

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

Wouter M. Kouw

Abstract—We propose an adaptive model-predictive controller that balances driving the system to a goal state and seeking system observations that are informative with respect to the parameters of a nonlinear autoregressive exogenous model. The controller’s objective function is derived from an expected free energy functional and contains information-theoretic terms expressing uncertainty over model parameters and output predictions. Experiments illustrate how parameter uncertainty affects the control objective and evaluate the proposed controller for a pendulum swing-up task.

I. INTRODUCTION

Nonlinear autoregressive exogenous (NARX) models are useful black-box descriptions of dynamical systems [1]. However, a model-predictive controller that relies on a NARX model that has yet to accurately capture dynamics can behave erratically, which may lead to damage during operation [2]. To avoid such risks, in practice, model parameters are often first estimated using a designed control signal (i.e., offline system identification). However, this is time-consuming. Here we present a model-predictive controller (also referred to as an agent) that seeks control inputs for rapid online system identification, when its parameter uncertainty and predictive variance are high. When these are low, it seeks control inputs that lead to the goal output.

The proposed agent utilizes variational Bayesian inference, a form of statistical inference that optimizes an analytic approximation to posterior distributions [3]. Specifically, it minimizes an expected free energy (EFE) functional, a quantity proposed in the neuroscience community that describes perception and action in terms of information-theoretic quantities (also known as active inference) [4]. Free energy minimization is a relatively new framework for control, but it has strong ties to stochastic optimal control and encompasses many classical methods, such as proportional-integrative-derivative and linear-quadratic Gaussian control [5], [6]. So far, EFE-based controllers have relied on (partially) known dynamics or neural networks [7], [8], [9]. We propose an EFE minimizing agent based on a polynomial NARX model. NARX models do not require knowledge of system dynamics and require far fewer parameters than neural networks. Furthermore, free energy minimization in autoregressive models has been shown to produce superior parameter estimates at small sample sizes with more stable k -step ahead output predictions [10], [11], [12]. Our work falls within the scope of information-theoretic model-predictive

control [13], [14], but we focus on information gathering specifically for online NARX model identification.

Our contributions include:

- The derivation of parameter update rules for a conjugate prior to the Gaussian NARX likelihood.
- The derivation of a location-scale T-distributed posterior predictive distribution for outputs given control inputs.
- The derivation of an objective function for an adaptive model-predictive controller that dynamically balances information-seeking and goal-seeking behaviour.

II. PROBLEM STATEMENT

Consider an agent, operating in discrete time, that receives noisy outputs $y_k \in \mathbb{R}$ from a system and sends control inputs $u_k \in \mathbb{R}$ back. It must drive the system to a desired output y_* without knowledge of the system’s dynamics.

III. MODEL SPECIFICATION

We propose an agent with a probabilistic model that factorizes according to:

$$p(y_{1:N}, u_{1:N}, \theta, \tau) = p(\theta, \tau) \prod_{k=1}^N p(y_k | \theta, \tau, u_k) p(u_k), \quad (1)$$

for N time steps. The likelihood is auto-regressive in nature,

$$p(y_k | \theta, \tau, u_k) \triangleq \mathcal{N}(y_k | \theta^\top \phi(x_k, u_k), \tau^{-1}), \quad (2)$$

with coefficients $\theta \in \mathbb{R}^D$ and measurement noise precision $\tau \in \mathbb{R}^+$ [11]. The variable x_k represents the collection of M_y past outputs and M_u past control inputs, $x_k \triangleq (y_{k-1}, \dots, y_{k-M_y}, u_{k-1}, \dots, u_{k-M_u})$. The function ϕ performs a polynomial basis expansion, $\phi : \mathbb{R}^M \rightarrow \mathbb{R}^D$, where $M = M_y + M_u + 1$ and D is the number of terms in the polynomial. For example, a second-order polynomial basis without cross-terms with $M_y = M_u = 1$ would produce $\phi(y_{k-1}, u_{k-1}, u_k) = [1 \ y_{k-1} \ u_{k-1} \ u_k \ y_{k-1}^2 \ u_{k-1}^2 \ u_k^2]$.

The prior distribution on the parameters is a multivariate Gaussian - univariate Gamma distribution [15, ID: D5]:

$$p(\theta, \tau) \triangleq \mathcal{NG}(\theta, \tau | \mu_0, \Lambda_0, \alpha_0, \beta_0) \quad (3)$$

$$= \mathcal{N}(\theta | \mu_0, (\tau \Lambda_0)^{-1}) \mathcal{G}(\tau | \alpha_0, \beta_0). \quad (4)$$

This choice of parameterization allows for inferring an exact posterior distribution (Lemma III.1) and an exact posterior predictive distribution (Lemma III.2).

The prior distributions over control inputs are assumed to be independent over time:

$$p(u_k) \triangleq \mathcal{N}(u_k | 0, \eta^{-1}), \quad (5)$$

with precision parameter η . This choice has a regularizing effect on the inferred controls (Sec. IV-B).

WM Kouw is with the Bayesian Intelligent Autonomous Systems laboratory, Department of Electrical Engineering, TU Eindhoven, Postal box 513, 5600 MB Eindhoven, Netherlands. @: w.m.kouw@tue.nl

IV. INFERENCE

A. Parameter estimation

First, we note that, at time k , the control u_k has been executed and is known to the agent. Henceforth, we shall use \hat{u}_k and \hat{y}_k to differentiate observed variables from unobserved ones, i.e., u_k and y_k .

