

Distribution-Free Guarantees for Systems with Decision-Dependent Noise

Heling Zhang¹, Lillian J. Ratliff² and Roy Dong³

Abstract—In many real-world dynamical systems, obtaining precise models of system uncertainty remains a challenge. It may be difficult to estimate noise distributions or robustness bounds, especially when the distributions/robustness bounds vary with different control inputs in unknown ways. Addressing this challenge, this paper presents a novel iterative method tailored for systems with decision-dependent noise without prior knowledge of the distributions. Our approach finds the open-loop control law that minimizes the worst-case loss, given that the noise induced by this control lies in its $(1 - p)$ -confidence set for a predetermined p . At each iteration, we use a quantile method inspired by conformal prediction to empirically estimate the confidence set shaped by the preceding control law. These derived confidence sets offer distribution-free guarantees on the system's noise, guiding a robust control formulation that targets worst-case loss minimization. Under specific regularity conditions, our method is shown to converge to a near-optimal open-loop control. While our focus is on open-loop controls, the adaptive, data-driven nature of our approach suggests its potential applicability across diverse scenarios and extensions.

I. INTRODUCTION

In the study of dynamical systems, the divergence between ideal models and real-world scenarios often manifests as a disturbance. This disturbance, compounded by the lack of knowledge about its statistical properties, presents significant challenges for control tasks. Traditional strategies, including stochastic and robust control, are theoretically effective but typically depend on prior knowledge or assumptions about the disturbance and its effect on the system, which are often unavailable in practice. Modern learning-based methods require less information about the distribution, but usually lack rigorous guarantees.

To address these challenges, we introduce an iterative method that gives distribution-free guarantees. That is, we assume no prior knowledge on the disturbance, even for the case where the disturbance is decision-dependent. The goal of our method is as follows: for a given noisy system with an unknown (potentially decision-dependent) disturbance, we aim to find the open-loop control law that minimizes the worst-case loss, given that the disturbance induced by this control lies in its $(1 - p)$ -confidence set for a predetermined p .

Our method works as follows: at each iteration, we use a quantile method inspired by conformal prediction to empirically estimate the confidence set shaped by the preceding control law. These derived confidence sets offer empirical constraints on the system disturbance, guiding a robust control formulation that targets worst-case loss minimization. By solving this robust control problem, we obtain a refined control law which enables us to re-estimate the confidence set of the disturbance. Our theoretical analysis shows that this method converges to a near-optimal open-loop control law under specific regularity conditions.

Devising guarantees is a common theme in the study of noisy systems. When the noise is known, classic reachability analysis such as [12] provide provable guarantees of safety. When the noise distribution is unknown, various empirical methods are often used. In [4], the authors derive safety guarantees by learning the worst-case disturbance online. In [1], the authors use a Gaussian process model to learn the system's unknown dynamics and approximates the maximal safe set. Our method adds to this list, providing a new way of deriving safety guarantees from empirical observations.

A key component of our method is inspired by *conformal prediction*, a method widely used for uncertainty quantification in machine learning ([2], [13]). In recent years, the application of conformal prediction in control tasks has also found great success. In [7], conformal prediction is used to construct predictors with probabilistic guarantees, which is then used to design model predictive controllers. In [9], the authors use conformal prediction to calibrate measures of uncertainty in trajectory forecasting models, which are then used to generate reachable sets that facilitates downstream navigation and planning tasks. In [6], the authors use conformal prediction to verify neural reachable tubes derived with DeepReach ([3]). In [8], the authors provide two predictive runtime verification algorithms that use conformal prediction to quantify prediction uncertainty and compute the probability that the system trajectory violates a signal temporal logic specification. In this work, we use a quantile method inspired by conformal prediction to compute confidence sets of the noise induced by the current control law, which is then used to formulate a robust control problem that leads to a more refined control law.

Methodologically, our work is inspired by the works on *performative prediction* [11], which is the study of decision-dependent distribution shift in learning. Conceptually, our problem shares several similarities with performative prediction: 1) the solutions to our problems alter the noise we encounter, and 2) we both seek the solution that gives the best performance in the noise it induces. As a result, our

¹Heling Zhang is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Champaign, IL 61801, USA hzhng120@illinois.edu +1 217-898-9702

²Lillian J. Ratliff is with the Department of Electrical and Computer Engineering, University of Washington, Seattle, WA 98195, USA ratliff1@uw.edu +1 206-543-6244

³Roy Dong is with the Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Champaign, IL 61801, USA roydong@illinois.edu +1 217-265-4062

method adopts a similar iterative structure, in which each iteration builds upon the noise induced by the result from the previous iteration, but eventually converges to a fixed-point under certain regularity conditions.

