} = 1 - (b_i + d_i)$. In the simplest case, a linear discount function can be chosen to calibrate uncertainty as follows.

$$u_i = (\delta/\delta_{max}) \times u_{max} \tag{2}$$

Here, δ/δ_{max} is the distance relative to the maximum distance, that maps onto the (0,1) domain.

In problems where uncertainty should be scaled only to a subset of the problem space, we use a threshold $0 < \tau \in \Delta < 1$ to accelerate the discounting of certainty and reach u_{max} in τ. Then, Equation 2 is adopted as follows.

$$u_i = \begin{cases} 1/\tau \times \delta/\delta_{max} \times u_{max} & \text{if } \delta \le (\tau \times \delta_{max}); \\ u_{max} & \text{if } \delta > (\tau \times \delta_{max}). \end{cases} \tag{3}$$

3) Computing disbelief (d) and belief (b) to form an opinion: Finally, the belief b_i and disbelief d_i can be computed, and the parameters of the opinion ω_i can be set. As per the parameter constraints explained in Sec. II-C, $u_i + b_i + d_i = 1$, and $0 \le u_i, b_i, d_i \le 1$, so the remaining weight of $1 - u_i$ is distributed between b_i and d_i based on the value of the advice α. Given n, the length of the advice value scale, and the value of the advice j (in ascending order of confidence from least to most), the calculation of b_i and d_i is given as

$$b_i = (j-1/n-1) \times (1 - u_i) \mid j \in \{1..n\} \tag{4}$$
$$d_i = (1 - u_i) - b_i. \tag{5}$$

By using Equations 1–5, the opinion ω_i is compiled as

$$\omega_i = \left(\frac{j-1}{n-1} \times (1 - u_i), \frac{n-j}{n-1} \times (1 - u_i), u_i, \frac{1}{|A|} \right). \tag{6}$$

The first parameter is the belief, b_i as formulated in Equation 4, and the second parameter is the disbelief d_i as formulated in Equation 5. The third parameter u_i is the uncertainty calculated automatically by an appropriate metric. The final parameter a_i is the base rate calculated by the dynamics of the problem space, as described in Equation 1.

Consider the advice from line 4 of Listing 6. Using an advice scale with length $n = 5$, the advisor indicates that

the goal is beneficial $\alpha = ([3,3], 2)$. Advice value $v = 2$ corresponds with $j = 5$. The agent can take one of the four actions in each cell in Fig. 1, so $a_i = 1/4 = 0.25$. This advisor has limited understanding of the environment, and is assigned uncertainty $u_i = 0.5$. Using Equation 6, we calculate ω_i as

$$\begin{aligned} \omega_i &= \left(\frac{j-1}{n-1} \times (1 - u_i), \frac{n-j}{n-1} \times (1 - u_i), u_i, \frac{1}{|A|} \right) \\ &= \left(\frac{5-1}{5-1} \times (1 - 0.5), \frac{5-5}{5-1} \times (1 - 0.5), 0.5, \frac{1}{|4|} \right) \\ &= (0.5, 0, 0.5, 0.25). \end{aligned}$$

C. Policy shaping via opinion fusion

Next, in the fully automated steps 4.3 and 4.4, the agent's policy is shaped by the advisor's opinions as follows.

1) Transforming Agent's Policy into certainty domain: As described in Sec. II-A, the agent maintains a policy $\pi(a|s)$, which gives the conditional probability of choosing action $a \in A$ while in state $s \in S$. Since the entries are probabilities, $\forall s \in S, a \in A : 0 \le \pi(a|s) \le 1$. Before exploration, the probability of choosing any action from a state is equal, $\forall s \in S, a_i, a_j \in A, i \ne j : \pi(a_i|s) = \pi(a_j|s)$.

To fuse opinions described in the certainty domain, with a policy described in the probability domain, we choose to transform the agent's policy to the certainty domain (step 4.3). Our choice is motivated by the well-defined and sound fusion operators that are readily available in the toolbox of subjective logic [34]. This translation is achieved as follows: $f_{P \to \Omega}(\pi) : S \times A \to \Omega \mid P$. As explained in Sec. II-C, probability p translates to opinion $\omega : p \mapsto (p, 1 - p, 0, p)$.

2) Policy shaping by opinion fusion: Next, in step 4.4, the advisor's opinions are fused with the agent's policy in the certainty domain. As described in Sec. II-C, combined opinions can be fused into one joint opinion. By fusion, we are referring to a mapping $\Omega \times \Omega \to \Omega$. For fusion, we require two elements: locating the policy entries that are affected by the advice (in the running example: the neighboring states of the advised state), and choosing an appropriate fusion operator from the toolbox of subjective logic.

a) Locating policy entries to be fused: Since the advisor gives advice about the value of a particular state being beneficial, this advice will affect the agent's policy in states from which the advised state is reachable. The following definitions are required to formalize this step.

Definition 1 (Neighboring states): States $s_i, s_j \in S$ are said to be neighbors if there exists a state-action pair in the policy for which $\pi(a|s_i) \mapsto s_j$, i.e., choosing action $a \in A$ in state s_i will transition the agent to state s_j.

Definition 2 (Neighborhood (of a state)): $N(s_i \in S) = \bigcup_{s_j \in S} \exists a \in A : \pi(a|s_i) \mapsto s_j$.

