Nonlinear Moving Horizon State Estimation (NMHE)

James B. Rawlings

Department of Chemical Engineering

University of Freiburg
Freiburg, Germany
July 19–20, 2022

Outline

1. Full information estimation
   - Motivational example
   - Disturbances and robust stability
2. Moving horizon estimation
   - Arrival cost and prior weighting
   - Robust stability
3. Concluding remarks

State Estimation

Now turn to the general problem of estimating the state of a **noisy dynamic system** given **noisy measurements**:

$$\begin{aligned} x^+ &= f(x, w) \\ y &= h(x) + v \end{aligned} \tag{1}$$

in which the process disturbance, w, measurement disturbance, v, and system initial state, $x(0)$, are independent random variables with stationary probability densities.

Can consider constraints on state, measurement, process disturbance, and measurement disturbance.

$$x \in \mathbb{X}, \quad y \in \mathbb{Y}, \quad w \in \mathbb{W}, \quad v \in \mathbb{V}$$

Full Information Estimation

- Full information estimation will prove to have the best theoretical properties in terms of stability and optimality.
- Unfortunately, it will also prove to be computationally intractable except for the simplest cases, such as a linear system model.
- One method for practical estimator design therefore is to come as close as possible to the properties of full information estimation while maintaining a tractable online computation (MHE).

Variables

Notation required to distinguish the system variables from the estimator variables:

	System variable	Decision variable	Optimal decision
state	x	χ	$\hat{x}$
process disturbance	w	ω	$\hat{w}$
measured output	y	η	$\hat{y}$
measurement disturbance	v	ν	$\hat{v}$

The relationships between these variables:

$$\begin{aligned} x^+ &= f(x, w) & y &= h(x) + v & & \\ \chi^+ &= f(\chi, \omega) & y &= h(\chi) + \nu & \eta &= h(\chi) \\ \hat{x}^+ &= f(\hat{x}, \hat{w}) & y &= h(\hat{x}) + \hat{v} & \hat{y} &= h(\hat{x}) \end{aligned}$$

Note that $h(x) - h(\hat{x}) = -(v - \hat{v})$ (useful later)

Full Information Estimation

The **full information** objective function is

$$V_T(\chi(0), \boldsymbol{\omega}) = \ell_x\big(\chi(0) - \overline{x}_0\big) + \sum_{i=0}^{T-1} \ell_i(\omega(i), \nu(i)) \tag{2}$$

subject to

$$\chi^+ = f(\chi, \omega) \qquad y = h(\chi) + \nu$$

in which T is the current time, $y(i)$ is the measurement at time i, and $\overline{x}_0$ is the prior information on the initial state.
The full information estimator is then defined as the solution to

$$\min_{\chi(0), \boldsymbol{\omega}} V_T(\chi(0), \boldsymbol{\omega}) \tag{3}$$

Constraints

- For estimation problem, some physically known facts should also be enforced such as:
  - Concentrations of impurities must be nonnegative,
  - Fluxes of mass must have the correct sign given concentration gradients.
  - Fluxes of energy must have the correct sign given temperature gradients.
  - ...
- However, unlike the regulator, the estimator has no way to enforce these constraints on the *system*.
- It is important that any constraints imposed on the estimator are satisfied by the system generating the measurements.

Motivational Example—Is online optimization useful?

MHE Performance

Consider the following set of reversible reactions taking place in a well-stirred, isothermal, gas-phase batch reactor

$$A \underset{k_{-1}}{\overset{k_1}{\rightleftharpoons}} B + C \qquad 2B \underset{k_{-2}}{\overset{k_2}{\rightleftharpoons}} +C$$

The states are the concentrations of species in mol/L and the measurement is the reactor pressure in atm

$$x = \begin{bmatrix} c_A \\ c_B \\ c_C \end{bmatrix} \qquad y = RT \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} x$$

We assume the ideal gas law in modeling the pressure

Material balances

Material balances lead to the following nonlinear state space model:

$$\frac{d}{dt}\begin{bmatrix} c_A \\ c_B \\ c_C \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 1 & -2 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} k_1 c_A - k_{-1} c_B c_C \\ k_2 c_B^2 - k_{-2} c_C \end{bmatrix}$$

$$\frac{dx}{dt} = f_c(x)$$

Note that $c_B c_C$ and c_B^2 are the only nonlinearities

Discrete time model

Model the system plus disturbances with the following discrete time model

$$x^{+} = f(x) + w$$
$$y = Cx + v$$

in which f is the solution of the ODEs over the sample time, Δ, i.e, if $s(t, x_0)$ is the solution of $\frac{dx}{dt} = f_c(x)$ with initial condition $x(0) = x_0$ at $t = 0$, then $f(x) = s(\Delta, x)$.

The state and measurement disturbances, w and v, are assumed to be zero-mean independent normals with constant covariances Q and R.