II. PROBLEM STATEMENT

Consider a system that follows the underlying dynamics given by

$$x(t+1) = f(t, x(t), u(t), \zeta(t)) \quad t = 0, 1, \dots, T-1, \quad (2.1)$$

where $x(t) \in X_t \subseteq \mathbb{R}^n$ is the system state, $u(t) \in U_t \subseteq \mathbb{R}^m$ is the control input, and $\zeta(t) \in \mathbb{R}^n$ is a noise whose distribution depends on $x(t)$ and $u(t)$, i.e. $\zeta(t) \sim \mathcal{D}(x(t), u(t))$. We assume that we do not have direct knowledge of the system dynamics, but instead, for any given sequence of control law $(u(0), \dots, u(T-1))$ and any initial state $x(0)$, we can take independent sample trajectories $(x_i(0), \dots, x_i(t))_{i=1}^M$. Out of these sample trajectories, we can then construct a unbiased estimator of $f(t, x(t), u(t), \zeta(t))$ using various methods. We write such unbiased estimator as $\hat{f}(t, x(t), u(t))$. Let $w(t) := f(t, x(t), u(t), \zeta(t)) - \hat{f}(t, x(t), u(t))$, we can see that $w(t)$ is a random variable whose distribution depends on $x(t)$ and $u(t)$. Then the system dynamics can be written as

$$x(t+1) = \hat{f}(t, x(t), u(t)) + w(t) \quad t = 0, 1, \dots, T-1. \quad (2.2)$$

We are mainly concerned with developing open-loop control for the system described by Equation 2.2, and assume $\hat{f}$ is a known function from this point on.

Without loss of generality, we further assume that the system starts at a fixed state $x(0) = x_0$. With a slight abuse of notation, we use x, u, w to denote associated trajectories

$$\begin{aligned} \mathbf{x} &= (x(0), \dots, x(T)) \in \mathbb{R}^{n(T+1)}, \\ \mathbf{u} &= (u(0), \dots, u(T-1)) \in \mathbb{R}^{mT}, \\ \mathbf{w} &= (w(0), \dots, w(T-1)) \in \mathbb{R}^{nT}. \end{aligned}$$

The objective is to determine an open-loop control law $\mathbf{u} = (u(0), \dots, u(T-1))$ that optimally governs the system's behavior. To quantify the performance of a control law, we introduce a loss function J dependent on the entire trajectory of states and control inputs. We define J as

$$J(\mathbf{x}, \mathbf{u}) := \phi(x(T)) + \sum_{t=0}^{T-1} \varphi(x(t), u(t)).$$

However, given the stochastic nature of the noise $w(t)$, the loss function J becomes a random object. This makes it impossible to directly minimize the value of the loss function. Traditional approaches either aim to minimize the expected loss, requiring knowledge of the noise distribution, or minimize the worst-case loss for bounded noise, requiring knowledge of the extreme values of the noise. Both approaches require some prior knowledge of the behavior of the noise. This work addresses scenarios where no such prior knowledge is available.

When sampling is feasible, we can construct confidence sets for the noise with probabilistic guarantees using techniques similar to conformal prediction. The primary focus of this

work is to minimize the worst-case loss given these confidence sets. Formally, for a designated probability level p , let $C_t(\mathbf{u})$ be the confidence set for $w(t)$ such that $w(t) \in C_t(\mathbf{u})$ with probability at least $1 - p/T$. Applying the union bound, this would give us $\mathbf{w} \in C_0(\mathbf{u}) \times \dots \times C_{T-1}(\mathbf{u})$ with probability at least $1 - p$. We are interested in the following optimization problem

$$\begin{aligned} \min_{\mathbf{u}} \quad & \max_{\mathbf{w}} J(\mathbf{x}(\mathbf{u}, \mathbf{w}), \mathbf{u}) \\ \text{s.t.} \quad & \mathbf{w} \in C_0(\mathbf{u}) \times \dots \times C_{T-1}(\mathbf{u}) \\ & \mathbf{u} \in U_0 \times \dots \times U_{T-1}. \end{aligned} \quad (2.3)$$

For clarity, we can also equivalently express the loss as $L(\mathbf{w}, \mathbf{u}) = J(\mathbf{x}(\mathbf{u}, \mathbf{w}), \mathbf{u})$, emphasizing its dependence on the noise and control. Further, let

$$\begin{aligned} C(\mathbf{u}) &:= C_0(\mathbf{u}) \times \dots \times C_{T-1}(\mathbf{u}) \\ U &:= U_0 \times \dots \times U_{T-1}, \end{aligned}$$

then we can write the optimization problem in a more concise form

$$\begin{aligned} \min_{\mathbf{u}} \quad & \max_{\mathbf{w} \in C(\mathbf{u})} L(\mathbf{w}, \mathbf{u}) \\ \text{s.t.} \quad & \mathbf{u} \in U. \end{aligned}$$

III. METHODS

In this section, we introduce our main method, empirical iterative refinement of performative control (E-IRPC). We call the control performative since it induces changes in the system noise distribution.

A. Empirical Iterative Refinement of Performative Control

The core of our methodology is an iterative process designed to progressively refine the control law to achieve optimal performance in the presence of noise. Our method consists of the following steps:

- 1) **Initialization:** We start by finding the optimal open-loop control for the nominal system (i.e. the system without noise) with loss function J :

$$\begin{aligned} \min_{\mathbf{u}} \quad & J(\mathbf{x}(\mathbf{u}), \mathbf{u}) \\ \text{s.t.} \quad & x(t+1) = \hat{f}(t, x(t), u(t)) \\ & \mathbf{u} \in U_0 \times \dots \times U_{T-1}. \end{aligned} \quad (3.4)$$

We denote the resulting control law as $\mathbf{u}_0 = (u_0(0), \dots, u_0(T-1))$.