The parameter posterior distribution is obtained by Bayesian filtering [16]:

$$p(\theta, \tau | \mathcal{D}_k) = \underbrace{\frac{p(\hat{y}_k | \theta, \tau, \hat{u}_k)}{p(\hat{y}_k | \hat{u}_k, \mathcal{D}_{k-1})}}_{\text{evidence}} \underbrace{p(\theta, \tau | \mathcal{D}_{k-1})}_{\text{prior}}. \quad (6)$$

where $\mathcal{D}_k = \{\hat{y}_i, \hat{u}_i\}_{i=1}^k$ is the data up to time k . The evidence (a.k.a. marginal likelihood) is

$$p(\hat{y}_k | \hat{u}_k, \mathcal{D}_{k-1}) = \int p(\hat{y}_k | \theta, \tau, \hat{u}_k) p(\theta, \tau | \mathcal{D}_{k-1}) d(\theta, \tau). \quad (7)$$

This integral is typically intractable. In earlier work, we obtained an approximate posterior distribution through mean-field variational inference [11]. In this model we obtain an exact posterior distribution using the multivariate Gaussian - univariate Gamma prior distribution specified in Eq. 3.

Lemma 4.1: The Gaussian distributed NARX likelihood in Eq. 2 combined with a multivariate Gaussian - univariate Gamma prior distribution over AR coefficients θ and measurement precision τ ,

$$p(\theta, \tau | \mathcal{D}_{k-1}) = \mathcal{NG}(\theta, \tau | \mu_{k-1}, \Lambda_{k-1}, \alpha_{k-1}, \beta_{k-1}), \quad (8)$$

yields a multivariate Gaussian - univariate Gamma posterior distribution

$$p(\theta, \tau | \mathcal{D}_k) = \mathcal{NG}(\theta, \tau | \mu_k, \Lambda_k, \alpha_k, \beta_k). \quad (9)$$

The parameters of the posterior distribution are

$$\begin{aligned} \mu_k &= (\phi_k \phi_k^\top + \Lambda_{k-1})^{-1} (\phi_k \hat{y}_k + \Lambda_{k-1} \mu_{k-1}), \\ \Lambda_k &= \phi_k \phi_k^\top + \Lambda_{k-1}, \quad \alpha_k = \alpha_{k-1} + \frac{1}{2}, \\ \beta_k &= \beta_{k-1} + \frac{1}{2} (\hat{y}_k^2 - \mu_k^\top \Lambda_k \mu_k + \mu_{k-1}^\top \Lambda_{k-1} \mu_{k-1}), \end{aligned} \quad (10)$$

where $\phi_k \triangleq \phi(\hat{x}_k, \hat{u}_k)$.

The proof is in Appendix A. The marginal posterior distributions are Gamma distributed and multivariate location-scale T-distributed [15, ID: P36]:

$$p(\tau | \mathcal{D}_k) = \int p(\theta, \tau | \mathcal{D}_k) d\theta = \mathcal{G}(\tau | \alpha_k, \beta_k), \quad (11)$$

$$p(\theta | \mathcal{D}_k) = \int p(\theta, \tau | \mathcal{D}_k) d\tau = \mathcal{T}_{2\alpha_k}(\theta | \mu_k, \frac{\beta_k}{\alpha_k} \Lambda_k^{-1}). \quad (12)$$

The subscript under \mathcal{T} refers to its degrees of freedom.

B. Control estimation

In order to effectively drive the system to the goal output, the agent must make accurate predictions for future outputs. We express the probability of the output and parameters at time $t = k + 1$ given a future control input as

$$p(y_t, \theta, \tau | u_t) = p(y_t | \theta, \tau, u_t) p(\theta, \tau | \mathcal{D}_k). \quad (13)$$

The dependence on \mathcal{D}_k on the left-hand side is left implicit for the sake of brevity. Note that this probability distribution does not yet include a goal output. Without it, the agent will drive the system towards *any* output that leads to minimal prediction error. To incorporate the goal output, we first decompose the joint distribution above to:

$$p(y_t, \theta, \tau | u_t) = p(\theta, \tau | y_t, u_t) p(y_t). \quad (14)$$

Then, we constrain the marginal prior distribution over future output, $p(y_t)$, to a specific functional form with chosen parameters:

$$p(y_t | y_*) \triangleq \mathcal{N}(y_t | m_*, v_*). \quad (15)$$

This constrained distribution on the future output is known as a *goal prior* distribution [4].

Using Eqs. 14, 15 and Bayes' rule, the posterior distribution for the controls is:

$$p(u_t | y_t, \theta, \tau) = \frac{p(\theta, \tau | y_t, u_t) p(y_t | y_*) p(u_t)}{\int p(\theta, \tau | y_t, u_t) p(y_t | y_*) p(u_t) du_t}. \quad (16)$$

The integral in the denominator is challenging to solve due to the basis function applied to u_t inside the likelihood. Instead, we introduce a variational distribution $q(u_t)$ to approximate the posterior. The approximation error is characterized with an expected free energy functional [4], [17],

$$\mathcal{F}_t[q] \triangleq \mathbb{E}_{q(y_t, \theta, \tau, u_t)} \left[\ln \frac{p(\theta, \tau | \mathcal{D}_k) q(u_t)}{p(\theta, \tau | y_t, u_t) p(y_t | y_*) p(u_t)} \right], \quad (17)$$

where the expectation is over

$$q(y_t, \theta, \tau, u_t) \triangleq p(y_t, \theta, \tau | u_t) q(u_t). \quad (18)$$

