

Reputation-driven Decision-making in Networks of Stochastic Agents

David Maoujoud¹ and Gavin Rens²

¹ Katholieke Universiteit Leuven, Belgium
david.maoujoud@student.kuleuven.be, david.maoujoud@hotmail.com

² Katholieke Universiteit Leuven, Belgium
gavin.rens@kuleuven.be

Abstract. This paper studies multi-agent systems that involve networks of self-interested agents. We propose a Markov Decision Process-derived framework, called RepNet-MDP, tailored to domains in which agent reputation is a key driver of the interactions between agents. The fundamentals are based on the principles of RepNet-POMDP, a framework developed by *Rens et al.* [11] in 2018, but addresses its mathematical inconsistencies and alleviates its intractability by only considering fully observable environments. We furthermore use an online learning algorithm for finding approximate solutions to RepNet-MDPs. In a series of experiments, RepNet agents are shown to be able to adapt their own behavior to the past behavior and reliability of the remaining agents of the network. Finally, our work identifies a limitation of the framework in its current formulation that prevents its agents from learning in circumstances in which they are not a primary actor.

Keywords: Uncertainty · Planning · Reputation · MDP · POMDP.

1 Introduction

Decision-making and learning in multi-agent settings is a multi-faceted area of research [4, 2, 3, 14, 7, 6, 1]. Frameworks used for fully cooperative networks of agents differ vastly from those used for networks of self-interested agents. A primary concern when dealing with self-centered agents is that it makes multi-agent learning inherently more complex than single-agent learning [6, 1]. In fact, each agent needs to take into account the behavior of the entire network of agents when learning its own behavior. Additionally, agent behavior tends to be ever-changing. This *non-stationarity* of agent behavior leads to the loss of policy *convergence* properties that can often be found in single-agent formalisms [1].

In 2018, *Rens et al.* [11] proposed a mathematical framework, called *RepNet-POMDP*, designed to handle partially observable environments in which an agent’s reputation among other agents dictates its behavior. The framework was subject to several mathematical inconsistencies, had no working implementation, and had a highly intractable planning algorithm.

Nonetheless, the framework does present some ideas we believe are worth pursuing. Hence, in this paper, we provide an updated version of the framework, called RepNet-MDP. We address the mathematical inconsistencies of the original framework and alleviate its intractability by only considering fully observable environments. We furthermore make use of an online learning algorithm for finding approximate solutions to RepNet-MDPs. The viability of the framework is tested in a series of experiments designed to highlight its strengths and shortcomings.

Section 2 summarizes the relevant background required. Section 3 gives an overview of the work related to our framework. Section 4 provides an intuitive introduction to RepNet-MDPs. Section 5 covers the formal definition of the framework. Section 6 covers planning for RepNet-MDPs. The experimental setup and results are given in Section 7.

2 Background - Markov Decision Processes

A Markov Decision Process (MDP) describes a process for modeling decision-making in stochastic environments [13]. An agent is assumed to move about in an environment, described by a set of states \mathcal{S} , by applying actions in \mathcal{A} to the environment. The transition rules of the environment are dictated by the transition model $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$, that is, $\mathcal{T}(s, a, s')$ returns the probability of the agent transitioning to state s' upon performing action a in state s . Each action applied to the environment results in a reward for the agent, dictated by the reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, that is, $\mathcal{R}(s, a)$ returns the reward received by the agent when performing action a in state s .

The objective of an MDP agent is to maximize its long-term cumulative reward, called *utility*. The utility U of a *finite* state-action sequence, sometimes called *episode*, $E = \langle s_0, a_0, s_1, a_1, \dots, s_T, a_T \rangle$ is defined as [10]:

$$U(E) = \sum_{t=0}^T \gamma^t \mathcal{R}(s_t, a_t),$$

where $\gamma \in [0, 1]$ is called the *discount factor*. An agent advances in the environment by following a *policy* $\pi : \mathcal{S} \times \mathbb{N} \rightarrow \mathcal{A}$ that maps each environment state and remaining time-steps to the action the agent should take.

The *expected utility*, or *value*, of being in any state s_t at time-step t , while following policy π , with d time-steps remaining, is defined as:

$$V^\pi(s_t, d) = \mathbb{E}[U(E_t) \mid s_t, \pi] = \mathbb{E}\left[\sum_{k=t}^{t+d} \gamma^{k-t} \mathcal{R}(s_k, a_k) \mid s_t, \pi\right],$$

where E_t is the sub-sequence of E starting at time-step t . An optimal policy π^* is a policy such that

$$\forall s \in \mathcal{S}, \forall d \in \mathbb{N}, \forall \pi : V^*(s, d) \geq V^\pi(s, d),$$

where $V^* : \mathcal{S} \times \mathbb{N} \rightarrow \mathbb{R}$ is the value function associated with optimal policy π^* . This policy satisfies the *optimality equations*, also known as the *Bellman equations* ($\forall s \in \mathcal{S}$):

$$\begin{cases} V^*(s, d) := \max_{a \in \mathcal{A}} \left\{ \mathcal{R}(s, a) + \gamma \sum_{s' \in \mathcal{S}} \mathcal{T}(s, a, s') V^*(s', d-1) \right\} & d > 1 \\ V^*(s, 1) := \max_{a \in \mathcal{A}} \left\{ \mathcal{R}(s, a) \right\} \end{cases}$$

