



Distributed Safe Learning using an Invariance-based Safety Framework

Conference Paper**Author(s):**

Carron, Andrea ; Sieber, Jerome ; Zeilinger, Melanie

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000454950>

Rights / license:

[Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International](#)

Originally published in:

IFAC-PapersOnLine 54(9), <https://doi.org/10.1016/j.ifacol.2021.06.067>

Funding acknowledgement:

SEED-19 18-2 - Collaborative Exploration-Exploitation: Distributed Decision-making and Estimation in Robotic Networks (ETHZ)

Distributed Safe Learning using an Invariance-based Safety Framework^{*}

Andrea Carron^{*} Jerome Sieber^{*} Melanie N. Zeilinger^{*}

^{*} *Institute for Dynamic Systems and Control, ETH Zurich, CH-8092 Zurich, Switzerland (e-mail: carrona@ethz.ch; jsieber@ethz.ch; mzeilinger@ethz.ch).*

Abstract: In large-scale networks of uncertain dynamical systems, where communication is limited and there is a strong interaction among subsystems, learning local models and control policies offers great potential for designing high-performance controllers. At the same time, the lack of safety guarantees, here considered in the form of constraint satisfaction, prevents the use of data-driven techniques to safety-critical distributed systems. This paper presents a safety framework that guarantees constraint satisfaction for uncertain distributed systems while learning. The framework considers linear systems with coupling in the dynamics and subject to bounded parametric uncertainty, and makes use of robust invariance to guarantee safety. In particular, a robust non-convex invariant set, given by the union of multiple ellipsoidal invariant sets, and a nonlinear backup control law, given by the combination of multiple stabilizing linear feedbacks, are computed offline. In presence of unsafe inputs, the safety framework applies the backup control law, preventing the system to violate the constraints. As the robust invariant set and the backup stabilizing controller are computed offline, the online operations reduce to simple function evaluations, which enables the use of the proposed framework on systems with limited computational resources. The capabilities of the safety framework are illustrated by three numerical examples.

Copyright © 2021 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Networked Control Systems, Linear Systems, Safe Learning.

1. INTRODUCTION

Large-scale safety-critical distributed systems are challenging to model and control due to their complexity, the interactions between subsystems, and their limited communication capabilities. Examples include power networks, automated warehouses, or vehicle platoons. Data-driven approaches for learning and control are appealing to overcome these challenges, and have shown superior performance compared to traditional control design methodologies via their capability of adapting the model and controller from data, see e.g. Busoni et al. (2008); Tuyls and Weiss (2012); Nguyen et al. (2018). However, despite the promising performance shown in research-driven applications Kar et al. (2013a,b); Varshavskaya et al. (2009), and Zhang et al. (2018), the applicability of data-driven techniques to safety-critical systems is severely limited by the lack of theoretical guarantees on the system behavior. Using learning-based controllers in closed-loop may lead the system to unwanted and dangerous behavior, e.g., in the form of constraint violation and instability. While safe learning algorithms for single systems have been proposed, e.g. in Gillula and Tomlin (2011); Berkenkamp and Schoellig (2015); Wabersich and Zeilinger (2018b); Wabersich et al. (2019); Fisac et al. (2019), the problem of guaranteeing constraint satisfaction while learning in multi-

agent systems has received less attention, see e.g. Larsen et al. (2017), Muntwiler et al. (2019), and Chen et al. (2019). This paper introduces a distributed framework for safe learning, in the form of an algorithm that prevents constraint violation at all times based on local conditions and controllers. We consider constrained distributed linear systems subject to bounded parametric uncertainties, and employ structured robust invariance at the core of the framework. The set of safe states is defined as the union of multiple ellipsoidal robust invariant sets, where each ellipsoidal robust invariant set is associated with a stabilizing linear feedback law. An input is considered unsafe if the one-step prediction leads the system state outside of the safe region. In the presence of unsafe inputs, the safety framework applies a combination of backup control laws to prevent the system from violating state and input constraints. Both, ellipsoidal sets and linear feedback laws are jointly precomputed offline via a distributed convex optimization problem. This reduces the online operation to a simple function evaluation, rendering the approach, feasible for real-time application.

Related Work: In recent years, the interest in safe learning has grown, see e.g. García and Fernández (2015) for a comprehensive overview in the context of reinforcement learning. Reachability analysis for dynamic systems is adopted in Gillula and Tomlin (2011) and Gillula and Tomlin (2012) to construct a safe set and backup controller. Safety frameworks based on Model Predictive Control (MPC) have been presented in Wabersich and Zeilinger (2018a,b);

^{*} The research of A. Carron was supported by the Swiss National Centre of Competence in Research NCCR Digital Fabrication. The research of J. Sieber was supported by the ETH Career Seed Grant 19-18-2.

Wabersich et al. (2019), and Koller et al. (2018), where robust and stochastic MPC formulations are used to guarantee constraint satisfaction at all times.

Safety frameworks for distributed linear systems with additive uncertainty were introduced in Larsen et al. (2017) and Muntwiler et al. (2019). The work in Larsen et al. (2017) is based on the computation of structured ellipsoidal robust invariant sets. The approach presented in Muntwiler et al. (2019) is instead based on a distributed MPC formulation, which is especially beneficial if the optimal operation point of the network is close to the border of the state constraints, but at the cost of a higher computational burden. A multi-agent Lyapunov-based safe feedback is also presented in Wieland et al. (2007), and the proposed scheme is applied to a collision avoidance problem. The safety framework in Chen et al. (2019) exploits control barrier functions and robust invariant sets to guarantee constraint satisfaction, where the dynamic interaction among subsystems is reduced using assume-guarantee contracts.

Contributions: The contribution of this paper is twofold. First, we propose a safety framework that makes use of robust invariant sets and stabilizing linear feedback controllers for distributed constrained linear systems with parametric uncertainty. The computation of the invariant sets is based on the work presented in Conte et al. (2016). Differently from the distributed safety framework introduced in Larsen et al. (2017), we consider parametric uncertainty instead of additive uncertainty, and make use of multiple ellipsoidal robust invariant sets, thereby reducing conservatism. Moreover, the optimization problem for computing the invariant sets is a linear semidefinite program (SDP) instead of a bilinear SDP. This renders the optimization problem free of tuning parameters and allows the use of standard SDP solvers, such as MOSEK or SDPT3, see MOSEK ApS (2019) and Tütüncü et al. (2003), respectively. The second contribution is an explicit description of the safe set through the union of ellipsoidal invariant sets. Differently from Muntwiler et al. (2019), the proposed algorithm moves most of the computational burden offline, and the online operations reduce to simple function evaluations, rendering the approach feasible for real-time application.