That is, neighborhood N of a state is the set of all neighbors. In the running example, the state space is topological; thus, the neighborhood of a cell comprises the cells that can be reached from the given cell through one of the agent's actions.

Advice α about state s_i is introduced to policy π by $\forall a \in N(s_i).A : \omega(\alpha) \odot \omega(a)$. That is, opinion $\omega(\alpha)$ formed from

the advice is fused with every opinion formed from actions ($\omega(a)$) that lead from its neighboring states $N(s_i)$ to s_i.

b) Choosing a fusion operator: To fuse opinions, Jøsang [34] defines several fusion operators that are classified based on the situation. In this work, the agent and advisor formulate their opinions separately, and it is fair to assume they will not compromise. Thus, we use the Belief Constraint Fusion (BCF) operator, as it is most appropriate when compromise is not sought, and those giving opinions have made up their minds. Joint opinion $\omega^\circ = (b^\circ, d^\circ, u^\circ, a^\circ)$ under BCF is calculated by first finding the harmony and conflict of opinions: *Harmony* $= b_1 u_2 + b_2 u_1 + b_1 b_2$; *Conflict* $= b_1 d_2 + b_2 d_1$. Then, the components of the joint opinion are calculated as follows.

$$b^\circ = \frac{Harmony}{(1 - Conflict)}; \qquad u^\circ = \frac{u_1 u_2}{(1 - Conflict)};$$
$$d^\circ = 1 - (b^\circ + u^\circ); \qquad a^\circ = \frac{a_1(1 - u_1) + a_2(1 - u_2)}{2 - u_1 - u_2}.$$

As described in Sec. II-C, more than two opinions may be fused into a joint opinion. Due to this and the fact that fusion operators are commutative, our approach does not limit the number of advisors that may provide advice to the agent.

D. Transforming opinion-infused policies from the certainty domain to the probability domain

Finally, in automated step ④.5, the agent's policy—now shaped by advisor opinions—is transformed from the certainty domain back to the probability domain. First, as per Sec. II-C, every opinion is transformed into a probability as $p : \omega \mapsto b + au$. Second, the policy needs to be normalized. State-specific policy entries form a probability space, and thus the invariant $\forall s \in S : \sum_{a \in A} p(a|s) = 1$ must hold. Normalization is defined as $\forall s \in S, a \in A : p(a|s) := p(a|s) \times \mathbf{1}/\sum_{a \in A} p(a|s)$.

The agent's policy now reflects the advice, and the agent may begin the exploration process.

VI. EVALUATION

In this section, we evaluate our approach by answering the following research question: **how does human guidance affect the performance of RL in inferring MTs?**

A. Evaluation setup

We conduct a comparative study in which, first, we drive MTs by an unadvised RL agent, and then, we provide advice at various levels of uncertainty. We use the 12×12 scaled-up, randomized version of the Frozen Lake environment (Fig. 1). The map contains 20% *Hole* tiles, which is a standard ratio in OpenAI Gym's Frozen Lake implementation. The exact map is visualized in the replication package. We assess the performance of RL agents by the cumulative reward they collect. We measure both the amount of the cumulative reward and the pace at which the RL agent collects it. Our hypothesis is that human guidance has a measurable positive impact on both aspects, even at moderate levels of uncertainty.

TABLE I: RL hyperparameters and settings

Parameter	Value
RL method	Discrete policy gradient
Learning rate (α)	0.9
Discount factor (γ)	1.0
Number of episodes	10 000
Maximum steps per episode	100
State-action space	$12^2 \times 4$

TABLE II: Experiment parameters

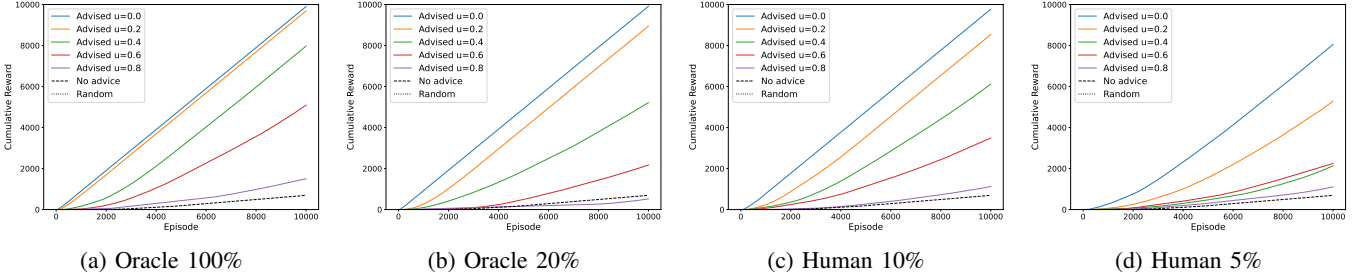| Parameter	Value
Agent type	{Random, Unadvised, Advised}
Source of advice	{Oracle, Single human}
Fusion operator	BCF
Advice quota – Oracle	{100%, 20%}
Advice quota – Single human	{10%, 5%}
Uncertainty (oracle and single human)	{$0.2k \mid k \in 0..4$}