Partially Observable Markov Decision Processes are a common extension of classic MDPs that deal with the problem of partial observability of the environment [13]. To address the agent’s inability to observe the exact state of the environment, the observation function $\mathcal{O} : \mathcal{A} \times \mathcal{S} \times \Omega \rightarrow [0, 1]$, where Ω is the set of observations, is introduced. $\mathcal{O}(a, s', o)$ returns the probability of the agent making observation o after performing action a and the environment transitioning to state s' .

Instead of working with the actual states of the environment, the POMDP agents make use of the notion of belief state $b \in \Delta(\mathcal{S})$ ³, which is a probability distribution over the possible states of the environment. Belief states are updated using the *state estimation function* SE defined as follows:

$$b' := SE(b, a, o) := \left\{ (s', p) \mid s' \in \mathcal{S} \wedge p = \frac{\mathcal{O}(a, s', o) \sum_s \mathcal{T}(s, a, s') b(s)}{P(o|b, a)} \right\},$$

where $P(o|b, a) = \sum_{s' \in \mathcal{S}} \mathcal{O}(a, s', o) \sum_{s \in \mathcal{S}} \mathcal{T}(s, a, s') b(s)$ is a normalizing constant. We refer to [13] for an extensive overview of POMDPs.

3 Related MDP-based frameworks

Early multi-agent frameworks, such as Multi-agent Markov Decision Processes (MMDPs) [4] and Decentralized Partially Observable MDPs (Dec-POMDPs) [2, 3], operate under the assumption that the agents are selfless and have a common goal. Consequently, planning can be centralized, that is, each agent’s policy can be computed by central unit, before being distributed amid the agents for execution [14]. Dec-POMDPs furthermore differ from MMDPs in that states are no longer fully observable, meaning that each agent is in possession of its own set of *local observations*.

In 2005, *Gmytrasiewicz et al.* formalized an extension of POMDPs to multi-agent settings, called Interactive-POMDP (I-POMDP) [7]. I-POMDPs are designed for reasoning in networks of selfish agents. I-POMDP agents update their beliefs not only over physical states of the environment but also over models of the other agents in the network. The difficulty of solving I-POMDPs lies in the recursive nature of the models. Consider agent g ’s belief update in a network inhabited by another agent, say h . A model of agent h may consist of the belief function of said agent h over physical states and models of all other agents.

³ $\Delta(\mathcal{E})$ is the set of probability distributions over the elements of set \mathcal{E} .

These models may, in turn, consist of belief functions of their own. This nesting of beliefs could theoretically be infinite, but is overcome by bounding the nesting depth by a finite number n , and solving the problem as a set of POMDPs.

The RepNet-MDP framework [11] simplifies the notion of model by focusing in on key concepts such as behavioral habits and reputation of other agents. While this reduces the insights RepNet agents can have into other agents' behavior, it makes the framework arguably more intuitive. The key, novel notion in the RepNet framework is that of subjective transitions, which have a dependence on the reputation of the agent performing the action.

4 Developing an intuition for RepNet-MDPs

To develop an intuition for the RepNet-MDP framework, parallels between the concepts found in classic POMDPs and RepNet-MDPs can be drawn. In a POMDP, a single agent, placed in a partially observable environment, applies an action a^* it deems optimal as per its current policy π^* , and is sent back an observation o . The state estimation function SE can be thought of as a way of extracting information from said observation o , and storing it in a belief state b' . More specifically, o contains information about the actual state of the environment. The POMDP loop is depicted in Fig. 1a.

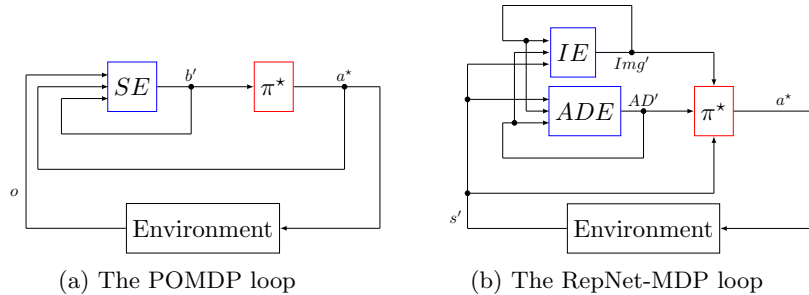


Fig. 1: POMDP and RepNet-MDP loops.

