Control design via Bayesian Optimization with safety constraints

Xavier Bombois (CNRS, France),
Marco Forgione (IDSIA, Lugano, Switzerland)

CT Identif, 19 January 2023

Introduction

We wish to determine a parametrized controller $K(\theta)$ for an unknown true system $\mathcal{S}$

Control design: determine $K(\theta)$ minimizing a certain objective function $V(\theta)$ for the loop $[K(\theta)\ \mathcal{S}]$

Since the true system $\mathcal{S}$ is unknown, $V(\theta)$ is unknown

Model-based control design: a model of $\mathcal{S}$ is used to approximate $V(\theta)$

Experiment-based control design: a model of $V(\theta)$ is learnt by performing multiple experiments on $\mathcal{S}$ with different $K(\theta)$

Introduction

We wish to determine a parametrized controller $K(\theta)$ for an unknown true system $\mathcal{S}$

Control design: determine $K(\theta)$ minimizing a certain objective function $V(\theta)$ for the loop $[K(\theta)\ \mathcal{S}]$

Since the true system $\mathcal{S}$ is unknown, $V(\theta)$ is unknown

Model-based control design: a model of $\mathcal{S}$ is used to approximate $V(\theta)$

Experiment-based control design: a model of $V(\theta)$ is learnt by performing multiple experiments on $\mathcal{S}$ with different $K(\theta)$

Introduction

We wish to determine a parametrized controller $K(\theta)$ for an unknown true system $\mathcal{S}$

Control design: determine $K(\theta)$ minimizing a certain objective function $V(\theta)$ for the loop $[K(\theta)\ \mathcal{S}]$

Since the true system $\mathcal{S}$ is unknown, $V(\theta)$ is unknown

Model-based control design: a model of $\mathcal{S}$ is used to approximate $V(\theta)$

Experiment-based control design: a model of $V(\theta)$ is learnt by performing multiple experiments on $\mathcal{S}$ with different $K(\theta)$

Introduction

Experiment-based control design: a model of $V(\theta)$ is learnt by performing multiple experiments on $\mathcal{S}$ with different $K(\theta)$

An experiment with $K(\theta)$ leads to a measure $\tilde{V}(\theta)$ of $V(\theta)$

From the data $(\theta_k, \tilde{V}(\theta_k))$ $(k = 1, ..., n_{it})$, a model of $V(\theta)$ can be determined/identified using e.g., a Gaussian Process

The successive experiments with θ_k $(k = 1, ..., n_{it})$ are generally determined via global optimization algorithms such as Bayesian Optimization

Introduction

Some of these multiple experiments may however pose a safety issue (e.g., instability)

The issue of safe experimentation is generally addressed via regularity assumptions on $V(\theta)$

Here we propose an alternative approach based on a nominal model allowing to determine a safety zone for experimentation

Introduction

Some of these multiple experiments may however pose a safety issue (e.g., instability)

The issue of safe experimentation is generally addressed via regularity assumptions on $V(\theta)$

Here we propose an alternative approach based on a nominal model allowing to determine a safety zone for experimentation

Classical Bayesian Optimization (without safety constraints)

Bayesian optimization minimizes $V(\theta)$ via a number n_{it} of iterations

At iteration i,

1. **Execute** experiment with $K(\theta_i)$, measure $\tilde{V}(\theta_i)$
2. **Compute** the Gaussian Process model $\mathcal{V}_i(\theta)$ of $V(\theta)$ (and its uncertainty) using $(\theta_k, \tilde{V}(\theta_k))$ $(k = 1, .., i)$
3. **Construct** acquisition function $A_i(\theta)$ using $\mathcal{V}_i(\theta)$ and its uncertainty
4. **Maximize** $A_i(\theta)$ to obtain next query point θ_{i+1}

θ_{i+1} is generally a point where $\mathcal{V}_i(\theta)$ is small and/or where the uncertainty of $\mathcal{V}_i(\theta)$ is large (i.e., in regions of the θ-space where few θ has been tested).

