Online Parameter Estimation

General framework (review):

- $J_t : \mathbb{R}^k \to \mathbb{R}$ C^1, convex
- at time t, know only $J_s(\cdot)$, $0 \le s \le t$
- construct $\hat{\theta}(t)$ based on this info (online)

Goal: asymptotically no-regret

$$\int_0^t J_s(\hat{\theta}(s))\,ds - \inf_{\tilde{\theta} \in \mathbb{R}^k} \int_0^t J_s(\tilde{\theta})\,ds \le C < \infty$$

for all $t \ge 0$

Running example: scalar gain

$u \to \boxed{} \to y$ $\qquad y(t) = \theta u(t) \qquad \theta \in \mathbb{R}$ unknown

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

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

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

($\gamma > 0$ adaptation gain)

If $u, \dot{u}$ are bdd, then $J_t(\hat{\theta}(t)) \to 0$ as $t \to \infty$
(Barbalat)

Moreover, if $u(\cdot)$ has Persistent Excitation property, then $\hat{\theta}(t) \to \theta$ as $t \to \infty$ exponentially fast.

Question: Can we still have useful guarantees without assuming bdd $u, \dot{u}$?

Yes — with a simple normalization trick.

$y(t) = \theta u(t)$ $\qquad$ $u, \dot{u}$ no longer assumed bdd

Rescaled (normalized) gradient flow:

$$\dot{\hat{\theta}}(t) = -\frac{1}{m^2(t)} \nabla J_t(\hat{\theta}(t)) \qquad J_t(\hat{\theta}) = \frac{1}{2}(\hat{\theta}-\theta)^2 u^2(t)$$

where $m(\cdot)$ is a nonzero signal to be chosen.

Go through the analysis:

$V(\hat{\theta}) := \frac{1}{2}(\hat{\theta}-\theta)^2$ $\qquad$ θ - true gain

$J_t(\theta) = 0 \quad \forall t \geq 0$

$$\underbrace{J_t(\theta)}_{=0} - J_t(\hat{\theta}(t)) \geq \nabla J_t(\hat{\theta}(t))^T(\theta - \hat{\theta}(t))$$

$$\dot{V} = (\hat{\theta}-\theta)^T \dot{\hat{\theta}}$$
$$= -(\hat{\theta}-\theta)^T \frac{\nabla J_t(\hat{\theta})}{m^2}$$

$$\Rightarrow \quad 0 \geq J_t(\theta) - J_t(\hat{\theta}(t))$$
$$= m^2(t)\dot{V}(\hat{\theta}(t))$$

$$\dot{V}(\hat{\theta}(t)) \leq 0 \quad \Rightarrow \quad |\hat{\theta}(t)-\theta|^2 \leq |\hat{\theta}(0)-\theta|^2$$

$\hat{\theta}$ bdd

$$J_t(\hat{\theta}(t)) \leq -m^2(t)\dot{V}(\hat{\theta}(t))$$

$$\frac{1}{m^2(t)} J_t(\hat{\theta}(t)) \leq -\dot{V}(\hat{\theta}(t))$$

$$\int_0^t \frac{1}{m^2(s)} J_s(\hat{\theta}(s))\,ds \leq V(\hat{\theta}(0)), \quad \forall t \geq 0$$

Normalized cost: $\bar{J}_t(\hat{\theta}) := \frac{1}{m^2(t)} J_t(\hat{\theta})$

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

Normalized input: $\bar{u}(t) := \frac{u(t)}{m(t)}$

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

$$\hat{\theta}\,\bar{u}(t) = \hat{\theta}\,\frac{u(t)}{m(t)} = \frac{\hat{y}(t)}{m(t)} =: \hat{\bar{y}}(t)$$

$$\bar{y}(t) = \frac{y(t)}{m(t)} = \theta\,\frac{u(t)}{m(t)} = \theta\,\bar{u}(t)$$

We would like to apply Barbalat:

$$0 \le \int_0^t \bar{J}_s(\hat{\theta}(s))\,ds \le V(\hat{\theta}(0)), \qquad \forall t \geq 0$$

$\int_0^\infty \bar{J}_t(\hat{\theta}(t))\,dt$ exists and is finite.