RL-specific parameters: The RL-specific parameters of our experiments are shown in Tab. I. For the RL method, we choose discrete policy gradient [4] as policy-based methods explicitly represent the policy and allow for the direct investigation of the effects of advice. *Learning rate* α affects how much the agent updates its policy parameters during training. The closer it is to 0, the smaller the policy parameter updates, and the closer it is to 1, the larger the policy parameter updates. *Discount factor* γ determines how much the agent values future rewards. The closer γ is to 0, the more the agent prioritizes immediate rewards, and the closer γ is to 1, the more the agent values long-term rewards. Each experiment runs for a *number of episodes*. An episode starts by resetting the environment and placing the agent in the initial position. Within the same experiment, the policy is continuously maintained. An episode lasts until the agent either finds a terminating state (hole or goal) or reaches the number of *maximum steps per episode*. To mitigate threats to validity, we tuned these hyperparameters to favor the learning dynamics of the unadvised agent.

Experiment parameters: Experiment parameters are shown in Tab. II. We conduct 30 experiments for combinations of parameters and calculate their average to identify systemic tendencies. We compare the performance of different *agent types*. The random agent chooses available actions randomly. The advised agent maintains and updates a policy. The unadvised agents also maintain and update a policy, and before exploration advice is incorporated into the policy.

We guide advised agents by different *sources of advice* and *advice quotas*. To form a ground truth, we first experiment with an idealized oracle with full information of the problem space, which provides advice based on pre-programmed rules. The oracle has two quotas; one where it provides advice about *all* states (100%) and another where it provides advice about the *holes and goal* on the map (20%).

In the *single human mode*, one of the authors acted as the advisor with full information of the problem space. This mode has two advice quotas: one where advice is provided about 10% of the state space, and the other where advice is

Linear scale

(a) Oracle 100% (b) Oracle 20% (c) Human 10% (d) Human 5%

Log scale

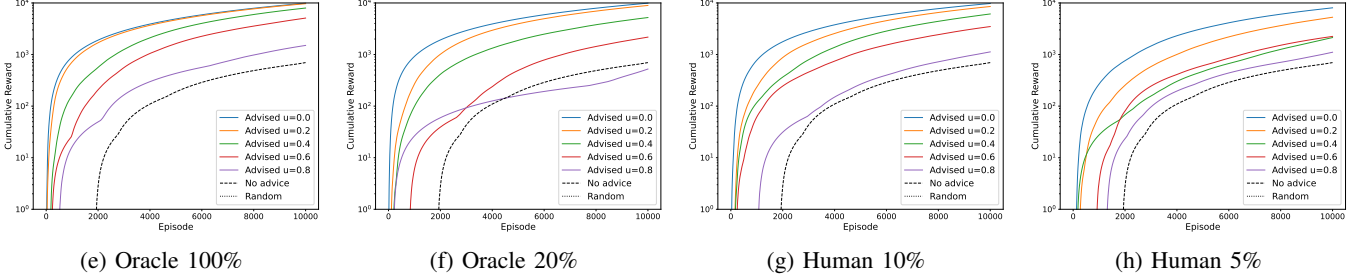(e) Oracle 100% (f) Oracle 20% (g) Human 10% (h) Human 5%

Fig. 6: Performance comparison of synthetic oracles and real human advisors in terms of cumulative reward

TABLE III: Cumulative rewards: Oracle and Single Human
(Random = 0.033, Unadvised = 698.266.)

	Oracle		Single human	
u	100% (All)	20% (H&G)	10%	5%
0.0	9 900.100	9 914.900	9 768.000	8 051.300
0.2	9 685.900	8 948.933	8 538.266	5 287.833
0.4	7 974.066	5 216.433	6 121.033	2 134.966
0.6	5 094.333	2 177.633	3 488.700	2 246.733
0.8	1 502.500	523.633	1 126.300	1 108.666

provided about 5% of the state space. For both the *oracle* and *single human* advisors, the *degree of uncertainty* is modulated synthetically by setting $u = 0.0, 0.2, 0.4, 0.6, 0.8$.

Note about additional experiments: The replication package contains additional data from experiments with two cooperating humans guiding the RL agent. Due to space restrictions, those results are not presented, but we note that **agents advised by cooperating humans outperform the unadvised agent** and **cooperating humans with lower quotas obtain similar performance to a single human advisor**.

B. Results and key takeaways

The results are shown in Tab. III and are visualized in Fig. 6. In addition, we evaluate statistical significance between each advice mode at different levels of uncertainty. We use independent samples t-tests for the pairwise comparison. The replication package contains the complete list of t-test results.

We observe that **advised agents perform better than unadvised ones**, as in all but one case, advised agents accumulate more reward than the unadvised agent (mean reward 698.266). We also note that higher advice quotas tend to result in better performance. The main takeaway of this result is that **guidance, even if uncertain, improves the performance of the RL agent learning complex MTs**.

We observe that the oracle with 100% quota has the best overall performance, accumulating the most cumulative reward for all but one $u = 0.0$, where the oracle with 20% quota is slightly better. The **oracles with 100% and 20% quota, and the human advisor with 10% have similar performance** for low uncertainty $u = 0.0$ and $u = 0.2$. This is indicated by the cumulative reward and the similar slope in the log plots. While the oracle with 100% advice quota (Fig. 6a) significantly outperforms the human advisor with 10% quota (Fig. 6c) at low levels of uncertainty ($u \leq 0.4$), this difference is not significant at higher levels of uncertainty ($u \geq 0.6$). That is, **in uncertain situations, a human advisor, even with a lower advice quota, does not perform systematically worse than a hypothetical oracle**. This finding highlights the utility of human advice in uncertain situations. This trend is even more pronounced when the oracle is given less advice quota (Fig. 6b). With 20% quota, the oracle does not significantly outperform a human advisor with 10% quota when $u > 0.0$.

