

CaRT: Certified Safety and Robust Tracking in Learning-based Motion Planning for Multi-Agent Systems

Hiroyasu Tsukamoto*, Benjamin Rivière*, Changrak Choi†, Amir Rahmani†, and Soon-Jo Chung*

Abstract—The key innovation of our analytical method, CaRT, lies in establishing a new hierarchical, distributed architecture to guarantee the safety and robustness of a given learning-based motion planning policy. First, in a nominal setting, the analytical form of our CaRT safety filter formally ensures safe maneuvers of nonlinear multi-agent systems, optimally with minimal deviation from the learning-based policy. Second, in off-nominal settings, the analytical form of our CaRT robust filter optimally tracks the certified safe trajectory, generated by the previous layer in the hierarchy, the CaRT safety filter. We show using contraction theory that CaRT guarantees safety and the exponential boundedness of the trajectory tracking error, even under the presence of deterministic and stochastic disturbance. Also, the hierarchical nature of CaRT enables enhancing its robustness for safety just by its superior tracking to the certified safe trajectory, thereby making it suitable for off-nominal scenarios with large disturbances. This is a major distinction from conventional safety function-driven approaches, where the robustness originates from the stability of a safe set, which could pull the system over-conservatively to the interior of the safe set. Our log-barrier formulation in CaRT allows for its distributed implementation in multi-agent settings. We demonstrate the effectiveness of CaRT in several examples of nonlinear motion planning and control problems, including optimal, multi-spacecraft reconfiguration.

I. INTRODUCTION

Learning-based control has been a subject of intense study for solving large-scale and complex problems that conventional approaches fail to handle; one example of which is to construct a real-time, nonlinear motion planning and control algorithm for multi-agent robotic and aerospace systems. The safety and robustness of the machine-learning approaches, however, highly depend on a number of factors such as structures of approximation models, tools used to generate training and test data, data pre-conditioning, learning duration, and training schemes. This makes it difficult to obtain a universal result applicable to any learning approach, treating some parts of the performance as a black box.

The purpose of this paper is to provide control theoretical safety and robust tracking guarantees to given learned motion planning policies for nonlinear multi-agent systems, independently of the performance of the learning approaches used in designing the learned policy. Our approach called CaRT (Certified Safety and Robust Tracking) is two-fold as follows.

Contributions: First, assuming that the dynamics model is free of disturbance and uncertainty, we construct an optimal safety filter for generating a safe target trajectory of nonlinear multi-agent systems in this ideal setting. We explicitly derive an analytical form of the optimal control input processed by the safety filter, which minimizes its deviation from the control input of a given learned motion planning policy while ensuring safety. We utilize the log-barrier formulation [1] so that the global safety violation can be decomposed as the sum of the local safety violations, allowing for the distributed implementation of our analytical safety filter in a multi-agent setting.

Second, we hierarchically design a robust filter based on contraction theory [2], [3] for guaranteeing exponential tracking of the safe target trajectory, even when the nominal dynamics model is subject to unknown deterministic and stochastic perturbation. This is because, although the safety filter mentioned earlier (including the CLF-CBF control [4], see Sec. II-B.3) also possesses a robustness guarantee, it is based primarily on the repelling force of its control input to push the system state back to the safe set (i.e., asymptotic/exponential stability of the safe set). This can pull the system over-conservatively to the interior of the safe set, e.g., when implementing the filter in real-world systems involving discretization of the control policy and dynamics. Adding a robust filter hierarchically, based on incremental stability of system trajectories, explicitly separates disturbances from safety violations and unknown dynamics, thus alleviating the burden of the safety filter in robustly dealing with the disturbance (see Sec. II-B and IV-A for details). This filter also provides an analytical form of the control input for real-time implementation. The application of these two filters in CaRT results in provable safety and robust tracking guarantees augmented on top of the given learned motion planning policy, without prior knowledge of its original learning performance. Note that the analytical solution processed by these filters is still differentiable, thereby allowing for the end-to-end learning of the safe control policy as in [5].

As for the nominal dynamics model, we start our discussion with the Lagrangian dynamical systems [6, p. 392], which can be used to describe a wide range of robotic and aerospace systems. We then extend the safety and robustness results to general control-affine nonlinear systems. Our approach is demonstrated in several nonlinear multi-agent motion planning and control problems.

Related Work: The high-level comparison between our approach, CaRT, and the methods outlined in this section

* Division of EAS, Caltech, Pasadena, CA, {htsukamoto, briviere, sjchung}@caltech.edu.

† Jet Propulsion Laboratory (JPL), Caltech, Pasadena, CA {changrak.choi, amir.rahmani}@jpl.nasa.gov.

This project is in part funded by DARPA LINC and JPL.

will be revisited in Sec. II-B in more detail.

The nonlinear robustness and stability can be analyzed using a Lyapunov function, which gives a finite tracking error with respect to a given target trajectory under the presence of external disturbances, including approximation errors of given learned motion planning policies. This framework thus provides a one way to ensure safety and robustness in learning-based control methods (see, e.g., [3], [7], [8] and references therein), which depends on the knowledge and the size of the approximation error of a given learned motion planning policy. Such information could be conservative for previously unseen data or available only empirically.

Control barrier functions, in contrast, guarantee the safety of nonlinear systems in real time without any knowledge of the learned motion planning errors at all. Formulating a CLF-CBF Quadratic Program (QP) [4], [9], [10] also gives some guarantees on robustness. As shall be elaborated in Sec. II-B.3, such guarantees are based on the repelling force of its control input to exponentially/asymptotically push the system state back to the stable safe set, which could lead to unnecessarily large/conservative control inputs for safety.

The purpose of our work is to propose a hierarchical approach to combine the best of both of these methods for safety and robustness, by performing contraction theory-based robust tracking of a provably safe trajectory, generated by a safety filter with the learned policy. The additional tracking-based robust filter is for reducing the burden in dealing with disturbances (see Sec. II-B and IV-A).

Notation: For $A \in \mathbb{R}^{n \times n}$, we use $A > 0$, $A \geq 0$, $A < 0$, and $A \leq 0$ for the positive definite, positive semi-definite, negative definite, negative semi-definite matrices, respectively. For $x \in \mathbb{R}^n$ and $A \in \mathbb{R}^{n \times m}$, we let $\|x\|$, $\|x\|_\Xi$, $\|A\|$, $\|A\|_F$ denote the Euclidean norm, weighted 2-norm (i.e., $\|x\|_\Xi = \sqrt{x^\top \Xi x}$ for $\Xi > 0$), induced 2-norm, and Frobenius norm, respectively. Also, \mathbb{E} denotes the expected value operator.

