

Learning Robustness with Bounded Failure: An Iterative MPC Approach

Monimoy Bujarbaruah,^{†,*} Akhil Shetty,^{†,*}
Kameshwar Poolla,^{*} and Francesco Borrelli^{*}

May 6, 2020

Abstract

We propose an approach to design a Model Predictive Controller (MPC) for constrained Linear Time Invariant systems performing an iterative task. The system is subject to an additive disturbance, and the goal is to learn to satisfy state and input constraints robustly. Using disturbance measurements after each iteration, we construct *Confidence Support* sets, which contain the true support of the disturbance distribution with a given probability. As more data is collected, the Confidence Supports converge to the true support of the disturbance. This enables design of an MPC controller that avoids conservative estimate of the disturbance support, while simultaneously bounding the probability of constraint violation. The efficacy of the proposed approach is then demonstrated with a detailed numerical example.

1 Introduction

As data-driven decision making and control becomes ubiquitous [1–3], system identification methods are being integrated with control algorithms for control of uncertain dynamical systems. The uncertainty in these systems can be typically attributed to two factors: (i) model uncertainty (eg. modeling mismatch and inaccuracies), and (ii) exogenous disturbances (eg. sensor noise). For such uncertain systems subject to state and input constraints, Model Predictive Control (MPC) [4–6] is a commonly used approach for ensuring robust constraint satisfaction.

The field of Adaptive MPC [7–11] deals with learning the model uncertainty to improve controller performance over time. These methods rely upon Set Membership approaches, which assume known set based bounds on the exogenous disturbances. As these disturbance supports are actually unknown in practice, conservative over-approximations are used for control design. This results in the controller either being infeasible, or incurring higher costs by following highly sub-optimal trajectories. This motivates learning the disturbance support over time in order to improve controller performance. In such cases, it is necessary

^{††} authors contributed equally to this work. The authors are with University of California Berkeley, Berkeley, CA, USA; E-mails: {monimoyb, shetty.akhil, poolla, fborrelli}@berkeley.edu.

to allow the possibility of *failure*, i.e, violation of imposed constraints. Such violations are acceptable for certain non safety critical robotic applications.

To that end, numerous works in MPC literature have considered constructing probabilistic approximations of both the model uncertainty and disturbance support [12–16], allowing room for violations of imposed constraints with a certain probability. Methods such as [13–16], utilize Gaussian Process (GP) Regression to model and update the uncertainty in the system. However, they have no theoretical bounds for rate of constraint violations by the closed loop system over time.

Assuming the presence of *only* exogenous disturbances, [12] addresses this issue by constructing disturbance support sets offline using the scenario approach [17, Chapter 12]. This approach involves solving a scenario program with potentially large number of samples, which is computationally expensive. Moreover, the rate of constraint violation is dependent on the number of disturbance samples available offline. In certain settings (for eg., iterative tasks), it is often the case that one starts the controller having observed no samples apriori. While learning the disturbance support over time in such cases, it is desirable to have a user-specified upper bound for probability of failure over all time. The approach in [12] is unable to satisfy such an upper bound at all times, since the required number of samples could be unavailable during operation.

In this paper, we present an approach to design an MPC controller for constrained LTI systems performing an iterative task [18]. Like [12] we consider an additive disturbance in the system, under no uncertainty in the system matrices. Instead of considering a conservative over-approximation of the disturbance support such as [7, 8, 19], we learn this set from observed disturbance samples. While doing so, we guarantee a user-specified upper bound on the probability of failure over all iterations. Our main contributions can be summarized as:

- We introduce the notion of a *Confidence Support*, which is guaranteed to contain the true disturbance support with a specified probability. Constructing and updating the Confidence Supports after each iteration is computationally cheap, unlike [12].
- Using these Confidence Supports, we attempt robust MPC design and demonstrate satisfaction of desired upper bound on probability of failure in each iteration. For any value of user-specified upper bound on probability of failure, the controller is able to learn robust satisfaction of imposed constraints asymptotically, without suffering conservatism that is inherent to existing approaches [7, 8, 19].

2 Problem Formulation

We consider uncertain linear time-invariant systems of the form:

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad (1)$$

where $x_t \in \mathbb{R}^d$ is the state at time step t , $u_t \in \mathbb{R}^m$ is the input, and A and B are known system matrices of appropriate dimensions. At each time step t , the system is affected by an independently and identically distributed (i.i.d.) random disturbance $w_t \stackrel{\text{iid}}{\sim} \mathcal{P}$ with a

convex and compact support $\mathbb{W} \subset \mathbb{R}^d$. We aim to satisfy state and input constraints on the system robustly. We define $H_x \in \mathbb{R}^{s \times d}$, $h_x \in \mathbb{R}^s$, $H_u \in \mathbb{R}^{o \times m}$ and $h_u \in \mathbb{R}^o$. We can then write the imposed constraints for all time steps $t \geq 0$ as:

$$\mathbb{Z} := \{(x, u) : H_x x \leq h_x, H_u u \leq h_u\}. \quad (2)$$

Throughout the paper, we assume that system (1) performs the same task repeatedly for J number of times. Each task execution is referred to as *iteration*. Our goal is to design a controller that, at each iteration j , solves the finite horizon robust optimal control problem:

$$\begin{aligned} V^{j,\star}(x_S) = & \min_{u_0^j, u_1^j(\cdot), \dots} \sum_{t=0}^{T-1} \ell(\bar{x}_t^j, u_t^j(\bar{x}_t^j)) \\ \text{s.t.} \quad & x_{t+1}^j = Ax_t^j + Bu_t^j(x_t^j) + w_t^j, \\ & \bar{x}_{t+1}^j = A\bar{x}_t^j + Bu_t^j(\bar{x}_t^j), \\ & H_x x_t^j \leq h_x, \\ & H_u u_t^j \leq h_u, \\ & \forall w_t^j \in \mathbb{W}, \\ & x_0^j = x_S, \quad t = 0, 1, \dots, (T-1), \end{aligned} \quad (3)$$

where x_t^j , u_t^j and w_t^j denote the realized system state, control input and disturbance at time t of the j^{th} iteration respectively, and $(\bar{x}_t^j, u_t^j(\bar{x}_t^j))$ denote the disturbance-free nominal state and corresponding nominal input. Notice that (3) minimizes the nominal cost over a time horizon of length $T \gg 0$ in any j^{th} iteration with $j \in [J]$. Here we use $[J]$ to denote the set $\{1, 2, \dots, J\}$. We point out that, as system (1) is uncertain, the optimal control problem (3) consists of finding $[u_0^j, u_1^j(\cdot), u_2^j(\cdot), \dots]$, where $u_t^j : \mathbb{R}^d \ni x_t^j \mapsto u_t^j = u_t^j(x_t^j) \in \mathbb{R}^m$ are state feedback policies. As task duration $T \gg 0$, for computational tractability we try to approximate a solution to the optimal control problem (3), by solving a simpler constrained optimal control problem with prediction horizon $N \ll T$ in a receding horizon fashion.