Parameter values

The following parameter values are used in the simulations

$$RT = 32.84\,\mathrm{mol}\cdot\mathrm{atm/L}$$

$$\Delta = 0.25 \qquad k_1 = 0.5 \qquad k_{-1} = 0.05 \qquad k_2 = 0.2 \qquad k_{-2} = 0.01$$

$$P(0) = (0.5)^2 I \qquad Q = (0.001)^2 I \qquad R = (0.25)^2$$

$$\bar{x}(0) = \begin{bmatrix} 1 \\ 0 \\ 4 \end{bmatrix} \qquad x(0) = \begin{bmatrix} 0.5 \\ 0.05 \\ 0 \end{bmatrix}$$

The prior density for the initial state, $N(x(0), P(0))$, is deliberately chosen to poorly represent the actual initial state to model a large initial disturbance to the system.

Examine how MHE recovers from this large unmodeled disturbance.

Example Results

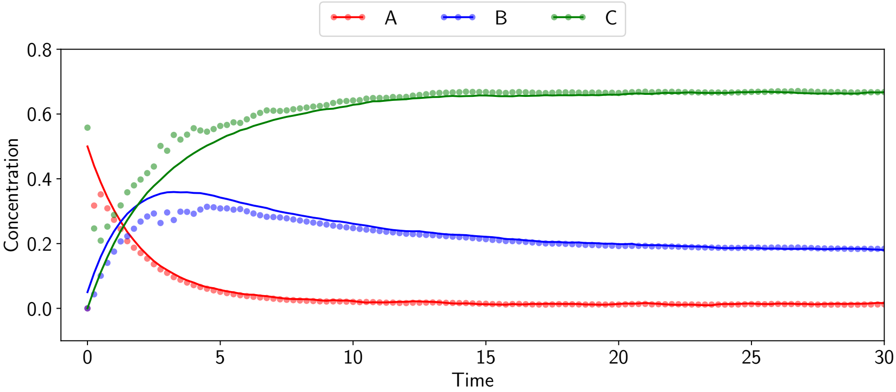

Figure 1: Example results for MHE on batch reactor system.

EKF Performance

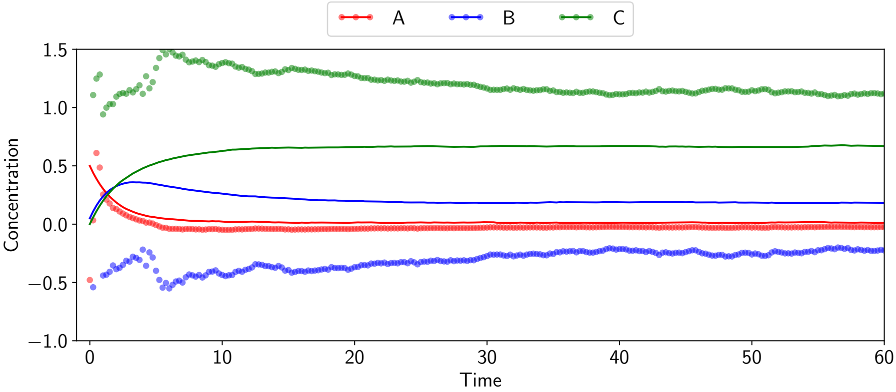

Figure 2: Poor performance of EKF on batch reactor system.

EKF Performance with Clipping

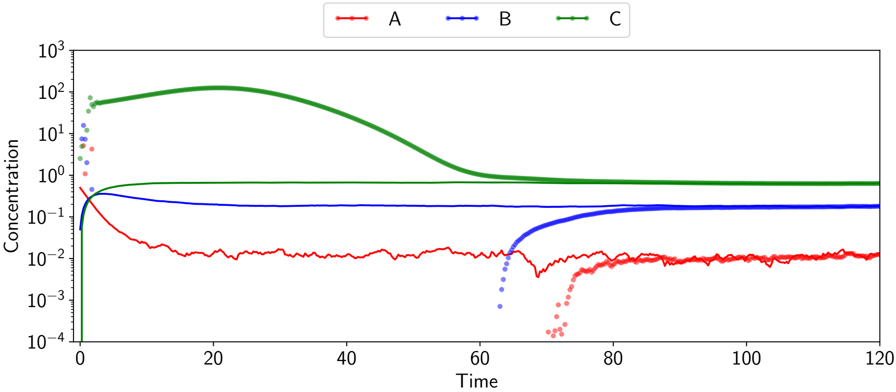

Figure 3: Poor performance of EKF with clipping on batch reactor system. Any negative concentration estimates are clipped to 0.

