

Multi-agent assignment via state augmented reinforcement learning

Leopoldo Agorio

University of Pittsburgh, Pittsburgh, PA, USA

LCA31@PITT.EDU

Sean Van Alen

University of Pittsburgh, Pittsburgh, PA, USA

SGV13@PITT.EDU

Miguel Calvo-Fullana

Universitat Pompeu Fabra, Barcelona, Spain

MIGUEL.CALVO@UPF.EDU

Santiago Paternain

Rensselaer Polytechnic Institute, Troy, NY, USA

PATERS@RPI.EDU

Juan Andrés Bazerque

University of Pittsburgh, Pittsburgh, PA, USA

JUANBAZERQUE@PITT.EDU

Editors: A. Abate, M. Cannon, K. Margellos, A. Papachristodoulou

Abstract

We address the conflicting requirements of a multi-agent assignment problem through constrained reinforcement learning, emphasizing the inadequacy of standard regularization techniques for this purpose. Instead, we recur to a state augmentation approach in which the oscillation of dual variables is exploited by agents to alternate between tasks. In addition, we coordinate the actions of the multiple agents acting on their local states through these multipliers, which are gossiped through a communication network, eliminating the need to access other agent states. By these means, we propose a distributed multi-agent assignment protocol with theoretical feasibility guarantees that we corroborate in a monitoring numerical experiment.

Keywords: Constrained reinforcement learning, multi-agent assignment, monitoring problem.

1. Introduction

Reinforcement Learning (RL) has experienced growing interest as a methodology for designing optimal controllers in cases where system dynamics are intractable or unknown [Kahn et al. \(2017\)](#). It has attained remarkable achievements in recent years, with agents capable of playing Go [Silver et al. \(2016\)](#) and Poker [Brown and Sandholm \(2019\)](#), applications in robotics [Kober et al. \(2013\)](#), and in training recommendation systems and natural language models [Christiano et al. \(2017\)](#).

Incorporating constraints into the RL problem offers extra capabilities, such as ensuring safety in robot navigation [Miryoosefi et al. \(2019\)](#); [Paternain et al. \(2022\)](#), imposing physical conditions [Li et al. \(2019\)](#); [Gao et al. \(2020\)](#); [Porteiro et al. \(2024\)](#), satisfy natural language rules [Ouyang et al. \(2022\)](#), or in general, setting problem specifications. Constrained RL is commonly addressed through regularization to balance conflicting requirements, weighting the rewards to create a multi-objective RL problem [Kahn et al. \(2017\)](#). In this approach, selecting the appropriate regularization weights via duality [Paternain et al. \(2019\)](#), Bayesian rules [Hutter \(2002\)](#), or simply heuristics is crucial for achieving good performance [Tessler et al. \(2018\)](#).

However popular, there are problems for which regularization approaches are insufficient to ensure feasibility [Calvo-Fullana et al. \(2023\)](#). In this paper, we focus on one such family of problems, assignment problems [Bertsekas \(1981\)](#); [Mills-Tettey et al. \(2007\)](#); [Chopra et al. \(2017\)](#), in a

dynamic multi-agent setup in which several tasks are assigned to a team of agents. We will argue that regularized formulations of these multi-agent assignment RL problems fail to produce policies satisfying the conflicting constraints, lacking convergence. To overcome this challenge, we leverage a recent state augmentation approach, where the Lagrangian weights are reinterpreted as states of an augmented Markov Decision Process (MDP) Calvo-Fullana et al. (2023). Under this model, the multipliers are not driven to convergence to a dual optimum. Instead, they are allowed to oscillate to induce alternating policies that satisfy the conflicting tasks sequentially.

We further demonstrate that these multipliers can be used for coordinating multiple agents, especially when many conflicting tasks cannot be achieved by a single agent or several ones acting independently, a state-of-the-art challenge in multi-agent reinforcement learning Zhang et al. (2021). In this context, we propose an assignment protocol to coordinate a team of agents running pre-trained distributed policies, where each agent might have only local information about the state of the system. For a fully distributed algorithm, we use the gossiping framework Aysal et al. (2009) to share the multipliers across a communication network established by the agents, and we prove consensus in finite time, taking advantage of the binary nature of the assignment variables.

Our state augmentation and multi-agent assignment protocol come with theoretical guarantees of almost sure feasibility, which we corroborate with numerical experiments involving a team of robots that must patrol an area of interest alternating between different zones.

2. Problem formulation

Consider N agents interacting with an environment that is modeled as an MDP with transition probabilities $\mathbb{P}(S_{t+1} | S_t, A_t)$. For each time index $t = 0, 1, 2, \dots$, variables $S_t = (S_{t1}, \dots, S_{tN})$ and $A_t = (A_{t1}, \dots, A_{tN})$ collect the states $S_{tn} \in \mathcal{S}$ and actions $A_{tn} \in \mathcal{A}$ for all agents. The team receives M rewards $r_m(A_t, S_t)$ associated with M different tasks that need to be satisfied collectively. We consider a family of assignment problems in which rewards are given by

$$r_m(S_t, A_t) = \max_{n=1, \dots, N} \mathbb{1}[S_{tn} \in \mathcal{S}_m] = \begin{cases} 1 & \text{if } \exists n \in \{1, \dots, N\} : S_{tn} \in \mathcal{S}_m \\ 0 & \text{otherwise} \end{cases}, \quad (1)$$

defined in terms of M regions $\mathcal{S}_m \subset \mathcal{S}$ of the state space. Then, we set the requirements as M thresholds $(c_1, \dots, c_M) \in [0, 1]^M$ on the average rewards and formulate a constrained Markov decision problem with constraints $V_m(\pi) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{S_t, A_t \sim \pi} \left[\sum_{t=0}^{T-1} r_m(S_t, A_t) \right] \geq c_m$, $m = 1, \dots, M$. The average is taken over the transition probabilities and the probability distribution defining the control policy $\pi(A_t | S_t)$ of the actions of all agents given their states. To avoid the exponential growth of the policy with the number of agents, we impose a separate design $\pi(A_t | S_t) = \prod_{n=1}^N \pi_n(A_{tn} | S_{tn})$, in which each agent decides its action taking only into account its local state and not the state of other agents. It is important to remark that even if the policies are local, the agents are still coupled by their dynamics and joint rewards.

