Data-driven Reachability using Christoffel Functions and Conformal Prediction

Abdelmouaiz Tebjou, Goran Frehse, Faicel Chamroukhi

IRT SYSTEMX , ENSTA Paris France

COPA 2023

Introduction

- **Objective**: Real-time detection of deviations in deployed AI components from expected or safe behavior using online monitoring.
- **Envelope-based Models**: Dynamic models that delineate possible future states based on past observations. Used for checking if current AI prediction is within defined boundaries or envelopes.
- **Goal**: Develop a methodological approach suitable for real-time monitoring systems.

Reachability analysis

- **Reachability Analysis**: Determine future states of a system from an initial state using a transition function.
- **Focus**: Estimating the image of f for a set of initial states, specifically the reachable set:

$$S = \{f(\vec{x}) : \vec{x} \in I\}.$$

- **Challenges**: For complex or unknown f, we aim for an approximation $\hat{S}$ that mostly covers S.
- **Solution Approach**: Represent set S with probability measure μ whose support is S.

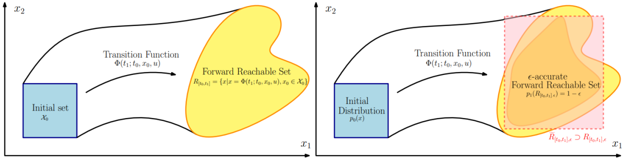

Preliminaries

- When working with n dimensions, we denote the number of monomials of degree less or equal to d with:

$$s(d) = \binom{n+d}{n}.$$

- Let $\vec{v}_d(\vec{x}) \in \mathbb{R}^{s(d)}$ be the vector of monomials of degree up to d. For $d = 2$ and $n = 2$: $\vec{v}_d(\vec{x}) = [1 \ x_1 \ x_2 \ x_1 x_2 \ x_1^2 \ x_2^2]$.

Christoffel function

- Christoffel functions are a class of functions associated with a finite measure and a parameter degree $d \in \mathbb{N}$ with a strong connection to approximation theory.
- **Moment Matrix**: For a finite measure μ on $\mathbb{R}^n$ and integer degree d,

$$\mathbf{M}_d = \int_{\mathbb{R}^n} \vec{v}_d(\vec{x})\vec{v}_d(\vec{x})^\top d\mu(\vec{x}).$$

- **Christoffel Function**: Defined as

$$\Lambda_{\mu,d}(\vec{x}) = \left(\mathbf{v}_d(\vec{x})^T \mathbf{M}_d^{-1} \mathbf{v}_d(\vec{x})\right)^{-1}.$$

Christoffel function

Since S is unknown and therefore μ is also unkown.

- **Empirical Measure**: μ approximated as :

$$\hat{\mu} = \frac{1}{N}\sum_{i=1}^{N} \delta_{\vec{x}^i},$$

  leading to the empirical moment matrix

$$\widehat{\mathbf{M}}_d = \frac{1}{N}\sum_{i=1}^{N} \mathbf{v}_d(\mathbf{x^i})\mathbf{v}_d(\mathbf{x^i})^T.$$

- **Empirical Christoffel Function**: Given by

$$\Lambda_{\hat{\mu},d}^{-1}(\vec{x}) = \mathbf{v}_d(\vec{x})^T \widehat{\mathbf{M}_d}^{-1} \mathbf{v}_d(\vec{x}).$$

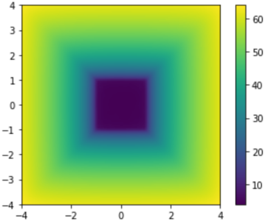

Set Approximation with Christoffel Functions

- [?] proposed various thresholding schemes for approximating the support of a probability measure using the Christoffel function or, more precisely, its empirical counterpart.
- This idea was applied by [?] to approximate the reachable set S with the superlevel sets of the Christoffel function.

[Thm. 1 in [?]] Given a training set of i.i.d samples $D = \{\vec{x}^1, \dots, \vec{x}^N\}$ from S, let

$$\hat{S} = \{\vec{x} \in \mathbb{R}^n \mid \Lambda_{\hat{\mu},d}^{-1}(\vec{x}) \leq \max_i \Lambda_{\hat{\mu},d}^{-1}(\vec{x}^i)\}. \quad (1)$$

If $N \geq \frac{5}{\epsilon}\left(\log \frac{4}{\delta} + \binom{n+2d}{n} \log \frac{40}{\epsilon}\right)$, then $\mathbb{P}\left(\mu\left(\hat{S}\right) \geq 1 - \epsilon\right) \geq 1 - \delta$.

Reach Set Approximation with Conformal Prediction

Intuitively, the Christoffel polynomial takes high values where the density is low and low values where the density is high, which makes it a good candidate for a nonconformity function.