UKF Performance

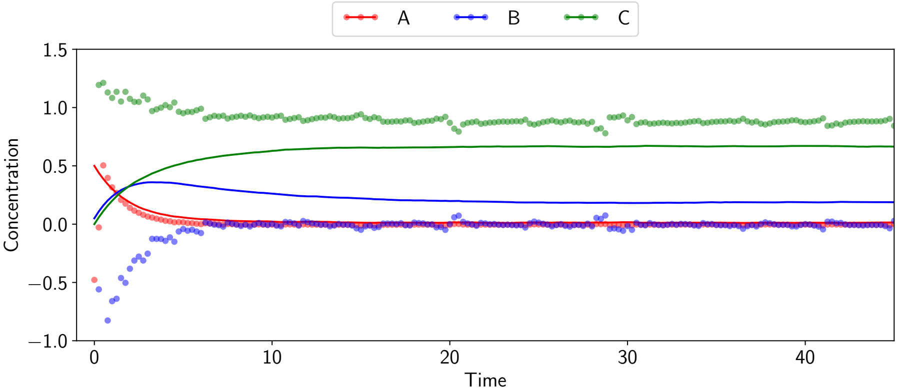

Figure 4: Poor performance of UKF on batch reactor system.

UKF Performance with Scaling

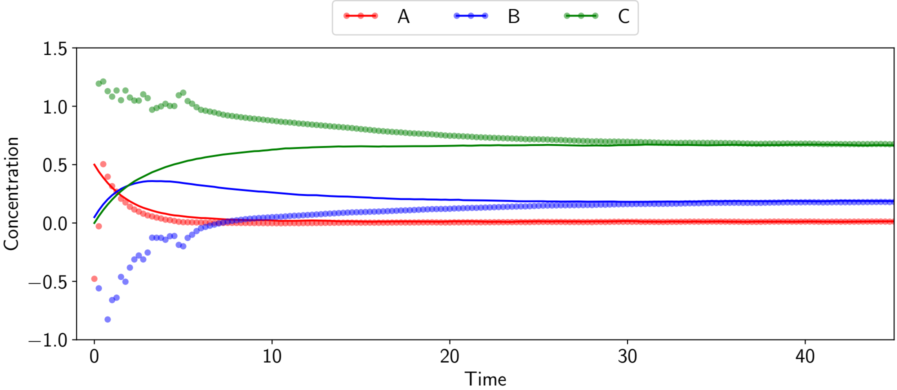

Figure 5: Mediocre performance of UKF with scaling on batch reactor system. When possible, negative concentrations at sigma points are scaled to be nonnegative.

Let's back up and review the linear theory

- Linear model, $x^+ = Ax + Gw, \quad y = Cx + v$.
  Full information estimation is equivalent to the Kalman filter.
- If (A, C) is *detectable* and (A, G) is stabilizable, $Q, R > 0$, then there exists $P^- \geq 0$ satisfying the discrete algebraic Riccati equation (DARE)

$$P^- = G'QG + AP^-A' - AP^-C'(CP^-C' + R)^{-1}CP^-A'$$

- The matrix $A - \tilde{L}C$ is a stable matrix

$$\tilde{L} = AP^-C'(CP^-C' + R)^{-1}$$

- The steady-state estimator takes the form

$$\hat{x}^+ = A\hat{x} + \tilde{L}(y - C\hat{x}) \qquad \hat{x}(0) = \overline{x}_0$$

Estimate error

- Subtract the estimator from the system

$$
\begin{aligned}
x^+ &= Ax + Gw \quad y = Cx + v \\
\hat{x}^+ &= A\hat{x} + \tilde{L}(y - C\hat{x}) \\
(x - \hat{x})^+ &= A(x - \hat{x}) + Gw - \tilde{L}(Cx + v - C\hat{x}) \\
(x - \hat{x})^+ &= (A - \tilde{L}C)(x - \hat{x}) + Gw - \tilde{L}v \\
\tilde{x}^+ &= A_L\tilde{x} + Gw - \tilde{L}v \qquad \tilde{x}(0) = x(0) - \overline{x}_0
\end{aligned}
$$

  and $A_L = A - \tilde{L}C$ is a stable matrix.
- Solve the linear system for estimate error

$$
\tilde{x}(k) = A_L^k(x(0) - \overline{x}_0) + \sum_{j=0}^{k-1} A_L^{k-j-1}(Gw(j) - Lv(j))
$$

Bound on estimate error—Linear case

- Since A_L is stable, there exist $c > 0, \lambda < 1$

$$\left| A_L^k \right| \leq c\lambda^k \qquad \text{for all } k \geq 0$$

- Take norm of estimate error solution

$$|\tilde{x}(k)| \leq c\lambda^k |x(0) - \overline{x}_0| + c\sum_{j=0}^{k-1}(|G|\,|w(j)| + |L|\,|v(j)|)\lambda^{k-j-1}$$

- Taking the largest disturbance terms outside and performing the sum then gives

$$|\tilde{x}(k)| \leq c\lambda^k |x(0) - \overline{x}_0| + \frac{c}{1-\lambda}\left[\,|G|\,\|\boldsymbol{w}\|_{0:k-1} + |L|\,\|\boldsymbol{v}\|_{0:k-1}\,\right]$$

using the sup norm over the sequence, $\|\boldsymbol{w}\|_{0:k-1} := \max_{j \in 0:k-1} |w(j)|$

$$\boxed{|\tilde{x}(k)| \leq c\lambda^k |x(0) - \overline{x}_0| + c_w \|\boldsymbol{w}\|_{0:k-1} + c_v \|\boldsymbol{v}\|_{0:k-1}}$$

What happens in the nonlinear theory?