After n_{it} iterations, θ_{opt} is then chosen as the tested θ_i leading to the smallest value of the final model $\mathcal{V}_{n_{it}}(\theta)$ (or the smallest $\tilde{V}(\theta_i)$)

Classical Bayesian Optimization (without safety constraints)

Bayesian optimization minimizes $V(\theta)$ via a number n_{it} of iterations

At iteration i,

1. **Execute** experiment with $K(\theta_i)$, measure $\tilde{V}(\theta_i)$
2. **Compute** the Gaussian Process model $\mathcal{V}_i(\theta)$ of $V(\theta)$ (and its uncertainty) using $(\theta_k, \tilde{V}(\theta_k))$ $(k = 1, .., i)$
3. **Construct** acquisition function $A_i(\theta)$ using $\mathcal{V}_i(\theta)$ and its uncertainty
4. **Maximize** $A_i(\theta)$ to obtain next query point θ_{i+1}

θ_{i+1} is generally a point where $\mathcal{V}_i(\theta)$ is small and/or where the uncertainty of $\mathcal{V}_i(\theta)$ is large (i.e., in regions of the θ-space where few θ has been tested).

After n_{it} iterations, θ_{opt} is then chosen as the tested θ_i leading to the smallest value of the final model $\mathcal{V}_{n_{it}}(\theta)$ (or the smallest $\tilde{V}(\theta_i)$)

Classical Bayesian Optimization (without safety constraints)

Bayesian optimization minimizes $V(\theta)$ via a number n_{it} of iterations

At iteration i,

1. **Execute** experiment with $K(\theta_i)$, measure $\tilde{V}(\theta_i)$
2. **Compute** the Gaussian Process model $\mathcal{V}_i(\theta)$ of $V(\theta)$ (and its uncertainty) using $(\theta_k, \tilde{V}(\theta_k))$ $(k = 1, .., i)$
3. **Construct** acquisition function $A_i(\theta)$ using $\mathcal{V}_i(\theta)$ and its uncertainty
4. **Maximize** $A_i(\theta)$ to obtain next query point θ_{i+1}

θ_{i+1} is generally a point where $\mathcal{V}_i(\theta)$ is small and/or where the uncertainty of $\mathcal{V}_i(\theta)$ is large (i.e., in regions of the θ-space where few θ has been tested).

After n_{it} iterations, θ_{opt} is then chosen as the tested θ_i leading to the smallest value of the final model $\mathcal{V}_{n_{it}}(\theta)$ (or the smallest $\tilde{V}(\theta_i)$)

Classical Bayesian Optimization: example

iteration 4

$\mathcal{V}_4(\theta)$ and uncertainty

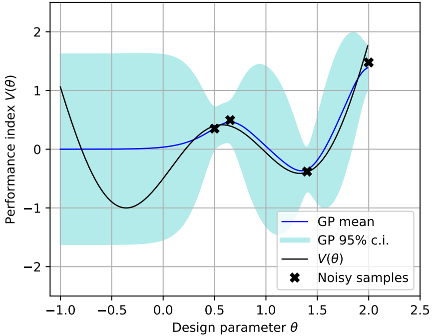

$$\Longrightarrow \theta_5 = -0.1$$

Classical Bayesian Optimization: example

iteration 6

$\mathcal{V}_6(\theta)$ and uncertainty

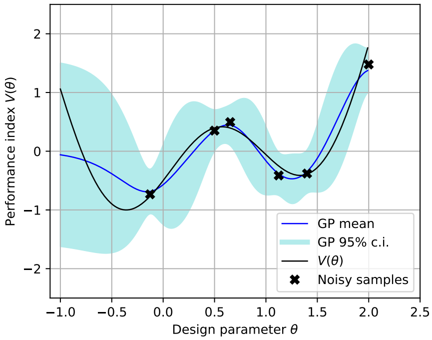

$\Longrightarrow \theta_7 = -0.4$

Imposing safety constraints

The BO algorithm may test $K(\theta)$ that pose safety concerns

How to avoid that?

We will here use a nominal model M of the unknown dynamics $\mathcal{S}$ to determine a safety zone Θ_{safe}

and we will ensure that only $\theta \in \Theta_{safe}$ will be effectively tested on $\mathcal{S}$

To define Θ_{safe}, we use the rationale:

If $K(\theta)$ gives poor performance on M, then it should not be applied on $\mathcal{S}$

Imposing safety constraints

The BO algorithm may test $K(\theta)$ that pose safety concerns

How to avoid that?