At low levels of uncertainty ($u \leq 0.4$) a human advisor with 10% quota significantly outperforms the human advisor with 5% quota. However, at higher levels of uncertainty ($u \geq 0.6$), the difference is not significant. These findings suggest that **different uncertainty levels require different advice strategies from human advisors**. Specifically, when advisors are certain, more advice is justified; however, when uncertain, advisors should be more frugal with their advice.

At low levels of uncertainty ($u \leq 0.4$), all advising modes (Fig. 6a–6d) significantly outperform the unadvised agent. However, as uncertainty increases ($u \geq 0.6$), the differences between advised and unadvised agents become statistically insignificant. These findings suggest that **guidance is beneficial with low-to-moderate uncertainty, but its utility is discounted as uncertainty increases**.

> Our observations indicate that human guidance contributes to significant performance improvement in RL-based inference of complex MTs, even at moderate-to-high uncertainty. Human advisors, even with substantially lower advice quotas, can perform at the level of an oracle, and might even outperform it. This results in faster and more efficient inference of complex MTs.

C. Assumptions, limitations, and threats to validity

1) Assumptions and limitations: We implicitly assumed that the reward structure could be effectively formulated. In some MDE applications, this is true. For example, Barriga *et al.* [9] formulate reward by a user-preferred quality characteristic, such as maintainability and understandability. Still, formulating reward might be challenging in open problems, such as design space exploration, in which reward often must be asserted to hypothetical, yet-to-encounter model configurations.

2) Threats to validity: A fortunately aligned reward structure may result in better, albeit non-systemic performance of the RL agent, threatening the *internal validity* of our study. To mitigate this threat, (i) we have chosen the least supportive reward function in which the goal reward is minimal, there are no intermediate rewards, and terminating states do not provide negative feedback; and (ii) tuned the hyperparameters to favor the learning dynamics of the unadvised agent. To assess the *external validity* of our study, we remark that we only tested policy-based RL. While it is reasonable to expect the method to be applicable to other RL methods (e.g., value-based RL, a popular choice in MDE [6, 9, 35]), our study does not provide evidence for this. Similarly, a minor threat to external validity stems from the technological choices, especially the MT platform. We rely on the reactive capabilities of VIATRA, and it is reasonable to expect that MT frameworks with support for similar MT execution semantics can adopt our approach.

VII. DISCUSSION

We now discuss the implications of our work and outline areas of interest for MDE researchers.

A. Towards RL-enabled MT platforms

The result of the RL process is a policy that encodes beneficial MT executions in specific graph morphisms. To utilize this information, we foresee two main directions.

1) Runtime policy exploitation in MT chains: The RL engine can be used as a facility to resolve MT activations based on the policy. The policy asserts higher probability to specific MTs in a model state due to the higher positive reward the MT yielded in previous situations of exhibiting the model state. Such scenarios can be supported through appropriate conflict resolves that factor in the policy information, as shown in Fig. 3 and Fig. 4. The benefit of this approach is that the RL engine can continuously maintain the complex MT and guide the execution of complex MTs in a non-intrusive way.

2) Generating standalone MT chains: Alternatively, standalone MT chains can be generated to be used independently from the RL engine. In such scenarios, the RL engine is used only for training purposes and does not maintain the MT chain.

We anticipate MT platform developers to adopt techniques such as ours and promote RL-enabled MTs to a first-class citizens in their platforms. While the former approach leverages our approach to a greater extent, the latter one is less intrusive and requires a smaller leap for the user base of MT platforms.

B. Typical applications

To better contextualize our approach within MDE, we explain how some characteristic MDE problems map onto RL and can make use of our approach.

1) Model synchronization and inconsistency management: These problems are related and both often require identifying complex MTs [36–38]. In these settings, the RL **state** encapsulates morphisms over the union of the source and target models, and the **reward** is typically implied by the negative distance between the source and target models. The lower the distance, the better the solution. Termination semantics are asserted to states that minimize the distance or, equivalently, maximize the reward. Distance can be measured at the syntactic level (e.g., by counting inconsistent model elements) [39] or at the semantic level (e.g., by counting inconsistent properties that are not being jointly satisfied by the two models) [40].

The role of our approach: **Actions** are typically expressed as domain-specific MTs, but are limited in complexity. Our approach allows for inferring complex MTs that allow for more sophisticated synchronization semantics between models. **Human guidance** is of particular value in these problems as multi-domain settings challenge the understanding of how exactly consistency should be restored [24, 41].

2) Model repair and refactoring: This class of problems is similar to model synchronization; however, there are some key differences. RL **state** in this case encapsulates only a single model's morphisms, and **reward** cannot be automatically calculated due to the lack of a suitable oracle. That is, it cannot be structurally defined when a refactoring is successful or sufficient. Instead, the reward structure must be defined by the domain expert. Typically, qualitative metrics are used for this purpose, either defined at the general syntactic level (e.g., number of dangling edges or isolated components) [42] or at the domain-specific syntactic level [30]. Occasionally, personalized metrics are considered [9].