II. PROBLEM FORMULATION

We consider the following multi-agent Lagrangian dynamical system, perturbed by deterministic disturbance $d^i(x, t)$ with $\sup_{x,t} \|d^i(x, t)\| = \bar{d}^i \in [0, \infty)$ and Gaussian white noise of a Wiener process $\mathcal{W}(t)$ with $\sup_{x,t} \|\Gamma^i(x, t)\|_F = \bar{\gamma}^i \in [0, \infty)$:

$$\begin{aligned} M^i(p^i)dv^i + (C^i(p^i, v^i)v^i + G^i(p^i) + D^i(p^i, v^i))dt \\ = (u^i + d^i(x^i, t))dt + \Gamma^i(x^i, t)d\mathcal{W}^i(t), \quad i = 1, \dots, N \end{aligned} \quad (1)$$

where $t \in \mathbb{R}_+$, $x^i = [p^{i\top}, v^{i\top}]^\top$, i is the index of the i th agent, $p^i : \mathbb{R}_+ \mapsto \mathbb{R}^n$ and $v^i : \mathbb{R}_+ \mapsto \mathbb{R}^n$ are the generalized position and velocity of the i th agent ($\dot{p}^i = v^i$), $u^i \in \mathbb{R}^m$ ($m = n$ in this case) is the system control input, M^i , C^i , G^i , and D^i are known smooth functions that define the Lagrangian system, d^i and Γ^i are unknown bounded functions for external disturbances, \mathcal{W}^i is a w -dimensional Wiener process, and we consider the case where $\bar{d}^i, \bar{\gamma}^i \in [0, \infty)$ are given. We have $M^i(p^i) > 0$ and that the matrix $C^i(p^i, \dot{p}^i)$ is selected to make $\dot{M}^i(p^i) - 2C^i(p^i, v^i)$ skew-symmetric. Hence, $z^\top (\dot{M}^i(p^i) - 2C^i(p^i, v^i))z = 0$, $\forall z \in \mathbb{R}^n$ [6, p. 392]. We also consider the following general control-

affine nonlinear system:

$$\begin{aligned} dv^i = (f^i(p^i, v^i, t) + B(p^i, v^i, t)u^i)dt \\ + d^i(x^i, t)dt + \Gamma^i(x^i, t)d\mathcal{W}^i(t), \quad i = 1, \dots, N \end{aligned} \quad (2)$$

where f^i , and B^i are known smooth functions, $u^i \in \mathbb{R}^m$ is the system control input, and the other notations are consistent with that of (1). We assume the existence and uniqueness conditions of the solutions of (1) and (2) as in [11, p. 105].

The nonlinear motion planning problem of our interest is defined as follows:

$$u_{\text{global}}(t) = \arg \min_{\{u^i(t) \in \mathbb{R}^m\}_{i=1}^N} \int_0^{t_f} c(x(\tau), u(\tau), \tau) d\tau \quad (3)$$

s.t. (1) or (2) with $d^i = 0$ and $\Gamma^i = 0$, $x^i(t_f) = x_f^i$, $x^i(0) = x_0^i$

$$\|p^i(t) - p^j(t)\| \geq r_s, \quad \forall t, \forall i, j \neq i$$

where $x(t) = \{x^i(t)\}_{i=1}^N$, $u(t) = \{u^i(t)\}_{i=1}^N$, $c(x(t), u(t), t)$ is a user-specified cost at time t , t_f is the terminal time, x_0^i and x_f^i are the initial and terminal state, respectively, r_s is the minimal safe distance between i th agent and other objects, and j is the index denoting other agents and obstacles. The control policy u_{global} generates the reference trajectory x_{global} of Fig. 1 to be discussed in this section.

A. Learned Distributed Motion Planning Policy

Let \mathcal{N} denote the set of all the N agents, \mathcal{M} denote the set of all the M static obstacles, and o^i denote the local observation of the i th agent given as follows:

$$o^i = (x^i, \{x^j\}_{j \in \mathcal{N}^i}, \{p^j\}_{j \in \mathcal{M}^i}) \quad (4)$$

where x^i is the state of the i th agent, $\{x^j\}_{j \in \mathcal{N}^i}$ are the states of neighboring agents defined with $\mathcal{N}^i = \{j \in \mathcal{N} \mid \|p^i - p^j\| \leq r_{\text{sen}}\}$, $\{p^j\}_{j \in \mathcal{M}^i}$ are the positions of neighboring static obstacles defined with $\mathcal{M}^i = \{j \in \mathcal{M} \mid \|p^i - p^j\| \leq r_{\text{sen}}\}$, and r_{sen} is the sensing radius.

In this paper, we assume that we have access to a learned distributed motion planning policy $u_\ell^i(o^i)$ obtained by [5]. In particular, (i) we generate demonstration trajectories by solving the global nonlinear motion planning (3) and (ii) extract local observations from them for deep imitation learning to construct $u_\ell^i(o^i)$. Our differentiable safety and robust filters to be seen in Sec. III and IV can be used also in this phase to allow for end-to-end policy training as in [5].

B. Augmenting Learned Policy with Safety and Robustness

As discussed in Sec. I, directly using the learned motion planning policy has the following two issues in practice: (i) even in nominal settings without external disturbance in (1) and (2), the system solution trajectories computed with the learned motion planning policy could violate safety requirements due to learning errors, and (ii) the learned policy lacks formal mathematical guarantees of safety and robustness under the presence of external disturbance. Before going into details, let us see how we address these two problems analytically in real-time for the general systems (1) and (2), optimally and independently of the performance of the learning method used in the learned motion planning policy.

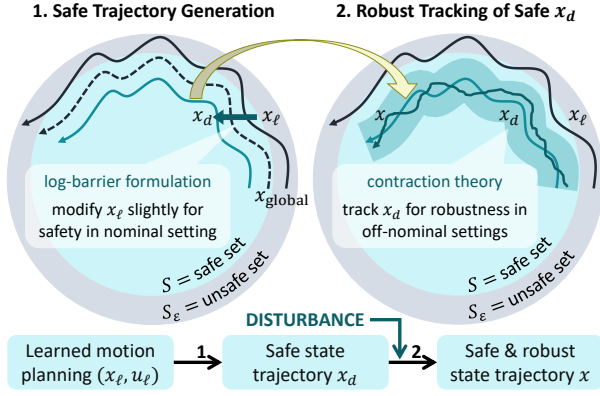


Fig. 1. Conceptual illustration of CaRT for a single agent, showing the hierarchical combination of our safety filter and robust filter, where S is the safe set, S_e is some fictitious unsafe set containing learned trajectory x_ℓ with learning error $\varepsilon > 0$, x_{global} is a reference trajectory given by global motion planner (see (3)), x_d is CaRT's target safe trajectory, and x is CaRT's actual state trajectory subject to disturbance. Note that we use a log-barrier formulation for the safe trajectory generation, which allows for the distributed and analytical implementation of CaRT.