- For $i \in \{0, ..., N\}$, the conformal region is defined as

$$C_D^{\frac{i}{N}} = \left\{ \vec{x} \in \mathbb{R}^n \;\middle|\; p_{value}(\vec{x}) \geq \frac{i}{N} \right\}$$

Theorem

Suppose that the nonconformity function $r(\vec{x})$ is continuous. $\forall \delta \in (0, 1)$,

$$\mathbb{P}\left[\exp\left(\frac{\log(1-\delta)}{N}\right) \geq \mu\left(C_{D_{\mathrm{cal}}}^{\frac{1}{N}}\right) \geq \exp\left(\frac{\log(\delta)}{N}\right)\right] \geq 1 - 2\delta. \tag{2}$$

Reach Set Approximation with Conformal Prediction(without outliers)

\# $D_{\text{train}} = \{\vec{x}^{N+1}, \ldots, \vec{x}^{M}\}$ and $D_{\text{cal}} = \{\vec{x}^{1}, \ldots, \vec{x}^{N}\}$

1. Compute the empirical moment matrix $\widehat{\mathbf{M}}_d$ and its inverse
    1. $\widehat{\mathbf{M}}_d = \frac{1}{M-N} \sum_{i=N+1}^{M} \mathbf{v}_d\left(\vec{x}^i\right) \mathbf{v}_d\left(\vec{x}^i\right)^\top$, with $\vec{x}^i \in D_{\text{train}}$
    2. Compute $\widehat{\mathbf{M}}_d^{-1}$.
2. Calculate the threshold α: $\alpha = \max_{i=1,\ldots,N} \mathbf{v}_d\left(\vec{x}^i\right)^\top \widehat{\mathbf{M}}_d^{-1} \mathbf{v}_d\left(\vec{x}^i\right)$, with $\vec{x}^i \in D_{\text{cal}}$
3. Given the returned $\widehat{\mathbf{M}}_d^{-1}$ and α, record the conformal region:

$$C_{D_{\text{cal}}}^{\frac{1}{N}} = \hat{S} = \left\{\vec{x} \in \mathbb{R}^n \,\middle|\, \mathbf{v}_d(\vec{x})^\top \widehat{\mathbf{M}}_d^{-1} \mathbf{v}_d(\vec{x}) \leq \alpha\right\}$$

Experiments

Let the transition function $f : \mathbb{R}^2 \to \mathbb{R}^2$ be

$$f(x, y) = (1 + sign(x) \cdot x^2, 1 + sign(y) \cdot y^2)$$

and let the initial set be $I = [-1, 1]^2$. The reachable set consists of four squares, i.e.,

$$S = [-3, -1]^2 \cup [-3, -1] \times [1, 3] \cup [1, 3] \times [-3, -1] \cup [1, 3]^2.$$

Figure: Reach set approximations (outlined in purple) with a sample size of $M = 1000$, of which $N = 200$ (red dots) are used as a calibration set.

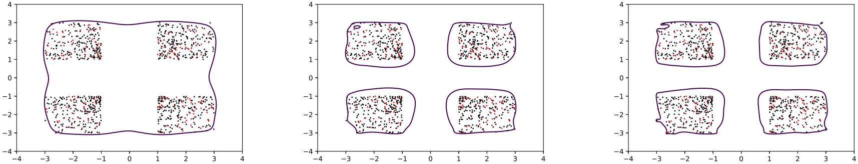

(a) $d = 6$ (b) $d = 10$ (c) $d = 15$

Avoiding the Calibration Set (Transductive CP)

- The calibration set is taken to be the entire training set plus the point at which the function is evaluated, in other words a new non conformity is modulated by the data point.
- we need to compute a new moment matrix and invert it for each evaluation which is computationally expensive, on the order of $\mathcal{O}(s(d)^3)$.
- To avoid this, we compute the new inverse moment matrix incrementally using the Sherman-Morrison formula :

$$(A + uu^T)^{-1} = A^{-1} - \frac{A^{-1}uu^TA^{-1}}{1 + u^TA^{-1}u}$$

.

Robustness to outliers

- One may have to work with a calibration set containing outliers without knowing which data point is an outlier and which one isn't.
- The presence of outliers in the training set does not affect the theoretical guarantees obtained using conformal prediction theory, though it will affect the tightness of the approximated reachable set.

Theorem

Consider a set of points $D = \{\vec{x}^1, \vec{x}^2, ..., \vec{x}^N\}$ containing no more than p outliers, with $2p + 1 < N$, and where the rest of samples are i.i.d from a probability measure μ. Then for any i.i.d vector $\vec{x}$ sampled from μ and $\epsilon \in (0, 1)$,

$$\mathbb{P}\left(\mu\left(C_D^{\frac{p+1}{N}}\right) \geq 1 - \epsilon\right) \geq \sum_{i=p+1}^{N-p} \binom{N-p}{i} \epsilon^i (1-\epsilon)^{N-p-i} \qquad (3)$$

Figure: An approximation of the reach set (purple outline) obtained with using a Christoffel polynomial of degree 15, on a data set with 10% outliers. The training set is shown in black, the calibration set in red.