- Linear model $\rightarrow$ nonlinear model, $x^+ = f(x, w) \quad y = h(x) + v$
- Detectable (A, C) $\rightarrow$ i-IOSS (Sontag and Wang, 1997)
- Stabilizable (A, G) $\rightarrow$ incremental stabilizability
- $w'Q^{-1}w + v'R^{-1}v \rightarrow \ell(w, v)$
- $Q, R > 0 \rightarrow$ stage cost: $\ell(w, v)$ underbounded by K-function
- Establishing the linear system's bound on estimate error $\rightarrow$ Establishing the RGAS property for the nonlinear system

Robust global asymptotic stability

Definition 1 (Robustly globally asymptotically stable estimation)

The estimate is based on the *noisy* measurement $\boldsymbol{y} = h(\boldsymbol{x}(x_0, \boldsymbol{w})) + \boldsymbol{v}$.
The estimate is RGAS if there exist functions $\alpha(\cdot) \in \mathcal{KL}$ and $\delta_w(\cdot) \in \mathcal{K}$ such that for all x_0 and $\overline{x}_0$, and bounded $(\boldsymbol{w}, \boldsymbol{v})$, the following holds for all $k \in \mathbb{I}_{\geq 0}$

$$|x(k; x_0, \boldsymbol{w}) - x(k; \hat{x}(0|k), \widehat{\boldsymbol{w}}_k)| \leq \alpha(|x_0 - \overline{x}_0|, k) + \delta_w(\|(\boldsymbol{w}, \boldsymbol{v})\|_{0:k-1})$$

Compare the linear system result

$$\boxed{|\tilde{x}(k)| \leq c\lambda^k |x(0) - \overline{x}_0| + c_w \|\boldsymbol{w}\|_{0:k-1} + c_v \|\boldsymbol{v}\|_{0:k-1}}$$

Recall the on-line optimization problem

The **full information** objective function is

$$V_T(\chi(0), \boldsymbol{\omega}) = \ell_x\big(\chi(0) - \overline{x}_0\big) + \sum_{i=0}^{T-1} \ell_i(\omega(i), \nu(i)) \qquad (4)$$

subject to

$$\chi^+ = f(\chi, \omega) \qquad y = h(\chi) + \nu$$

The full information estimator is then defined as the solution to

$$\min_{\chi(0), \boldsymbol{\omega}} V_T(\chi(0), \boldsymbol{\omega}) \qquad (5)$$

What class of nonlinear models?

Assumption 2 (Continuity)

The functions $f(\cdot)$, $h(\cdot)$, $\ell_x(\cdot)$, and $\ell(\cdot)$ are continuous, $\ell_x(0) = 0$, and $\ell(0, 0) = 0$. The sets $\mathbb{X}$ and $\mathbb{W}$ are closed.

Detectability of nonlinear systems and i-IOSS

What class of systems have a stable state estimator?

- Assume system observability?
  Too restrictive for even linear systems (recall the definition of detectability).
- We need a similar detectability definition for nonlinear systems – **i-IOSS**:

Definition 3 (i-IOSS)

The system $x^+ = f(x, w), y = h(x)$ is *incrementally input/output-to-state stable* (i-IOSS) if there exists some $\beta(\cdot) \in \mathcal{KL}$ and $\gamma_1(\cdot)$, $\gamma_2(\cdot) \in \mathcal{K}$ such that for every two initial states z_1 and z_2, and any two disturbance sequences $\boldsymbol{w}_1$ and $\boldsymbol{w}_2$

$$|x(k; z_1, \boldsymbol{w}_1) - x(k; z_2, \boldsymbol{w}_2)| \leq \beta(|z_1 - z_2|, k) + \gamma_1\left(\|\boldsymbol{w}_1 - \boldsymbol{w}_2\|_{0:k-1} \right) + \gamma_2\left(\|h(\boldsymbol{x}_1) - h(\boldsymbol{x}_2)\|_{0:k} \right)$$

Stabilizability of nonlinear systems

Definition 4 (Incremental Stabilizability with respect to stage cost $L(\cdot)$)

A nonlinear system $x^+ = f(x, u)$ is said to be *incrementally stabilizable with respect to stage cost* $L(\cdot)$ if there exists $\mathcal{K}$-function $\overline{\alpha}$ such that for every two initial conditions $x_1, x_2 \in \mathbb{X}$ and control sequence $\boldsymbol{w}_1 \in \mathbb{W}^\infty$, another control sequence $\boldsymbol{w}_2 \in \mathbb{W}^\infty$ exists such that

$$\sum_{k=0}^{\infty} L(x_1(k), x_2(k), w_1(k), w_2(k)) \leq \overline{\alpha}(|x_1 - x_2|)$$

System stabilizability and detectability

Assumption 5 (Stabilizability)

The system (1) is stabilizable with respect to the stage cost $L(x_1, x_2, w_1, w_2) := \ell(w_2 - w_1, h(x_1) - h(x_2))$.

