Aircraft trajectory optimization under unknown dynamics

C. Rommel[1,2], J. F. Bonnans[1],
B. Gregorutti[2] and P. Martinon[1]

CMAP Ecole Polytechnique - INRIA[1]
Safety Line[2]

PGMODays - November 21st 2018
Optimal control and applications session

Motivation - Optimal Control

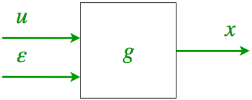

$$\dot{x}(t) = g(\boldsymbol{u}(t), \boldsymbol{x}(t)) + \varepsilon(t)$$

Motivation - Optimal Control

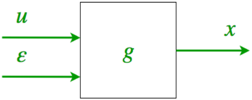

$$\dot{x}(t) = g(\boldsymbol{u}(t), \boldsymbol{x}(t)) + \varepsilon(t)$$

Optimal Control Problem

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t), \boldsymbol{x}(t))dt,$$

$$\text{s.t.} \begin{cases} \dot{\boldsymbol{x}}(t) = g(\boldsymbol{u}(t), \boldsymbol{x}(t)) + \varepsilon(t), & \text{for a.e. } t \in [0, t_f], \\ \text{Other constraints...} \end{cases} \tag{OCP}$$

Motivation - Optimal Control

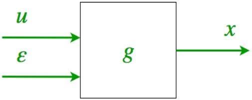

$$\dot{x}(t) = g(\boldsymbol{u}(t), \boldsymbol{x}(t)) + \varepsilon(t)$$

Optimal Control Problem

$$\min_{(\boldsymbol{x},\boldsymbol{u}) \in \mathbb{X} \times \mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t), \boldsymbol{x}(t)) dt,$$

$$\text{s.t.} \begin{cases} \dot{\boldsymbol{x}}(t) = g(\boldsymbol{u}(t), \boldsymbol{x}(t)) + \varepsilon(t), & \text{for a.e. } t \in [0, t_f], \\ \text{Other constraints...} \end{cases} \tag{OCP}$$

Motivation - Optimal Control

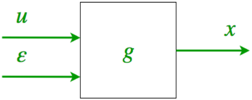

$$\dot{x}(t) = g(\boldsymbol{u}(t), \boldsymbol{x}(t)) + \varepsilon(t)$$

Optimal Control Problem

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t), \boldsymbol{x}(t))dt, \tag{OCP}$$

$$\text{s.t.} \quad \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t), \boldsymbol{x}(t)), & \text{for a.e. } t \in [0, t_f], \\ \text{Other constraints...} \end{cases}$$

Motivation - Optimal Control

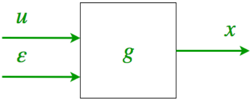

$$\dot{x}(t) = g(\boldsymbol{u}(t), \boldsymbol{x}(t)) + \varepsilon(t)$$

Optimal Control Problem

$$\min_{(\boldsymbol{x},\boldsymbol{u}) \in \mathbb{X} \times \mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t), \boldsymbol{x}(t)) dt, \tag{OCP}$$
$$\text{s.t.} \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t), \boldsymbol{x}(t)), & \text{for a.e. } t \in [0, t_f], \\ \text{Other constraints...} \end{cases}$$

Use of past data to learn how to control a system efficiently

Motivation - Optimal Control

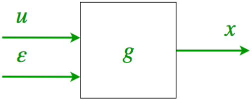

$$\dot{x}(t) = g(\boldsymbol{u}(t), \boldsymbol{x}(t)) + \varepsilon(t)$$

Optimal Control Problem

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t), \boldsymbol{x}(t))dt,$$

$$\text{s.t.} \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t), \boldsymbol{x}(t)), \quad \text{for a.e. } t \in [0, t_f], \\ \text{Other constraints...} \end{cases} \tag{OCP}$$

Use of past data to learn how to control a system efficiently

"Model-based reinforcement learning" - [Recht, 2018]