Inferring the optimal control at time t refers to minimizing the free energy functional with respect to the variational distribution $q(u_t)$:

$$q^*(u_t) = \arg \min_{q \in Q} \mathcal{F}_t[q]. \quad (19)$$

where Q represents the space of candidate distributions. We can re-arrange the free energy functional to simplify the variational minimization problem:

$$\mathbb{E}_{q(y_t, u_t, \theta, \tau)} \left[\ln \frac{p(\theta, \tau | \mathcal{D}_k) q(u_t)}{p(\theta, \tau | y_t, u_t) p(y_t | y_*) p(u_t)} \right] = \quad (20)$$

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

Now we define the expected free energy *function*:

$$\mathcal{J}(u_t) \triangleq \mathbb{E}_{p(y_t, \theta, \tau | u_t)} \left[\ln \frac{p(\theta, \tau | \mathcal{D}_k)}{p(\theta, \tau | y_t, u_t) p(y_t | y_*)} \right]. \quad (21)$$

Using $\mathcal{J}(u_t) = \ln(1/\exp(-\mathcal{J}(u_t)))$, the expected free energy functional can be expressed as a Kullback-Leibler divergence

$$\mathcal{F}_t[q] = \mathbb{E}_{q(u_t)} \left[\ln \frac{q(u_t)}{\exp(-\mathcal{J}(u_t))p(u_t)} \right], \quad (22)$$

which is minimal when [3]:

$$q^*(u_t) = \exp(-\mathcal{J}(u_t))p(u_t). \quad (23)$$

Thus, we have an optimal approximate posterior distribution over controls.

To proceed, we must solve the expectation in (21). We have access to the parametric forms of all distributions involved except for the distribution over parameters given the future output and control. It can be tied to known distributions through Bayes' rule:

$$p(\theta, \tau | y_t, u_t) = \frac{p(y_t | \theta, \tau, u_t) p(\theta, \tau | \mathcal{D}_k)}{\int p(y_t | \theta, \tau, u_t) p(\theta, \tau | \mathcal{D}_k) d(\theta, \tau)}. \quad (24)$$

The distribution that results from the marginalization in the denominator is the posterior predictive distribution $p(y_t | u_t)$ and can be derived analytically within our model.

Lemma 4.2: The marginalization of the NARX likelihood (Eq. 2) over the parameter posterior distribution (Eq. 9) yields a location-scale T-distribution:

$$p(y_t | u_t) = \int p(y_t | \theta, \tau, u_t) p(\theta, \tau | \mathcal{D}_k) d(\theta, \tau) \quad (25)$$

$$= \int \mathcal{N}(y_t | \theta^\top \phi_t, \tau^{-1}) \mathcal{NG}(\theta, \tau | \mu_k, \Lambda_k, \alpha_k, \beta_k) d(\theta, \tau) \quad (26)$$

$$= \mathcal{T}_{\nu_t}(y_t | m_t, s_t^2), \quad (27)$$

where $\phi_t \triangleq \phi(\hat{x}_t, u_t)$ and

$$\nu_t \triangleq 2\alpha_k, \quad m_t \triangleq \mu_k^\top \phi_t, \quad s_t^2 \triangleq \frac{\beta_k}{\alpha_k} (\phi_t^\top \Lambda_k^{-1} \phi_t + 1). \quad (28)$$

The proof is in Appendix B.

If we replace $p(\theta, \tau | y_t, u_t)$ in Eq. 21 with the right-hand side of Eq. 24 and use Eq. 13, then the EFE function can be split into two components:

$$\mathcal{J}(u_t) = \mathbb{E}_{p(y_t, \theta, \tau | u_t)} \left[\ln \frac{1}{p(y_t | y_*)} \right] \quad (29)$$

$$+ \mathbb{E}_{p(y_t, \theta, \tau | u_t)} \left[\ln \frac{p(\theta, \tau | \mathcal{D}_k) p(y_t | u_t)}{p(y_t | \theta, \tau, u_t) p(\theta, \tau | \mathcal{D}_k)} \right] \\ = \mathbb{E}_{p(y_t | u_t)} \left[-\ln p(y_t | y_*) \right] \quad (30)$$

$$- \mathbb{E}_{p(y_t, \theta, \tau | u_t)} \left[\ln \frac{p(y_t, \theta, \tau | u_t)}{p(\theta, \tau | \mathcal{D}_k) p(y_t | u_t)} \right].$$

One may recognize the first term as a cross-entropy, describing the dissimilarity between the predictive distribution and the goal prior distribution [3]. The second term is the mutual information between the parameter posterior and the posterior predictive distribution. It describes how much information is gained on the parameters upon measuring a system output [3]. Solving the expectations yields a compact objective:

Theorem 4.3: The expected free energy function in Eq. 30 evaluates to:

$$\mathcal{J}(u_t) = \frac{1}{2\nu_*} \left((\mu_k^\top \phi_t - m_*)^2 + \frac{\beta_k}{\alpha_k} (\phi_t^\top \Lambda_k^{-1} \phi_t + 1) \frac{2\alpha_k}{2\alpha_k - 2} \right) \\ - \frac{1}{2} \ln \left(\phi_t^\top \Lambda_k^{-1} \phi_t + 1 \right) + \text{constants}. \quad (31)$$

The proof is found in Appendix C.

We consider MAP estimation of the control policy because, although it is informative to quantify uncertainty over controls, we can ultimately only execute one action. The MAP estimate can be expressed as a minimization over a negative logarithmic transformation of $q^*(u_t)$ (Eq. 23):

$$\hat{u}_t = \arg \max_{u_t \in \mathcal{U}} q^*(u_t) = \arg \min_{u_t \in \mathcal{U}} \mathcal{J}(u_t) - \ln p(u_t). \quad (32)$$

The \mathcal{U} refers to the space of affordable controls and serves to incorporate practical constraints such as torque limits.