Example 1 *As a practical example of the assignment problem described by (1), consider a team of N robots that are required to monitor $M > N$ areas, visiting each area \mathcal{S}_m a fraction of time c_m . These conflicting requirements may not be achievable by a single agent, particularly when $\sum_{m=1}^M c_m > 1$, and not even by multiple agents acting independently following the same optimal single-agent policy. Therefore, coordination is essential to solve this problem successfully.*

A task reward $r_0(S_t, A_t)$ can be added, together with its corresponding optimization objective $V_0(\pi) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{S_t, A_t \sim \pi} \left[\sum_{t=0}^{T-1} r_0(S_t, A_t) \right]$. This would be useful, for instance, to set a secondary goal, such as spending the least amount of energy on the multi-robot surveillance task in Example 1. However, in the following, we will focus on the primary goal of satisfying the specifications, setting $V_0(\pi) = r_0 = 0$. Specifically, the feasibility problem we aim to solve is

$$P^* = \max_{\pi_1, \dots, \pi_N} 0, \quad \text{subject to } \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{S_t, A_t \sim \pi} \left[\sum_{t=0}^{T-1} r_m(S_t, A_t) \right] \geq c_m, \quad m = 1, \dots, M, \quad (2)$$

with $r_m(S_t, A_t)$ as in (1) and separate policy structure $\pi(A_t | S_t) = \prod_{n=1}^N \pi_n(A_{tn} | S_{tn})$.

3. State augmented Markov decision process

We start this section by defining the Lagrangian associated with problem (2).

$$\mathcal{L}(\pi, \lambda) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{S_t, A_t \sim \pi} \left[\sum_{t=0}^{T-1} \sum_{m=1}^M \lambda_m (r_m(S_t, A_t) - c_m) \right], \quad (3)$$

where in the previous expression $\lambda \in \mathbb{R}_+^M$ are the dual variables. The previous expression can be simplified by defining the following reward parameterized by λ

$$r_\lambda(S_t, A_t) := \sum_{m=1}^M \lambda_m (r_m(S_t, A_t) - c_m). \quad (4)$$

Under the reward structure of (1), standard regularized methodologies are unable to produce feasible policies (Calvo-Fullana et al., 2023, Proposition 1). Hence, we aim to solve problem (2) by a two-step process. First, in an offline training stage, we will learn the family of policies parameterized by λ that solve the regularized version of (2), i.e.,

$$\Pi(\lambda) = \operatorname{argmax}_{\pi_1, \dots, \pi_N} \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{S_t, A_t \sim \pi} \left[\sum_{t=0}^{T-1} r_\lambda(S_t, A_t) \right]. \quad (5)$$

Then, in the deployment stage, each agent will sample from policies $\Pi(\lambda)$, with λ following dual gradient descent. To be formal, let us define the dual function as $d(\lambda) := \max_{\pi} \mathcal{L}(\pi(\lambda), \lambda)$, where $\pi(\lambda) \in \Pi(\lambda)$. For the optimal multipliers, we minimize $d(\lambda)$ over $\lambda_m \geq 0$. By Danskin's Theorem (Danskin (1966)), we substitute the constraints evaluated at the optimal policy for the gradient of $d(\lambda)$ to obtain the dual update $\lambda^{k+1} = [\lambda^k - \eta (V(\pi(\lambda^k)) - c)]_+$, where $[\cdot]_+$ denotes the projection on the non-negative orthant and $c = (c_1, \dots, c_M)^T$. Section 5 will incorporate a truncated stochastic version of these dual dynamics into an augmented MDP with state (S_t, λ^k) in two timescales.

4. Offline training

In this section, we parameterize the policy of each agent by a vector $\theta_n \in \mathbb{R}^{p_n}$ to deal with the continuous augmented state-space $\mathcal{S} \times \mathbb{R}_+^M$. Hence, each agent's action is drawn from a distribution $\pi_{\theta_n}(A_{tn} | S_{tn}, \lambda)$. As in Section 2, the joint probability distribution is $\prod_{n=1}^N \pi_{\theta_n}(A_{tn} | S_{tn}, \lambda)$.

Thus, the policy gradient Theorem [Sutton et al. \(2000\)](#) takes a form that allows for distributed estimation (see [Proposition 1](#)). Before stating this result, we define some key concepts.

For fixed multipliers and parameters $\theta \in \mathbb{R}^p$, where $p = \sum_{n=1}^N p_n$, we define the following occupancy measure for the joint state and actions of all agents, $\rho_{\theta, \lambda}(s, a) = \lim_{t \rightarrow \infty} \mathbb{P}(S_t = s, A_t = a)$. Likewise, $Q_{\lambda}^{\pi_{\theta}}(s, a) = \sum_{t=0}^{\infty} \mathbb{E}_{S_t, A_t \sim \pi_{\theta}} [r_{\lambda}(S_t, A_t) - \mathcal{L}(\pi_{\theta}, \lambda) | S_0 = s, A_0 = a]$, defines the state-action value function, where $r_{\lambda}(s, a)$ and $\mathcal{L}(\theta, \lambda)$ are the functions defined in [\(4\)](#) and [\(3\)](#) respectively. Notice that we omitted the dependency of π_{θ} on the multipliers λ to simplify the notation. We further define the state-action value function for agent n as

$$Q_{n, \lambda}^{\pi_{\theta}}(s, a) = \mathbb{E}_{(S_t, A_t) \sim \rho_{\theta, \lambda}} [Q_{\lambda}^{\pi_{\theta}}(S_t, A_t) | S_{tn} = s, A_{tn} = a]. \quad (6)$$

From the perspective of agent n , this quantity is the average of the expected return, assuming that all other agents are following policy π_{θ} . From this perspective, $Q_{n, \lambda}^{\pi_{\theta}}(s, a)$ is estimated per agent by averaging the global rewards that result from their local states and actions, or modeled as a parametric function of these local entries. We are now in conditions of providing the specific form that the policy gradient Theorem takes under the conditions described in this section.