Flight optimization

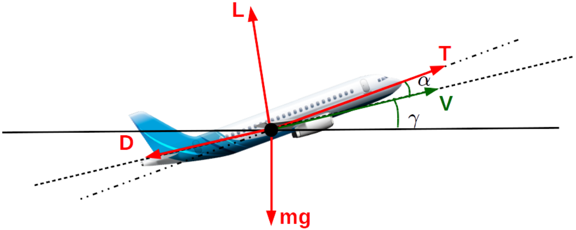

Flight optimization

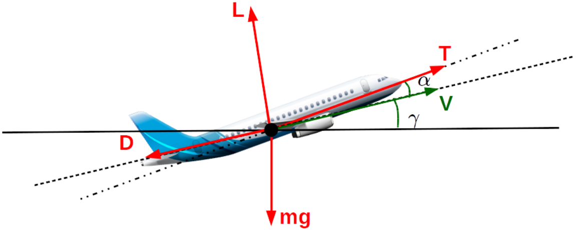

CO_2

Dynamics are learned from QAR data

Black box

Dynamics are learned from QAR data

Black box

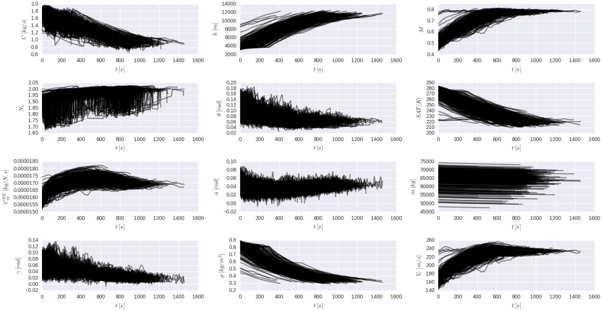

Recorded flights = functional data

Trajectory acceptability

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t),\boldsymbol{x}(t))dt, \qquad \text{(AOCP)}$$

$$\text{s.t.} \quad \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t),\boldsymbol{x}(t)), \quad \text{a.e. } t\in[0,t_f], \\ \text{Other constraints...} \end{cases}$$

Trajectory acceptability

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t),\boldsymbol{x}(t))dt, \qquad \text{(AOCP)}$$
$$\text{s.t.} \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t),\boldsymbol{x}(t)), & \text{a.e. } t \in [0,t_f], \\ \text{Other constraints...} \end{cases}$$

$$\Rightarrow \hat{\boldsymbol{z}} = (\hat{\boldsymbol{x}},\hat{\boldsymbol{u}}) \text{ solution of (AOCP).}$$

Trajectory acceptability

$$
\begin{aligned}
&\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t),\boldsymbol{x}(t))dt, \\
&\text{s.t.} \quad \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t),\boldsymbol{x}(t)), \quad \text{a.e. } t\in[0,t_f], \\ \text{Other constraints...} \end{cases}
\end{aligned}
\tag{AOCP}
$$

$$\Rightarrow \hat{\boldsymbol{z}} = (\hat{\boldsymbol{x}}, \hat{\boldsymbol{u}}) \text{ solution of (AOCP).}$$

- Is $\hat{\boldsymbol{z}}$ inside the validity region of the dynamics model $\hat{g}$?

Trajectory acceptability

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t),\boldsymbol{x}(t))dt, \qquad \text{(AOCP)}$$

$$\text{s.t.} \quad \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t),\boldsymbol{x}(t)), \quad \text{a.e. } t \in [0,t_f], \\ \text{Other constraints...} \end{cases}$$

$$\Rightarrow \hat{\boldsymbol{z}} = (\hat{\boldsymbol{x}}, \hat{\boldsymbol{u}}) \text{ solution of (AOCP).}$$

- Is $\hat{\boldsymbol{z}}$ inside the validity region of the dynamics model $\hat{g}$?
- Does it look like a real trajectory ?

Trajectory acceptability

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t),\boldsymbol{x}(t))dt, \qquad \text{(AOCP)}$$

$$\text{s.t.} \quad \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t),\boldsymbol{x}(t)), & \text{a.e. } t \in [0,t_f], \\ \text{Other constraints...} \end{cases}$$

$$\Rightarrow \hat{\boldsymbol{z}} = (\hat{\boldsymbol{x}},\hat{\boldsymbol{u}}) \text{ solution of (AOCP).}$$

- Is $\hat{\boldsymbol{z}}$ inside the validity region of the dynamics model $\hat{g}$?
- Does it look like a real trajectory ?