Fig. 2. Different sources of robustness in conventional safety-driven approaches (left, e.g., CLF-CBF [4]) and CaRT (right), where $\alpha_h, \alpha_v > 0$, H is some safety function, S is its safe set, V_r are respective Lyapunov functions for robustness, x_d is a target trajectory, and M is a contraction metric [2], [3]. Safety-driven approaches are robust due to the stability of the safe set, while CaRT is robust due to the incremental stability of the closed-loop system with respect to the safe target trajectory x_d (see Sec. II-B).

1) CaRT Safety Filter and Built-in Robustness: Given the learned policy u_ℓ^* , we slightly modify it using our CaRT distributed safety filter to ensure the agents' safe operation. This is achieved by imposing a log-barrier safety constraint to guarantee safety under the presence of learning error $\varepsilon > 0$, as in the left-hand side of Fig. 1. Intuitively, since the learning error ε is expected to be small empirically, the contribution required for the safety filter is also expected to be small in a *nominal* setting in practice. Note that we use the log-barrier formulation for the safe of real-time, distributed implementation of the safety filter in nonlinear multi-agent settings.

However, the robustness of such safety-driven approaches results from the stability of the safe set S [9], as shown in the left-hand side of Fig. 2. This yields the pulling force to the set S that could be undesirably large in *off-nominal* settings, which could then lead to an unnecessarily large tracking error. This is especially true, e.g., in real-world scenarios involving the discretization of the control policy and dynamics.

2) CaRT Robust Filter and Tracking-based Robustness: Instead of handling both safety and robustness just by the safety filter, we can further utilize our CaRT robust filter

hierarchically to take over the role of disturbance attenuation in *off-nominal* settings. As depicted in the right-hand side of Fig. 1, this is achieved by contraction theory-based robust tracking [2], [3], which guarantees the *off-nominal* system state to stay in a bounded tube around the safe target trajectory x_d . Again, x_d is a slight modification of the learned trajectory x_ℓ in a *nominal* setting, processed through the CaRT safety filter.

We still use the Lyapunov formulation as in the safety filter for robustness, but now the Lyapunov function is defined incrementally as $V_r = (x - x_d)^T M (x - x_d)$, where x is the off-nominal system state and $M > 0$ is the contraction metric [2], [3]. As shown in the right-hand side of Fig. 2, improving the robustness performance here will simply result in superior tracking of the safe trajectory x_d , and thus can be achieved without losing too much information of the learned trajectory x_ℓ . In other words, safety is handled directly with $x \in S$ in the robust filter, unlike safety-driven approaches with indirect derivative safety constraints $\dot{H} \geq -\alpha_H H$.

These observations imply that

- when the learning error is much larger than the size of the external disturbance, then we can use our CaRT safety filter and its built-in robustness, and
- when the learning error is much smaller than the size of the external disturbance, which is often the case, then we can 1) modify the learned policy slightly with our CaRT safety filter in a *nominal* setting, and 2) handle disturbance primarily with our CaRT tracking-based robust filter in *off-nominal* settings, hierarchically on top of the safety filter.

3) Relationship with CLF-CBF: The CLF-CBF control [4] also considers safety and robustness. In our context, it constructs an optimal control input by solving a QP to minimize its deviation from the learned motion planning policy, subject to the safety constraint and the stability constraint. Since safety is its priority, the stability constraint has to be relaxed to $\dot{V}_r \leq -\alpha_v V_r + \rho$, where ρ is for QP feasibility and the Lyapunov function V_r is now defined as $V = (x - x_\ell)^T M (x - x_\ell)$ for the learned trajectory x_ℓ of Fig. 1. We list key differences between the CLF-CBF controller and our approach, CaRT:

- The primary distinction is the direction in which the respective tracking component steers the system. The CLF-CBF stability component steers towards the learned trajectory, which, because of learning error, might not be safe. This means that tracking stability can be compromised for CLF-CBF by prioritizing safety over stability and robustness. In contrast, CaRT's robust filter steers the system toward a certified safe trajectory, generated by the previous layer in the hierarchy, CaRT's safety filter. Because of this distinction, whereas CaRT can safely reject large disturbances with a large tracking gain α_v , this strategy is not practical for the CLF-CBF controller, which is forced to reject disturbances with large safety gain α_H . This could pull the system over-conservatively towards the interior of the safety set as shown in Fig. 2.
- The secondary distinction is that CLF-CBF requires solving a QP with a given Lyapunov function, while CaRT provides an explicit way to construct the incremental

Lyapunov function using contraction theory, and gives an analytical solution for the optimal control input in a distributed manner. This makes CaRT end-to-end trainable with a faster computation evaluation time.

The trade-off of Sec. II-B.2 and the strengths implied in Sec. II-B.3 will be demonstrated in Sec. V.

III. ANALYTICAL FORM OF OPTIMAL SAFETY FILTER

In this section, we derive an analytical way to design a safety filter that guarantees safe operations of the systems (1) and (2) when $d^i = 0$ and $\Gamma^i = 0$. One of the benefits of the log-barrier formulation in the following is that the global safety violation can be decomposed as the sum of the local safety violations, allowing for the distributed implementation of our analytical safety filter in a multi-agent setting.

Although (5) considers collision-free operation as the objective of safety to show one example of its use, we remark that all the proofs to be discussed work also with general notions of safety with a slight modification with (7). Collision avoidance is just one of the most critical safety requirements in a multi-agent setting.

A. Multi-Agent Safety Certificates

We use the following for local safety certificates:

$$\psi^i(o^i) = -\log \prod_{j \in S^i} h(p^{ij}), \quad h(p^{ij}) = \frac{\|p_{ij}\|_{\Xi} - (r_s + \Delta r_s)}{r_{\text{sen}} - (r_s + \Delta r_s)} \quad (5)$$

where $p^{ij} = p^j - p^i$, S^i is the set of all the neighboring objects (i.e., agents and obstacles), r_s is the minimal safe distance between i th agent and other objects, r_{sen} is the sensing radius, and $\Delta r_s > 0$ is a positive scalar parameter to account for the external disturbance to be formally defined in Sec. IV. We use the weighted 2-norm $\|\cdot\|_{\Xi}$ with the weight $\Xi > 0$ to consider a collision boundary defined by an ellipsoid. The parameters are selected to satisfy $\sup_{\|p_{ij}\| \leq r_{\text{sen}}} \|p_{ij}\|_{\Xi} \leq r_{\text{sen}}$ and $r_{\text{sen}} - (r_s + \Delta r_s) > 0$ to ensure $\psi^i(o^i) \geq 0$ always when the distance to the j th object is less than r_{sen} . Having a negative value of h implies a safety violation. The idea for our CaRT safety filter is first to construct a safe target velocity given as

$$v_d^i(o^i) = -k_p \nabla_{p^i} \psi^i(o^i) \quad (6)$$

where $k_p > 0$, and then track it optimally using the knowledge of Lagrangian systems and contraction theory.

