Week 2. Continuous Controlled Dynamics and Value Functions

Primary reference. H. V. Tran, *Hamilton–Jacobi Equations: Theory and Applications*, §2.1, with forward links to §2.2–2.3. In this course, we intentionally begin with continuous optimal control as the motivating framework, and only later derive the Hamilton–Jacobi equation from dynamic programming.

1. Overview

In Week 1, we studied discrete optimal control and saw that value functions may fail to be smooth even in very simple models. We now pass to continuous time. The first objective is to define the continuous control problem rigorously: admissible controls, controlled trajectories, running costs, discounted total costs, and the associated value function.

This week is therefore foundational. We do **not** yet prove the Hamilton–Jacobi–Bellman equation in full detail. Instead, we establish the control-theoretic structure that will later lead to the dynamic programming principle and the HJB equation.

2. Learning Objectives

By the end of this week, students should be able to:

1. formulate a deterministic continuous-time control system;
2. state the standard assumptions on the controlled vector field b;
3. prove the basic trajectory estimates using Gronwall's inequality;
4. define the discounted cost functional and the value function;
5. explain the role of discounting in infinite-horizon control;
6. compute the value function in several basic examples;
7. understand why value functions may fail to be differentiable;
8. interpret the regularity theorem for value functions as a consequence of the balance between instability of trajectories and discounting.

3. Continuous Controlled Dynamics

We consider a deterministic control system in continuous time:

$$y'(s) = b(y(s), v(s)), \qquad y(0) = x,$$

where:

- $x \in \mathbb{R}^n$ is the initial state,
- V is a compact metric space of control values,
- $v(\cdot) : [0, \infty) \to V$ is a measurable control.

This is the continuous-time analogue of the discrete control system studied in Week 1. Formally, if

$$y_{k+1} \approx y_k + \Delta t\, b(y_k, v_k),$$

then passing to the limit $\Delta t \to 0$ leads to the above ODE.

4. Standing Assumptions on the Dynamics

We assume that

$$b \in C(\mathbb{R}^n \times V; \mathbb{R}^n),$$

and that there are constants $M_b, L_b > 0$ such that for all $x, x_1, x_2 \in \mathbb{R}^n$ and $v \in V$,

$$|b(x,v)| \le M_b, \qquad |b(x_1,v) - b(x_2,v)| \le L_b|x_1 - x_2|.$$

These are the standard hypotheses in Tran's setup. They imply:

- **uniform boundedness of the velocity**, so trajectories cannot move infinitely fast;
- **uniform Lipschitz continuity in the state variable**, which ensures well-posedness and stability of trajectories.

For a given initial point x and control $v(\cdot)$, we write the corresponding trajectory as

$$y_{x,v(\cdot)}(\cdot),$$

or simply $y_x(\cdot)$ when the control is understood.

5. Existence and Uniqueness of Trajectories

Under the above assumptions, for every initial condition $x \in \mathbb{R}^n$ and every measurable control $v(\cdot)$, the state equation admits a unique absolutely continuous trajectory

$$y_{x,v(\cdot)} : [0,\infty) \to \mathbb{R}^n.$$

This follows from the standard Carathéodory theory for ODEs with measurable time dependence and Lipschitz dependence on the state variable. The boundedness of b prevents finite-time blow-up, so the solution is defined for all $s \ge 0$.

6. Basic Trajectory Estimates

The following estimates are fundamental.

Lemma 6.1 (Tran, Lemma 2.2)

Fix a control $v(\cdot)$. Let $y_x(\cdot)$ and $y_z(\cdot)$ be the trajectories starting from x and z, respectively. Then for all $t, s \ge 0$,

