Machine learning and Optimal Control

Donato Vásquez-Varas

Universidad de Chile

Departamento de Ingeniería Matemática

January 8, 2026

Table of content

Tabla de contenidos

This course focuses on the application of machine learning techniques to optimal control. Specifically, we will study the problem of construction a smooth (approximately) optimal Feedback law.
The general content of the course is the following:

- A brief introduction to optimal control: dynamic programming, Pontryaguin maximum principle, Hamilton Jacobi-Bellman equation.
- Learning problems for the synthesis of optimal Feedback laws.
- Convergence analysis of the cost functional (stability and consistency).
- Approximation of semiconcave functions.

Introduction

Some question arise:

- Why should we use Machine Learning at all?
- Are performance guarantees? Convergence?
- Which of the multiple methods available is more suitable? Does it depends on the parameters of the problem?

We will try to answer these question during the course.

Tabla de contenidos

Optimal control

We start by defining the type of control problem we will work with:

$$(P)_{y_0} \min_{u} J(u, y_0) := \int_0^{\infty} \left(\ell(y(t)) + \frac{\beta}{2} |u(t)|^2 \right) dt$$

where $y \in H^1((0,\infty); \mathbb{R}^d)$ is the unique solution of

$$y'(t) = f(y(t)) + Bu(t), \quad t \in (0, \infty), \quad y(0) = y_0.$$

with

- $\beta > 0$,
- $\Omega \subset \mathbb{R}^d$ open, bounded y convex.
- $\ell \in C^1([t_0, T]; \overline{\Omega})$, $\ell \geq 0$, $\ell(0) = 0$, $\nabla \ell(0) = 0$,
- $f \in C^1(\overline{\Omega})$, $B \in \mathbb{R}^{d \times m}$.

Optimal control

We start by defining the type of control problem we will work with:

$$(P)_{y_0} \min_u J(u, y_0) := \int_0^\infty \left(\ell(y(t)) + \frac{\beta}{2}|u(t)|^2 \right) dt$$

where $y \in H^1((0,\infty);\mathbb{R}^d)$ is the unique solution of

$$y'(t) = f(y(t)) + Bu(t), \quad t \in (0,\infty), \quad y(0) = y_0.$$

Remark 1

- *This is a infinite horizon problem which aims to stabilize the system to* 0*, in the sense that* $\ell(0) = 0$*.*
- *The methods that we will see can be applied to more general type of problems, but for the ease of the presentation this one.*

Optimal Control

We are interested in finding a solution in form of a Feedback law, that is, we seek for a function $\hat{u} : \mathbb{R}^d \mapsto \mathbb{R}^m$ such that

$$\begin{cases} y'(t) = f(y(t)) + B\hat{u}(y(t)), \text{ for all } t > 0 \text{ and} \\ y(0) = y_0 \end{cases} \quad (1)$$

has a solution $\hat{y} \in H_1((0, \infty); \mathbb{R}^d)$ and $u^* = \hat{u} \circ \hat{y}$ is a solution of $(P)_{y_0}$.

Optimal Control

There are classical methods to construct an optimal Feedback law. They are based on

- Pontryagin Maximum Principle
- Dynamic Programming

We point out, that the Pontryagin Maximum Principle does not deliver a Feedback-law, but is important to understand the methods we will work with, in particular for the ML based methods.

Optimal Control

To understand these topics we need to bring up some definitions:

- Lagrangian:
  $$\mathcal{L}(y,u,p) = \int_0^\infty \left\{ \ell(y) + \frac{\beta}{2}|u|^2 + p^\top \left(y' - f(y) - Bu \right) \right\} dt$$
- Adjoint equation and state:
  $$-p' + \nabla \ell(y) - Df^\top(y)p = 0$$
- The Hamiltonian:
  $$H(y,u,q) = \ell(y) + \frac{\beta}{2}|u|^2 + q^\top \cdot (f(y) + Bu)$$

Optimal Control

The Pontryagin maximum principle refers to the following optimality conditions of $(P)_{y_0}$:

$$(PM) \begin{cases} u^* = \frac{1}{\beta} B^\top p^*, \; u \in L^2((0,\infty); \mathbb{R}^m) \\ -\frac{d}{dy} p^* - Df^\top(y^*) p^* + \nabla \ell(y^*) = 0 \\ \frac{d}{dt} y^* = f(y^*) + Bu^*, \; y^*(0) = y_0. \end{cases}$$

The first condition is equivalent to:

$$H(y^*, u^*, -p^*) = \min_{u \in \mathbb{R}^m} H(y^*, u, p^*) = 0.$$

Remark 2

It is important to note that no transversality condition is consider. If (y^*, u^*, p^*) are regular enough, then we may expect

$$\lim_{t \to \infty} p(t) = 0,$$

but this will not hold in general.

Remark 3

If the control problem were convex, solving the optimality conditions would be enough, but this is not the general case. Further, the systems PM is not easily solvable unless the problems is very simple or there is a good guess.

Optimal Control

We now turn to dynamic programming. The value function of $(P)_{y_0}$ is defined by

$$V(y_0) = \min_{u \in L^2((0,\infty);\mathbb{R}^d)} J(u, y_0).$$

The core of dynamic programming is the *dynamic programming principle* or *Bellman principle*:

$$\begin{aligned} V(y_0) &= \min_{u \in L^2((0,T);\mathbb{R}^d)} V(y(T)) + \int_0^T (\ell(y) + \frac{\beta}{2}|u|^2)dt \\ &= V(y^*(T)) + \int_0^T (\ell(y^*) + \frac{\beta}{2}|u^*|^2)dt \end{aligned}$$

where y is the state associated to u and T is any time horizon in $(0, \infty)$.

Optimal Control

Using the dynamic programming principle and assuming that V is C^1 we can arrange the terms to obtain:

$$0 = \min_{u \in L^2((0,T);\mathbb{R}^m)} \frac{1}{T}\left(V(y(T)) - V(y_0)\right) + \frac{1}{T}\int_0^T (\ell(y) + \frac{\beta}{2}|u|^2)dt$$

Optimal Control

Using the dynamic programming principle and assuming that V is C^1 we can arrange the terms to obtain:

$$0 = \min_{u \in L^2((0,T);\mathbb{R}^m)} \frac{1}{T}\left(V(y(T)) - V(y_0)\right) + \frac{1}{T}\int_0^T (\ell(y) + \frac{\beta}{2}|u|^2)dt$$

taking $T \to \infty$

$$0 = \min_{u \in \mathbb{R}^m} \nabla V(y_0)^\top (f(y_0) + Bu) + (\ell(y_0) + \frac{\beta}{2}|u|^2)$$

solves the so called Hamilton-Jacobi-Bellman equation:

$$-\min_{u\in\mathbb{R}^m} H(y_0, u, \nabla v(y_0)) = 0, \text{ for all } y_0 \in \mathbb{R}^d. \tag{HJB}$$

solves the so called Hamilton-Jacobi-Bellman equation:

$$-\min_{u\in\mathbb{R}^m} H(y_0, u, \nabla v(y_0)) = 0, \text{ for all } y_0 \in \mathbb{R}^d. \tag{HJB}$$

An implication of this fact, is the verification formula, namely, u^* is an optimal control with associated state y^* and adjoint state p^* if an only if

$$p^*(t) = -\nabla V(y^*(t)), \ u^*(t) = -\frac{1}{\beta} B^\top \nabla V(y^*).$$

solves the so called Hamilton-Jacobi-Bellman equation:

$$-\min_{u\in\mathbb{R}^m} H(y_0, u, \nabla v(y_0)) = 0, \text{ for all } y_0 \in \mathbb{R}^d. \tag{HJB}$$

An implication of this fact, is the verification formula, namely, u^* is an optimal control with associated state y^* and adjoint state p^* if an only if

$$p^*(t) = -\nabla V(y^*(t)), \; u^*(t) = -\frac{1}{\beta}B^\top \nabla V(y^*).$$

This gives the following formula for the optimal feedback law:

$$\hat{u}(y) = -\frac{1}{\beta}B^\top \nabla V(y).$$

Optimal Control

Classical approaches relies on solving the HJB equation Classical approaches relies on:

- Solving the (HJB) equation. this requires to discretize the problem!
- Use the Verification Formula to construct and optimal Feedback.

Classical approaches relies on solving the HJB equation Classical approaches relies on:

- Solving the (HJB) equation. this requires to discretize the problem!
- Use the Verification Formula to construct and optimal Feedback.

But they suffer from the curse of dimensionality!

Optimal Control

- We can consider for example the Semi-Lagrangian approach.

Optimal Control

- We can consider for example the Semi-Lagrangian approach.
- In this method the control problem is discretized first:

$$y_{i+1} = y_i + h(f(y_i) + Bu_i)$$

$$J_h(u, y_0) = \sum_{i=1}^{\infty} \left(\ell(y_i) + \frac{\beta}{2}|u_i|^2 \right) h$$

Optimal Control

- We can consider for example the Semi-Lagrangian approach.
- In this method the control problem is discretized first:

$$y_{i+1} = y_i + h(f(y_i) + Bu_i)$$

$$J_h(u, y_0) = \sum_{i=1}^{\infty} \left(\ell(y_i) + \frac{\beta}{2}|u_i|^2 \right) h$$

- From which we obtain a discrete version of the Bellman principle:

$$V_h(y_0) = \min_{u \in \mathbb{R}^d} h(\ell(y_0) + \frac{\beta}{2}|u|^2) + V_h(y_0 + h(f(y_0) + Bu)).$$

Optimal Control

- We can try to solve this problem by considering a grid ${y_i}^n$ of $\overline{\Omega}$ of size Δy and considering a set of basis functions $\mathcal{B}_n = \{\phi_i\}_{i=1}^n$ for which the grid is unisolvent, the equations is replaced by

$$V_{h,n}(y_i) = \min_{u \in \mathbb{R}^d} h(\ell(y_i) + \frac{\beta}{2}|u|^2) + V_{h,n}(y_i + h(f(y_i) + Bu))$$

where $V_{h,n}$ lives in the space generated by $\mathcal{B}_n$.

Optimal Control

- We can try to solve this problem by considering a grid ${y_i}^n$ of $\overline{\Omega}$ of size Δy and considering a set of basis functions $\mathcal{B}_n = \{\phi_i\}_{i=1}^n$ for which the grid is unisolvent, the equations is replaced by

$$V_{h,n}(y_i) = \min_{u \in \mathbb{R}^d} h(\ell(y_i) + \frac{\beta}{2}|u|^2) + V_{h,n}(y_i + h(f(y_i) + Bu))$$

  where $V_{h,n}$ lives in the space generated by $\mathcal{B}_n$.
- Unisolvent means that for all $\phi_1, \phi_2 \in$ the space generated by $\mathcal{B}_n$ we have

$$\phi_1(y_i) = \phi_2(y_i) \text{ for all } i \Rightarrow \phi_1 = \phi_2$$

Optimal Control

- Since the grid is unisolvent, this equation determines the coefficientes of $V_{h,n}$. In turns, this could be solve by Picard iterations:

$$V_{h,n,j+1}(y_i) = \min_{u \in \mathbb{R}^d} h(\ell(y_i) + \frac{\beta}{2}|u|^2) + V_{h,n,j}(y_i + h(f(y_i) + Bu))$$

Optimal Control

- Since the grid is unisolvent, this equation determines the coefficientes of $V_{h,n}$. In turns, this could be solve by Picard iterations:

$$V_{h,n,j+1}(y_i) = \min_{u\in\mathbb{R}^d} h(\ell(y_i) + \frac{\beta}{2}|u|^2) + V_{h,n,j}(y_i + h(f(y_i) + Bu))$$

- In the case of finite elements we obtain the following error bound:

$$\|V - V_{h,n}\|_{C(\overline{\Omega})} \leq C\left(h^{\frac{1}{2}} + \frac{\Delta y}{h}\right).$$

Optimal Control

- Since the grid is unisolvent, this equation determines the coefficientes of $V_{h,n}$. In turns, this could be solve by Picard iterations:

$$V_{h,n,j+1}(y_i) = \min_{u \in \mathbb{R}^d} h(\ell(y_i) + \frac{\beta}{2}|u|^2) + V_{h,n,j}(y_i + h(f(y_i) + Bu))$$

- In the case of finite elements we obtain the following error bound:

$$\|V - V_{h,n}\|_{C(\overline{\Omega})} \leq C\left(h^{\frac{1}{2}} + \frac{\Delta y}{h}\right).$$

- This seems to be a nice bound, however, in a high dimensional context, if we want to bound the space error by $\varepsilon > 0$, we need at least $\left(\frac{1}{\varepsilon}\right)^d$ points! This issue is called the *curse of dimensionality*.

The curse of dimensionality is our main motivation to use machine learning techniques for solving control problems, although they have some limitations as well. Some new questions arise

1. Which methods exist?
2. Can we give any performance guarantee of the Feedback? In the case of classical methods there is an error bound but also growths badly with the dimension.

Before continuing we need to address another important issue we were forgetting, the solutions of (HJB) equation are not C^1. We need to work we *viscosity solutions* instead of classical solutions.

