

Online system identification in a Duffing oscillator by free energy minimisation

Wouter M. Kouw

Bayesian Intelligent Autonomous Systems lab
TU Eindhoven, Eindhoven, 5612AP Netherlands
w.m.kouw@tue.nl

Abstract. Online system identification is the estimation of parameters of a dynamical system, such as mass or friction coefficients, for each measurement of the input and output signals. Here, the nonlinear stochastic differential equation of a Duffing oscillator is cast to a generative model and dynamical parameters are inferred using variational message passing on a factor graph of the model. The approach is validated with an experiment on data from an electronic implementation of a Duffing oscillator. The proposed inference procedure performs as well as offline prediction error minimisation in a state-of-the-art nonlinear model.

Keywords: Online system identification · Duffing oscillator · Free energy minimisation · Variational message passing · Forney factor graphs

1 Introduction

Natural agents are believed to develop an internal model of their motor system by generating actions in muscles and observing limb movements [11]. It has been suggested that forming this internal model is analogous to a form of online system identification [24]. System identification, i.e. estimating dynamical parameters from observed input and output signals, has a rich history in engineering. But there might still be much to gain from considering biologically-plausible procedures. Here, I explore online system identification using a leading theory of how brains process information: free energy minimisation [8,3].

To test free energy minimisation for use in engineering applications, I consider a specific benchmark¹ problem called a Duffing oscillator. Duffing oscillators are relatively well-behaved nonlinear differential equations, making them excellent toy problems for methodological research. Its differential equation is cast to a generative model, with a corresponding factor graph. The factor graph admits a recursive parameter estimation procedure through message passing [14,12]. Specifically, variational message passing minimises free energy [5,13,18]. Here, I infer the parameters of a Duffing oscillator using online variational message passing. Experiments show that it performs as well as a nonlinear ARX model with parameters trained offline using prediction error minimisation [2].

¹ <http://nonlinearbenchmark.org/>

2 System

Consider a rigid frame with two prongs facing rightwards (see Figure 1 left). A steel beam is attached to the top prong. If the frame is driven by a periodic forcing term, the beam will displace horizontally as a driven damped harmonic oscillator. Two magnets are attached to the bottom prong, with the steel beam suspended in between. These act as a nonlinear feedback term on the beam's position, attracting or repelling it as it gets closer [15].




Fig. 1: (Left) Example of a physical implementation of a Duffing oscillator. (Right) Example of input and output signals.

Let $y(t)$ be the observed displacement, $x(t)$ the true displacement, and $u(t)$ the observed driving force. The position of the beam is described as follows [25]:

$$m \frac{d^2 x(t)}{dt^2} + c \frac{dx(t)}{dt} + ax(t) + bx^3(t) = u(t) + w(t) \quad (1a)$$

$$y(t) = x(t) + v(t) , \quad (1b)$$

where m is mass, c is damping, a the linear and b the nonlinear spring stiffness coefficient. Both the state transition as well as the observation likelihood contain noise terms, which are assumed to be Gaussian distributed: $w(t) \sim \mathcal{N}(0, \tau^{-1})$ (process noise) and $v(t) \sim \mathcal{N}(0, \xi^{-1})$ (measurement noise). The challenge is to estimate m , c , a , b , τ and ξ such that the output of the system can be predicted as accurately as possible.

3 Identification

First, I discretise the state transition of Equation 1 using a central difference for the second derivative and a forward difference for the first derivative. Rearranging to form an expression in terms of x_{t+1} yields:

$$x_{t+1} = \frac{2m + c\delta - a\delta^2}{m + c\delta} x_t + \frac{-b\delta^2}{m + c\delta} x_t^3 + \frac{-m}{m + c\delta} x_{t-1} + \frac{\delta^2}{m + c\delta} (u_t + w_t), \quad (2)$$

where δ is the sample time step. Secondly, to ease inference at a later stage, I perform the following variable substitutions:

$$\theta_1 = \frac{2m+c\delta-a\delta^2}{m+c\delta}, \quad \theta_2 = \frac{-b\delta^2}{m+c\delta}, \quad \theta_3 = \frac{-m}{m+c\delta}, \quad \eta = \frac{\delta^2}{m+c\delta}, \quad \gamma = \frac{\tau(m+c\delta)^2}{\delta^4}, \quad (3)$$

where the square in the numerator for γ stems from absorbing the coefficient into the noise term ($\mathbb{V}[\eta w_t] = \eta^2 \mathbb{V}[w_t]$). Note that the mapping between $\phi = (m, c, a, b, \tau)$ and $\psi = (\theta_1, \theta_2, \theta_3, \eta, \gamma)$ can be inverted to recover point estimates:

$$m = \frac{-\theta_3\delta^2}{\eta}, \quad c = \frac{(1+\theta_3)\delta}{\eta}, \quad a = \frac{1-\theta_1-\theta_3}{\eta}, \quad b = \frac{-\theta_2}{\eta}, \quad \tau = \gamma\eta^2. \quad (4)$$

Thirdly, the state transition can be cast to a multivariate first-order form:

$$\underbrace{\begin{bmatrix} x_{t+1} \\ x_t \end{bmatrix}}_{z_t} = \underbrace{\begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}}_S \underbrace{\begin{bmatrix} x_t \\ x_{t-1} \end{bmatrix}}_{z_{t-1}} + \underbrace{\begin{bmatrix} 1 \\ 0 \end{bmatrix}}_s g(\theta, z_{t-1}) + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \eta u_t + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \tilde{w}_t, \quad (5)$$

where $g(\theta, z_{t-1}) = \theta_1 x_t + \theta_2 x_t^3 + \theta_3 x_{t-1}$ and $\tilde{w}_t \sim \mathcal{N}(0, \gamma^{-1})$. The system is now a nonlinear autoregressive process. Lastly, integrating out \tilde{w}_t and v_t produces a Gaussian state transition and a Gaussian likelihood, respectively:

$$z_t \sim \mathcal{N}(f(\theta, z_{t-1}, \eta, u_t), V) \quad (6a)$$

$$y_t \sim \mathcal{N}(s^\top z_t, \xi^{-1}), \quad (6b)$$

where $f(\theta, z_{t-1}, \eta, u_t) = S z_{t-1} + s g(\theta, z_{t-1}) + s \eta u_t$ and $V = [\gamma^{-1} \ 0; 0 \ \epsilon]$. The number ϵ represents a small noise injection to stabilise inference [6].

To complete the generative model description, priors must be defined. Mass m and process precision τ are known to be strictly positive parameters, while the damping and stiffness coefficients can be both positive and negative. By examining the variable substitutions, it can be seen that $\theta_1, \theta_2, \theta_3$ and η can be both positive and negative, but γ can only be positive. As such, the following parametric forms can be chosen for the priors:

$$\theta \sim \mathcal{N}(m_\theta^0, V_\theta^0), \quad \eta \sim \mathcal{N}(m_\eta^0, v_\eta^0), \quad \gamma \sim \Gamma(a_\gamma^0, b_\gamma^0), \quad \xi \sim \Gamma(a_\xi^0, b_\xi^0). \quad (7)$$

3.1 Free energy minimisation

Given the generative model, a free energy functional with a recognition model q can be formed as follows:

$$-\log p(\mathbf{y}, \mathbf{u}) \leq \iint q(\psi, \mathbf{z}) \frac{q(\psi, \mathbf{z})}{p(\mathbf{y}, \mathbf{u}, \mathbf{z}, \psi)} d\mathbf{z} d\psi = \mathcal{F}[q] \quad (8)$$

where $\mathbf{z} = (z_1, \dots, z_T)$, $\mathbf{y} = (y_1, \dots, y_T)$ and $\mathbf{u} = (u_1, \dots, u_T)$. I assume the states factor over time and that the parameters are largely independent:

$$q(\psi, \mathbf{z}) = q(\theta)q(\eta)q(\gamma)q(\xi) \prod_{t=1}^T q(z_t). \quad (9)$$

All recognition densities are Gaussian distributed, except for $q(\gamma)$ and $q(\xi)$, which are Gamma distributed. In free energy minimisation, the parameters of the recognition distributions depend on each other and are iteratively updated.