So far, we have only considered a 1-step ahead prediction. Generalizing to $t > k + 1$ is, in principle, straightforward. Due to the independence assumptions on the prior $p(u_t)$ and the variational control posteriors $q(u_t)$, the joint variational control posterior distribution factorizes over time:

$$q^*(u_t, \dots, u_{t+T}) = \prod_{t=1}^T p(u_t) \exp(-\mathcal{J}(u_t)). \quad (33)$$

If we apply the same negative logarithmic transformation as in Eq. 32 to MAP estimation for the joint variational control posterior, then the final control optimization problem becomes:

$$\hat{u}^{\text{EFE}} = \arg \min_{u \in \mathcal{U}^T} \sum_{t=1}^T \frac{1}{2\nu_*} (\mu_k^\top \phi(\hat{x}_t, u_t) - m_*)^2 \quad (34) \\ + \frac{\beta_k}{\nu_* (2\alpha_k - 2)} \left((\phi(\hat{x}_t, u_t)^\top \Lambda_k^{-1} \phi(\hat{x}_t, u_t) + 1) \right) \\ - \frac{1}{2} \ln \left(\phi(\hat{x}_t, u_t)^\top \Lambda_k^{-1} \phi(\hat{x}_t, u_t) + 1 \right) + \eta u_t^2.$$

where $u = (u_t, \dots, u_T)$ and \mathcal{U}^T is the Cartesian product of T spaces over affordable controls.

However, this generalization to $t > k + 1$ assumes that all elements of x_t have been observed. The predicted y_t is typically incorporated into x_{t+1} as if it were an observation [1], [18]. However, from a Bayesian perspective, y_t is not a number but a random variable. This causes a problem because the delay vector x_t would also become a random variable for $t > k + 1$, and would interfere with the conditioning and marginalization operations described earlier. We avoid this problem by collapsing the posterior predictive distribution for y_t (Eq. 27) into its most probable value $\mu_k^\top \phi(\hat{x}_t, \hat{u}_t)$, and incorporating that number into x_{t+1} . This is a suboptimal solution because the predictive variance will not accumulate over the horizon.

The optimization problem in Eq. 34 can be solved by iterative methods with automatic differentiation for obtaining the gradient. The box constraints over affordable controls \mathcal{U} can be incorporated by utilizing an interior-point method.

V. EXPERIMENTS

A. Seeking informative observations

The more informative a system input-output observation, the faster the agent is able to accurately estimate parameters and make accurate predictions. We shall compare an agent based on the EFE objective in Eq.34 (referred to as EFE) with an agent that uses the typical quadratic cost with regularization objective (referred to as QCR) [2]:

$$\hat{u}^{\text{QCR}} = \arg \min_{u \in \mathcal{U}^T} \sum_{t=1}^T (\mu_k^T \phi(\hat{x}_t, u_t) - m_*)^2 + \eta u_t^2. \quad (35)$$

QCR can be described as purely goal-seeking, because it only seeks to minimize the difference between the predictions and the goal. When comparing EFE with QCR we observe the effect of adding parameter uncertainty and predictive variance to the objective function.

Consider a system that evolves according to $y_k = \theta_1^* y_{k-1} + \theta_2^* u_k$ where $\theta_1^* = 0.5$ and $\theta_2^* = -0.5$. The task is to drive this system to $m_* = 0.5$ (for EFE, $v_* = 1$) under $\mathcal{U} = [-1, 1]$ for $T = 1$. We used Optim.jl to minimize the objective functions [19]. The models' prior parameters are set to shape $\alpha_0 = 10$, rate $\beta_0 = 1$, mean $\mu_0 = [1, 1]$, precision $\Lambda_0 = \frac{1}{2}I$ and control precision $\eta = 0$. The top row in Figure 1 compares the two objective functions (QCR left, EFE right) for the first action u_1 , with $\hat{u}_1^{\text{QCR}} = 0.5$ and $\hat{u}_1^{\text{EFE}} = 0.96$. This difference is caused by the mutual information term (dashed gray), which indicates that u_t 's further from 0 are more informative (note Eq. 34 uses *negative* mutual information, meaning the objective is smaller for more informative controls). After executing the respective \hat{u}_1 's, the systems return \hat{y}_1 's and the parameter distributions are updated. The middle row in Figure 1 shows contour levels of the marginal prior and posterior parameter distributions (see Eq. 12) over the $[\theta_1, \theta_2]$ plane. Note that the posterior under the EFE objective has moved closer to the system coefficients θ^* (i.e., $p(\theta^* | \mathcal{D}_1)$ is higher). The bottom row in Figure 1 show the posterior predictive distributions for y_2 and indicates that EFE's predictions are closer to the systems' \hat{y}_2 .

The information-seeking quality depends on parameter certainty, as can be seen by replicating the previous example with different values of Λ_0 . Figure 2 (left) shows that the mutual information term is flatter for $\Lambda_0 = 2I$. It has less of an effect on the overall EFE objective (compare with $\Lambda_0 = \frac{1}{2}I$ shown in Figure 1 top right) and the optimal control input shrinks to $u_1 = 0.75$. For $\Lambda_0 = 100I$ (see Figure 2 right), the mutual information term is so flat that the EFE objective resembles the QCR objective (see Figure 1 top left), and is also minimal at $u_1 = 0.5$.