The role of our approach: The challenge is similar to the previous class of problems: **actions** are usually expressed as domain-specific MTs with limited complexity. Here again, an approach such as ours would allow for inferring complex MTs that, in turn, enable more complex refactorings and model

repair actions. **Human guidance** helps prevent models from growing out of size under the premise of repair.

3) Design space exploration: Design space exploration (DSE) is the process of identifying favorable candidate designs under various constraints and soft objectives [43]. Starting from an initial design, a DSE process gradually applies MTs to uncover alternative designs. In such scenarios, the **state** encapsulates a specific graph morphism in a model instance. **Actions** are typically defined in domain-specific terms [10] rather than domain-agnostic CRUD primitives. The **reward** structure is given through constraints and soft objectives.

a) The role of our approach: Optimal solutions are often found deep in the design space and are computationally demanding to identify. Our approach can aid the computational load of DSE by learning complex MTs that have led to beneficial design alternatives. These complex MTs, then, can be used as fixed points in the design space from which subsequent DSE executions may continue deeper explorations. **Human guidance** helps by providing known favorable design decisions and by that, drive the RL agent to favorable states and accelerate the inference of complex MTs. The challenge of uncertainty in DSE has been recognized before [44]. Our approach allows for the explicit modeling and management of uncertainty that stems from human guidance.

C. Research opportunities for the MDE community

There are some key challenges to be addressed in order to materialize the opportunities our approach creates. In the following, we elaborate on the ones that align particularly well with the expertise and knowledge of the MDE community.

1) Domain-specific languages for guidance: A key advantage of opinion-based guidance is the rapid pace at which opinions emerge. To leverage this advantage to its fullest, intuitive languages are needed to express opinions. Domain-specific languages (DLSs) [33] have been widely used to capture expert opinions at high levels of abstraction, narrowing the gap between the modeling language and human cognition. However, DSLs typically operate with reduced syntax to remain truly specific to the problem domain, and therefore, **automated methods for DSL engineering** are in high demand [45]. This is true for guided RL as well. We envision language workbenches that allow for the automated construction of DSLs through which advice can be provided by domain concepts. In the running example (Sec. III), one might want to use descriptive quantifiers, such as "hazardous" or "beneficial" instead of relying on an integer scale. Unfortunately, there is little research on **domain-specific languages for RL**, and existing languages mostly focus on programming RL procedures [46] rather than expressing opinions at high cognitive levels. We see plenty of room for research targeting domain-specific modeling languages for RL, inspired by the advancements in the model-driven engineering community [47, 48].

2) (Domain-specific) Tool support: Our approach can be aided by domain-specific tool support at key points. For example, the domain expert needs a **domain-specific editor** with an intuitive user interface to efficiently provide advice. We implemented the prototype of such an editor on top of Eclipse and Xtext, as discussed in Sec. V-A. Instead of case-by-case tool development, syntax-directed editors could be generated from high-level specifications of the RL problem and the application domain [49], and equipped with automatically generated specialized concrete syntax [50]. Adaptive modeling languages [51] offer an apt solution through a product line-inspired approach to customizing DSLs to different problems.

Additional tooling considerations relate to **model management** and RL execution. In our work, we executed the RL agent's learning process on an instance model that should likely be prevented from automated complex MTs until the RL agent is trained. Instead, an improved version of our tool could create replicas of the real model and treat those as sandbox models for the agent to train. The result of the training, i.e., the policy, can be applied on the real model. Such utility functionality will pave the way for adoption.

3) Interactive advice: In our experiments, we assumed that advice is provided once, before the agent begins exploring the environment. However, given adequate interaction protocols, our approach can accommodate advice at any point during the RL process. We see opportunities in developing more **interactive advice protocols** through advanced human-computer interactions (HCI) [52]. MDE can be of high value in such situations, as it can support **domain-specific interfaces** that integrate user models in the system design. This would allow for interactive advising systems that are tailored to the user's level of expertise about the problem space [53]. Blended modeling—the activity of interacting seamlessly with a single model (i.e., abstract syntax) through multiple notations (i.e., concrete syntaxes), allowing a certain degree of temporary inconsistencies [54]—can be of high utility in facilitating seamless advising workflows in multi-view settings.

Finally, interaction protocols amenable to **verification and validation** can be generated, e.g., from state chart models [55] or sequence diagrams [56].

VIII. Conclusion

We presented a method for engineering complex model transformation (MT) sequences by reinforcement learning (RL), guided by potentially uncertain human advice. Our results show that human advice, even if uncertain, can be a major contributor to more efficient MT development. Our results open the door for the RL-based human-in-the-loop treatment of key MDE problems, e.g., model synchronization and design-space exploration. We envision MT platform developers adopting techniques such as ours to promote RL-backed MTs to first-class citizens in MT platforms. This will enable efficient MT development, execution, and maintenance.

Future work will focus on evaluating our approach for different flavors of RL and developing domain-specific support for RL-driven human-in-the-loop MT engineering.

References

[1] K. Marussy, O. Semeráth, and D. Varró, "Incremental view model synchronization using partial models," in *Proceedings of the 21th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems*, ser. MODELS '18, Copenhagen, Denmark: ACM, 2018, pp. 323–333. DOI: 10.1145/3239372.3239412.