Assumption 6 (Detectability)

The system (1) is i-IOSS.

What stage and initial cost?

Assumption 7 (Positive-definite stage cost)

The stage cost $\ell(\cdot)$ satisfies

$$\sigma_w(|\omega|) + \sigma_v(|\nu|) \leq \ell(\omega, \nu) \leq \overline{\sigma}_w(|\omega|) + \overline{\sigma}_v(|\nu|)$$

for all $\omega \in \mathbb{W}, \nu \in \mathbb{V}$ for some $\mathcal{K}_\infty$-functions $\overline{\sigma}_w$ and $\overline{\sigma}_v$. [a]
Furthermore, we have that

$$\sigma_x(|\chi - \overline{x}_0|) \leq \ell_x(\chi - \overline{x}_0) \leq \overline{\sigma}_x(|\chi - \overline{x}_0|)$$

for all $\chi, \overline{x}_0 \in \mathbb{X}$ for some $\mathcal{K}_\infty$-functions σ_x and $\overline{\sigma}_x$.

[a] Note: the lower bounding $\mathcal{K}_\infty$-functions σ_w and σ_v come from the i-IOSS Lyapunov function implied by assuming i-IOSS.

Stability

Zero error system

First consider the zero estimate error solution for all $k \geq 0$ (initial state is equal to the estimator's prior and zero disturbances). In this case, the optimal solution is:

$$\begin{aligned} \hat{x}(0|T) &= \overline{x}_0 \\ \hat{w}(i|T) &= 0 \quad \text{for all} \quad 0 \leq i \leq T, T \geq 1 \\ h(\hat{x}(i|T)) &= y(i) \quad \text{for all} \quad 0 \leq i \leq T, T \geq 1 \end{aligned}$$

The **perturbation** to this solution are: the system's initial state (distance from $\overline{x}_0$), and the process and measurement disturbances.
We next define stability properties so that:

- *asymptotic stability* considers the case $x_0 \neq \overline{x}_0$ with zero disturbances.
- *robust stability* considers the case in which $(w(i), v(i)) \neq 0$.

Nominal estimator stability—zero disturbances

Theorem 8 (Stability of full information estimation)

Let Assumptions 2, 7, 5, and 6 hold. Then full information estimation is GAS.

In other words

$$|x(k; x_0, 0) - x(k; \hat{x}(0|k), \widehat{\boldsymbol{w}}_k)| \leq \alpha(|x_0 - \overline{x}_0|, k)$$

for all $x_0, \overline{x}_0 \in \mathbb{X}$ and $k \geq 0$. Note: $(\boldsymbol{w}, \boldsymbol{v}) = 0$.

Effect of disturbances on stability of FIE

We adjust the assumption on the stage cost and stabilizability assumption to account for the nonzero (w, v) (Rawlings, Mayne, and Diehl, 2020, p. 284). We then have the following result

Theorem 9 (Robust stability of full information estimation)

1. *Let Assumptions 2, 6, and modified stage cost and incremental stabilizability assumptions hold. Then full information estimation is RGAS.*
2. *Let Assumptions 2 and exponential versions of stage cost, stabilizability, and detectability assumptions hold. Then full information estimation is RGES.*

The proof for RGES is given in (Allan and Rawlings, 2021, Theorem 3.16). The considerably more involved proof for RGAS is given in (Allan, 2020, Theorem 5.18).

From full information to MHE

- The full information problem becomes intractable as time T increases.
- So we use a moving horizon approximation with horizon N to bound the computation.

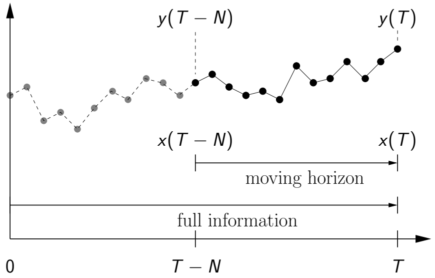

MHE problem statement

- In MHE we consider only **the N most recent measurements**, $\boldsymbol{y}_N(T) = (y(T-N), y(T-N+1), \dots y(T-1))$.
- For $T > N$, the MHE problem is defined to be:

$$\min_{\chi(T-N),\boldsymbol{\omega}} \hat{V}_T(\chi(T-N),\boldsymbol{\omega}) = \Gamma_{T-N}(\chi(T-N)) + \sum_{i=T-N}^{T-1} \ell_i(\omega(i),\nu(i)) \tag{6}$$

subject to:

$$\begin{aligned} \chi^+ &= f(\chi,\omega) \\ y &= h(\chi) + \nu \\ \boldsymbol{\omega} &= (\omega(T-N),\dots,\omega(T-1)) \end{aligned}$$

Prior Information

- The designer chooses the prior weighting $\Gamma_k(\cdot)$ for $k > N$ until the data horizon is full.
- For times $T \leq N$, we generally **define** the MHE problem to be *the full information problem*.