When dealing with general safety $h(x^i, x^j) \geq 0$ for each i and j , we can also use the local safety function (5) modified as

$$\psi^i(o^i) = -\log \prod_{j \in S^i} h(x^i, x^j). \quad (7)$$

As mentioned earlier, our framework to be discussed can handle general notions of safety with just a slight modification using (7) instead of (5).

Remark 1. The velocity (6) renders the single integrator system (i.e., $\dot{p}^i = v^i$) safe [5]. Note that instead of the condition $\dot{\psi} \leq 0$ in [5], [12], we could use $\dot{\psi} \leq \alpha(h)$ to increase the available set of control inputs [4], where ψ is

a barrier function, h is a safety function associated with ψ , and α is a class \mathcal{K} function [13, p. 144]. This requires an additional global Lipschitz assumption on α as seen in [10].

B. Optimal Safety Filter for Lagrangian Systems

Given the learned motion planning policy $u_{\ell}^i(o^i)$ of Sec. II-A for the system (1), we design a control policy u_d^i processed by our CaRT safety filter as follows:

$$u_d^i(o^i) = u_{\ell}^i(o^i) - \begin{cases} \frac{e_v^i (u_{\ell}^i(o^i) - \bar{u}_d^i)^{\top} e_v^i}{\|e_v^i\|^2} & \text{if } (u_{\ell}^i(o^i) - \bar{u}_d^i)^{\top} e_v^i > 0 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where $e_v^i = v^i - v_d^i(o^i)$ for $v_d^i(o^i)$ of (6), \bar{u}_d^i is given as

$$\bar{u}_d^i = M^i \hat{v}_d^i + C^i v_d^i + G^i + D^i + v_d^i - k_v M^i e_v^i \quad (9)$$

with its arguments omitted and $k_p, k_v > 0$ being design parameters. The controller (8) is well-defined even with the division by $\|e_v^i\|$ as the relation $(u_{\ell}^i(o^i) - \bar{u}_d^i)^{\top} e_v^i = 0 \leq 0$ holds when $\|e_v^i\| = 0$. We have the following for the safety guarantee.

Theorem 1. Consider the following optimization problem:

$$u_{\text{opt}}^i = \arg \min_{u^i \in \mathbb{R}^m} \|u^i - u_{\ell}^i(o^i)\|^2 \text{ s.t. } (u^i - \bar{u}_d^i)^{\top} e_v^i \leq 0. \quad (10)$$

where \bar{u}_d^i is given by (9). Suppose that there exists a control input u^i that satisfies the constraint of (10) for each i . The safety of the system (1) is then guaranteed when $d^i = 0$ and $\Gamma^i = 0$, i.e., all the agents will not collide with the other objects when there is no external disturbance.

Also, the problem (10) is always feasible and its solution is given by $u_{\text{opt}}^i = u_d^i(o^i)$ for u_d^i of (8), thereby minimizing the deviation of the safe control input from $u_{\ell}^i(o^i)$.

Proof. Let us consider the following Lyapunov-type function:

$$\mathcal{V}_s = k_p \psi(X) + \sum_{i=1}^N \frac{\|v^i - v_d^i\|_{M^i(p^i)}^2}{2}$$

where v_d^i is given in (8) and ψ is given as

$$\psi(X) = -\log \left\{ \left(\prod_{i=1}^N \prod_{\substack{j>i \\ j \in \mathcal{N}^i}} h(p^{ij}) \right) \left(\prod_{i=1}^N \prod_{j \in \mathcal{M}^i} h(p^{ij}) \right) \right\} \quad (11)$$

for $X = \{x^1, \dots, x^N\}$. By the definition of h in (5), the safe operation of the system (1) is guaranteed as long as ψ^i of (5) is bounded. Taking the time derivative of \mathcal{V}_s , we get

$$\dot{\mathcal{V}}_s = \sum_{i=1}^N k_p \nabla_{p^i} \psi^i \top v_d^i - k_v \|e_v^i\|_{M^i}^2 + (u^i - \bar{u}_d^i)^{\top} e_v^i \quad (12)$$

by using (1) for $d^i = 0$ and $\Gamma^i = 0$ with the relation $z^{\top}(\dot{M}^i - 2C^i)z = 0$, $\forall z \in \mathbb{R}^n$, where $e_v^i = v^i - v_d^i$ and the arguments are omitted. Having $(u^i - \bar{u}_d^i)^{\top} e_v^i \leq 0$ as in the constraint of (10) gives $\dot{\mathcal{V}}_s \leq \sum_{i=1}^N -k_p^2 \|\nabla_{p^i} \psi^i\|^2 - k_v \|e_v^i\|_{M^i}^2$, which guarantees the boundedness of V and then ψ^i for all i , implying no safety violation as long as the system is initially safe. Also, the constraint $(u^i - \bar{u}_d^i)^{\top} e_v^i \leq 0$ is always feasible for u^i given as $u^i = \bar{u}_d^i$. Finally, applying the KKT condition [1, pp. 243-244] to (10) results in $u_{\text{opt}}^i = u_d^i(o^i)$ for $u_d^i(o^i)$ given in (8). \square

C. Optimal Safety Filter for General Nonlinear Systems

This section generalizes the results of Theorem 1 for nonlinear systems given by (2).

1) *Incremental Lyapunov Function:* We use the following incremental Lyapunov function inspired by contraction theory [2], [3], which leads to a safety analysis analogous to that of LTV systems and Lagrangian systems (1):

$$\mathcal{E}^i(o^i, t) = \frac{\|v^i - v_d^i(o^i)\|_{M^i(x^i, t)}^2}{2} \quad (13)$$

$$\text{s.t. } \dot{M}^i + M^i A_d^i + A_d^{i\top} M^i - 2M^i B^i R^{i-1} B^{i\top} M^i \leq -k_v M^i \quad (14)$$

where $M^i(x^i, t) > 0$, k_v is a design parameter, and $A_d^i(o^i, t)$ is a nonlinear state-dependent coefficient matrix for the system of (2) defined as follows:

$$A_d^i(o^i, t)(v^i - v_d^i(o^i)) = f^i(p^i, v^i, t) - f^i(p^i, v_d^i(o^i), t).$$

The arguments are omitted in (14) and the notations (13) and (14) are intentionally consistent with the ones of Theorem 1 to imply the analogy between the methods in Sec. III-B and Sec. III-C. Note that the nonlinear matrix A_d always exists when f^i is continuously differentiable [14].

The underlying benefit of using contraction theory here is that the problem of finding M^i in (14) can be expressed as a convex optimization problem for optimal disturbance attenuation [15]. Although we use one of the simple versions of a contraction metric in this paper for simplicity of discussion, we can consider more general types of Lyapunov functions and contraction metrics for the sake of the broader applicability of our approach. Reviewing how to find M^i in (14) is beyond the scope of this paper, but those interested in knowing more about it can refer to [3], [8] and references therein.