The remainder of the paper is organized as follows: Section 2 introduces preliminaries and the problem formulation. The distributed safety framework is presented in Section 3, and Section 4 discusses the numerical results. Finally, Section 5 concludes the paper.

2. PRELIMINARIES

2.1 Notation

Let $\text{col}(x_i)$ define a column vector obtained by stacking the vectors x_i . The symbol \geq (\leq) used with matrices indicates a positive (negative) semi-definite matrix. The notation $[\cdot]_l$ applied to a vector or matrix, denotes the l -th row of the vector or matrix. Given a matrix $A \in \mathbb{R}^n$, the trace of A is indicated with $\text{tr}(A)$.

2.2 Problem Formulation

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected communication graph, where $\mathcal{V} = \{1, \dots, N\}$ represents the nodes of the network, and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ contains the unordered pairs of nodes $\{i, j\}$ which can communicate with each other. We denote with $\mathcal{N}_i = \{j \mid i, j \in \mathcal{E}\} \subseteq \mathcal{V}$ the set of neighbors of node i . We associate with each node $i \in \mathcal{V}$ a discrete-time dynamical linear system with parametric uncertainty

$$x_{i,k+1} = A_i(\theta_i)x_{\mathcal{N}_i,k} + B_i(\theta_i)u_{i,k}, \quad (1)$$

where $x_i \in \mathbb{R}^{n_i}$, $x_{\mathcal{N}_i} \in \mathbb{R}^{n_{\mathcal{N}_i}}$, $u_i \in \mathbb{R}^{m_i}$ are the state of agent i , the state of agent i and its neighbors, and the input of agent i , respectively. The uncertain vector θ_i belongs to a polytopic set $\Theta_i \subset \mathbb{R}^{p_i}$, i.e. $\theta_i \in \Theta_i$. Each subsystem (1) is subject to polytopic state and input constraints containing the origin in their interior

$$\begin{aligned} \mathbb{X}_i &= \{x_{\mathcal{N}_i} \mid H_i x_{\mathcal{N}_i} \leq h_i\}, \\ \mathbb{U}_i &= \{u_i \mid O_i u_i \leq o_i\}, \end{aligned} \quad (2)$$

where $H_i \in \mathbb{R}^{n_{h_i} \times n_{\mathcal{N}_i}}$, $h_i \in \mathbb{R}^{n_{h_i}}$, $O_i \in \mathbb{R}^{n_{o_i} \times m_i}$, and $o_i \in \mathbb{R}^{n_{o_i}}$. Note that the state constraints can couple neighboring subsystems, while the input constraints only act locally. The dynamics of the overall network is given by

$$x_{k+1} = A(\theta)x_k + B(\theta)u_k, \quad (3)$$

where $x = \text{col}(x_i)$, $u = \text{col}(u_i)$, and $\theta = \text{col}(\theta_i)$ represent the state, the input, and the uncertain parameter vector, respectively. The matrix $A \in \mathbb{R}^{n \times n}$ is an appropriate composition of the blocks A_i , and the input matrix is block diagonal, i.e., $B = \text{diag}(B_1, \dots, B_N)$. The state constraints for system (3) are given by $\mathbb{X} = \{x \mid x_{\mathcal{N}_i} \in \mathbb{X}_i \forall i \in \mathcal{V}\}$. Taking the cartesian product of all local input constraints, we denote the global input constraint set as $\mathbb{U} = \mathbb{U}_1 \times \dots \times \mathbb{U}_N$. Additionally, we denote with $T_i \in \{0, 1\}^{n_i \times n}$ and $W_i \in \{0, 1\}^{n_{\mathcal{N}_i} \times n}$ two lifting matrices such that $x_i = T_i x$ and $x_{\mathcal{N}_i} = W_i x$.

In the context of this paper, we refer to a learning algorithm as a generic procedure used to either retrieve information from data during closed-loop operation, e.g., to learn a model, see e.g. Chiuso and Pillonetto (2019), or to train learning-based controllers, e.g., via Q-learning, Watkins and Dayan (1992), or Proximal Policy Optimization (PPO), Schulman et al. (2017). The learning algorithm would apply an input u_i^L to each subsystem in order to achieve its goal. Most available learning algorithms, however, do not guarantee the satisfaction of state and input constraints. Constraint satisfaction is particularly challenging and crucial in distributed safety-critical applications, where the interactions among subsystems can render the global system unstable, see Blanchini et al. (2015).

In this paper, we propose a safety framework for distributed systems that can augment any learning algorithm to guarantee satisfaction of state and input constraints. The proposed safety framework verifies if an input is safe to apply and, if it is not, provides a backup control law that guarantees constraint satisfaction for all future times. Formally this is captured by the following definition of a safety framework, a similar definition was introduced in Wabersich and Zeilinger (2018b).

Definition 1. A distributed safety framework, defined by the functions $g_i(x_{\mathcal{N}_i}, u_i) : \mathbb{R}^{n_{\mathcal{N}_i}} \times \mathbb{R}^{m_i} \rightarrow \mathbb{R}^{m_i}$ for

$i \in \mathcal{V}$, certifies an input $u_{i,\bar{k}}^L$ as safe for system (1) at iteration \bar{k} , if $u_{i,\bar{k}}^L = g_i(x_{\mathcal{N}_i,\bar{k}}, u_{i,\bar{k}}^L)$. The application of $u_{i,k} = g_i(x_{\mathcal{N}_i,k}, u_{i,k}^L)$ for all $i \in \mathcal{V}$ and for $k \geq \bar{k}$ guarantees the satisfaction of constraints (2) at all times.

3. DISTRIBUTED INVARIANCE-BASED SAFETY FRAMEWORK

In this section, we present the proposed distributed safety framework for constrained linear systems subject to bounded parametric uncertainty, defined in (1). We address the problem by first introducing a safety framework that implicitly defines the safe states and computes an input that keeps the system safe by solving an optimization problem. In order to address real-time requirements, we then derive an explicit approximation of the safety framework. A safe region of the state space and a backup control law are explicitly computed in the form of a union of robust invariant sets and stabilizing linear feedback laws.