Pilots acceptance

Air Traffic Control[1]

Trajectory acceptability

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t),\boldsymbol{x}(t))dt, \qquad \text{(AOCP)}$$

$$\text{s.t.} \quad \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t),\boldsymbol{x}(t)), & \text{a.e. } t \in [0, t_f], \\ \text{Other constraints...} \end{cases}$$

$$\Rightarrow \hat{\boldsymbol{z}} = (\hat{\boldsymbol{x}}, \hat{\boldsymbol{u}}) \text{ solution of (AOCP).}$$

- Is $\hat{\boldsymbol{z}}$ inside the validity region of the dynamics model $\hat{g}$?
- Does it look like a real trajectory ?

Pilots acceptance

Air Traffic Control[1]

How can we quantify the closeness from the optimized trajectory to the set of real flights?

Optimized trajectory likelihood

Assumption: We suppose that the real flights are observations of the same functional random variable $Z = (Z_t)$ valued in $\mathcal{C}(\mathbb{T}, E)$, with E compact subset of $\mathbb{R}^d$ and $\mathbb{T} = [0, t_f]$.

How likely is it to draw the optimized trajectory from the law of Z ?

How to apply this to functional data?

Problem: Computation of probability densities in infinite dimensional space.

How to apply this to functional data?

Problem: Computation of probability densities in infinite dimensional space.

- Standard approach in Functional Data Analysis: use Functional Principal Component Analysis to decompose the data in a small number of coefficients

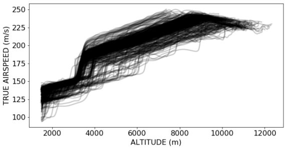

$\Rightarrow$

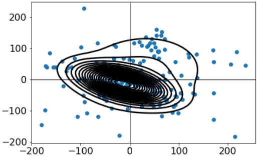

How to apply this to functional data?

Problem: Computation of probability densities in infinite dimensional space.

- Standard approach in Functional Data Analysis: use Functional Principal Component Analysis to decompose the data in a small number of coefficients
- **Or: we can use the marginal densities**

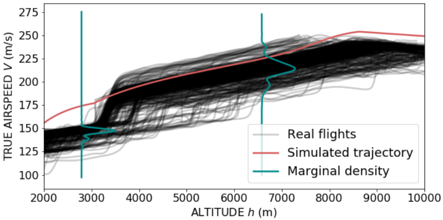

How do we aggregate the marginal likelihoods?

- f_t marginal density of Z, i.e. probability density function of Z_t,
- $\boldsymbol{y}$ new trajectory,
- $f_t(\boldsymbol{y}(t))$ marginal likelihood of $\boldsymbol{y}$ at t, i.e. likelihood of observing $Z_t = \boldsymbol{y}(t)$.

How do we aggregate the marginal likelihoods?

- f_t marginal density of Z, i.e. probability density function of Z_t,
- $\boldsymbol{y}$ new trajectory,
- $f_t(\boldsymbol{y}(t))$ marginal likelihood of $\boldsymbol{y}$ at t, i.e. likelihood of observing $Z_t = \boldsymbol{y}(t)$.

Mean marginal likelihood

$$\boxed{\mathrm{MML}(Z, \boldsymbol{y}) = \frac{1}{t_f} \int_0^{t_f} \psi[f_t, \boldsymbol{y}(t)] dt,}$$

where $\psi : L^1(E, \mathbb{R}_+) \times \mathbb{R} \to [0; 1]$ is a continuous scaling map,

How do we aggregate the marginal likelihoods?

- f_t marginal density of Z, i.e. probability density function of Z_t,
- $\boldsymbol{y}$ new trajectory,
- $f_t(\boldsymbol{y}(t))$ marginal likelihood of $\boldsymbol{y}$ at t, i.e. likelihood of observing $Z_t = \boldsymbol{y}(t)$.

Mean marginal likelihood

$$\boxed{\mathrm{MML}(Z, \boldsymbol{y}) = \frac{1}{t_f} \int_0^{t_f} \psi[f_t, \boldsymbol{y}(t)]dt,}$$

where $\psi : L^1(E, \mathbb{R}_+) \times \mathbb{R} \to [0; 1]$ is a continuous scaling map, because marginal densities may have really different shapes.

How do we aggregate the marginal likelihoods?

Possible scalings are the normalized density

$$\psi[f_t, \boldsymbol{y}(t)] := \frac{f_t(\boldsymbol{y}(t))}{\max\limits_{z \in E} f_t(z)},$$

How do we aggregate the marginal likelihoods?