Proposition 1 *Assume that the policy of each agent is parameterized by a vector $\theta_n \in \mathbb{R}^{p_n}$ and that $A_n \sim \pi_{\theta_n}(\cdot | S_n, \lambda)$. Let $\mathcal{L}(\pi_{\theta}, \lambda)$ and $Q_{n, \lambda}^{\pi_{\theta}}(S_n, A_n)$ be the functions defined in [\(3\)](#) and [\(6\)](#) respectively. Then, it follows that $\nabla_{\theta_n} \mathcal{L}(\pi_{\theta}, \lambda) = \mathbb{E}_{S_n, A_n \sim \rho_{\theta, \lambda}} [\nabla_{\theta_n} \log \pi_{\theta_n}(A_n | S_n, \lambda) Q_{n, \lambda}^{\pi_{\theta}}(S_n, A_n)]$.*

Leveraging on the local gradients in [Proposition 1](#), we optimize a parametric version of [\(5\)](#)

$$(\pi_{\theta_1}, \dots, \pi_{\theta_N}) = \operatorname{argmax}_{\pi_{\theta}} \mathbb{E}_{S_t, A_t \sim \pi_{\theta}} \left[\frac{1}{T_0} \sum_{t=0}^{T_0-1} r_{\lambda}(S_t, A_t) \right]. \quad (7)$$

were the optimization variable was written as $\pi_{\theta} = (\pi_{\theta_1}, \dots, \pi_{\theta_N})$ to reduce notation. Notice that, in addition to the parametric expansion of the policies, we introduced a second theoretical variation by removing the limit and adhering to a finite time horizon. This modification is key for the practical implementation of the algorithm. For convergence analysis, we still link it to our original formulation in [\(2\)](#) using the following assumption.

Assumption 1 *Estimation consistency.* The truncating error in [\(8\)](#) satisfies $\lim_{T_0 \rightarrow \infty} \|e_{T_0}\| = 0$.

$$\mathbb{E} \left[\frac{1}{T_0} \sum_{t=kT_0}^{(k+1)T_0} (r(S_t, A_t) - c) \right] = \mathbb{E} \left[\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} (r(S_t, A_t) - c) \right] + e_{T_0}. \quad (8)$$

Remark Training for [\(7\)](#) requires that all agents interact to learn from rewards [\(4\)](#) that are jointly activated via [\(1\)](#). This training phase is designed to run offline, with multiple agents and a common λ drawn randomly at the start of an episode and fixed until the end of it. Once each policy $\pi_{\theta_n}(A_{tn} | S_{tn}, \lambda)$ in [\(7\)](#) has been trained, its online execution depends only on the local state S_{tn} of the agent, and a common vector λ which varies as agents enter and exit the zones. Thus, the online coordination between agents depends on sharing the multipliers λ only, as described next.




Figure 1: Gossip protocol between two agents n and i communicating through the link ϵ_{ni} . Both agents aim to know if the shaded zone \mathcal{S}_m is occupied at time $\tau = 0$, as defined by $\max\{\mathbb{1}[S_{\tau i} \in \mathcal{S}_m], \mathbb{1}[S_{\tau n} \in \mathcal{S}_m]\}$. Since agent n is in \mathcal{S}_m at time $t = 0$, it knows the actual reward. Hence $\hat{R}_{1,\tau,m,t} = 1$ for all $t \geq 0$. Instead, agent i is outside \mathcal{S}_m at $t = 0$ so that $R_{i,\tau,m,0} = 0$ is incorrect, but is corrected at time $t = 1$ by gossiping from agent n , so that $R_{i,\tau,m,t} = 1$ for all $t \geq 1$.

5. Distributed online execution

Aiming for a practical online implementation of a feasible policy for (2), we introduced a parametric training strategy in (7) and truncated the episode. There are still two practical issues to be addressed before running the policies (7) in the execution phase. These are computing the expectation over the transition dynamics in the dual update, and assuming that all agents can see the same λ^k .

We resort to a stochastic approximation of the gradient to address the first issue. As introduced in Section 3, a dual gradient descent update over the Lagrange multipliers for problem (2) amounts to $\lambda_m^{k+1} = \left[\lambda_m^k - \eta \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{S_t, A_t \sim \pi} \left[\sum_{t=0}^{T-1} (r_m(S_t, A_t) - c_m) \right] \right]_+$. For a data-driven implementation, we run stochastic gradient descent Dvoretzky (1955), dropping the expectation and replacing it with reward samples. As we did for training, we also set a finite horizon T_0 ,

$$\hat{\lambda}_m^{k+1} = \left[\hat{\lambda}_m^k - \frac{\eta}{T_0} \sum_{t=kT_0}^{(k+1)T_0-1} (r_m(S_t, A_t) - c_m) \right]_+. \quad (9)$$

The trade-off for choosing the roll-out horizon T_0 entails keeping the induced error in (8) sufficiently small while ensuring that the multipliers are updated sufficiently often.

We propose a gossiping protocol to solve the second issue. Notice that agents must know the global rewards $r_m(S_t, A_t)$ to compute the update in (9). The following analysis aims for a distributed implementation of (9) in which each agent updates a local copy of $\lambda^{(k)}$ using their own views of the global rewards, which are agreed with the team through local exchanges with neighboring agents over a communication network. To formalize these ideas, consider agents $n \in \mathcal{V} = \{1, \dots, N\}$ as nodes of a *connected* graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with communication links \mathcal{E} such that $e_{n,i} \in \mathcal{E}$ if agents n and i are within direct communication range. Define the neighborhood $\mathcal{N}_n = \{i \in \mathcal{V} : e_{n,i} \in \mathcal{E}\}$ of agents communicating with n , the distance $d_{n,i}^*$ as the minimum number of hops required to communicate two remote nodes n and i , and the graph diameter $d(\mathcal{G}) := \max_{n,m \in \mathcal{V}} d_{n,m}^*$ as the maximum distance between nodes across the graph.