3.1 Implicit Formulation

The safety framework is based on the concepts of robust invariance and structured invariant sets, proposed in Conte et al. (2016). The algorithm has two key features: 1) it provides an implicit description of the set of safe states and safe inputs through a convex optimization problem, 2) it allows for distributed computation.

The safety framework works as follows: Given a state $x_{i,k} \forall i \in \mathcal{V}$, we certify a learning input $u_{i,k}^L$ as safe, by verifying that the uncertain one-step prediction lies within a robust invariant set of the form $\mathcal{X} = \{x \mid x^T P x = \sum_{i=1}^N x_i^T P_i x_i \leq 1\} \subseteq \mathbb{X}$ for system (1) under a stabilizing control law $u_{i,k} = K_i x_{\mathcal{N}_i}$, where $K_i \in \mathbb{R}^{m_i \times n_{\mathcal{N}_i}}$, and $P = \text{diag}(P_1, \dots, P_N)$. The following conditions are sufficient for the existence of such a set:

$$(A(\theta) + B(\theta)K)^T P (A(\theta) + B(\theta)K) - P \leq 0, \quad (4)$$

$$\sum_{i=1}^N x_{i,k+1}^T P_i x_{i,k+1} \leq 1, \quad (5)$$

$$P_i > 0 \quad \forall i \in \mathcal{V}, \quad (6)$$

$$x_{i,k+1} = A_i(\theta_i) x_{\mathcal{N}_i,k} + B_i(\theta_i) u_{i,k}^L, \quad (7)$$

for all possible uncertain parameter values θ_i , where K is an appropriate composition of the structured linear feedback gains K_i . Note that (4) requires a matrix P which makes \mathcal{X} a robust ellipsoidal invariant set for the uncertain autonomous linear system $A(\theta) + B(\theta)K$. The LMI (5) verifies whether the learning inputs $u_{i,k}^L$ are safe, i.e., the one-step predictions satisfy $x_{i,k+1} \in \mathcal{X}_i \forall i \in \mathcal{V}$.

Remark 1. Condition (4) is a global condition, i.e., it has to hold for the whole network of systems. However, this condition can be distributed by upper bounding the global Lyapunov decrease with a structured symmetric matrix. More details are given in the proof of Proposition 1 and can be found in Conte et al. (2016).

In order to define a distributed safety framework that interferes as little as possible with the learning algorithm, we propose the following approach. We introduce an optimization problem that ensures feasibility of (4) - (7) for

a safe input that is as close as possible to the learning input $u_{i,k}^L$. This can be achieved by defining the safe input as $u_{i,k} = u_{i,k}^L + \Delta u_{i,k}$, and minimizing $\Delta u_{i,k}$ subject to (4) - (7), and $u_{i,k} \in \mathbb{U}_i$.

The following proposition formalizes the distributed safety framework described above.

Proposition 1. Consider system (1) subject to (2). Assume $x_{i,0} \forall i \in \mathcal{V}$ to be a feasible initial condition for the following optimization problem

$$\min_{E, Y, S, \Delta u} \|\Delta u_{i,k}\|^2 \quad (8a)$$

$$\text{s.t. } \forall l \in \{1, \dots, n_{h_i}\}, \forall e \in \{1, \dots, n_{o_i}\}, \quad (8b)$$

$$\forall i \in \mathcal{V}, \theta_i \in \Theta_i, \quad (8c)$$

$$E_i > 0, \quad (8c)$$

$$\begin{bmatrix} \bar{E}_i + \mathcal{S}_{\mathcal{N}_i} & \star \\ A_{\mathcal{N}_i}(\theta_i) E_{\mathcal{N}_i} + B_i(\theta_i) Y_i & E_i \end{bmatrix} \geq 0, \quad (8d)$$

$$\sum_{r \in \mathcal{N}_i} T_i W_r^T S_{\mathcal{N}_r} W_r T_i^T \leq 0, \quad (8e)$$

$$\begin{bmatrix} 1/N & x_{i,k+1}^T \\ x_{i,k+1} & E_i \end{bmatrix} \geq 0, \quad (8f)$$

$$O_i (\Delta u_{i,k} + u_{i,k}^L) \leq o_i, \quad (8g)$$

$$\begin{bmatrix} [h_i]_l^2 & [H_i]_{l,*} E_i \\ E_i [H_i]_{l,*}^T & E_i \end{bmatrix} \geq 0, \quad (8h)$$

$$\begin{bmatrix} [o_i]_e^2 & [O_i]_{e,*} Y_i \\ Y_i^T [O_i]_{e,*}^T & E_{\mathcal{N}_i} \end{bmatrix} \geq 0, \quad (8i)$$

$$x_{i,k+1} = A_{\mathcal{N}_i}(\theta_i) x_{\mathcal{N}_i} + B_i(\theta_i) (u_{i,k}^L + \Delta u_{i,k}) \quad (8j)$$

where $E_i = P_i^{-1}$, $E_{\mathcal{N}_i} = W_i E W_i^T$, $\bar{E}_i = W_i T_i^T E_i T_i W_i^T$, $\mathcal{S}_{\mathcal{N}_i} = E_{\mathcal{N}_i} \Gamma_{\mathcal{N}_i} E_{\mathcal{N}_i}$ and $Y_i = K_i E_{\mathcal{N}_i}$. Then

$$g_i(x_{\mathcal{N}_i,k}, u_{i,k}^L) = u_{i,k}^L + \Delta u_{i,k} \quad (9)$$

defines a distributed safety framework according to Definition 1, where $\Delta u_{i,k}$ is the solution of (8).

Proof. The proof is structured in three parts proving that *i*) inequalities (8c)-(8e) define a structured robust invariant set \mathcal{X} , *ii*) inequalities (8f)-(8g) determine a safe input u_k , and *iii*) inequalities (8h)-(8i) ensure the feasibility of the states and the inputs.

- i) Using the distribution of the Lyapunov decrease condition proposed in Conte et al. (2016), we reformulate (4) as a set of local inequalities coupled by a global condition

$$(A_{\mathcal{N}_i}(\theta_i) + B_i(\theta_i) K_i)^T P_i (A_{\mathcal{N}_i}(\theta_i) + B_i(\theta_i) K_i) - W_i T_i^T P_i T_i W_i^T \leq \Gamma_{\mathcal{N}_i} \quad \forall i \in \mathcal{V}, \quad (10)$$

$$\sum_{r=0}^N W_r^T \Gamma_{\mathcal{N}_i} W_r \leq 0. \quad (11)$$

Using the Schur complement on (10) and relaxing condition (11) by upper bounding it with block diagonal matrices $S_{\mathcal{N}_i}$, we obtain the LMIs (8c)-(8e).