2) *Augmenting General Learned Policy with Safety:* Let us first introduce the following assumption for generalizing the result of Theorem 1.

Assumption 1. Consider a Lyapunov-type function for (2) defined as follows:

$$\mathcal{V}_s^i(o^i, t) = k_p \psi^i(o^i) + \frac{\|v^i - v_d^i(o^i)\|_{M^i(x^i, t)}^2}{2}$$

for $k_p > 0$, $\psi^i(o^i)$ of (5), $v_d^i(o^i)$ of (6), and $M^i(x^i, t)$ of (14). We assume that

$$\begin{aligned} e_v^i &= B(p^i, v^i, t)^\top \nabla_{v^i} \mathcal{V}_s^i = B(p^i, v^i, t)^\top M^i(v^i - v_d^i) = 0 \\ \Rightarrow \dot{\mathcal{V}}_s^i &\leq -k_p^2 \|\nabla_{p^i} \psi^i(o^i)\|^2 - k_v \mathcal{E}^i(o^i, t) \end{aligned} \quad (15)$$

where $e_v^i = v^i - v_d^i$ for v_d^i of (6), $\psi^i(o^i)$ is given in (5), α is given in (6), $\mathcal{E}^i(o^i, t)$ is given in (13), and the arguments of \mathcal{V}_s and M^i are omitted for notational simplicity.

This assumption simply says that the system naturally satisfies the safety condition (16) when the velocity displacements are in the directions orthogonal to the span of the actuated directions as discussed in [16].

Remark 2. Assumption 1 always holds for fully and over-actuated systems s.t. $B(p^i, v^i, t)B(p^i, v^i, t)^\dagger = I_{n \times n}$, where $B(p^i, v^i, t)^\dagger$ the Moore-Penrose pseudo inverse. Furthermore, even when the system is under-actuated and Assumption 1

does not hold, the error associated with the under-actuation can be treated robustly as to be seen in Sec. IV.

Given the learned motion planning policy $u_\ell^i(o^i)$ of Sec. II-A for the system (2), we design a control policy u_d^i processed by our CaRT safety filter as (8), where e_v^i is now given by (15) and \bar{u}_d^i is defined as

$$\bar{u}_d^i = \begin{cases} \frac{e_v^i (e_v^{i\top} M^i (\dot{v}_d^i - f_d^i) - k_p e_v^{i\top} \nabla_{p^i} \psi^i)}{\|e_v^i\|^2} & \text{if } \|e_v^i\| \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

with $f_d^i = f^i(p^i, v_d^i, t)$ and k_p being a design parameter. The controller (8) is well-defined with the division by $\|e_v^i\|$ under Assumption 1, because the relation $(u_\ell^i(o^i) - \bar{u}_d^i)^\top e_v^i = 0 \leq 0$ holds when $\|e_v^i\| = 0$.

Theorem 2. Consider the optimization problem (10), where u_ℓ^i is now given by the learned motion planning policy for (3) with (2), e_v^i is by (15), \bar{u}_d^i is by (17). Suppose that Assumption 1 holds and that there exists a control input u^i that satisfies the constraint of (10) for each i . The safety of the system (2) is then guaranteed when $d^i = 0$ and $\Gamma^i = 0$, i.e., all the agents will not collide with the other objects when there is no external disturbance.

Also, the problem (10) is always feasible and its solution is given by $u_{\text{opt}}^i = u_d^i(o^i)$ for u_d^i of (8), thereby minimizing the deviation of the safe control input from $u_\ell^i(o^i)$.

Proof. Let us consider a Lyapunov-type function $\mathcal{V}_s = \psi(X) + \sum_{i=1}^N \mathcal{E}^i(o^i, t)$, where ψ is given in (11), v_d^i is given in (8), and $\mathcal{E}^i(o^i, t)$ is given in (13). Using the relation (14) and Assumption 1, we have $\dot{\mathcal{V}}_s \leq \sum_{i=1}^N -k_p^2 \|\nabla_{p^i} \psi^i\|^2 - k_v \mathcal{E}^i(o^i, t) + (u^i - \bar{u}_d^i)^\top e_v^i$ when $d^i = 0$ and $\Gamma^i = 0$ in (2) as in the proof of Theorem 1, where $\psi^i(o^i)$ is given in (5), α is given in (6), and $\mathcal{E}^i(o^i, t)$ is given in (13). The rest follows from the proof of Theorem 1 below (12). \square

Remark 3. The safety filter of Theorem 2 minimizes the deviation of the safe control input from the learned motion planning input of the general system (2), which implies that it instantaneously minimizes the contribution of the under-actuation error when Assumption 1 does not hold. This error can be then treated robustly as discussed in Remark 2.

IV. ROBUSTNESS TO EXTERNAL DISTURBANCE

The results in Sec. III depend on the assumption that $d^i = 0$ and $\Gamma^i = 0$ in (1) and (2). This section discusses the safety of these systems in the presence of external disturbance.

A. Revisiting Built-in Robustness of Safety Filter

Due to its Lyapunov-type formulation in Theorems 1 and 2, our safety filter inherits the robustness properties discussed in, e.g., [13]. It also inherits the robustness of the barrier function of [4], [9], [10] ([17], [18] for stochastic disturbance) as seen in Sec. II-A with Fig. 1 and 2.

This section is for hierarchically augmenting such built-in robustness with the tracking-based robustness as in the tube-based motion planning [19], [20] to lighten the burden of the safety filter in dealing with the disturbance. Given a safety

condition $h \geq 0$, these two ways of augmenting the learned motion planning with robustness as in Fig. 1 are achieved by

- (a) changing our safety filter parameters (e.g., making k_p and k_v larger in Theorems 1 and 2)
- (b) tracking a safe trajectory that satisfies $h \geq 0$, ensuring the deviation from the perturbed trajectory is finite.

The first approach (a) could lead to a large repelling force due to the stability of the safe set originating from the use of \dot{h} , especially when we use the log-barrier formulation with the dynamics discretization (i.e., we get a larger control input as the agents get closer to the safety boundary, implying a large discretization error). In contrast, (b) does not involve such behavior as it simply attempts to track the safe trajectory satisfying $h \geq 0$. As illustrated in Fig. 2, there are the following two sources of robustness in our approach:

- (a) asymptotic/exponential stability of the safe set
- (b) incremental asymptotic/exponential stability of the system trajectory with respect to a safe target trajectory

and this section is about (b), which significantly reduces the responsibility of the safety filter (a) in meeting the robustness requirement, allowing the safe set to be less stable (i.e., the unsafe set to be less repelling, meaning more freedom in choosing the safety filter parameters). These observations will be more appreciable in the numerical simulations in Sec. V.