$$|y_x(t) - y_z(t)| \le e^{L_b t}|x - z|,$$

and

$$|y_{x,v(\cdot)}(t) - y_{x,v(\cdot)}(s)| \le M_b|t - s|.$$

Proof

Using the integral form of the ODE,

$$y_x(t) - y_z(t) = x - z + \int_0^t (b(y_x(r), v(r)) - b(y_z(r), v(r)))\, dr.$$

Hence

$$|y_x(t) - y_z(t)| \le |x - z| + \int_0^t L_b\, |y_x(r) - y_z(r)|\, dr.$$

By Gronwall's inequality,

$$|y_x(t) - y_z(t)| \le e^{L_b t}|x - z|.$$

Next, if $t > s$, then

$$y_x(t) - y_x(s) = \int_s^t b(y_x(r), v(r))\, dr,$$

so

$$|y_x(t) - y_x(s)| \le \int_s^t |b(y_x(r), v(r))|\, dr \le M_b|t - s|.$$

This proves the second estimate. □

Interpretation

The first estimate shows that the trajectory depends Lipschitz continuously on the initial state, although the constant may grow exponentially in time. The second estimate says that every trajectory is Lipschitz continuous in time, with slope uniformly bounded by M_b.

The following example separates the actual trajectory gap from the Gronwall upper bound. Both trajectories use the same input in the bounded, 1-Lipschitz system

$$y' = \tanh(y) + 0.25\sin(1.4t).$$

The exponential curve is an a priori bound, not a prediction that the actual separation must grow exponentially.

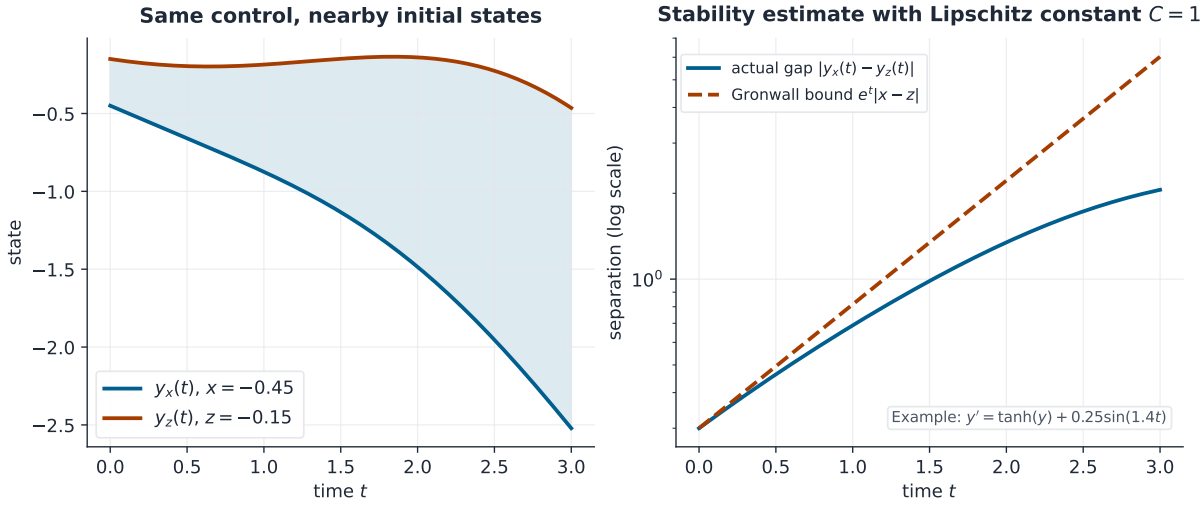


Figure 2.1. Nearby initial states remain controlled by $|y_x(t) - y_z(t)| \le e^{L_b t}|x - z|$. The estimate is intentionally one-sided: the actual gap may be substantially smaller than the universal bound.

7. Running Cost and Discounted Total Cost

We now assign a cost to each trajectory. Let

$$f : \mathbb{R}^n \times V \to \mathbb{R}$$

be the running cost, and assume

$$f \in C(\mathbb{R}^n \times V), \qquad |f(x, v)| \le M_f, \qquad |f(x_1, v) - f(x_2, v)| \le L_f|x_1 - x_2|.$$

Given an admissible control $v(\cdot)$, we define the discounted infinite-horizon cost by

$$J(x, v(\cdot)) = \int_0^\infty e^{-\lambda s} f(y_{x,v(\cdot)}(s), v(s))\, ds, \qquad \lambda > 0.$$

The factor $e^{-\lambda s}$ is called the **discount factor**. It plays two important roles:

1. it guarantees that the integral remains finite when f is only bounded;
2. it is precisely what later produces the zeroth-order term λu in the stationary HJB equation.