- ii) Using the Schur complement on (5) we obtain

$$\begin{bmatrix} 1/N & x_{i,k+1}^T \\ x_{i,k+1} & E_i \end{bmatrix} \geq 0, \quad \forall i \in \mathcal{V},$$

where we use the fact that P and therefore $E = P^{-1}$ is block-diagonal. Condition (8f) guarantees that the un-

certain one-step prediction lies within the structured robust invariant set \mathcal{X}_i , while (8g) verifies that $u_{i,k}$ satisfies the input constraints.

- iii) Last, we consider the feasibility of the states and inputs, which is guaranteed if

$$\mathcal{X} \subseteq \{x \mid O_i K_i x_{\mathcal{N}_i} \leq o_i, \forall i \in \mathcal{V}\} \cap \mathbb{X}.$$

Using standard techniques for containment of ellipsoids in polyhedra, see e.g. Boyd et al. (1994), this can be expressed as (8h) and (8i), see e.g. Gao (2001). \square

Note that the objective of (8) is to minimize the modification of the learning inputs $u_{i,k}^L$ given by $\Delta u_{i,k}$. Problem (8) is convex and structured, which renders the optimization problem amenable to distributed optimization, see e.g. Bertsekas and Tsitsiklis (1989) for an overview. Checking all possible realizations of the uncertain parameter space Θ_i could be computationally intractable. However, given the fact that we are considering linear systems, it is sufficient to check the extremal points of the polytopic sets Θ_i to compute (8) $\forall \theta_i \in \Theta_i$, see e.g. Kharitonov (1978).

The presented safety framework requires to solve (8) at every iteration, which can be computationally demanding. In the following section, we introduce an approximation that reduces the online computational burden to simple function evaluations.

3.2 Explicit Approximation

The explicit safety framework derived in this section takes inspiration from explicit model predictive control (MPC), Bemporad et al. (2002), and locally weighted projection regression (LWPR), Vijayakumar and Schaal (2000). Explicit MPC removes the need to solve an optimization problem online by pre-computing the control law offline, such that the online operations reduce to simple function evaluations. LWPR achieves nonlinear function approximation in high dimensional spaces by partitioning the feature space using data, and fitting local linear models to each partition.

The proposed explicit safety framework consists of two phases: the initialization and the learning phase. The former is an offline phase, where the safe set is explicitly approximated. First, the state space is divided into M regions \mathcal{R}^j with $j \in \mathcal{M} = \{1, \dots, M\}$, ensuring coverage of the whole state space. Subsequently, a structured ellipsoidal robust invariant set $\mathcal{X}^j = \{x \mid \sum_{i=1}^N x_i^T P_i^j x_i \leq 1\} \subset \mathcal{X}$, such that $\mathcal{X}^j \cap \mathcal{R}^j \neq \emptyset$, and a corresponding structured linear stabilizing controller $K_i^j \in \mathbb{R}^{m_i \times n_{\mathcal{N}_i}}$, $\forall i \in \mathcal{V}$, is computed for every region \mathcal{R}^j . The overall safe set is defined as the union of all robust invariant sets \mathcal{X}^j . By partitioning the state space and requiring non-empty intersection of each partition and the corresponding robust invariant set, as well as by computing a linear feedback law for each invariant set, the resulting robust invariant sets will have different shapes and orientations, see e.g. Figure 1. The union of ellipsoidal robust invariant sets will thereby result in a larger safe set than achievable with a single robust ellipsoidal invariant set.

The state space, delimited by the state constraints (2), is divided into polytopic regions of the form $\mathcal{R}^j = \{x \mid$

$A_{\mathcal{R}^j} x \leq b_{\mathcal{R}^j}\}$ using Voronoi partitioning, see e.g. Cortes et al. (2004). The structured robust invariant sets \mathcal{X}^j are computed in a distributed manner by modifying (8). The non-empty intersection between the robust invariant set \mathcal{X}^j and the region \mathcal{R}^j is enforced with the following convex constraint

$$x^T P^j x \leq 1, \quad (12a)$$

$$A_{\mathcal{R}^j} x \leq b_{\mathcal{R}^j}. \quad (12b)$$

The computation of the structured invariant sets and the stabilizing control law is formalized by the following proposition.

Proposition 2. Consider system (1) subject to (2). Given M convex regions of the form $\mathcal{R}^j = \{x \mid A_{\mathcal{R}^j} x \leq b_{\mathcal{R}^j}\}$, M structured robust invariant sets $\mathcal{X}^j = \{x \mid \sum_{i=1}^N x_i^T P_i^j x_i \leq 1\}$ such that $\mathcal{X}^j \cap \mathcal{R}^j \neq \emptyset$, and M structured linear stabilizing control laws $u_i^j = K_i^j x_{\mathcal{N}_i}$ $\forall i \in \mathcal{V}$ with $j \in \mathcal{M}$ can be computed with the following optimization problem:

$$\min_{E, Y, S, x} \sum_{i=1}^N \text{tr} \left(E_i^j \right) \quad (13a)$$

$$\text{s.t. } \forall l \in \{1, \dots, n_{h_i}\}, \forall e \in \{1, \dots, n_{o_i}\}, \quad (13b)$$

$$\forall i \in \mathcal{V}, \forall j \in \{1, \dots, M\}, \theta_i \in \Theta_i, \quad (13c)$$

$$E_i^j > 0, \quad (13c)$$

$$\begin{bmatrix} \bar{E}_i^j + S_{\mathcal{N}_i}^j & \star \\ A_{\mathcal{N}_i}(\theta_i) E_{\mathcal{N}_i}^j + B_i(\theta_i) Y_i^j & E_i^j \end{bmatrix} \geq 0, \quad (13d)$$

$$\sum_{r \in \mathcal{N}_i} T_i W_r^T S_{\mathcal{N}_r}^j W_r T_i^T \leq 0, \quad (13e)$$

$$\begin{bmatrix} [h_i]_l^2 & [H_i]_{l,*} E_{\mathcal{N}_i}^j \\ E_{\mathcal{N}_i}^j [H_i]_{l,*}^T & E_{\mathcal{N}_i}^j \end{bmatrix} \geq 0, \quad (13f)$$

$$\begin{bmatrix} [o_i]_e^2 & [O_i]_{e,*} Y_i^j \\ (Y_i^j)^T [O_i]_{e,*}^T & E_{\mathcal{N}_i}^j \end{bmatrix} \geq 0, \quad (13g)$$

$$\begin{bmatrix} 1/N x_i^T \\ x_i & E_i^j \end{bmatrix} \geq 0, \quad (13h)$$

$$W_i A_{\mathcal{R}^j} W_i^T x_{\mathcal{N}_i} \leq W_i b_{\mathcal{R}^j}, \quad (13i)$$

where $E_i^j = (P_i^j)^{-1}$, $E_{\mathcal{N}_i}^j = W_i^T E^j W_i$, $\bar{E}_i^j = W_i T_i^T E_i^j T_i W_i^T$, and $Y_i^j = K_i^j E_{\mathcal{N}_i}^j$.

