Linear-Quadratic Estimation

Jenny Hong   Stephen Boyd

EE103
Stanford University

November 26, 2015

Outline

Least-squares state estimation

Missing measurements

State estimation

- dynamic system model:

$$x_{t+1} = Ax_t + Bw_t, \quad y_t = Cx_t + v_t, \quad t = 1, 2, \ldots$$

- x_t is *state* (n-vector)
- y_t is *measurement* (p-vector)
- v_t is *measurement noise* or *measurement residual* (p-vector)
- w_t is *input* or *process noise* (m-vector)
- we know A, B, C, and measurements $y_1, \ldots, y_T$
- w_t, v_t are unknown, but assumed small
- *state estimation*: estimate/guess $x_1, \ldots, x_T$
- some common variations:
  - x_1 is known
  - y_t not known for some t's
  - matrices A, B, C are time-varying

Least-squares state estimation

- choose estimates $\hat{x}_t$, $\hat{w}_t$, $\hat{v}_t$ by solving linearly constrained least-squares problem

$$
\begin{array}{ll}
\text{minimize} & \sum_{t=1}^{T} \|v_t\|_2^2 + \lambda \|w_t\|_2^2 \\
\text{subject to} & x_{t+1} = Ax_t + Bw_t, \quad t = 1, \ldots, T-1 \\
& y_t = Cx_t + v_t, \quad t = 1, \ldots, T-1
\end{array}
$$

- variables are $x_1, \ldots, x_T,\ w_1, \ldots, w_T,\ v_1, \ldots, v_T$
- $\lambda > 0$ is a parameter, trades off measurement and process errors

Running example

- we'll use a running example of a vehicle moving in 2-D
- $x_t = (p_t, v_t)$
  - 2-vector p_t is the 2-D position
  - 2-vector v_t is the 2-D velocity
- dynamics

$$p_{t+1} = p_t + v_t, \quad v_{t+1} = v_t + w_t$$

- so w_t is the force on the vehicle, which we generate
- $y_t = p_t + v_t$, where we generate measurement noise v_t *randomly*
- $T = 100$
- since we generate the data, we know the true value of $x_1, x_2, \ldots$ (in real applications, of course, you don't know the true state)

Measurements and true positions

- y_t shown as red crosses
- Cx_t ('true' position) shown in green

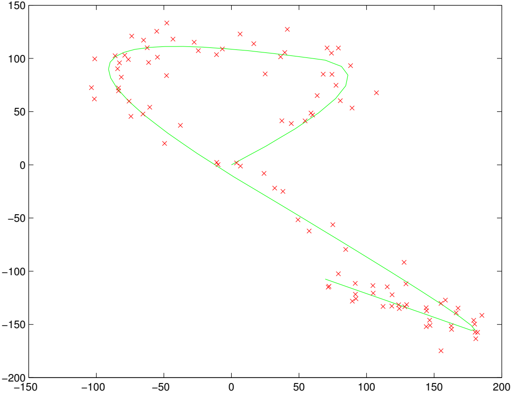

Least-squares state estimation

- blue curve is $C\hat{x}_t$, for $\lambda = 0.07$

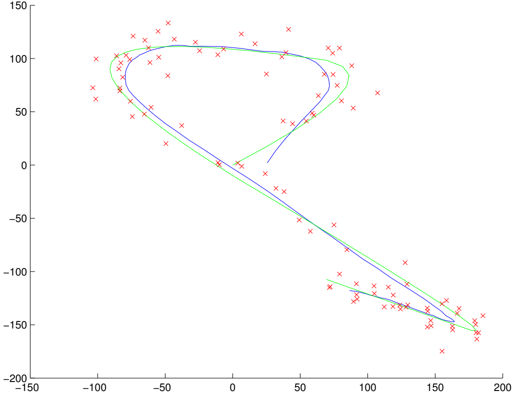

True and estimated state trajectories

- solid lines show true trajectory $(x_t)_i$
- dashed lines shows estimates $(\hat{x}_t)_i$