Zero Prior Weighting

When we choose $\Gamma_i(\cdot) = 0$ for all $i \geq N$:

$$\hat{V}_T(\chi(T-N), \boldsymbol{\omega}) = \sum_{i=T-N}^{T-1} \ell_i(w(i), v(i))$$

- Because it discounts the past data completely, this form of MHE must be able to asymptotically reconstruct the state using only the most recent N measurements.

Nonzero Prior Weighting

There are two drawbacks to zero prior weighting:

- The system has to be assumed *observable* rather than detectable to ensure existence of the solution to the MHE problem.
- A large horizon N may be required to obtain performance comparable to full information estimation.

We address these disadvantages by using **nonzero prior weighting**:

$$\min_{\chi(T-N),\boldsymbol{\omega}} \hat{V}_T(\chi(T-N),\boldsymbol{\omega}) = \Gamma_{T-N}(\chi(T-N)) + \sum_{i=T-N}^{T-1} \ell_i(\omega(i),\nu(i))$$

Full information arrival cost

Definition 10 (Full information arrival cost)

The full information arrival cost is defined as

$$Z_T(p) = \min_{\chi(0),\boldsymbol{\omega}} V_T(\chi(0), \boldsymbol{\omega}) \tag{7}$$

subject to

$$\chi^+ = f(\chi, \omega) \qquad y = h(\chi) + \nu \qquad \chi(T; \chi(0), \boldsymbol{\omega}) = p$$

Here forward DP is used to decompose the full information problem exactly into the MHE problem (6) in which $\Gamma(\cdot)$ is chosen as arrival cost.

Lemma 11 (MHE and full information estimation)

The MHE problem (6) *is equivalent to the full information problem* (5) *for the choice* $\Gamma_k(\cdot) = Z_k(\cdot)$ *for all* $k > N$ *and* $N \geq 1$.

Prior weighting

Assumption 12 (MHE Prior weighting bounds)

We assume that $\Gamma_k(\cdot)$ is continuous and satisfies the following upper and lower bounds uniformly in k for all $\chi, \hat{x}(k|k) \in \mathbb{X}$

$$\underline{c}_\Gamma \left| p - \hat{x}(k|k) \right|^\sigma \leq \Gamma_k(p) \leq \overline{c}_\Gamma \left| p - \hat{x}(k|k) \right|^\sigma \qquad (8)$$

in which $\sigma \geq 1$

Prior weighting—Sufficient condition for stability

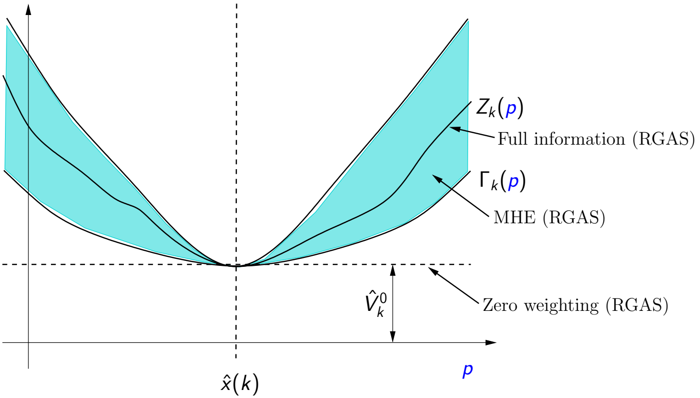

Robust stability of MHE—Recent results

Theorem 13 (MHE is RGES)

Let Assumptions 2 and exponential versions of stage cost, stabilizability, and detectability assumptions hold. Let Assumption 12 on the prior weighting hold.
Then there exists a horizon length $\underline{N}$ such that MHE is RGES for all $N \geq \underline{N}$.

See (Allan and Rawlings, 2021, Theorem 4.2) for precise statement and proof of this result.

Some open questions

- Is MHE RGAS if the system is asymptotically (rather than exponentially) i-IOSS?
- What are the best methods to update the MHE initial penalty, $\Gamma_k(\cdot)$ to obtain an accurate estimator with a small horizon N for computational efficiency?

The statistical viewpoint

- System model

$$x^+ = f(x, w) \qquad y = h(x) + v$$

- Disturbances, w and v, initial state, $x(0)$, modeled as random variables. The central limit theorem justifies using (zero mean) normal distributions for w, v. Obtaining variances part of the modeling/identification problem!
- We observe $\mathbf{y}(k) := y(0), y(1), \ldots, y(T)$ and wish to estimate $x(k)$.
- Statistically optimal estimate; maximize conditional density

$$\max_{x(k)} p(x(k) \mid \mathbf{y}(k))$$

- Computing p usually intractable for nonlinear models

State estimation — Why use optimization?

