Model Predictive Control Under Uncertainty

(Guest Lecture for COMP0211 Control 2 at UCL)

Kai Wang

Department of Aeronautical and Automotive Engineering

Loughborough University

Overview

Backgrounds:

- Model Predictive Control (MPC)
- Types of Uncertainty
- Types of Constraints
- Types of Cost Functions

MPC Under Uncertainty:

- Robust MPC
- Stochastic MPC
- Distributionally Robust MPC
- Tube MPC

Part I. Backgrounds

Model Predictive Control

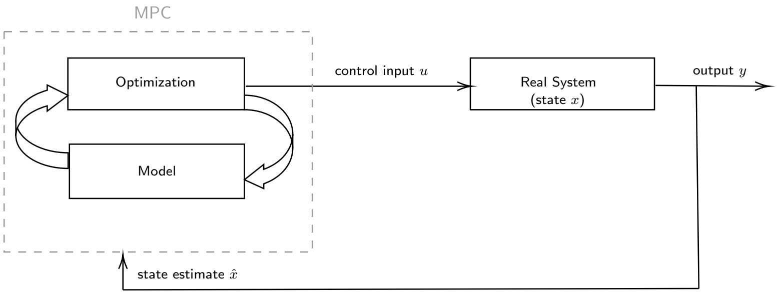

Main idea:

- Once the state estimate $\hat{x}$ is received, MPC obtains the control action u by solving an optimization problem, and then the process is repeated.
- Constraints and a performance index are incorporated in the model.
- In what follows, the state x is assumed to be perfectly known ($\hat{x} = x$). If not, the problem falls under the scope of Output Feedback MPC.

Model Predictive Control

MPC Formulation:

$$V_N^\star(x) := \underset{\mathbf{d}_N}{\text{minimize}} \sum_{k=0}^{N-1} \ell(x_k, u_k) + V_f(x_N)$$

$$\text{subject to} \begin{cases} \forall k \in \mathbb{N}_{N-1},\ x_{k+1} = f(x_k, u_k), \\ \forall k \in \mathbb{N}_{N-1},\ (x_k, u_k) \in \mathcal{Y}, \\ x_0 = x,\ \text{and } x_N \in \mathcal{X}_f, \end{cases}$$

with $\mathbf{d}_N := (x_0, \ldots, x_N, u_0, \ldots, u_{N-1})$.

- Control system: $x^+ = f(x, u)$
- Constraint: $(x, u) \in \mathcal{Y}$
- Stage cost: $\ell : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$
- Terminal cost: $V_f : \mathbb{R}^n \to \mathbb{R}$
- Terminal set: $\mathcal{X}_f$
- Prediction horizon: N
- Current state: x

Technical Principle:

- Under mild technical conditions (e.g., continuity, compactness/closedness), the optimization problem is well-posed.
- The value function $V_N^\star(\cdot)$, the optimizer $\mathbf{d}_N^\star(\cdot)$, and the feasible domain $\mathcal{C}_N$.
- The component $u_0^\star(x)$ of the optimizer $\mathbf{d}_N^\star(x)$ is employed such that the MPC controlled dynamics is given by

$$\forall x \in \mathcal{C}_N,\ x^+ = f(x, \kappa_N(x)),\ \text{with } \kappa_N(x) = u_0^\star(x).$$

Model Predictive Control

MPC Controlled Dynamics: $\forall x \in \mathcal{C}_N, \ x^+ = f(x, \kappa_N(x))$.

Desirable Control-Theoretic Properties:

- Positive invariance of $\mathcal{C}_N$ (i.e., $\forall x \in \mathcal{C}_N \implies x^+ = f(x, \kappa_N(x)) \in \mathcal{C}_N$), also known as *recursive feasibility* of the underlying optimization problem, is guaranteed *indirectly* by imposing positive invariance on $\mathcal{X}_f$, i.e.,

$$\forall x \in \mathcal{X}_f, \ \exists \kappa_f(x) \in \mathcal{U}(x) \text{ such that } f(x, \kappa_f(x)) \in \mathcal{X}_f.$$

- Asymptotic stability of the origin for the MPC-controlled dynamics is guaranteed *indirectly* via the condition that $V_f(\cdot)$ is a Lyapunov function over $\mathcal{X}_f$, satisfying

$$\forall x \in \mathcal{X}_f, \ V_f(f(x, \kappa_f(x))) - V_f(x) \leq -\ell(x, \kappa_f(x)).$$

Limitations:

- In the presence of disturbances, desirable properties may deteriorate, and constraint violations may occur.
- In safety-critical applications, it becomes necessary to **explicitly** model uncertainty in the MPC formulation, e.g., robust MPC, stochastic MPC, distributionally robust MPC.