We will here use a nominal model M of the unknown dynamics $\mathcal{S}$ to determine a safety zone Θ_{safe}

and we will ensure that only $\theta \in \Theta_{safe}$ will be effectively tested on $\mathcal{S}$

To define Θ_{safe}, we use the rationale:

If $K(\theta)$ gives poor performance on M, then it should not be applied on $\mathcal{S}$

Imposing safety constraints

The BO algorithm may test $K(\theta)$ that pose safety concerns

How to avoid that?

We will here use a nominal model M of the unknown dynamics $\mathcal{S}$ to determine a safety zone Θ_{safe}

and we will ensure that only $\theta \in \Theta_{safe}$ will be effectively tested on $\mathcal{S}$

To define Θ_{safe}, we use the rationale:

If $K(\theta)$ gives poor performance on M, then it should not be applied on $\mathcal{S}$

Safety zone Θ_{safe}

If $K(\theta)$ gives poor performance on M, then it should not be applied on $\mathcal{S}$

$\Longrightarrow$

$$\Theta_{safe} = \{ \; \theta \; | \; V_M(\theta) < V_{safe} \; \}$$

with $V_M(\theta)$ the performance measure corresponding to $V(\theta)$, but for the loop $[K(\theta) \; M]$ and V_{safe} a scalar defining an acceptable $V_M(\theta)$

Bayesian Optimization with safety constraints

We wish to minimize $V(\theta)$, but by testing only $K(\theta)$ with

$$\theta \in \Theta_{safe} = \{ \ \theta \mid V_M(\theta) < V_{safe} \ \}$$

Modified optimization problem:

$$\arg\min_{\theta \in \Theta} J(\theta) \qquad J(\theta) = \begin{cases} V(\theta) & \text{if } \ V_M(\theta) < V_{safe} \\ V_M(\theta) & \text{if } \ V_M(\theta) \geq V_{safe} \end{cases}$$

If Bayesian Optimization is used to tackle this modified optimization problem, the GP will model $J(\theta)$ instead of $V(\theta)$.

Bayesian Optimization with safety constraints

We wish to minimize $V(\theta)$, but by testing only $K(\theta)$ with

$$\theta \in \Theta_{safe} = \{ \ \theta \mid V_M(\theta) < V_{safe} \ \}$$

Modified optimization problem:

$$\arg\min_{\theta \in \Theta} J(\theta) \qquad J(\theta) = \begin{cases} V(\theta) & \text{if } \ V_M(\theta) < V_{safe} \\ V_M(\theta) & \text{if } \ V_M(\theta) \geq V_{safe} \end{cases}$$

If Bayesian Optimization is used to tackle this modified optimization problem, the GP will model $J(\theta)$ instead of $V(\theta)$.

Bayesian Optimization with safety constraints

$$\arg\min_{\theta\in\Theta} J(\theta) \qquad J(\theta) = \begin{cases} V(\theta) \text{ if } V_M(\theta) < V_{safe} \\ V_M(\theta) \text{ if } V_M(\theta) \geq V_{safe} \end{cases}$$

The GP will model $J(\theta)$ instead of $V(\theta)$.

For this purpose, at iteration i, $(\theta_i, \tilde{J}(\theta_i))$ will be collected with

$$\tilde{J}(\theta_i) = \begin{cases} \tilde{V}(\theta_i) \text{ if } V_M(\theta_i) < V_{safe} \\ V_M(\theta_i) \text{ if } V_M(\theta_i) \geq V_{safe} \end{cases}$$

If the to-be-tested $\theta_i \notin \Theta_{safe}$, the *unsafe* experiment is replaced by the computation of $V_M(\theta_i)$

Bayesian Optimization with safety constraints

In order to stress even more the safety constraint, we will in fact consider the following constrained BO optimization problem:

$$\arg\min_{\theta\in\Theta} J(\theta)$$
$$\text{s.t. } V_M(\theta) < V_{safe}.$$

which can be solved via a similar procedure (the constraint is also modeled as a GP)