Possible scalings are the normalized density

$$\psi[f_t, \boldsymbol{y}(t)] := \frac{f_t(\boldsymbol{y}(t))}{\max\limits_{z \in E} f_t(z)},$$

or the confidence level

$$\psi[f_t, \boldsymbol{y}(t)] := \mathbb{P}\left(f_t(Z_t) \leq f_t(\boldsymbol{y}(t))\right).$$

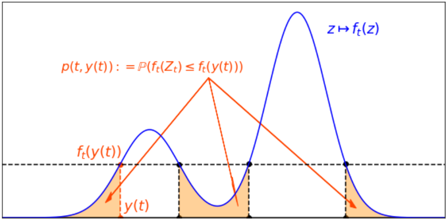

How do we deal with sampled curves?

In practice, the m trajectories are sampled at variable discrete times:

$$\mathcal{T}^D := \{(t_j^r, z_j^r)\}_{\substack{1\leq j\leq n \\ 1\leq r\leq m}} \subset \mathbb{T}\times E, \qquad z_j^r := \boldsymbol{z}^r(t_j^r),$$

$$\mathcal{Y} := \{(\tilde{t}_j, y_j)\}_{j=1}^{\tilde{n}} \subset \mathbb{T}\times E, \qquad y_j := \boldsymbol{y}(\tilde{t}_j).$$

How do we deal with sampled curves?

In practice, the m trajectories are sampled at variable discrete times:

$$\mathcal{T}^D := \{(t_j^r, z_j^r)\}_{\substack{1 \leq j \leq n \\ 1 \leq r \leq m}} \subset \mathbb{T} \times E, \qquad z_j^r := \boldsymbol{z}^r(t_j^r),$$

$$\mathcal{Y} := \{(\tilde{t}_j, y_j)\}_{j=1}^{\tilde{n}} \subset \mathbb{T} \times E, \qquad y_j := \boldsymbol{y}(\tilde{t}_j).$$

Hence, we approximate the MML using a Riemann sum which aggregates consistent estimators $\hat{f}_{\tilde{t}_j}^m$ of the marginal densities $f_{\tilde{t}_j}$:

$$\boxed{\mathrm{EMML}_m(\mathcal{T}^D, \mathcal{Y}) := \frac{1}{t_f} \sum_{j=1}^{\tilde{n}} \psi[\hat{f}_{\tilde{t}_j}^m, y_j] \Delta \tilde{t}_j.}$$

How can we estimate marginal densities?

How can we estimate marginal densities?

- In practice, the altitude plays the role of time, so we can't assume the same sampling for each trajectory;

How can we estimate marginal densities?

- In practice, the altitude plays the role of time, so we can't assume the same sampling for each trajectory;
- Assume sampling times $\{t_j^r : j = 1, \ldots, n; r = 1, \ldots, m\}$ to be i.i.d. observations of a r.v. T, indep. Z;

How can we estimate marginal densities?

- In practice, the altitude plays the role of time, so we can't assume the same sampling for each trajectory;
- Assume sampling times $\{t_j^r : j = 1, \ldots, n; r = 1, \ldots, m\}$ to be i.i.d. observations of a r.v. T, indep. Z;
- Our problem can be seen as a conditional probability density learning problem with $(X, Y) = (T, Z_T)$, where f_t is the density of $Z_t = (Z_T | T = t) = (Y|X)$.

How can we estimate marginal densities?

- In practice, the altitude plays the role of time, so we can't assume the same sampling for each trajectory;
- Assume sampling times $\{t_j^r : j = 1, \ldots, n; r = 1, \ldots, m\}$ to be i.i.d. observations of a r.v. T, indep. Z;
- Our problem can be seen as a conditional probability density learning problem with $(X, Y) = (T, Z_T)$, where f_t is the density of $Z_t = (Z_T | T = t) = (Y|X)$.

1. We can apply SOA conditional density estimation techniques, such as LS-CDE [Sugiyama et al., 2010],

How can we estimate marginal densities?

- In practice, the altitude plays the role of time, so we can't assume the same sampling for each trajectory;
- Assume sampling times $\{t_j^r : j = 1, \ldots, n; r = 1, \ldots, m\}$ to be i.i.d. observations of a r.v. T, indep. Z;
- Our problem can be seen as a conditional probability density learning problem with $(X, Y) = (T, Z_T)$, where f_t is the density of $Z_t = (Z_T | T = t) = (Y|X)$.