Types of Uncertainty

Let w represent process noise and/or model-plant mismatch.

i. Set-membership description: $w \in \mathcal{W}$

- The set $\mathcal{W}$ is bounded, meaning that only worst-case geometric bounds of the uncertainty are available.

ii. Stochastic description: $w \sim \omega$

- The probability distribution/measure ω is known and supported by $\mathcal{W}$, i.e., $\omega(\mathcal{W}) = 1$.
- The uncertainty w is bounded if $\mathcal{W}$ is bounded, and unbounded if $\mathcal{W}$ is unbounded.

iii. Distributionally robust description: $w \sim \omega$, with $\omega \in \Omega$

- The probability distribution ω is not known exactly, but belongs to an ambiguity set Ω.
- The ambiguity set Ω is known, which is a family of possible probability distributions.
- E.g., only the first two moments of ω are known.

iii $\Rightarrow$ ii (when Ω is a singleton). iii $\Rightarrow$ i (when Ω is specified by $\Omega = \{\omega : \omega(\mathcal{W}) = 1\}$)

Types of Constraints

Consider the uncertain system: $x^+ = f(x, u, w)$

- Interpretation 1. At any time $k \in \mathbb{N}$, when the decision on u_k is taken, the state x_k is known, while the current disturbance w_k and all future disturbances w_j, $j > k$ are not known but are guaranteed to obey the set-membership, stochastic, or distributionally robust descriptions.
- Interpretation 2. Even through the state of the system at current time is known, the future state cannot be predicted accurately as the current and the future disturbances are unknown to us.

Consider the state, control and disturbance processes: $\{x_k\}_{k=0}^{\infty}$, $\{u_k\}_{k=0}^{\infty}$ and $\{w_k\}_{k=0}^{\infty}$

- If we know exactly $w_k \sim \omega$, then we suppose $x_k \sim \rho_k$ and $u_k \sim \xi_k$.
- If we know $w_k \sim \omega$, with $\omega \in \Omega$, then we suppose $x_k \sim \rho_k \in \mathcal{P}_k$ and $u_k \sim \xi_k \in \Xi_k$.

Types of Constraints

i. Hard Constraints:

$$(x_k, u_k) \in \mathcal{X} \times \mathcal{U},\ \forall k \in \mathbb{N}.$$

- The disturbance constraint set $\mathcal{W}$ is typically required to be bounded.
  (E.g., If $\mathcal{X}$ and $\mathcal{U}$ are bounded, an unbounded disturbance $w \in \mathcal{W}$ in $x^+ = f(x, u) + w$ would make satisfying hard constraints impossible.)
- While highly desirable in practice, hard constraint satisfaction is not always feasible.

ii. Expectation Constraints:

$$\int_{\mathcal{X}} r(x)\rho_k(\mathrm{d}x) \geq b,\ \ \forall k \in \mathbb{N},$$

where $r : \mathbb{R}^n \to \mathbb{R}$ is a measurable reward function, and $b \in \mathbb{R}$ is a scalar.

- The goal is to ensure that the average value of $r(x)$ over $\mathcal{X}$ remains no less than b.

Types of Constraints

iii. Chance (Probabilistic) Constraints:

$$\Pr(x_k \in \mathcal{X}) = \int_{\mathcal{X}} 1 \,\mathrm{d}\rho_k = \int \mathbf{1}_{\mathcal{X}}(x)\, \rho_k(\mathrm{d}x) = \rho_k(\mathcal{X}) \geq p, \ \forall k \in \mathbb{N},$$

where $\mathbf{1}_{\mathcal{X}} : \mathbb{R}^n \to \{0, 1\}$ is the indicator function of the set $\mathcal{X}$, defined as $\mathbf{1}_{\mathcal{X}}(x) = 1$ if $x \in \mathcal{X}$ and 0 otherwise, and $p \in (0, 1)$ is a user-specified tolerance probability.

- The system state x_k should remain within the set $\mathcal{X}$ with a probability of at least p.

iv. Worst-Case Chance Constraints:

$$\min_{\rho_k \in \mathcal{P}_k} \Pr(x_k \in \mathcal{X}) \geq p, \ \forall k \in \mathbb{N}$$

- It ensures that the chance constraint holds even in the most unfavorable distribution within a specified ambiguity set $\mathcal{P}_k$.

iii and iv are difficult to handle due to complexity of the set $\mathcal{X}$, discontinuity of the indicator function $\mathbf{1}_{\mathcal{X}}$, and typical difficulties in evaluating integral.