Let us now consider a fully observable environment made up of 3 selfish agents, of which the behavior of the first is dictated by the RepNet-MDP framework. The willingness of the RepNet agent to engage with agents 2 or 3 is to be conditioned by their reputation and behavioral habits. The first agent once again applies action a^* , as per its policy π^* . The environment returns its new state s' . In an effort to make well-informed decisions, the RepNet agent should extract the other agents' behavior from s' .

Two functions, analogous to the state estimation function SE in POMDPs, are used to this end: The action distribution estimation function ADE extracts information regarding other agents' behavioral habits. The image estimation function IE informs the RepNet agent on the image all the agents have of each other. The RepNet-MDP loop is shown in Fig. 1b.

Closely tied to the concept of image is the notion of reputation. Specifically, the reputation of any agent in the framework can be seen as a summary of the

information encapsulated by the image. Unlike POMDPs, RepNet-MDPs feature two types of actions, and by extension two types of transition models:

- *Objective* actions, which, when performed, have a real effect on the environment. These actions can be seen as equivalent to actions as they exist in MDPs. The associated transition model is called the objective transition model OT and describes the rules of the environment as they apply to the RepNet agent.
- *Subjective* actions, which, unlike objective actions, are never actually applied to the environment. Instead, they are associated with another transition model called the subjective transition model ST : This transition model describes a RepNet agent’s subjective perception of the rules of the environment. This perception is a function of said agent’s reputation, and can be used by the agent to aid in its decision-making.

5 Formal definition of RepNet-MDPs

In this section, we will formalize the RepNet-MDP framework⁴ introduced in Section 4. A RepNet-MDP \mathcal{M} is defined as a pair of tuples $\mathcal{M} := \langle \Sigma, \Gamma \rangle$, where Σ is called the System tuple and incorporates aspects of the network that apply to all agents, and Γ is called the Agents tuple and contains each agent’s subjective understanding of the environment it operates in.

Specifically, a System in a RepNet-MDP Σ is formally defined as a tuple

$$\Sigma := \langle \mathcal{G}, \mathcal{S}, \mathcal{A}, \mathcal{I}, \mathcal{U}, OT \rangle,$$

where:

- \mathcal{G} is the set of agents that can interact with the environment.
- \mathcal{S} is the set of possible states of the environment.
- \mathcal{A} is the set of possible actions, both *objective* and *subjective*. Formally, $\mathcal{A} := \mathcal{A}^o \cup \mathcal{A}^s$, $\mathcal{A}^o \cap \mathcal{A}^s := \emptyset$. The concept of *subjective actions* will be discussed in Section 5.2.
- $\mathcal{I} : \mathcal{G} \times \mathcal{G} \times \mathcal{S} \times \mathcal{A} \rightarrow [-1, 1]$ is called the *impact function*. $\mathcal{I}(g, h, s, a)$ returns the impact on agent g that is due to agent h performing action a in state s . This function can be thought of as analogous to a Markov Decision Process’s immediate reward function \mathcal{R} .
- $\mathcal{U} : [-1, 1] \times [-1, 1] \rightarrow [-1, 1]$ is called the *image update function*. Given a current value v , of the image an agent has of another agent, to be updated, and a new expected total impact i , of which the definition will be given shortly, $\mathcal{U}(v, i)$ returns an updated value of the image v' . Many instantiations of this function are possible, two of which are presented in [11]. We will use the following instantiation:

$$\mathcal{U}(v, i) := \begin{cases} v + (1 - v)i & \text{if } i \geq 0 \\ v + (1 + v)i & \text{if } i < 0 \end{cases}$$

⁴ The implementation of the RepNet-MDP framework can be found at <https://github.com/davidmaoujoud/RepNet-MDP>

- $OT : \mathcal{G} \times \mathcal{S} \times \mathcal{A}^o \times \mathcal{S} \rightarrow [0, 1]$ is called the *objective transition model*. $OT(h, s, a, s')$ returns the probability of the environment transitioning from state s to state s' when objective action a is taken by agent h .

