Information-seeking polynomial NARX model-predictive control through expected free energy minimization

Wouter M. Kouw

American Control Conference 2024

Bayesian Intelligent Autonomous Systems lab | Department of Electrical Engineering

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

What is this paper about?

An intelligent autonomous agent that balances ..

1. **.. driving the system to a goal and ..**
2. **.. acquiring data for maximally efficient system identification.**

Why?

- **Autonomous calibration of control system.**
- **Cautious exploration of input-output function.**
- **Robustness to time-varying disturbances.**

Problem statement

Consider a system with control input $u_k \in \mathcal{U} \subseteq \mathbb{R}$ and measured output $y_k \in \mathbb{R}$.

- The system's dynamics are unknown but are assumed to be continuous.
- Noise is unknown but is assumed to be continuous and symmetric.
- Disturbances are unknown but are assumed to be continuous.

Goal: reach neighbourhood of a point: $y_* \sim \mathcal{N}(m_*, v_*)$

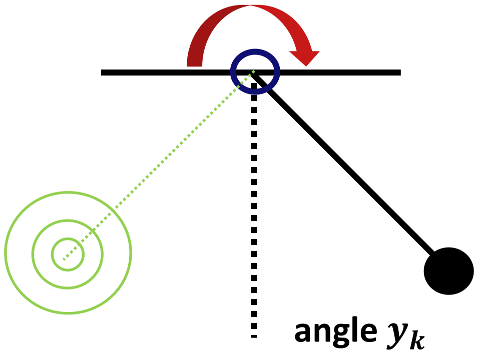

Demonstration

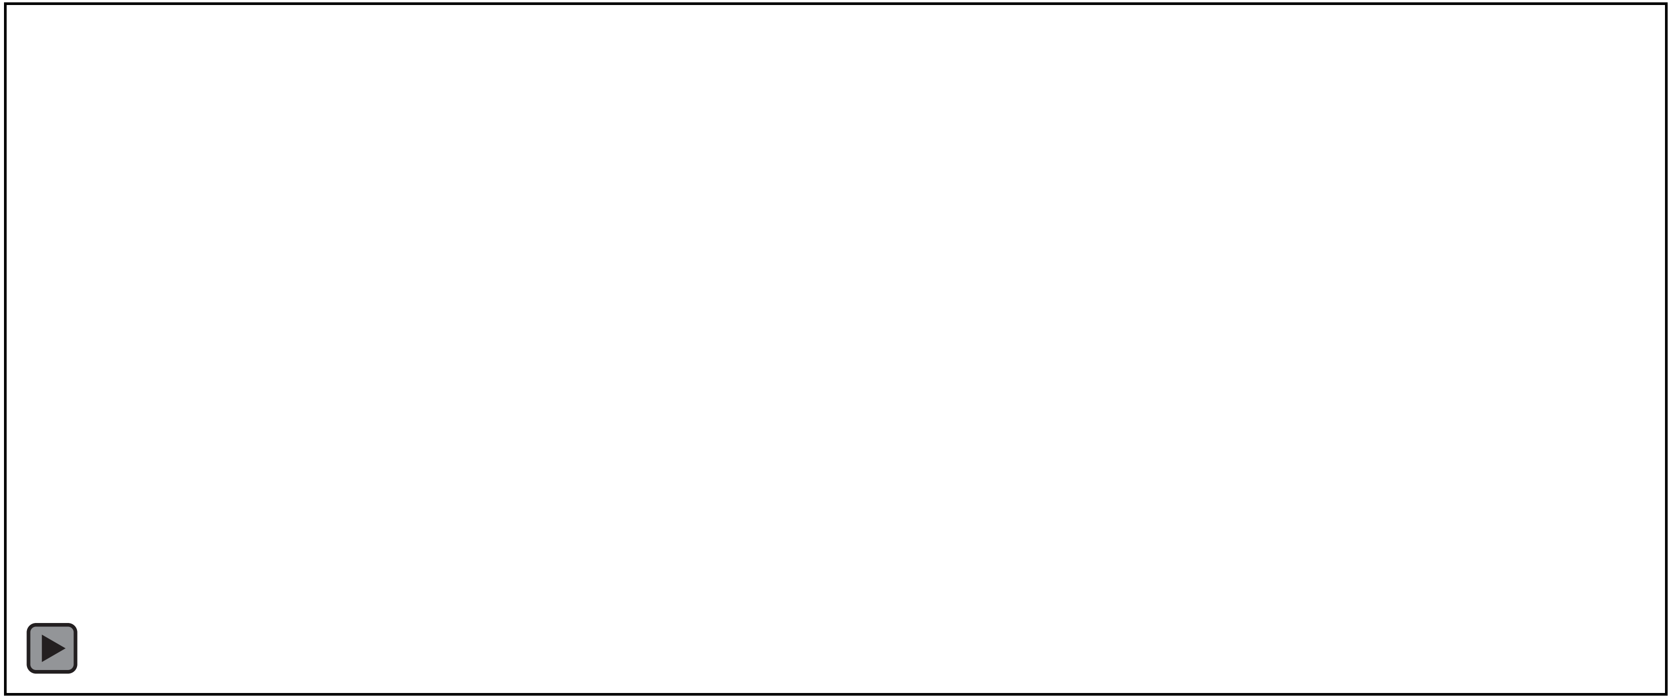

Model Specification

We adopt a nonlinear autoregressive with exogenous input (NARX) model:

$$p(y_k \mid \theta, \tau, u_k) = \mathcal{N}(y_k \mid \theta^T \phi(x_k, u_k), \tau^{-1})$$

Autoregression coefficients (unknown)

Basis expansion function (known)

Data buffer (known)

Noise precision parameter (unknown)

Example data buffer, for $M_y = 2, M_u = 1$:

$$x_k = [y_{k-1} \;\; y_{k-2} \;\; u_{k-1}]$$

Example basis function, order-2 polynomial without cross-terms:

$$\phi(x_k, u_k) = [1 \;\; y_{k-1} \;\; y_{k-2} \;\; u_{k-1} \;\; y_{k-1}^2 \;\; y_{k-2}^2 \;\; u_{k-1}^2]$$

Inference: parameter estimation

Prior distribution over autoregression coefficients θ and noise precision τ:

$$p(\theta, \tau) = p(\theta \mid \tau) p(\tau) = \mathcal{N}(\theta \mid \mu_0, (\tau\Lambda)^{-1})\, \mathcal{G}(\tau \mid \alpha, \beta)$$

Mean vector — Precision matrix — Shape parameter — Rate parameter

Bayes' rule to calculate posterior distribution:

$$p(\theta, \tau \mid \mathcal{D}_k) = \frac{p(y_k \mid \theta, \tau, u_k)}{p(y_k \mid u_k, \mathcal{D}_{k-1})} p(\theta, \tau \mid \mathcal{D}_{k-1})$$

Recursive parameter update rules:

$$\mu_k = \Lambda_k^{-1}(\phi_k y_k + \Lambda_{k-1}\mu_{k-1}) \qquad \alpha_k = \alpha_{k-1} + \frac{1}{2}$$

$$\Lambda_k = \Lambda_{k-1} + \phi_k \phi_k^T \qquad \beta_k = \beta_{k-1} + \frac{1}{2}\left(y_k^2 - \mu_k^T \Lambda_k \mu_k + \mu_{k-1}^T \Lambda_{k-1} \mu_{k-1}\right)$$

Inference: control estimation

Transitioning the probabilistic model forward in time gives:

$$p(y_t, u_t, \theta, \tau \mid \mathcal{D}_k) = p(y_t \mid u_t, \theta, \tau)p(\theta, \tau \mid \mathcal{D}_k)p(u_t)$$

We want to infer a marginal posterior distribution for u_t, but exact inference is not possible.

We shall adopt an expected free energy functional:

$$\mathcal{F}_k[q] = \int q(y_t, u_t, \theta, \tau) \ln \frac{q(y_t, u_t, \theta, \tau)}{p(y_t, u_t, \theta, \tau \mid \mathcal{D}_k)} d(y_t, u_t, \theta, \tau)$$

with variational model:

$$q(y_t, u_t, \theta, \tau) = p(y_t \mid u_t, \theta, \tau)p(\theta, \tau \mid \mathcal{D}_k)q(u_t)$$

TU/e

Inference: control estimation

We adopt MAP estimation, i.e., the most probable value under the control posterior:

$$u_t^* = \arg\max\ q(u_t)$$

If we aim for MAP, then the expected free energy function can be written as (see paper):

$$\mathcal{J}(u_t) = \mathbb{E}_{q(y_t,u_t,\theta,\tau)}\left[\ln \frac{q(y_t,u_t,\theta,\tau)}{p(\theta,\tau \mid \mathcal{D}_k)p(y_t \mid u_t)}\right] + \mathbb{E}_{p(y_t \mid u_t)}[\ln p(y_t \mid y_*)]$$