Problem (13) is convex and can be solved in a distributed manner using the same techniques as described in Subsection 3.1.

We now describe the learning phase. Given the state $x_{i,k}$ $\forall i \in \mathcal{V}$, the safety framework applies the learning input $u_{i,k}^L$ $\forall i \in \mathcal{V}$ if the uncertain one-step prediction lies within one of the robust invariant sets \mathcal{X}^j with $j \in \mathcal{M}$. Otherwise, the safety framework applies one of the precomputed structured linear stabilizing controllers $u_{i,k} = K_i^{b^*} x_{\mathcal{N}_i,k}$ $\forall i \in \mathcal{V}$, where b^* is the index of the robust invariant set that contains the current state and results in the smallest modification of the learning input, i.e.,

$$b^* = \arg \min_b \sum_{i=1}^N \|u_{i,k}^L - K_i^b x_{\mathcal{N}_i,k}\| \quad (14)$$

s.t. $x_{i,k} \in \mathcal{X}_i^b, \forall i \in \mathcal{V}.$

We formalize the safety framework outlined above in the following proposition.

Proposition 3. Consider system (1) subject to (2), given the initial condition $x_{i,0} \in \mathcal{X}_i^j \forall i \in \mathcal{V}$ for some $j \in \mathcal{M}$, where \mathcal{X}_i^j are computed using (13). Then

$$g_i(x_{\mathcal{N}_i,k}, u_{i,k}^L) = \begin{cases} u_{i,k}^L & \text{if } \exists j \in \mathcal{M} : x_{i,k+1} \in \mathcal{X}_i^j, \forall i \in \mathcal{V} \\ K_i^{b^*} x_{\mathcal{N}_i,k} & \text{otherwise,} \end{cases} \quad (15)$$

$\forall i \in \mathcal{V}$, with b^* in (14), defines a distributed safety framework according to Definition 1.

Proof. We prove the two distinct cases in (15) separately. The first case follows directly from the fact that $\mathcal{X}^j \subset \mathbb{X}$ and constraint satisfaction in the future follows from robust invariance of \mathcal{X}^j . The second case is trivial since $K_i^{b^*}$ is one of the safe control laws K_i^j , thus $K_i^{b^*} x_{\mathcal{N}_i,k}$ is always a safe input. Therefore, (15) defines a safety framework according to Definition 1. \square

Note that the safety framework defined by (15) can be computed in a distributed fashion. The containment in the invariant set can be cast as an averaging consensus problem, where the cost to average is

$$J(x) = \frac{1}{N} x^T P x = \frac{1}{N} \sum_{i=1}^N x_i^T P_i x_i \forall j \in \mathcal{M}.$$

The containment is guaranteed if and only if the consensus algorithm converges to a steady state smaller than or equal to $\frac{1}{N}$, for more information on consensus problems see Olfati-Saber et al. (2007). Similarly, the optimization problem (14) can be solved using min-consensus where the containment in the invariant set can be verified as before. The explicit formulation of the safety framework is summarized in Algorithm 1.

Algorithm 1 Explicit Distributed Invariance-based Safety Framework

- 1: **Initialization Phase** (offline)
 - 2: Compute the partitions \mathcal{R}^j
 - 3: Solve optimization problem (13)
 - 4:
 - 5: **Learning Phase** (online)
 - 6: **for** $k = 1, 2, \dots$ **do**
 - 7: **for** $\forall i \in \mathcal{V}$ **do**
 - 8: Measure $x_{i,k}$
 - 9: Exchange $x_{i,k}$ with neighbors
 - 10: Obtain learning input $u_{i,k}^L$
 - 11: Apply $u_{i,k} = g_i(x_{\mathcal{N}_i,k}, u_{i,k}^L)$, defined by (15)
 - 12: **end for**
 - 13: **end for**
-

4. NUMERICAL RESULTS

In this section, we show the effectiveness of the proposed safety framework with three numerical examples. First, we illustrate the application of the explicit approximation presented in Subsection 3.2 in combination with a learning algorithm on a small-scale system. Subsequently, we also show the application to a large-scale system and end the section showing the quality of the explicit approximation

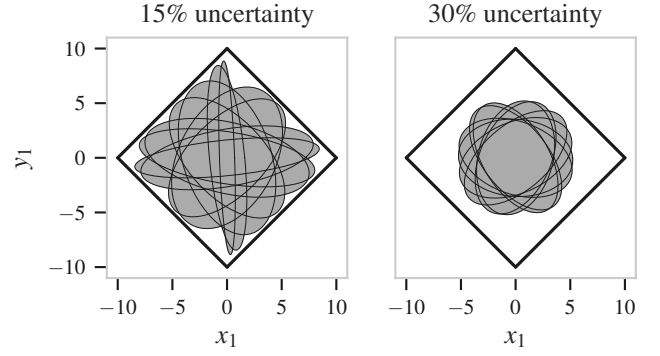


Fig. 1. Invariant sets for one agent, using 10 partitions with an uncertainty level of 15% and 30%.

compared to the implicit formulation presented in Subsection 3.1. All examples are implemented in Python using CVXPY and OpenAI Gym, see Diamond and Boyd (2016) and Brockman et al. (2016), respectively, and solved using MOSEK, see MOSEK ApS (2019). The examples were run on a machine equipped with an Intel Core i9 (3.1 GHz) CPU and 32 GB of RAM.