In addition to the global information stored in Σ , each RepNet agent’s subjective knowledge is stored in the Agents tuple Γ , formally defined as ⁵

$$\Gamma := \langle \{ST_g\}, \{AD_g\}, \{Img_g\} \rangle,$$

where:

- $ST_g : \mathcal{G} \times \mathcal{S} \times \mathcal{A}^s \times \mathcal{S} \times [-1, 1] \rightarrow [0, 1]$ is called the *subjective transition model* of agent g . $ST_g(h, s, a, r_h, s')$ returns the probability, *as perceived by agent g* , of the environment transitioning from state s to state s' if agent h were to perform subjective action a , and has a reputation r_h according to agent g .
- $AD_g : \mathcal{G} \times \mathcal{S} \rightarrow \Delta(\mathcal{A})$ is called the *action distribution* according to agent g . $AD_g(h, s)$ returns a probability distribution over actions in \mathcal{A} for agent h in state s , according to agent g .
- $Img_g : \mathcal{G} \times \mathcal{G} \rightarrow [-1, 1]$ is called the *image function* according to agent g . $Img_g(h, i)$ returns the image agent i has of agent h according to agent g . Said differently, it returns what g thinks i thinks of h .

As introduced in Section 4, every agent bases its decision-making on the image it believes all agents to have of each other, as well as each agent’s behavioral habits. Let g be an agent, whose image at time t is Img_g , and action distribution is AD_g . At time $t + 1$, these constructs are updated via the image estimation function IE and action distribution estimation function ADE respectively, to produce Img'_g and AD'_g .

5.1 Image and Reputation

This subsection builds towards the formal definition of the image estimation function IE . To this end, we introduce the notion of *expected total impact*. Consider two agents h and i . In any given state, agent h can perform one of several actions which may or may not have an impact on agent i . Likewise, agent i can be expected to have an impact on agent h when performing an action. The *expected total impact* should be thought of as a way of assigning a numerical value to the *bidirectional* impact these two agents can be expected to have on each other. Additionally, one direction of the impact may be perceived as more important than the other and thus be weighed differently. According to an observing agent, say g , the total impact h is expected to have on, as well as perceive from, i when the environment is in state s , is defined as

$$ETI_g(h, i, s, AD_g) := \sum_{a \in \mathcal{A}} \left[\delta AD_g(i, s)(a) \mathcal{I}(h, i, s, a) + (1 - \delta) AD_g(h, s)(a) \mathcal{I}(i, h, s, a) \right],$$

⁵ $\{X_g\}$ is used as the shorthand notation for $\{X_g \mid g \in \mathcal{G}\}$

where $\delta \in [0, 1]$ weighs the importance of impact due to agent h and impact perceived by h .

Agent g 's image of other agents, as well as the image it believes all agents to have of each other changes as it observes the agents' behavior. Let Img_g be the current image function of agent g . Concretely, we wish to *update* the image any agent i has of any other agent h , according to the observing agent g ($= Img_g(h, i)$) on the basis of the impact h is expected to have on i ($= ETI_g(h, i, s, AD_g)$). The updated image function Img'_g is computed as follows:

$$\begin{aligned}
 Img'_g &:= IE(g, Img_g, \alpha, s, AD_g) \\
 &:= \left\{ (h, i, t) \mid h, i \in \mathcal{G} \wedge t = \mathcal{U}(Img_g(h, i), ETI_g(h, i, s, AD_g)) \right\},
 \end{aligned}$$

where s is the current state of the environment, IE is called the *image estimation function*, and \mathcal{U} is the *image update function*.

Finally, the notion of reputation as it is understood in this framework can be thought of as a way of summarizing the information encapsulated by the image.

Say agent g wishes to estimate the reputation of agent h in a network made up of several other agents. It can, to this end, use the image each agent i has of agent h ($= Img_g(h, i)$) as a guiding principle. A first idea might be to take agent h 's reputation to be equal to its average image in the network. If, however, some agent i has a poor image of agent h ($Img_g(h, i) < 0$), but agent g has a poor image of agent i ($Img_g(i, g) < 0$), it may be unreasonable for agent g to assume that agent i 's opinion of agent h is indicative of agent h 's reputation being poor. These concerns are addressed by weighing the image each agent i has of h by the image g has of i . As such, if both images are negative, the resulting reputation of h will not be affected negatively ($Img_g(h, i) \times Img_g(i, g) > 0$). Formally, the reputation of an agent h , according to agent g , is defined as

$$REP_g(h, Img_g) := \frac{1}{|\mathcal{G}'|} \sum_{i \in \mathcal{G}'} Img_g(h, i) \times Img_g(i, g),$$

where $Img_g(i, i) = 1 \forall i \in \mathcal{G}$, and $\mathcal{G}' = \mathcal{G}$ if $h \neq g$ and $\mathcal{G}' = \mathcal{G} \setminus \{g\}$ if $h = g$. Recall that in the RepNet framework, reputation influences subjective transition probabilities, which, in turn, influence a RepNet agent's planning.