Optimal Control

Definition 1

For $v \in C(\overline{\Omega})$ *we consider*

$$D^+v := \{q \in \mathbb{R}^d : \limsup_{h \to 0} \frac{v(y+h) - v(y) - q\top \cdot h}{|h|} \leq 0\}$$

is called the upper-differential of v *at* y *and*

$$D^-v := \{q \in \mathbb{R}^d : \liminf_{h \to 0} \frac{v(y+h) - v(y) - q\top \cdot h}{|h|} \geq 0\}$$

is called the sub-differential of v *at* y

Some important facts:

- D^+v and D^-v are not empty almost everywhere.
- v is differnetiable at y if and only if

$$D^+v(y) = D^-v(y) = \{\nabla v(y)\}.$$

- $D^+v(y)$ and $D^-v(y)$ are convex and closed sets for all y.
- D^+v is upper semi-continuous, that is, if $y_n \to y$ and $p_n \in D^+v(y_n)$ is such $p_n \to p$, then

$$p \in D^+v(y).$$

Definition 2

Let us consider $F : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$, *we say that* $v \in C(\overline{\Omega})$ *is a viscosity solution of*

$$F(y, \nabla v(y)) = 0, \ \forall y \in \overline{\Omega}$$

if for all $y \in \overline{\Omega}$

$$F(y, q) \leq 0, \ \forall q \in D^{+}v(y) \ \text{(sub-solution)}$$

and

$$F(y, q) \geq 0, \ \forall q \in D^{-}v(y) \ \text{(super-solution)}.$$

In the case of HJB equation, it is well known that V is the unique viscosity solution of (HJB):

$$-\min_{u\in\mathbb{R}^m} H(y,u,\nabla V(y)) = 0$$

that is

$$\min_{u\in\mathbb{R}^m} H(y,u,p) \geq 0 \text{ for all } p \in D^+V(y)$$

and

$$\min_{u\in\mathbb{R}^m} H(y,u,p) \leq 0 \text{ for all } p \in D^-V(y).$$

As we mentioned before, V is not differentiable in general, nevertheless it does have a regularity notion which is connected to the concept of viscosity solutions:

Definition 3

We say that a function $v \in C(\overline{\Omega})$ is C-semiconcave if $x \mapsto v(x) - \frac{C}{2}|x|^2$ is concave.

A natural regularity assumption for the value function (and the solutions of HJB) is the semiconcavity.

In fact, a semiconcave functions is always kind of a value function:

Theorem 1

$v \in C(\overline{\Omega})$ is $C-$semi-concave if and only if there exists $\{\phi_i\}_{i=1}^{\infty} \subset C^2(\overline{\Omega})$ with $\|\phi_i\|_{C(\overline{\Omega};R^{d\times d})} \leq C$ such that

$$v(y) = \inf_{i\in\mathbb{N}} \phi_i(y).$$

This will be important to contruct a parametrization-approximation for semiconcave functions!

Optimal Control

Semiconcave functions satisfies the following properties:

1. v is locally Lipschitz in Ω.
2. The upper differential of v is never empty and

$$p \in D^+v(y) \Leftrightarrow v(h+y) - v(y) + p^\top \cdot h \leq C|h|^2,$$

   $\forall |h| < dist(x, \partial\Omega)$.

3. v is semiconcave if and only if the largest eigenvalue of $\nabla^2 v$ is bounded by C in the sense of distributions:

$$\int_{\mathbb{R}^d} x^\top \nabla^2 \phi(y) x v(y) dy \leq C|x|^2 \text{ for all } x \in \mathbb{R}^d.$$

4. v is C^1 and twice differentiable almost everywhere.

There is a connection with HJB:

Theorem 2

If v is semiconcave and v satsifies (HJB) almost everywhere, i.e.,

$$\min_{u \in \mathbb{R}^m} H(y, u, \nabla v(y)) = 0 \text{ for almost all } y \in \Omega$$

then v is a viscosity solution of (HJB) in Ω.

The semiconcavity of the value function will allow us to study the convergence of machine learning methods. It also helps to construct a parametrization.

Tabla de contenidos

ML Algorithms

There are many methods in the literature that attempt to solve HJB equations and/or construct a optimal feedback laws. In the rest of the course we will consider the followings:

1. **Regression along trajectories**: In this method the optimal feedback law is obtained by solving the control problem for many initial conditions and then fit a ML model to approximate the optimal control, adjoint state or the value function along the trajectories.

ML Algorithms

There are many methods in the literature that attempt to solve HJB equations and/or construct a optimal feedback laws. In the rest of the course we will consider the followings:

1. **Regression along trajectories**: In this method the optimal feedback law is obtained by solving the control problem for many initial conditions and then fit a ML model to approximate the optimal control, adjoint state or the value function along the trajectories.
2. **Averaged method**: This method lies in the category of unsupervised ML algorithms. It consists in replacing J by an averaged version (with respecto to the intial conditions) and parametrized the control.

ML Algorithms

There are many methods in the literature that attempt to solve HJB equations and/or construct a optimal feedback laws. In the rest of the course we will consider the followings:

1. **Regression along trajectories**: In this method the optimal feedback law is obtained by solving the control problem for many initial conditions and then fit a ML model to approximate the optimal control, adjoint state or the value function along the trajectories.
2. **Averaged method**: This method lies in the category of unsupervised ML algorithms. It consists in replacing J by an averaged version (with respecto to the intial conditions) and parametrized the control.
3. **PINNS**: This also can be considered a unsupervised method. This method tries to directly solve the HJB equation. The fact that we look for a viscosity solution makes the analysis of this problem very challenging.

ML algorithms

To describe the methods we need the concept of setting

Definition 4

We define a setting as a tupple $(\Theta, v, \mathcal{P})$ *in which:*

- Θ *is a finite dimensional Banach space, we call this the space of parameters or parametric space.*
- $v : \Theta \mapsto C^2(\overline{\Omega})$ *is a continuous function. We call it the parametrization.*
- $\mathcal{P} : \Theta \mapsto [0, \infty)$ *a continuous coersive function, that is,*

$$\lim_{\|\theta\|_\Theta \to \infty} \mathcal{P}(\theta) = \infty.$$

We call it the penalty function.

We can consider the following examples of settings:

- **Polynomials**: $\phi_i(x) = x^i$, $\Theta_n = \mathbb{R}^{n+1}$, $v_n(\theta)(x) = \sum_{i=1}^{n+1} \phi_{i-1}(x)\theta_i$, $\mathcal{P}_n(\theta) = \alpha_1|\theta|^2 + \alpha_2|\theta|_1$.
- **d-Polynomials** For $\alpha \in \mathbb{N}^d$, $\phi_\alpha(x) = \prod_{i=1}^{d} x^{\alpha_i}$. Defining

$$\Lambda_n = \{\alpha \in \mathbb{N}^d : |\alpha|_\infty \leq n\},$$

$$\Theta_n = \mathbb{R}^{\Lambda_n}$$

,

$$v_n(\theta)(x) = \sum_{\alpha \in \Lambda_n} \theta_\alpha \phi_\alpha(x)$$

$\mathcal{P}_n(\theta) = \alpha_1|\theta|^2 + \alpha_2|\theta|_1$.

Another important case is the neural network parametrization. We will only consider the case of Shallow Neural Networks. For $\theta = (a_1, b_1, a_0, b_0)$ with $a_1 \in \mathbb{R}^n$, $a_0 \in \mathbb{R}^{n\times d}$, $b_0 \in \mathbb{R}^n$ and $b_1 \in \mathbb{R}$, we set consider the following parametrization

$$v(\theta) = b_1 + \sum_{i=1}^{n} a_{1,i}\phi(a_{0,i}^\top \cdot x + b_{0,1}\cdot).$$

with $\phi : \mathbb{R} \mapsto \mathbb{R}$ being the activation function. For example

$$\phi(x) = \max(x,0)^2, \phi(x) = \frac{1}{1+\exp(-x)}, \phi(x) = \tanh(x)$$

ML algorithms

In this case the parametric space is $\Theta = \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}^{n\times d} \times \mathbb{R}^n$. As for the penalty functions, typically, the euclidean norm is used. Nevertheless, there other options, for example

$$\mathcal{P}_n(\theta) = |b_1| + \sum_{i=1}^{n} |a_1|(|a_{0,i}| + |b_i|).$$

This is norm is connected to the *Barron* space associated to the activation function.

ML algorithm

We know recall the verification formula, which states that the optimal Feedback must be of the form

$$u = -\frac{1}{\beta} B^\top \nabla V$$

with V the value function of the control problem. The idea is to use the parametrization to replace the value function in the verification formula, which delivers the following parametrization for the optimal feedback law:

$$u(\theta) = -\frac{1}{\beta} B^\top \nabla v(\theta)$$

In general, ∇ represent the gradient with respect to the state variables.

ML algorithm

The three methods we are about to discuss consists in finding a parameter θ^* as the solution of a *Continuous Learning problem* of the following form

$$\min_{\theta \in \Theta} \int_{\Omega} L(\theta, x) f(x) dx + \alpha \mathcal{P}(\theta)$$

where L is called the loss function and f is a probability density function over Ω. For simplicity we will assume $f(x) = \frac{1}{|\Omega|}$.
For each method, we consider a different choice of loss function, with the hope that $v(\theta^*)$ delivers an optimal feedback law by means of the verification formula.

ML algorithm

Of course, in practice we cannot solve such a problem, because the integrals involved are intractable. Instead we solve a discrete version. That is, for a training set $\mathcal{Y} = \{y_i\}_{i=1}^N$ we consider the Monte Carlo approximation of the learning problem:

$$\min_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^{N} L(\theta, y_i) + \alpha \mathcal{P}(\theta)$$

where the training set is sampled independently from the distribution associated to f.

ML algorithm

The convergence of the discrete problem towards the continues one is ensured by:

1. The Uniform Law of Large Numbers, which states that if $g : X \times Y \mapsto \mathbb{R}$ is continuous in x, X is a compact subset of $\mathbb{R}^d$, and integrable with respect to Y, then for a iid sample $\{y_i\}_{i=1}^{\infty}$ we have

$$\lim_{N\to\infty} \sup_{x\in X} \left| \frac{1}{N} \sum_{i=1}^{N} g(x, y_i) - \mathbb{E}_y(g(x, \cdot)) \right| = 0$$

   almost surely.

2. The set of solutions of the continuous learning problem is compact due to the coersivity of $\mathcal{P}$.

Tabla de contenidos

Regression Along Trajectories

A first attempt to construct an optimal feedback law could be to find $\theta^* \in \Theta$ by minimizing the distance between the optimal feedback law and our parametrization:

$$\min_{\theta \in \Theta} \frac{1}{|\Omega|} \int_\Omega \frac{1}{\beta} |B^\top (\nabla V(x) - \nabla v(\theta)(x))|^2 dx + \alpha \mathcal{P}(\theta)$$

Of course, we do not have access to V, but we could consider a set of initial conditions $\mathcal{Y}_{train} = \{y_i\}_{i=1}^N$ in Ω and approximate this problem by Monte Carlo (or any other integration method):

$$\min_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \frac{1}{\beta} |B^\top (\nabla V(y_i) - \nabla v(\theta)(y_i))|^2 + \alpha \mathcal{P}(\theta)$$

ML algorithm

Again, we do not have direct access to V, but if we solve the optimal control problem, we know that

$$V(y_0) = J(u^*, y_0), \;\nabla V(y_0) = -p^*(0).$$

- Therefore this approach implies that we need to solve N optimization problems
- After that, this the Learning problems becomes a simple regression problem.
- In some applications N must be very large, and the control problem could be very expensive to solve.
- We can leverage the dynamic programming principle.

Regression Along Trajectories

By the dynamic programming principle we have that if u^* is an optimal control and y^* is the corresponding optimal trajectory, then $u^*|_{[T,\infty)}$ is an optimal control for $P_{y^*(T)}$. Consequently we have

$$V(y^*(T)) = J(u^*|_{[T,\infty)}, y^*(T)), \; \nabla V(y^*(T)) = -p^*(T),$$

that is, we know the value function and its gradient along the optimal trajectory.

Regression Along Trajectories

In the regression along trajectories approach the idea is to utilize the fact that we know the value function along the trajectory:

$$\min_{\theta\in\Theta} \frac{1}{|\omega|}\int_{\omega}\int_0^T \frac{1}{\beta}|B^\top(\nabla V(y^*(t;y_0)) - \nabla v(\theta)(y^*(t;y_0)))|^2 dt dy_0 + \alpha\mathcal{P}(\theta)$$

where $T > 0$ is a finite time horizon time horizon and $y^*(\cdot, y_0)$ is the optimal trajectory emanating from y_0. In practice we have

$$\min_{\theta\in\Theta} \frac{h}{N}\sum_{i=1}^{N}\sum_{j=1}^{M}\frac{1}{\beta}|B^\top(\nabla V(y_{i,j}) - \nabla v(\theta)(y_{i,j}))|^2 + \alpha\mathcal{P}(\theta)$$

where $h = T/M$ and $\{y_{i,j}\}_{j=1}^{M}$ is a discretization of $y^*(\cdot, y_i)$ at times steps $t_j = jh$.

Regression Along Trajectories

Now we arrive to the following question: How can we solve the control control?

1. Solve the boundary value problem stemming from the Pontryagin principle up to time $T > 0$:

$$\begin{cases} \dfrac{d}{dt}y = f(y) + Bu, \;\; y(0) = y_0 \\ -\dfrac{d}{dt}p - Df^\top(y)p + \nabla\ell(y) = 0, \;\; p(T) = 0 \\ u = \dfrac{1}{\beta}B^\top p. \end{cases}$$

2. Directly solve the optimization problem by a gradient-Newton type method.

In both cases we must choose a discretization method for the involved equations.

We need a discretization method which ensures stability and precision, for instance the Crank-Nicolson method which is an implicit method second order method.