The engineering viewpoint

- Reduce goal from statistical optimality to practical real time algorithm.
- Choose a merit function

$$V_T(x(0), \boldsymbol{w}) = \ell_x\big(x(0) - \overline{x}_0\big) + \sum_{i=0}^{T-1} \ell_i(w(i), v(i))$$

subject to model

$$x^+ = f(x, w) \qquad y = h(x) + v$$

- Full information estimator is

$$\min_{x(0), \boldsymbol{w}} V_T(x(0), \boldsymbol{w})$$

- Can easily add knowledge (constraints) on w, v, x to the formulation.

Given all this extra generality, what do we lose?

- Lose the analytical, recursive solution.
  Least squares problem $\rightarrow$ nonconvex optimization problem
- Lose the strict equivalence between the the full information problem and its moving horizon approximation.
  But expect similar behavior for large horizon, N
- Because of this online computational complexity, the big divide in state estimation is between linear and nonlinear models

Current status of nonlinear etiimation

- The theory perspective
  - We have reasonable theory for full information estimation.
  - By *extension*, we can expect reasonable theoretical properties for MHE.
  - The "theory" for EKF remains: if the system is (almost) linear, the EKF works. But the practice is more encouraging.
- The practice perspective
  - Try the EKF! It usually (sometimes, never) works!
  - If that fails, try MHE, it usually (always, sometimes) works! But it's a slog to compute.
  - If that fails, try full information, it always (usually, sometimes) works! Good luck managing the computation.

Another computational example

Predator-Prey Model

Let x_1 = number of untagged prey, x_2 be number of tagged prey, and x_3 number of predator.

$$\frac{dx_1}{dt} = -ax_3x_1 + b(x_1 + x_2) - ux_1$$
$$\frac{dx_2}{dt} = -ax_3x_2 + ux_1 \qquad\qquad y = x_1/x_2$$
$$\frac{dx_3}{dt} = cx_3(x_1 + x_2) - dx_3$$

Parameters: predation rate a, prey reproduction rate b, predator reproduction rate c, and predator death rate d. The control u is the rate at which prey are tagged.
We measure the *fraction* of tagged prey that are observed, i.e., not total number.

MHE Example Results

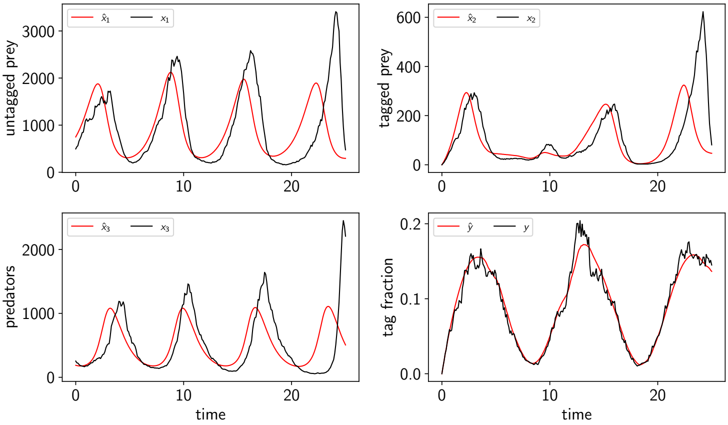

Figure 6: Example results for MHE on predator-prey dynamics.

EKF Example Results

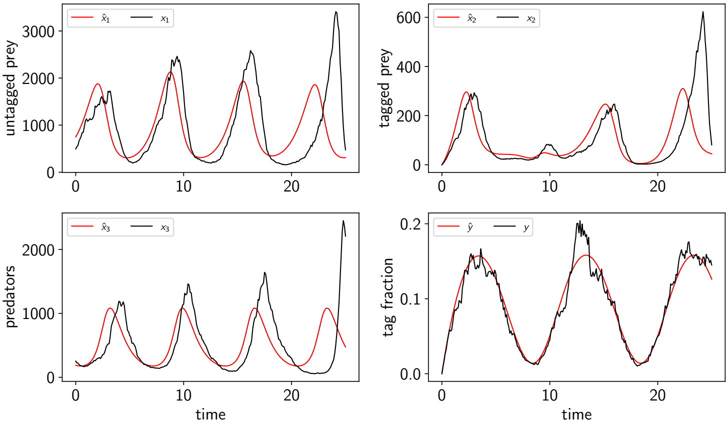

Figure 7: Example results for EKF on predator-prey dynamics.