Linear simulation example

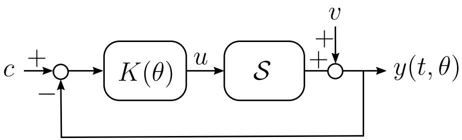

True system $\mathcal{S}$ described by $P(s) = \frac{10}{(s+10)(s+1)}$ with white noise v ($\sigma_v = 0.01$)

To be controlled by PID $K(s,\theta) = \frac{k(s+p_1)(s+p_2)}{s(s+4.2)}$ ($\theta = (k, p_1, p_2)$)

Model reference objective: $y(t,\theta)$ for unit step $c(t)$ close to:

$$y_{des}(t) = \frac{9}{s^2 + 4.2s + 9}\, c(t) \qquad (t_r = 1\ s \text{ and } D = 5\%)$$

Initial model $M(s) = \frac{12}{(s+8)(s+2)}$

Linear simulation example: $\tilde{V}(\theta)$ and Θ_{safe}

Experiment of 5 seconds on $[K(s,\theta)\ P(s)]$ with unit step c ($T_s = 0.025\ s$)

$$\Longrightarrow \tilde{V}(\theta) = T_s \sum_{n=1}^{200} (y(nT_s,\theta) - y_{des}(nT_s))^2$$

which is a measure of $V(\theta) = E\tilde{V}(\theta)$

Safety zone $\Theta_{safe} = \{\ \theta \mid V_M(\theta) < V_{safe}\ \}$ defined with $V_{safe} = 0.1$ and

$$V_M(\theta) = \begin{cases} T_s \sum_{n=1}^{200} (y_M(nT_s,\theta) - y_{des}(nT_s))^2 & \text{if } [K(s,\theta)\ M(s)] \text{ stable} \\ 100\ V_{safe} & \text{otherwise.} \end{cases}$$

where $y_M(t,\theta) = \frac{M(s)K(s,\theta)}{1+M(s)K(s,\theta)} c(t)$

Linear simulation example: Bayesian Optimization with safety constraints

$$\arg\min_{\theta\in\Theta} J(\theta) \qquad J(\theta) = \begin{cases} V(\theta) \text{ if } V_M(\theta) < 0.1 \\ V_M(\theta) \text{ if } V_M(\theta) \geq 0.1 \end{cases}$$

$$\text{s.t. } V_M(\theta) < 0.1.$$

We tackle this optimization problem by running 300 iterations of Bayesian Optimization initialized with $\theta_{\text{init}} = (0.75, 8, 2)$

$K(\theta_{\text{init}})$ achieves $y_M(t, \theta_{\text{init}}) = y_{des}(t)$ i.e., $V_M(\theta_{\text{init}}) = 0$

Linear simulation example: results

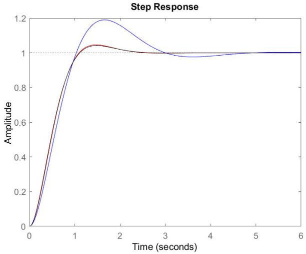

- $y(t, \theta_{init})$ (blue) far from $y_{des}(t)$ (red)
- Optimized output $y(t, \theta_{opt})$ (black) almost equal to $y_{des}(t)$

From the 300 values of θ selected by the BO algorithm, 10% were outside Θ_{safe} and were thus not tested on $\mathcal{S}$ $\Longrightarrow$

- Experiments with unstable controllers avoided
- Average $\tilde{V}(\theta)$ for non-tested θ would have been 400x larger than for the tested θ

Linear simulation example: results

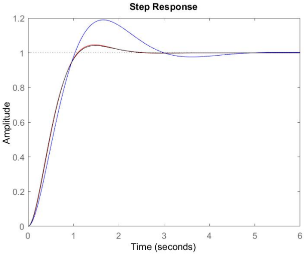

- $y(t, \theta_{init})$ (blue) far from $y_{des}(t)$ (red)
- Optimized output $y(t, \theta_{opt})$ (black) almost equal to $y_{des}(t)$

From the 300 values of θ selected by the BO algorithm, 10% were outside Θ_{safe} and were thus not tested on $\mathcal{S} \Longrightarrow$

- Experiments with unstable controllers avoided
- Average $\tilde{V}(\theta)$ for non-tested θ would have been 400x larger than for the tested θ