1. Discretization of the cost:

$$J_M(u, y_0) = \sum_{j=0}^{M-1} \frac{h}{2} \left(\frac{\beta}{2}(|u_j|^2 + |u_{j+1}|^2) + \ell(y_j) + \ell(y_{j+1}) \right)$$

2. Discretization of the dynamics:

$$y_{j+1} = y_j + \frac{h}{2}\left(f(y_j) + f(y_{j+1}) + B(u_j + u_{j+1})\right).$$

Regression Along Trajectories

In this case, the discrete adjoint state satisfies the following equation:

$$\frac{1}{h}(p_{j-1} - p_j) + \nabla\ell(y_j) - Df(y_j)^\top(p_{j-1} + p_j) = 0 \text{ for } j < M$$

$$p_{M-1} = -\frac{h}{2}(I_{d\times d} - \frac{h}{2}Df^\top(y_M))^{-1}\nabla\ell(y_M)$$

Regression Along Trajectories

The regression along trajectories can be seen as a **weighted regression**. To see this, we note that by using the following change of variables for $t > 0$ fixed

$$z = y^*(t, y_0), dz = |det(Dy^*(t; y_0^*))| dy_0$$

we have

$$\int_\omega \int_0^T \frac{1}{\beta} |B^\top (\nabla V(y^*(t; y_0)) - \nabla v(\theta)(y * (t; y_0)))|^2 dt dy_0$$

$$= \int_{\{y^*(t;y_0): t \in [0,T]; y_0 \in \Omega\}} \frac{1}{\beta} |B^\top (\nabla V(z) - \nabla v(\theta)(z)|^2 g(z) dz$$

with

$$g(z) = \int_{\{t \in [0,T]: \exists y_0 \in \omega,\ z = y^*(t, y_0)\}} \frac{1}{|det(Dy(t, y_0(z)))|} dt.$$

Let us see some examples. We start by considering a modified version of the Van der Pol oscillator:

$$\min \int_0^\infty \left(\frac{1}{2}|y|^2 + \frac{\beta}{2}|u|^2 \right) dt$$

$$\frac{d^2}{dt^2}y = \mu(1-y^2)\frac{d}{dt}y - y + \gamma y^3 + u \qquad (2)$$

$$y(0) = y_0, \frac{d}{dt}y(0) = v.$$

with $\mu = \frac{3}{2}$, $\gamma = \frac{4}{5}$. We will consider a reference domain $\Omega = [-10, 10]^2$ (state and velocity).

To measure the performance of this approach we will consider the following metrics:

$$VRMAE = \sum_{i=1}^{N} J(\hat{u}_i, y_i) / \sum_{i=1}^{N} V(y_i)$$

and

$$CRMSE = \sum_{i=1}^{N} \|\hat{u}_i - u_i^*\|_{L^2((0,T);\mathbb{R}^m)} / \sum_{i=1}^{N} \|u_i^*\|_{L^2((0,T);\mathbb{R}^m)}$$

where $\{y_i\}_{i=1}^N$ is a set of initial conditions, u_i^* an approximation of the optimal controls and $\hat{u}_i$ the control given by the Feeedback law applied to the intial condition y_i.

The Van der Pol oscillator is a 2-dimensional example, maybe we can consider something with a larger dimensionality, for example, the discretization of a PDE control problem.

Regression Along Trajectories

We can consider the stabilziation of the Allen Chan equation by a finite number of actuators:

$$
\begin{aligned}
\min_{u \in L^2((0,\infty);\mathbb{R}^m)} \quad & \frac{1}{2}\int_0^\infty \|y\|_{L^2(-1,1)} + \frac{\beta}{2}\sum_{i=1}^m \int_0^\infty |u_i|^2 dt \\
s.a. \quad & \frac{d}{dt}y = \nu \frac{d^2}{dx^2}y + y(1-y^2) + \sum_{i=1}^m \chi_{\omega_i} u_i \\
& \frac{d}{dx}y(t,-1) = \frac{d}{dx}y(t,1) = 0, \;\; y(0,x) = y_0(x)
\end{aligned} \tag{3}
$$

Regression Along Trajectories

- The uncontrolled system has 3 steady states of interest: $y = -1$ and $y = 1$ (stable) and $y = 0$ (unstable).

To discretize the PDE we use the Chebyshev Spectral Collocation method. This is not the focus of these lectures, hence we will only sketch the method.

- Chebyshev polynomials $\{\phi_i\}_{i=1}^{\infty}$ are a Ortogonal basis of $L^2_\mu(-1,1)$ which is the space of square integrable functions using the following measure

$$\mu = \frac{1}{\sqrt{1-x^2}}.$$

- They are given by the following recusive formulas

$$\phi_0(x) = 1,\ \phi_1(x),\ \phi_{i+1}(x) = 2x\phi_i(x) - \phi_{i-1}(x)$$

- They are used together with the Chebyshev points:

$$x_{i,N} = -\cos\left(\pi \frac{i}{N}\right)$$

for $i \in \{1, \ldots, N\}$. These points are unisolvent for the the Chebyshev polynomials of degree smaller than N.

Regression Along Trajectories

- The state is replaced by its Chebyshev truncation:

$$y(t,x) \approx \sum_{0=1}^{N} \phi_i(x)\mathcal{Y}_i(t)$$

- The equation is approximated by evaluating it at the Chebyshev points, which deliver the following finite dimensional system:

$$\frac{d}{dt}\mathcal{Y} = \nu A\mathcal{Y}_i + Bu + (1-\mathcal{Y}^2)\mathcal{Y}$$

with

$$B_{i,j} = \begin{cases} 1 & x_{i,n} \in \omega_j \\ 0 & x_{i,n} \notin \omega_j \end{cases}$$

Regression Along Trajectories

- For the setting of the learning problem we consider Polynomials and NN.

Regression Along Trajectories

- For the setting of the learning problem we consider Polynomials and NN.
- But we cannot use the full basis of polynomials! ($|\mathcal{B}_n| = n^d$)

Regression Along Trajectories

- For the setting of the learning problem we consider Polynomials and NN.
- But we cannot use the full basis of polynomials! ($|\mathcal{B}_n| = n^d$)
- Instead we consider the Hyperbolic cross basis:

$$\Gamma_n = \{\alpha \in \mathbb{N}^d : \prod_{i=1}^{d}(\alpha_i + 1) \leq n\}$$

$$\tilde{\mathcal{B}}_n = \{\phi_\alpha : \alpha \in \Gamma_n\}$$

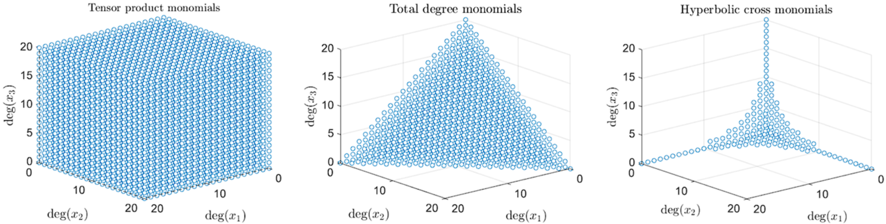

Figure: Polynomials basses.

Tabla de contenidos

AFLS

We now describe the Averaged Feedback Law Scheme (AFLS), which consists in:

1. Parametrizing the control by the Verification Formula
$$u = -\frac{1}{\beta} B^\top \nabla v(\theta)$$
2. Minimize a truncated averaged version of the cost:
$$\mathcal{J}_T(v) = \int_\omega \mathcal{V}_T(y_0; v) d_y 0$$
with
$$\mathcal{V}_T(y_0; v) = \int_0^T \left(\ell(y(t; y_0, v)) + \frac{1}{2\beta} |B^\top \nabla v(y(t; y_0, v))| \right) dt$$
and $y(t; y_0, v)$ being the (unique) solution of
$$\frac{d}{dt} y = f(y) - \frac{1}{\beta} BB^\top \nabla v(y), \; y(0) = y_0.$$

The learning problem in the AFLS method is the following:

$$\min_{\theta\in\Theta} \frac{1}{|\omega|}\mathcal{J}_T(v(\theta)) + \alpha\mathcal{P}(\theta).$$

The discrete version of this problem is given by

$$\min_{\theta\in\Theta} \frac{1}{N}\sum_{i=1}^{N}\mathcal{V}_{T,M}(y_i; v(\theta)) + \alpha\mathcal{P}(\theta).$$

where $\mathcal{V}_{T,M}$ is obtained, for example, by applying the Crank-Nicolson scheme.

AFLS

We have then that $\mathcal{V}_{T,M}$ is given by

$$\mathcal{V}_{T,M}(y_0; v(\theta)) = \sum_{j=1}^{M} \frac{h}{2} \left(\ell(y_j) + \ell(y_{j+1}) + \frac{1}{2\beta} |\nabla v(y_{j+1})|^2 + |\nabla v(y_{j+1})|^2 \right)$$

and

$$y_{j+1} = y_j + \frac{h}{2} \left(f(y_j) + f(y_{j+1}) - \frac{1}{2} BB^\top (\nabla v(y_i) + \nabla v(y_{i+1+})) \right)$$

with $h = \frac{T}{M}$.

The corresponding adjoint state is given by

$$\frac{1}{h}(p_{j-1} - p_j) + \nabla\ell(y_j) - \frac{1}{2}Df(y_j)^\top(p_j + p_{j-1})$$
$$+\nabla^2 v(\theta)(y_j)^2 BB^\top(\frac{1}{2}(p_{j-1} + p_j) + \nabla v(y_j)) = 0$$

for $j < M$ and

$$\frac{h}{2}\left(\nabla\ell(y_M) + \frac{1}{\beta}\nabla^2 v(y_M)BB^\top\nabla v(y_M)\right)$$
$$-\frac{h}{2}\left(Df^\top(y_M) + \frac{1}{\beta}\nabla v^2(y_M)BB^\top\right)p_{M-1} + p_{M-1} = 0$$

AFLS

In the continuous case the adjoint is given by

$$-\frac{d}{dt}p + \nabla \ell(y) - Df(y)^\top p + \frac{1}{\beta}\nabla^2 v(y)BB^\top(p + \nabla v) = 0 \text{ and } p(T) = 0.$$

With the help of the adjoint we have the derivative of the objective function is

$$\frac{d}{d\theta}\mathcal{J}(v(\theta)) = \frac{1}{|\omega|\beta}\int_\omega \frac{d}{d\theta}\nabla v(\theta)BB^\top(p + \nabla v(\theta))dtdy_0$$

From this we can see that the optimality conditions of the learning problem implies that the solution of the learning problems θ^* also solves

$$\min_{\theta\in\Theta} \frac{1}{2|\omega|\beta}\int_\omega |B^\top(p + \nabla v(\theta))|^2 dtdy_0 + \alpha\mathcal{P}(\theta)$$

which is very similar to the Learning problem of RAT method!

AFLS

Let us see the performance of this method in same examples as in the case of the RAT method.

Tabla de contenidos

PINNS

- The Physics Informed Neural Networks approach (PINNS) try to solve directly the HJB equation

PINNS

- The Physics Informed Neural Networks approach (PINNS) try to solve directly the HJB equation
- it is less direct than the other methods, since we do not try to obtain the feedback from the resolution of a learning problem.

PINNS

- The Physics Informed Neural Networks approach (PINNS) try to solve directly the HJB equation
- it is less direct than the other methods, since we do not try to obtain the feedback from the resolution of a learning problem.
- We first solve the HJB equation with the hope that the solution approximates the value function of the control problem.

PINNS

We need to recall that the value function solves the HJB equation:

$$\mathcal{H}(y, \nabla v(y)) := -\min_{u \in \mathbb{R}^m} H(y, u, \nabla v(y))$$

$$= -\ell(y) - \nabla v(y) \cdot f(y) + \frac{1}{2\beta} |B^\top \nabla v(y)|^2 = 0$$

... in the viscosity sense.

PINNS

Typically, the PINNS method tries to minimize the square of the equations residuals. In this case, than means:

$$\min_{\theta \in \Theta} \frac{1}{|\Omega|} \int_{\Omega} (\mathcal{H}(u, \nabla v(\theta)(y)))^2 dy + \alpha \mathcal{P}(\theta).$$

However this do not deliver (necessarily) a viscosity solution, instead, it gives a generalized solution (satisfies HJB a.e.).

It is important to bear in mind that generalized solutions are not unique, for instance, the distance function to the boundary of Ω is unique viscosity solution of

$$-|\nabla v|^2 = -1 \text{ in } \Omega \ , v = 0 \text{ on } \partial\Omega,$$

and consequently a generalized solutions. Additionally, $\tilde{v} = -v$ is also a generalized solution, since it satisfies the equation *a.e.*.

There are two strategies to remedy this problem

- To regularize the equation by adding a viscosity term $-\varepsilon \nabla v^2$. This makes the solution of the problem unique, but its performance is not clear.

There are two strategies to remedy this problem

- To regularize the equation by adding a viscosity term $-\varepsilon \nabla v^2$. This makes the solution of the problem unique, but its performance is not clear.
- To modify the problem using that the viscosity solution is given by:

$$V(y) = \sup_{\phi \in C^2 \text{ and is a sub-solution of HJB}} \phi(x).$$

Following the second approach we solve:

$$\min_{\theta\in\Theta}\frac{1}{|\Omega|}\int_{\Omega}(-v(x)+\gamma\max(\mathcal{H}(x,\nabla v(\theta)(x)),0)^2)+\alpha\mathcal{P}(\theta)$$

As γ tend to infinity, the function $v(\theta^*_\gamma)$ (θ^*_γ solution of the problem) the term $\mathcal{H}(x,\nabla v(\theta^*_\gamma)(x))$ is becoming negative for almost all $x\in\Omega$.

It is interesting to note that this problem is convex if the parametrization is linear in the parameters, since

$$\mathcal{H}(x, \nabla v(\theta)(x)) = -\ell(x) - \nabla v(\theta)(x) \cdot f(x) + \frac{1}{2\beta}|B^\top \nabla v(\theta)(x)|^2$$

which would be a quadratic functions in the parameters.

The discrete version of this problem is therefore given by

$$\min_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^{N} \left(-v(x_i) + \gamma \max(\mathcal{H}(x_i, \nabla v(\theta)(x_i)), 0)^2\right) + \alpha \mathcal{P}(\theta)$$

PINNS

Let us see know the performance of this approach in the same examples as before.

Summary

Method	Needs data?	Complexity	Convexity	Opt Feedback
RAT	Yes	$N \times M$	True	Yes
AFLS	No	$N \times M$	False	Yes
PINSS	No	N	True	False

Summary

- RAT and AFLS method behaves similarly and they are able to find a feedback-law.
- PINNs approach is not able to find an optimal Feedback law.
- Why? Is this true for any control problem?