Recommended exercises

- Observability, detectability, i-IOSS. Exercises 4.1, 4.2, 4.3, 4.4, 4.5, 4.8, 4.9, 4.11, 4.14.[1]
- Estimator convergence. Exercises 4.18, 4.19.[1]

[1] Rawlings, Mayne, and Diehl (2020, Chapter 4). Downloadable from `engineering.ucsb.edu/~jbraw/mpc`.

Further reading I

D. A. Allan. *A Lyapunov-like Function for Analysis of Model Predictive Control and Moving Horizon Estimation*. PhD thesis, University of Wisconsin–Madison, August 2020. URL https://engineering.ucsb.edu/~jbraw/jbrweb-archives/theses/allan.pdf.

D. A. Allan and J. B. Rawlings. Robust stability of full information estimation. *SIAM J. Cont. Opt.*, 59(5):3472–3497, 2021.

J. B. Rawlings, D. Q. Mayne, and M. M. Diehl. *Model Predictive Control: Theory, Design, and Computation*. Nob Hill Publishing, Santa Barbara, CA, 2nd, paperback edition, 2020. 770 pages, ISBN 978-0-9759377-5-4.

E. D. Sontag and Y. Wang. Output-to-state stability and detectability of nonlinear systems. *Sys. Cont. Let.*, 29:279–290, 1997.

Simulation exercise: Hot air balloon autopilot using NMHE/NMPC

Example 14

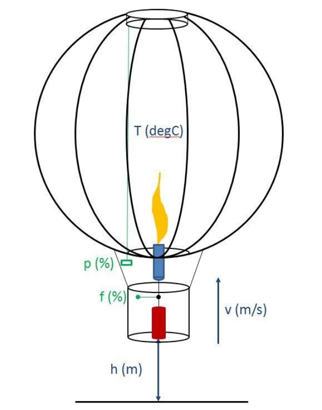

We consider the problem of designing an auto-pilot for a hot-air balloon. For this problem the MVs are fuel valve position $f(\%)$ and top vent position $p(\%)$. The CVs are altitude $h(m)$, vertical velocity $v(m/s)$, and the temperature of the air inside the balloon $T(^{\circ}C)$.

Hot air balloon model

Here we use a first-principles hot-air balloon model based on a vertical force balance, a standard model of the atmosphere, and an energy balance on the balloon. This model is used to simulate the process and is also used directly by the NMHE and NMPC algorithms.

$$m_b \frac{d^2 h}{dt^2} = (\rho_s - \rho) V g - m_b g - k_d \left(\frac{dh}{dt} \right) \left| \frac{dh}{dt} \right|$$

$$\rho = \frac{MP}{RT}; T_s = T_0 - ah; \frac{\rho_s}{\rho_0} = \left(\frac{T_s}{T_0} \right)^{\gamma - 1}; \frac{P_s}{P_0} = \left(\frac{T_s}{T_0} \right)^{\gamma}$$

$$\rho V C_p \frac{dT}{dt} = -UA(T - T_s) + \epsilon \Delta H_c f - C_p (T - T_s) k_v p$$

Hot air balloon exercise - Closed-loop flight

Now we'll fly the balloon with a NMHE/NMPC autopilot.

1. Start up Octave and run mpcsim_setup.m
2. Run hab_mpcsim.m
3. Toggle the Controller switch to ON and press the Play button.
4. On the CV menu set the h setpoint to 2000m.
5. Let the simulation run until the altitude stabilizes at the new setpoint. Note that the autopilot adjusts the fuel valve and the vent valve simultaneously to achieve the desired altitude setpoint. And note again that the the fuel valve oscillates to take out the oscillation in h.
6. Increase the h setpoint to 4000m. Note that the fuel required for this second increment of 2000m in altitude is significantly more than the fuel required for the first increment. Likewise, the increase in balloon temperature is significantly higher when going from 2000m to 4000m than it was when going from the ground to 4000m.

Hot air balloon exercise - Closed-loop, cont.

1. Increase the h setpoint to 6000m. The autopilot is unable to reach the new altitude setpoint because the balloon temperature has reached it's maximum limit of 120 degC. Note that this limit is now respected both dynamically and at steady-state. This is one of the main advantages of using NMHE/NMPC for the autopilot.
2. On the CV menu, set the minimum limit of v to -5m/s. Now try landing the balloon by entering a h setpoint of 0m. Note that the autopilot closes the fuel valve f gradually to maintain a constant descent of -5m/s, followed by a gentle landing.
3. Now try a flight on your own. Happy landings!

Hot air balloon exercise - Closed-loop, discrete fuel

Now we will try the NMHE/NMPC autopilot with a discrete fuel valve.

1. Start up a fresh copy of the mpcsim hot air balloon simulation
2. On the Options menu, select Fuel increment and enter a value of 1. This tells the autopilot that the valve can only be fully open or closed.
3. Toggle the Controller switch to ON and press the Play button.
4. As we did before, enter a h setpoint of 2000m.
5. Let the simulation run until the altitude stabilizes at the new setpoint. Note the discrete nature of the f signal.
6. Continue the flight on your own.