4.1 Explicit Approximation in Combination with a Learning Algorithm

In the first example, we make use of Proximal Policy Optimization (PPO), see Schulman et al. (2017), to learn a control policy for a three-agent mass-damper system, where the individual agent’s dynamics are given by

$$\begin{aligned} m_i \ddot{x}_i &= a_i \dot{x}_i + d_{ij} (\dot{x}_j - \dot{x}_i) + u_i, \\ m_i \ddot{y}_i &= a_i \dot{y}_i + d_{ij} (\dot{y}_j - \dot{y}_i) + u_i \end{aligned} \quad \forall j \in \mathcal{N}_i, \quad (16)$$

where d_{ij} is the damping coefficient, a_i is a parameter damping the agent’s local dynamics, and x_i and y_i are the lateral and longitudinal axes in the two dimensional plane, respectively. The parameter vector θ_i consists of the parameter a_i and the parameters d_{ij} , which couple agent i to all neighboring agents. The uncertainty set reads as

$$\Theta_i = [(1 - \gamma) \theta_i^{nominal}, (1 + \gamma) \theta_i^{nominal}],$$

where γ is the uncertainty level, and $\theta_i^{nominal}$ is the true parameter vector, which is also used in the simulation model. The agents’ state constraints are of the form (2) and only constrain the local positions, while the local velocities are unconstrained. As the communication graph \mathcal{G} we use again a line graph and the numerical values of the example are given in Table 1.

To measure the performance of the explicit approximation presented in Subsection 3.2, we compute the fraction of the constrained state space covered by the union of pre-computed robust invariant sets, i.e., $\text{vol}(\cup_{j=1}^N \mathcal{X}^j) / \text{vol}(\cup_{i=1}^N \mathbb{X}_i)$. Figure 1 shows the union of the robust invariant sets for the first agent for two levels of uncertainty. The shown cross section illustrates the plane at zero velocity in both directions x and y . The plot shows that a large portion of the state space is covered, and that the use of multiple robust invariant sets is beneficial to cover a larger area of the state space. Figure 2 shows the fraction of the constrained state space covered by the union of invariant sets as a function of the number of regions in a partition and the uncertainty level γ . Since the partitioning affects




Fig. 2. Fraction of the constrained state space covered by the union of pre-computed robust invariant sets for different numbers of partitions and different uncertainty levels.

the computation of the robust invariant sets, the coverage results in Figure 2 are averaged over 20 Voronoi partitions obtained from random initial conditions.

To demonstrate the ability of the explicit approximation to ensure constraint satisfaction, we consider the aforementioned system in a learning context, where the goal is to learn a controller, which regulates the state of each system back to the origin. We set the uncertainty level to 20% and the number of partitions M to 15. The learning setting is episodic, and at the start of each episode, each system is set to a random position in the two dimensional plane within the state constraints. For the learning algorithm we use PPO, which is also applicable in a multi-agent setting, see e.g. Srinivasan et al. (2018). The learning curve for PPO with and without the explicit approximation of the distributed safety framework is depicted in Figure 3. In this example, violating the state constraints would result in a negative reward, which never happens if the safety framework is active. In both cases, the learning algorithm converges in about the same number of iterations, therefore the safety framework has no negative affect on the learning rate in this example.

Table 1. Numerical values for the first example.

m_i	a_i	d_{ij}	H_i	h_i	O_i	o_i
1	0.1	0.5	$\begin{bmatrix} 10 & 0 & 10 & 0 \\ 10 & 0 & -10 & 0 \\ -10 & 0 & 10 & 0 \\ -10 & 0 & -10 & 0 \end{bmatrix}$	$\begin{bmatrix} 100 \\ 100 \\ 100 \\ 100 \end{bmatrix}$	$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$	$\begin{bmatrix} 5 \\ -5 \end{bmatrix}$

4.2 Guaranteed Safety on a Large-Scale System

In the second example, we use a Gaussian Processes-based exploration algorithm, see e.g. Todescato et al. (2017), in combination with the proposed safety framework to safely explore the state space of a large-scale system. The Gaussian Process is exploited to estimate the uncertainty at any point in the state space and to select the new input locations in order to reduce the maximum of the posterior covariance. We consider a 25-agent mass-spring-damper system, where the individual agent's dynamics are described by

$$m_i \ddot{x}_i = k_{ij} (x_j - x_i) + d_{ij} (\dot{x}_j - \dot{x}_i) + u_i \quad \forall j \in \mathcal{N}_i, \quad (17)$$

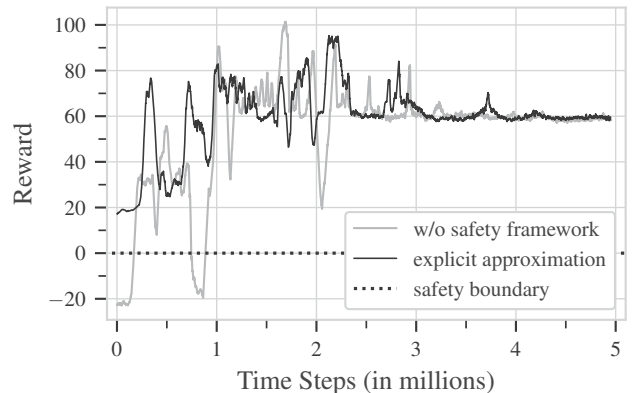


Fig. 3. Learning curve for a simple regulation task, using PPO with and without the explicit approximation of the safety framework.

where k_{ij} and d_{ij} are the spring and damping coefficients, respectively. Additionally, the dynamics (17) are subject to constraints of the form (2) and the communication graph \mathcal{G} is a line graph. Table 2 provides the numerical values of both, the dynamical parameters and the constraint matrices. Figure 4 shows the ability of the safety framework to keep the system safe, despite an input from the exploration algorithm, which drives some agents close to or beyond the state constraints, which is highlighted for a subset of the 25 agents. Note that the constraints shown in Figure 4 are the state constraints of the individual agents and do not include the additional constraints imposed by the dynamic coupling, i.e., the restrictions on the state of agent i due to the state of its neighbors, as these result in time-varying local constraints.