In this work, we consider the support $\mathbb{W}$ of disturbance w_t^j to be an unknown convex and compact set. We estimate $\mathbb{W}$ using observed disturbance samples. At the start of iteration j , the estimated support is $\hat{\mathbb{W}}^j$.

3 Iterative MPC Problem

The MPC controller solves a finite horizon optimal control problem at each time step t in the j^{th} iteration. Since the disturbance support $\mathbb{W}$ is unknown and is estimated with $\hat{\mathbb{W}}^j$ built from data, robust satisfaction of (2) along the iteration is not guaranteed. This implies that the closed loop task execution might fail. We will formally define this notion of *failure* after defining the closed loop controller in this section.

We attempt to design a robust MPC controller in the j^{th} iteration with our best estimate

$\hat{\mathbb{W}}^j$ of disturbance support $\mathbb{W}$, by solving the following optimal control problem:

$$\begin{aligned}
V_{t \rightarrow t+N}^{\text{MPC},j}(x_t^j, \hat{\mathbb{W}}^j, \hat{\mathcal{X}}_N^j) := \\
\min_{U_t^j(\cdot)} \quad & \sum_{k=t}^{t+N-1} \ell(\bar{x}_{k|t}^j, v_{k|t}^j) + Q(\bar{x}_{t+N|t}^j) \\
\text{s.t.} \quad & x_{k+1|t}^j = Ax_{k|t}^j + Bu_{k|t}^j + w_{k|t}^j, \\
& \bar{x}_{k+1|t}^j = A\bar{x}_{k|t}^j + Bv_{k|t}^j, \\
& u_{k|t}^j = \sum_{l=t}^{k-1} M_{k,l|t}^j w_{l|t}^j + v_{k|t}^j, \\
& H_x x_{k|t}^j \leq h_x, \\
& H_u u_{k|t}^j \leq h_u, \\
& x_{t+N|t}^j \in \hat{\mathcal{X}}_N^j, \\
& \forall w_{k|t}^j \in \hat{\mathbb{W}}^j, \\
& \forall k = \{t, \dots, t+N-1\}, \\
& x_{t|t}^j = \bar{x}_{t|t}^j = x_t^j,
\end{aligned} \tag{4}$$

where in the j^{th} iteration, x_t^j is the measured state at time t , $x_{k|t}^j$ is the prediction of state at time k , obtained by applying predicted input policies $U_t^j(\cdot) = [u_{t|t}^j, \dots, u_{k-1|t}^j]$ to system (1) and $\{\bar{x}_{k|t}^j, v_{k|t}^j\}$ with $v_{k|t}^j = u_{k|t}^j(\bar{x}_{k|t}^j)$ denote the disturbance-free nominal state and corresponding input respectively. The MPC controller minimizes the cost over the predicted disturbance free nominal trajectory $\left\{ \{\bar{x}_{k|t}^j, v_{k|t}^j\}_{k=t}^{t+N-1}, \bar{x}_{t+N|t}^j \right\}$, which comprises of the positive definite stage cost $\ell(\cdot, \cdot)$, and the terminal cost $Q(\cdot)$. Notice, the above uses affine disturbance feedback parametrization [20] of input policies. We use state feedback to construct terminal set $\hat{\mathcal{X}}_N^j = \{x \in \mathbb{R}^d : \hat{Y}^j x \leq \hat{z}^j, \hat{Y}^j \in \mathbb{R}^{r^j \times d}, \hat{z}^j \in \mathbb{R}^{r^j}\}$, which is the $(T-N)$ step robust reachable set [6, Chapter 10] to set of state constraints in (2), obtained with a state feedback controller $u = Kx$, dynamics (1) and constraints (2). This set has the properties:

$$\begin{aligned}
\hat{\mathcal{X}}_N^j &\subseteq \{x | (x, Kx) \in \mathbb{Z}\}, \\
H_x((A+BK)^i x + \sum_{\tilde{i}=0}^{i-1} (A+BK)^{i-\tilde{i}-1} w_{\tilde{i}}) &\leq h_x, \\
H_u(K((A+BK)^i x + \sum_{\tilde{i}=0}^{i-1} (A+BK)^{i-\tilde{i}-1} w_{\tilde{i}})) &\leq h_u, \\
\forall x \in \hat{\mathcal{X}}_N^j, \forall w_i \in \hat{\mathbb{W}}^j, \forall i = 1, 2, \dots, (T-N).
\end{aligned} \tag{5}$$

After solving (4), in closed loop, we apply

$$u_t^j = v_{t|t}^{j,*} \tag{6}$$

to system (1). We then resolve the problem (4) again at the next $(t+1)$ -th time step, yielding a receding horizon strategy.

Remark 1. Computing sets such as (5) can become expensive in certain scenarios, where for example the number of constraints in $\mathbb{Z}$, or the dimension d of states is too large. In such cases one may opt for data driven methods such as [18, 21] to construct these terminal sets.

Assumption 1 (Well Posedness). We assume that given an initial state x_S , optimization problem (4) is feasible at all times $0 \leq t \leq T-1$ with true uncertainty support $\hat{\mathbb{W}}^j = \mathbb{W}$ for all iterations $j \in [J]$.

Since $\mathbb{W}$ is unknown and is being estimated with $\hat{\mathbb{W}}^j$ in the j^{th} iteration, we might lose the feasibility of (4) during $0 \leq t \leq T-1$. We formalize this with the following definition:

Definition 1 (State Constraint Failure). A State Constraint Failure at time step t in iteration j is the event

$$[\text{SCF}]_t^j : H_x x_t^j > h_x. \quad (7)$$

That is, a State Constraint Failure implies the violation of imposed constraints (2) by system (1) in closed loop with MPC controller (6).

Remark 2. Let $T^j < T$ denote the time step in the j^{th} iteration when a State Constraint Failure occurs. In that case, problem (4) becomes infeasible at T^j . We then stop the j^{th} iteration and update $\hat{\mathbb{W}}^j \xrightarrow{\text{update}} \hat{\mathbb{W}}^{j+1}$. When $T^j = T$, it denotes a successful iteration without any State Constraint Failure.