3.2 Factor graphs and message passing

In online system identification, parameter estimates should be updated at each time-step. That puts time constraints on the inference procedure. Message passing is an ideal inference procedure due to its efficiency in factorised generative models [12]. Figure 2 is a graphical representation of the generative model, with nodes for factors and edges for variables. Square nodes with Greek letters represent stochastic operations while \square and \equiv represent deterministic operations. The node marked "NLARX" represents the state transition described in Equation 6a.

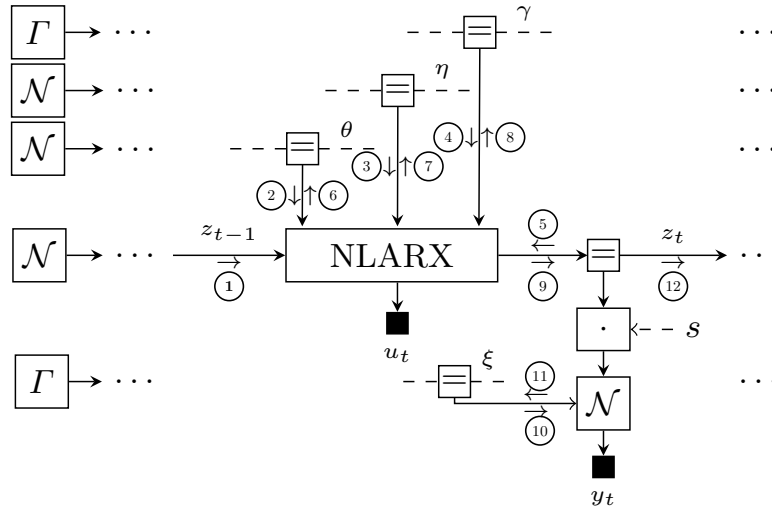


Fig. 2: Forney-style factor graph of the generative model of a Duffing oscillator. Nodes represent conditional distributions and edges represent variables. Nodes send messages to connected edges. When two messages on an edge collide, the marginal belief q for the corresponding variable is updated. Each belief update reduces free energy. By iterating message passing, free energy is minimised.

The terminal nodes on the left represent the initial priors for the states and dynamical parameters. Inference starts when these nodes pass messages. The subgraph - separated by columns of dots - represents the structure of a single time step, recursively applied. Messages ①, ②, ③, ④ and ⑩ represent beliefs q from previous time-steps. Message ⑤, arriving at the state transition node, originates from the likelihood node attached to observation y_t . Messages

⑥, ⑦, ⑧, ⑨ and ⑪ combine priors from previous time steps and likelihoods of observations, and are used to update beliefs q . Message ⑫ is the current state belief and becomes message ① in the next time step.

The graph actually contains more messages, such as those sent by equality nodes. I have hidden them to avoid complicating the figure. Their form has been extensively described in the literature and can be looked up easily [12,14]. Modern message passing toolboxes, such as Infer.NET and ForneyLab.jl, automatically incorporate them. However, the NLARX node is new. Its messages can be computed with²:

$$\textcircled{6} \vec{\nu}(\theta) = \exp\left(\mathbb{E}_{q(z_t)q(z_{t-1})q(\eta)q(\gamma)}[\log \mathcal{N}(f(\theta, z_{t-1}, \eta, u_t), V)]\right) \quad (10a)$$

$$\textcircled{7} \vec{\nu}(\eta) = \exp\left(\mathbb{E}_{q(z_t)q(z_{t-1})q(\theta)q(\gamma)}[\log \mathcal{N}(f(\theta, z_{t-1}, \eta, u_t), V)]\right) \quad (10b)$$