Tabla de contenidos

Convergence Analysis

As is usual in numerical analsysi, there are two concepts which are important regarding the convergence of these methods:

1. **Stability**: For a given $T > 0$, $v \in C^2(\overline{\Omega})$, we say that v is a stable feedback if $y([0, T]; \omega, v) \subset \Omega_\delta$ (distance $\delta > 0$ to the boundary).
2. **Consistency**: We say that a sequence of Feedback laws $\phi_n \in C^2(\overline{\Omega})$ is consistent if for a given $p \in [1, \infty]$ there exists $0 < T_n \to \infty$ such that

$$\lim_{n \to \infty} \|\mathcal{V}_{T_n}(\cdot; \phi_n) + V(y(T_n; \cdot, \phi_n)) - V\|_{L^p(\omega)} = 0.$$

Convergence Analysis

We want to analyze the convergence the methods in the following sense.

Definition 5

Let $S_n = (\Theta_n, v_n, \mathcal{P}_n)$ *be a sequence of setting and consider* $\theta_n^* \in \Theta_n$ *a sequence of solution for one of the methods with the setting* S_n. *We want to elucidate under which conditions there exists* $0 < T_n \to \infty$ *such that*

$$\lim_{n\to\infty} \|\mathcal{V}_{T_n}(\cdot; v_n(\theta_n^*)) - V\|_{L^p(\omega)} = 0$$

for given $p \in [1, \infty]$ *and* $\omega \Subset \Omega$.

As we will see, for this is enough that $v_n(\theta_n^*)$ to be **stable** and **consistent**.

Convergence Analysis

For carrying out the convergence analysis we need the following ingredients

- Escape time estimates from the reference domain Ω.
- Consistency error estimates.

In both cases we will look for reasonable hypotheses (achievable).

Remark 4

The escape time estimates are crucial for ensuring that the trajectories stay in a domain where we can provide a local approximation of the value function. This in turns will allow us to use the consistency error estimates which are of local nature.

Tabla de contenidos

Escape time estimates

To be precise the stability condition or hypothesis is the following:

Hypothesis 1

For the tuple (T, ϕ, δ, y_0) *with* $T > 0$, $\phi \in C^2(\overline{\Omega})$, $\delta > 0$, *and* $y_0 \in \Omega_\delta$, $y(\cdot; y_0, \phi)$ *exists on* $[0, T]$ *and* $y(t; y_0, \phi) \in \Omega_\delta$ *for all* $t \in [0, T]$.

What happen if we approximate a stable feedback in ω, is it also stable?

Escape time estimates

Lemma 3

Let $T > 0$ *be a time horizon,* $y_0 \in \omega$, $g_1 \in C^2(\Omega)$, $g_2 \in C^2(\Omega)$ *and* $\delta > 0$, *be such that* $(T, u_{g_2}, \delta, y_0)$ *satisfies Hypothesis 1. Assume that*

$$\frac{|B|^2 \|\nabla g_1 - \nabla g_2\|_{L^\infty(\Omega_{\frac{\delta}{4}};\mathbb{R}^d)}}{a\beta}\left(e^{Ta} - 1\right) \leq \frac{\delta}{2},$$

where

$$a = \|Df\|_{L^\infty(\Omega_{\frac{\delta}{4}};\mathbb{R}^{d\times d})} + \frac{|B|^2}{\beta}\|\nabla^2 g_2\|_{L^\infty(\Omega_{\frac{\delta}{4}};\mathbb{R}^{d\times d})},$$

and Df *stands for the derivative of* f. *Then* $(T, u_{g_1}, \frac{\delta}{2}, y_0)$ *satisfies Hypothesis 1.*

- If the value function is $C^2(\overline{\Omega})$ as satisfies Hypothesis 1, then any sufficiently good approximation of V in $C^1(\overline{\Omega})$ will provide a stable feedback.

Escape time estimates

- If the value function is $C^2(\overline{\Omega})$ as satisfies Hypothesis 1, then any sufficiently good approximation of V in $C^1(\overline{\Omega})$ will provide a stable feedback.
- Further more, the escape time from Ω for the approximation is bounded from below by

$$\hat{T}_\varepsilon \geq \frac{1}{a} \log\left(\frac{a\beta\delta}{|B|^2}\varepsilon^{-1} + 1\right)$$

where ε is the error in $C^1(\overline{\Omega})$.

- What happen if we cannot control the $C^1(\overline{\Omega})$ norm of the approximation, even worse, if the value function is not $C^1(\overline{\Omega})$?

Escape time estimates

- What happen if we cannot control the $C^1(\overline{\Omega})$ norm of the approximation, even worse, if the value function is not $C^1(\overline{\Omega})$?
- We can weaken the smoothness of V in exchange of a stronger stability assumption.

Escape time estimates

Hypothesis 2

There exist $\tilde{\delta} > 0$, *and* $w \in C^1(\Omega)$ *such that for*

$$\omega_{\tilde{\delta}} := \{y \in \Omega : w(y) < \sup_{y_0 \in \omega} w(y_0) + \tilde{\delta}\},$$

we have that $\omega \subset \omega_{\tilde{\delta}}$, $\overline{\omega_{\tilde{\delta}}} \subset \Omega$, *and* $\partial\omega_{\tilde{\delta}}$ *is of class* C^1. *Moreover* $\phi \in C(\Omega)$ *is a viscosity super solution of*

$$-\nabla w(y)^\top (f(y) - \frac{1}{\beta} BB^\top \nabla\phi(y)) = 0 \text{ in } \omega_{\tilde{\delta}},$$

i.e. for every $\bar{y} \in \omega_{\tilde{\delta}}$ *and every* $q \in D^-\phi(\bar{y})$ *the following inequality holds*

$$\nabla w(\bar{y})^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top(\bar{y}) q) \leq 0.$$

Escape time estimates

For the next estimates we need to define the following quantities:

$$\sigma_\varepsilon^1 = \sup_{x\in\omega_\delta, y\in B(x,\varepsilon)} \left| -\nabla w(y)^\top f(y) + \nabla w(x)^\top f(x) \right|$$

and

$$\sigma_\varepsilon^2 = \sup_{x\in\omega_\delta, y\in B(x,\varepsilon)} \left| B(x)^\top \nabla w(x) - B(y)^\top \nabla w(y) \right|,$$

for $\varepsilon > 0$ and $w \in C^1(\overline{\Omega})$.

Lemma 4

Let $\omega \subset \Omega$, $\phi \in C(\Omega)$ and $\delta > 0$ such that they satisfy Hypothesis 2.

(a) *If $\phi \in C^1(\Omega)$, consider $\hat{\phi} \in C^2(\Omega)$, and let $\hat{T}$ be the maximum $T > 0$ such that $y([0,T];\omega,\hat{\phi}) \subset \omega_\delta$. Then the following holds*

$$\hat{T}\frac{|B|^2}{\beta}\left(\|\nabla\phi - \nabla\hat{\phi}\|_{C(\omega_\delta;\mathbb{R}^d)}\|\nabla w\|_{C(\omega_\delta;\mathbb{R}^d)}\right) \geq \delta.$$

Lemma 4

Let $\omega \subset \Omega$, $\phi \in C(\Omega)$ and $\delta > 0$ such that they satisfy Hypothesis 2.

(a) *If $\phi \in C^1(\Omega)$, consider $\hat{\phi} \in C^2(\Omega)$, and let $\hat{T}$ be the maximum $T > 0$ such that $y([0,T];\omega,\hat{\phi}) \subset \omega_\delta$. Then the following holds*

$$\hat{T}\frac{|B|^2}{\beta}\left(\|\nabla\phi - \nabla\hat{\phi}\|_{C(\omega_\delta;\mathbb{R}^d)}\|\nabla w\|_{C(\omega_\delta;\mathbb{R}^d)}\right) \geq \delta.$$

(b) *If $\phi \in Lip(\Omega)$, set $\phi_\varepsilon = \phi * \rho_\varepsilon$ a mollification of ϕ, and let T_ε be the maximum T such that $y([0,T];\omega,\phi_\varepsilon) \subset \omega_\delta$. Then there exists ε_0 such that all $\varepsilon \in (0,\varepsilon_0)$ we have*

$$T_\varepsilon\left(\sigma^1_\varepsilon + \frac{\sigma^2_\varepsilon}{\beta}\|B^\top\nabla\phi\|_{L^\infty(\omega_\delta;\mathbb{R}^d)}\right) \geq \delta.$$

where σ^1_ε is defined in (79) and σ^2_ε in (79).

Escape time estimates

- In a) the escape time estimates depends on the C^1 norm of the approximation, but it is not require the C^2 norm to be bounded.
- In b), for the case where the approximation is given by the mollification, the lack of regularity of the feedback is compensated by the extra-smoothnes of the Lyapunov function w.

Tabla de contenidos

Consistency estimates

We start by assuming that V is C^1:

Lemma 5

Let $v \in C^1(\Omega)$ *be a super-solution of* (HJB). *Consider* $\bar{v} \in C^{1,1}(\Omega)$, $y_0 \in \Omega$, *and* T *such that* $y(\cdot; y_0, \bar{v})$ *exists on* $[0, T]$ *and satisfies* $y([0, T]; y_0, u_{\bar{v}}) \subset \Omega$. *Then we have*

$$\mathcal{V}_T(y_0; \bar{v}) + v(y(T; y_0, \bar{v})) - v(y_0)$$
$$\leq \frac{|B|^2}{\beta} \int_0^T |\nabla v(y(t; y_0, u_{\bar{v}})) - \nabla \bar{v}(y(t; y_0, u_{\bar{v}}))|^2 dt$$

Consistency estimates

For the ease of the presentation we write $\bar{y} = y(\cdot; y_0, \bar{v})$, $\bar{u} = -\frac{1}{\beta} B^\top \nabla \bar{v}(\bar{y})$ and $u = -\frac{1}{\beta} B^\top \nabla v(\bar{y})$.

- We have

$$\ell(\bar{y}) + \frac{\beta}{2}|\bar{u}|^2 + \nabla \bar{v}(\bar{y})(f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y}))$$
$$\leq \ell(\bar{y}) + \frac{\beta}{2}|u|^2 + \nabla \bar{v}(\bar{y})(f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla v(\bar{y}))$$

Consistency estimates

For the ease of the presentation we write $\bar{y} = y(\cdot; y_0, \bar{v})$, $\bar{u} = -\frac{1}{\beta} B^\top \nabla \bar{v}(\bar{y})$ and $u = -\frac{1}{\beta} B^\top \nabla v(\bar{y})$.

- We have

$$\ell(\bar{y}) + \frac{\beta}{2}|\bar{u}|^2 + \nabla \bar{v}(\bar{y})(f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y}))$$
$$\leq \ell(\bar{y}) + \frac{\beta}{2}|u|^2 + \nabla \bar{v}(\bar{y})(f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla v(\bar{y}))$$

- Using that v is a subsolution of HJB we get

$$\ell(\bar{y}) + \frac{\beta}{2}|\bar{u}|^2 + \nabla \bar{v}(\bar{y})^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y}))$$
$$\leq -\nabla v(\bar{y})^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla v(\bar{y}))$$
$$+\nabla \bar{v}(\bar{y})^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla v(\bar{y}))$$

Consistency estimates

- Rearranging the terms:

$$\ell(\bar{y}) + \frac{\beta}{2}|\bar{u}|^2 + \nabla \bar{v}(\bar{y})^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y}))$$

$$\leq (\nabla \bar{v}(\bar{y}) - \nabla v(\bar{y}))^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla v(\bar{y}))$$

Consistency estimates

- Rearranging the terms:

$$\ell(\bar{y}) + \frac{\beta}{2}|\bar{u}|^2 + \nabla \bar{v}(\bar{y})^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y}))$$

$$\leq (\nabla \bar{v}(\bar{y}) - \nabla v(\bar{y}))^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla v(\bar{y}))$$

- Adding and subtracting

$$(\nabla \bar{v}(\bar{y}) - \nabla v(\bar{y}))^\top \cdot \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y})$$

in the left hand side we get

$$\ell(\bar{y}) + \frac{\beta}{2}|\bar{u}|^2 + \nabla \bar{v}(\bar{y})^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y}))$$

$$\leq \frac{1}{\beta} |B^\top (\nabla v(\bar{y}) - \nabla \bar{v}(\bar{y}))|^2$$

$$+ (\nabla \bar{v}(\bar{y}) - \nabla v(\bar{y}))^\top \left(f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y}) \right)$$

Consistency estimates

- Rearranging the terms we obtain

$$\ell(\bar{y}) + \frac{\beta}{2}|\bar{u}|^2 + \nabla v(\bar{y})^\top (f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y}))$$
$$\leq \frac{1}{\beta}|B^\top(\nabla v(\bar{y}) - \nabla \bar{v}(\bar{y}))|^2$$

Consistency estimates

- Rearranging the terms we obtain

$$\ell(\bar{y}) + \frac{\beta}{2}|\bar{u}|^2 + \nabla v(\bar{y})^\top \left(f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y})\right)$$
$$\leq \frac{1}{\beta}|B^\top(\nabla v(\bar{y}) - \nabla \bar{v}(\bar{y}))|^2$$

- recalling that

$$\frac{d}{dt}\bar{y} = f(\bar{y}) - \frac{1}{\beta} BB^\top \nabla \bar{v}(\bar{y})$$

and integrating we arrive at

$$\mathcal{V}_T(y_0; \bar{v}) + v(\bar{y}(T)) - v(y_0)$$
$$\leq \int_0^T \frac{1}{\beta}|B^\top(\nabla v(\bar{y}) - \nabla \bar{v}(\bar{y}))|^2 dt$$

Consistency estimates

If the value function is not C^1, but semi-concave we can still do something. We recall that a function $v \in C(\overline{\Omega})$ is C-semiconcave if and only if for all $d \in \mathbb{R}^d$

$$d\nabla v^2(x)d \leq C|d|^2$$

in the sense of the distributions.