The probability of State Constraint Failure $[\text{SCF}]_t^j$ is a function of the sets $\hat{\mathbb{W}}^j$. In certain safety critical applications, it is necessary to keep the probability of $[\text{SCF}]_t^j$ very low, whereas in other applications a higher probability can be tolerated. However, it is not enough to focus on probability of $[\text{SCF}]_t^j$ alone. For example, a low probability of $[\text{SCF}]_t^j$ can be achieved by considering worst-case apriori estimates for $\mathbb{W}$ but it results in deteriorated controller “performance”. Thus, it is desirable to not only keep probability of $[\text{SCF}]_t^j$ low, but also maintain satisfactory controller performance during successful iterations (as defined in Remark 2). Let the closed loop cost of a successful iteration j be denoted by

$$\hat{\mathcal{V}}^j(x_S, w^{1:j}) = \sum_{t=0}^{T-1} \ell(x_t^j, v_{t|t}^{j,*}). \quad (8)$$

where notation $w^{1:j}$ denotes the set $\bigcup_{i=1}^j \bigcup_{t=0}^{T-1} w_t^i$. We use average closed loop cost $\mathbb{E}[\hat{\mathcal{V}}^j(x_S, w^{1:j})]$ to quantify controller performance. The goal is to lower the *performance loss* defined as

$$[\text{PL}]^j = \mathbb{E}[\hat{\mathcal{V}}^j(x_S, w^{1:j})] - \mathbb{E}[\mathcal{V}^*(x_S, w^{1:j})], \quad (9)$$

where $\mathbb{E}[\mathcal{V}^*(x_S, w^{1:j})]$ denotes the average closed loop cost of the j^{th} iteration if $\mathbb{W}$ had been known, i.e., $\hat{\mathbb{W}}^j = \mathbb{W}$ for all $j \in [J]$.

In the next section, we introduce two design specifications (D1) and (D2) to formalize this joint focus on lowering probability of State Constraint Failure and maintaining satisfactory controller performance. We then show how the sets $\hat{\mathbb{W}}^j$ are constructed according to these specifications.

4 Learning Robustness with Bounded Failure

We consider the following design specifications:

- (D1) Closed loop MPC control law (6) ensures that system (1) in the j^{th} iteration satisfies a user specified upper bound α on probability of State Constraint Failure (Definition 1),
- (D2) Minimize $[\text{PL}]^j$ (as defined in (9)) over all iterations $j \in [J]$ while satisfying (D1).

For satisfaction of (D1) we require,

$$\mathbb{P}(H_x x_t^j > h_x) \leq \alpha. \quad (10)$$

Since the above probability is difficult to compute, we consider an alternative notion of failure in order to upper bound the probability of State Constraint Failure.

Definition 2 (Disturbance Support Failure). *A Disturbance Support Failure at time step t in iteration j is the event*

$$[\text{DSF}]_t^j : w_t^j \notin \hat{\mathbb{W}}^j. \quad (11)$$

As the MPC controller (4) is robust to all $w_t^j \in \hat{\mathbb{W}}^j$, we have $[\text{SCF}]_t^j \subseteq [\text{DSF}]_t^j$. Therefore, probability of Disturbance Support Failure is an upper bound for probability of State Constraint Failure, i.e., $\mathbb{P}([\text{SCF}]_t^j) \leq \mathbb{P}([\text{DSF}]_t^j)$. Therefore, we focus on the following specification:

$$\mathbb{P}(w_t^j \notin \hat{\mathbb{W}}^j) \leq \alpha. \quad (12)$$

In the next few sections, we discuss how such sets $\hat{\mathbb{W}}^j$ can be constructed based on disturbance samples observed during the iterative task.

4.1 Need for Distributional Assumption on $\mathcal{P}$

Consider i.i.d. samples $Z_{1:n} = (Z_1, \dots, Z_n)$ from an unknown distribution $\mathcal{P}$. All we know about the distribution is that its support $\mathbb{S}$ is convex and compact. Our objective is to find an estimate $\hat{\mathbb{S}}(Z_{1:n})$ for the support $\mathbb{S}$ such that for a user specified failure probability α ,

$$\mathbb{P}(\bar{Z} \notin \hat{\mathbb{S}}(Z_{1:n})) \leq \alpha, \quad (13)$$

where $\bar{Z}$ is an i.i.d. draw from $\mathcal{P}$. The convex hull $C^{\text{hull}}(Z_{1:n})$ of observed samples $(Z_1, \dots, Z_n)$ is an intuitive estimator for the support $\mathbb{S}$. It is clear that $C^{\text{hull}}(Z_{1:n}) \subseteq \mathbb{S}$. Let $\mathcal{A} \setminus \mathcal{B}$ denote the set $\{y \mid y \in \mathcal{A} \text{ and } y \notin \mathcal{B}\}$. It turns out that $\mathbb{P}(\mathbb{S} \setminus C^{\text{hull}}(Z_{1:n})) \rightarrow 0$ as $n \rightarrow \infty$ [22], i.e., $C^{\text{hull}}(Z_{1:n})$ asymptotically converges to the support $\mathbb{S}$. However, $C^{\text{hull}}(Z_{1:n})$ may not satisfy (13) for an arbitrary user specified failure probability α . In order to do so, $C^{\text{hull}}(Z_{1:n})$ may need to be scaled up in a suitable manner. We illustrate through a simple example that an upper bound on failure probability cannot be guaranteed without additional assumptions on the distribution $\mathcal{P}$.

Consider an unknown univariate distribution $\mathcal{P}$ with support $\mathbb{S} \subset \mathbb{R}$. Suppose we observe i.i.d. samples $Z_{1:4} = \{-1, 0.5, 1, -0.2\}$ from this distribution. The objective is to find $\hat{\mathbb{S}}(Z_{1:4})$ that satisfies (13) with $\alpha = 0.1$. As we know that $\mathbb{S}$ is convex and compact, it is clear that $C^{\text{hull}}(Z_{1:4}) = [-1, 1] \subseteq \mathbb{S}$. However, it is unclear whether $\hat{\mathbb{S}} = C^{\text{hull}}(Z_{1:4})$ would satisfy (13) with $\alpha = 0.1$. Consider two potential distributions $\mathcal{P}_1, \mathcal{P}_2$ with densities $p_1(\cdot), p_2(\cdot)$ respectively such that

$$\begin{aligned} p_1(z) &= 0.4\mathbb{I}\{|z| \leq 1\} + 0.1\mathbb{I}\{1 \leq |z| \leq 2\}, \\ p_2(z) &= 0.4\mathbb{I}\{|z| \leq 1\} + 0.01\mathbb{I}\{1 \leq |z| \leq 11\}, \end{aligned}$$

where $\mathbb{I}\{\cdot\}$ denotes the indicator function. Note that both these distributions are equally likely to generate the observed samples as they have the same distribution on $C^{\text{hull}}(Z_{1:4}) = [-1, 1]$. Observe that $\hat{\mathbb{S}} = 1.5C^{\text{hull}}(Z_{1:4})$ satisfies (13) for $\mathcal{P}_1$, whereas $\hat{\mathbb{S}}$ has to be set to $6C^{\text{hull}}(Z_{1:4})$ to get the same probability of failure for $\mathcal{P}_2$. Thus, without any additional assumption about the distribution, it is not possible to give any probability of failure guarantees just based on sets constructed from observed samples.