- 2) **Sampling and Confidence Set Construction:** In this step, we fix the control law and compute the confidence sets for the induced noise. Specifically, with the control law from the previous step, we sample N_1 independent trajectories of system states, represented as $\mathbf{x}_1^{(j)}$ for $j = 1, \dots, N_1$. Since $\hat{f}$ is known, this gives us N_1 trajectories of the noise, $\mathbf{w}_1^{(j)}$ for $j = 1, \dots, N_1$. Notice that since we are using the same control input, at each time step t , $(w_1^{(j)}(t))_{j=1}^{N_1}$ are i.i.d. We can then use quantile regression to construct sets $\hat{C}_t(\mathbf{u})$ such that $\mathbb{P}(w(t) \in \hat{C}_t(\mathbf{u})) \geq 1 - \frac{p}{T}$. Again, for simplicity, we

denote $\hat{C}(\mathbf{u}) := \hat{C}_0(\mathbf{u}) \times \cdots \times \hat{C}_{T-1}(\mathbf{u})$. The details of such construction will be outlined in the next subsection.

- 3) **Robust Control Problem Formulation:** In this step, we fix the confidence sets and vary the control law. Specifically, with the confidence sets computed in the previous step, we can formulate a robust control problem similar to Equation 2.3:

$$\begin{aligned} \min_{\mathbf{u}} \quad & \max_{\mathbf{w}} J(\mathbf{x}(\mathbf{u}, \mathbf{w}), \mathbf{u}) \\ \text{s.t} \quad & \mathbf{w} \in \hat{C}_0(u_0) \times \cdots \times \hat{C}_{T-1}(u_0) \\ & \mathbf{u} \in U_0 \times \cdots \times U_{T-1}. \end{aligned} \quad (3.5)$$

We solve this problem to obtain a new control law $\mathbf{u}_1$.

- 4) **Iterative Refinement:** The new control law is then used as the initial control law in Step 2, and the process repeats.

In summary, for $i = 0, 1, \dots$, we fix the control law $\mathbf{u}_i$, sample N_i trajectories $(\mathbf{x}_{i+1}^{(j)})_{j=1}^{N_i}$ and $(\mathbf{w}_{i+1}^{(j)})_{j=1}^{N_i}$ to construct a confidence set $\hat{C}(\mathbf{u}_i)$, then we fix $\hat{C}(\mathbf{u}_i)$ and formulate a robust control problem to obtain $\mathbf{u}_{i+1}$. To simplify our notation, define $\hat{\mathcal{A}} : U \rightarrow U$ as

$$\hat{\mathcal{A}}(\mathbf{u}) \in \arg \min \left\{ \max \{L(\mathbf{w}, \mathbf{v}) : \mathbf{w} \in \hat{C}(\mathbf{u})\} : \mathbf{v} \in U \right\}, \quad (3.6)$$

then our method is summarized in Algorithm 1.

The iterative nature of this method is essential due to the performative aspect of the control: the control law influences the distribution of the noise, necessitating repeated refinement. Additionally, without prior knowledge of how different controls influence this noise, we need to re-collect samples to estimate bounds on the noise whenever we change the control.

Algorithm 1 E-IRPC

```

 $u_0 \leftarrow$  Solution to Problem 3.4
while  $i > 0$  do
  Sample  $(\mathbf{x}_{i+1}^{(j)})_{j=1}^{N_i}$  and  $(\mathbf{w}_{i+1}^{(j)})_{j=1}^{N_i}$ 
  Construct  $\hat{C}(\mathbf{u}_{i-1})$  from samples
   $\mathbf{u}_i \leftarrow \hat{\mathcal{A}}(\mathbf{u}_{i-1})$   $\triangleright \hat{\mathcal{A}}$  defined in Equation 3.6.
end while

```

B. Confidence Set Construction

In this section, we will outline the method we use to construct the confidence sets in Step 2. As described in the previous section, for each time step t , we have N_i samples of $w(t)$. Let's sort their Euclidean norm in the increasing order as follows: $|w^{(\eta_1)}(t)| \leq \cdots \leq |w^{(\eta_{N_i})}(t)|$. Since these samples are i.i.d, we can easily show that the Euclidean norm of another i.i.d sample $w(t)$ is equally likely to fall between that of any existing samples [2], this gives us

$$\mathbb{P}(|w(t)| \leq |w^{(\eta_k)}(t)|) = \frac{k}{N_i + 1} \quad (3.7)$$

for any $k = 1, \dots, N_i$. This allows us to construct the confidence sets with desired probabilistic guarantee.

Specifically, for the overall probability level p , pick $k = \lceil (N_i + 1)(1 - \frac{p}{T}) \rceil$ and define

$$\hat{C}_t(\mathbf{u}) := \left\{ w(t) : |w(t)| \leq |w^{(\eta_k)}(t)| \right\}, \quad (3.8)$$

by Equation 3.7, we have

$$\mathbb{P}(w(t) \in \hat{C}_t(\mathbf{u})) = \frac{\lceil (N_i + 1)(1 - \frac{p}{T}) \rceil}{N_i + 1} \geq 1 - \frac{p}{T},$$

which means that Equation 3.8 gives us the desired confidence set. It is worth noting that when constructing these confidence sets, we can replace the Euclidean norm with any other score function. This will allow us to construct confidence sets with the same probabilistic guarantee but different shapes. For example, if we can chose $s(w(t)) := \sqrt{w^\top(t)Hw(t)}$ for some positive definite H , and rank the sample noises according to the score $s(w^{(\eta_1)}(t)) \leq \cdots \leq s(w^{(\eta_{N_i})}(t))$, the resulting confidence set given by $\tilde{C}_t(\mathbf{u}) = \{w(t) : s(w(t)) \leq s(w^{(\eta_k)}(t))\}$ will still have the same probabilistic guarantee, but takes the shape of an ellipsoid instead of a ball.