[2] C. E. Mokaddem, H. Sahraoui, and E. Syriani, "Recommending model refactoring rules from refactoring examples," in *Proc. of the 21th ACM/IEEE Intl. Conference on Model Driven Engineering Languages and Systems*, ser. MODELS '18, Copenhagen, Denmark: ACM, 2018, pp. 257–266. DOI: 10.1145/3239372.3239406.

[3] H. Abdeen, D. Varró, H. Sahraoui, A. S. Nagy, C. Debreceni, Á. Hegedüs, and Á. Horváth, "Multi-objective optimization in rule-based design space exploration," in *Proc. of the 29th ACM/IEEE Intl Conference on Automated Software Engineering*, ACM, 2014, pp. 289–300. DOI: 10.1145/2642937.2643005.

[4] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[5] A. Barriga, A. Rutle, and R. Heldal, "Personalized and automatic model repairing using reinforcement learning," in *2019 ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*, 2019, pp. 175–181. DOI: 10.1109/MODELS-C.2019.00030.

[6] M. Eisenberg, H.-P. Pichler, A. Garmendia, and M. Wimmer, "Towards reinforcement learning for in-place model transformations," in *2021 ACM/IEEE 24th Intl Conf. on Model Driven Engineering Languages and Systems (MODELS)*, 2021, pp. 82–88.

[7] M. Dehghani, S. Kolahdouz-Rahimi, M. Tisi, and D. Tamzalit, "Facilitating the migration to the microservice architecture via model-driven reverse engineering and reinforcement learning," *Soft. Sys. Mod.*, vol. 21, no. 3, pp. 1115–1133, 2022. DOI: 10.1007/s10270-022-00977-3.

[8] A. Najar and M. Chetouani, "Reinforcement learning with human advice: A survey," *Frontiers in Robotics and AI*, vol. 8, 2021, ISSN: 2296-9144. DOI: 10.3389/frobt.2021.584075.

[9] A. Barriga, R. Heldal, A. Rutle, and L. Iovino, "Parmorel: A framework for customizable model repair," *Soft. Sys. Mod.*, vol. 21, no. 5, pp. 1739–1762, 2022.

[10] Á. Hegedüs, Á. Horváth, and D. Varró, "A model-driven framework for guided design space exploration," *Autom Softw Eng*, vol. 22, no. 3, pp. 399–436, 2015. DOI: 10.1007/s10515-014-0163-1.

[11] K. Dagenais, "Towards model repair by human opinion–guided reinforcement learning," in *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems*, ser. MODELS Companion '24, Linz, Austria: ACM, 2024, pp. 192–195. DOI: 10.1145/3652620.3676878.

[12] M. L. Puterman, "Markov decision processes," *Handbooks in operations research and management science*, vol. 2, pp. 331–434, 1990.

[13] M. Moradi, B. J. Oakes, M. Saraoglu, A. Morozov, K. Janschek, and J. Denil, "Exploring fault parameter space using reinforcement learning-based fault injection," in *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*, 2020, pp. 102–109. DOI: 10.1109/DSN-W50199.2020.00028.

[14] I. David and E. Syriani, "DEVS model construction as a reinforcement learning problem," in *2022 Annual Modeling and Simulation Conference (ANNSIM)*, 2022, pp. 30–41. DOI: 10.23919/ANNSIM55834.2022.9859369.

[15] F. Cruz, P. Wüppen, S. Magg, A. Fazrie, and S. Wermter, "Agent-advising approaches in an interactive reinforcement learning scenario," in *2017 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, 2017, pp. 209–214. DOI: 10.1109/DEVLRN.2017.8329809.

[16] A. Najar, O. Sigaud, and M. Chetouani, "Training a robot with evaluative feedback and unlabeled guidance signals," in *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2016, pp. 261–266. DOI: 10.1109/ROMAN.2016.7745140.

[17] A. C. Tenorio-Gonzalez, E. F. Morales, and L. Villaseñor-Pineda, "Dynamic reward shaping: Training a robot by voice," in *Advances in Artificial Intelligence – IBERAMIA 2010*, Springer, 2010, pp. 483–492, ISBN: 978-3-642-16952-6.

[18] A. L. Thomaz and C. Breazeal, "Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance," in *Proc of the 21st National Conf on Artificial Intelligence - Volume 1*, ser. AAAI'06, Boston, Massachusetts: AAAI Press, 2006, pp. 1000–1005, ISBN: 9781577352815.

[19] S. Griffith, K. Subramanian, J. Scholz, C. L. Isbell, and A. L. Thomaz, "Policy shaping: Integrating human feedback with reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 26, 2013.

[20] A. Jøsang, *Subjective logic*. Springer, 2016, vol. 3.

[21] L. Burgueño, M. F. Bertoa, N. Moreno, and A. Vallecillo, "Expressing confidence in models and in model transformation elements," in *Proc of the 21th ACM/IEEE Intl Conference on Model Driven Engineering Languages and Systems*, ACM, 2018, pp. 57–66.