Assumption 2. *We assume that the unknown distribution $\mathcal{P}$ defined in Section 2 belongs to a finite dimensional parametric family $\{\mathcal{P}_\theta : \theta \in \Theta, \Theta \subseteq \mathbb{R}^l\}$.*

We next explore how to construct the sets $\hat{\mathbb{W}}^j$ using Assumption 2, so that design specification (D1) is satisfied. For that purpose, we introduce the notion of Confidence Supports which are closely related to the notion of confidence intervals in classical statistics. Subsequently in Section 4.3 we present our algorithm.

4.2 Confidence Support of a Distribution

Consider i.i.d. samples $Z_{1:n} = (Z_1, \dots, Z_n)$ from a distribution $\mathcal{P}_\theta$ parametrized by $\theta \in \mathbb{R}$, i.e., $Z_i \stackrel{\text{iid}}{\sim} \mathcal{P}_\theta$. In classical statistics, the notion of confidence interval provides a convenient way to characterize the uncertainty of parameter θ from the observed samples $Z_{1:n}$.

Definition 3 (Confidence Interval). *A set $\mathcal{C}(Z_{1:n})$ is a $(1 - \alpha)$ -confidence interval for the parameter θ of distribution $\mathcal{P}_\theta$ if*

$$\mathbb{P}(\theta \notin \mathcal{C}(Z_{1:n})) \leq \alpha. \quad (14)$$

If $\theta \in \mathbb{R}^d$, $d > 1$, then the term *confidence region* is used for the set $\mathcal{C}(Z)$ as defined above.

Remark 3. *Note that $\mathcal{C}(Z)$ is a random set as it is a function of the collection of random samples $Z_{1:n}$, whereas θ is an unknown deterministic parameter. We refer the reader to [23, Chapter 9] for an introduction to confidence intervals and methods to compute them.*

We now introduce an analogous definition for the support of a distribution.

Definition 4 (Confidence Support). *A set $\mathcal{S}(Z_{1:n})$ is a $(1 - \alpha)$ -Confidence Support of a distribution $\mathcal{P}_\theta$ with support $\mathbb{S}_\theta$ if*

$$\mathbb{P}(\mathbb{S}_\theta \subseteq \mathcal{S}(Z_{1:n})) \geq 1 - \alpha, \quad (15)$$

i.e., $\mathcal{S}(Z_{1:n})$ contains the support $\mathbb{S}_\theta$ of $\mathcal{P}_\theta$ with probability greater than or equal to $(1 - \alpha)$.

Using the above notion of Confidence Supports, we now demonstrate how the disturbance support estimates $\hat{\mathbb{W}}^j$ (as defined in iterative MPC problem (4)) can be computed based on observed disturbance samples.

4.3 Computing $\hat{\mathbb{W}}^j$

Consider i.i.d. disturbance samples $w_t^j \sim \mathcal{P}_\theta$, $\theta \in \mathbb{R}^l$ with support $\mathbb{W}$. Let $w_t^j(q)$ denote the q^{th} element of $w_t^j \in \mathbb{R}^d$. Let $w^{1:j}$ denote the set $\bigcup_{i=1}^j \bigcup_{t=0}^{T^i} w_t^i$. Recall that $[d]$ denotes the set $\{1, 2, \dots, d\}$. We make the following simplifying assumption:

Assumption 3. *The elements of random vector $w_t^i \in \mathbb{R}^d$ are independently distributed,*

$$w_t^j(q) \sim \mathcal{P}_{\theta_q}^q, \quad q \in [d], \quad (16)$$

where $\theta = (\theta_1, \dots, \theta_d)$ and $\{\mathcal{P}_{\theta_q}^q : \theta_q \in \Theta_q, \Theta_q \subset \mathbb{R}^{l/d}\}$ is the corresponding parametric family for the q^{th} element. Remark 4 contains a discussion about the general case.

At the start of the j^{th} iteration, the collection of samples $w^{1:j-1}$ would have been observed. As the uncertainty distribution $\mathcal{P}_\theta$ is completely specified by θ , we can compute a $(1 - \alpha)$ -Confidence Support $\hat{\mathbb{W}}^j(w^{1:j-1})$ by computing confidence regions for the individual parameters $(\theta_1, \dots, \theta_d)$. Note that the confidence regions and supports are functions of the observed disturbance samples $w^{1:j-1}$. For notational convenience, we represent such sets without explicitly showing this dependence.

Lemma 1. *Let $\hat{\Theta}_q^j$ be a $(1 - \alpha_q)$ -confidence region for θ_q . Consider $\hat{\mathbb{W}}_q^j = \bigcup_{\bar{\theta}_q \in \hat{\Theta}_q^j} \text{Supp}(\mathcal{P}_{\bar{\theta}_q}^q)$, where $\text{Supp}(\mathcal{P}_{\bar{\theta}_q}^q)$ denotes the support of distribution $\mathcal{P}_{\bar{\theta}_q}^q$. Then, $\hat{\mathbb{W}}^j = \hat{\mathbb{W}}_1^j \times \dots \times \hat{\mathbb{W}}_d^j$ is a $(1 - \sum_q \alpha_q)$ -Confidence Support of $\mathcal{P}_\theta$.*

Proof. By definition, $\mathbb{W} = \text{Supp}(\mathcal{P}_{\theta_1}^1) \times \dots \times \text{Supp}(\mathcal{P}_{\theta_d}^d)$. As $\hat{\mathbb{W}}^j = \hat{\mathbb{W}}_1^j \times \dots \times \hat{\mathbb{W}}_d^j$, we have

$$\begin{aligned} \mathbb{P}(\mathbb{W} \not\subseteq \hat{\mathbb{W}}^j) &= \mathbb{P}\left(\bigcup_{q=1}^d \text{Supp}(\mathcal{P}_{\theta_q}^q) \not\subseteq \hat{\mathbb{W}}_q^j\right) \\ &= \mathbb{P}\left(\bigcup_{q=1}^d \theta_q \notin \hat{\Theta}_q^j\right), \\ &\leq \sum_{q=1}^d \mathbb{P}(\theta_q \notin \hat{\Theta}_q^j), \end{aligned} \quad (17)$$

$$\leq \sum_{q=1}^d \alpha_q, \quad (18)$$

where (17) follows from the union bound and (18) follows from $\hat{\Theta}_q^j$ being a $(1 - \alpha_q)$ -confidence region for θ_q . $\square$

Thus, a $(1 - \alpha)$ -Confidence Support can be constructed using $(1 - \alpha_q)$ -confidence regions by setting $\alpha_q = \frac{\alpha}{d}$. We now show that such a Confidence Support has a bounded probability of Disturbance Support Failure, as defined in (11).

