Optimization-based frequentist confidence intervals for functionals in constrained inverse problems: Resolving the Burrus conjecture

Pau Battle[1] Pratik Patil[2] Mike Stanley[3] Houman Owhadi[1] Mikael Kuusala[3]
[1]California Institute of Technology [2]University of California, Berkeley [3]Carnegie Mellon University

Motivation

In safety-based applications, we need Uncertainty Quantification (UQ) that is:

- **Certifiable**: With a clear and transparent set of assumptions and outputs
- **Well-calibrated**: Precisely adjustable to the desired level of coverage

If it **undercovers**, we incur **unnecessary risk**, if it **overcovers**, we incur **economic cost**

Inverse problems setup

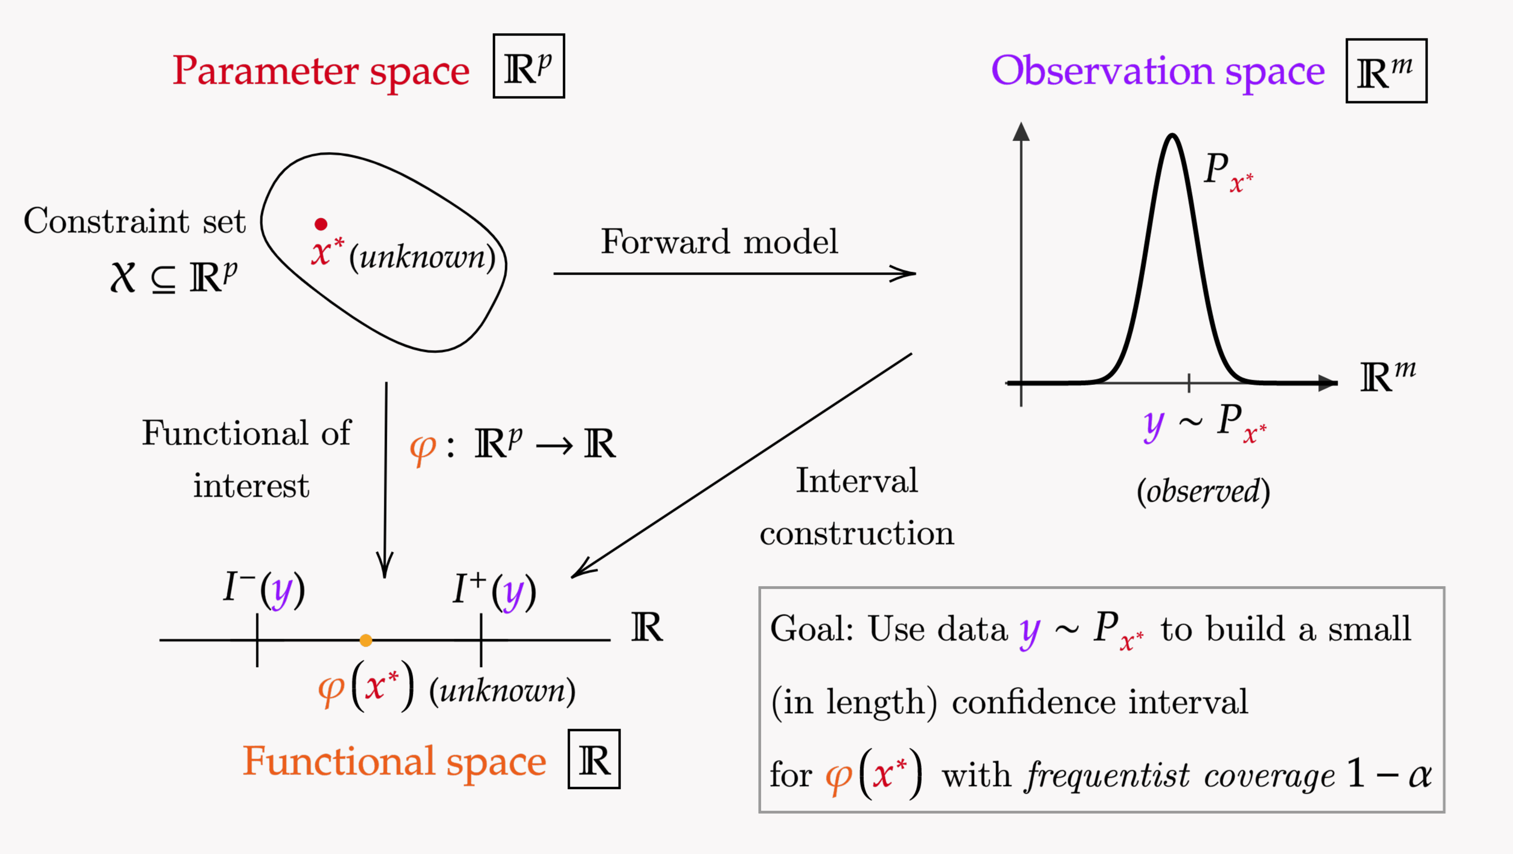

Goal

Find an interval **as small as possible** with **finite sample guarantees** that contains $\varphi(x^*)$ with probability $1-\alpha$:

$$\inf_{x \in \mathcal{X}} \mathbb{P}_{y \sim P_x}\left(\varphi(x) \in [I^-(y), I^+(y)]\right) \geq 1-\alpha$$

Stronger notion of intervals than Bayesian credible intervals and only assumes likelihood and constraints

Previous approach

A general method to build intervals with correct coverage is the **simultaneous approach** [Stark '92, '94]:

(1) Find a $1-\alpha$ confidence set $\mathcal{C}(y)$ for x^*
(2) Intersect it with the constraint set $\mathcal{X}$
(3) Map through $\varphi(x)$

It overcovers, since $\mathcal{C}(y)$ does not need to be $1-\alpha$ for the interval to be $1-\alpha$

Linear Gaussian model with linear constraints

$$y = Kx^* + \varepsilon, \quad \mathcal{X} = \{x : Ax \leq b\}$$

with $y \in \mathbb{R}^m$, $x^* \in \mathbb{R}^p$, $\varepsilon \sim \mathcal{N}(0, I_m)$, and $\varphi(x) = h^T x$

- Generally ill-posed, the constraint allows finite-length intervals
- Studied in unfolding gamma-ray and neutron spectra applications
- Most studied constrained inference problem

Burrus conjecture (1965)

Consider the model $y = Kx^* + \varepsilon, \varphi(x) = h^T x$,
A$x^* \leq b, \varepsilon \sim \mathcal{N}(0, I_m)$
A valid $1-\alpha$ interval for $\varphi(x^*)$ has extremes given by:

$$\begin{aligned} \min_x / \max_x \quad & h^T x \\ \text{s.t.} \quad & \|y - Kx\|_2^2 \leq \psi_\alpha^2 \\ & Ax \leq b \end{aligned}$$

With ψ_α^2 much smaller than the simultaneous approach

Contributions

- New **optimization-based approach** to obtain $1-\alpha$ confidence intervals with frequentist guarantees
- Comes with an **algorithm that beats previous approaches** in toy and real problems
- **Theorem: The Burrus conjecture is false**

Numerical example

Our method fixes the Burrus conjecture intervals when they undercover

$x, y \in \mathbb{R}^3, y = x + \varepsilon, x^* \geq 0, \varphi(x) = x_1 + x_2 - x_3, \varepsilon \sim \mathcal{N}(0, I)$