The global occupation of zone \mathcal{S}_m at time τ takes a binary value $r_m(S_\tau, A_\tau) \in \{0, 1\}$. Agent n computes a sequence of local estimates of $r_m(S_\tau, A_\tau)$ for times $t \geq \tau$. As shown in Fig. 1, these estimates $\hat{R}_{m,\tau,n,t}$ are initialized by agent n at time $t = \tau$ using its local occupancy index; i.e., $\hat{R}_{m,\tau,n,\tau} = \mathbb{1}[S_{\tau n} \in \mathcal{S}_m]$, and then agreed with the neighbor agents $i \in \mathcal{N}_n$ including n itself via $\hat{R}_{m,\tau,n,t} = \max_{i \in \mathcal{N}_n \cup \{n\}} \hat{R}_{m,\tau,n,t-1}$ at each subsequent time $t > \tau$. Through this gossip algorithm, binary data $\mathbb{1}[S_{\tau n} \in \mathcal{S}_m]$ percolates across the network via local exchanges, reaching consensus within a time interval no longer the network diameter, as shown next

Lemma 2 *The gossip estimate $\hat{R}_{m,\tau,n,t}$ underestimates $r_m(S_\tau, A_\tau)$; i.e., $\hat{R}_{m,\tau,n,t} \leq r_m(S_\tau, A_\tau)$ with equality if the gossiping time t is large enough so that remote data arrive through local exchanges across the network; i.e., $\hat{R}_{m,\tau,n,t} = r_m(S_\tau, A_\tau)$, for all $n \in \mathcal{V}$ and $t \geq \tau + d(\mathcal{G})$.*

Algorithm 1 Distributed multi-agent dual update and policy execution

for $k = 0, 1, \dots$ **do**

 Set local policy $\pi_n = \pi_\theta(\lambda_n^k)$ as in (7) pre-trained offline using local gradients.

for $t = kT_0 : (k+1)T_0 - 1$ **do**

 Take action, transition, collect local rewards $\mathbb{1}\{S_{tn} \in \mathcal{S}_m\}$

 Network gossiping: Update $\hat{R}_{m,\tau,n,t}$ according to (10)

end

Compute the stochastic dual gradients

$$g_{nmk} = \frac{1}{T_0} \sum_{\tau=(k-1)T_0}^{kT_0-1} \left(\hat{R}_{m,\tau,n,(k+1)T_0} - c_m \right), \quad g_{nm(k+1)} = \frac{\eta}{T_0} \sum_{\tau=kT_0}^{(k+1)T_0-1} \left(\hat{R}_{m,\tau,n,(k+1)T_0} - c_m \right)$$

 Update the multipliers $\lambda_{m,n}^k = [\lambda_{m,n}^{k-1} - \eta g_{nmk}]_+$, $\lambda_{m,n}^{k+1} = [\lambda_{m,n}^k - \eta g_{nm(k+1)}]_+$
end

This lemma is to be used for convergence analysis, but also for algorithm construction, since it confirms that the estimates $\hat{R}_{m,\tau,n,t}$ do not need to be updated beyond $t \geq \tau + d$, where $d \geq d(\mathcal{G})$ is any overestimate of the network diameter. Thus, putting $t = \tau$, $t \in \{\tau + 1, \tau + d - 1\}$ and $t > \tau + d$ together, we update the distributed reward estimates as

$$\hat{R}_{m,\tau,n,t} = \begin{cases} \hat{R}_{m,\tau,n,t-1}, & \tau = (k-1)T_0, \dots, t - d - 1 \\ \max_{p \in \mathcal{N}_n \cup \{n\}} \hat{R}_{m,\tau,p,t-1}, & \tau = t - d, \dots, t - 1 \\ \mathbb{1}\{S_{nt} \in \mathcal{S}_m\}, & \tau = t \end{cases} \quad (10)$$

We can substitute them into the dual update (9) to obtain a distributed algorithm for the multipliers. Algorithm 1 not only indicates how to update the local estimates $\lambda_{m,n}$ via network gossiping of the rewards but also gives each agent the procedure to run its fully distributed policy $\pi_\theta(\lambda_n^k)$ in the execution phase. The reason to keep two copies of $\lambda_{m,n}$ by Algorithm 1 is to wait until agents reach consensus on the rewards via Lemma 2; see the proof of Proposition 3 for more details. In the proof of Lemma 2 is shown that $\hat{\lambda}_{m,n,t}^k$ reaches steady state and can be discarded once $t \geq t_{k+1}$ since, according to (12), it is not needed for recalculating $\hat{\lambda}_{m,n}^{k+1}$ any more.

Next, we show that running the gossiping and distributed policies ensures that the constraints in (2) are satisfied almost surely. Building up to this main result, we first characterize the multipliers mismatch, and then add two working assumptions.

Proposition 3 *Let $T_{k+1} := T_0(k+1)$ the time at the end of rollout $k+1$. Assume that $T_0 \geq d(\mathcal{G})$. Then, the error on the multipliers is bounded by $\lambda_m^{k+1} - \lambda_{m,n}^{k+1} \leq \frac{\eta}{T_0} d(\mathcal{G})$*

Assumption 2 No repelling forces *The underlying Markov decision process is such that, given m , there exists a feasible policy π_n satisfying $r_m(S_t, A_t) = 1$ for all t .*

The previous assumption implies that an agent can be stationed at a particular zone $S_{nt} \in \mathcal{S}_m$. This will be used in the proof of Lemma 6.

Assumption 3 Representation capacity The policy parameterization is dense enough; i.e., $\exists \beta > 0$, such that for each $\pi(\lambda) \in \Pi(\lambda)$ there is a $\pi_\theta(\lambda)$ satisfying $\left| \sum_{m=1}^M \lambda^T (V_m(\pi(\lambda)) - V_m(\pi_\theta(\lambda))) \right| \leq \beta \|\lambda\|$, where $V_m(\pi(\lambda))$ and $V_m(\pi_\theta(\lambda))$ are the values resulting from $\pi(\lambda)$ and $\pi_\theta(\lambda)$, respectively.