Proposition 1. *Let $\hat{\mathbb{W}}^j$ be a $(1 - \alpha)$ -Confidence Support of $\mathcal{P}_\theta$ computed using samples $w^{1:j-1}$. Then, we have*

$$\mathbb{P}(w_t^j \notin \hat{\mathbb{W}}^j) \leq \alpha, \quad 0 \leq t \leq T - 1. \quad (19)$$

Proof. Note that both w_t^j and $\hat{\mathbb{W}}^j$ are random. Using the law of total probability, we have

$$\begin{aligned} \mathbb{P}(w_t^j \notin \hat{\mathbb{W}}^j) &= \mathbb{P}(w_t^j \notin \hat{\mathbb{W}}^j | \mathbb{W} \subseteq \hat{\mathbb{W}}^j) \mathbb{P}(\mathbb{W} \subseteq \hat{\mathbb{W}}^j) \\ &\quad + \mathbb{P}(w_t^j \notin \hat{\mathbb{W}}^j | \mathbb{W} \not\subseteq \hat{\mathbb{W}}^j) \mathbb{P}(\mathbb{W} \not\subseteq \hat{\mathbb{W}}^j), \\ &= \mathbb{P}(w_t^j \notin \hat{\mathbb{W}}^j | \mathbb{W} \not\subseteq \hat{\mathbb{W}}^j) \mathbb{P}(\mathbb{W} \not\subseteq \hat{\mathbb{W}}^j), \\ &\leq \mathbb{P}(\mathbb{W} \not\subseteq \hat{\mathbb{W}}^j), \end{aligned} \quad (20)$$

$$\leq \alpha, \quad (21)$$

where (21) follows from the fact that $\hat{\mathbb{W}}^j$ is a $(1 - \alpha)$ -Confidence Support of $\mathcal{P}_\theta$. $\square$

Remark 4. *The Confidence Supports constructed in this section also hold in the case that the elements of w_t^j are dependent. However, as we are not exploiting the correlations across dimensions, the above approach would yield a hyper-rectangle outer-approximation to the actual support as iteration j goes to infinity. Confidence regions for the parameter θ rather than individual elements θ_q are needed in such a case to converge to the true support, but such regions are in general difficult to compute.*

Remark 5. *As long as the confidence regions $\hat{\Theta}_q^j$ converge to the true parameter θ_q in probability, the Confidence Supports asymptotically converge to the true uncertainty support, i.e., $\hat{\mathbb{W}}^j \rightarrow \mathbb{W}$ in probability. The MPC controller (6) thus asymptotically learns to satisfy (2) robustly.*

4.4 The LRBF Algorithm

We present our Learning Robustness from Bounded Failure (LRBF) algorithm which uses Confidence Supports $\hat{\mathbb{W}}^j$ from Section 4.3 in MPC optimization problem (4). This guarantees satisfaction of (10) (i.e., design requirement (D1)) by system (1) in closed loop with controller (6).

Remark 6. *We assume that for all iterations $j \in [J]$, at time step $t = 0$, MPC problem (4) is feasible with disturbance supports $\hat{\mathbb{W}}^j$ constructed in Algorithm 1. This guarantees that we are able to collect at least one data point in each iteration to update Confidence Support $\hat{\mathbb{W}}^j$ while satisfying (10). In case such an assumption is not satisfied, $\hat{\mathbb{W}}^j$ can be scaled down (for eg., by increasing α).*

Remark 7. *The convergence of $\hat{\mathbb{W}}^j$ to the true support $\mathbb{W}$ can be sped up by keeping the iteration running until time step T despite State Constraint Failure. This can be done by introducing slack variables in MPC problem (4). Details can be found in the Appendix.*

4.5 Case Studies

We now demonstrate our approach for two parametric distribution families: (i) uniform distribution, and (ii) truncated normal distribution.

Algorithm 1: Learning Robustness with Bounded Failure (LRBF)

Inputs: $\mathbb{Z}, \hat{\mathbb{W}}^1, x_S$.
for $j = 2, \dots, J$ **do**
Computing Confidence Support $\hat{\mathbb{W}}^j$
for $q = 1, \dots, d$ **do**
 Compute $(1 - \frac{\alpha}{d})$ -confidence region $\hat{\Theta}_q^j$ for θ_q
 Compute $\hat{\mathbb{W}}_q^j = \cup_{\bar{\theta}_q \in \hat{\Theta}_q^j} \text{Supp}(\mathcal{P}_{\bar{\theta}_q}^q)$
end for
 Set $\hat{\mathbb{W}}^j = \hat{\mathbb{W}}_1^j \times \dots \times \hat{\mathbb{W}}_d^j$
Solving MPC problem (4) using $\hat{\mathbb{W}}^j$
for $t = 0, 1, \dots, T - 1$ **do**
 Apply $v_{t|t}^{j,*}$ from (6) with $\hat{\mathbb{W}}^j$ as uncertainty
end for
end for

4.5.1 Uniform Distribution.

Consider the uniform distribution with hyper-rectangle support $\mathbb{W} = [-\theta_1, \theta_1] \times \dots \times [-\theta_d, \theta_d]$. Then we have,

$$\mathcal{P}_{\theta_q}^q = \text{Unif}(-\theta_q, \theta_q), \quad q \in [d].$$

Let $\bar{w}^j(q) = \max_{\bar{w} \in w^{1:j-1}} |\bar{w}|$, $q \in [d]$ and let $\mathcal{T}^j = \sum_{i=1}^{j-1} T^i$. The following set turns out to be a $(1 - \frac{\alpha}{d})$ -confidence interval for θ_q ,

$$\hat{\Theta}_q^j = \left[\bar{w}^j(q), \frac{\bar{w}^j(q)}{(\frac{\alpha}{d})^{1/\mathcal{T}^j}} \right].$$

A derivation of the above confidence interval can be found in the Appendix. Using Lemma 1, we have the $(1 - \alpha)$ -Confidence Support $\hat{\mathbb{W}}^j = \hat{\mathbb{W}}_1^j \times \dots \times \hat{\mathbb{W}}_d^j$, where

$$\hat{\mathbb{W}}_q^j = \left[-\frac{\bar{w}^j(q)}{(\frac{\alpha}{d})^{1/\mathcal{T}^j}}, \frac{\bar{w}^j(q)}{(\frac{\alpha}{d})^{1/\mathcal{T}^j}} \right]. \quad (22)$$

Remark 8. This can be extended to the asymmetric case with $\mathcal{P}_{\theta_q}^q = \text{Unif}(-\theta_q^1, \theta_q^2)$. In this case, there is no analytical expression for the Confidence Support but it can be computed numerically.