C. Ideal Iterative Refinement of Performative Control

So far, we have detailed our primary method, which inherently relies on estimating specific properties of the noise distribution using finite samples. Naturally, the algorithm's performance is influenced by the number of samples taken. In this section, we introduce an idealized variant of our method, termed the *Ideal Iterative Refinement of Performative Control* (I-IRPC). This version sidesteps the variability introduced by finite sample trajectories. As we will discuss, the I-IRPC can be conceptualized as the infinite sample counterpart of our E-IRPC.

Formally, let $\mathbf{u}$ be the control input for the system defined by 2.2, and $w(t)$ be the noise at time t whose distribution is given by $\mathcal{D}(x(t), u(t))$. We define the *ideal confidence set* for $w(t)$ as

$$C_t(\mathbf{u}) = \{w(t) : |w(t)| \leq Q_t(\mathbf{u})\}, \quad (3.9)$$

where $Q_t(\mathbf{u})$ is the quantile function of $|w(t)|$ defined as

$$Q_t(\mathbf{u}) := \inf \left\{ r : \mathbb{P}(|w(t)| \leq r) \geq 1 - \frac{p}{T} \right\} \quad (3.10)$$

where p is the overall probability level. The set $C_t(\mathbf{u})$ is the idealized counterpart of $\hat{C}_t(\mathbf{u})$ defined in 3.8, which is impossible to obtain without knowing the exact distribution of $w(t)$. In the I-IRPC, we simply assume that there is an oracle that gives $C_t(\mathbf{u})$ for any $\mathbf{u}$ we choose.

Notice that Q_t depends solely on $\mathbf{u}$ because $\mathbf{u}$ uniquely defines the distribution of $w(t)$. Now, let $C(\mathbf{u}) := C_0(\mathbf{u}) \times \cdots \times C_{T-1}(\mathbf{u})$ and define $\mathcal{A} : U \rightarrow U$ as

$$\mathcal{A}(\mathbf{u}) \in \arg \min \left\{ \max \{L(\mathbf{w}, \mathbf{u}) : \mathbf{w} \in C(\mathbf{u})\} : \mathbf{u} \in U \right\}. \quad (3.11)$$

Then the I-IRPC is given by Algorithm 2. The I-IRPC gives us a starting point for analysis. Specifically, due to its deterministic nature, it allows us to make the following key definitions.

Algorithm 2 I-IRPC

```

 $u_0 \leftarrow$  Solution to Problem 3.4
while  $i > 0$  do
  Get  $C(\mathbf{u})$  from the oracle
   $\mathbf{u}_i \leftarrow \mathcal{A}(\mathbf{u}_{i-1})$   $\triangleright \mathcal{A}$  defined in Equation 3.11.
end while

```

Definition 1 (Performatively Stable Control). We say $\mathbf{u}_{PS} \in U$ is a *performatively stable control* if $\mathcal{A}(\mathbf{u}_{PS}) = \mathbf{u}_{PS}$.

Definition 2 (Performatively Optimal Control). We say $\mathbf{u}_{PO} \in U$ is a *performatively optimal control* if

$$\mathbf{u}_{PO} \in \arg \min_{\mathbf{u} \in U} \max_{\mathbf{w} \in C(\mathbf{u})} L(\mathbf{w}, \mathbf{u}),$$

where $C(\mathbf{u})$ is defined in Equation 3.9.

Conceptually, the performatively stable control is the control law that performs optimally in the noise it induces. The performatively optimal control, on the other hand, gives the best performance among all feasible control laws in the presence of self-induced noise. Notice that the performative optimal control, in contrast to performative stable control, may not be the optimal control law in the noise that it induces. As we will see shortly, these concepts play a central role in our theoretical analysis of our proposed method.

Now, let's revisit E-IRPC from the lens of I-IRPC. Define the empirical quantile function by

$$\begin{aligned} \hat{Q}_t(\mathbf{u}) &:= \inf \left\{ r : \frac{1}{N_i} \sum_{j=1}^{N_i} \mathbf{1} \left\{ |w^{(j)}(t)| \leq r \right\} \geq 1 - \frac{p}{T} \right\} \\ &= |w^{(\eta_k)}(t)|. \end{aligned}$$

As the name suggests, the empirical quantile function is a finite sample approximation of the actual quantile function defined in Equation 3.10. The confidence sets constructed in Section III-B is then given by

$$\begin{aligned} \hat{C}_t(\mathbf{u}) &= \left\{ w(t) : |w(t)| \leq \hat{Q}_t(\mathbf{u}) \right\}, \\ \hat{C}_u &= \hat{C}_0(\mathbf{u}) \times \cdots \times \hat{C}_{T-1}(\mathbf{u}). \end{aligned}$$

In this sense, the E-IRPC algorithm is a finite-sample approximation of I-IRPC. This connection is essential to the upcoming analysis.

