Online Optimization in Continuous Time

Recall: $\quad f : \mathbb{R}^n \to \mathbb{R} \qquad C^1$, convex

$$f(\lambda x + (1-\lambda)x') \le \lambda f(x) + (1-\lambda) f(x')$$
for all $x, x' \in \mathbb{R}^n$, all $0 \le \lambda \le 1$

$$\iff f(x') \ge f(x) + \nabla f(x)^T (x' - x) \quad \forall x, x' \in \mathbb{R}^n$$

Gradient descent: $\quad \dot{x} = -\nabla f(x) \qquad x(0) \in \mathbb{R}^n$

ideally, want $x(t) \to x^*$ (global min.)

Adaptive control: $\quad u(\cdot) \longrightarrow \boxed{\theta} \longrightarrow y(\cdot)$

- at time t, we know $(u(s), y(s))$ for $s \in [0, t]$
- estimator: $\{(u(s), y(s)) : s \in [0,t]\} \longrightarrow \hat{\theta}(t)$
- output prediction error: $\hat{y}(t) - y(t)$, where $\hat{y}(t)$ is determined by a model

$$y(t) = F_t(\{u(s) : s \in [0,t]\}; \theta) \qquad \forall t \ge 0$$

$\hat{\theta}$: $\quad \hat{y}(t) = F_t(\{u(s) : s \in [0,t]\}; \hat{\theta})$

Tune $\hat{\theta}(t)$ to make $\hat{y}(t) - y(t)$ small.

Estimation error at time t:

$$J_t(\hat{\theta}) := \frac{1}{2}(\hat{y}(t) - y(t))^2,$$
where $\hat{y}(t) = F_t(u_{[0,t]}; \hat{\theta})$

We would like the optimization to "perform well" over time —

do almost as well, as $t \to \infty$, as the best estimator of θ based on $u_{[0,t]}, y_{[0,t]}$.

Integrated prediction error:

$\vec{\theta} \in \mathbb{R}^k \longrightarrow \int_0^t J_s(\vec{\theta})\,ds$

(k-parameter dim.)

An estimation strategy — law for updating $\hat{\theta}(t)$:

$$\int_0^t J_s(\hat{\theta}(s))\,ds \qquad \hat{\theta}(s) \leftarrow u_{[0,s]}, y_{[0,s]}$$

Regret at time t:

$$R(t) := \int_0^t J_s(\hat{\theta}(s))\,ds - \min_{\vec{\theta} \in \mathbb{R}^k} \int_0^t J_s(\vec{\theta})\,ds$$

- depends on estimation law and on the plant, input, and output

Goal: $\frac{1}{t} R(t) \to 0$ as $t \to \infty$

Online optimization (in cont. time)

$J_t : \mathbb{R}^k \to \mathbb{R}$ $\quad C^1$, convex

- at time t, we know only $J_{[0,t]} := (J_s : s \in [0,t])$
- goal: come up with a strategy for generating $\hat{\theta}_t \in \mathbb{R}^k$ s.t.

$$\int_0^t J_s(\hat{\theta}(s))\,ds - \inf_{\vec{\theta} \in \mathbb{R}^k} \int_0^t J_s(\vec{\theta})\,ds$$

is small $[o(t)]$

This is possible with simple gradient descent rules, under some assumptions on $J_t(\cdot)$ for all $t \geq 0$.

Example ①

Assume all $J_t(\cdot)$ are C^1, convex, and uniformly Lipschitz continuous: $\exists L > 0$

$$|J_t(\hat\theta) - J_t(\hat\theta')| \le L\,|\hat\theta - \hat\theta'| \qquad \forall \theta, \theta' \in \mathbb{R}^k$$

$$\Leftrightarrow \quad |\nabla J_t(\hat\theta)| \le L \qquad \forall \theta \in \mathbb{R}^n.$$

Claim: $\dot{\hat\theta} = -\gamma \nabla J_t(\hat\theta)$ [$\gamma > 0$ – fixed adaptation gain]

suffices for $R(t) = O(1)$

Analysis: based on Lyapunov functions

- we will look at $\int_0^t J_s(\hat\theta(s))ds - \int_0^t J_s(\theta)ds$

  for an arbitrary point $\theta \in \mathbb{R}^k$.