4.5.2 Truncated Normal Distribution.

Consider the truncated normal distribution with mean μ_q , variance σ_q^2 , and support $[\mu_q - 3\sigma_q, \mu_q + 3\sigma_q]$, i.e.,

$$\mathcal{P}_{\theta_q}^q = \mathcal{N}_{\text{trunc}}(\mu_q, \sigma_q^2, 3), \quad q \in [d].$$

As the distribution is fully specified by μ_q and σ_q , we have $\theta_q = [\mu_q, \sigma_q]^\top$. Although it is difficult to derive exact confidence intervals in this case, approximate confidence intervals for μ_q and σ_q can be computed via the Bootstrap [24, Chapter 13]. Let $[\mu_{\min}^j(q), \mu_{\max}^j(q)]$ and $[\sigma_{\min}^j(q), \sigma_{\max}^j(q)]$ denote the $(1 - \frac{\alpha}{2d})$ -Bootstrap confidence intervals for μ_q and σ_q respectively. By union bound, we have the following approximate $(1 - \frac{\alpha}{d})$ -confidence interval for θ_q ,

$$\begin{aligned}\hat{\Theta}_q^j &= \{[\mu, \sigma]^\top \mid \mu \in [\mu_{\min}^j(q), \mu_{\max}^j(q)], \\ &\quad \sigma \in [\sigma_{\min}^j(q), \sigma_{\max}^j(q)]\},\end{aligned}$$

which gives us an approximate $(1 - \alpha)$ -Confidence Support $\hat{\mathbb{W}}^j = \hat{\mathbb{W}}_1^j \times \dots \times \hat{\mathbb{W}}_d^j$, where

$$\hat{\mathbb{W}}_q^j = [\mu_{\min}^j(q) - 3\sigma_{\max}^j(q), \mu_{\max}^j(q) + 3\sigma_{\max}^j(q)]. \quad (23)$$

5 Numerical Simulations

In this section we find approximate solutions to the following iterative optimal control problem in receding horizon:

$$\begin{aligned}V^{j,\star}(x_S) &= \\ \min_{u_0^j, u_1^j(\cdot), \dots} &\sum_{t=0}^{T-1} 10 \|\bar{x}_t^j - x_{\text{ref}}\|_2^2 + 2 \|u_t^j(\bar{x}_t^j)\|_2^2 \\ \text{s.t.} & \\ x_{t+1}^j &= Ax_t^j + Bu_t^j(x_t^j) + w_t^j, \\ \bar{x}_{t+1}^j &= A\bar{x}_t^j + Bu_t^j(\bar{x}_t^j), \\ \begin{bmatrix} -30 \\ -30 \\ -40 \end{bmatrix} &\leq \begin{bmatrix} x_t^j \\ u_t^j(\bar{x}_t^j) \end{bmatrix} \leq \begin{bmatrix} 30 \\ 30 \\ 40 \end{bmatrix}, \forall w_t^j \in \mathbb{W}, \\ x_0^j &= x_S, \quad t = 0, 1, \dots, T-1.\end{aligned}$$

We consider two parametric distributions:

$$\mathcal{P}_{\theta_q}^q = \text{Unif}(-3, 3), \quad (24a)$$

$$\mathcal{P}_{\theta_q}^q = \mathcal{N}_{\text{trunc}}(0, 1, 3), \quad (24b)$$

with $q \in \{1, 2\}$. In both cases, $\mathbb{W} = [-3, 3] \times [-3, 3]$. We construct Bootstrap confidence intervals for the truncated normal case by resampling 1000 times. System matrices $A = \begin{bmatrix} 1.2 & 1.3 \\ 0 & 1.5 \end{bmatrix}$ and $B = [0, 1]^\top$ are known. We solve the above optimization problem with the initial state $x_S = [0, 0]^\top$ and reference point $x_{\text{ref}} = [27, 27]^\top$ for task duration $T = 20$ steps over $J = 30$ iterations. Algorithm 1 is implemented with a control horizon of $N = 4$, and the feedback gain K in (6) is chosen to be the optimal LQR gain for system $x^+ = (A + BK)x$ with parameters $Q_{\text{LQR}} = 10I_{2 \times 2}$ and $R_{\text{LQR}} = 2$. The goal is to show:

- Design specification (D1) is satisfied. Consequently, a lower probability of Disturbance Support Failure across all iterations using support $\hat{\mathbb{W}}^j$ from Algorithm 1, compared to that from the convex hull support estimate $C^{\text{hull}}(w^{1:j-1})$.
- The performance loss $[\text{PL}]^j$ rapidly approaches 0 within the first few iterations. However, in the initial iterations, there is a significant trade-off between a desired upper bound α on probability of State Constraint Failure and average closed loop cost $\mathbb{E}[\hat{\mathcal{V}}^j(x_S, w^{1:j})]$ (defined in (8)). That is, lower the upper bound α , higher is the average closed loop cost in the initial iterations. This suggests the need for tailoring the confidence level $(1 - \alpha)$ in Algorithm 1 according to the application at hand.

5.1 Bounding the Probability of Failure (D1)

In this section, we demonstrate satisfaction of design specification (D1) by Algorithm 1 and compare the probability of Disturbance Support Failure $\mathbb{P}(w_t^j \notin \hat{\mathbb{W}}^j)$ for any timestep t in the j^{th} iteration, with $\hat{\mathbb{W}}^j$ obtained using Algorithm 1 and $\hat{\mathbb{W}}^j = C^{\text{hull}}(w^{1:j-1})$. This probability is estimated by averaging over 100 Monte Carlo draws of disturbance samples $w^{1:J}$, i.e.,

$$\mathbb{P}(w_t^j \notin \mathbb{W}^j) \approx \frac{1}{100} \sum_{\tilde{m}=1}^{100} (\mathbf{1}_{\mathcal{F}}(w_t^j))^{\star \tilde{m}},$$

where

$$(\mathbf{1}_{\mathcal{F}}(w_t^j))^{\star \tilde{m}} = \begin{cases} 1, & \text{if } w_t^j \notin (\hat{\mathbb{W}}^j)^{\star \tilde{m}} | (w^{1:j-1})^{\star \tilde{m}}, \\ 0, & \text{otherwise,} \end{cases}$$

and $(\cdot)^{\star \tilde{m}}$ represents the $\tilde{m}^{\text{th}}$ Monte Carlo sample. Fig. 1 shows this comparison for uniformly distributed disturbance (24a). Using LRBF to construct Confidence Supports $\hat{\mathbb{W}}^j$ allows for lowering $\mathbb{P}(w_t^j \notin \hat{\mathbb{W}}^j)$, i.e., probability of $[\text{DSF}]_t^j$ as defined in (11) below a user specified bound α , as opposed to simply utilizing $\hat{\mathbb{W}}^j = C^{\text{hull}}(w^{1:j-1})$. We plot the probability of $[\text{DSF}]_t^j$ for 2 different values of $\alpha = 0.05$ and $\alpha = 0.70$. We see that for $\alpha = 0.05$ the probability of $[\text{DSF}]_t^j$ with LRBF is on average 94% smaller than that from the convex hull support estimate for all iterations $j \in [30]$. Similarly for $\alpha = 0.70$, the probability of $[\text{DSF}]_t^j$ is on average 61% lower than that with the convex hull support estimate across all $j \in [30]$.