Non-linear simulation example

The only difference is the true system $\mathcal{S}$ which is now an Hammerstein system:

$$y(t) = P(s)\ sat(u(t))$$

with the same $P(s) = \frac{10}{(s+10)(s+1)}$ and

$$sat(u(t)) = \begin{cases} u(t) & \text{if } u(t) \in [-1.5\ 1.5] \\ 1.5 & \text{if } u(t) > 1.5 \\ -1.5 & \text{if } u(t) < -1.5 \end{cases}$$

$\tilde{V}(\theta)$ will therefore be determined with an experiment on $[K(s,\theta)\ \mathcal{S}]$

Safe BO algorithm run for 300 iterations

Non-linear simulation example: results

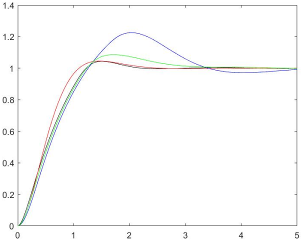

- $y(t, \theta_{init})$ (blue) far from $y_{des}(t)$ (red)
- $y(t, \theta_{opt})$ (black) close to $y_{des}(t)$
- $y(t, \theta_{opt})$ better than with ideal linear design using $P(s)$ i.e., ignoring saturation (green)

From the 300 values of θ selected by the BO algorithm, 13% were outside Θ_{safe} and were thus not tested on $\mathcal{S}$ $\Longrightarrow$

- Experiments with controllers yielding a limit cycle avoided
- Average $\tilde{V}(\theta)$ for non-tested θ would have been 9x larger than for the tested θ

Non-linear simulation example: results

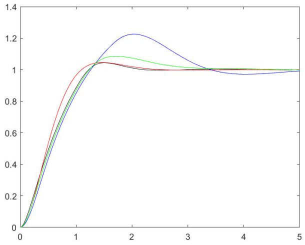

- $y(t, \theta_{init})$ (blue) far from $y_{des}(t)$ (red)
- $y(t, \theta_{opt})$ (black) close to $y_{des}(t)$
- $y(t, \theta_{opt})$ better than with ideal linear design using $P(s)$ i.e., ignoring saturation (green)

From the 300 values of θ selected by the BO algorithm, 13% were outside Θ_{safe} and were thus not tested on $\mathcal{S}$ $\Longrightarrow$

- Experiments with controllers yielding a limit cycle avoided
- Average $\tilde{V}(\theta)$ for non-tested θ would have been 9x larger than for the tested θ

Conclusions

Safety constraint added to an experiment-driven controller calibration approach to guarantee safe exploration

Safety constraint based on a nominal model knowledge

Current/future work

- Adaptation of the model during the iterative process

More details at *https://hal.archives-ouvertes.fr/hal-03559979*

Conclusions

Safety constraint added to an experiment-driven controller calibration approach to guarantee safe exploration

Safety constraint based on a nominal model knowledge

Current/future work

- Adaptation of the model during the iterative process

More details at *https://hal.archives-ouvertes.fr/hal-03559979*

2023 Spring School with M. Schoukens as guest lecturer

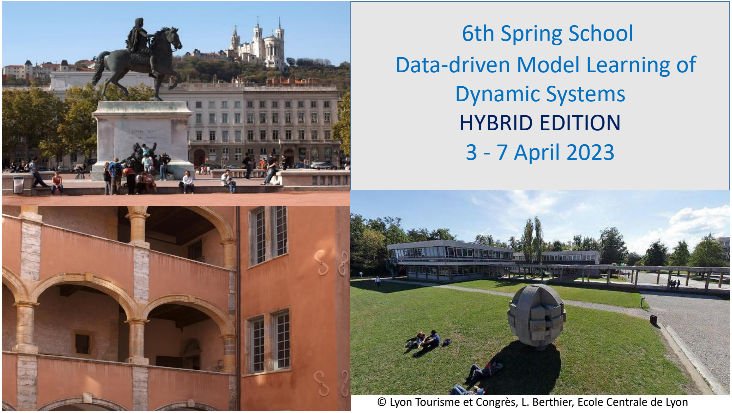

© Lyon Tourisme et Congrès, L. Berthier, Ecole Centrale de Lyon

More details at *https://spring-id-2023.sciencesconf.org/*