1. We can apply SOA conditional density estimation techniques, such as LS-CDE [Sugiyama et al., 2010],
2. **We can use a fine partitioning of the time domain.**

Partition based marginal density estimation

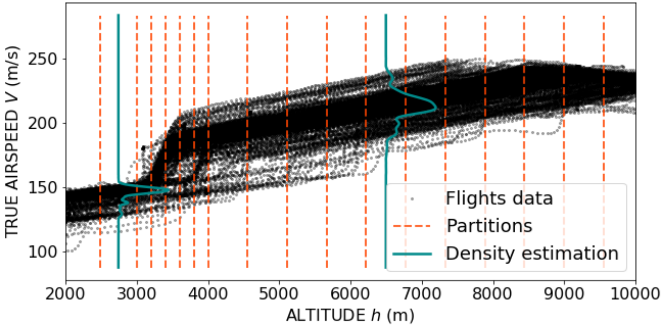

Idea: to average in time the marginal densities over small bins by applying classical multivariate density estimation techniques to each subset.

Consistency

We denote by:

- $\Theta : \mathcal{S} \to L^1(E, \mathbb{R}_+)$ multivariate density estimation statistic,
- $\mathcal{S} = \{(z_k)_{k=1}^N \in E^N : N \in \mathbb{N}^*\}$ set of finite sequences,

Consistency

We denote by:

- $\Theta : \mathcal{S} \to L^1(E, \mathbb{R}_+)$ multivariate density estimation statistic,
- $\mathcal{S} = \{(z_k)_{k=1}^N \in E^N : N \in \mathbb{N}^*\}$ set of finite sequences,
- m the number of random curves;
- $\mathcal{T}_t^m$ subset of data points whose sampling times fall in the bin containing t;

Consistency

We denote by:

- $\Theta : \mathcal{S} \to L^1(E, \mathbb{R}_+)$ multivariate density estimation statistic,
- $\mathcal{S} = \{(z_k)_{k=1}^N \in E^N : N \in \mathbb{N}^*\}$ set of finite sequences,
- m the number of random curves;
- $\mathcal{T}_t^m$ subset of data points whose sampling times fall in the bin containing t;
- $\hat{f}_t^m := \Theta[\mathcal{T}_t^m]$ estimator trained using $\mathcal{T}_t^m$.

Consistency

Assumption 1 - Positive time density

$\nu \in L^\infty(E, \mathbb{R}_+)$ density function of T, s.t.

$$\nu_+ := \operatorname*{ess\,sup}_{t \in \mathbb{T}} \nu(t) < \infty, \qquad \nu_- := \operatorname*{ess\,inf}_{t \in \mathbb{T}} \nu(t) > 0.$$

Consistency

Assumption 1 - Positive time density
$\nu \in L^\infty(E, \mathbb{R}_+)$ density function of T, s.t.

$$\nu_+ := \operatorname*{ess\,sup}_{t \in \mathbb{T}} \nu(t) < \infty, \qquad \nu_- := \operatorname*{ess\,inf}_{t \in \mathbb{T}} \nu(t) > 0.$$

Assumption 2 - Lipschitz in time
Function $(t, z) \in \mathbb{T} \times E \mapsto f_t(z)$ is continuous and

$$|f_{t_1}(z) - f_{t_2}(z)| \leq L|t_1 - t_2|, \qquad L > 0.$$

Consistency

Assumption 1 - Positive time density
$\nu \in L^\infty(E, \mathbb{R}_+)$ density function of T, s.t.

$$\nu_+ := \operatorname*{ess\,sup}_{t\in\mathbb{T}} \nu(t) < \infty, \qquad \nu_- := \operatorname*{ess\,inf}_{t\in\mathbb{T}} \nu(t) > 0.$$

Assumption 2 - Lipschitz in time
Function $(t, z) \in \mathbb{T} \times E \mapsto f_t(z)$ is continuous and

$$|f_{t_1}(z) - f_{t_2}(z)| \leq L|t_1 - t_2|, \qquad L > 0.$$

Assumption 3 - Shrinking bins
The homogeneous partition $\{B_\ell^m\}_{\ell=1}^{q_m}$ of $[0; t_f]$, with binsize b_m, is s.t.

$$\lim_{m\to\infty} b_m = 0, \qquad \lim_{m\to\infty} mb_m = \infty.$$

Consistency

Assumption 4 - i.i.d. consistency

- $\mathcal{G}$ arbitrary family of probability density functions on E, $\rho \in \mathcal{G}$,
- S_ρ^N i.i.d sample of size N drawn from ρ valued in $\mathcal{S}$.