Consistency estimates

Lemma 6

Let $\omega \Subset \Omega$, *let* $\bar{v} \in C^2(\overline{\Omega}; \mathbb{R}^M)$ *be such that there exist a positive constant* $C > 0$ *satisfying*

$$\frac{1}{\beta} tr(BB^\top \nabla^2 \bar{v}) \leq C \text{ in } \Omega_1,$$

and $\bar{v}$ *is stable. Then for all* $\phi \in C(\overline{\Omega}; \mathbb{R}^+)$

$$\int_\omega \int_0^T \phi(y(t; y_0, \bar{v}))dtdy_0 \leq \frac{e^{KT} - 1}{K} \int_\Omega \phi(z)dz$$

holds, where

$$K = d\left(\frac{|B|^2}{\beta} C + \|f\|_{Lip(\overline{\Omega};\mathbb{R}^d)}\right).$$

Consistency estimates

For proving this result we just need to use a change of variables an Fubini Theorem:

- Using the transformation $z = y(t; y_0, \bar{v})$, we obtain for a fixed time $t \in [0, T]$ that

$$\int_{\omega} \phi(y(t; y_0, \bar{v}))dy_0 = \int_{y(t;\omega)} \phi(z) \frac{dz}{|D_{y_0} y(t; y^{-}1(t; z, \phi), \phi)|}$$

Consistency estimates

For proving this result we just need to use a change of variables an Fubini Theorem:

- Using the transformation $z = y(t; y_0, \bar{v})$, we obtain for a fixed time $t \in [0, T]$ that

$$\int_{\omega} \phi(y(t; y_0, \bar{v}))dy_0 = \int_{y(t;\omega)} \phi(z) \frac{dz}{|D_{y_0} y(t; y^-1(t; z, \phi), \phi)|}$$

- The Jacobi formula implies that

$$|D_{y_0} y(t; y_0, \phi)| =$$
$$\exp\left(\int_0^t tr\left(Df(y(s; y_0, \phi)) - \frac{1}{\beta} BB^\top \nabla \bar{v}(y(s; y_0, \phi))\right) ds\right)$$
$$\geq \exp\left(-t\left(d\|f\|_{Lip(\overline{\Omega};\mathbb{R}^d)} + \frac{d|B|^2}{\beta}\|\nabla^2 \bar{v}\|_{L^\infty(\Omega;\mathbb{R}^{d\times d})}\right)\right)$$

Plugging this into the change of variables formula we obtain

$$\int_{\omega} \phi(y(t; y_0, \bar{v}))dy_0 \leq \exp(tK) \int_{y(t;\omega,\phi)} \phi(z)dz \ .$$

Since $y(t;\omega,\phi) \subset \Omega$ for all $t \in [0,T]$ we conclude that

$$\int_0^T \int_{\omega} \phi(y(t; y_0, \bar{v}))dy_0 dt \leq \frac{\exp(KT) - 1}{K} \int_{\Omega} \phi(z)dz$$

Theorem 7

Let $v \in Lip(\overline{\Omega_1})$ *be a super-solution of* (HJB) *in* Ω*, and let* $\bar{v} \in C^2(\overline{\Omega})$ *be such that for some constant* $C > 0$

$$\frac{1}{\beta} tr(BB^\top \nabla^2 \bar{v}(v)) \leq C \text{ for all } y \in \Omega.$$

Let $\omega \Subset \Omega$ *and Hypothesis 1 holds true with* $v = \bar{v}$*, then the following inequality holds*

$$\|(\mathcal{V}_T(\cdot;\bar{v}) + v(y(T;\cdot,\bar{v})) - v)^+\|_{L^1(\omega)}$$
$$\leq \frac{|B|^2}{\beta}\left(\frac{e^{KT}-1}{K}\right)\|\nabla v - \nabla \bar{v}\|^2_{L^2(\Omega;\mathbb{R}^d)}$$

Consistency estimates

1. The previous result states that if we construct a semiconcave approxmiation $V_n \in C^2(\overline{\Omega})$ of V such that ∇V_n converges to ∇V in $L^2(\Omega; \mathbb{R}^d)$ and it is stable, then we have that it is consistent.

Consistency estimates

1. The previous result states that if we construct a semiconcave approxmiation $V_n \in C^2(\overline{\Omega})$ of V such that ∇V_n converges to ∇V in $L^2(\Omega; \mathbb{R}^d)$ and it is stable, then we have that it is consistent.
2. This justifies the name of the concept (consistency), since if the feedback approximation converges in $H^1(\Omega)$, then the cost function also converges to the value function.

Consistency estimates

1. The previous result states that if we construct a semiconcave approxmiation $V_n \in C^2(\overline{\Omega})$ of V such that ∇V_n converges to ∇V in $L^2(\Omega; \mathbb{R}^d)$ and it is stable, then we have that it is consistent.
2. This justifies the name of the concept (consistency), since if the feedback approximation converges in $H^1(\Omega)$, then the cost function also converges to the value function.
3. In particular, we can use these result to prove that the mollification of V provides an stable and consistent feedback law.

Consistency estimates

1. The previous result states that if we construct a semiconcave approxmiation $V_n \in C^2(\overline{\Omega})$ of V such that ∇V_n converges to ∇V in $L^2(\Omega; \mathbb{R}^d)$ and it is stable, then we have that it is consistent.
2. This justifies the name of the concept (consistency), since if the feedback approximation converges in $H^1(\Omega)$, then the cost function also converges to the value function.
3. In particular, we can use these result to prove that the mollification of V provides an stable and consistent feedback law.
4. What about the L^1 convergence of the cost?

Consistency estimates

To prove the convergence of a consistent sequence of feedback laws we split the error of the cost functional in two terms:

$$\int_\omega |\mathcal{V}_{T_n}(y_0; v_n) - V(y_0)| dy_0 =$$

$$\int_\omega (\mathcal{V}_{T_n}(y_0; v_n) - V(y_0))^+ dy_0 + \int_\omega (V(y_0) - \mathcal{V}_{T_n}(y_0; v_n))^+ dy_0$$

Consistency estimates

To prove the convergence of a consistent sequence of feedback laws we split the error of the cost functional in two terms:

$$\int_\omega |\mathcal{V}_{T_n}(y_0; v_n) - V(y_0)| dy_0 =$$

$$\int_\omega (\mathcal{V}_{T_n}(y_0; v_n) - V(y_0))^+ dy_0 + \int_\omega (V(y_0) - \mathcal{V}_{T_n}(y_0; v_n))^+ dy_0$$

For the first term in the right-hand side we have, from the monotonicity of the positive part that:

$$\int_\omega (\mathcal{V}_{T_n}(y_0; v_n) - V(y_0))^+ dy_0 \leq$$

$$\int_\omega (\mathcal{V}_{T_n}(y_0; v_n) + V(y(T_n; y_0, \phi_n)) - V(y_0))^+ dy_0$$

which converges to 0 from the consistency of the sequence.

For the second term we have the following result:

Proposition 1

Let $v_n \in C^2(\overline{\Omega})$*. Then for any* $0 < T_n \to \infty$

$$\lim_{n\to\infty} (V(y_0) - \mathcal{V}_{T_n}(y_0; v_n))^+ = 0$$

for all $y_0 \in \Omega$.

Consistency estimates

We prove the porposition by contradiction:

- There exists $y_{0,n} \to y_0$ such that

$$V(y_{0,n}) - \mathcal{V}_{T_n}(y_{0,n}; v_n) > \varepsilon > 0.$$

Consistency estimates

We prove the porposition by contradiction:

- There exists $y_{0,n} \to y_0$ such that

$$V(y_{0,n}) - \mathcal{V}_{T_n}(y_{0,n}; v_n) > \varepsilon > 0.$$

- This means that $u_n = \frac{1}{\beta}\nabla v_n(y(\cdot; y_{0,n}, v_n))$ satisfies

$$\mathcal{V}_{T_n}(y_0; v_n) \leq V(y_{0,n}) - \varepsilon$$

Consistency estimates

We prove the porposition by contradiction:

- There exists $y_{0,n} \to y_0$ such that

$$V(y_{0,n}) - \mathcal{V}_{T_n}(y_{0,n}; v_n) > \varepsilon > 0.$$

- This means that $u_n = \frac{1}{\beta}\nabla v_n(y(\cdot; y_{0,n}, v_n))$ satisfies

$$\mathcal{V}_{T_n}(y_0; v_n) \leq V(y_{0,n}) - \varepsilon$$

- By compactness, we have that u_n^* converges to u^* weakly in $L^2_{loc}((0,\infty);\mathbb{R}^m)$ and $y(;y_{0,n},u_n^*)$ strongly in $C_{loc}([0,\infty);\mathbb{R}^d)$ to $y(;y_0,u^*)$.

Consistency estimates

We prove the porposition by contradiction:

- There exists $y_{0,n} \to y_0$ such that

$$V(y_{0,n}) - \mathcal{V}_{T_n}(y_{0,n}; v_n) > \varepsilon > 0.$$

- This means that $u_n = \frac{1}{\beta}\nabla v_n(y(\cdot; y_{0,n}, v_n))$ satisfies

$$\mathcal{V}_{T_n}(y_0; v_n) \leq V(y_{0,n}) - \varepsilon$$

- By compactness, we have that u_n^* converges to u^* weakly in $L^2_{loc}((0,\infty);\mathbb{R}^m)$ and $y(; y_{0,n}, u_n^*)$ strongly in $C_{loc}([0,\infty);\mathbb{R}^d)$ to $y(; y_0, u^*)$.
- This implies that

$$J(u^*; y_0) \leq V(y_0) - \varepsilon < V(y_0)$$

which contradicts definition of V.

Consistency estimates

This implies the convergence of a consistent and stable feedback law:

Theorem 8

Let $v_n \in C^2(\overline{\Omega})$ *be a stable in* $\omega \subset \Omega$ *and consistent sequence of feedback laws for* $0 < T_n \to \infty$. *Then*

$$\lim_{n\to\infty} \|\mathcal{V}_{T_n}(\cdot; v_n) - V\|_{L^1(\omega)} = 0.$$

Consistency estimates

Corollary 1

Let us assume that V is semi-concave and V satisfies Hypothesis 2. Then $V_\varepsilon = V * \rho_\varepsilon$ (mollification of V) is stable and consistent. Consequently, there exists $0 < T_\varepsilon \to \infty$ as $\varepsilon \to 0^+$ and

$$\lim_{\varepsilon \to 0^+} \|\mathcal{V}_{T_\varepsilon}(\cdot; v_\varepsilon) - V\|_{L^1(\omega)} = 0.$$

Since mollification preserves semi-concavity, this proves that, at least, there is one stable and consistent Feedback law.

Convergence of the AFLS method

How we use this for the convergence of the methods?

- We need a universal approximation property for the sequence of settings:

$$\text{for all } g \in C^2(\overline{\Omega}) : \ \exists\, \theta_n \in \Theta_n,\ \alpha_n > 0,\ \lim_{n\to\infty} \|v_n(\theta_n) - V\|_{C^2(\overline{\Omega})} = 0.$$

Convergence of the AFLS method

How we use this for the convergence of the methods?

- We need a universal approximation property for the sequence of settings:

$$\text{for all } g \in C^2(\overline{\Omega}) : \; \exists\, \theta_n \in \Theta_n, \; \alpha_n > 0, \; \lim_{n\to\infty} \|v_n(\theta_n) - V\|_{C^2(\overline{\Omega})} = 0.$$

- This directly proves that the solutions of the AFLS method satisfies

$$\lim_{n\to\infty} \int_\Omega \mathcal{V}_{T_n}(y_0; v_n(\theta_n^*))dy_0 + \alpha_n \mathcal{P}_n(\theta_n^*) = 0$$

which implies that

$$\lim_{n\to\infty} \int_\Omega |\mathcal{V}_{T_n}(y_0; v_n(\theta_n^*)) - V(y_0)|dy_0 = 0$$

Convergence of the AFLS method

Theorem 9

- *Hypothesis 1* $\forall g \in C^1(\overline{\Omega})\ \exists \theta_n \in \Theta_n$ *satisfying* $\lim_{n\to\infty} \|g - v_n(\theta_n)\|_{C^1(\overline{\Omega})} + \gamma_n \cdot \mathcal{P}_n(\theta_n) = 0,$

Convergence of the AFLS method

Theorem 9

- *Hypothesis 1* $\forall g \in C^1(\overline{\Omega})$ $\exists \theta_n \in \Theta_n$ *satisfying* $\lim_{n\to\infty} \|g - v_n(\theta_n)\|_{C^1(\overline{\Omega})} + \gamma_n \cdot \mathcal{P}_n(\theta_n) = 0$,
- *Hypothesis 2 (key):* $\exists\ 0 < T_n \to \infty$ *and a consistent sequence* $V_n \in C^2$ *feedback-laws with* $p = 1$, *i.e.,*

$$\lim_{n\to\infty} \|\mathcal{V}_{T_n}(\cdot, u_{V_n}) - V\|_{L^1(\omega)} = 0,$$

satisfying $\overline{y([0, T_n]; \omega, u_{V_n})} \subset \Omega$,

Theorem 9

- *Hypothesis 1* $\forall g \in C^1(\overline{\Omega})$ $\exists \theta_n \in \Theta_n$ *satisfying* $\lim_{n\to\infty} \|g - v_n(\theta_n)\|_{C^1(\overline{\Omega})} + \gamma_n \cdot \mathcal{P}_n(\theta_n) = 0$,
- *Hypothesis 2 (key):* $\exists\ 0 < T_n \to \infty$ *and a consistent sequence* $V_n \in C^2$ *feedback-laws with* $p = 1$, *i.e.*,

$$\lim_{n\to\infty} \|\mathcal{V}_{T_n}(\cdot, u_{V_n}) - V\|_{L^1(\omega)} = 0,$$

  satisfying $\overline{y([0, T_n]; \omega, u_{V_n})} \subset \Omega$,
- *the result:* *then there exists a* $k : \mathbb{N} \mapsto \mathbb{N}$ *and* γ_n *such that the method converges with settings* $S_n = (\Theta_{k(n)}, \mathcal{P}_{k(n)}, v_{k(n)})$ *and penalty* $\alpha = \gamma_n$.