Types of Constraints

Key Insights:

1. Hard constraints:
   - pros: guarantee feasibility and safety in all scenarios.
   - cons: may be overly conservative or infeasible under large disturbances.
2. Expectation constraints:
   - pros: balance performance and constraint satisfaction on average.
   - cons: do not prevent rare, significant violations of constraints.
3. Chance constraints:
   - pros: allow flexibility with probabilistic safety guarantees.
   - cons: require knowledge of the distribution; computationally expensive for complex systems.
4. Worst-case chance constraints:
   - pros: ensure robustness against distributional uncertainty.
   - cons: highly conservative and computationally expensive; requires defining ambiguity set

Types of Cost Functions

Define Control Policy:

$$\boldsymbol{\mu} := \{\mu_0(\cdot), \mu_1(\cdot), \ldots, \mu_{N-1}(\cdot)\},$$

where $x \mapsto \mu_k(x)$ represents a control law in the prediction horizon.

Define function:

$$J_N(x, \boldsymbol{\mu}, \boldsymbol{w}) := \sum_{k=0}^{N-1} \ell(x_k, \mu_k(x_k)) + V_f(x_N),$$

where $\boldsymbol{w} := (w_0, w_1, \ldots, w_{N-1})$, and $x_{k+1} = f(x_k, \mu_k(x_k), w_k)$ with $x_0 = x$.

Types of Cost Functions

i. Nominal Cost:

$$V_N(x, \boldsymbol{\mu}) = J_N(x, \boldsymbol{\mu}, \mathbf{0}) = \sum_{k=0}^{N-1} \ell(x_k, \mu_k(x_k)) + V_f(x_N).$$

- For simplicity, the nominal cost is sometimes employed; here $\mathbf{0}$ is defined to be the disturbance sequence $\boldsymbol{w} = (0, 0, \dots, 0)$.

ii. Worst-Case Cost:

$$V_N(x, \boldsymbol{\mu}) = \max_{\boldsymbol{w} \in \mathcal{W}^N} J_N(x, \boldsymbol{\mu}, \boldsymbol{w}) = \max_{\boldsymbol{w} \in \mathcal{W}^N} \sum_{k=0}^{N-1} \ell(x_k, \mu_k(x_k)) + V_f(x_N).$$

- The system's performance against the most adverse disturbances.
- It may lead to overly conservative control actions, sacrificing optimality.

Types of Cost Functions

iii. Expectation Cost:

$$V_N(x, \boldsymbol{\pi}) = \underset{w_k \sim \omega}{\mathbb{E}} \left(J_N(x, \boldsymbol{\mu}, \boldsymbol{w}) \right) = \underset{x_k \sim \rho_k}{\mathbb{E}} \left(\sum_{k=0}^{N-1} \ell(x_k, \mu_k(x_k)) + V_f(x_N) \right).$$

- Require accurate knowledge of disturbance distribution ω.
- Evaluate the average cost over possible disturbance realizations.
- The evaluation requires sampling methods or analytical approaches if possible.

iv. Worst-Case Expectation Cost:

$$V_N(x, \boldsymbol{\pi}) = \max_{\omega \in \Omega} \underset{w_k \sim \omega}{\mathbb{E}} \left(J_N(x, \boldsymbol{\mu}, \boldsymbol{w}) \right) = \max_{\rho_k \in \mathcal{P}_k} \underset{x_k \sim \rho_k}{\mathbb{E}} \left(\sum_{k=0}^{N-1} \ell(x_k, \mu_k(x_k)) + V_f(x_N) \right).$$

- Expectation with respect to the adversarial distribution from the ambiguity set.
- Addresses scenarios where exact probability distributions are partially known.