5.2 Subjective actions and the subjective transition model

In this section, we describe the use of subjective actions and subjective transition models in the RepNet framework. As introduced in Section 4, we make a distinction between the purpose of an *objective* transition model, which describes the actual rules of the environment as they apply to the RepNet agent, and that of a *subjective* transition model, which describes that agent's *subjective* perception of the rules of the environment, this perception being influenced by the reputation of the RepNet agent. To illustrate this further, we will make use of a simple trading example between two agents A and B . Agent A wishes to

trade with agent B , who can either accept or refuse the trade offer. The environment is made up of the set of states $\mathcal{S} = \{s_0, s_1, s_a, s_r\}$. s_0 is the initial state, prior to any trade, s_1 is the state in which agent B is made aware of agent A 's trade offer, s_a is the accept state, and s_r is the refuse state. The set of *objective* actions at the disposal of both agents is given by

$$\mathcal{A}^o = \{\text{trade_with_A}, \text{trade_with_B}, \text{accept}, \text{refuse}, \text{wait}\}.$$

The transition model of the environment assumed to be deterministic, is given in Fig. 2. *In the eyes of* agent A , agent B 's response to a trade offer, characterized




Fig. 2: Transition model of the environment (trading example). Each transition has two *objective* actions, the first action represents agent A 's *objective* action, the second action represents agent B 's *objective* action.

by transitions $s_1 \rightarrow s_a$ and $s_1 \rightarrow s_r$, depends on A 's reputation. The action taken by agent A during these transitions is **wait**. To make use of the notion of subjective actions, the set of *subjective* actions \mathcal{A}^s will contain the counterpart⁶ of **wait** in its *subjective* form, that is, $\mathcal{A}^s = \{\text{wait_s}\}$.

The way agent A makes use of actions in \mathcal{A}^o and \mathcal{A}^s can now be detailed. When *planning* to maximize its expected impact, agent A will make use of the *objective* transition model whenever the action currently investigated has no subjective counterpart in \mathcal{A}^s . For instance, the transition probability when investigating action **trade_with_B** is given by $OT(A, s_0, \text{trade_with_B}, s_1)$. When an action in \mathcal{A}^o has a counterpart in \mathcal{A}^s , agent A will make use of the subjective transition model. For instance, the transition probability when investigating action **wait/wait_s** is given by $ST_A(A, s_1, \text{wait_s}, s_a, r_A)$. As such, the reputation of agent A is accounted for when agent A *plans* to maximize its expected impact.

5.3 Action distribution

The next step in the formalization of RepNet-MDPs consists in redefining the updating scheme of the action distribution AD_g of each agent g . Say an environment hosting two agents g and h is currently in state s . Agent g has an *a priori* notion of the probability of agent h picking an action a in state s , $P_g(a|h, s, r_h)$. Following agent h performing action a in this state, the environment transitions

⁶ We define *counterpart* as a partial mapping $\mathcal{C} : \mathcal{A}^o \rightarrow \mathcal{A}^s$. If $\mathcal{C}(a)$ is not defined, then a has no counterpart in \mathcal{A}^s .

from state s to state s' . The *a posteriori* probability of agent h performing that same action a in state s in the future is now computed using Bayes' rule:

$$P_g(a|h, s, r_h, s') = \frac{P_g(s'|h, s, r_h, a)P_g(a|h, s)}{\sum_{a'} P_g(s'|h, s, r_h, a')P_g(a'|h, s)}.$$

One can add a *smoothing* technique called *Laplace smoothing* to the present result in an effort to avoid undesirable side effects related to the use of deterministic transition models. The definition for the *action distribution estimation*, obtained after replacing the probabilities by the terms defined previously and applying the smoothing technique, is then given by:

$$\begin{aligned} AD'_g &:= ADE(g, s', AD_g, Img_g) \\ &:= \left\{ (h, s, a, p) \mid h \in \mathcal{G} \wedge s \in \mathcal{S} \wedge a \in \mathcal{A} \wedge r_h = REP_g(h, Img_g) \right. \\ &\quad \left. \wedge p = \frac{T_g(h, s, a, s', r_h)AD_g(h, s)(a) + \eta}{\sum_{a'} (T_g(h, s, a', s', r_h)AD_g(h, s)(a') + \eta)} \right\}, \end{aligned}$$

where *ADE* is called the *action distribution estimation function*, s' is the state the environment transitions to, and η is the *Laplace smoothing parameter*. We refer to the accompanying extensive version of this paper [9] for more details.

Note that to simplify the notation, we combined the *objective* and *subjective* transition models into a single model T_g , called the *global transition model* and formally defined as

$$T_g(h, s, a_h, s', r_h) := \begin{cases} ST_g(h, s, a_h, s', r_h) & \text{if } a_h \in \mathcal{A}^s \\ OT(h, s, a_h, s') & \text{if } a_h \in \mathcal{A}^o \end{cases} \quad (1)$$

6 Planning in the RepNet framework

We now describe optimal behavior in the context of RepNet-MDPs, for finite horizon look-ahead. To simplify the notation, we can define a construct called *epistemic state*. The epistemic state θ_g of agent g is formally defined as a tuple $\theta_g := \langle s, AD_g, Img_g \rangle$, where s is the current state of the environment, AD_g is the current action distribution of agent g , and Img_g is the current image function of agent g . $\theta_g \in \Theta_g$, and Θ_g is called the *epistemic state space*. This set contains every possible combination of physical states of the environment, action distributions, and image functions of agent g .