B. Control under unknown dynamics

In this experiment, the controllers must swing up a damped pendulum and maintain an upright position¹. The system parameters are mass $m = 1.0$, length $\ell = 0.5$, friction $c = 0.01$ and time-step $\Delta t = 0.1$. The angle θ_t is observed

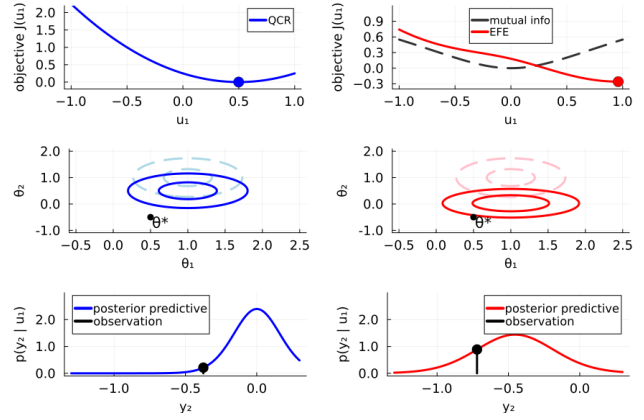


Fig. 1. Example of information-seeking. (Top row) Control objective functions for quadratic cost (QCR, left) vs proposed cost (EFE, right). Black dashed line displays mutual information, which causes $\hat{u}_1^{\text{EFE}} > \hat{u}_1^{\text{QCR}}$. (Middle row) Marginal prior (dashed) and posterior (solid) distributions for coefficients θ obtained after executing \hat{u}_1^{QCR} (left) and \hat{u}_1^{EFE} (right), and observing the resulting \hat{y}_1 's. Note that the posteriors under EFE controls are closer to the system AR coefficients θ^* . (Bottom row) Posterior predictive distributions for y_2 under QCR (left) and EFE (right) controls. Note that \hat{y}_2 's probability is larger under EFE, i.e., the prediction error is smaller.




Fig. 2. Examples of how parameter certainty affects the mutual information term and the resulting EFE control objective. For $\Lambda_0 = 2I$ (left; medium certainty) the minimizer shrinks to $u_1 = 0.75$ (compared to $\Lambda_0 = \frac{1}{2}I$, Figure 1 top right). For $\Lambda_0 = 100I$ (right; high certainty), the minimizer shrinks all the way to $u_1 = 0.5$ as in the QCR objective (Figure 1 top left).

under zero-mean Gaussian noise with a standard deviation of 0.001. The agent's NARX model consists of a second-order polynomial basis without cross-terms and with delays $M_y = 2$, $M_u = 2$. The prior distribution's parameters are shape $\alpha_0 = 10$, rate $\beta_0 = 0.1$, mean $\mu_0 = 10^{-8} \cdot [1, \dots, 1]$, coefficient precision $\Lambda = \frac{1}{2}I$ and control prior precision $\eta = 0.001$. These parameters would be regarded as weakly informative for noise precision τ and non-informative for coefficients θ . The goal prior has mean $m_* = \pi$ (i.e., upward) and variance $v_* = 0.5$. The top row of Figure 3 visualizes




Fig. 3. Pendulum angle over time during swing-up task under unknown dynamics. (Top left) QCR-based controller, (top right) EFE-based controller. (Bottom) Trials for PILCO-based controller.

¹Code at <https://github.com/biaslab/LCSS2024-NARXEFE>

the paths the two controllers took. The QCR-based controller does not initially move much as it cannot predict the effect of the control inputs under the non-informative prior for the coefficients. When it starts, it swings a few times before it reaches the goal and settles. The EFE-based controller moves away from zero earlier (i.e., it is seeking system input-output observations that produce more accurate output predictions) and swings up the pendulum sooner.

For the sake of comparison, we also report the performance of PILCO [20]. It is a Gaussian Process-based identification method and although it is one of the most data-efficient reinforcement learning methods, it still requires several trials before it finds a successful policy. The bottom row of Figure 3 shows its progress over the first 4 trials.

VI. DISCUSSION

Close inspection of Eq. 31 reveals that the balance between the mutual information, cross-entropy and prior distribution is also affected by the goal prior variance v_* . It may be possible to optimize this dynamically. Furthermore, note that the optimization problem in Eq. 34 is quadratic in ϕ_t . This means that for all nonlinear polynomials (i.e., degree larger than 1), the optimization becomes non-convex in u_t .

VII. CONCLUSION

We have proposed an adaptive NARX model-predictive controller that seeks informative observations for online system identification when its parameter uncertainty is high, and seeks control inputs that drive the system to a goal when its parameter uncertainty is low. This is achieved through a control objective with information-theoretic quantities, derived from the free energy minimization framework.

APPENDIX A

Proof: The posterior is proportional to the likelihood times the prior distribution:

$$p(\theta, \tau | \mathcal{D}_k) \propto \mathcal{N}(\hat{y}_k | \theta^\top \phi_k, \tau^{-1}) \mathcal{NG}(\theta, \tau | \mu_{k-1}, \Lambda_{k-1}, \alpha_{k-1}, \beta_{k-1}) \quad (36)$$

$$\propto \tau^{1/2} \exp\left(-\frac{\tau}{2}(\hat{y}_k - \theta^\top \phi_k)^2\right) \quad (37)$$

$$\tau^{D/2 + \alpha_{k-1} - 1} \exp\left(-\frac{\tau}{2}((\theta - \mu_{k-1})^\top \Lambda_{k-1}(\theta - \mu_{k-1}) + 2\beta_{k-1})\right) \\ = \tau^{\alpha_{k-1} + (D+1)/2 - 1} \exp\left(-\frac{\tau}{2}((\hat{y}_k - \theta^\top \phi_k)^2 + (\theta - \mu_{k-1})^\top \Lambda_{k-1}(\theta - \mu_{k-1}) + 2\beta_{k-1})\right). \quad (38)$$