$$\textcircled{8} \vec{\nu}(\gamma) = \exp\left(\mathbb{E}_{q(z_t)q(z_{t-1})q(\theta)q(\eta)}[\log \mathcal{N}(f(\theta, z_{t-1}, \eta, u_t), V)]\right) \quad (10c)$$

$$\textcircled{9} \vec{\nu}(z_t) = \exp\left(\mathbb{E}_{q(z_{t-1})q(\theta)q(\eta)q(\gamma)}[\log \mathcal{N}(f(\theta, z_{t-1}, \eta, u_t), V)]\right), \quad (10d)$$

where I use a first-order Taylor expansion to approximate the expected value of the nonlinear autoregressive function $g(\theta, z_{t-1})$.

Loeliger et al. (2007) have written an accessible introduction on message passing in factor graphs [14]. Variational message passing in autoregressive processes has been described in detail as well [5,19].

4 Experiment

The Duffing oscillator has been implemented in an electronic system called Silverbox [25]. It consists of $T = 131702$ samples, gathered with a sampling frequency of 610.35 Hz. Figure 3 shows the time-series, plotted at every 80 time steps. There are two regimes: the first 40000 samples are subject to a linearly increasing amplitude in the input (left of the black line in Figure 3) and the remaining samples are subject to a constant amplitude but contain only odd harmonics (right of the black line). The second regime is used as a training data set, where both input and output data were given and parameters needed to be inferred. The first regime is used as a validation data set, where the inferred parameters are fixed and the model needs to make predictions for the output signal.

I performed two experiments³: a 1-step ahead prediction error and a simulation error setting. I used ForneyLab.jl, with NLARX as a custom node, to run the message passing inference procedure [4]. I call the model above FEM-NLARX, for Nonlinear Latent Autoregressive model with eXogenous input using Free Energy Minimisation. I implemented two baselines: the first is NLARX without the nonlinearity (i.e. the nonlinear spring coefficient $b = 0$), dubbed FEM-LARX.

² Derivations at <https://github.com/biaslab/IWAI2020-onlinesysid>

³ Experiment notebooks at <https://github.com/biaslab/IWAI2020-onlinesysid>




Fig. 3: Silverbox data set, sampled at every 80 time steps for visualisation. The black line splits it into validation data (left) and training data (right).

The second is a standard NARX model, implemented using MATLAB’s System Identification Toolbox. I modelled the static nonlinearity with a sigmoid network of 4 units (in line with the 4 coefficients used by NLARX and LARX). Parameters were inferred offline using Prediction Error Minimisation. Hence, this baseline is called PEM-NARX.

I chose uninformative priors for the coefficients θ and η : Gaussians centred at 1 with precisions of 0.1. The authors of Silverbox indicate that the signal-to-noise ratio at measurement time was high [25]. I therefore chose informative priors for the noise parameters: $a_{\xi}^0 = 1e8$ and $a_{\gamma}^0 = 1e3$ (shape parameters) and $b_{\xi}^0 = 1e3$ and $b_{\gamma}^0 = 1e1$ (scale parameters).

4.1 1-step ahead prediction error

At each time-step in the validation data, the models were given the previous output signal y_{t-1}, y_{t-2} and the current input signal u_t and had to infer the current output y_t . It is a relatively easy task, which is reflected in all three models’ performance. The top row in Figure 4 shows the predictions of all three models in purple and their squared error with respect to the true output signal in black. The left column shows the offline NARX baseline (PEM-NARX), the middle column the linear online latent autoregressive baseline (FEM-LARX) and the right column the nonlinear online latent autoregressive model (FEM-NLARX). Note that the errors in the top row seem completely flat. The bottom row in the figure plots the errors on a log-scale. PEM-NARX has a mean squared error of $5.831e-5$, FEM-LARX one of $5.945e-5$ and FEM-NLARX one of $5.830e-5$.