Since $|f| \le M_f$,

$$|J(x, v(\cdot))| \le \int_0^\infty e^{-\lambda s} M_f \, ds = \frac{M_f}{\lambda}.$$

Thus the total discounted cost is always well defined.

8. The Value Function

The **value function** is defined by minimizing over all admissible controls:

$$u(x) = \inf_{v(\cdot)} J(x, v(\cdot)).$$

It represents the infimal total cost when the system starts from the state x. This is the central object of the theory. The main question of optimal control is:

> What can be said about the function u, and what equation does it satisfy?

Later we will show that u solves a Hamilton–Jacobi–Bellman equation in the viscosity sense.

From the bound on J, we immediately obtain

$$|u(x)| \le \frac{M_f}{\lambda} \qquad \text{for all } x \in \mathbb{R}^n.$$

9. First Regularity Heuristic for the Value Function

Even before the dynamic programming principle is introduced, one can already see why u should inherit continuity from the data.

Fix a control $v(\cdot)$. Then

$$\begin{aligned} |J(x, v(\cdot)) - J(z, v(\cdot))| &\le \int_0^\infty e^{-\lambda s} |f(y_x(s), v(s)) - f(y_z(s), v(s))| \, ds \\ &\le L_f \int_0^\infty e^{-\lambda s} |y_x(s) - y_z(s)| \, ds \\ &\le L_f |x - z| \int_0^\infty e^{-(\lambda - L_b)s} \, ds. \end{aligned}$$

Hence, if $\lambda > L_b$,

$$|J(x, v(\cdot)) - J(z, v(\cdot))| \le \frac{L_f}{\lambda - L_b} |x - z|,$$

and therefore

$$|u(x) - u(z)| \le \frac{L_f}{\lambda - L_b} |x - z|.$$

So strong discounting implies Lipschitz continuity of the value function. This argument is only a first illustration; the full regularity theorem is more precise and depends on the dynamic programming principle.

10. Supplementary Example W2.1: A Smooth Value Function

Consider the one-dimensional control problem

$$n = 1, \qquad V = [-1, 1], \qquad b(x, v) = v, \qquad f(x, v) = \frac{1}{2} x^2, \qquad \lambda = 1.$$

This example is used as an explicit supplementary computation. Since $f(x,v) = \frac{1}{2}x^2$ is unbounded, it lies outside the bounded running-cost assumptions imposed in Tran §2.1.

Then the dynamics are

$$y'(s) = v(s), \qquad |v(s)| \le 1.$$

Since the running cost penalizes only the state x, the optimal strategy is intuitively clear: move toward the origin as fast as possible, then stay there. Thus the optimal feedback is

$$v^*(x) = -\operatorname{sgn}(x) \quad \text{for } x \neq 0, \qquad v^*(0) = 0.$$

If $x \ge 0$, the optimal trajectory is

$$y^*(s) = \begin{cases} x - s, & 0 \le s \le x, \\ 0, & s \ge x. \end{cases}$$

Therefore

$$u(x) = \frac{1}{2}\int_0^x e^{-s}(x-s)^2\,ds = \frac{1}{2}(x^2 - 2x + 2 - 2e^{-x}), \qquad x \ge 0.$$

By symmetry,

$$u(x) = \frac{1}{2}(|x|^2 - 2|x| + 2 - 2e^{-|x|}), \qquad x \in \mathbb{R}.$$

This example shows that value functions can be explicitly computable and can be rather regular. In fact, this function is C^2 at the origin, although it is not C^3. Thus optimal control problems do not always produce singularities; nonsmoothness depends on the geometry of the problem, not merely on the presence of optimization.

11. A Necessary Warning: the Naive Infinite-Horizon Eikonal Model

At first glance, one may try to consider

$$b(x,v) = v, \qquad |v| \le 1, \qquad f \equiv 1, \qquad \lambda = 1,$$

and hope to obtain a nontrivial eikonal-type value function.