The same trend is seen in Fig. 2 for truncated normal distribution (24b), where probability of $[\text{DSF}]_t^j$ is at least 99% and 96% lower than convex hull support estimate for $\alpha = 0.05$ and $\alpha = 0.70$ respectively until iteration $j = 3$, and reaches a value of 0 for both values of α afterwards. The above trend in probability of $[\text{DSF}]_t^j$ is explained by Proposition 1, which relates the desired confidence $(1 - \alpha)$ for support $\hat{\mathbb{W}}^j$ to the probability of $[\text{DSF}]_t^j$. Moreover, from Fig. 1 and Fig. 2 we see that in practice probability of $[\text{DSF}]_t^j$ is always at least 60% – 80% lower than corresponding chosen α . This highlights satisfaction of (D1) and also the conservatism in Proposition 1 arising from the upper bound in (20).




Figure 1: Probability of Disturbance Support Failure vs Iteration Number for Uniformly Distributed Disturbance on $\mathbb{W}$.

5.2 Performance Loss Reduction Over Iterations

In Fig. 3 and Fig. 4, we approximate the average closed loop cost $\mathbb{E}[\hat{\mathcal{V}}^j(x_S, w^{1:j})]$ of the j^{th} iteration by taking an empirical average over 100 Monte Carlo draws of $w^{1:J}$ as,

$$\mathbb{E}[\hat{\mathcal{V}}^j(x_S, w^{1:j})] \approx \frac{1}{100} \sum_{\tilde{m}=1}^{100} \hat{\mathcal{V}}^j(x_S, (w^{1:j})^{\star \tilde{m}}), \quad (25)$$

for $\alpha = 0.05$, and $\alpha = 0.70$. The cost values are normalized by $\mathcal{V}^*(x_S)$, which denotes the empirical average closed loop cost of the j^{th} iteration if $\mathbb{W}$ had been known, i.e., $\hat{\mathbb{W}}^j = \mathbb{W}$. For both cases of α , we see that in Fig. 3 and Fig. 4 the average closed loop cost rapidly approaches $\mathcal{V}^*(x_S)$. For (24a) in Fig. 3, cost (25) approaches to within 0.5% of $\mathcal{V}^*(x_S)$ after just 5 iterations whereas for (24b) in Fig. 4, it is within 3% of $\mathcal{V}^*(x_S)$ in the same duration.

However, the average closed loop cost incurred in earlier iterations has a trade-off with desired α . This trade-off is also highlighted in Fig. 3 and Fig. 4 for (24a) and (24b) respectively. We see from Fig. 3 and Fig. 4 that for lower value of probability of $[\text{SCF}]_t^j$ with $\alpha = 0.05$, we pay a maximum of 13% higher average closed loop cost for (24a), and a maximum of 10% higher average closed loop cost for (24b) compared to $\mathcal{V}^*(x_S)$ until iteration $j = 5$. Allowing for higher probability of $[\text{SCF}]_t^j$ with $\alpha = 0.70$ proves to be cost-efficient, where we only pay a maximum of 0.3% higher average closed loop cost for (24a), and a maximum of 4% higher average closed loop cost for (24b) compared to $\mathcal{V}^*(x_S)$ in the same duration. This essentially reflects the key trade-off between specifications (D1) and




Figure 2: Probability of Disturbance Support Failure vs Iteration Number for Truncated Normal Distribution of Disturbance on $\mathbb{W}$.

(D2) in the initial iterations. Thus, the upper bound α of $[\text{SCF}]_t^j$ must be chosen in an application-specific manner.

Acknowledgement

We thank Hamidreza Tavafoghi for helpful discussions and reviews. This work was partially funded by Office of Naval Research grant ONR-N00014-18-1-2833, National Science Foundation under grants EAGER-1549945 and CPS-1646612, and by the National Research Foundation of Singapore under a grant to the Berkeley Alliance for Research in Singapore.

References

- [1] B. Recht, “A tour of reinforcement learning: The view from continuous control,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, pp. 253–279, 2019.
- [2] M. Tanaskovic, L. Fagiano, C. Novara, and M. Morari, “Data-driven control of nonlinear systems: An on-line direct approach,” *Automatica*, vol. 75, pp. 1–10, 2017.
- [3] U. Rosolia, X. Zhang, and F. Borrelli, “Data-driven predictive control for autonomous systems,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 259–286, 2018.




Figure 3: Normalized Average Closed Loop Cost (25): Uniform Disturbance.

- [4] M. Morari and J. H. Lee, “Model predictive control: past, present and future,” *Computers & Chemical Engineering*, vol. 23, no. 4, pp. 667–682, 1999.
- [5] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, “Constrained model predictive control: Stability and optimality,” *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.
- [6] F. Borrelli, A. Bemporad, and M. Morari, *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.
- [7] M. Tanaskovic, L. Fagiano, R. Smith, and M. Morari, “Adaptive receding horizon control for constrained MIMO systems,” *Automatica*, vol. 50, no. 12, pp. 3019–3029, 2014.
- [8] M. Lorenzen, M. Cannon, and F. Allgöwer, “Robust MPC with recursive model update,” *Automatica*, vol. 103, pp. 461 – 471, 2019.
- [9] J. Köhler, E. Andina, R. Soloperto, M. A. Müller, and F. Allgöwer, “Linear robust adaptive model predictive control : Computational complexity and conservatism,” *arXiv preprint arXiv:1909.01813*, 2019.
- [10] M. Bujarbaruah, S. Nair, and F. Borrelli, “A semi-definite programming approach to robust adaptive MPC under state dependent uncertainty,” *arXiv preprint arXiv:1910.04378*, 2019.




Figure 4: Normalized Average Closed Loop Cost (25): Truncated Normal Disturbance.