Convergence of the AFLS method

The meaning of each hypothesis:

- The first conditions can be seen as a Universal Approxmiation Property: we need to have a setting which is able to express sufficiently well any $C^2(\overline{\Omega})$ function.

The meaning of each hypothesis:

- The first conditions can be seen as a Universal Approxmiation Property: we need to have a setting which is able to express sufficiently well any $C^2(\overline{\Omega})$ function.
- The second one is the Stability-Consistency property: It must exists a stable and consistent sequence of Feedback-laws. This is ensured if V is semiconcave and is stable.

Convergence of the AFLS method

The meaning of each hypothesis:

- The first conditions can be seen as a Universal Approxmiation Property: we need to have a setting which is able to express sufficiently well any $C^2(\overline{\Omega})$ function.
- The second one is the Stability-Consistency property: It must exists a stable and consistent sequence of Feedback-laws. This is ensured if V is semiconcave and is stable.
- But what can we say about the RAT method?

Convergence Regression along trajectories

- There exists a Banach space $(\Theta, \|\cdot\|_\Theta)$ such that $V \in \Theta \cap C^2(\Omega)$ and Θ is compactly embedded in $C^1(\overline{\Omega})$.

- There exists a Banach space $(\Theta, \|\cdot\|_\Theta)$ such that $V \in \Theta \cap C^2(\Omega)$ and Θ is compactly embedded in $C^1(\overline{\Omega})$.
- $\overline{y([0,\infty); u_V, \omega)} \subset \Omega$.

Convergence Regression along trajectories

- There exists a Banach space $(\Theta, \|\cdot\|_\Theta)$ such that $V \in \Theta \cap C^2(\Omega)$ and Θ is compactly embedded in $C^1(\overline{\Omega})$.
- $\overline{y([0,\infty); u_V, \omega)} \subset \Omega$.
- $\forall g \in \Theta$ there exists a sequence $\theta_n \in \Theta_n$ satisfying

$$\lim_{n\to\infty} \|v_n(\theta_n) - g\|_\Theta = 0, \ \sup_{n\in\mathbb{N}} \mathcal{P}_n(\theta_n) < \infty.$$

Convergence Regression along trajectories

- There exists a Banach space $(\Theta, \|\cdot\|_\Theta)$ such that $V \in \Theta \cap C^2(\Omega)$ and Θ is compactly embedded in $C^1(\overline{\Omega})$.
- $\overline{y([0,\infty); u_V, \omega)} \subset \Omega$.
- $\forall g \in \Theta$ there exists a sequence $\theta_n \in \Theta_n$ satisfying

$$\lim_{n\to\infty} \|v_n(\theta_n) - g\|_\Theta = 0, \ \sup_{n\in\mathbb{N}} \mathcal{P}_n(\theta_n) < \infty.$$

- there exist constants $C > 0$ and $\sigma > 0$ such that for all $n \in \mathbb{N}$ and $\theta \in \Theta_n$ it holds that

$$\|v_n(\theta)\|_\Theta \leq C\mathcal{P}_n(\theta)^\sigma.$$

Convergence Regression along trajectories

- Under the previous there exist a sequence $\theta_n \in \Theta$ and $C > 0$ such that

$$\|B^\top(\nabla v_n(\theta_n) - \nabla V)\|_{C(\overline{\Omega};\mathbb{R}^m)} = 0 \text{ and } \sup_{n\in\mathbb{N}} \mathcal{P}_n(\theta_n) \leq C$$

- Under the previous there exist a sequence $\theta_n \in \Theta$ and $C > 0$ such that

$$\|B^\top(\nabla v_n(\theta_n) - \nabla V)\|_{C(\overline{\Omega};\mathbb{R}^m)} = 0 \text{ and } \sup_{n\in\mathbb{N}} \mathcal{P}_n(\theta_n) \leq C$$

- The result: Setting $\gamma_n = \|B^\top(\nabla v_n(\theta_n) - \nabla V)\|^2_{L^2(\Omega;\mathbb{R}^m)}$ we have that the solution of the regression problem θ_n^* with $S = (\Theta_n, \mathcal{P}_n, v_n)$ and penalty $\alpha = \gamma_n$ satisfies

$$\lim_{n\to\infty} \|\mathcal{V}_{T_n}(\cdot, u_{v_n(\theta_n)}) - V\|_{L^\infty(\omega)} = 0.$$

Comparison

- The AFLS method will converge in $L^1(\omega)$ if the value function is smooth and the optimal system is stable.
- The RAT method convergences in $L^\infty(\omega)$ if the value function is smooth enough, but it is unknown what happen if it is only semiconcave.

Tabla de contenidos

Introduction of the example

- Objective: To avoid the ball $B(z, \sigma) \subset \mathbb{R}^2$

Introduction of the example

- Objective: To avoid the ball $B(z, \sigma) \subset \mathbb{R}^2$
- We can control the velocity: $y' = u$.

- Objective: To avoid the ball $B(z, \sigma) \subset \mathbb{R}^2$
- We can control the velocity: $y' = u$.
- $\alpha > 0$

$$\ell_\alpha(y) = \frac{1}{2}|y|^2 \left(1 + \alpha\psi\left(\frac{|y-z|}{\sigma}\right)\right)$$

Introduction of the example

- Objective: To avoid the ball $B(z,\sigma) \subset \mathbb{R}^2$
- We can control the velocity: $y' = u$.
- $\alpha > 0$

$$\ell_\alpha(y) = \frac{1}{2}|y|^2 \left(1 + \alpha\psi\left(\frac{|y-z|}{\sigma}\right)\right)$$

-

$$\psi(s) = \begin{cases} \exp\left(-\frac{1}{1-s^2}\right) & \text{if } |s| < 1 \\ 0 & \text{if } |s| \geq 1 \end{cases},$$

Introduction of the example

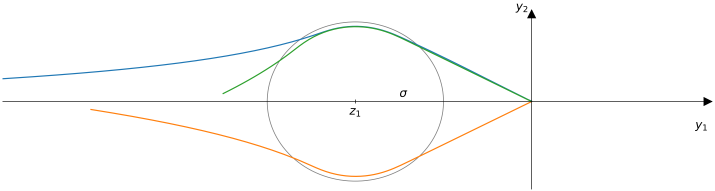

Figure: Obstacle problem for α large.

Introduction of the example

For $\alpha \in [0, \infty)$, we consider the following control problem

$$\min_{\substack{u \in L^2((0,\infty);\mathbb{R}^2), \\ y' = u,\ y(0) = y_0}} \int_0^\infty \ell_\alpha(y(t))dt + \frac{\beta}{2} \int_0^\infty |u(t)|^2 dt.$$

The value function of this problem is denoted by V_α.

Value function

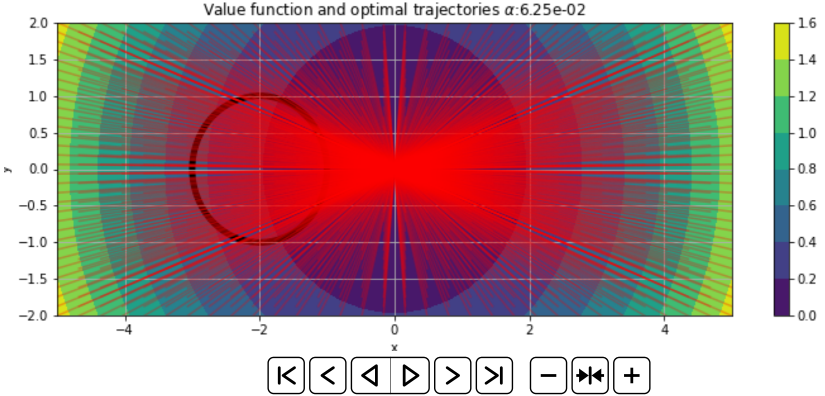

Properties

There are some important properties of this family of problems

1. The value function of this problem is locally Semiconcave.

Properties

There are some important properties of this family of problems

1. The value function of this problem is locally Semiconcave.
2. The optimal trajectories are uniformly asymptotically stable.

Properties

There are some important properties of this family of problems

1. The value function of this problem is locally Semiconcave.
2. The optimal trajectories are uniformly asymptotically stable.
3. $\exists \alpha_s$ such that for all $\alpha < \alpha_s$, V_α is C^∞.

Properties

There are some important properties of this family of problems

1. The value function of this problem is locally Semiconcave.
2. The optimal trajectories are uniformly asymptotically stable.
3. $\exists \alpha_s$ such that for all $\alpha < \alpha_s$, V_α is C^∞.
4. $\exists \alpha_{ns}$ such that for all $\alpha > \alpha_{ns}$, V_α is not differentiable.

Properties

There are some important properties of this family of problems

1. The value function of this problem is locally Semiconcave.
2. The optimal trajectories are uniformly asymptotically stable.
3. $\exists \alpha_s$ such that for all $\alpha < \alpha_s$, V_α is C^∞.
4. $\exists \alpha_{ns}$ such that for all $\alpha > \alpha_{ns}$, V_α is not differentiable.
5. In the following we compere the two approaches in order to observe in practice the relevance of the smoothness of the value function.

Setting

- For the numerical experiments we have consider $\omega = (-5, 5) \times (-2, 2)$.

Setting

- For the numerical experiments we have consider $\omega = (-5, 5) \times (-2, 2)$.
- Different α are considered ranging from the smooth to the noon-smooth cases.

Setting

- For the numerical experiments we have consider $\omega = (-5, 5) \times (-2, 2)$.
- Different α are considered ranging from the smooth to the noon-smooth cases.
- For the parametrization of the feedback-laws we use an orthogonal polynomial basis of $H^1_{mix}(\omega) = H^1(-5, 5) \otimes H^1(-2, 2)$ of degree 20.

Discretization

- The state-space integral is discretize by a regular grid of 10×10. We call this set of initial conditions the training set.

Discretization

- The state-space integral is discretize by a regular grid of 10×10. We call this set of initial conditions the training set.
- We also consider a finer grid of 20×20. We name this grid as the test set.

Discretization

- The state-space integral is discretize by a regular grid of 10×10. We call this set of initial conditions the training set.
- We also consider a finer grid of 20×20. We name this grid as the test set.
- The time integral is approximated by the trapezoid rule and the differential equations are discretized by Crank Nicholson method.

Discretization

- The state-space integral is discretize by a regular grid of 10×10. We call this set of initial conditions the training set.
- We also consider a finer grid of 20×20. We name this grid as the test set.
- The time integral is approximated by the trapezoid rule and the differential equations are discretized by Crank Nicholson method.
- The problem is solved by using the training set and then the performance of approaches is compere in the test set.

Performance comparison

In order to compere the performance of the two approaches we consider the following error measures in both the training and test sets.

- $$NMAE_V = \frac{\sum_{i=1}^{N} |\mathcal{V}(y_0^i, u_{v(\theta)}) - V(y_0^i)|}{\sum_{i=1}^{N} V(y_0^i)}$$
- $$NRMSE_u = \frac{\sum_{i=1}^{N} \|u_{v(\theta)}(y(\cdot; y_0^i, u_{v(\theta)})) - u_i^*\|_{L^2((0,T);\mathbb{R}^2)}}{\sum_{i=1}^{N} \|u_i^*\|_{L^2((0,T);\mathbb{R}^2)}}$$

where u_i^* are the optimal controls.

Performance comparison $NMAE_V$

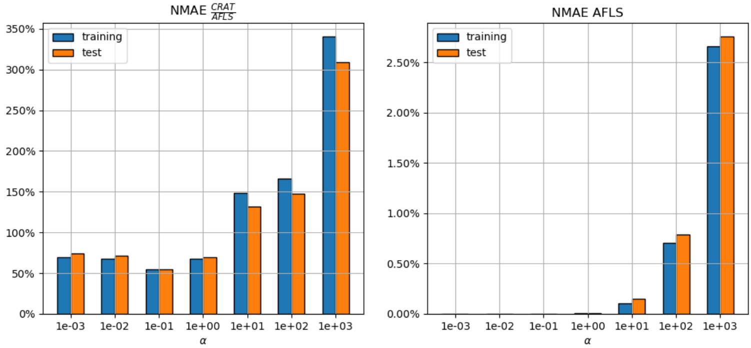

Figure: $NMAE_V$

Performance comparison $NMRSE_u$

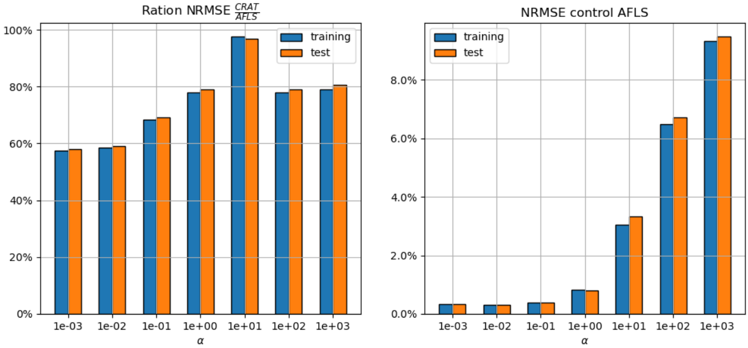

Figure: $NRMSE_u$

Tabla de contenidos

Representation Theorem

Theorem

A function $v \in C(\overline{\Omega})$ $L-$Lipschitz is $C-$semi- continuous over $\overline{\Omega}$ if and only if there exists a family $\{\phi_i\}_{i\in\mathcal{I}} \subset C^2(\overline{\Omega})$ satisfying

$$\sup_{i\in I} \|\nabla^2 \phi_i\|_{C(\overline{\Omega};\mathbb{R}^{d\times d})} \leq C \text{ and } \sup_{i\in I} \|\nabla \phi_i\|_{C(\overline{\Omega};\mathbb{R}^{d})} \leq L.$$

and

$$v(x) = \min_{i\in I} \phi_i(x), \ \forall x \in \Omega.$$

Remark: We can assume that I is countable (but not necessarily finite).