Theorem 4 Let Assumptions 1–3 hold. Consider specifications $c_{\max} = \max_{m=1, \dots, M} c_m < 1$, and $\sum_{m=1}^M c_m \leq N - 1$ so that $\delta_c = (1 - c_{\max})/\sqrt{M} > 0$. If $\beta + \frac{M}{T_0} d(\mathcal{G})\eta + \epsilon_{T_0} + \eta/2 < \delta_C$, the trajectories generated by Algorithm 1 are feasible with probability one, i.e., for all $m = 1, \dots, M$

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} r_m(S_t, A_t) = c_m, \text{ a.s.} \quad (11)$$

The condition in Theorem 4 can be satisfied by designing a sufficiently powerful parameterization or neural network, a sufficiently long time horizon, and then choosing a small step-size.

6. Convergence analysis

Define $\hat{g}^k = \frac{1}{T_0} \sum_{\tau=(k-1)T_0}^{kT_0-1} (r(S_\tau, A_\tau) - c)$. This vector contains the average constraint violation during execution between two updates of the Lagrange multipliers λ^k . In the following lemma, we establish that \hat{g}^k is in expectation similar to the gradient of the dual function.

Lemma 5 With \mathcal{F}_k being the filtration for the process \hat{g}^k , assumptions 1–3 ensure that for every $\epsilon_{T_0} > 0$, there exists a T_0 such that $(\lambda^k)^T (\nabla_\lambda d(\lambda^k) - \mathbb{E}_{\pi_\theta} [\hat{g}^k | \mathcal{F}_k]) \leq \|\lambda^k\| \left(\beta + \frac{M}{T_0} d(\mathcal{G})\eta + \epsilon_{T_0} \right)$.

Proof Write $(\lambda^k)^T (\nabla_\lambda d(\lambda^k) - \mathbb{E}_{\pi_\theta} [\hat{g}^k | \mathcal{F}_k]) = (\lambda^k)^T (V - V_N + V_N - V_{T_0} + V_{T_0} - V_\theta)$ starting from $V = V(\pi_1(\lambda^k), \dots, \pi_N(\lambda^k)) = \nabla_\lambda d(\lambda^k)$ with policies in (5), then adding and subtracting $V_N = V(\pi_1(\lambda_1^k), \dots, \pi_N(\lambda_N^k))$ such that the agents optimize using distributed local copies of λ^k , and V_{T_0} where additionally the objective is finite horizon. The value $V_\theta = \mathbb{E}_{\pi_\theta} [\hat{g}^k | \mathcal{F}_k] = V(\pi_{\theta_1}(\lambda_1^k), \dots, \pi_{\theta_N}(\lambda_N^k))$ on the right involves the algorithmic policies such that the agents optimize parametric policies over a finite horizon with distributed copies of λ^k . The proof follows from the bounds in Assumption 3, Proposition 3, and Assumption 1. For the difference $(\lambda^k)^T (V - V_N)$ we need to accumulate the bounds in Proposition 3 per zone, and then use Lipschitz continuity with respect to the multipliers, with unit constant. The continuity can be argued as in the proof of Lemma 6 using Assumption 2 and noticing that when the multipliers change their order the assignment is altered, and the difference in the weighted value is equal to the difference of the multipliers. ■

Lemma 6 Assumption 2 with δ_C as in Theorem 4 results in $\sum_{i=1}^M \lambda_m (V_m - c_m) \geq \delta_c \|\lambda\|$.

Proof Without loss of generality, let us sort the multipliers $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_M$. Then, one policy that optimizes the Lagrangian would sit one agent at each zone $m = 1 \dots, N$. Therefore $V_m = 1$ for $m \leq N$ and $V_m = 0$ for $m > N$. Hence $\sum_{m=1}^M \lambda_m (V_m - c_m) = \lambda_1 (1 - c_1) + \sum_{m=2}^N \lambda_m (1 - c_m) + \sum_{m=N+1}^M \lambda_m (0 - c_m) \geq (1 - c_{\max})\lambda_1 + \sum_{m=2}^M \lambda_n (1 - c_m) + \sum_{m=N+1}^M \lambda_n (0 - c_m) = (1 - c_{\max})\|\lambda\|_\infty + \left(N - 1 - \sum_{m=2}^M c_m \right) \lambda_n \geq (1 - c_{\max})\|\lambda\|_\infty \geq (1 - c_{\max}) \frac{\|\lambda\|}{\sqrt{M}}$. ■

With this result, we establish conditions in the following lemma that the norm of the multipliers is a supermartingale. This will be instrumental in proving that the multipliers are bounded, and from there, that the constraints are satisfied.

Lemma 7 *If $\beta + \frac{M}{T_0}d(\mathcal{G})\eta + \epsilon_{T_0} + \eta/2 < \delta_C$ and $\|\lambda^k\| \geq 1$, then $\|\lambda^k\|^2$ is a supermartingale.*

Proof Using the update of the multipliers given in (9) and using the non-expansiveness of the projection, it follows that $\Lambda^{k+1} \leq \|\lambda^k\|^2 - 2\eta(\lambda^k)^T \hat{g}^k + \eta^2 \|\hat{g}^k\|^2 \leq \Lambda^k - 2\eta(\lambda^k)^T \hat{g}^k + \eta^2$. where we defined $\Lambda^k = \|\lambda^k\|^2$ and used $\|\hat{g}^k\|^2 \leq 1$. Therefore, to ensure that the process Λ^k is a super-martingale, we need to prove that $\mathbb{E} \left[(\lambda^k)^T \hat{g}^k | \mathcal{F}_k \right] = (\lambda^k)^T \mathbb{E} [\hat{g}^k | \mathcal{F}_k] > \eta/2$, where the expectation is taken with respect to the distributed policy (7) being executed. To this end, consider $(\lambda^k)^T \mathbb{E} [\hat{g}^k | \mathcal{F}_k] = (\lambda^k)^T (\nabla d(\lambda_k) + \mathbb{E} [\hat{g}^k | \mathcal{F}_k] - \nabla d(\lambda_k)) \geq (\lambda^k)^T \nabla d(\lambda_k) - \left(\beta + \frac{M}{T_0}d(\mathcal{G})\eta + \epsilon_{T_0} \right) \|\lambda^k\|$, where the inequality follows from the Cauchy-Schwartz inequality and the bound in Lemma 5. Combining the previous inequality along with the result of Lemma 5, and the hypothesis of this lemma, i.e., $\beta + \frac{M}{T_0}d(\mathcal{G})\eta + \epsilon_{T_0} + \eta/2 < \delta_c$, it follows for all λ^k $(\lambda^k)^T \mathbb{E}_{\pi^*} [\hat{g}^k | \mathcal{F}_k] \geq \eta/2 \|\lambda^k\| \geq \eta/2$, where we use the hypothesis $\|\lambda^k\| \geq 1$. ■