Part II. MPC Under Uncertainty

Model Predictive Control Under Uncertainty

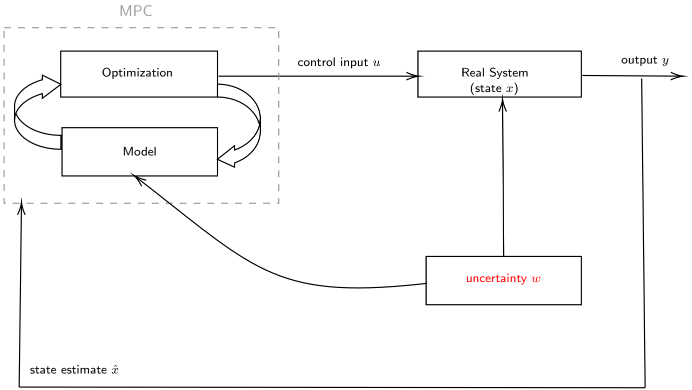

We will explicitly account for uncertainty using robust, stochastic and distributionally robust MPC frameworks, tailored to different uncertainty models

Robust MPC

Robust MPC Formulation:

$$V_N^\star(x) = \min_{\boldsymbol{\mu}} \max_{\boldsymbol{w} \in \mathcal{W}^N} \sum_{k=0}^{N-1} \ell(x_k, \mu_k(x_k)) + V_f(x_N)$$

$$\text{subject to} \begin{cases} \forall k \in \mathbb{N}_{N-1},\ x_{k+1} = f(x_k, \mu_k(x_k), w_k), \\ \forall k \in \mathbb{N}_{N-1},\ (x_k, \mu_k(x_k)) \in \mathcal{X} \times \mathcal{U}, \\ x_0 = x,\ \text{and } x_N \in \mathcal{X}_f. \end{cases}$$

Key Ideas:

- The disturbance set $\mathcal{W}$ must be bounded.
- The terminal set $\mathcal{X}_f$ is required to be robust positively invariant, while V_f is a robust Lyapunov function over $\mathcal{X}_f$.
- The *inner optimization* (max) evaluates the worst-case disturbance, while the *outer optimization* (min) computes the control policy to minimize the worst-case cost.
- Guarantees of robust positive invariance and asymptotic stability are well-established.

Stochatstic MPC

Stochastic MPC Formulation:

$$V_N^\star(x) = \min_{\boldsymbol{\mu}} \mathop{\mathbb{E}}_{x_k \sim \rho_k} \left(\sum_{k=0}^{N-1} \ell(x_k, \mu_k(x_k)) + V_f(x_N) \right)$$

$$\text{subject to} \begin{cases} \forall k \in \mathbb{N}_{N-1},\ x_{k+1} = f(x_k, \mu_k(x_k), w_k), \\ \forall k \in \mathbb{N}_{N-1},\ \Pr(x_k \in \mathcal{X}) \geq p \\ x_0 = x,\ \text{and } x_N \in \mathcal{X}_f. \end{cases}$$

Key Ideas:

- The disturbance distribution $w_k \sim \omega$ is required to be known.
- Hard constraints or chance constraints on the control input $u_k = \mu_k(x_k)$ can be incorporated into the formulation.
- The terminal set $\mathcal{X}_f$ must exhibit positive invariance in a probabilistic sense, and V_f should act as a stochastic Lyapunov function within $\mathcal{X}_f$.
- Guarantees for recursive feasibility and stochastic stability require further research.

Distributionally Robust MPC

Distributionally Robust MPC Formulation:

$$V_N^\star(x) = \min_{\boldsymbol{\mu}} \max_{\rho_k \in \mathcal{P}_k} \mathop{\mathbb{E}}_{x_k \sim \rho_k} \left(\sum_{k=0}^{N-1} \ell(x_k, \mu_k(x_k)) + V_f(x_N) \right).$$

$$\text{subject to} \begin{cases} \forall k \in \mathbb{N}_{N-1},\ x_{k+1} = f(x_k, \mu_k(x_k), w_k), \\ \forall k \in \mathbb{N}_{N-1},\ \min\limits_{\rho_k \in \mathcal{P}_k} \Pr(x_k \in \mathcal{X}) \geq p \\ x_0 = x,\ \text{and } x_N \in \mathcal{X}_f. \end{cases}$$

Key Ideas:

- The exact disturbance distribution ω is unknown, but the ambiguity set Ω needs to be predetermined.
- Hard constraints or worst-case chance constraints on the control input $u_k = \mu_k(x_k)$ can be incorporated into the formulation.
- The terminal set $\mathcal{X}_f$ must exhibit positive invariance in a probabilistic sense with respect to the worst-case disturbance distribution, and V_f should act as a stochastic Lyapunov function within $\mathcal{X}_f$.
- Guarantees of recursive feasibility and stability are even more challenging to establish.