Expanding the squares within the exponential and collecting terms involving θ yields:

$$(\hat{y}_k - \theta^\top \phi_k)^2 + (\theta - \mu_{k-1})^\top \Lambda_{k-1}(\theta - \mu_{k-1}) \\ = \theta^\top (\Lambda_{k-1} + \phi_k \phi_k^\top) \theta - (\mu_{k-1}^\top \Lambda_{k-1} + \hat{y}_k \phi_k^\top) \theta \\ - \theta^\top (\Lambda_{k-1} \mu_{k-1} + \phi_k \hat{y}_k) + \mu_{k-1}^\top \Lambda_{k-1} \mu_{k-1} + \hat{y}_k^2. \quad (39)$$

Then, using $\Lambda_k \triangleq \Lambda_{k-1} + \phi_k \phi_k^\top$ and $\xi_k \triangleq \Lambda_{k-1} \mu_{k-1} + \phi_k \hat{y}_k$, one may isolate a quadratic form:

$$\theta^\top \Lambda_k \theta - \xi_k^\top \theta - \theta^\top \xi_k \\ = (\theta - \Lambda_k^{-1} \xi_k)^\top \Lambda_k (\theta - \Lambda_k^{-1} \xi_k) - \xi_k^\top \Lambda_k^{-1} \xi_k. \quad (40)$$

Plugging the above back into Eq. 38, reveals a multivariate Gaussian - univariate Gamma distribution:

$$p(\theta, \tau | \mathcal{D}_k) \propto \tau^{\alpha_{k-1} + (D+1)/2 - 1} \quad (41)$$

$$\exp\left(-\frac{\tau}{2}((\theta - \Lambda_k^{-1} \xi_k)^\top \Lambda_k (\theta - \Lambda_k^{-1} \xi_k) - \xi_k^\top \Lambda_k^{-1} \xi_k + \mu_{k-1}^\top \Lambda_{k-1} \mu_{k-1} + \hat{y}_k^2 + 2\beta_{k-1})\right) \\ = \mathcal{NG}(\theta, \tau | \mu_k, \Lambda_k, \alpha_k, \beta_k), \quad (42)$$

with Λ_k defined as above (40), $\alpha_k \triangleq \alpha_{k-1} + 1/2$ and

$$\mu_k \triangleq \Lambda_k^{-1} \xi_k = (\Lambda_{k-1} + \phi_k \phi_k^\top)^{-1} (\Lambda_{k-1} \mu_{k-1} + \phi_k \hat{y}_k), \\ \beta_k \triangleq \beta_{k-1} + \frac{1}{2} (\hat{y}_k^2 - \xi_k^\top \Lambda_k^{-1} \xi_k + \mu_{k-1}^\top \Lambda_{k-1} \mu_{k-1}). \quad (43)$$

Note that $\xi_k = \Lambda_k \mu_k$, so $\xi_k^\top \Lambda_k^{-1} \xi_k = \mu_k^\top \Lambda_k \mu_k$. ■

APPENDIX B

Proof: We start by combining the Gaussian likelihood with the conditional Gaussian distribution for θ :

$$\mathcal{N}(y_t | \theta^\top \phi_t, \tau^{-1}) \mathcal{N}(\theta | \mu_k, (\tau \Lambda_k)^{-1}) = \quad (44) \\ \mathcal{N}\left(\begin{bmatrix} \theta \\ y_t \end{bmatrix} \middle| \begin{bmatrix} \mu_k \\ \mu_k^\top \phi_t \end{bmatrix}, \begin{bmatrix} (\tau \Lambda_k)^{-1} & (\tau \Lambda_k)^{-1} \phi_t \\ \phi_t^\top (\tau \Lambda_k)^{-1} & \phi_t^\top (\tau \Lambda_k)^{-1} \phi_t + \tau^{-1} \end{bmatrix}\right).$$

Marginalizing (44) over θ yields [15, ID: P35]:

$$p(y_t | \tau, u_t) = \mathcal{N}(y_t | \mu_k^\top \phi_t, \tau^{-1} (\phi_t^\top \Lambda_k^{-1} \phi_t + 1)). \quad (45)$$

Let $m_t \triangleq \mu_k^\top \phi_t$ and $\lambda_t \triangleq (\phi_t^\top \Lambda_k^{-1} \phi_t + 1)^{-1}$. Then

$$p(y_t | u_t) = \int p(y_t | \tau, u_t) p(\tau) d\tau \quad (46)$$

$$= \int \mathcal{N}(y_t | m_t, (\tau \lambda_t)^{-1}) \mathcal{G}(\tau | \alpha_k, \beta_k) d\tau \quad (47)$$

$$= \int \left(\frac{\lambda_t}{2\pi}\right)^{1/2} \tau^{1/2} \exp\left(-\frac{\tau \lambda_t}{2} (y_t - m_t)^2\right) \quad (48)$$

$$\frac{\beta_k^{\alpha_k}}{\Gamma(\alpha_k)} \tau^{\alpha_k - 1} \exp(-\tau \beta_k) d\tau$$

$$= \left(\frac{\lambda_t}{2\pi}\right)^{1/2} \frac{\beta_k^{\alpha_k}}{\Gamma(\alpha_k)} \quad (49)$$

$$\int \tau^{\alpha_k + 1/2 - 1} \exp\left(-\tau \left(\beta_k + \frac{\lambda_t}{2} (y_t - m_t)^2\right)\right) d\tau$$