B. Optimal Robust Filter for Lagrangian Systems

Let us consider the following control policy $u_r^i(o^i, t)$ for the CaRT robust filter of the Lagrangian system (1):

$$u_r^i(o^i, t) = u_d^i(t) - \begin{cases} \frac{s^i (u_d^i(t) - \bar{u}_r^i)^T s^i}{\|s^i\|^2} & \text{if } (u_d^i(t) - \bar{u}_r^i)^T s^i > 0 \\ 0 & \text{otherwise} \end{cases} \quad (18)$$

where $s^i = (v^i - v_d^i(t)) + \Lambda_r^i(p^i - p_d^i(t))$ for a positive definite position control gain $\Lambda_r^i > 0$, $p_d^i(t)$, $v_d^i(t)$, and $u_d^i(t)$ are the safe target position, velocity, and control input computed by integrating (1) with the safe control input (8) assuming $d^i = 0$ and $\Gamma^i = 0$, respectively, and \bar{u}_r^i is given as

$$\bar{u}_r^i = M^i \dot{s}^i + C^i s^i + G^i + D^i - k_r^i M^i s^i$$

with its arguments omitted for simplicity and $k_r^i > 0$ being a scalar control gain for the composite state error s^i . We have the following result. Note that $x_d^i = [p_d^i{}^T, v_d^i{}^T]^T$ corresponds to x_d of the conceptual illustration in Fig. 1.

Theorem 3. Suppose that the system (1) is controlled by (18) and that the target safe trajectory of each agent is expressed with $p_d^i(t)$, $v_d^i(t)$, and $u_d^i(t)$ of (18). If there exist bounded positive constants \underline{m}^i , \bar{m}^i , \bar{d}_s^i , $\bar{m}_{x_1}^i$, and $\bar{m}_{x_2}^i$ satisfying $\underline{m}^i \mathbb{I} \leq M^i \leq \bar{m}^i \mathbb{I}$, $\|(M^i)^{-1} \Gamma^i\|_F^2 \leq \bar{d}_s^i$, $\|\partial M^i / \partial x_k\| \leq \bar{m}_{x_1}^i$, and $\|\partial^2 M^i / (\partial x_k \partial x_\ell)\| \leq \bar{m}_{x_2}^i$, $\forall x, t$, then there exists an appropriate set of the control parameters and the positive scalar Δr_s of (5), which guarantee a safe operation of the system (1) at time t with a finite probability, even under the presence of external disturbance.

Also, the controller (18) is the optimal solution to

$$u_{\text{opt}}^i = \arg \min_{u^i \in \mathbb{R}^m} \|u^i - u_d^i(t)\|^2 \text{ s.t. } (u^i - \bar{u}_r^i)^T s^i \leq 0.$$

i.e., $u_{\text{opt}}^i = u_r^i(o^i, t)$, which is always feasible with $(u^i - \bar{u}_r^i)^T s^i \leq 0$ representing an incremental exponential stability condition, and thus it minimizes the deviation of the robust control input from the learned safe control input $u_d^i(t)$.

Proof. The first part follows from the incremental stability analysis with the Lyapunov function $\mathcal{V}_r^i = s^{iT} M(p^i) s^i$. Applying the weak infinitesimal operator \mathcal{A} [17, p. 9] for analyzing the time evolution of (1), we get

$$\mathcal{A} \mathcal{V}_r^i \leq -(k_r^i - \underline{m}^{i-1} ((\bar{d}^i + \bar{d}_s^i \bar{m}_{x_1}^i) \varepsilon_d^{i-1} + \bar{d}_s^i \bar{m}_{x_2}^i / 2)) \mathcal{V}_r^i + \underline{m}^i C_d^i$$

where \bar{d}^i is given in (1), $\varepsilon_d^i > 0$ is a design parameter, and $C_d^i = (\bar{d}_s^i \bar{m}^i + \varepsilon_d^i (\bar{d}^i + \bar{d}_s^i \bar{m}_{x_1}^i)) / \underline{m}^i$. If we select k_r^i and ε_d^i to have $2\bar{k}_r^i \leq (k_r^i - \underline{m}^{i-1} ((\bar{d}^i + \bar{d}_s^i \bar{m}_{x_1}^i) \varepsilon_d^{i-1} + \bar{d}_s^i \bar{m}_{x_2}^i / 2))$ for some $\bar{k}_r^i > 0$, then the application Dynkin's formula [17, p. 10] and the Gronwall-type lemma [21] gives

$$\mathbb{E}[\|s^i\|] \leq \sqrt{\frac{C_d^i}{2\bar{k}_r^i}} + \sqrt{\left[\mathbb{E}[\mathcal{V}_r^i|_{t=0}] - \frac{C_d^i}{2\bar{k}_r^i} \right]^+} e^{-\bar{k}_r^i t} = a^i + b^i e^{-\bar{k}_r^i t}$$

where $[\cdot]^+ = \max\{0, \cdot\}$. The hierarchical structure of the controller (18) along with Markov's inequality [22, pp. 311-312] yields $\mathbb{P}[\|p^i(t) - p_d^i(t)\| > D_s^i] \leq D_{\mathbb{B}}^i / D_s^i$, where $D_s^i > 0$, $D_{\mathbb{B}}^i = \mathbb{E}[\|e_d^i(0)\|^2] e^{-\lambda_r^i t} + a^i (1 - e^{-\lambda_r^i t}) / \lambda_r^i + b^i (e^{-\bar{k}_r^i t} - e^{-\lambda_r^i t}) / (\lambda_r^i - \bar{k}_r^i)$, and $\lambda_r^i \mathbb{I} \leq \Lambda_r^i$. Since $p_d(t)$ is a safe trajectory, selecting Δr_s of (5) as $\Delta r_s = D_s^i$ guarantees safety at time t with probability $[1 - D_{\mathbb{B}}^i / D_s^i]^+$.

The optimality and feasibility argument follows as in the proof of Theorem 1. \square

Remark 4. The safety argument of Theorem 3 is based on the assumption that the target trajectory is safe for all t . Therefore, each agent updates the safe target trajectory as in [23] if new obstacles/agents are reported or their goal region is changed. We can also incorporate the uncertainty associated with this fact by modifying the minimal safe distance r_s of (5) larger than the actual one, or by formulating and solving stochastic motion planning problems in training the learned motion planning policy $u_e^i(o^i)$ of Sec. II-A. If the disturbance and the uncertainty are too large to be treated robustly, then we can further generalize our approach using nonlinear adaptive control and system identification for nonlinear systems (see, e.g., Sec. 8 and 9 of [3]).

C. Optimal Robust Filter for General Nonlinear Systems

Due to the existence of the matrix function M^i in (14), we can readily obtain a robustness result of Theorem 3 even for general nonlinear systems (2), simply by replacing each u of (18) with u as we derived the results of Theorem 2 analogously to Theorem 1. We thus omit the derivation of the robustness guarantee in this paper due to space limitation without repeating the argument, but those interested can refer to the existing literature such as [3], [15], [24] which explicitly discusses the error bound as the one of Theorem 3.