Proof [of Theorem 4] Due to space limitations we provide a sketch of the proof. Since $\Lambda^k = \|\lambda^k\|^2$ is a supermartingale, which is a generalization of a non-increased sequence, the multipliers must be bounded. More formally, the sequence $p(\lambda^k | \lambda^0)$ is *tight*. That is, for any $\delta > 0$, it exists a compact set \mathcal{K}_δ such that for every $k \geq 0$, $\mathbb{P}(\lambda^k \in \mathcal{K}_\delta) > 1 - \delta$. For simplicity in the sketch of this proof, we will consider bounded multipliers instead of tight ones and show how this implies feasibility. For a similar argument on how Λ^k being a supermartingale implies tightness and feasibility see Lemma 4 and Proposition 3 in Calvo-Fullana et al. (2023). Following the case of bounded multipliers, recall the dual descent update in (9) $\hat{\lambda}_m^{k+1} = \left[\hat{\lambda}_m^k - \frac{\eta}{T_0} \sum_{t=kT_0}^{(k+1)T_0} (r_m(S_t, A_t) - c_m) \right]_+ \geq \hat{\lambda}_m^k - \frac{\eta}{T_0} \sum_{t=kT_0}^{(k+1)T_0} (r_m(S_t, A_t) - c_m)$. Applying this inequality recursively, it follows that $\hat{\lambda}_m^{k+1} \geq \hat{\lambda}_m^0 - \frac{\eta}{T_0} \sum_{t=0}^{(k+1)T_0} (r_m(S_t, A_t) - c_m)$. Taking the limsup in both sides of the previous inequality, and using that the multipliers are bounded it implies that $\limsup_{k \rightarrow \infty} - \sum_{t=0}^{(k+1)T_0} (r_m(S_t, A_t) - c_m) = L < \infty$. Defining $T = (k+1)T_0$ the previous limit is equivalent to the desired (11). ■

7. Numerical experiments

We consider a simulated experiment with $N = 2$ agents monitoring the $M = 4$ circular zones in Fig. 2(e). The position S_{tn} of each agent evolves as $S_{(t+1)n} = S_{tn} + T_s A_{tn}$ with $T_s = 0.5$ within the area $\mathcal{S} = [0, 10] \times [0, 10]$ measured in squared meters. At time t , a reward $r_m(S_t, A_t) = 1$ is obtained by the team if there is an agent inside the m -th region and $r_m(S_t, A_t) = 0$ otherwise. By specification, each region must be visited at least 30% of the time. As the requirements add up to more than 100% of the time, a single agent cannot satisfy them by acting alone. The agent policies are parameterized as linear combinations of Gaussian kernels and trained offline, running 150,000 episodes of Actor-Critic updates using the local gradients in Proposition 1. Then, in the execution phase, the agents gossip their rewards within a communication range of 2.5 meters, coordinating via Algorithm 1 during 200,000 time steps, and updating the multipliers every $T_0 = 1,000$ steps.

In this first experiment, we weaken the assumption of a connected graph, allowing for stochastic links. The agents asymptotically satisfy the requirements for all regions (Fig. 2(c)), despite being within communication range only 30% of the time (Fig. 2(d)). Coordination is achieved via the dual variables, which differ for each agent, alternating as they incorporate the rewards collected directly

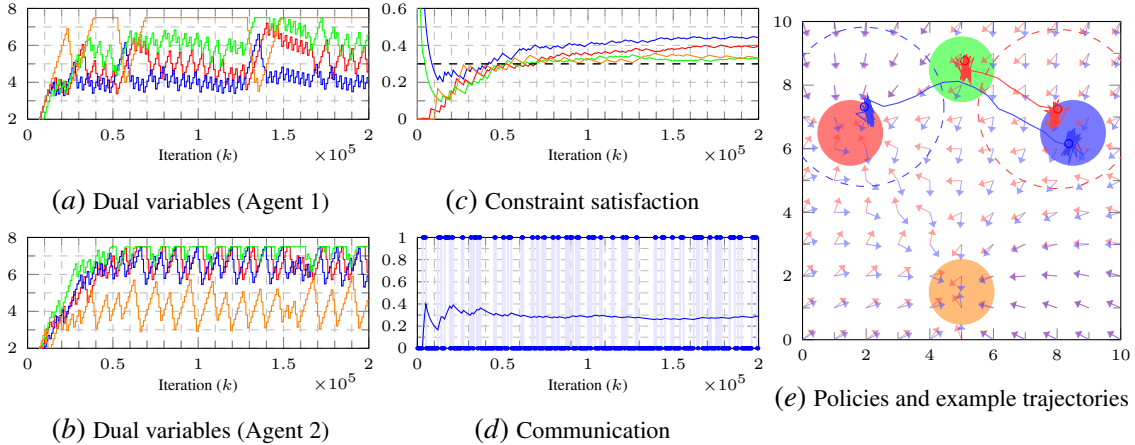


Figure 2: Simulation results obtained after executing for 200,000 iterations, a two-agent policy trained to monitor the regions shown in (e), with requirements $c = [0.3, 0.3, 0.3, 0.3]$. Colors in (a)–(c) correspond to the matching colored region in (e). A subset policy for $\lambda = [5, 2.5, 0, 5]$ and 100 steps of the runtime trajectory corresponding to $k = (128,00, 128,100)$ are shown in (e), where the communication range of agents is shown by dashed lines.

or via gossiping. Fig. 2(e) shows the policies of the two agents as a function of the spatial variable for a fixed $\lambda = [5, 2.5, 0, 5]$. Both red and orange regions produce the same weighted rewards ($\lambda_1 = \lambda_4 = 5$), yet the agents are directed to different regions, acquiring both rewards simultaneously. The trajectories of the agents are also shown between iterations $k = (128,000, 128,100)$. Here, the red agent moves from the right region to the center region, while the blue agent moves from the left region to the right region. As this movement occurs, the agents enter the communication range, exchanging their stored rewards and updating their local dual variables (Fig. 2(a)–2(b)). This update causes the blue agent to avoid the middle region (occupied now by the red agent) and to continue to the region on the right, thus acquiring a higher reward as a team.