4.2 Simulation error

In this experiment, the models were not given the previous output signal, but had to use their predictions from the previous time-step. This is a much harder task, because errors will accumulate. The top row in Figure 5 again shows the predictions of all three models (purple) and their squared error (black). It can already be seen that the errors increase as the input signal’s amplitude rises. The bottom row plots the errors on a log-scale. PEM-NARX has a mean squared error of $1.000e-3$, FEM-LARX one of $1.002e-3$ and FEM-NLARX one of $0.926e-3$.




Fig. 4: 1-step ahead prediction errors. (Left) Offline NARX model with sigmoid net (PEM-NARX), (middle) online linear model (FEM-LARX) and (right) on-line nonlinear model (FEM-NLARX). (Top) Predictions (purple) and squared error (black). (Bottom) Squared prediction errors in log-scale.

5 Discussion

The experimental results seem to justify looking to nature for inspiration. Free energy minimisation, in the form of variational message passing, seems a generally applicable and well-performing inference technique. The difficulties mostly lie in deriving variational messages (i.e. Equations 10).

Improvements in the proposed procedure could be made with a richer approximation of the nonlinear autoregressive function (e.g. unscented transform) [20]. Alternatively, a hierarchy of latent Gaussian filters or autoregressive processes could be used to obtain time-varying noise parameters or time-varying coefficients [22,19]. Furthermore, instead of discretising such that an auto-regressive model is obtained, one could express the evolution of the states in generalised coordinates. Lastly, black-box models could be explored for further performance improvements.

A natural next step is for an active inference agent to determine the control signal regime (i.e. optimal design). Unfortunately, this is not straightforward: the current formulation relies on variational free energy which does not produce an epistemic term in the objective. The epistemic term is needed to encourage exploration; i.e. try sub-optimal inputs to reduce uncertainty. To arrive at an epistemic term, one would need to work with expected free energy [17]. But it is unclear how expected free energy could be incorporated into factor graphs.

5.1 Related work

Online system identification procedures typically employ recursive least-squares or maximum likelihood inference, with nonlinearities modelled by basis expan-




Fig. 5: Simulation errors. (Left) Offline NARX model with sigmoid net (PEM-NARX), (middle) online linear model (FEM-LARX) and (right) online nonlinear model (FEM-NLARX). (Top) Predictions (purple) and squared error (black). (Bottom) Squared prediction errors in log-scale.

sions or neural networks [16,23,7]. Online Bayesian identification procedures come in two flavours: sequential Monte Carlo samplers [10,1] and online variational Bayes [26,9]. This work is novel in the use of variational message passing as an efficient implementation of online variational Bayes and its application to a nonlinear autoregressive model.

6 Conclusion

I have presented a free energy minimisation procedure for online system identification. Experimental results showed comparable performance to a state-of-the-art nonlinear model with parameters estimated offline. This indicates that the procedure performs well enough to be deployed in engineering applications.

Future work should test variational message passing in more challenging nonlinear identification settings, such as a Wiener-Hammerstein benchmark [21]. Furthermore, problems with time-varying dynamical parameters, such as a robotic arm picking up objects with mass, would be interesting for their connection to natural agents.

7 Acknowledgements

The author thanks Magnus Koudahl, Albert Podusenko and Thijs van de Laar for insightful discussions and the reviewers for their constructive feedback.