An agent should perform actions according to the *perceived immediate impact* they have on the agent itself. The perceived immediate impact on agent g resulting from performing action a in state s is defined as

$$PI_g(s, AD_g, a) := \frac{1}{|\mathcal{G}|} \left[\mathcal{I}(g, g, s, a) + \sum_{h \in \mathcal{G} \setminus \{g\}} \sum_{a' \in \mathcal{A}} \mathcal{I}(g, h, s, a')AD_g(h, s)(a') \right],$$

where AD_g is the current action distribution of agent g . The first term describes the immediate self-impact as a consequence of agent g performing action a , while the second term describes the expected immediate impact that the network (i.e., the remaining agents) has on agent g .

Analogously to regular MDPs, a RepNet-MDP agent g strives to maximize its expected discounted perceived impact $\mathbb{E}[\sum_{t=0}^k \gamma^t PI_{g,t}]$, where γ is the *discount factor* and $PI_{g,t}$ is agent g 's perceived immediate impact at time-step t . This is accomplished by computing the optimal value function $V_g : \Theta_g \times \mathbb{N} \rightarrow \mathbb{R}$ (in a *finite-horizon* setting). It satisfies the *optimality equations*, which are defined as ($\forall \theta_g \in \Theta_g$)

$$\begin{cases} V_g(\theta_g, k) := \max_{a \in \mathcal{A}} \left\{ PI_g(s, AD_g, a) + \gamma \sum_{s' \in \mathcal{S}} T_g(g, s, a, s', r_g) V_g(\theta'_g, k - 1) \right\} \\ V_g(\theta_g, 1) := \max_{a \in \mathcal{A}} \left\{ PI_g(s, AD_g, a) \right\} \end{cases} \quad (2)$$

where $r_g = REP_g(g, Img_g)$, $\theta_g = \langle s, AD_g, Img_g \rangle$, and $\theta'_g = \langle s', ADE(g, s', AD_g, Img_g), IE(g, Img_g, \alpha, s, AD_g) \rangle$.

In this work, we implement (approximate) online planning [12] instead of exact planning. The general principle of model-based online planning can be described as the interleaving of two phases, the *planning phase*, in which the (PO)MDP performs a look-ahead search of a given depth D , starting at the current environment state, the goal being to determine the most suitable action, and the *execution phase*, in which this action is applied to the environment [12]. In this paper, we make use of an implementation of this approximate technique for the RepNet-MDP framework. We refer to [9] for details on the implementation.

7 Experiments

The goal of the experiments is to showcase the strengths and shortcomings of the framework. To this end, the experimental setup consists of 2 trading scenarios, for which several experiments are conducted. All experiments were conducted with *look-ahead depth* $D = 3$, *Laplace smoothing parameter* $\eta = 0.1$, and *discount factor* $\gamma = 0.7$. Note that these experiments serve as a proof of concept for the RepNet framework and, as such, are not designed to reflect the framework's applicability to problems of realistic scale.

7.1 Experiment 1: Trading between two agents

Let A and B be two agents. Agent A plays the role of the buyer, agent B the role of the seller. Agent A can engage in a trade with agent B , and B can accept or refuse the trade offer. Furthermore, agent A can, prior to making a trade offer, do a good deed in an effort to improve its image in the eyes of agent B .

In this series of experiments, Agent A is managed by the RepNet algorithm. Agent B is run by a simple algorithm that accepts or rejects trade offers made by

agent A according to a set schedule. In particular, agent B is asked to reject trade offers for the 20 first time-steps, accept them for the 60 subsequent time-steps, and finally reject them for the last 20 time-steps.

Two series of experiments are conducted, the first one without making use of subjective actions, the second one by modeling the action of agent A awaiting agent B 's response to a trade offer as a subjective action, meaning the outcome of agent A 's planning will be influenced by its reputation. A well-designed subjective transition model, schematized in Fig. 3, that realistically reflects how the reputation of agent A may influence the willingness of agent B to accept A 's trade offers is put to the test. The variables tracked are the action distribution, image, and by extension the reputation of both agents *in the eyes of agent A*, and frequency at which agent A makes trade offers.






Fig. 3: Perceived probability of agent B accepting and refusing the trade offers, as a function of the self-reputation of agent A .