Mutual information:
predicted future output and parameter beliefs

Cross entropy:
difference predictions and goal

goal-seeking

Working out expectations yield:

$$\mathcal{J}(u_t) = C - \frac{1}{2}\ln\left(\phi_t^T \Lambda_k^{-1} \phi_t + 1\right) + \frac{1}{2v_*}\left(\frac{\beta_k}{\alpha_k - 1}\left(\phi_t^T \Lambda_k^{-1} \phi_t + 1\right) + \left(\mu_k^T \phi_t - m_*\right)^2\right)$$

information-seeking

Experiments

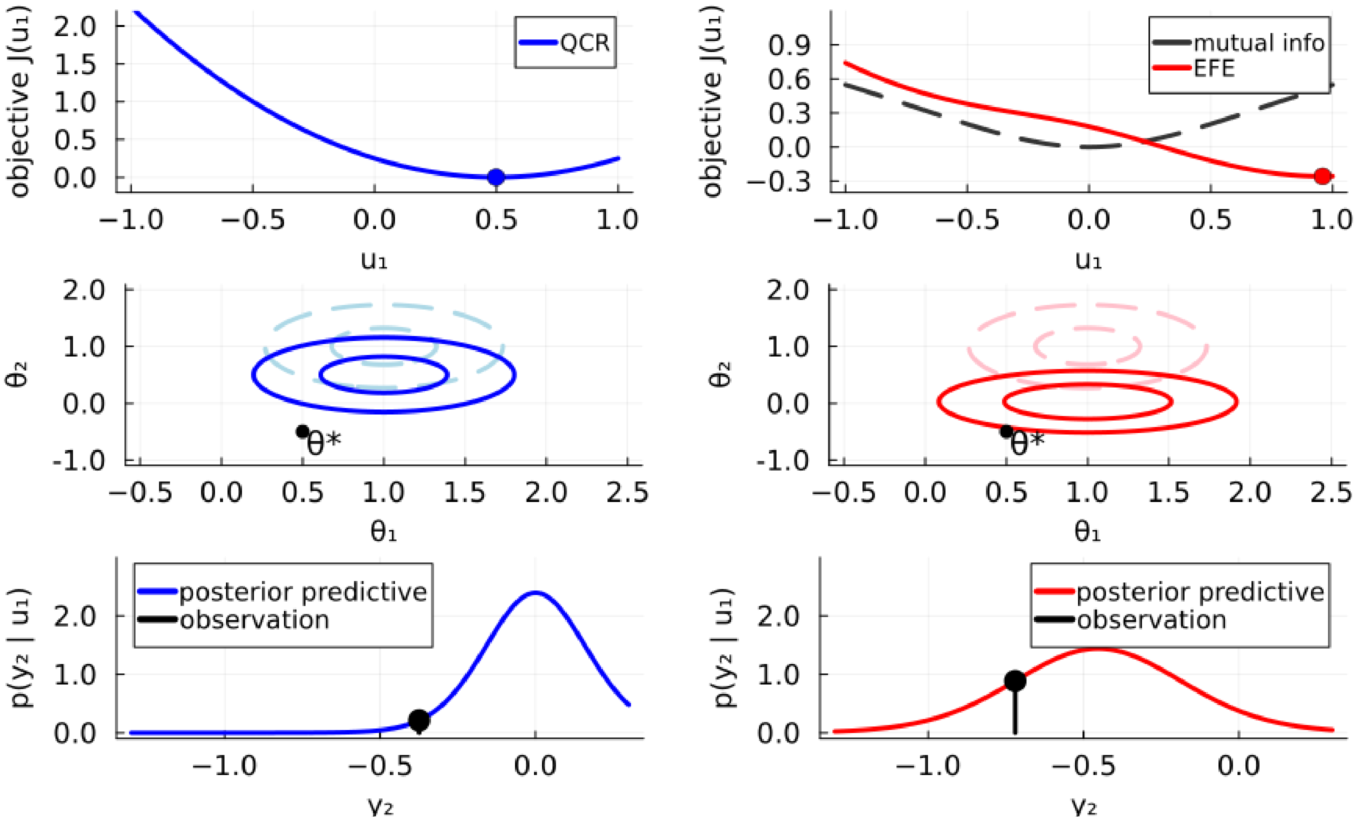

Discussion

Value:

- **Goal prior variance balances information- and goal-seeking.**
- **Caution: large values of control signal are avoided when parameter uncertainty is high.**

Limitations:

- **Variance parameter is dropped when filling the buffer x_t, which means uncertainty does not accumulate in the feedforward model.**

TU/e

Take-aways

Given a probabilistic model, one can derive an information-theoretic control objective that ..

1. **.. plans a control policy (given parameterized prediction) to reach a goal and ..**
2. **.. plans an action to maximize mutual information between parameters and predictions.**

Github repository

BIASlab