Table 2. Numerical values for the second and third example.

m_i	k_{ij}	d_{ij}	H_i	h_i	O_i	o_i
1	2	1	$\begin{bmatrix} 1 & 0 \\ -1 & 0 \\ 0 & 1 \\ 0 & -1 \end{bmatrix}$	$\begin{bmatrix} 1 \\ -1 \\ 3 \\ -3 \end{bmatrix}$	$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$	$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$

4.3 Quality of the Explicit Approximation

In the third example, we use deep double Q-learning, see van Hasselt et al. (2015), to learn a control policy that regulates a three-agent mass-spring-damper system back to the origin. The goal of this example is to evaluate the quality of the explicit approximation by comparing it against the implicit safety framework. We again use the model in (17) with constraints (2) as in the second example but reduce the number of agents, in order to facilitate the use of a reinforcement algorithm. Figure 5 shows the learning curve in terms of the cumulative reward for the system without a safety framework, for the safety framework, and for its explicit approximation. We tuned the hyper-parameters of the Q-learning algorithm on the first setting and retrained it with the same hyper-parameters in the latter two settings. In all three settings, negative rewards indicate constraint violations. We can see from the figure that both, the distributed safety framework as well as the explicit approximation show comparable behavior and guarantee constraint satisfaction at all times.

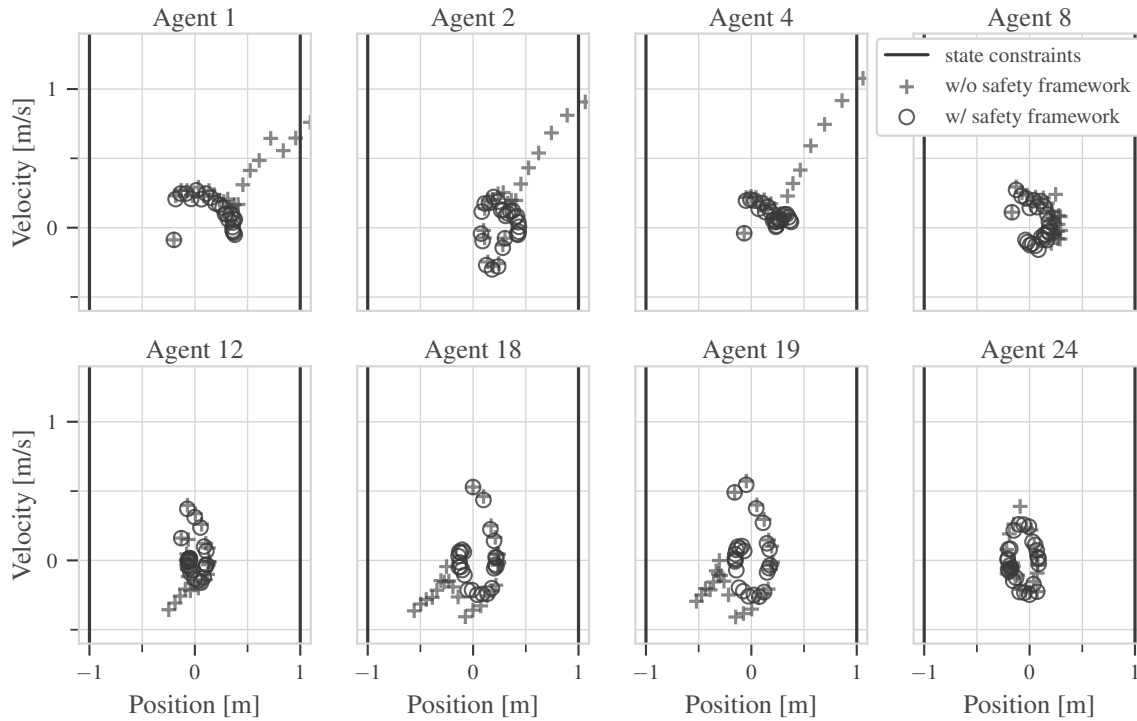


Fig. 4. State evolution of eight selected agents driven by the unsafe exploration algorithm with and without the safety framework.

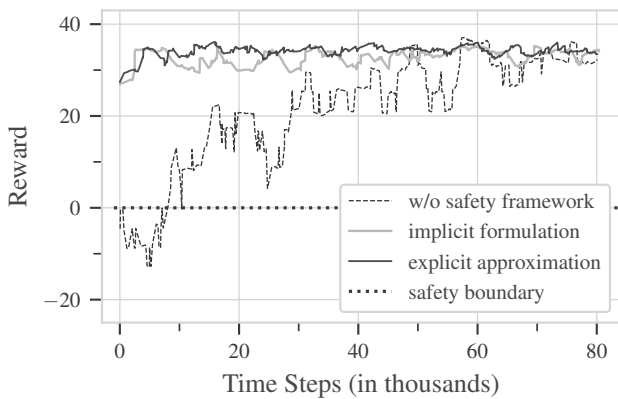


Fig. 5. Learning curve for deep double Q-learning on a second-order system with three agents, using no safety framework, the implicit formulation of the safety framework, and its explicit approximation.

5. CONCLUSION

In this paper, we introduced an explicit safety framework, which guarantees constraint satisfaction for distributed uncertain linear systems based on local conditions. We showed with three numerical examples the effectiveness of the proposed algorithm in ensuring constraint satisfaction while learning.

REFERENCES

Bemporad, A., Morari, M., Dua, V., and Pistikopoulos, E.N. (2002). The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1), 3 – 20. doi: [https://doi.org/10.1016/S0005-1098\(01\)00174-1](https://doi.org/10.1016/S0005-1098(01)00174-1).

Berkenkamp, F. and Schoellig, A.P. (2015). Safe and robust learning control with gaussian processes. In *2015 European Control Conference (ECC)*, 2496–2501. doi: 10.1109/ECC.2015.7330913.

Bertsekas, D.P. and Tsitsiklis, J.N. (1989). *Parallel and distributed computation: numerical methods*, volume 23. Prentice hall Englewood Cliffs, NJ.

Blanchini, F., Franco, E., and Giordano, G. (2015). Network-decentralized control strategies for stabilization. *IEEE Transactions on Automatic Control*, 60(2), 491–496. doi:10.1109/TAC.2014.2331415.

Boyd, S., El Ghaoui, L., Feron, E., and Balakrishnan, V. (1994). *Linear Matrix Inequalities in System and Control Theory*, volume 15 of *Studies in Applied Mathematics*. SIAM, Philadelphia, PA.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). OpenAI Gym.

Busoniu, L., Babuska, R., and De Schutter, B. (2008). A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2), 156–172. doi:10.1109/TSMCC.2007.913919.

Chen, Y., Anderson, J., Kalsi, K., Ames, A.D., and Low, S.H. (2019). Safety-Critical Control Synthesis for network systems with Control Barrier Functions and Assume-Guarantee Contracts. *arXiv e-prints*, arXiv:1911.03452.

