ECE 586 – Markov Decision Processes and Reinforcement Learning

Ordinary Differential Equations: Some Concepts

Instructor: Dimitrios Katselis *Acknowledgment: R. Srikant's Notes and Discussions*

1 Basic Concepts

Consider an autonomous ordinary differential equation (ODE)

$$\dot{\theta}(t) = h(\theta(t)), \quad \theta(0) = \theta_0, \quad \theta(t) \in \mathbb{R}^d, \quad t \in \mathbb{R}. \tag{1}$$

The ODE would be non-autonomous if the driving vector field h had an explicit time dependence, i.e., if the equation were $\dot{\theta}(t) = h(\theta(t), t)$ instead.

Terminology Origin: Originates from mechanical examples in which an explicit time dependence corresponds to an external force imposed on the system. When no such force exists, the system is **self-governing**, or **autonomous**.

The ODE is well-posed, i.e., for each initial condition $\theta_0 \in \mathbb{R}^d$ it has a unique solution $\theta(t), t \geq 0$ and the map associating an initial condition θ_0 to its corresponding solution $\theta(t) \in C([0,\infty), \mathbb{R}^d)$ is continuous for the topology of uniform convergence on compacts.

Note: $C([0,\infty), \mathbb{R}^d)$ is the space of all continuous functions $f : [0,\infty) \to \mathbb{R}^d$.

Sufficient Condition such that (1) is well-posed: h is Lipschitz, i.e., $\exists L > 0$ such that

$$\|h(x) - h(y)\| \leq L\|x - y\|, \quad \forall x, y \in \mathbb{R}^d.$$

Invariant Set: A closed set $\mathcal{O} \subset \mathbb{R}^d$ is an invariant set of (1) if any trajectory $\theta(t), t \in (-\infty, \infty)$ with $\theta(0) \in \mathcal{O}$ satisfies $\theta(t) \in \mathcal{O}$ for all $t \in \mathbb{R}$. $\mathcal{O}$ is positively (respectively negatively) invariant if $\theta(0) \in \mathcal{O}$ implies that $\theta(t) \in \mathcal{O}$ for all $t \geq 0$ (respectively for all $t \leq 0$). Considering positive invariance, this property of $\mathcal{O}$ means that once the trajectory $\theta(t)$ enters $\mathcal{O}$, it never leaves.

Internally Chain Transitive Set: A closed set $\mathcal{O} \subset \mathbb{R}^d$ such that for any $x, y \in \mathcal{O}$ and any $\epsilon > 0, T > 0$, there exists an integer $n \geq 1$ and a sequence of points $x_0 = x, x_1, \ldots, x_{n-1}, x_n = y$ with $x_i \in \mathcal{O}, \forall i$ such that the trajectory of (1) starting at x_i for some $0 \leq i < n$ meets with the ϵ-neighborhood of x_{i+1} after time $\geq T$. For $x = y$, we recover the definition of an **internally chain recurrent set** $\mathcal{O}$. Clearly, the notion of internal chain transitivity is stronger than the notion of internal chain recurrence. Also the sequence $\{x_0, \ldots, x_n\}$ is called ϵ-chain in $\mathcal{O}$ connecting x, y.

ω-limit set: Given a trajectory $\theta(t)$ of (1), the set $\Omega = \bigcap_{t>0} \overline{\{\theta(t') : t' > t\}}$, i.e., the set of its limit points as $t \to \infty$ is called its ω-limit set. Here, $\overline{\{\cdot\}}$ denotes the closure of a set. In general, the set depends on $\theta(0)$, i.e., on the actual trajectory. Clearly, Ω is an invariant set for (1). If $\Omega = \{\theta^*\}$, then $\theta(t) = \theta^*$ is a trajectory of the ODE and since it is a constant solution it should satisfy $h(\theta^*) = 0$. Conversely, if $h(\theta^*) = 0$, then $\theta(t) = \theta^*$ corresponds to a trajectory (a solution) of (1) for initial condition $\theta(0) = \theta^*$. Points satisfying $h(\theta^*) = 0$ correspond to equilibrium points of (1).

Attractor: A compact invariant set (or closed in more generality) $\mathcal{M}$ is an attractor if it has an open neighborhood $\mathcal{O}$ such that every trajectory in $\mathcal{O}$ remains in $\mathcal{O}$ and converges to $\mathcal{M}$. The largest such $\mathcal{O}$ is called the **domain of attraction** of $\mathcal{M}$.

Lyapunov Stable Set: A compact invariant set $\mathcal{M}$ such that for any $\epsilon > 0$, there exists a $\delta > 0$ with the property that every trajectory initiated in the δ-neighborhood of $\mathcal{M}$ remains in its ϵ-neighborhood.

Asymptotically Stable Set: A compact invariant set $\mathcal{M}$ that is both Lyapunov stable and an attractor. If $\mathcal{M} = \{\theta^*\}$, the equilibrium point θ^* is asymptotically stable.