IV. MAIN RESULTS

In this section, we provide a theoretical analysis of our proposed method. We start with an examination of the I-IRPC, where we give sufficient conditions that guarantees 1) the existence and uniqueness of the performatively stable control, 2) convergence to the performatively stable control and 3) the proximity of the performatively stable control to the performatively optimal control. We then extend these results to the E-IRPC.

A. Preliminaries

We start by presenting some standard assumptions on the loss function $L(\mathbf{w}, \mathbf{u})$ and an important assumption regarding the noise distribution.

Assumption 1 (λ -strong convexity in $\mathbf{u}$). We assume that the loss function $L(\mathbf{w}, \mathbf{u})$ is λ -strongly convex in $\mathbf{u}$.

Assumption 2 (β -smoothness in $\mathbf{w}$). We assume that the loss function $L(\mathbf{w}, \mathbf{u})$ is β -smooth in $\mathbf{w}$,

Assumption 3 (Lipschitz continuity of quantile functions). Let $Q_t(\mathbf{u})$ be the $(1-p/T)$ -quantile function for $|w(t)|$ when $\mathbf{u}$ is the control input and p is the predefined probability level (the definition of $Q_t(\mathbf{u})$ is given in Equation 3.10). We assume that $Q_t(\mathbf{u})$ is ϵ_t -Lipschitz, i.e. for all $\mathbf{u}, \mathbf{u}' \in U$, $|Q_t(\mathbf{u}) - Q_t(\mathbf{u}')| \leq \epsilon_t |\mathbf{u} - \mathbf{u}'|$.

B. Existence of fixed-points and optima

Let $g(\mathbf{u}, \mathbf{v}) = \max\{L(\mathbf{w}, \mathbf{v}) : \mathbf{w} \in C(\mathbf{u})\}$.

Theorem 1 (Existence of performatively stable control). Suppose the following assumptions hold:

- $L(\mathbf{w}, \mathbf{u})$ is strictly convex in $\mathbf{u}$,
- $g(\mathbf{u}, \mathbf{v})$ jointly continuous in $\mathbf{u}$ and $\mathbf{v}$,
- the quantile functions $Q_t(\mathbf{u})$ are ϵ_t in $\mathbf{u}$ for all $t = 0, \dots, T-1$ (Assumption 3), and
- U is compact and convex.

Then there exists at least one performatively stable control.

Proof. Since $g(\mathbf{u}, \mathbf{v})$ jointly continuous, by the Maximum Theorem (see, e.g. [10, Chapter E.3]), we know that the set valued mapping given by $\mathbf{u} \mapsto \arg \min\{g(\mathbf{u}, \mathbf{v}) : \mathbf{v} \in U\}$ is upper hemicontinuous. Further, by strict convexity of $L(\mathbf{w}, \mathbf{u})$ in $\mathbf{u}$, $g(\mathbf{u}, \mathbf{v})$ is also strictly convex in $\mathbf{u}$. Therefore, the set $\arg \min\{g(\mathbf{u}, \mathbf{v}) : \mathbf{v} \in U\}$ contains only one element, which is $\mathcal{A}(\mathbf{u})$. This implies that $\mathcal{A}(\mathbf{u})$ is continuous. The existence of a fixed points then follows by the Brouwer's Fixed-Point Theorem. $\square$

Uniqueness of the performatively stable control requires a bit more conditions, as we will see shortly. The existence and uniqueness of the performatively optimal control, on the other hand, is much easier to establish.

Theorem 2 (Existence and Uniqueness of performatively optimal control). Suppose the following assumptions hold:

- $L(\mathbf{w}, \mathbf{u})$ is jointly continuous in $\mathbf{u}$ and
- U is compact and convex.

Then there exists a unique performatively optimal control.

Proof. Assumption 3 implies that the set-valued map $\mathbf{u} \mapsto C(\mathbf{u})$ is continuous. Define $h(\mathbf{u}) := \max_{\mathbf{w} \in C(\mathbf{u})} L(\mathbf{w}, \mathbf{u}) = g(\mathbf{u}, \mathbf{u})$. Notice the performative optimal control is the minimizer of $h(\mathbf{u})$ over U . Since $C(\mathbf{u})$ is continuous, by the Maximum Theorem, $h(\mathbf{u})$ is continuous in $\mathbf{u}$. The existence of the performative optimal control then follows from the Weierstrass Extreme Value Theorem. $\square$

C. Convergence

We are now ready to investigate the convergence properties of the I-IRPC algorithm. Our analysis relies on the following additional assumption.

Definition 3. We say that the worst case noise align for the loss function L and the confidence set mapping C if for any $\mathbf{v}, \mathbf{u}, \mathbf{u}' \in U$, align if there exists $\mathbf{w}_1^* \in \arg \max_{\mathbf{w} \in C(\mathbf{u})} L(\mathbf{w}, \mathbf{v})$ and $\mathbf{w}_2^* \in \arg \max_{\mathbf{w} \in C(\mathbf{u}')} L(\mathbf{w}, \mathbf{v})$ such that $\mathbf{w}_1^*/|\mathbf{w}_1^*| = \mathbf{w}_2^*/|\mathbf{w}_2^*|$.

Assumption 4 (Alignment of worst-case noises). We assume that the worst-case noises align.