$$= \left(\frac{\lambda_t}{2\pi}\right)^{1/2} \frac{\beta_k^{\alpha_k}}{\Gamma(\alpha_k)} \frac{\Gamma(\alpha_k + \frac{1}{2})}{\left(\beta_k + \frac{1}{2} \lambda_t (y_t - m_t)^2\right)^{\alpha_k + 1/2}} \quad (50)$$

$$= \left(\frac{\lambda_t}{2\pi}\right)^{1/2} \frac{\Gamma(\alpha_k + \frac{1}{2})}{\Gamma(\alpha_k)} \left(\frac{\beta_k}{\beta_k + \frac{1}{2} \lambda_t (y_t - m_t)^2}\right)^{\alpha_k + 1/2} \beta_k^{-1/2} \quad (51)$$

$$= \left(\frac{\alpha_k \lambda_t}{\alpha_k \beta_k 2\pi}\right)^{1/2} \frac{\Gamma(\frac{2\alpha_k + 1}{2})}{\Gamma(\frac{2\alpha_k}{2})} \left(1 + \frac{\alpha_k \lambda_t (y_t - m_t)^2}{2\alpha_k \beta_k}\right)^{-(2\alpha_k + 1)/2} \quad (52)$$

where the integrand in (49) is an unnormalized Gamma distribution and evaluates to its normalization constant in (50). Defining $\nu_t \triangleq 2\alpha_k$ and $s_t^2 \triangleq \left(\frac{\alpha_k}{\beta_k} \lambda_t\right)^{-1}$, reveals a location-scale T-distribution:

$$p(y_t | u_t) = \frac{\Gamma(\frac{\nu_t + 1}{2})}{\sqrt{\pi \nu_t s_t^2} \Gamma(\frac{\nu_t}{2})} \left(1 + \frac{(y_t - m_t)^2}{\nu_t s_t^2}\right)^{-(\nu_t + 1)/2}. \quad (53)$$

■

APPENDIX C

Proof: The cross-entropy in 30 is:

$$\begin{aligned} & \mathbb{E}_{\mathcal{T}_{v_t}(y_t | m_t, s_t^2)} \left[-\ln \mathcal{N}(y_t | m_*, v_*) \right] \\ &= \frac{1}{2} \ln(2\pi v_*) + \frac{1}{2v_*} \left((m_t - m_*)^2 + s_t^2 \frac{v_t}{v_t - 2} \right). \quad (54) \\ &= \text{constants} + \quad (55) \\ & \frac{1}{2v_*} \left((\mu_k^\top \phi_t - m_*)^2 + \frac{\beta_k}{\alpha_k} (\phi_t^\top \Lambda_k^{-1} \phi_t + 1) \frac{2\alpha_k}{2\alpha_k - 2} \right). \end{aligned}$$

The mutual information is with respect to the joint distribution in Eq. 13. Marginalizing $p(y_t, \theta, \tau | u_t)$ over y_t produces $p(\theta, \tau | \mathcal{D}_k)$ and marginalizing it over θ, τ produces $p(y_t | u_t)$ (Lemma 4.2). As such, we can factorize the mutual information into three differential entropies [3]:

$$\begin{aligned} & \mathbb{E}_{p(y_t, \theta, \tau | u_t)} \left[\ln \frac{p(y_t, \theta, \tau | u_t)}{p(\theta, \tau | \mathcal{D}_k)p(y_t | u_t)} \right] \quad (56) \\ &= -H[y_t, \theta, \tau | u_t] + H[\theta, \tau] + H[y_t | u_t]. \end{aligned}$$

For the first entropy, note that $p(y_t, \theta, \tau | u_t)$ consists of Eq. 44 with a Gamma distribution:

$$p(y_t, \theta, \tau | u_t) = \mathcal{N} \left(\begin{bmatrix} \theta \\ y_t \end{bmatrix} \middle| \begin{bmatrix} \mu_k \\ \mu_k^\top \phi_t \end{bmatrix}, (\tau \bar{\Lambda}_t)^{-1} \right) \Gamma(\tau | \alpha_k, \beta_k) \quad (57)$$

where the precision matrix is

$$\bar{\Lambda}_t = \begin{bmatrix} \Lambda_k^{-1} & \Lambda_k^{-1} \phi_t \\ \phi_t^\top \Lambda_k^{-1} & \phi_t^\top \Lambda_k^{-1} \phi_t + 1 \end{bmatrix}^{-1}. \quad (58)$$

The differential entropy of a multivariate Normal - univariate Gamma distribution is [15, ID: P238]:

$$\begin{aligned} H[y_t, \theta, \tau | u_t] &= \frac{D+1}{2} \ln(2\pi) + \frac{D+1}{2} - \frac{1}{2} \ln |\bar{\Lambda}_t| \quad (59) \\ &+ \alpha_k + \ln \Gamma(\alpha_k) - \frac{D-1+2\alpha_k}{2} \psi(\alpha_k) + \frac{D-1}{2} \ln \beta_k, \end{aligned}$$

where $\psi(\cdot)$ refers to a digamma function. The determinant of the precision matrix can be simplified to:

$$|\bar{\Lambda}_t| = (|\Lambda_k^{-1}| |\phi_t^\top \Lambda_k^{-1} \phi_t + 1 - \phi_t^\top \Lambda_k^{-1} \Lambda_k \Lambda_k^{-1} \phi_t|)^{-1} = |\Lambda_k|. \quad (60)$$

If we plug this result back into Eq. 59, then we see that this differential entropy term actually does not depend on u_t . The differential entropy of the parameter posterior distribution (Eq. 9) evaluates to [15, ID: P238]:

$$\begin{aligned} H[\theta, \tau] &= \frac{D}{2} \ln(2\pi) + \frac{D}{2} - \frac{1}{2} \ln |\Lambda_k| \quad (61) \\ &+ \alpha_k + \ln \Gamma(\alpha_k) - \frac{D-2+2\alpha_k}{2} \psi(\alpha_k) + \frac{D-2}{2} \ln \beta_k. \end{aligned}$$