Next, we include two additional simulations as steps towards our future work, which aims to implement Algorithm 1 in a real-life environment with autonomous vehicles. In the first experiment, shown in Fig. 3(a), we tested the policy trained for the experiment in Fig. 2(e) in a Gazebo world, under simulated dynamics of two TurtleBot robots, including collision-avoidance. For that purpose, we coded the distributed Algorithm 1 as a ROS2 node communicating with the TurtleBot navigator. The second experiment in Fig. 3(b) shows the trajectories of two agents that use our multi-agent Algorithm 1 in an environment describing an actual floorplan, with an L-shaped corridor connecting three offices and three labs. In this case, the agents had to address more complex dynamics and obstacles to reach the marked areas in this floorplan. But part of these dynamics is handled by a low-level SLAM and navigation controller capable of guiding each robot to a desired destination, avoiding collisions on the way. Thus, our policies only need to assign the $M = 3$ zones to the $N = 2$ agents via actions in a finite space $A_{tn} \in \{1, \dots, M\}$. The next state observed by the policy is the position reached by the agent after navigating $T_s = 1s$ towards the goal. At this moment, the assignment may be changed by the policy, directing the car to an alternate zone. Fig. 3(c) shows how constraints $c = [0.25, 0.25, 0.25, 0.25]$ and $c = [0.5, 0.4, 0.3]$, which are unattainable by a single agent, can still be satisfied by our multi-agent assignment policies in these more realistic cases, reaching the specified time requirements for all the zones.

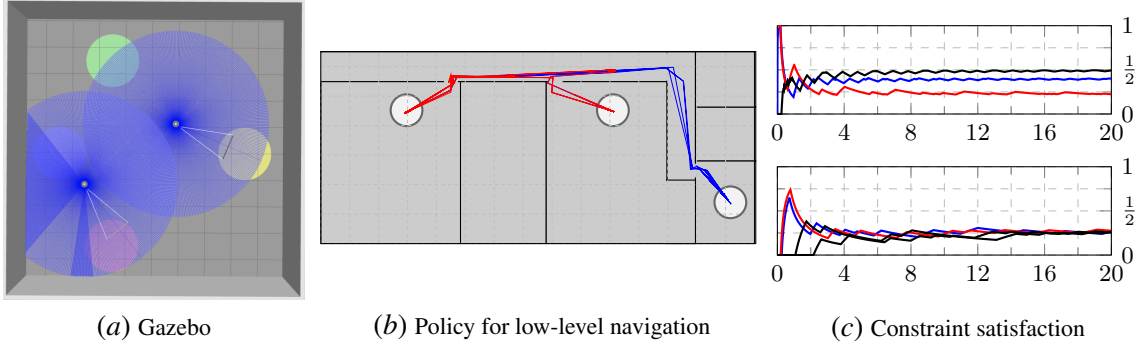


Figure 3: (a) ROS2 Implementation of the multi-agent assignment Algorithm 1 in a Gazebo TurtleBot environment; (b) floorplan patrol with low-level navigation control; (c) constraint satisfaction as a function of the time step (in thousands) for the floorplan (top) and the TurtleBot (bottom).

8. Conclusion

We showcased the efficacy of constrained RL in handling conflicting requirements via a state-augmented approach with oscillating dual variables. By introducing a multi-agent system that coordinates through gossiped binary variables, we offer a solution for multi-agent assignment establishing almost sure feasibility guarantees in Theorem 4, corroborated in a monitoring numerical experiment involving communication intermittence and realistic robot navigation dynamics.

Appendix

Proof [Proposition 1] The structure $\pi_\theta(A_t | S_t) = \prod_{n=1}^N \pi_{\theta_n}(A_{tn} | S_{tn})$ yields $\nabla_{\theta_n} \log \pi_\theta(S_t, A_t) = \nabla_{\theta_n} \log \pi_{\theta_n}(A_{tn} | S_{tn})$. Substituted in the policy gradient Sutton et al. (2000) yields $\nabla_{\theta_n} \mathcal{L}(\pi_\theta, \lambda) = \mathbb{E}_{(S_t, A_t) \sim \rho_{\theta, \lambda}} [\nabla_{\theta_n} \log \pi_{\theta_n}(A_{tn} | S_{tn}) Q_\lambda^{\pi_\theta}(S_t, A_t)]$. By the law of total expectation on S_{tn}, A_{tn} , $\nabla_{\theta_n} \mathcal{L}(\pi_\theta, \lambda) = \mathbb{E}_{(S_{tn}, A_{tn})} \left[\mathbb{E}_{(S_t, A_t) \sim \rho_{\theta, \lambda}} [\nabla_{\theta_n} \log \pi_{\theta_n}(A_{tn} | S_{tn}) Q_\lambda^{\pi_\theta}(S_t, A_t) | S_{tn}, A_{tn}] \right]$ $= \mathbb{E}_{S_{tn}, A_{tn}} \left[\nabla_{\theta_n} \log \pi_{\theta_n}(A_{tn} | S_{tn}) \mathbb{E}_{(S_t, A_t) \sim \rho_{\theta, \lambda}} [Q_\lambda^{\pi_\theta}(S, A) | S_{tn}, A_{tn}] \right]$ and substitute (6). ■