[22] E. Bagheri and A. A. Ghorbani, "A belief-theoretic framework for the collaborative development and integration of para-consistent conceptual models," *J. Sys. & Soft.*, vol. 82, no. 4, pp. 707–729, 2009, ISSN: 0164-1212. DOI: https://doi.org/10.1016/j.jss.2008.10.012.

[23] J. Troya, N. Moreno, M. F. Bertoa, and A. Vallecillo, "Uncertainty representation in software models: A survey," *Soft. Sys. Mod.*, vol. 20, no. 4, pp. 1183–1213, 2021, ISSN: 1619-1374. DOI: 10.1007/s10270-020-00842-1.

[24] R. Jongeling and A. Vallecillo, "Uncertainty-aware consistency checking in industrial settings," in *ACM/IEEE International Conference on Model Driven Engineering Languages and Systems*, 2023, pp. 73–83.

[25] L. Iovino, B. Angela, R. Adrian, and H. Rogardt, "Model repair with quality-based reinforcement learning.," *JOT*, vol. 19, no. 2, 17:1, 2020. DOI: 10.5381/jot.2020.19.2.a17.

[26] M. Eisenberg and M. Wimmer, "From single-objective to multi-objective reinforcement learning-based model transformation," *Soft. Sys. Mod.*, 2024, ISSN: 1619-1374. DOI: 10.1007/s10270-024-01233-6.

[27] F. Basciani, D. Di Ruscio, M. D'Emidio, D. Frigioni, A. Pierantonio, and L. Iovino, "A tool for automatically selecting optimal model transformation chains," in *Proceedings of the 21st ACM/IEEE Intl. Conf. on Model Driven Engineering Languages and Systems: Companion Proceedings*, ser. MODELS '18, Copenhagen, Denmark: ACM, 2018, pp. 2–6. DOI: 10.1145/3270112.3270123.

[28] M. Eisenberg, A. Sahay, D. Di Ruscio, L. Iovino, M. Wimmer, and A. Pierantonio, "Multi-objective model transformation chain exploration with MOMoT," *Inf Softw Technol*, vol. 174, p. 107 500, 2024, ISSN: 0950-5849.

[29] R. Heckel, "Graph transformation in a nutshell," *Electronic Notes in Theoretical Computer Science*, vol. 148, no. 1, pp. 187–198, 2006, ISSN: 1571-0661. DOI: https://doi.org/10.1016/j.entcs.2005.12.018.

[30] E. Syriani, R. Bill, and M. Wimmer, "Domain-specific model distance measures," *Journal of Object Technology*, vol. 18, no. 3, 3:1–19, 2019, The 12th International Conference on Model Transformations, ISSN: 1660-1769. DOI: 10.5381/jot.2019.18.3.a3.

[31] G. Bergmann, I. David, Á. Hegedüs, Á. Horváth, I. Ráth, Z. Ujhelyi, and D. Varró, "VIATRA 3: A reactive model transformation platform," in *Theory and Practice of Model Transformations - 8th International Conference, ICMTSTAF 2015, L'Aquila, Italy. Proceedings*, ser. LNCS, vol. 9152, Springer, 2015, pp. 101–110. DOI: 10.1007/978-3-319-21155-8_8.

[32] K. Dagenais and I. David, "Opinion-guided reinforcement learning," Tech. Rep., 2024. arXiv: 2405.17287 [cs.LG].

[33] D. C. Schmidt, "Model-driven engineering," *Computer-IEEE Computer Society*, vol. 39, no. 2, p. 25, 2006.

[34] A. Jøsang, *Subjective Logic*. Springer International Publishing, 2016, ISBN: 9783319423371. DOI: 10.1007/978-3-319-42337-1.

[35] M. Sharbaf, B. Zamani, and G. Sunyé, "Automatic resolution of model merging conflicts using quality-based reinforcement learning," *Journal of Computer Languages*, vol. 71, p. 101 123, 2022. DOI: https://doi.org/10.1016/j.cola.2022.101123.

[36] T. Mens and R. Van Der Straeten, "Incremental resolution of model inconsistencies," in *Recent Trends in Algebraic Development Techniques*, Springer, 2007, pp. 111–126, ISBN: 978-3-540-71998-4.

[37] M. El Hamlaoui, S. Bennani, M. Nassar, S. Ebersold, and B. Coulette, "A mde approach for heterogeneous models consistency," in *Proc of the 13th Intl Conference on Evaluation of Novel Approaches to Software Engineering*, ser. ENASE 2018, Funchal, Madeira, Portugal: SCITEPRESS, 2018, pp. 180–191. DOI: 10.5220/0006774101800191.

[38] I. David, J. Denil, K. Gadeyne, and H. Vangheluwe, "Engineering process transformation to manage (in)consistency," in *Proceedings of the 1st International Workshop on Collaborative Modelling in MDE (COMMitMDE 2016) co-located with ACM/IEEE 19th International Conference on Model Driven Engineering Languages and Systems (MoDELS 2016), St. Malo, France, October 4, 2016*, ser. CEUR Workshop Proceedings, vol. 1717, CEUR-WS.org, 2016, pp. 7–16. [Online]. Available: http://ceur-ws.org/Vol-1717/paper5.pdf.

[39] L. Marchezan, M. Homolka, A. Blokhin, W. K. G. Assunção, E. Herac, and A. Egyed, "A tool for collaborative consistency checking during modeling," in *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems*, ser. MODELS Companion '24, Linz, Austria: ACM, 2024, pp. 655–659, ISBN: 9798400706226. DOI: 10.1145/3652620.3688558.