- $V(\hat\theta) := \frac{1}{2}|\hat\theta - \theta|^2$

$$\dot V(\hat\theta(t)) = (\hat\theta(t) - \theta)^T \dot{\hat\theta}(t)$$

$$= -\gamma\,(\hat\theta(t) - \theta)^T \nabla J_t(\hat\theta(t))$$

- At each s, by convexity,

$$J_s(\theta) - J_s(\hat\theta(s)) \ge \nabla J_s(\hat\theta(s))^T(\theta - \hat\theta(s))$$

$$= \frac{1}{\gamma}\dot V(\hat\theta(s))$$

$$\Rightarrow \int_0^t \left[J_s(\hat\theta(s)) - J_s(\theta)\right]ds \le -\frac{1}{\gamma}\int_0^t \dot V(\hat\theta(s))\,ds$$

$$\int_0^t J_s(\hat\theta(s))ds - \int_0^t J_s(\theta)ds \le -\frac{1}{\gamma}\left[V(\hat\theta(t)) - V(\hat\theta(0))\right]$$

$$= \frac{1}{\gamma}V(\hat\theta(0)) - \frac{1}{\gamma}V(\hat\theta(t))$$

$$\le \frac{1}{\gamma}V(\hat\theta(0)) \qquad [V(\cdot) \ge 0]$$

Regret bound: for any $T \geq 0$

$$\underbrace{\int_0^T J_t(\hat{\theta}(t))\,dt - \inf_{\theta} \int_0^T J_t(\theta)\,dt}_{= R(T)} \leq \frac{1}{\gamma} V(\hat{\theta}(0))$$

$R(T) \leq C < \infty$ for all T

$$\dot{V}(\hat{\theta}(t)) = -\gamma(\hat{\theta}(t) - \theta)^T \nabla J_t(\hat{\theta}(t))$$
$$\leq \gamma\left[J_t(\theta) - J_t(\hat{\theta}(t))\right]$$

Assume $\theta^* \in \mathbb{R}^k$ is the global minimum of <u>all</u> $J_t(\cdot) \Rightarrow$

$J_t(\theta^*) - J_t(\hat{\theta}(t)) \leq 0 \quad \forall t \qquad \left[V(\hat{\theta}) = \frac{1}{2}|\hat{\theta} - \theta^*|^2\right]$

$\dot{V}(\hat{\theta}(t)) \leq 0 \qquad \boxed{V(\hat{\theta}(t)) \leq V(\hat{\theta}(0))}$

$|\hat{\theta}(t) - \theta^*| \leq |\hat{\theta}(0) - \theta^*|$

$\Rightarrow \hat{\theta}(t)$ is bdd

$\dot{\hat{\theta}} = -\gamma \nabla J_t(\hat{\theta}) \qquad |\nabla J_t(\cdot)| \leq L < \infty$
(by Lipschitz assumption)

$|\dot{\hat{\theta}}(t)| = \gamma|\nabla J_t(\hat{\theta}(t))| \leq \gamma L$ — bdd

$J_t(\hat{\theta}(t)) - J_t(\theta^*) \geq 0 \qquad [\theta^* - \text{global min.}]$

$\int_0^\infty \left[J_t(\hat{\theta}(t)) - J_t(\theta^*)\right] dt$ exists, is finite

Can we conclude that

$$\lim_{t \to \infty} \left[J_t(\hat{\theta}(t)) - J_t(\theta^*)\right] = 0 \quad ?$$

Recall <u>Barbalat's lemma</u>: Let $g(t)$ be uniformly cont. (as a fcn of t), and suppose
$\int_0^\infty g(t)\,dt$ exists, is finite.

Then $g(t) \to 0$ as $t \to \infty$.

We have $g(t) = J_t(\hat\theta(t)) - J_t(\theta^*)$
– need info on continuity in t.

$|g'(t)| \le K < \infty$ for all $t \geq 0$ is sufficient.

$$\frac{d}{dt} g(t) = \frac{\partial}{\partial t} J_t(\theta)\Big|_{\theta=\hat\theta(t)} + \dot{\hat\theta}(t)^T \nabla J_t(\hat\theta(t))$$

$$= \frac{\partial}{\partial t} J_t(\theta)\Big|_{\theta=\hat\theta(t)} - \underbrace{\frac{1}{\gamma}\dot{\hat\theta}(t)^T \dot{\hat\theta}(t)}_{\text{bdd}}$$

<u>Ex.</u> $u(\cdot) \longrightarrow \boxed{} \longrightarrow y(\cdot)$ $\qquad y(t) = \theta u(t)$

$$J_t(\hat\theta) := \frac{1}{2}\big(\underbrace{\hat\theta u(t)}_{\hat y(t)} - \theta u(t)\big)^2$$

$$= \frac{1}{2}(\hat\theta - \theta)^2 u^2(t)$$

$$\nabla J_t(\hat\theta) = (\hat\theta - \theta)u^2(t)$$

$$\frac{\partial}{\partial t} J_t(\hat\theta) = (\hat\theta - \theta)^2 u(t)\dot u(t)$$

If $u, \dot u$ are bdd, then all conditions are satisfied, and so

$$\dot{\hat\theta}(t) = -\gamma \nabla J_t(\hat\theta(t))$$

has no-regret property, and $(\hat\theta(t) - \theta)^2 u^2(t) \to 0$.

$\hat{\theta}(t)$ does not necessarily converge to θ; we will need more conditions on u [Persistent Excitation property].