This entropy also does not depend on u_t . For $H[y_t | u_t]$, note that a location-scale T-distributed variable is equivalent to a linearly transformed standard student's T-distributed variable. Differential entropy is invariant with respect to translation, but follows $H[sy] = H[y] + \ln |s|$ with respect to scaling [3]. Thus, we have:

$$\begin{aligned} H[y_t | u_t] &= \frac{2\alpha_k + 1}{2} \left(\psi \left(\frac{2\alpha_k + 1}{2} \right) - \psi \left(\frac{2\alpha_k}{2} \right) \right) \quad (62) \\ &+ \ln \sqrt{2\alpha_k} B \left(\frac{2\alpha_k}{2}, \frac{1}{2} \right) + \frac{1}{2} \ln \left(\frac{\beta_k}{\alpha_k} (\phi_t^\top \Lambda_k^{-1} \phi_t + 1) \right), \end{aligned}$$

with $B(\cdot)$ the beta function. Only the last term here depends on u_t (through ϕ_t). Plugging Eqs. 59, 61 and 62 into Eq. 56 and plugging Eqs. 56 and 55 into Eq. 30 completes the proof:

$$\begin{aligned} \mathcal{J}(u_t) &= \text{constants} - \frac{1}{2} \ln (\phi_t^\top \Lambda_k^{-1} \phi_t + 1) \quad (63) \\ &+ \frac{1}{2v_*} \left((\mu_k^\top \phi_t - m_*)^2 + \frac{\beta_k}{\alpha_k} (\phi_t^\top \Lambda_k^{-1} \phi_t + 1) \frac{2\alpha_k}{2\alpha_k - 2} \right). \end{aligned}$$

REFERENCES

- [1] S. A. Billings, *Nonlinear system identification: NARMAX methods in the time, frequency, and spatio-temporal domains*. John Wiley & Sons, 2013.
- [2] L. Grüne, J. Pannek, L. Grüne, and J. Pannek, *Nonlinear model predictive control*. Springer, 2017.
- [3] D. J. MacKay, *Information theory, inference and learning algorithms*. Cambridge University Press, 2003.
- [4] K. Friston, F. Rigoli, D. Ognibene, C. Mathys, T. Fitzgerald, and G. Pezzulo, "Active inference and epistemic value," *Cognitive neuroscience*, vol. 6, no. 4, pp. 187–214, 2015.
- [5] M. Baioumy, P. Duckworth, B. Lacerda, and N. Hawes, "Active inference for integrated state-estimation, control, and learning," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 4665–4671.
- [6] T. van de Laar, A. Özçelikkale, and H. Wymeersch, "Application of the free energy principle to estimation and control," *IEEE Transactions on Signal Processing*, vol. 69, pp. 4234–4244, 2021.
- [7] L. Pio-Lopez, A. Nizard, K. Friston, and G. Pezzulo, "Active inference and robot control: a case study," *Journal of The Royal Society Interface*, vol. 13, no. 122, p. 20160616, 2016.
- [8] A. A. Meera and M. Wisse, "Free energy principle based state and input observer design for linear systems with colored noise," in *American Control Conference*, 2020, pp. 5052–5058.
- [9] J. Huebotter, S. Thill, M. v. Gerven, and P. Lanillos, "Learning policies for continuous control via transition models," in *International Workshop on Active Inference*. Springer, 2023, pp. 162–178.
- [10] K. Fujimoto and Y. Takaki, "On system identification for ARMAX models based on the variational Bayesian method," in *IEEE Conference on Decision and Control*, 2016, pp. 1217–1222.
- [11] W. M. Kouw, A. Podusenko, M. T. Koudahl, and M. Schoukens, "Variational message passing for online polynomial NARMAX identification," in *American Control Conference*, 2022, pp. 2755–2760.
- [12] A. Podusenko, S. Akbayrak, İ. Şenöz, M. Schoukens, and W. M. Kouw, "Message passing-based system identification for NARMAX models," in *IEEE Conference on Decision and Control*, 2022, pp. 7309–7314.
- [13] C. Leung, S. Huang, N. Kwok, and G. Dissanayake, "Planning under uncertainty using model predictive control for information gathering," *Robotics and Autonomous Systems*, vol. 54, no. 11, pp. 898–910, 2006.
- [14] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.
- [15] J. Soch, T. J. Faulkenberry, K. Petrykowski, and C. Allefeld, *The Book of Statistical Proofs*, 2021.
- [16] S. Särkkä, *Bayesian filtering and smoothing*. Cambridge University Press, 2013, vol. 3.
- [17] T. van de Laar, M. Koudahl, B. van Erp, and B. de Vries, "Active inference and epistemic value in graphical models," *Frontiers in Robotics and AI*, vol. 9, p. 794464, 2022.
- [18] D. Khandelwal, M. Schoukens, and R. Tóth, "On the simulation of polynomial NARMAX models," in *IEEE Conference on Decision and Control*, 2018, pp. 1445–1450.
- [19] P. K. Mogensen and A. N. Riseth, "Optim: A mathematical optimization package for Julia," *Journal of Open Source Software*, vol. 3, no. 24, p. 615, 2018.
- [20] M. Deisenroth and C. E. Rasmussen, "PILCO: A model-based and data-efficient approach to policy search," in *International Conference on Machine Learning*, 2011, pp. 465–472.