Fig. 4 shows the evolution of agent A 's action distribution for target agent B . Fig. 5 shows the evolution of agent A 's self-reputation during the experiment involving the subjective transition model. Note that A 's self-reputation, and more generally A 's image function, have no bearing on its decision-making if no subjective actions are used (see Equations 1 and 2, T_g makes use of the notion of reputation only for subjective actions). Finally, Fig. 6 shows the evolution of the frequency at which A makes trade offers.

In the first 20 time-steps, B refuses each trade offer. Regardless of the series of experiments, agent A is able to pick up on this via the action distribution. As a consequence, it quickly reduces the frequency at which it attempts to trade with B . In the 60 following time-steps, B is asked to change its behavior and accept each trade offer. Hesitant at first, A gradually increases the frequency at which it attempts to trade with B . Agent A is able to pick up on B reverting back to its old behavior during the final 20 steps.

Additionally making use of a well-designed subjective transition model noticeably improves the RepNet agent's performance. While the trajectories showcase the same key elements, the pace at which agent A is able to adapt improves greatly. The subjective transition model was designed such that agent A believes that its reputation must be good for B to be willing to trade with A (Fig. 3). As such, during the first 20 time-steps, A 's relatively poor-in-comparison self-reputation has an immediate negative effect on the *value* it associates with the `trade.with.B` action during the look-ahead search. It quickly becomes more *valuable* to stop trading with B . Similarly, A 's reputation needs to be high for it to start trading with B again, explaining the slow increase of the frequency of trade offers at the start of the second phase.


Fig. 4: Probability of B accepting A 's trade offers, according to A Fig. 5: Reputation of agent A , according to itself.

7.2 Experiment 2: Trading between three agents

Let A , B , and C be three agents. Each agent simultaneously plays the role of buyer and seller, and can thus engage in a trade with any other agent. Each agent can accept or refuse any trade offer made by any remaining agent.

The present scenario is used to verify the ability of a RepNet agent, say agent A , to manage its trades with the two remaining agents B and C , based not only on their behavior towards the agent of interest but also their behavior with each other.

In the first part, agent B is asked to refuse each trade offer made by agent A , while agent C is expected to accept each trade offer coming from A . This portion of the experiments assesses the ability of the RepNet agent (agent A) to accurately determine which agent it is more likely to successfully engage in trades with. In the second part, the roles are switched, and agent B accepts the trade offers, while agent C refuses them. This portion assesses the ability of the agent of interest to *unlearn* what it has learned and adapt its behavior accordingly. In the third and final part, the RepNet agent is asked to not trade with either B or C , that is, to only make use of the `wait` action. Said differently, the optimal action according to its planning, while tracked throughout the experiment, is not performed on the environment. All the while, agents B and C are asked to

Fig. 6: Frequency of the trade offers made by A , measured in 5 time-step intervals

engage in trades with each other. Agent B is asked to reject all trade offers, while agent C is asked to accept all trade offers. The variables tracked are the action distribution and reputation of B and C *in the eyes of agent A*, as well as the evolution of whom agent A would rather trade with. This portion of the experiment aims at testing the ability of the RepNet agent to draw conclusions on how it should act based on interactions it is not directly affected by.



Fig. 7 shows the evolution of the reputations of agents B and C . Fig. 8 displays the evolution of the probabilities of agents B and C accepting trade offers from agent A . Finally, Fig. 9 shows the evolution of whom agent A would rather trade with.

Agent B is told to refuse, and agent C to accept, each trade offer during the first 33 time-steps. In accordance with the results obtained in Section 7.1, agent A is able to pick up on the other agents' behavioral habits it is affected by. As a result, the reputation of B and its probability of accepting trade offers decrease. Similarly, the reputation of C and its probability of accepting trade offers increase. All the while, agent A chooses to conduct the majority of its trades with C . The following 33 time-steps reverse B 's and C 's roles. Similarly, agent A is able to adapt its behavior accordingly and ends up trading mostly with B . The reputation of B has increased, while the reputation of C has decreased.

During the last 33 time-steps, agents B and C are tasked with trading with one another while A plays the role of observer, that is, only makes use of the `wait` action. B is asked to refuse all trade offers, while C is asked to accept all trade offers. Interestingly, Fig. 9 shows that, based on its planning, agent A would prefer to keep trading with B , even though the reputation of B decreases and the reputation of C increases in the eyes of A . Said differently, as long as B does not refuse A 's offers directly, agent A will prefer to trade with B over C .

The explanation for this is twofold. Firstly, the subjective transition probability of a trade A might want to do with B is, in the eyes of A , conditioned only by A 's *own* reputation. As such, B 's falling or rising reputation has no bearing on A 's decision-making. Secondly, the probability of B accepting (or refusing) A 's trade offer, according to A , can only be updated through the direct experience it has with B . As such, the action distribution does not change and can thus not influence A decision-making either.