[40] I. David, H. Vangheluwe, and E. Syriani, "Model consistency as a heuristic for eventual correctness," *Journal of Computer Languages*, vol. 76, p. 101 223, 2023, ISSN: 2590-1184. DOI: https://doi.org/10.1016/j.cola.2023.101223.

[41] K. Dagenais and I. David, "Driving requirements evolution by engineers' opinions," in *ACM/IEEE International Conference on Model Driven Engineering Languages and Systems Companion, MODELS-C*, ACM, 2024. DOI: 10.1145/3652620.3688566.

[42] G. Taentzer, M. Ohrndorf, Y. Lamo, and A. Rutle, "Change-preserving model repair," in *Fundamental Approaches to Software Engineering*, Springer, 2017, pp. 283–299, ISBN: 978-3-662-54494-5.

[43] S. J. I. Herzig, S. Mandutianu, H. Kim, S. Hernandez, and T. Imken, "Model-transformation-based computational design synthesis for mission architecture optimization," in *2017 IEEE Aerospace Conference*, 2017, pp. 1–15. DOI: 10.1109/AERO.2017.7943953.

[44] M. Földiák, "Probabilistic graph queries for design space exploration under uncertainty," in *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems*, ser. MODELS Companion '24, Linz, Austria: ACM, 2024, pp. 142–148, ISBN: 9798400706226. DOI: 10.1145/3652620.3688199.

[45] A. Bucchiarone, J. Cabot, R. F. Paige, and A. Pierantonio, "Grand challenges in model-driven engineering: An analysis of the state of the research," *Soft. Sys. Mod.*, vol. 19, no. 1, pp. 5–13, 2020, ISSN: 1619-1374. DOI: 10.1007/s10270-019-00773-6.

[46] T. Molderez, B. Oeyen, C. De Roover, and W. De Meuter, "Marlon: A domain-specific language for multi-agent reinforcement learning on networks," in *Proc of the 34th ACM/SIGAPP Symposium on Applied Computing*, ACM, 2019, pp. 1322–1329. DOI: 10.1145/3297280.3297413.

[47] G. Kulagin, I. Ermakov, and L. Lyadova, "Ontology-based development of domain-specific languages via customizing base language," in *2022 IEEE 16th International Conference on Application of Information and Communication Technologies (AICT)*, 2022, pp. 1–6. DOI: 10.1109/AICT55583.2022.10013619.

[48] H. Wu and J. Gray, "Automated generation of testing tools for domain-specific languages," in *Proceedings of the 20th IEEE/ACM International Conference on Automated Software Engineering*, ACM, 2005, pp. 436–439. DOI: 10.1145/1101908.1101993.

[49] E. Syriani, D. Riegelhaupt, B. Barroca, and I. David, "Generation of custom textual model editors," *Modelling*, vol. 2, no. 4, pp. 609–625, 2021, ISSN: 2673-3951. DOI: 10.3390/modelling2040032.

[50] M. Ben Chaaben, O. Ben Sghaier, M. Dhaouadi, N. Elrasheed, I. Darif, I. Jaoua, B. Oakes, E. Syriani, and M. Hamdaqa, "Toward intelligent generation of tailored graphical concrete syntax," in *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems*, ser. MODELS '24, Linz, Austria: ACM, 2024, pp. 160–171. DOI: 10.1145/3640310.3674085.

[51] J. de Lara and E. Guerra, "Adaptive modelling languages: Abstract syntax and model migration," *ACM Trans. Softw. Eng. Methodol.*, vol. 34, no. 3, 2025, ISSN: 1049-331X. DOI: 10.1145/3702975.

[52] B. Shajari and I. David, *Bridging the silos of digitalization and sustainability by twin transition: A multivocal literature review*, 2025. arXiv: 2506.04267 [physics.soc-ph]. [Online]. Available: https://arxiv.org/abs/2506.04267.

[53] S. Abrahão, F. Bourdeleau, B. Cheng, S. Kokaly, R. Paige, H. Stöerrle, and J. Whittle, "User experience for model-driven engineering: Challenges and future directions," in *2017 ACM/IEEE 20th International Conference on Model Driven Engineering Languages and Systems (MODELS)*, 2017, pp. 229–236. DOI: 10.1109/MODELS.2017.5.

[54] I. David, M. Latifaj, J. Pietron, W. Zhang, F. Ciccozzi, I. Malavolta, A. Raschke, J.-P. Steghöfer, and R. Hebig, "Blended modeling in commercial and open-source model-driven software engineering tools: A systematic study," *Soft Sys Mod*, vol. 22, no. 1, pp. 415–447, 2023. DOI: 10.1007/s10270-022-01010-3.

[55] D. Harel, "Statecharts: A visual formalism for complex systems," *Science of Computer Programming*, vol. 8, no. 3, pp. 231–274, 1987. DOI: https://doi.org/10.1016/0167-6423(87)90035-9.

[56] D. Kundu, D. Samanta, and R. Mall, "Automatic code generation from unified modelling language sequence diagrams," *IET Software*, vol. 7, no. 1, pp. 12–28, 2013. DOI: https://doi.org/10.1049/iet-sen.2011.0080.