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

– need this to be uniformly cont. in t.

$$m(t) = \sqrt{1+u^2(t)}$$

$$\bar{u}(t) = \frac{u(t)}{\sqrt{1+u^2(t)}} \quad \text{bdd}$$

$$\dot{\hat{\theta}}(t) = -\frac{1}{m^2(t)}\nabla J_t(\hat{\theta}(t))$$

$$= -\nabla \bar{J}_t(\hat{\theta}(t))$$

$$= -(\hat{\theta}(t)-\theta)\,\bar{u}^2(t) \qquad \text{bdd}$$

$\Rightarrow \quad \bar{J}_t(\hat{\theta}(t)) \longrightarrow 0$ provided $\dot{\bar{u}}$ is bdd

$$\frac{d}{dt}\bar{J}_t(\hat{\theta}(t)) = \frac{d}{dt}\left\{\frac{1}{2}(\hat{\theta}(t)-\theta)^2\,\bar{u}^2(t)\right\}$$

$$= \underbrace{(\hat{\theta}(t)-\theta)\,\bar{u}^2(t)\,\dot{\hat{\theta}}(t)}_{\text{bdd}} + \underbrace{(\hat{\theta}(t)-\theta)^2\,\bar{u}(t)}_{\text{bdd}}\,\dot{\bar{u}}(t)$$

So far:

- $\tilde{\theta}(t) := \hat{\theta}(t) - \theta$ (parameter estimation error) is bdd

- $\bar{e}(t) := \hat{\bar{y}}(t) - \bar{y}(t) = (\hat{\theta}(t) - \theta)\,\bar{u}(t)$
$= \tilde{\theta}(t)\,\bar{u}(t)$ bdd

- $\int_0^\infty \bar{J}_t(\hat{\theta}(t))\,dt < \infty$ (shown above)

$$\int_0^\infty \bar{J}_t(\hat{\theta}(t))\,dt = \frac{1}{2}\int_0^\infty \bar{e}^2(t)\,dt, \qquad \bar{e} \in L_2 \cap L_\infty$$

- Slow adaptation: $\dot{\hat{\theta}} = -\nabla \bar{J}_t(\hat{\theta})$
$= -(\hat{\theta} - \theta)\,\bar{u}^2$
$= -\tilde{\theta}\,\bar{u}^2$

$$|\dot{\hat{\theta}}(t)| = |\tilde{\theta}(t)|\,\bar{u}^2(t) \quad \text{bdd} \Rightarrow \dot{\hat{\theta}} \in L_\infty$$

$$\left.\begin{aligned} \dot{\hat{\theta}} &= -(\hat{\theta} - \theta)\,\bar{u}^2 \\ &= -\underbrace{(\hat{\theta} - \theta)\,\bar{u}}_{\bar{e}} \cdot \bar{u} \end{aligned}\right\} \quad \dot{\hat{\theta}} = -\bar{e}\,\bar{u}$$

$$0 \le \int_0^t |\dot{\hat{\theta}}(s)|^2\,ds = \int_0^t \bar{e}^2(s)\,\bar{u}^2(s)\,ds$$
$$\le \|\bar{u}\|_\infty^2 \int_0^t \bar{e}^2(s)\,ds \le C < \infty \text{ for all } t$$

$\Rightarrow \int_0^\infty |\dot{\hat{\theta}}(t)|^2\,dt$ exists, finite

$\boxed{\dot{\hat{\theta}} \in L_\infty \cap L_2}$ – key property when dealing with unstable systems for ensuring closed-loop stability.