The simplest way of alleviating this shortcoming is to extend the subjective transition model. Adding the reputation of the agent at the receiving end of the trade offer (e.g., agent B) as a parameter to the subjective transition model would allow agent A to incorporate other agents' reputation in its decision-making process. As such, if the subjective transition probability of B accepting A 's trade offer were given by $ST_A(A, \text{offer_state}, \text{wait_s}, \text{accept_state}, r_A, r_B)$, where the newly introduced parameter r_B is B 's reputation, agent A could make use of r_B to assist with its decision-making. This comes with the drawback of increasing the complexity of designing the subjective transition model.

Fig. 7: Reputation of agents B and C , according to A Fig. 8: Probability of agents B and C accepting trade offers from agent A , according to A


8 Summary and future work

In this paper, we revised the multi-agent framework called RepNet introduced by *Rens et al.* [11], addressed its mathematical inconsistencies and proposed an online learning algorithm for finding approximate solutions. The viability of the framework was then tested in a series of experiments.

The current definition of *objective* transitions could be extended to incorporate the reputation of agents other than the RepNet agent. The experimental results showed that the RepNet agent is incapable of adapting its behavior to situations that do not directly affect it. Including the reputation of the agent at the receiving end of a directed action in the *directed* transition model is likely to lead to better-informed decision-making.

We did not address partially observable environments. Many real-world problems do not benefit from full observability, bringing the updated RepNet framework back to a partially observable setting should be considered for future work.

The small-scale experiments conducted in Section 7 served as a proof of concept for the RepNet framework. While applying the framework to problems

Fig. 9: Average action taken by agent A . Action $a = 0$ corresponds to trading with agent B , action $a = 1$ corresponds to trading with agent C .

of realistic size was beyond the scope of this paper, the absence of large-scale tests does raise questions as to the scalability of this approach. Real-world problems can easily become too complex for transition models to be designed by any one person without leveraging common state features [5]. A compact way to represent real-world state spaces can be achieved by introducing elements of relational logic [5]. From a *logic programming* point of view, a state space is hereby defined by a collection of relations, while a state is an *interpretation* of this collection [8]. Transition models and reward schemes are then represented by *probabilistic rules* [10].

References

1. Abbeel, P.: Learning for Robotics and Control - Value Iteration, CS294-40, University of California, Berkeley (2008), <https://inst.eecs.berkeley.edu/cs294-40/fa08/scribes/lecture2.pdf>
2. Becker, R., Zilberstein, S., Lesser, V., Goldman, C.: Solving Transition Independent Decentralized Markov Decision Processes. *J. Artif. Intell. Res. (JAIR)* **22**, 423–455 (07 2004)
3. Bernstein, D.S., Zilberstein, S., Immerman, N.: The Complexity of Decentralized Control of Markov Decision Processes. *CoRR* **abs/1301.3836** (2013)
4. Boutilier, C.: Planning, Learning and Coordination in Multiagent Decision Processes. In: *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge*. pp. 195–210. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1996)
5. Boutilier, C., Reiter, R., Price, B.: Symbolic Dynamic Programming for First-Order MDPs. pp. 690–700 (01 2001)
6. Busoniu, L., Babuska, R., De Schutter, B.: A Comprehensive Survey of Multiagent Reinforcement Learning. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* **38**, 156 – 172 (04 2008)
7. Doshi, P., Gmytrasiewicz, P.J.: A Framework for Sequential Planning in Multi-Agent Settings. *CoRR* **abs/1109.2135** (2011)
8. Joshi, S., Kersting, K., Khardon, R.: Self-taught decision theoretic planning with first order decision diagrams. pp. 89–96 (01 2010)
9. Maoujoud, D., Rens, G.: Reputation-driven Decision-making in Networks of Stochastic Agents (2020), <https://arxiv.org/abs/2008.11791>
10. Nitti, D., Belle, V., De Laet, T., De Raedt, L.: Planning in hybrid relational MDPs. *Machine Learning* **106**(12), 1905–1932 (Dec 2017)
11. Rens, G., Nayak, A., Meyer, T.: Maximizing Expected Impact in an Agent Reputation Network - Technical Report. *CoRR* **abs/1805.05230** (2018), <http://arxiv.org/abs/1805.05230>
12. Ross, S., Pineau, J., Paquet, S., Chaib-draa, B.: Online Planning Algorithms for POMDPs. *CoRR* **abs/1401.3436** (2014)
13. Russell, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*. Prentice Hall Press, Upper Saddle River, NJ, USA, 3rd edn. (2009)
14. Wiering, M., Otterlo, M.: *Reinforcement Learning: State-Of-The-Art*, vol. 12 (01 2012)