Example

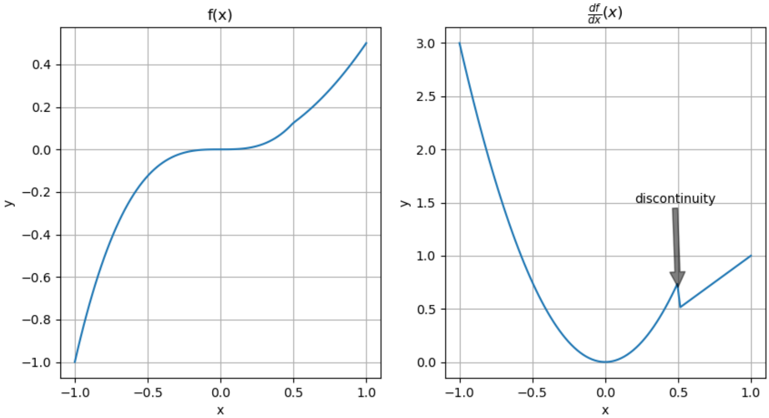

Figure: For example $f(x) = min\left(x^3, \frac{1}{2}x^2\right)$.

Tabla de contenidos

Parametrization

The idea is:

- truncation:
  $\psi_n(a) = \min_{i\in\{1,\dots,n\}} a_i \Rightarrow \psi_{i+1}(a) = a_{i+1} - (a_{i+1} - \psi_i(a))_+$

The idea is:

- truncation:
  $\psi_n(a) = \min_{i\in\{1,\dots,n\}} a_i \Rightarrow \psi_{i+1}(a) = a_{i+1} - (a_{i+1} - \psi_i(a))_+$
- smoothing $\varepsilon > 0$ $\psi_{1,\varepsilon}(a) = a_1$, $\psi_{i+1,\varepsilon}(a) = a_{i+1} - g_\varepsilon(a_{i+1} - \psi_{i,\varepsilon}(a))$.

Parametrization

The idea is:

- truncation:
  $\psi_n(a) = \min_{i \in \{1,\dots,n\}} a_i \Rightarrow \psi_{i+1}(a) = a_{i+1} - (a_{i+1} - \psi_i(a))_+$
- smoothing $\varepsilon > 0$ $\psi_{1,\varepsilon}(a) = a_1$, $\psi_{i+1,\varepsilon}(a) = a_{i+1} - g_\varepsilon(a_{i+1} - \psi_{i,\varepsilon}(a))$.
- $g_\varepsilon(s) = \begin{cases} 0 & \text{if } s < 0 \\ \frac{1}{2\varepsilon} s^2 & \text{if } s \in [0, \varepsilon) \\ s - \frac{\varepsilon}{2} & \text{if } s \geq \varepsilon \end{cases}$

Parametrization

The idea is:

- truncation:
  $\psi_n(a) = \min_{i\in\{1,\dots,n\}} a_i \Rightarrow \psi_{i+1}(a) = a_{i+1} - (a_{i+1} - \psi_i(a))_+$
- smoothing $\varepsilon > 0$ $\psi_{1,\varepsilon}(a) = a_1$, $\psi_{i+1,\varepsilon}(a) = a_{i+1} - g_\varepsilon(a_{i+1} - \psi_{i,\varepsilon}(a))$.
- $g_\varepsilon(s) = \begin{cases} 0 & \text{if } s < 0 \\ \frac{1}{2\varepsilon}s^2 & \text{if } s \in [0,\varepsilon) \\ s - \frac{\varepsilon}{2} & \text{if } s \geq \varepsilon \end{cases}$
- smoothing ϕ_i: $\xi : \Theta \mapsto C^2(\overline{\Omega})$ parametrization. $\phi_i \approx \xi(\theta_i)$

Parametrization

The idea is:

- truncation:
  $\psi_n(a) = \min_{i\in\{1,\dots,n\}} a_i \Rightarrow \psi_{i+1}(a) = a_{i+1} - (a_{i+1} - \psi_i(a))_+$
- smoothing $\varepsilon > 0$ $\psi_{1,\varepsilon}(a) = a_1$, $\psi_{i+1,\varepsilon}(a) = a_{i+1} - g_\varepsilon(a_{i+1} - \psi_{i,\varepsilon}(a))$.
- $g_\varepsilon(s) = \begin{cases} 0 & \text{if } s < 0 \\ \frac{1}{2\varepsilon}s^2 & \text{if } s \in [0,\varepsilon) \\ s - \frac{\varepsilon}{2} & \text{if } s \geq \varepsilon \end{cases}$
- smoothing ϕ_i: $\xi : \Theta \mapsto C^2(\overline{\Omega})$ parametrization. $\phi_i \approx \xi(\theta_i)$
- $V \approx \tilde{\psi}_{n,\varepsilon}(\theta) = \psi_{n,\varepsilon}(\xi(\theta_1),\dots,\xi(\theta_n))$ for some $\theta \in \Theta^n$.

- (Penalty) $\mathcal{P}_n : \Theta^n \mapsto \mathbb{R}$ be a continuous and coercive function together with a penalty parameter $\alpha > 0$.
- (loss function) $\mathcal{J}(v) = \dfrac{1}{|\Omega|} \displaystyle\int_\Omega \left(|V - v| + \sum_{i=1}^{d} \left| \frac{\partial V}{\partial x_i} - \frac{\partial v}{\partial x_i} \right| \right) dx$

Learning problem

For $\varepsilon > 0$, $\alpha > 0$, and a setting $S = (\Theta, \xi, \mathcal{P}_n)$, we want to solve:

$$\min_{\theta \in \Theta^n} \mathcal{J}(\tilde{\psi}_{n,\varepsilon}(\theta)) + \alpha \mathcal{P}_n(\theta)$$

Performance? We provide conditions on a sequence of settings $S_m = (\Theta_m, \xi_m, \mathcal{P}_{n,m})$.

Approximability hypothesis

There exists $\mu > 0$ and $C_{SC} > 0$ independent of m such that

$$\sup_{i=1,\dots,n} \|\xi_m(\theta_i)\|_{C^2(\overline{\Omega})} \leq C_{SC}\mathcal{P}_{n,m}(\theta)^{\mu}, \forall\theta \in \Theta^n,$$

and there exists $\tilde{\theta}_m \in \Theta_m^n$ and $C_V > 0$ satisfying

$$\sup_{m\in\mathbb{N}} \mathcal{P}_{n,m}(\tilde{\theta}_m) \leq C_V,$$

and

$$\lim_{m\to\infty} \|\psi_n(\xi_m(\theta_{m,1}), \dots, \xi_m(\theta_{m,n})) - V\|_{C(\overline{\Omega})} = 0.$$

Convergence Theorem 1

- Approximability hypothesis.
- $\varepsilon_m > 0$, $\lim_{m\to\infty} \varepsilon_m = 0$, $\alpha_m > 0$:

$$\lim_{m\to\infty} \alpha_m + \frac{1}{\alpha_m}\left(\varepsilon_m + \mathcal{J}(\tilde{\psi}_{n,m,\varepsilon_m}(\tilde{\theta}_m))\right) = 0$$

- $\hat{\theta}_m \in \Theta_m^n$ a sequence of solutions of the learning problem with $S = S_m$.
- $V_m = \tilde{\psi}_{n,m,\varepsilon_m}(\hat{\theta}_m)$

Then V_m is uniformly semi-concave and Lipschitz, and

$$\lim_{m\to\infty} V_m = V \text{ in } C(\overline{\Omega}).$$

Convergence Theorem 2

- Hypotheses of Convergence Theorem 1.
- Stability hypothesis (Lyapunov function for V).

Then V_m is consistent: for every $\omega \Subset \Omega$ there exists $T_m > 0$ such that $\lim_{n\to\infty} T_m = \infty$,

$$\overline{y([0, T_m]; \omega, V_m)} \subset \Omega \text{ for all } m \in \mathbb{N}$$

and

$$\lim_{m\to\infty} \|\mathcal{V}_{T_m}(\cdot, V_m) - V\|_{L^1(\omega)} = 0.$$

Tabla de contenidos

Legendre Learning Problem

For $d > 1$ and $\alpha \in \mathbb{N}^d$:

- Legendre polynomials:
  $(n+1)P_{n+1}(x) = (2n+1)xP_{n+1}(x) - nP_{n-1}(x), \; P_0 = 1, \; P_1 = x.$

Legendre Learning Problem

For $d > 1$ and $\alpha \in \mathbb{N}^d$:

- Legendre polynomials:
  $(n+1)P_{n+1}(x) = (2n+1)xP_{n+1}(x) - nP_{n-1}(x),\ P_0 = 1,\ P_1 = x.$
- $P_\alpha(x) = \prod_{i=1}^{d} P_{\alpha_i}(x_i),\ x \in [-1,1]^d.$

Legendre Learning Problem

For $d > 1$ and $\alpha \in \mathbb{N}^d$:

- Legendre polynomials:
  $(n+1)P_{n+1}(x) = (2n+1)xP_{n+1}(x) - nP_{n-1}(x),\ P_0 = 1,\ P_1 = x.$
- $P_\alpha(x) = \prod_{i=1}^{d} P_{\alpha_i}(x_i),\ x \in [-1,1]^d.$
- Orthogonality in L^2: $\phi \in L^2((-1,1)^d)$

$$\sum_{\alpha \in \mathbb{N}^d} c_\alpha(\phi) P_\alpha,\ c_\alpha(\phi) = \left(\prod_{i=1} \frac{2\alpha_i + 1}{2} \right) \int_{(-1,1)^d} P_\alpha(x)\phi(x)dx.$$

Legendre Learning Problem

For $d > 1$ and $\alpha \in \mathbb{N}^d$:

- Legendre polynomials:
  $(n+1)P_{n+1}(x) = (2n+1)xP_{n+1}(x) - nP_{n-1}(x),\ P_0 = 1,\ P_1 = x.$
- $P_\alpha(x) = \prod_{i=1}^{d} P_{\alpha_i}(x_i),\ x \in [-1,1]^d.$
- Orthogonality in L^2: $\phi \in L^2((-1,1)^d)$

$$\sum_{\alpha \in \mathbb{N}^d} c_\alpha(\phi) P_\alpha,\ c_\alpha(\phi) = \left(\prod_{i=1} \frac{2\alpha_i + 1}{2} \right) \int_{(-1,1)^d} P_\alpha(x)\phi(x)dx.$$

- for $\gamma \leq \alpha$: $\|D_\gamma P_\alpha\|_{C([-1,1]^d)} = \prod_{i=1}^{d} \frac{(\alpha_i + \gamma_i)!}{2^{\gamma_i}\gamma_i!(\alpha_i - \gamma_i)!}.$

For $N \in \mathbb{N}$, let us now consider the hyperbolic cross set of indexes:

$$\Gamma_N := \{\alpha \in \mathbb{N}^d : \pi_{HC}(\alpha) \leq N\}$$

where for $\alpha \in \mathbb{N}^d$

$$\pi_{HC}(\alpha) = \prod_{i=1}^{d}(1 + \alpha_i).$$

Remark: It has been proved that the cardinality of this set does not increases exponentially with the dimension.

Let $v_N = \sum_{\alpha \in \Gamma_N} c_\alpha(v)$, then we have

$$\begin{aligned} \|v_N - v\|_{C^m([-1,1]^d)} &\leq \sum_{\pi_{HC}(\alpha) > N} |c_\alpha(v)| \|P_\alpha\|_{C^m([-1,1]^d)} \\ &\leq \frac{1}{N^{k+\frac{3}{4}}} \|v\|_{\tilde{H}^{2m+2+k}_{mix}((-1,1))} \end{aligned}$$

where

$$\|v\|^2_{\tilde{H}^r_{mix}((-1,1))} = \sum_{\alpha \in \mathbb{N}^d} \left(\prod_{i=1}^{d} \frac{2\alpha_i + 1}{2} \right)^{2r+1} c_\alpha(v)^2.$$

We consider the following setting.

- $\Theta_N = \mathbb{R}^{|\Gamma(N)|}$ with $\Gamma(N) = \{\alpha^i\}_{i=1}^{|\Gamma(N)|}$.
- $\xi_N(\theta) = \sum_{i=1}^{|\Gamma(N)|} P_{\alpha^i}\theta_i$.
- $\mathcal{P}_{n,N}(\theta^1, \ldots, \theta^n) = \sum_{i=1}^{n} \sum_{j=1}^{|\Gamma(N)|} |\theta_i^j| \|P_{\alpha^i}\|_{C^2([-1,1]^d)}$.

We have

$$\|\xi_N(\theta^j)\|_{C^2([-1,1]^d)} \leq \sum_{i=1}^{|\Gamma(N)|} |\theta_i^j| \|P_{\alpha^i}\|_{C^2([-1,1]^d)} \leq \mathcal{P}_{n,N}(\theta)$$

Remark: if $V = \min_{i=1,\ldots,n} \phi_i$ and $\{\phi_i\}_{i=1}^n \subset \tilde{H}_{mix}^{6+k}((-1,1)^d)$ for some $k \geq 0$ we can obtain an efficient representation of V!