The estimator obtained by applying Θ to S_ρ^N, denoted by

$$\hat{\rho}^N := \Theta[S_\rho^N] \in L^1(E, \mathbb{R}_+),$$

is a (pointwise) consistent density estimator, uniformly in ρ:

For all $z \in E, \varepsilon > 0, \alpha_1 > 0$, there is $N_{\varepsilon,\alpha_1} > 0$ such that, for any $\rho \in \mathcal{G}$,

$$N \geq N_{\varepsilon,\alpha_1} \Rightarrow \mathbb{P}\left(\left|\hat{\rho}^N(z) - \rho(z)\right| < \varepsilon\right) > 1 - \alpha_1.$$

Consistency

Theorem 1
Under assumptions 1 to 4, for any $z \in E$ and $t \in \mathbb{T}$, $\hat{f}^m_{\ell^m(t)}(z)$ consistently approximates the marginal density $f_t(z)$ as the number of curves m grows:

$$\forall \varepsilon > 0, \quad \lim_{m \to \infty} \mathbb{P}\left(|\hat{f}_t^m(z) - f_t(z)| < \varepsilon\right) = 1.$$

Consistency

Theorem 1

Under assumptions 1 to 4, for any $z \in E$ and $t \in \mathbb{T}$, $\hat{f}^m_{\ell^m(t)}(z)$ consistently approximates the marginal density $f_t(z)$ as the number of curves m grows:

$$\forall \varepsilon > 0, \quad \lim_{m \to \infty} \mathbb{P}\left(|\hat{f}_t^m(z) - f_t(z)| < \varepsilon\right) = 1.$$

Note that:

- $m \to \infty \neq N \to \infty$,

Consistency

Theorem 1

Under assumptions 1 to 4, for any $z \in E$ and $t \in \mathbb{T}$, $\hat{f}^{m}_{\ell^m(t)}(z)$ consistently approximates the marginal density $f_t(z)$ as the number of curves m grows:

$$\forall \varepsilon > 0, \quad \lim_{m \to \infty} \mathbb{P}\left(|\hat{f}^{m}_{t}(z) - f_t(z)| < \varepsilon \right) = 1.$$

Note that:

- $m \to \infty \neq N \to \infty$,
- Number of samples $=$ random,

Consistency

Theorem 1

Under assumptions 1 to 4, for any $z \in E$ and $t \in \mathbb{T}$, $\hat{f}^m_{\ell^m(t)}(z)$ consistently approximates the marginal density $f_t(z)$ as the number of curves m grows:

$$\forall \varepsilon > 0, \quad \lim_{m \to \infty} \mathbb{P}\left(|\hat{f}^m_t(z) - f_t(z)| < \varepsilon\right) = 1.$$

Note that:

- $m \to \infty \neq N \to \infty$,
- Number of samples $=$ random,
- Training data not i.i.d.