- [11] J. Köhler, P. Kötting, R. Soloperto, F. Allgöwer, and M. A. Müller, “A robust adaptive model predictive control framework for nonlinear uncertain systems,” *arXiv preprint arXiv:1911.02899*, 2019.
- [12] X. Zhang, K. Margellos, P. Goulart, and J. Lygeros, “Stochastic model predictive control using a combination of randomized and robust optimization,” in *IEEE Conference on Decision and Control (CDC)*, Florence, Italy, 2013.
- [13] L. Hewing and M. N. Zeilinger, “Cautious model predictive control using gaussian process regression,” *arXiv preprint arXiv:1705.10702*, 2017.
- [14] R. Soloperto, M. A. Müller, S. Trimpe, and F. Allgöwer, “Learning-based robust model predictive control with state-dependent uncertainty,” in *IFAC Conference on Nonlinear Model Predictive Control*, Madison, Wisconsin, USA, Aug. 2018.
- [15] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, “Learning-based model predictive control for safe exploration,” in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 6059–6066.
- [16] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, “Learning-based nonlinear model predictive control to improve vision-based mobile robot path-tracking in challenging outdoor environments,” in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 4029–4036.

- [17] R. Tempo, G. Calafiore, and F. Dabbene, *Randomized algorithms for analysis and control of uncertain systems: with applications*. Springer Science & Business Media, 2012.
- [18] U. Rosolia and F. Borrelli, “Learning model predictive control for iterative tasks. a data-driven control framework.” *IEEE Transactions on Automatic Control*, 2017.
- [19] X. Lu and M. Cannon, “Robust adaptive tube model predictive control,” in *2019 American Control Conference (ACC)*, Jul. 2019, pp. 3695–3701.
- [20] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski, “Optimization over state feedback policies for robust control with constraints,” *Automatica*, vol. 42, no. 4, pp. 523–533, 2006.
- [21] K. P. Wabersich and M. N. Zeilinger, “Linear model predictive safety certification for learning-based control,” in *2018 IEEE Conference on Decision and Control (CDC)*, Dec 2018, pp. 7130–7135.
- [22] V.-E. Brunel *et al.*, “Methods for estimation of convex sets,” *Statistical Science*, vol. 33, no. 4, pp. 615–632, 2018.
- [23] R. W. Keener, *Theoretical statistics: Topics for a core course*. Springer, 2011.
- [24] B. Efron and R. Tibshirani, “Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy,” *Statistical science*, pp. 54–75, 1986.

Appendix

Derivation of Confidence Support (22)

Consider $w_t^j(q) \stackrel{\text{iid}}{\sim} \text{Unif}(-\theta_q, \theta_q)$. This implies that $\frac{|w_t^j(q)|}{\theta_q} \stackrel{\text{iid}}{\sim} \text{Unif}(0, 1)$. Let $\bar{w}^j(q) = \max_{\bar{w} \in w^{1:j-1}} |\bar{w}|$. Then, for any $c \in [0, 1]$ we have,

$$\begin{aligned}
 \mathbb{P}\left(\frac{\bar{w}^j(q)}{\theta_q} \leq c\right) &= \mathbb{P}\left(\bigcap_{\bar{w} \in w^{1:j-1}} \frac{|\bar{w}|}{\theta_q} \leq c\right), \\
 &= \Pi_{\bar{w} \in w^{1:j-1}} \mathbb{P}\left(\frac{|\bar{w}|}{\theta_q} \leq c\right), \\
 &= c^{\mathcal{T}^j},
 \end{aligned} \tag{26}$$

where (26) follows as disturbance samples $\bar{w} \in w^{1:j-1}$ are independent. Setting $c = \alpha_q$, we have

$$\mathbb{P}\left(\frac{\bar{w}^j(q)}{\theta_q} \leq \alpha_q^{\frac{1}{\mathcal{T}^j}}\right) = \alpha_q.$$

Therefore, we have

$$\mathbb{P}\left(\alpha_q^{\frac{1}{T^j}} \leq \frac{\bar{w}^j(q)}{\theta_q} \leq 1\right) = 1 - \alpha_q,$$

which gives us

$$\mathbb{P}\left(\bar{w}^j(q) \leq \theta_q \leq \frac{\bar{w}^j(q)}{\alpha_q^{\frac{1}{T^j}}}\right) = 1 - \alpha_q.$$

Setting $\alpha_q = \frac{\alpha}{d}$ and using Lemma 1 completes the derivation.

Speeding up Convergence of $\hat{\mathbb{W}}^j$

In order to speed up convergence of $\hat{\mathbb{W}}^j$ in Algorithm 1 to the true support $\mathbb{W}$, the following MPC optimization problem with slack variables is solved:

$$\begin{aligned} & \tilde{V}_{t \rightarrow t+N}^{\text{MPC},j}(x_t^j, \hat{\mathbb{W}}^j, \hat{\mathcal{X}}_N^j) := \\ & \min_{U_t^j(\cdot)} \sum_{k=t}^{t+N-1} \ell(\bar{x}_{k|t}^j, v_{k|t}^j) + Q(\bar{x}_{t+N|t}^j) + \Lambda \|\mathbf{s}_t^j\|_2^2 \\ & \text{s.t.} \quad x_{k+1|t}^j = Ax_{k|t}^j + Bu_{k|t}^j + w_{k|t}^j, \\ & \quad \bar{x}_{k+1|t}^j = A\bar{x}_{k|t}^j + Bv_{k|t}^j, \\ & \quad u_{k|t}^j = \sum_{l=t}^{k-1} M_{k,l|t}^j w_{l|t}^j + v_{k|t}^j, \\ & \quad H_x x_{k|t}^j \leq h_x + s_t^j, \\ & \quad H_u u_{k|t}^j \leq h_u, \\ & \quad \hat{Y}^j x_{t+N|t}^j \leq \hat{z}^j + \hat{s}_t^j, \\ & \quad \text{with } \hat{\mathcal{X}}_N^j = \{x : \hat{Y}^j x \leq \hat{z}^j\}, \\ & \quad \mathbf{s}_t^j = [(s_t^j)^\top, (\hat{s}_t^j)^\top]^\top \geq 0, \\ & \quad \forall w_{k|t}^j \in \hat{\mathbb{W}}^j, \\ & \quad \forall k = \{t, \dots, t+N-1\}, \\ & \quad x_{t|t}^j = \bar{x}_{t|t}^j = x_t^j, \Lambda \gg 0, \end{aligned} \tag{27}$$

with $\mathbf{s}_0^j = 0$ (from Remark 6), and then closed loop control law $u_t^j = v_{t|t}^{j,\star}$ is applied to system (1). By solving the relaxed optimization problem (27) which is feasible for all timesteps $0 \leq t \leq T-1$ in the j^{th} iteration, we ensure that after each iteration, a set of T additional samples are obtained for the update $\hat{\mathbb{W}}^j \xrightarrow{\text{update}} \hat{\mathbb{W}}^{j+1}$. From Section 4.3 we can infer that this speeds up the convergence of $\hat{\mathbb{W}}^j$.