V. NUMERICAL SIMULATION

This section demonstrates the effectiveness of CaRT for several motion planning and control problems.




Fig. 3. Position trajectories of the nonlinear system (2) with (19) with small and large disturbance, $\bar{d}^i = \bar{\gamma}^i = 2.0 \times 10^{-5}$ and $\bar{d}^i = \bar{\gamma}^i = 2.0 \times 10^{-2}$ in (2), respectively, which capture the trade-off discussed in (a) and (b) of Sec. II-B.2.

A. Illustrative Examples

1) *General Nonlinear System:* Let us first consider a nonlinear dynamical system (2) with f^i given as

$$f^i(p^i, v^i, t) = \begin{bmatrix} \cos(p_1^i) p_2^i - v_1^i + v_2^i \\ -\sin(p_1^i) p_1^i v_2^i + (v_1^i)^2 - v_2^i - 2v_1^i v_2^i \end{bmatrix} \quad (19)$$

where $N = 1$ (number of agents) and $M = 5$ (number of obstacles). Our safety filter and robust filter are constructed using Theorem 2 and the methods outlined in Sec. IV-C, where the state-dependent coefficient matrix is used to construct the contraction metric of (14) as in [3], [15]. This example does not have a clear physical interpretation, but finding the incremental Lyapunov function is non-trivial as the system is nonlinear and non-polynomial and the set of CLFs is non-convex [16].

The reference trajectory shown in the left-hand side of Fig. 3 is the one generated simply by a baseline contraction-based tracking control law [3], [15] with the target trajectory being the stationary origin $x_d(t) = [0, 0, 0, 0]^T$, $\forall t$. The initial condition is selected to be $x(0) = [1.0, 1.0, 1.0, 1.0]^T$. Because of the lack of safety consideration, it violates a safety constraint at the position indicated by \times . When the disturbance is small ($\bar{d}^i = \bar{\gamma}^i = 2.0 \times 10^{-5}$), our baseline safety filter works without the robust filter thanks to its built-in robustness even with the control interval $dt = 0.1$, as we discussed in (a) of Sec. II-B.2. This is as expected from the argument of Sec. II-B.3.

When \bar{d} and $\bar{\gamma}$ get larger ($\bar{d}^i = \bar{\gamma}^i = 2.0 \times 10^{-2}$) as in the right-hand side of Fig. 3, however, the baseline safety filter becomes too sensitive to the disturbance for the control interval $dt = 0.1$, leading to a large control input and safety violation indicated by \times . This situation can be avoided by using a smaller control interval $dt = 0.01$, but still, its control input for safety becomes more dominant than the ideal control input for the reference trajectory even in this case, taking longer to reach the target position as can be seen from the green trajectory of Fig. 3. The combination of the safety filter and robust filter, CaRT, indeed allows for considering robustness separately when applying the safety filter, thereby handling large disturbances without violating safety and losing too much of the reference control performance, as discussed in (b) of Sec. II-B.2.

2) *Thruster-based Spacecraft Simulators:* Such a trade-off is more evident in a practical multi-agent robotic system, where we cannot use a smaller control time interval due to hardware limitations. We next consider a nonlinear spacecraft simulator system given in [25] with $N = 6$ (number of agents), $M = 10$ (number of obstacles), and $r_{\text{sen}} = 2.0$ (m) (sensing

radius), where the control time interval dt is required to be $dt \geq 0.1$ (s). The dynamics parameters are normalized to 1.

The learned motion planning policy detailed in Sec. II-A is constructed using a neural network used in [26] with the training process outlined in [5]. It utilizes the centralized global solution data sampled by solving (3) using, e.g., the sequential convex programming, for the decentralized approximation by the neural network with the local observation (4). The initial and target states of the spacecraft and the positions of the circular static obstacles in Fig. 4 are randomized during the training and simulation. The cost function for the objective function is selected as $c(x(\tau), u(\tau), \tau) = \sum_{i=1}^N \|u^i(\tau)\|^2$ in (3).

As shown in Fig. 4, we can indeed observe the differences between each approach as discussed in Sec. II-B with Fig. 1 and 2. When the size of the disturbance is relatively smaller than the learning error ((ii) of Fig. 4), then our baseline safety filter works with its built-in robustness (see (a) of Sec. II-B.2). The loss of optimality in control results from the presence of disturbance and its decentralized implementation with distributed information. As the disturbance gets larger ((iii) – (v) of Fig. 4), the baseline safety filter starts to fail. Also, the CLF-CBF approach starts to yield excessively large control input even when its QP is solved with a control input constraint ($|u^i_k| \leq 1.00$, $k = 1, \dots, n$), which results in additional computational burden for each agent (see (a) – (b) of Sec. II-B.3 with Fig 2). Task failure is defined as the situation where at least one of the spacecraft does not reach the target state. Even when the size of the disturbance is relatively larger than the learning error ((vi) of Fig. 4), CaRT, the safety filter equipped with the robust filter, still works, retaining its control effort 4.96 times smaller than that of the CLF-CBF approach and 8.10 times greater than that of the optimal solution (see (b) of Sec. II-B.2 with Fig. 1 and 2).

B. Multi-Spacecraft Reconfiguration

As a part of JPL's CASTOR project, let us consider the optimal reconfiguration in Low Earth Orbit (LEO) with 10 spacecraft ($N = 10$ and $M = 0$). The distributed motion planning policy of Sec. II-A is trained again using [5] as discussed in Sec. V-A.2 with $r_{\text{sen}} = 2.0$ (m) for the sensing radius. Its dynamical system can be expressed as a Lagrangian system (1) as in [27]. Task success is defined as the situation where the agent safety reaches a given target terminal state x_f within a given time horizon. The success rate is computed as the percentage of successful trials in the total 50 simulations. The initial and target states are randomized during the training and simulation. Also, the cost function of (3) is again selected as $c(x(\tau), u(\tau), \tau) = \sum_{i=1}^N \|u^i(\tau)\|^2$ in (3).

As implied in Fig. 5 and as discussed in Sec. V-A.2, we can still see that the baseline safety filter works for small disturbances, and CaRT, the safety filter equipped with the robust filter, works for large disturbances. Such an observation can be corroborated by the results summarized in Table I. In particular, CaRT augments the learned motion planning policy with safety and robust tracking, resulting in its success rate of 100 % with its control effort 4.64 times greater than that of the optimal solution. Again, the loss of optimality in control results




Fig. 4. Position trajectories of the spacecraft simulator system [25] with various disturbances (small: $\bar{d}^i = \bar{\gamma}^i = 3.0 \times 10^{-4}$, medium: $\bar{d}^i = \bar{\gamma}^i = 5.0 \times 10^{-3}$, and large: $\bar{d}^i = \bar{\gamma}^i = 5.0 \times 10^{-2}$ in (2)), which capture the trade-off discussed in Sec. II-B.