Tabla de contenidos

- We only consider the case $\alpha = 10^{-10}$ for the sake of simplicity.
- Monte Carlo approximation, 10^4 samples (training set).
- 10^4 samples as test set.
- L^1 percentage error to asses the performance:

$$e_{L^1} = \sum_{y \in S} |V(y) - \hat{V}(y)| / \sum_{i \in S} |V(y)| \cdot 100$$

$$egrad_{L^1} = \sum_{y \in S} |\nabla V(y) - \hat{V}(y)|_1 / \sum_{i \in S} |\nabla V(y)|_1 \cdot 100$$

where S can be either the training or the test sets, $\hat{V}$ approximation and V target.

Exponential distance

For this example we consider a family of semi-concave functions:

$$v_d = \min_{i=1,\dots,d} \exp\left(-\frac{1}{2}|x - e_i|^2\right)$$

- e_i is the $i-$th element of the canonical basis.
- v_d is is semiconcave and globally Lipscthiz continuous.
- It is non differentiable in the set

$$D_d = \{x \in \mathbb{R}^d : |x - e_i| = |x - e_j| \text{ for } i \neq j\}.$$

Results

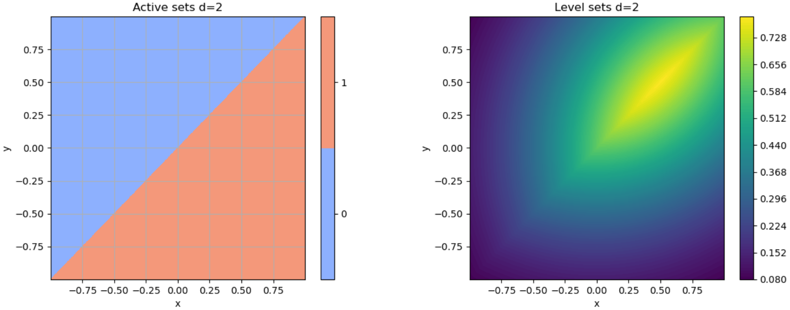

Figure: Active and level set sets $d = 2$.

Results

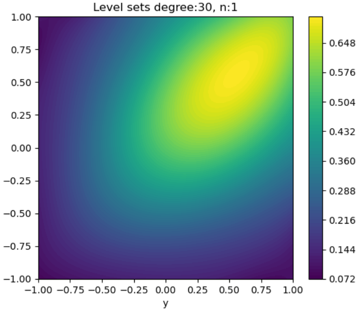

Figure: Level sets $d = 2$.

Results

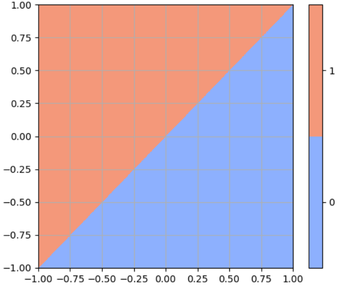

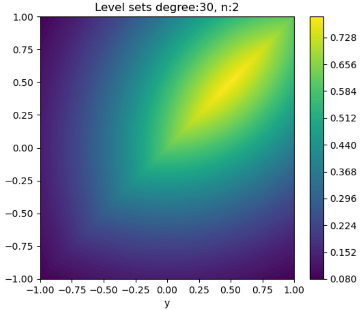

Figure: Active and level set sets $d = 2$, $n = 2$.

Results

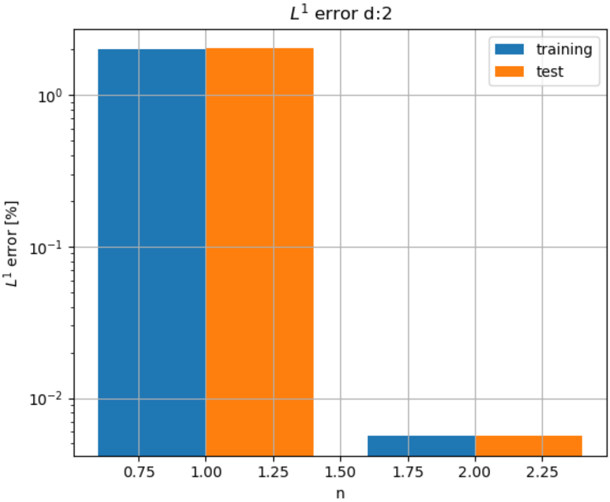

L^1 error d:2
training
test
10^0
10^{-1}
10^{-2}
L^1 error [%]
0.75
1.00
1.25
1.50
1.75
2.00
2.25
n

Results

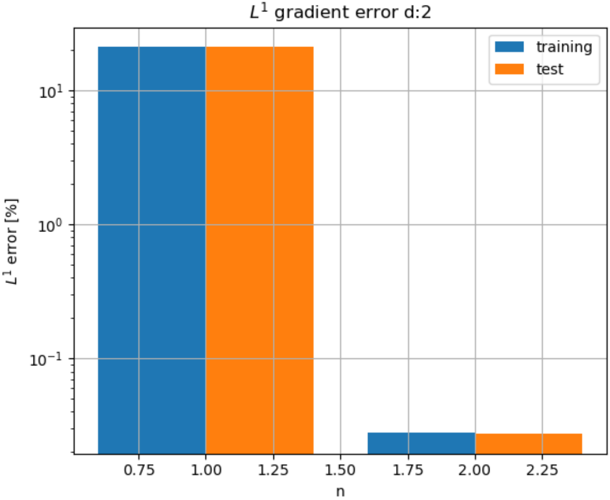

L^1 gradient error d:2
training
test
10^1
10^0
10^{-1}
L^1 error [%]
0.75
1.00
1.25
1.50
1.75
2.00
2.25
n

Obstacle problem

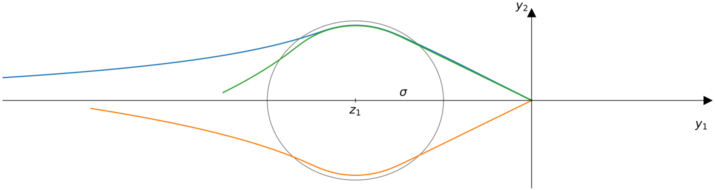

Figure: Obstacle problem for γ large.

- $\ell_\gamma(y) = \frac{|y|^2}{2}\left(1 + \gamma\psi\left(\frac{|y-z|}{\sigma}\right)\right)$
- $\psi(s) = \begin{cases} \exp\left(-\frac{1}{1-s^2}\right) & \text{if } |s| < 1 \\ 0 & \text{if } |s| \geq 1 \end{cases}$

Obstacle problem

$$V_\gamma(y_0) = \min_{\substack{u \in L^2((0,\infty);\mathbb{R}^2),\\ y' = u,\ y(0) = y_0}} \int_0^\infty \ell_\gamma(y(t))dt + \frac{\beta}{2}\int_0^\infty |u(t)|^2 dt.$$

Remark: V_γ is semiconcave but non-differentiable for $\gamma > 0$ large enough. In our experiments we take $\gamma = 100$

Results

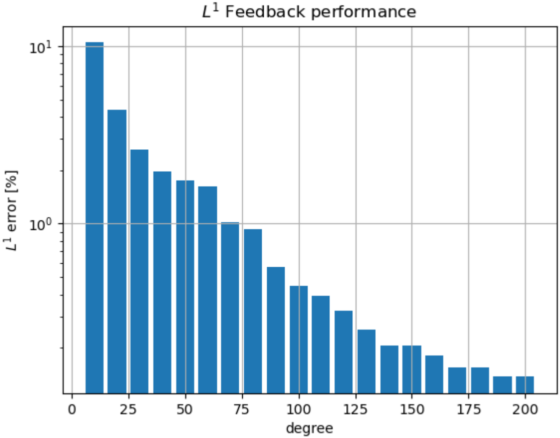

Figure: L^1 Feedback performance.

Optimal trajectories

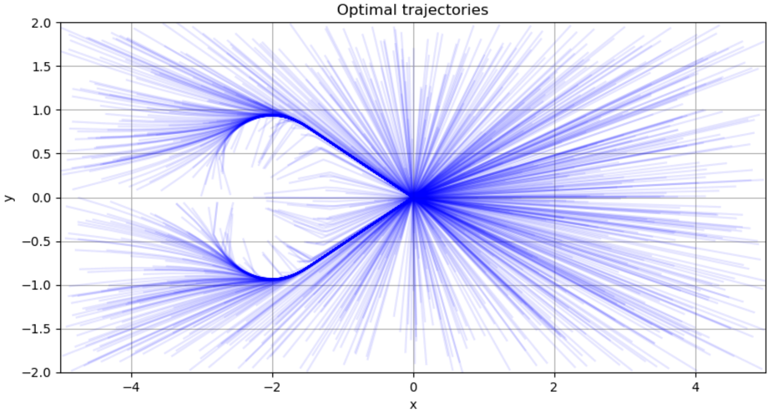

Figure: Optimal trajectories.

Learned trajectories

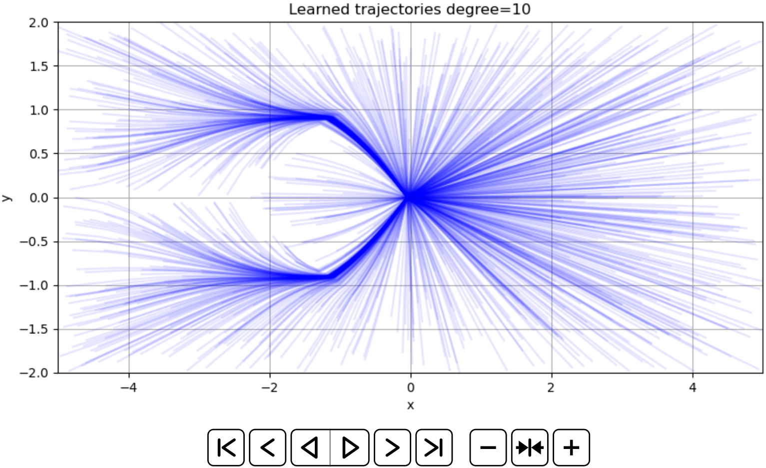

Learned active sets

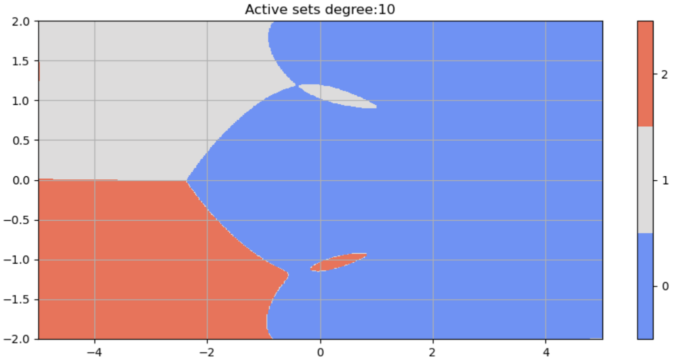

Connections with Inverse problems

- The most direct connection between control and inverse problems is Mortensen observers.
- There is a connections between inverse problems, semiconvex functions and control problems.
- As we will see, if we write the inverse problem as an optimziation problem, it is possible to use the value function of the inverse problems as in the case of inverse control.
- Further, in many ill-conditioned inverse problems the key is to find the correct regularization or penalty function, where semiconvex functions are relevant

Connections with Inverse problems

Let us consider the following inverse problem:

$$\min_{u \in U,\ z = A(u)} \frac{1}{2}\|B(z) - p\|^2 + G(u),$$

- $p \in \mathbb{R}^d$ is a vector of observations.
- $A : U \mapsto Z$ is the forward operator associated to the corresponding direct problem.
- $B : Z \mapsto \mathbb{R}^d$ is the observation operator.
- $G : U \mapsto \mathbb{R}_+$ is a regularizing function.
- $U \subset \mathbb{R}^m$ is an open and bounded set, $Z \subset \mathbb{R}^n$ is open and $P \subset \mathbb{R}^d$ is compact.

Connections with Inverse problems

We define the value function for this problem by

$$V(p) = \min_{u \in U,\ z = A(u)} \frac{1}{2} \|B(z) - p\|^2 + G(u).$$

The idea is to use the value function to obtain the solution mapping $u^* : P \mapsto U$.

Connections with Inverse problems

We have that the optimal solution $u^*(p)$ for a given observation p satisfies:

$$(BA)^\top \cdot (BAu^* - p) + \nabla G(u^*) = 0$$

Additionally, we can differentiate (formally) the value function to obtain

$$\nabla V(p) = (p - BAu^*).$$

Combining these two expressions we obtain

$$-(BA)^\top \cdot \nabla V(p) + \nabla G(u^*) = 0$$

Connections with Inverse problems

The condition

$$-(BA)^\top \cdot \nabla V(p) + \nabla G(u^*) = 0$$

is equivalent to

$$u^*(p) = \underset{u \in U, z = Au}{arg\,min} \{G(u) - \nabla V(p)^\top \cdot Bz(u)\}$$

which is the same formula that we have in the control problem case!

Connections with Inverse problems

We can use the previous formula to propose a parametrization for the solution mapping of the following form

$$u_\theta^*(p) = F_\theta((BA)^\top \nabla v_\theta(p))$$

and we can try to find the optimal paramter by solving an averaged version of the inverse problem:

$$\min_{\theta \in \Theta} \int_P \left\{ \frac{1}{2} \|BAu_\theta^*(p) - p\|^2 + G(u_\theta^*(p)) \right\} f(p) dp + \alpha \mathcal{P}(\theta).$$

Connections with Inverse problems

- In many inverse problems the key is to use the right regularization.
- That means that one should choose g wisely according the prior information realted to the problem.
- Maybe one can try to learn it!

Connections with Inverse problems

- In this case one assumes that G is parametrized by a function G_θ, for $\theta \in \Theta$ and we replace it to obtain a regularaized inverse problem:

$$\min_{u \in U, z = Au} \frac{1}{2}\|Bz - p\|^2 + G_\theta(\theta)(u)$$

- The regularization should preserves the convexity of the problem. If C is the smallest eigenvalue of $A^\top B^\top BA$, then G_θ should be $C-$semiconvex. We can use our approach for this.