Theorem 3 (Convergence of I-IRPC to $\mathbf{u}_{PS}$). Under Assumptions 1, 2, 3, and 4, the I-IRPC converges to a unique performative stable control $\mathbf{u}_{PS}$ at a linear rate: $|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \delta$ for all

$$i \geq \left(1 - \frac{\beta \sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda}\right)^{-1} \log\left(\frac{|\mathbf{u}_0 - \mathbf{u}_{PS}|}{\delta}\right).$$

Proof. Let $\mathbf{u}, \mathbf{u}' \in U$ be two different control inputs. Since $L(\mathbf{w}, \mathbf{u})$ is λ -strongly convex in $\mathbf{u}$, for any $\mathbf{u} \in U$, $g(\mathbf{u}, \mathbf{v})$ is λ -strongly convex in $\mathbf{u}$. Let $\partial_{\mathbf{u}} g(\mathbf{u}, \mathbf{v})$ denote the subgradient of $g(\mathbf{u}, \mathbf{v})$ with respect to $\mathbf{u}$. Take any $s_{\mathbf{u}} \in \partial_{\mathcal{A}(\mathbf{u})} g(\mathbf{u}, \mathcal{A}(\mathbf{u}))$ and $s_{\mathbf{u}'} \in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u}, \mathcal{A}(\mathbf{u}'))$, by strong convexity and definition of subgradients, we can derive $\lambda|\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}')|^2 \leq (\mathcal{A}(\mathbf{u}') - \mathcal{A}(\mathbf{u}))^\top (s_{\mathbf{u}'} - s_{\mathbf{u}})$. Notice again that this hold for any $s_{\mathbf{u}'} \in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u}, \mathcal{A}(\mathbf{u}'))$ and any $s_{\mathbf{u}} \in \partial_{\mathcal{A}(\mathbf{u})} g(\mathbf{u}, \mathcal{A}(\mathbf{u}))$. Now, take $\mathbf{w}_1^* \in \arg \max_{\mathbf{w} \in C(\mathbf{u})} L(\mathbf{w}, \mathcal{A}(\mathbf{u}'))$ and $\mathbf{w}_2^* \in \arg \max_{\mathbf{w} \in C(\mathbf{u}')} L(\mathbf{w}, \mathcal{A}(\mathbf{u}'))$. It's easy to show that $\nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_1^*, \mathcal{A}(\mathbf{u}')) \in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u}, \mathcal{A}(\mathbf{u}'))$ and $\nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_2^*, \mathcal{A}(\mathbf{u}')) \in \partial_{\mathcal{A}(\mathbf{u}')} g(\mathbf{u}', \mathcal{A}(\mathbf{u}'))$. Combining these, we have

$$\begin{aligned} & |\mathcal{A}(\mathbf{u}) - \mathcal{A}(\mathbf{u}')| \\ & \leq \frac{1}{\lambda} |\nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_1^*, \mathcal{A}(\mathbf{u}')) - \nabla_{\mathcal{A}(\mathbf{u}')} L(\mathbf{w}_2^*, \mathcal{A}(\mathbf{u}'))| \\ & \stackrel{(i)}{\leq} \frac{\beta}{\lambda} |\mathbf{w}_1^* - \mathbf{w}_2^*| \\ & \stackrel{(ii)}{=} \frac{\beta}{\lambda} \sqrt{\sum_{t=0}^{T-1} |Q_t(\mathbf{u}) - Q_t(\mathbf{u}')|^2} \\ & \stackrel{(iii)}{\leq} \frac{\beta \sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda} |\mathbf{u} - \mathbf{u}'|. \end{aligned}$$

where step (i) follows from β -smoothness (Assumption 2), step (ii) follows from Assumption 4 and step (iii) follows from ϵ_t -Lipschitz continuity (Assumption 3). Thus, when $\frac{\beta \sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda} < 1$, the I-IRPC converges as stated by the Contraction Mapping Theorem (see, for example, [5, Appendix B]). $\square$

D. Relating Performatively Stable Control to Performatively Optimal Control

So far, we have identified the conditions that allow the I-IRPC to converges to a unique performatively stable control,

$\mathbf{u}_{PS}$. However, what we really care about is the performatively optimal control, $\mathbf{u}_{PO}$, which is the solution to Problem 2.3. In this section, we will establish the sufficient condition that makes $\mathbf{u}_{PO}$ close to $\mathbf{u}_{PS}$. To do so, we first need to make the following additional assumption.

Assumption 5 (L_w -Lipschitz continuity in $\mathbf{w}$). We assume that $L(\mathbf{w}, \mathbf{u})$ is L_w -Lipschitz in $\mathbf{w}$, i.e. for any $\mathbf{w}, \mathbf{w}' \in \mathbb{R}^{nT}$ and $\mathbf{u} \in U$, $|L(\mathbf{w}, \mathbf{u}) - L(\mathbf{w}', \mathbf{u})| \leq L_w |\mathbf{w} - \mathbf{w}'|$.