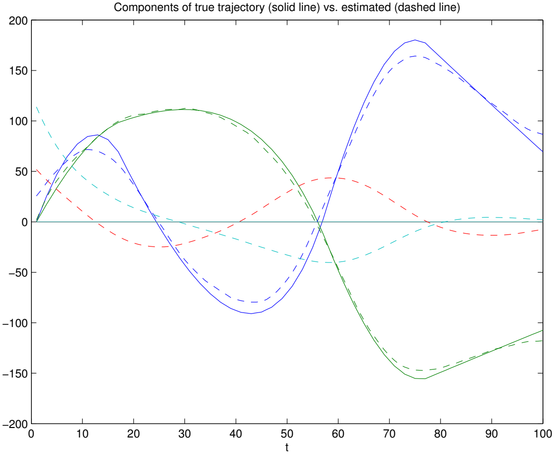

Varying λ

- repeat with $\lambda = 0.001, 0.01, 0.1, 1$

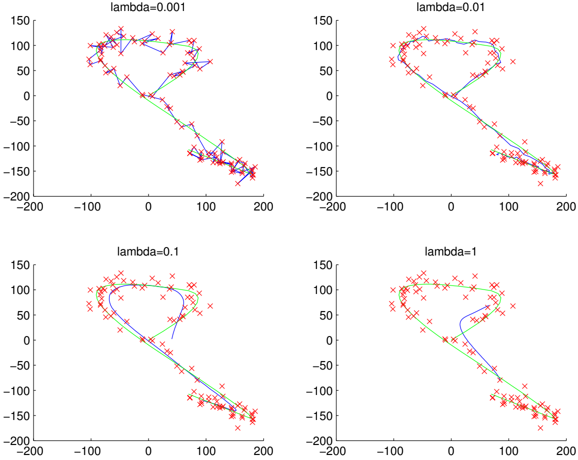

Varying λ

- objective is $J_1 + \lambda J_2$, $J_1 = \sum_{t=1}^{T} \|v_t\|_2^2$, $J_2 = \sum_{t=1}^{T} \|w_t\|_2^2$
- plot below is over range $0.01 \leq \lambda \leq 1$

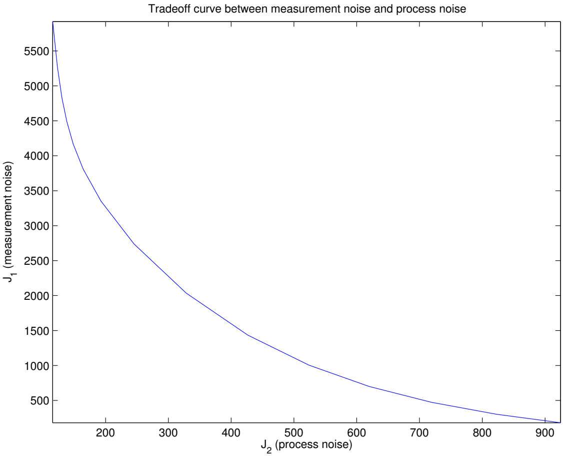

Outline

Least-squares state estimation

Missing measurements

Missing measurements

- suppose we only get measurements y_t for $t \in \mathcal{T} \subseteq \{1, \ldots, T-1\}$
- we form state estimate by solving problem

$$\begin{array}{ll} \text{minimize} & \sum_{t \in \mathcal{T}} \|v_t\|_2^2 + \lambda \sum_{t=1}^{T} \|w_t\|_2^2 \\ \text{subject to} & x_{t+1} = Ax_t + Bw_t, \quad t = 1, \ldots, T-1 \\ & y_t = Cx_t + v_t, \quad t \in \mathcal{T} \end{array}$$

  with variables $x_1, \ldots, x_T$, $w_1, \ldots, w_T$, and v_t for $t \in \mathcal{T}$
- a linearly constrained least-squares problem

Example

- same example, 20% of measurements removed
- blue curve shows $C\hat{x}_t$, $\lambda = 8$
- black curve shows estimate using all measurements

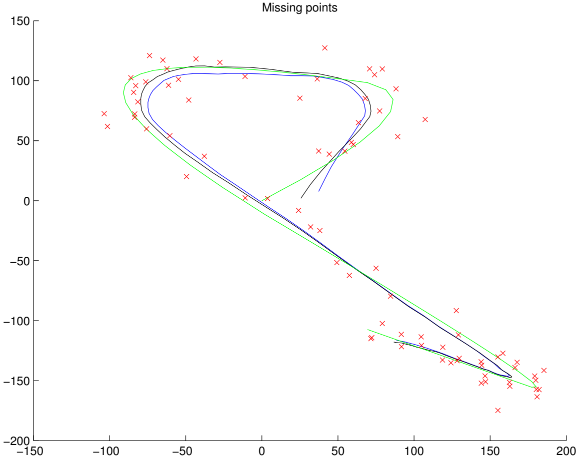

Cross validation

- randomly remove 20% (say) of the measurements and use as *test measurements*
- for many values of λ
  - carry out state estimation using other (*training*) measurements
  - evaluate RMS measurement residuals on test set
- choose λ to (approximately) minimize the RMS test residuals

Example

XXX plot shows training RMS and test RMS values versus lambda (which is on log plot)

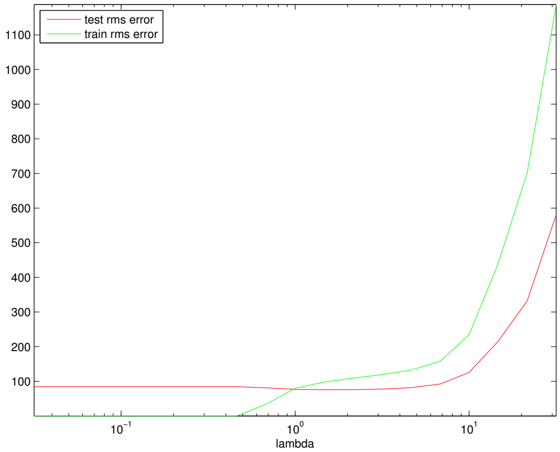

Example

- result shown by $\lambda = 4.64$
- training y_t shown as red crosses, test y_t shown as red circles