Lyapunov's Second Method to verify asymptotic stability of θ^*: If there exists a continuously differentiable function V defined in a neighborhood $\mathcal{O}$ of θ^* such that

$$\langle \nabla V(\theta), h(\theta) \rangle < 0, \quad \theta \in \mathcal{O}, \quad \theta \neq \theta^*$$

and $= 0$ for $\theta = \theta^*$, with $V(\theta) \to \infty$ as $\theta \to \partial\mathcal{O}$, then the asymptotic stability of θ^* follows by observing that for any trajectory $\theta(\cdot) \in \mathcal{O}$, we have $\frac{d}{dt}V(\theta(t)) \leq 0$ with equality only for $\theta(t) = \theta^*$.

Conversely, asymptotic stability of θ^* implies the existence of such a function.

Remarks:

1. $\partial\mathcal{O}$ is the boundary of $\mathcal{O}$.
2. This also generalizes to compact invariant sets $\mathcal{M}$ that are asymptotically stable.

Globally Asymptotically Stable θ^*: An asymptotically stable θ^* such that all trajectories of the ODE (1) converge to it. In this case, the neighborhood $\mathcal{O}$ of θ^* over which the Lyapunov function V is defined can be taken to be the whole space.

LaSalle Invariance Principle: If there exists a continuously differentiable function $V : \mathbb{R}^d \to \mathbb{R}$ such that $V(\theta) \to \infty$ when $\|\theta\| \to \infty$ and $\langle \nabla V(\theta), h(\theta) \rangle \leq 0, \forall\theta$, then all trajectories $\theta(\cdot)$ must converge to the largest invariant set contained in $\{\theta : \langle \nabla V(\theta), h(\theta) \rangle = 0\}$.

Not every equilibrium point is asymptotically stable: Consider the basic linear system $\dot{\theta} = A\theta$ and assume that A is nonsingular. Then, $\theta^* = 0$ is the only equilibrium point. The solution of this ODE is $\theta(t) = e^{At}\theta(0)$. If all eigenvalues of A have strictly negative real parts (i.e., A is stable or Hurwitz) then $\theta(t) \xrightarrow[t\to\infty]{} 0$ exponentially fast for any initial condition. Otherwise, it will converge to $\theta^* = 0$ only for those $\theta(0)$ that lie on the *stable subspace*, i.e., the eigenspace of those eigenvalues (if any) which have strictly negative real parts.

2 Gronwall and Discrete Gronwall Inequalities

To prove that the Lipschitz condition on h ensures that (1) is well-posed, the following results are used (we only provide these key inequalities in the context of dynamical systems with no further analysis):

Lemma 1. *(**Gronwall Inequality**) For continuous $x(\cdot), y(\cdot) \geq 0$ and scalars $C, K, T \geq 0$*

$$x(t) \leq C + K \int_0^t x(s)y(s)ds, \quad \forall t \in [0, T]$$

implies that

$$x(t) \leq Ce^{K\int_0^t y(s)ds}, \quad t \in [0, T].$$

Most common case: $y(t) = 1, \forall t$, leading to

$$x(t) \leq Ce^{Kt}.$$

A useful discrete analog of this inequality is:

Lemma 2. ***(Discrete Gronwall inequality)*** *Let $\{x_n\}_{n\geq 0}$ and $\{\varepsilon_n\}_{n\geq 0}$ be a nonnegative and a positive sequence, respectively. Suppose that $C, L \geq 0$ are scalars such that*

$$x_{n+1} \leq C + L \sum_{k=0}^{n} \varepsilon_k x_k, \quad \forall n.$$

Then,

$$x_{n+1} \leq C e^{L \sum_{k=0}^{n} \varepsilon_k}.$$

We will use these inequalities in deriving the ODE method result for stochastic approximation.

3 A Note on Stochastic Gradient Algorithms

Consider the stochastic gradient descent algorithm:

$$\theta_{n+1} = \theta_n + \varepsilon_n \left(-\nabla f(\theta_n) + M_n \right),$$

where $h(\theta) = -\nabla f(\theta)$ in our previous stochastic approximation file and M_n is a noise sequence. This algorithm aims at minimizing f, which is assumed to be continuously differentiable. The limiting ODE is

$$\dot{\theta} = -\nabla f(\theta). \tag{2}$$

For such a gradient flow dynamical system, f itself serves as a Lyapunov function:

$$\frac{d}{dt} f(\theta(t)) = \nabla f(\theta(t)) \frac{d}{dt}\theta(t) = -\|\nabla f(\theta(t))\|^2 \leq 0,$$

and the inequality is strict when $\nabla f(\theta(t)) \neq 0$. The set of equilibrium points is $\mathcal{O} = \{\theta \in \mathbb{R}^d : \nabla f(\theta) = 0\}$. LaSalle Invariance Principle implies that the only possible ω-limit sets for (2) are subsets of $\mathcal{O}$ and the ODE method yields the conclusion that the iterates converge almost surely to such an invariant set. Convergence to local maxima or saddle points is ruled out, since these correspond to unstable equilibrium points. If f has only isolated local minima, then $\{\theta_n\}$ converges almost surely to one of them.