Marginal density estimation results

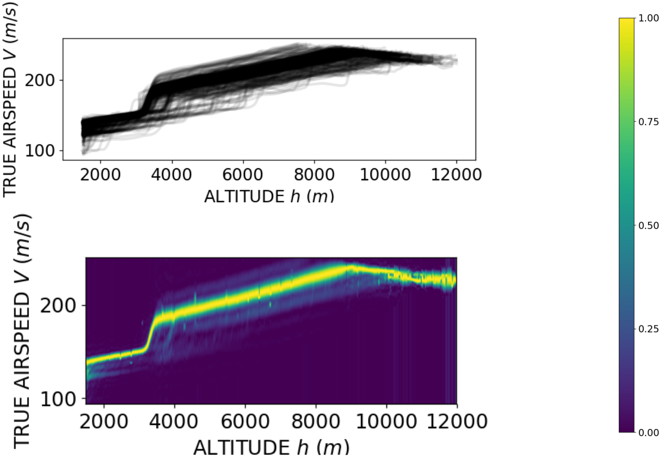

Marginal density estimation results

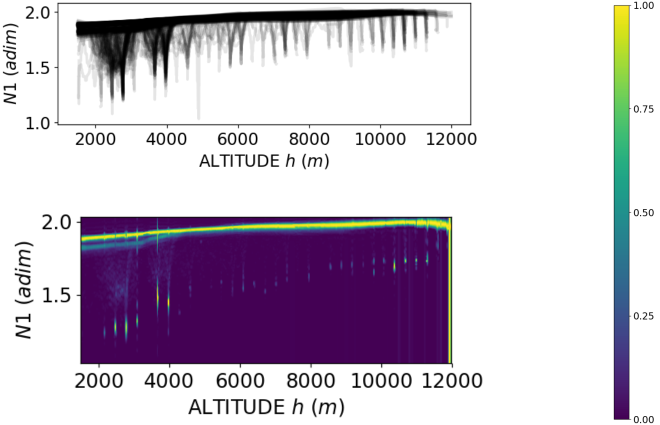

How good is it compared to other methods?

How good is it compared to other methods?

- Training set of $m = 424$ flights $\simeq$ 334 531 point observations,

How good is it compared to other methods?

- Training set of $m = 424$ flights $\simeq 334\,531$ point observations,
- Test set of 150 flights

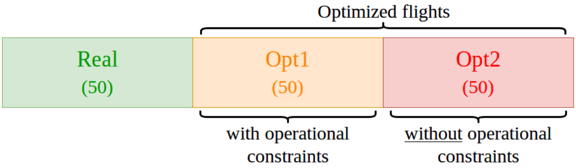

How good is it compared to other methods?

- Training set of $m = 424$ flights $\simeq$ 334 531 point observations,
- Test set of 150 flights

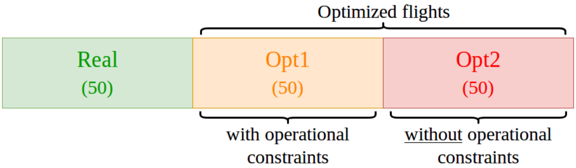

- Discrimination power comparison with (gmm-)FPCA and (integrated) LS-CDE:

Var.	Estimated Likelihoods		
	Real	Opt1	Opt2
MML	$\mathbf{0.63 \pm 0.07}$	$\mathbf{0.43 \pm 0.08}$	$\mathbf{0.13 \pm 0.02}$
FPCA	0.16 ± 0.12	6.4E-03 $\pm$ 3.8E-03	3.6E-03 $\pm$ 5.4E-03
LS-CDE	0.77 ± 0.05	0.68 ± 0.04	0.49 ± 0.06

How good is it compared to other methods?

- Training set of $m = 424$ flights $\simeq$ 334 531 point observations,
- Test set of 150 flights

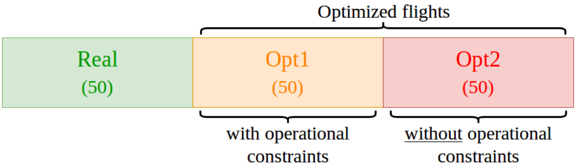

- Discrimination power comparison with (gmm-)FPCA and (integrated) LS-CDE:

Var.	Estimated Likelihoods			Tr. Time
	Real	Opt1	Opt2	
MML	**0.63 ± 0.07**	**0.43 ± 0.08**	**0.13 ± 0.02**	5s
FPCA	0.16 ± 0.12	6.4e-03 ± 3.8e-03	3.6e-03 ± 5.4e-03	20s
LS-CDE	0.77 ± 0.05	0.68 ± 0.04	0.49 ± 0.06	14h

MML penalty

The MML can be used not only to assess the optimization solutions, but also to penalize the optimization itself:

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t),\boldsymbol{x}(t))dt$$

$$\text{s.t.} \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t),\boldsymbol{x}(t)), \quad \text{a.e. } t \in [0, t_f], \\ \text{Other constraints...} \end{cases} \qquad (\qquad \text{AOCP})$$

MML PENALTY

The MML can be used not only to assess the optimization solutions, but also to penalize the optimization itself:

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t),\boldsymbol{x}(t))dt - \lambda\,\mathrm{MML}(Z,\boldsymbol{x}),$$

$$\text{s.t.} \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t),\boldsymbol{x}(t)), & \text{a.e. } t\in[0,t_f], \\ \text{Other constraints...} \end{cases} \tag{MML-AOCP}$$