Proof [Lemma 2] If no $S_{n\tau} \in \mathcal{S}_m$ at time τ then $\hat{R}_{m, \tau, n, \tau} = 0$ for all $n \in V$ and thus $\hat{R}_{m, \tau, n, \tau} = r_{m, \tau} = 0$ for all $t \geq \tau$ and for all $n \in V$. Otherwise, if it exists $n \in V$ such that $S_{n, \tau} \in \mathcal{S}_m$, then $\hat{R}_{m, \tau, n, \tau} = r_{m, \tau} = 1$ for all $t \geq \tau$. Hence $\hat{R}_{m, \tau, i, \tau+1} = r_{m, \tau} = 1$ for all $i \in \mathcal{N}_n$, and by induction $\hat{R}_{m, \tau, i, \tau+h} = r_{m, \tau} = 1$ for all h -hop neighbors of n , from which the consensus follows. ■

Proof [Proposition 3] Consider $\hat{\lambda}_{m, n, t}^{k+1} = \left[\hat{\lambda}_{m, n, t}^k - \frac{\eta}{T_0} \sum_{\tau=kT_0}^{(k+1)T_0} (\hat{R}_{m, \tau, n, t} - c_m) \right]_+$ as in Algorithm 1. Although this algorithm indicates updating the whole set of multipliers for all k at each time step t , a simplification can be done using Lemma 2. Indeed, the result in Lemma 2 ensures that the rewards $\hat{R}_{m, \tau, n, t}$ reach its steady state value $r_{m, \tau}$ when $t \geq \tau + d(\mathcal{G})$, so that $\hat{\lambda}_{m, n, t}^k = \hat{\lambda}_m^k$ for all $t \geq t_{k+1} := kT_0 + d(\mathcal{G})$. Thus, the dual update can be simplified to

$$\hat{\lambda}_{m, n, t}^{k+1} = \begin{cases} \left[\hat{\lambda}_{m, n, t}^k - \frac{\eta}{T_0} \sum_{\tau=kT_0}^{(k+1)T_0} (\hat{R}_{m, \tau, n, t} - c_m) \right]_+, & \text{for } t = kT_0, \dots, t_{k+1} - 1 \\ \hat{\lambda}_{m, n, t_{k+1}}^{k+1} = \hat{\lambda}_m^{k+1} & \text{for } t \geq t_{k+1} \end{cases} \quad (12)$$

With $t_{k+1} - kT_0 = d(\mathcal{G})$ binary $R_{m, \tau, n, t}$ in (12), the error in $\hat{\lambda}_{m, n, t}^{k+1}$ is $\eta \frac{d(\mathcal{G})}{T_0}$ at most. ■

Acknowledgments

This work was supported in part by Spain’s Agencia Estatal de Investigación under grant RYC2021-033549-I and the María de Maeztu Strategic Research Program under grant CEX2021-001195-M.

References

- Tuncer Can Aysal, Mehmet Ercan Yildiz, Anand D Sarwate, and Anna Scaglione. Broadcast gossip algorithms for consensus. *IEEE Transactions on Signal processing*, 57(7):2748–2761, 2009.
- Dimitri P Bertsekas. A new algorithm for the assignment problem. *Mathematical Programming*, 21(1):152–171, 1981.
- Noam Brown and Tuomas Sandholm. Superhuman ai for multiplayer poker. *Science*, 365(6456):885–890, 2019.
- Miguel Calvo-Fullana, Santiago Paternain, Luiz FO Chamon, and Alejandro Ribeiro. State augmented constrained reinforcement learning: Overcoming the limitations of learning with rewards. *IEEE Transactions on Automatic Control*, 2023.
- Smriti Chopra, Giuseppe Notarstefano, Matthew Rice, and Magnus Egerstedt. A distributed version of the hungarian method for multirobot assignment. *IEEE Transactions on Robotics*, 33(4):932–947, 2017.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- John M Danskin. The theory of max-min, with applications. *SIAM Journal on Applied Mathematics*, 14(4):641–664, 1966.
- Aryeh Dvoretzky. *On stochastic approximation*. Mathematics Division, Office of Scientific Research, US Air Force, 1955.
- Yuanqi Gao, Wei Wang, Jie Shi, and Nanpeng Yu. Batch-constrained reinforcement learning for dynamic distribution network reconfiguration. *IEEE Transactions on Smart Grid*, 11(6):5357–5369, 2020.
- Marcus Hutter. Self-optimizing and pareto-optimal policies in general environments based on bayes-mixtures. In *International Conference on Computational Learning Theory*, pages 364–379. Springer, 2002.
- Gregory Kahn, Adam Villaflor, Vitthay Pong, Pieter Abbeel, and Sergey Levine. Uncertainty-aware reinforcement learning for collision avoidance. *arXiv preprint arXiv:1702.01182*, 2017.
- Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- Hepeng Li, Zhiqiang Wan, and Haibo He. Constrained ev charging scheduling based on safe deep reinforcement learning. *IEEE Transactions on Smart Grid*, 11(3):2427–2439, 2019.

- G Ayorkor Mills-Tettey, Anthony Stentz, and M Bernardine Dias. The dynamic hungarian algorithm for the assignment problem with changing costs. *Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-07-27*, 2007.
- Sobhan Miryoosefi, Kianté Brantley, Hal Daume III, Miro Dudik, and Robert E Schapire. Reinforcement learning with convex constraints. *Advances in neural information processing systems*, 32, 2019.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744, 2022.
- Santiago Paternain, Luiz Chamon, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained reinforcement learning has zero duality gap. *Advances in Neural Information Processing Systems*, 32, 2019.
- Santiago Paternain, Miguel Calvo-Fullana, Luiz FO Chamon, and Alejandro Ribeiro. Safe policies for reinforcement learning via primal-dual methods. *IEEE Transactions on Automatic Control*, 68(3):1321–1336, 2022.
- Rodrigo Porteiro, Fernando Paganini, and Juan Andres Bazerque. Combined learning and optimal power flow for storage dispatch in grids with renewables. In *IEEE Conference on Innovative Smart Grid Technologies (ISGT)*, 2024.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Adv. in neural information processing systems*, pages 1057–1063, 2000.
- Chen Tessler, Daniel J Mankowitz, and Shie Mannor. Reward constrained policy optimization. *arXiv preprint arXiv:1805.11074*, 2018.
- Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of reinforcement learning and control*, pages 321–384, 2021.