However, in this infinite-horizon discounted setting,

$$J(x, v(\cdot)) = \int_0^\infty e^{-s}\,ds = 1$$

for every control and every initial state. Hence

$$u(x) \equiv 1.$$

So this model is mathematically correct but completely trivial from the point of view of value-function geometry. In particular, it does **not** generate the intended nonsmooth example. For a meaningful analogue of the shortest-path problem from Week 1, one must introduce a target set and stop paying cost once the target is reached.

12. Supplementary Example W2.2: Discounted Time-to-Target and Nonsmoothness

This exit-time problem is a supplementary model outside the whole-space infinite-horizon setup of Tran §2.1.

Let

$$n = 1, \qquad V = [-1, 1], \qquad b(x, v) = v, \qquad \lambda = 1,$$

and let the target set be

$$K = \{0\}.$$

For a control $v(\cdot)$, define its first hitting time by

$$\tau_0(x; v) := \inf\{s \geq 0 : y_{x,v(\cdot)}(s) = 0\}, \qquad \inf \emptyset := \infty.$$

Then

$$J_K(x, v(\cdot)) := \int_0^{\tau_0(x;v)} e^{-s}\, ds.$$

Thus we pay one unit of running cost per unit time until the state first reaches the target, and the problem terminates at that time. Equivalently, one may make K absorbing and set the running cost equal to zero on K. With this exit-time convention,

$$J_K(x, v(\cdot)) = \int_0^{\tau_0(x;v)} e^{-s}\, ds = 1 - e^{-\tau_0(x;v)},$$

where $e^{-\infty} := 0$. Since $|v| \leq 1$, the minimum hitting time is

$$T(x) := \inf_{v(\cdot)} \tau_0(x; v) = |x|,$$

attained by moving directly toward the origin at maximal speed. Hence the value function is

$$u(x) = 1 - e^{-|x|}.$$

This function is continuous on $\mathbb{R}$, satisfies $u(0) = 0$, and is **not differentiable at** $x = 0$. Indeed,

$$u'(x) = \begin{cases} e^{-x}, & x > 0, \\ -e^{x}, & x < 0. \end{cases}$$

Thus the right derivative at 0 is 1, while the left derivative is -1.

This is the key example of the week. The value function is perfectly meaningful and naturally arises from optimization, yet it fails to be classically differentiable at the target boundary. Thus even simple control problems naturally lead to nonsmooth value functions; the PDE here is posed away from the target together with a boundary condition on the target.

Away from the target, the associated Hamiltonian is

$$H(x, p) = \sup_{|v| \leq 1} \{-vp - 1\} = |p| - 1,$$

so the formal PDE is

$$u + |Du| - 1 = 0 \quad \text{in } \mathbb{R} \setminus \{0\}, \qquad u(0) = 0.$$

This is the continuous-time analogue of the nonsmooth shortest-path value function from Week 1.

13. Supplementary Example W2.3: Discounted Minimum-Time Interpretation

More generally, let $K \subset \mathbb{R}^n$ be a compact target set, and consider

$$y'(s) = b(y(s), v(s)), \qquad v(s) \in V,$$

and define the first hitting time and minimum arrival time by

$$\tau_K(x; v) := \inf\{s \geq 0 : y_{x,v(\cdot)}(s) \in K\}, \qquad \inf \emptyset := \infty, \qquad T_K(x) := \inf_{v(\cdot)} \tau_K(x; v).$$

Define the exit-time value function by

$$u(x) = \inf_{v(\cdot)} \int_0^{\tau_K(x;v)} e^{-\lambda s}\, ds = \inf_{v(\cdot)} \frac{1 - e^{-\lambda \tau_K(x;v)}}{\lambda} = \frac{1 - e^{-\lambda T_K(x)}}{\lambda},$$

where $e^{-\lambda \infty} := 0$.

Since the function $t \mapsto (1 - e^{-\lambda t})/\lambda$ is strictly increasing and continuous on $[0, \infty]$, taking the infimum of the discounted costs is equivalent to taking the infimum of the arrival times. Thus the discounted cost formulation is a monotone transformation of a minimum-time control problem. If one instead writes an infinite-horizon indicator running cost, one must additionally assume that K is absorbing, or otherwise ensure that no cost is incurred after the first hitting time.