MML PENALTY

The MML can be used not only to assess the optimization solutions, but also to penalize the optimization itself:

$$\min_{(\boldsymbol{x},\boldsymbol{u})\in\mathbb{X}\times\mathbb{U}} \int_0^{t_f} C(\boldsymbol{u}(t),\boldsymbol{x}(t))dt - \lambda\,\mathrm{MML}(Z,\boldsymbol{x}),$$

$$\text{s.t.} \quad \begin{cases} \dot{\boldsymbol{x}}(t) = \hat{g}(\boldsymbol{u}(t),\boldsymbol{x}(t)), \quad \text{a.e. } t\in[0,t_f], \\ \text{Other constraints...} \end{cases} \tag{MML-AOCP}$$

- λ sets trade-off between a fuel minimization and a likelihood maximization,

Penalty effect

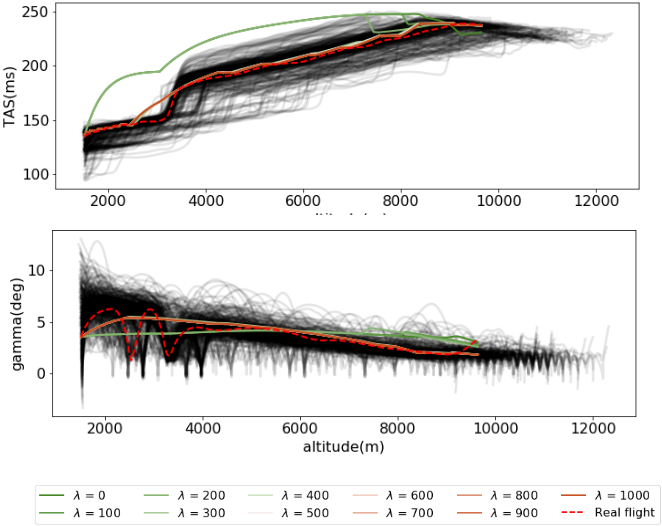

Trajectory acceptability conclusion

1. General probabilistic criterion using marginal densities to quantify the closeness between a curve and a set of random trajectories,

Trajectory acceptability conclusion

1. General probabilistic criterion using marginal densities to quantify the closeness between a curve and a set of random trajectories,
2. Class of consistent plug-in estimators, based on "histogram" of multivariate density estimators,

Trajectory acceptability conclusion

1. General probabilistic criterion using marginal densities to quantify the closeness between a curve and a set of random trajectories,
2. Class of consistent plug-in estimators, based on "histogram" of multivariate density estimators,
3. Applicable to the case of aircraft climb trajectories,

Trajectory acceptability conclusion

1. General probabilistic criterion using marginal densities to quantify the closeness between a curve and a set of random trajectories,
2. Class of consistent plug-in estimators, based on "histogram" of multivariate density estimators,
3. Applicable to the case of aircraft climb trajectories,
   - Competitive with other well-established SOA approaches,

Trajectory acceptability conclusion

1. General probabilistic criterion using marginal densities to quantify the closeness between a curve and a set of random trajectories,
2. Class of consistent plug-in estimators, based on "histogram" of multivariate density estimators,
3. Applicable to the case of aircraft climb trajectories,
   - Competitive with other well-established SOA approaches,
4. Particular Adaptive Kernel and Gaussian mixture implementation,

Trajectory acceptability conclusion

1. General probabilistic criterion using marginal densities to quantify the closeness between a curve and a set of random trajectories,
2. Class of consistent plug-in estimators, based on "histogram" of multivariate density estimators,
3. Applicable to the case of aircraft climb trajectories,
   - Competitive with other well-established SOA approaches,
4. Particular Adaptive Kernel and Gaussian mixture implementation,
   - Showed that it can be used in optimal control problems to obtain solutions close to optimal, and still realistic.

THANK YOU FOR YOUR ATTENTION

References

Recht, B. (2018). A tour of reinforcement learning: The view from continuous control. arXiv:1806.09460.

Sugiyama, M., Takeuchi, I., Suzuki, T., Kanamori, T., Hachiya, H., and Okanohara, D. (2010). Conditional density estimation via least-squares density ratio estimation. In Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pages 781–788.