Theorem 4 (Relating $\mathbf{u}_{PS}$ and $\mathbf{u}_{PO}$). In addition to assumptions made in Theorem 3, if Assumption 5 is also satisfied, then

$$|\mathbf{u}_{PS} - \mathbf{u}_{PO}| \leq \frac{2L_w \sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda}.$$

Proof. By the definition of $\mathbf{u}_{PO}$ and $\mathbf{u}_{PS}$, we have $g(\mathbf{u}_{PO}, \mathbf{u}_{PO}) \leq g(\mathbf{u}_{PS}, \mathbf{u}_{PS}) \leq g(\mathbf{u}_{PS}, \mathbf{u}_{PO})$. Since $\mathbf{u}_{PS}$ is a maximizer of $g(\mathbf{u}_{PS}, \mathbf{v})$, we know that $0 \in \partial_{\mathbf{v}} g(\mathbf{u}_{PS}, \mathbf{v})$. Therefore, by λ -strong convexity (Assumption 1), we have $g(\mathbf{u}_{PS}, \mathbf{u}_{PO}) - g(\mathbf{u}_{PS}, \mathbf{u}_{PS}) \geq \frac{\lambda}{2} |\mathbf{u}_{PS} - \mathbf{u}_{PO}|^2$. Furthermore, by Assumption 5, we also have $g(\mathbf{u}_{PS}, \mathbf{u}_{PO}) - g(\mathbf{u}_{PO}, \mathbf{u}_{PO}) \leq L_w |\mathbf{w}_1^* - \mathbf{w}_2^*|$. Combining above equations and applying Assumption 4, we get

$$|\mathbf{u}_{PS} - \mathbf{u}_{PO}|^2 \leq \left(\frac{2L_w}{\lambda} \sqrt{\sum_{t=0}^{T-1} \epsilon_t^2} \right) |\mathbf{u}_{PS} - \mathbf{u}_{PO}|.$$

Simplifying this completes the proof. $\square$

E. Finite Sample Results

So far, we have focused our analysis on the I-IRPC, which is the ideal version of our proposed method. Now we are finally ready to bridge the gap between them, and establish the convergence properties of our E-IRPC algorithm. The key here is to view the radius of confidence sets obtained in Section III-B as the empirical approximation of the quantile function defined by Equation 3.10.

Definition 4 (Hazard rate). Let X be a 1-dimensional random variable. The hazard rate of X evaluated at point x is defined by

$$h(x) = \lim_{\theta \rightarrow 0} \frac{\mathbb{P}(x \leq X \leq x + \theta | X \geq x)}{\theta} = \frac{f(x)}{1 - F(x)},$$

where f is the p.d.f (assuming that it exists) and F is the c.d.f.

Assumption 6 (Additional assumptions on the distribution of noise). We assume that for any $\mathbf{u} \in U$, we have

- the hazard rate of $|w|$ is positive and non-decreasing
- the p.d.f of $|w|$ exists and is continuously differentiable.

This assumption is the combination of [14, Assumption 1 and 2].

Theorem 5 (Convergence of E-IRPC to a neighborhood of $\mathbf{u}_{PS}$). Under Assumptions 1, 2, 3, 4, and 6, the E-IRPC converges to the neighborhood of a unique performative stable

control $\mathbf{u}_{PS}$ at a linear rate. Specifically, for any $\delta \in (0, 1)$, if we take

$$N_i = \mathcal{O} \left(\frac{4\lambda^2 T^3}{\beta^2 \delta^2 \sum_{t=0}^{T-1} \epsilon_t^2} \log \left(\frac{6p}{\pi^2 i^2 T^2} \right) \right)$$

samples in each iteration, then with probability $1 - p$, $|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \delta$ for all $i \geq (1 - 2\alpha_1)^{-1} \log \left(\frac{|\mathbf{u}_0 - \mathbf{u}_{PS}|}{\delta} \right)$ with $\alpha_1 := \frac{\beta \sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda}$.

Proof. For each timestep t , take the N_i samples of $w(t)$ and sort their Euclidean norm in the increasing order as follows $|w^{(\eta_1)}(t)| \geq \dots \geq |w^{(\eta_{N_i})}(t)|$. Let $k = \left\lfloor \frac{Np}{T} \right\rfloor$, then the empirical quantile function is given by

$$\begin{aligned} \hat{Q}_t(\mathbf{u}) &:= \inf \left\{ r : \frac{1}{N_i} \sum_{j=1}^{N_i} \mathbf{1} \{ |w^j(t)| \leq r \} \geq 1 - \frac{p}{T} \right\} \\ &= |w^{(\eta_k)}(t)|. \end{aligned}$$

Then for any $\mathbf{u}, \mathbf{u}' \in U$, we have $|\hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u}')| \leq |\hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u})| + |Q_t(\mathbf{u}) - Q_t(\mathbf{u}')|$. By [14, Theorem 2], for any $\epsilon > 0$, we have

$$\begin{aligned} &\mathbb{P} \left(|\hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u})| \geq \epsilon \right) \\ &\leq \exp \left(-\frac{\epsilon^2}{2(v^r + (c^r + \omega_n)\epsilon)} \right) + \exp \left(-\frac{\epsilon^2}{2(v^l + \omega_n\epsilon)} \right) \end{aligned}$$

where $v^r = \frac{2}{kL^2}$, $v^l = \frac{2(N_i - k + 1)}{(k-1)^2 L^2}$, $c^r = \frac{2}{kL}$ and $\omega_n = \frac{b}{N_i}$, with L and b finite for all $\mathbf{u} \in U$. With some algebraic manipulations, we get that with $N_i = \mathcal{O} \left(\frac{4\lambda^2 T^3}{\beta^2 \delta^2 \sum_{t=0}^{T-1} \epsilon_t^2} \log \left(\frac{6p}{\pi^2 i^2 T^2} \right) \right)$,

$$\mathbb{P} \left(\left| \hat{Q}_t(\mathbf{u}_i) - Q_t(\mathbf{u}_i) \right| \geq \frac{\beta \sqrt{\sum_{t=0}^{T-1} \epsilon_t^2}}{\lambda T} \delta \right) \leq \frac{6p}{\pi^2 i^2 T}.$$