Robust, Stochastic and Distributionally Robust MPC

Comparisons:

Feature	Robust MPC	Stochastic MPC	DR-MPC
Uncertainty Model	worst-case bounded disturbances	known probabilistic distribution	ambiguity set of distributions
Objective	worst-case cost	expected cost	worst-case expected cost
Constraint	hard constraint	chance constraint	worst-case chance constraint
Conservatism	high	low	medium to high

Remarks:

- In these three formulations, the nominal cost, $V_N(x, \boldsymbol{\mu}) = J_N(x, \boldsymbol{\mu}, \mathbf{0})$, is frequently used for simplicity and other practical considerations.
- Solving these optimization problems can be computationally challenging.
- Scenario-based and convex approximation methods are often employed.
- Advances in parameterizing state and control predictions have been widely adopted.
- Tube MPC has emerged as a leading paradigm.

Tube MPC

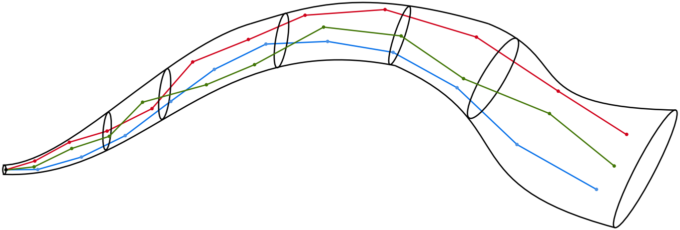

For robust MPC, the state (control) tube is a sequence of sets of all possible states (controls).

Tube MPC

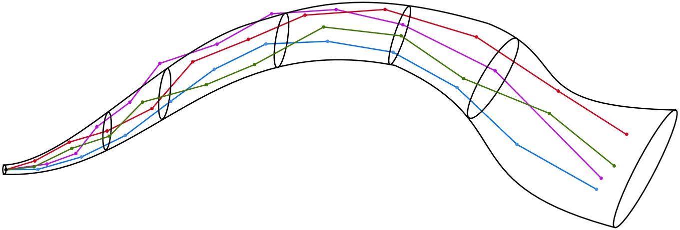

For stochastic MPC and distributionally robust MPC, the state (and/or control) tube is a sequence of sets of states (and/or controls), which stay within the tube with a probabilistic guarantee.

Tube MPC

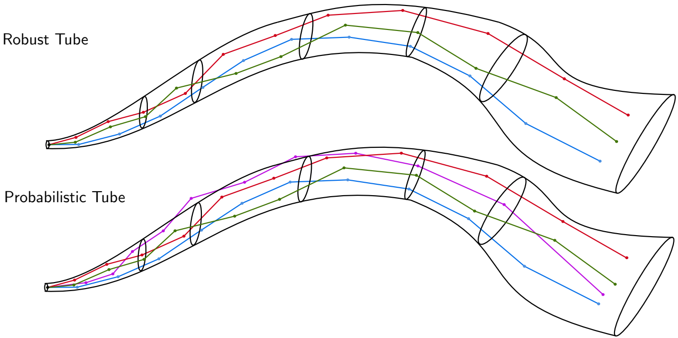

Ideas:

- State and control tubes replace traditional state and control sequences.
- Tubes are induced from the dynamics, uncertainty and employed control policy.
- Optimal tube is obtained via solving tube-based optimal control.
- Parameterization of tubes and control policy is of major importance.

Tube-Based Robust MPC

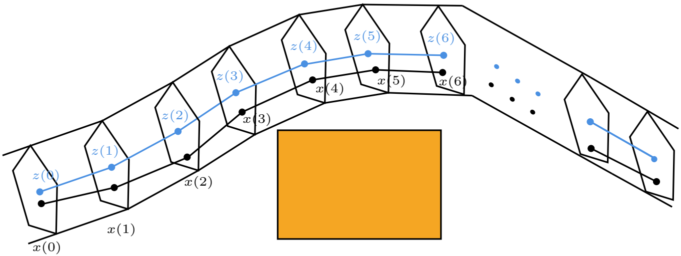

Settings:

- **Uncertain control system:** $x^+ = Ax + Bu + w$.
- **Constraints:** $w \in \mathcal{W}$ (bounded) and $(x, u) \in \mathcal{X} \times \mathcal{U}$.
- **State decomposition:** $x = z + s$, where

$$z^+ = Az + Bv \text{ and}$$
$$s^+ = (A + BK)s + w.$$

- **Affine control policy:** $u_k = v_k + K(x_k - s_k), \ \forall k \in \mathbb{N}_{N-1}$.