Normalization leads to "smaller steps" –

discretized gradient flow $\quad \dot{\hat{\theta}} = -\frac{1}{m^2}\nabla J_t(\hat{\theta})$

- pick a small $h > 0$

$$\hat{\theta}(t+h) = \hat{\theta}(t) - \frac{h}{m^2(t)}\nabla J_t(\hat{\theta}(t)) + o(h)$$

$$\underbrace{\hat{\theta}(t+h) - \hat{\theta}(t)}_{\text{step size}} \simeq -\frac{h}{m^2(t)}\nabla J_t(\hat{\theta}(t))$$

- if $m^2(t)$ is large, we take a smaller step

Example: dynamic plant

$$\dot{x} = -ax + bu \qquad u, x \in \mathbb{R}$$
$$y = x$$

Unknown parameters: $a > 0$, $b \in \mathbb{R}$ (stable case)

Here,

$$\begin{aligned} y(t) &= x(t) \\ &= e^{-at}x(0) + \int_0^t e^{-a(t-s)}u(s)\,ds \qquad \text{(variation of constants)} \\ &=: F_t(u_{[0,t]}; \theta) \qquad \theta := (a, b)^T \end{aligned}$$

(assume known i.c.)

$$\hat{y}(t) = F_t(u_{[0,t]}; \hat{\theta}) \qquad \hat{\theta} : (\hat{a}, \hat{b})^T$$

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

Instead, let's try a dynamic prediction/estimation strategy:

$$\dot{x} = -ax + bu \qquad \dot{\hat{x}} = -\lambda(\hat{x} - x) - \hat{a}x + \hat{b}u$$

($\lambda > 0$ - fixed parameter)

$\hat{a}, \hat{b}$ - generated dynamically

$e := \hat{x} - x$

$$\begin{aligned}
\dot{e} &= \dot{\hat{x}} - \dot{x} \\
&= -\lambda e - \hat{a}x + \hat{b}u + ax - bu \\
&= -\lambda e - \tilde{a}x + \tilde{b}u
\end{aligned}$$

$\tilde{a} := \hat{a} - a$

$\tilde{b} := \hat{b} - b$

$$e(t) = e^{-\lambda t} e(0) + \int_0^t e^{-\lambda(t-s)} \left\{ -\tilde{a}(s)x(s) + \tilde{b}(s)u(s) \right\} ds$$

Lyapunov fcn:

$$V(e, \tilde{a}, \tilde{b}) := \frac{1}{2}e^2 + \frac{1}{2}\tilde{a}^2 + \frac{1}{2}\tilde{b}^2$$

$$\begin{aligned}
\dot{V} &= e\dot{e} + \tilde{a}\dot{\tilde{a}} + \tilde{b}\dot{\tilde{b}} \\
&= e\dot{e} + \tilde{a}\dot{\hat{a}} + \tilde{b}\dot{\hat{b}} \\
&= e(-\lambda e - \tilde{a}x + \tilde{b}u) + \tilde{a}\dot{\hat{a}} + \tilde{b}\dot{\hat{b}} \\
&= -\lambda e^2 - \tilde{a}xe + \tilde{a}\dot{\hat{a}} + \tilde{b}ue + \tilde{b}\dot{\hat{b}}
\end{aligned}$$

$\dot{\tilde{a}} = \dot{\hat{a}}$

$\dot{\tilde{b}} = \dot{\hat{b}}$

a, b — constants

$\Rightarrow$ take $\dot{\hat{a}} = xe$, $\dot{\hat{b}} = -ue$

Closed-loop system:

$$\left.\begin{aligned}
\dot{x} &= -ax + bu \\
\dot{\hat{x}} &= -\lambda(\hat{x} - x) - \hat{a}x + \hat{b}u \\
\dot{\hat{a}} &= xe \\
\dot{\hat{b}} &= -ue
\end{aligned}\right\} \quad \dot{V} = -\lambda e^2 \leq 0$$

- can show $\hat{a}, \hat{b}$ bdd; $e \to 0$, $\tilde{a}, \tilde{b} \to 0$ under appropriate assumptions ($u, \dot{u}$ bdd)
- not a gradient flow type scheme can't give it an online optimization interpretation

What about unstable systems?

$\dot{x} = -ax + bu$ $\qquad\qquad$ $\theta = (a, b)^T$

no longer assume $a > 0$

If u is not bdd, what can we do?
Naive renormalization does not work!

Next lecture: linear parametrization (i/o)