Fig. 5. Position trajectories for the multi-spacecraft reconfiguration task in LEO [27] with various disturbances (small: $\bar{d}^i = \bar{\gamma}^i = 1.0 \times 10^{-2}$ and large: $\bar{d}^i = \bar{\gamma}^i = 5.0 \times 10^{-2}$ in (1)).

TABLE I

CONTROL PERFORMANCES FOR THE MULTI-SPACECRAFT RECONFIGURATION IN LEO AVERAGED OVER 50 SIMULATIONS, WHERE $J = \int_0^{t_f} \sum_{i=1}^N \|u^i\|^2 d\tau$.

| guidance and control methods | success rate (%) | control effort J |
|--------------------------------|------------------|--------------------|
| global solution (w/o disturb.) | 100 | 1.26 |
| safety filter (small disturb.) | 94.0 | 4.93 |
| CLF-CBF QP (lrg. disturb.) [4] | 31.6 | 10.36 |
| CaRT (large disturb.) | 100 | 5.86 |

*small disturb.: $\bar{d}^i = \bar{\gamma}^i = 1.0 \times 10^{-2}$ and large disturb.: $\bar{d}^i = \bar{\gamma}^i = 5.0 \times 10^{-2}$ in (1).

from the presence of disturbance and the CaRT's decentralized implementation with distributed information.

VI. CONCLUSION

In this paper, we present CaRT, a hierarchical, control-theoretic filter for augmenting machine-learning motion planning algorithms with certified safety and robust tracking guarantees, for a large class of multi-agent nonlinear dynamical systems. Unlike existing methods, CaRT uses a safety filter, which steers the system into the safe set, to certify the safety of the learned policy, then uses a robust filter, which steers the system into the safe trajectory, to reject deterministic and stochastic disturbances. As demonstrated in numerical experiments, this hierarchical nature allows CaRT to guarantee safety and robustness under much larger disturbances in off-nominal settings. This makes a major distinction from conventional safety-driven approaches including CLF-CBF, where the robustness could originate from the over-conservative pulling force to the interior of the safe set.

REFERENCES

- [1] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, Mar. 2004.
- [2] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for nonlinear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [3] H. Tsukamoto, S.-J. Chung, and J.-J. E. Slotine, "Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview," *Annual Reviews in Control*, vol. 52, pp. 135–169, 2021.
- [4] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, 2017.
- [5] B. Rivière, W. Hönig, Y. Yue, and S.-J. Chung, "GLAS: Global-to-local safe autonomy synthesis for multi-robot motion planning with end-to-end learning," *IEEE Robot. Automat. Lett.*, vol. 5, no. 3, pp. 4249–4256, 2020.
- [6] J.-J. E. Slotine and W. Li, *Applied Nonlinear Control*. Upper Saddle River, NJ: Pearson, 1991.
- [7] A. J. Taylor, V. D. Dorobantu, H. M. Le, Y. Yue, and A. D. Ames, "Episodic learning with control Lyapunov functions for uncertain robotic systems," in *IEEE/RSJ IROS*, 2019, pp. 6878–6884.
- [8] C. Dawson, S. Gao, and C. Fan, "Safe control with learned certificates: A survey of neural Lyapunov, barrier, and contraction methods for robotics and control," *IEEE Trans. Robot.*, pp. 1–19, 2023.
- [9] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, "Robustness of control barrier functions for safety critical control," *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015.
- [10] A. J. Taylor, P. Ong, T. G. Molnar, and A. D. Ames, "Safe backstepping with control barrier functions," in *IEEE CDC*, 2022, pp. 5775–5782.
- [11] L. Arnold, *Stochastic Differential Equations: Theory and Applications*. Wiley, 1974.
- [12] P. F. Hokayem, D. M. Stipanović, and M. W. Spong, "Coordination and collision avoidance for Lagrangian systems with disturbances," *Applied Mathematics and Computation*, vol. 217, no. 3, pp. 1085–1094, 2010.
- [13] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ: Prentice-Hall, 2002.
- [14] H. Tsukamoto, S.-J. Chung, and J.-J. E. Slotine, "Neural stochastic contraction metrics for learning-based control and estimation," *IEEE Control Syst. Lett.*, vol. 5, no. 5, pp. 1825–1830, 2021.
- [15] H. Tsukamoto and S.-J. Chung, "Robust controller design for stochastic nonlinear systems via convex optimization," *IEEE Trans. Autom. Control*, vol. 66, no. 10, pp. 4731–4746, 2021.
- [16] I. R. Manchester and J.-J. E. Slotine, "Control contraction metrics: Convex and intrinsic criteria for nonlinear feedback design," *IEEE Trans. Autom. Control*, vol. 62, no. 6, pp. 3046–3053, Jun. 2017.
- [17] H. J. Kushner, *Stochastic Stability and Control*. Academic Press New York, 1967.
- [18] A. Clark, "Control barrier functions for stochastic systems," *Automatica*, vol. 130, p. 109688, 2021.
- [19] D. Q. Mayne, E. C. Kerrigan, E. J. van Wyk, and P. Falugi, "Tube-based robust nonlinear model predictive control," *Int. J. Robust Nonlinear Control*, vol. 21, no. 11, pp. 1341–1353, 2011.
- [20] S. Singh, H. Tsukamoto, B. T. Lopez, S.-J. Chung, and J.-J. E. Slotine, "Safe motion planning with tubes and contraction metrics," in *IEEE CDC*, 2021, pp. 2943–2948.
- [21] Q.-C. Pham, "A variation of Gronwall's lemma," arXiv:0704.0922, 2007.
- [22] D. S. Geoffrey R. Grimmett, *Probability and Random Processes*, 3rd ed. United Kingdom: Oxford University Press, 2001.
- [23] S. Singh, A. Majumdar, J.-J. E. Slotine, and M. Pavone, "Robust online motion planning via contraction theory and convex optimization," in *IEEE ICRA*, May 2017, pp. 5883–5890.
- [24] H. Tsukamoto, S.-J. Chung, J.-J. Slotine, and C. Fan, "A theoretical overview of neural contraction metrics for learning-based control with guaranteed stability," in *IEEE CDC*, 2021, pp. 2949–2954.
- [25] Y. Nakka, R. Foust, E. Lupu, D. Elliott, I. Crowell, S.-J. Chung, and F. Hadaegh, "A six degree-of-freedom spacecraft dynamics simulator for formation control research," in *Astrodyn. Special Conf.*, 2018.
- [26] H. Tsukamoto and S.-J. Chung, "Learning-based robust motion planning with guaranteed stability: A contraction theory approach," *IEEE Robot. Automat. Lett.*, vol. 6, no. 4, pp. 6164–6171, 2021.
- [27] S.-J. Chung, U. Ahsun, and J.-J. E. Slotine, "Application of synchronization to formation flying spacecraft: Lagrangian approach," *J. Guid. Control Dyn.*, vol. 32, no. 2, pp. 512–526, 2009.