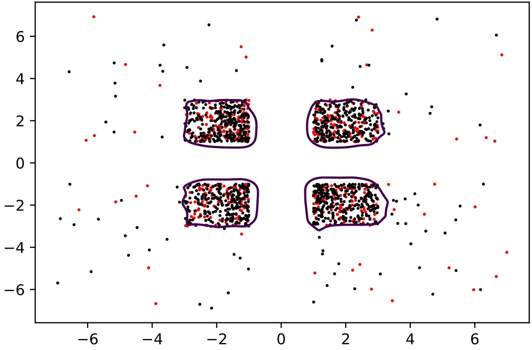

Experiments

Figure: Reach set approximations (purple outline) using one-class SVM, Isolation Forest, and Local Outlier Factor (LOF) as nonconformity functions, for a common training set of size 800 (black dots) and calibration set of size 200 (red dots).

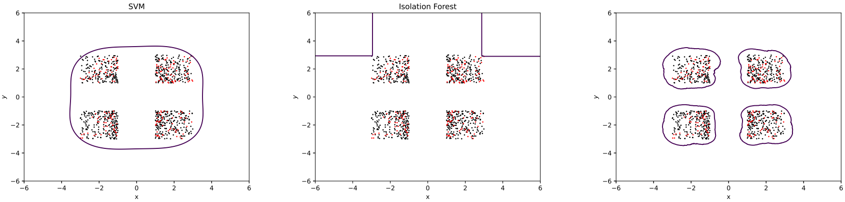

(a) One-class SVM (b) Isolation Forest (c) LOF

Under Outliers

Figure: Comparison of reachable set approximations for the empirical Christoffel polynomial (degree 10) and LOF with the region $[-1, 1]^2$ as the target. The training set, containing outliers, is represented by black dots, while the calibration set is shown in red. The plot highlights the performance differences and robustness of both methods in the presence of outliers, demonstrating how the empirical Christoffel polynomial is far more robust.

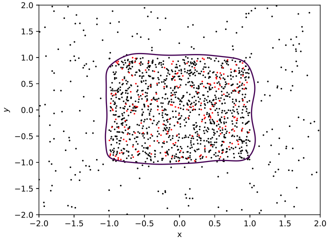

(a) Christoffel polynomial

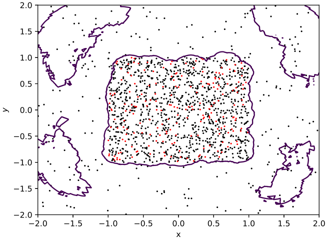

(b) LOF

Experiments

Under Outliers

Figure: A comparison of reach set approximation (purple outline) using the Christoffel polynomial with degree 15 and LOF which targets a star-shaped region. Training set samples are in black and calibration set in red. This plot highlights the performance and robustness of both methods when encountering outliers in a complex geometric scenario, illustrating the effectiveness of the empirical Christoffel polynomial under the presence of outliers.

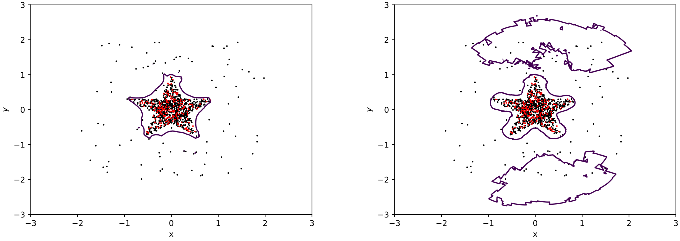

(a) Christoffel polynomial

(b) LOF

Conclusion and future work

- Introduced the use of the Christoffel function in reach set approximation within dynamical systems, enhancing sample efficiency and robustness against outliers.
- Bypassed data splitting by utilizing an incremental form of the Christoffel function for transductive conformal prediction.
- Demonstrated the effectiveness of our approach through extensive numerical experiments.
- Future work will delve into the impact of numerical errors on the computation of the Christoffel function and extend the application to other domains.

Thank You!

Alex Devonport, Forest Yang, Laurent El Ghaoui, and Murat Arcak.
Data-driven reachability analysis with christoffel functions.
In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 5067–5072, 2021.

Jean-Bernard Lasserre and Edouard Pauwels.
The empirical christoffel function with applications in data analysis.
Advances in Computational Mathematics, 45(3):1439–1468, 2019.