References

1. Abdesslem, A.B., Dervilis, N., Wagg, D., Worden, K.: Identification of nonlinear dynamical systems using approximate Bayesian computation based on a sequential Monte Carlo sampler. In: International Conference on Noise and Vibration Engineering (2016)
2. Aguirre, L.A., Letellier, C.: Modeling nonlinear dynamics and chaos: a review. *Mathematical Problems in Engineering* **2009** (2009)
3. Buckley, C.L., Kim, C.S., McGregor, S., Seth, A.K.: The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology* **81**, 55–79 (2017)
4. Cox, M., van de Laar, T., de Vries, B.: Forneylab.jl: Fast and flexible automated inference through message passing in Julia. In: International Conference on Probabilistic Programming (2018)
5. Dauwels, J.: On variational message passing on factor graphs. In: IEEE International Symposium on Information Theory. pp. 2546–2550 (2007)
6. Dauwels, J., Eckford, A., Korl, S., Loeliger, H.A.: Expectation maximization as message passing - Part I: Principles and Gaussian messages. arXiv:0910.2832 (2009)
7. Engel, Y., Mannor, S., Meir, R.: The kernel recursive least-squares algorithm. *IEEE Transactions on Signal Processing* **52**(8), 2275–2285 (2004)
8. Friston, K., Kilner, J., Harrison, L.: A free energy principle for the brain. *Journal of Physiology* **100**(1-3), 70–87 (2006)
9. Fujimoto, K., Satoh, A., Fukunaga, S.: System identification based on variational Bayes method and the invariance under coordinate transformations. In: IEEE Conference on Decision and Control and European Control Conference. pp. 3882–3888 (2011)
10. Green, P.L.: Bayesian system identification of a nonlinear dynamical system using a novel variant of simulated annealing. *Mechanical Systems and Signal Processing* **52**, 133–146 (2015)
11. de Klerk, C.C., Johnson, M.H., Heyes, C.M., Southgate, V.: Baby steps: Investigating the development of perceptual–motor couplings in infancy. *Developmental Science* **18**(2), 270–280 (2015)
12. Korl, S.: A factor graph approach to signal modelling, system identification and filtering. Ph.D. thesis, ETH Zurich (2005)
13. van de Laar, T., Cox, M., Senoz, I., Bocharov, I., de Vries, B.: Forneylab: a toolbox for biologically plausible free energy minimization in dynamic neural models. In: Conference on Complex Systems (2018)
14. Loeliger, H.A., Dauwels, J., Hu, J., Korl, S., Ping, L., Kschischang, F.R.: The factor graph approach to model-based signal processing. *Proceedings of the IEEE* **95**(6), 1295–1322 (2007)
15. Moon, F., Holmes, P.J.: A magnetoelastic strange attractor. *Journal of Sound and Vibration* **65**(2), 275–296 (1979)
16. Paleologu, C., Benesty, J., Ciochina, S.: A robust variable forgetting factor recursive least-squares algorithm for system identification. *IEEE Signal Processing Letters* **15**, 597–600 (2008)
17. Parr, T., Friston, K.J.: Generalised free energy and active inference. *Biological Cybernetics* **113**(5-6), 495–513 (2019)
18. Parr, T., Markovic, D., Kiebel, S.J., Friston, K.J.: Neuronal message passing using mean-field, Bethe, and marginal approximations. *Scientific reports* **9**(1), 1–18 (2019)

19. Podusenko, A., Kouw, W.M., de Vries, B.: Online variational message passing in hierarchical autoregressive models. In: IEEE International Symposium on Information Theory. pp. 1343–1348 (2020)
20. Särkkä, S.: Bayesian filtering and smoothing, vol. 3. Cambridge University Press (2013)
21. Schoukens, M., Noël, J.P.: Three benchmarks addressing open challenges in nonlinear system identification. *IFAC-PapersOnLine* **50**(1), 446–451 (2017)
22. Senoz, I., Podusenko, A., Kouw, W.M., de Vries, B.: Bayesian joint state and parameter tracking in autoregressive models. In: Conference on Learning for Dynamics and Control. pp. 1–10 (2020)
23. Tangirala, A.K.: Principles of system identification: theory and practice. CRC Press (2018)
24. Tin, C., Poon, C.S.: Internal models in sensorimotor integration: perspectives from adaptive control theory. *Journal of Neural Engineering* **2**(3), S147 (2005)
25. Wigren, T., Schoukens, J.: Three free data sets for development and benchmarking in nonlinear system identification. In: European Control Conference (ECC). pp. 2933–2938 (2013)
26. Yoshimoto, J., Ishii, S., Sato, M.a.: System identification based on online variational Bayes method and its application to reinforcement learning. In: Artificial Neural Networks and Neural Information Processing, pp. 123–131. Springer (2003)