This viewpoint will become important later when we study front propagation and geometric Hamilton–Jacobi equations.

14. Preview: the Hamiltonian

The Hamiltonian associated with the infinite-horizon discounted control problem is defined by

$$H(x, p) = \sup_{v \in V}\{-b(x, v) \cdot p - f(x, v)\}.$$

This is the sign convention used throughout Tran and throughout this course. Under the standing assumptions, H is continuous, convex in p, and Lipschitz in p, with Lipschitz constant M_b. Moreover, it satisfies the estimate

$$|H(x, p) - H(y, p)| \leq (L_f + L_b|p|)|x - y|.$$

The stationary HJB equation associated with the value function is

$$\lambda u + H(x, Du) = 0.$$

We will derive this equation systematically next week from the dynamic programming principle.

15. Regularity of the Value Function

The following theorem is one of the first major regularity results for value functions.

Let λ_0 denote the uniform Lipschitz constant of b with respect to x. If $D_x b$ exists and is essentially bounded, this can be written as

$$\lambda_0 = \|D_x b\|_{L^\infty(\mathbb{R}^n \times V)}.$$

Theorem 15.1 (Tran, Theorem 2.9)

Let u be the value function of the discounted infinite-horizon problem. Then:

- if $\lambda > \lambda_0$, then $u \in \mathrm{Lip}(\mathbb{R}^n)$;
- if $\lambda = \lambda_0$, then $u \in C^{0,\alpha}(\mathbb{R}^n)$ for every $0 < \alpha < 1$;
- if $0 < \lambda < \lambda_0$, then $u \in C^{0,\lambda/\lambda_0}(\mathbb{R}^n)$;
- in all cases, $u \in \mathrm{BUC}(\mathbb{R}^n)$.

Interpretation

This theorem expresses a competition between two effects:

- perturbations of the initial condition grow under the dynamics at a rate controlled by λ_0;
- future cost is damped by the discount factor $e^{-\lambda s}$.

If discounting is stronger than the instability of the dynamics, Lipschitz regularity survives. If it is weaker, one still retains Hölder continuity, but with a smaller exponent.

In the book, the proof uses the dynamic programming principle. Since our course discusses DPP only in Week 3, we state this theorem here as a structural preview and return to its proof later.

16. Concluding Summary

This week introduced the basic framework of deterministic continuous optimal control.

- A measurable control generates a unique absolutely continuous trajectory.
- Controlled trajectories satisfy stability estimates with respect to both the initial state and time.
- A bounded running cost together with discounting yields a well-defined infinite-horizon total cost.
- Minimizing over all admissible controls defines the value function.
- Value functions may be smooth in some problems, but in geometrically natural models they can fail to be differentiable.
- This failure of classical smoothness is not pathological; it is intrinsic to the theory and motivates the study of viscosity solutions.

Thus the logical path of the course is now clear:

$$\text{controlled dynamics} \longrightarrow \text{cost functional} \longrightarrow \text{value function} \longrightarrow \text{DPP} \longrightarrow \text{HJB equation}.$$

17. Exercises

Exercise 2.1

Show directly from the state equation that for every trajectory,

$$|y_x(t) - x| \le M_b t \qquad \text{for all } t \ge 0.$$

Exercise 2.2

Assume $|f| \le M$. Prove that

$$|u(x)| \le \frac{M}{\lambda} \qquad \text{for all } x \in \mathbb{R}^n.$$

Exercise 2.3

In Supplementary Example W2.1, prove rigorously that the optimal strategy is to move to the origin at maximal speed and then stay there.

Exercise 2.4

In Supplementary Example W2.2, compute the one-sided derivatives of

$$u(x) = 1 - e^{-|x|}$$

at the origin and verify formally that

$$u + |Du| - 1 = 0 \qquad \text{for } x \neq 0.$$

Exercise 2.5

Assume $\lambda > L_b$. Complete the proof that u is Lipschitz continuous by comparing the costs corresponding to the same control for two nearby initial points.