Chiuso, A. and Pillonetto, G. (2019). System identification: A machine learning perspective. *Annual Review of Control, Robotics, and Autonomous Systems*, 2(1), 281–304. doi:10.1146/annurev-control-053018-023744.

Conte, C., Jones, C.N., Morari, M., and Zeilinger, M.N. (2016). Distributed Synthesis and Stability of Coop-

- erative Distributed Model Predictive Control for Linear Systems. *Automatica*, 69(January), 117–125. doi:10.1016/j.automatica.2016.02.009.
- Cortes, J., Martinez, S., Karatas, T., and Bullo, F. (2004). Coverage control for mobile sensing networks. *Robotics and Automation, IEEE Transactions on*, 20(2), 243–255.
- Diamond, S. and Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83), 1–5.
- Fisac, J.F., Akametalu, A.K., Zeilinger, M.N., Kaynama, S., Gillula, J., and Tomlin, C.J. (2019). A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7), 2737–2752. doi:10.1109/TAC.2018.2876389.
- Gao, L. (2001). *Ellipsoidal Approximation to Polytopes and Computational Study of Lenstra’s Algorithm*. Ph.D. thesis, Rice University.
- García, J. and Fernández, F. (2015). A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16, 1437–1480. URL <http://jmlr.org/papers/v16/garcia15a.html>.
- Gillula, J.H. and Tomlin, C.J. (2011). Guaranteed safe online learning of a bounded system. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2979–2984.
- Gillula, J.H. and Tomlin, C.J. (2012). Guaranteed safe online learning via reachability: tracking a ground target using a quadrotor. In *2012 IEEE International Conference on Robotics and Automation*, 2723–2730.
- Kar, S., Moura, J.M.F., and Poor, H.V. (2013a). QD-learning: A collaborative distributed strategy for multi-agent reinforcement learning through consensus + innovations. *IEEE Transactions on Signal Processing*, 61(7), 1848–1862.
- Kar, S., Moura, J.M.F., and Poor, H.V. (2013b). Distributed reinforcement learning in multi-agent networks. In *2013 5th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 296–299.
- Kharitonov, V.L. (1978). Asymptotic stability of an equilibrium position of a family of systems of differential equations. *Differential Equations*, 14, 1483–1485.
- Koller, T., Berkenkamp, F., Turchetta, M., and Krause, A. (2018). Learning-based Model Predictive Control for Safe Exploration and Reinforcement Learning. URL <http://arxiv.org/abs/1803.08287>.
- Larsen, R.B., Carron, A., and Zeilinger, M.N. (2017). Safe Learning for Distributed Systems with Bounded Uncertainties. *IFAC-PapersOnLine*, 50(1), 2536–2542. doi:10.1016/j.ifacol.2017.08.069.
- MOSEK ApS (2019). *MOSEK Optimizer API for Python. Version 9.1*. URL <https://docs.mosek.com/9.1/pythonapi/index.html>.
- Muntwiler, S., Wabersich, K.P., Carron, A., and Zeilinger, M.N. (2019). Distributed Model Predictive Safety Certification for Learning-based Control. *arXiv preprint*. URL <https://arxiv.org/abs/1911.01832>.
- Nguyen, T.T., Nguyen, N.D., and Nahavandi, S. (2018). Deep reinforcement learning for multi-agent systems: A review of challenges, solutions and applications. *CoRR*, abs/1812.11794. URL <http://arxiv.org/abs/1812.11794>.
- Olfati-Saber, R., Fax, J.A., and Murray, R.M. (2007). Consensus and Cooperation in Networked Multi-Agent Systems. *Proceedings of the IEEE*, 95(1), 215–233. doi:10.1109/JPROC.2006.887293. URL <http://ieeexplore.ieee.org/document/4118472/>.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint*. URL <http://arxiv.org/abs/1707.06347>.
- Srinivasan, S., Lanctot, M., Zambaldi, V., Perolat, J., Tuyls, K., Munos, R., and Bowling, M. (2018). Actor-critic policy optimization in partially observable multi-agent environments. In *Advances in Neural Information Processing Systems 31*, 3422–3435.
- Todescato, M., Carron, A., Carli, R., Pillonetto, G., and Schenato, L. (2017). Multi-robots gaussian estimation and coverage control: From client-server to peer-to-peer architectures. *Automatica*, 80, 284–294.
- Tütüncü, R.H., Toh, K.C., and Todd, M.J. (2003). Solving semidefinite-quadratic-linear programs using SDPT3. *Mathematical Programming*, 95(2), 189–217. doi:10.1007/s10107-002-0347-5. URL <https://doi.org/10.1007/s10107-002-0347-5>.
- Tuyls, K. and Weiss, G. (2012). Multiagent learning: Basics, challenges, and prospects. *Ai Magazine*, 33, 41–52. doi:10.1609/aimag.v33i3.2426.
- van Hasselt, H., Guez, A., and Silver, D. (2015). Deep Reinforcement Learning with Double Q-learning. URL <http://arxiv.org/abs/1509.06461>.
- Varshavskaya, P., Kaelbling, L.P., and Rus, D. (2009). *Efficient Distributed Reinforcement Learning through Agreement*, 367–378. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Vijayakumar, S. and Schaal, S. (2000). Locally weighted projection regression: An $O(n)$ algorithm for incremental real time learning in high dimensional space. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, 288–293.
- Wabersich, K., Hewing, L., Carron, A., and N. Zeilinger, M. (2019). Probabilistic model predictive safety certification for learning-based control. *arXiv preprint 1906.10417*.
- Wabersich, K. and Zeilinger, M.N. (2018a). Linear model predictive safety certification for learning-based control. In *2018 IEEE Conference on Decision and Control (CDC)*, 7130–7135. doi:10.1109/CDC.2018.8619829.
- Wabersich, K. and Zeilinger, M.N. (2018b). Safe exploration of nonlinear dynamical systems: A predictive safety filter for reinforcement learning. *arXiv preprint 1812.05506*.
- Watkins, C.J.C.H. and Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3-4), 279–292. doi:10.1007/BF00992698. URL <http://link.springer.com/10.1007/BF00992698>.
- Wieland, P., Ebenbauer, C., and Allgower, F. (2007). Ensuring task-independent safety for multi-agent systems by feedback. In *2007 American Control Conference*, 3880–3885.
- Zhang, K., Yang, Z., Liu, H., Zhang, T., and Basar, T. (2018). Fully Decentralized Multi-Agent Reinforcement Learning with Networked Agents. 1–45. URL <http://arxiv.org/abs/1802.08757>.