Tube-Based Robust MPC

$x = z + s$

$z^+ = Az + Bv$

$s^+ = (A + BK)s + w$

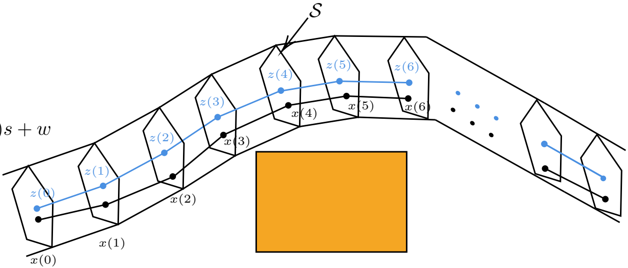

Key Ideas:

- The set $\mathcal{S}$ is robust positively invariant, i.e.,
  $\forall s \in \mathcal{S} \Rightarrow s^+ = (A + BK)s + w \in \mathcal{S}$, $\forall w \in \mathcal{W}$.
- State and control tubes are given by $z(t) \oplus \mathcal{S}$ and $v(t) \oplus K\mathcal{S}$, for all $t \in \mathbb{N}$.
- The selection of optimal tube reduces to the selection of the tube centers $z(t)$, $t \in \mathbb{N}$.

Tube-Based Robust MPC

Other Tube MPC Schemes:

Rigid tube MPC:

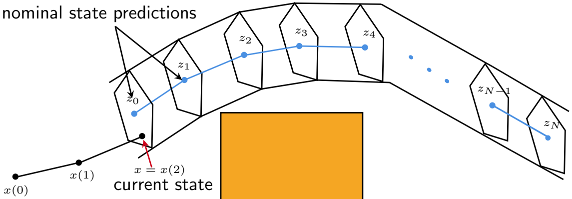

Time-Varying Tube MPC:

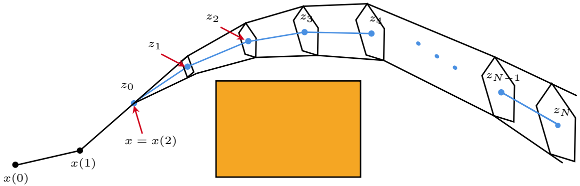

Tube-Based Robust MPC

Other Tube MPC Schemes:

Tube MPC with Time-Varying Cross-Sections:

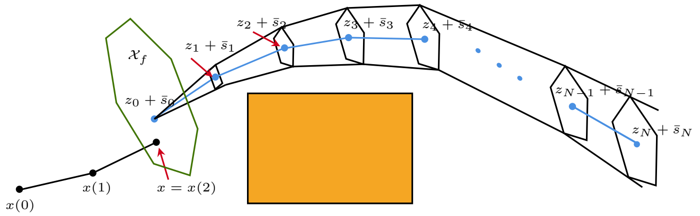

Remarks:

- Tube MPC is computationally simple and has guaranteed control-theoretic properties.
- Stochastic and Distributionally Robust MPC can be addressed similarly using the Tube MPC framework.

Closing Remarks

- We have discussed three uncertainty models:
  (set-membership, stochastic and distributionally robust uncertainty.)
- We have discussed four types of constrains:
  (hard, expectation, chance and worse-cast chance constrains.)
- We have discussed four types of cost functions:
  (nominal, worst-case, expectation, worse-cast expectation costs.)
- We have discussed three types of MPC under uncertainty:
  (robust, stochastic and distributionally robust MPC.)
- Tube MPC is a natural and leading framework to solve MPC under uncertainty.

Recommended Reading

1. Foundational Concepts and In-Depth Understanding of MPC
   - Rawlings et al. "Model Predictive Control: Theory and Design (2nd Edition)." (MPC bible)
2. Robust MPC
   - Raković. "Robust Model Predictive Control," in *Encyclopedia of systems and control*, 2021.
   - Houska and Villanueva. "Robust Optimization for MPC," *Book chapter*, 2019
3. Stochastic MPC
   - Kouvaritakis. "Model Predictive Control: Classical, Robust and Stochastic," *Book*, 2016.
   - Mesbah. "Stochastic model predictive control: an overview and perspectives for future research," *IEEE Control Systems Magazine*, 2016.
4. Distributionally MPC
   - Wu et al. "Ambiguity Tube MPC," *Automatica*, 2022.
   - Van Parys. "Distributionally robust control of constrained stochastic systems," *TAC*, 2016.
5. Tube MPC
   - Mayne et al. "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, 2005.
   - Wang et al. "Tube MPC with time-varying cross-sections," *TAC*, 2024.

As always, follow-up discussions are welcome!

Contact: `k.wang@lboro.ac.uk`