Following the same arguments as in the proof of Theorem 3, with probability at least $1 - \frac{6p}{\pi^2 i^2}$, we have

$$|\hat{\mathcal{A}}(\mathbf{u}_i) - \mathcal{A}(\mathbf{u}_i)| \leq \frac{\beta}{\lambda} \sqrt{\sum_{t=0}^{T-1} \left| \hat{Q}_t(\mathbf{u}) - Q_t(\mathbf{u}) \right|^2} \leq \alpha_1 \delta.$$

Now, applying Theorem 3 we get that when $|\mathbf{u}_i - \mathbf{u}_{PS}| \geq \delta$, $|\mathbf{u}_{i+1} - \mathbf{u}_{PS}| \leq 2\alpha_1 |\mathbf{u}_i - \mathbf{u}_{PS}|$. Therefore, when $\alpha_1 < \frac{1}{2}$, applying the union bound, with probability at least $1 - \sum_{i=0}^{\infty} \frac{6p}{\pi^2 i^2} = 1 - p$, for all $i = 0, 1, \dots$ we have $|\mathbf{u}_i - \mathbf{u}_{PS}| \leq \max \left\{ (2\alpha_1)^i |\mathbf{u}_0 - \mathbf{u}_{PS}|, \delta \right\}$. This completes the proof. $\square$

To summarize our analysis, Theorem 1 establishes the existence of a performative stable control, which means the I-IRPC algorithm has a fixed point. Theorem 2 establishes the existence of a performative optimal control, in which case the problem we investigate is meaningful. Theorems 3 gives a set of conditions for the I-IRPC algorithm to converge to a performative control, and Theorems 4 examines when the corresponding performative control is close to optimal. Finally,

Theorems 5 extends the convergence results in Theorems 3 to our main method E-IRPC.

V. CONCLUSION

In summary, we present a novel way to approach control tasks with noisy dynamical systems, requiring no prior knowledge of the noise. In particular, we consider settings where the noise depends on our control in unknown ways, which requires us to re-estimate the uncertainty whenever a different control is used. We derive convergence bounds under certain regularity conditions. Note that although we mainly focus on open-loop control in this paper, we believe it is possible to extend our method to closed-loop control, and this will be our primary direction for future research.

REFERENCES

- [1] A. K. Akametalu, S. Kaynama, J. F. Fisac, M. N. Zeilinger, J. H. Gillula, and C. J. Tomlin, “Reachability-Based Safe Learning with Gaussian Processes,” *53rd IEEE Conference on Decision and Control*, pp. 1424–1431, 2014.
- [2] A. N. Angelopoulos and S. Bates, *A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification*, 2022. arXiv: [2107.07511 \[cs.LG\]](#).
- [3] S. Bansal and C. J. Tomlin, “DeepReach: A Deep Learning Approach to High-Dimensional Reachability,” *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1817–1824, 2020.
- [4] J. H. Gillula and C. J. Tomlin, “Reducing Conservativeness in Safety Guarantees by Learning Disturbances Online: Iterated Guaranteed Safe Online Learning,” in *Robotics: Science and Systems VIII*, The MIT Press, Jul. 2013.
- [5] H. Khalil, *Nonlinear Systems*. Prentice Hall, 2002.
- [6] A. Lin and S. Bansal, *Verification of Neural Reachable Tubes via Scenario Optimization and Conformal Prediction*, 2023. arXiv: [2312.08604 \[cs.RO\]](#).
- [7] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, “Safe Planning in Dynamic Environments using Conformal Prediction,” *IEEE Robotics and Automation Letters*, 2023.
- [8] L. Lindemann, X. Qin, J. V. Deshmukh, and G. J. Pappas, “Conformal prediction for stl runtime verification,” *2023 59th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1–1, 2022.
- [9] A. Muthali, H. Shen, S. Deglurkar, et al., *Multi-Agent Reachability Calibration with Conformal Prediction*, 2023. arXiv: [2304.00432 \[eess.SY\]](#).
- [10] E. Ok, *Real Analysis with Economic Applications*. Princeton University Press, 2007.
- [11] J. Perdomo, T. Zrnic, C. Mendler-Dünnner, and M. Hardt, “Performative Prediction,” in *International Conference on Machine Learning*, PMLR, 2020, pp. 7599–7609.

- [12] S. Prajna, A. Jadbabaie, and G. J. Pappas, “A Framework for Worst-Case and Stochastic Safety Verification Using Barrier Certificates,” *IEEE Transactions on Automatic Control*, vol. 52, no. 8, pp. 1415–1428, 2007.
- [13] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*, 2nd ed. Springer International Publishing, 2022.
- [14] M. Zhang and C. S. Ong, “Quantile Bandits for Best Arms Identification,” in *International Conference on Machine Learning*, PMLR, 2021, pp. 12 513–12 523.