

ADAPTIVE, CAUTIOUS, PREDICTIVE CONTROL WITH GAUSSIAN PROCESS PRIORS

Roderick Murray-Smith * Daniel Sbarbaro **
Carl Edward Rasmussen *** Agathe Girard ****

* *Department of Computing Science, University of Glasgow, Glasgow
G12 8QQ, Scotland, UK. & Hamilton Institute, NUI Maynooth, Ireland
E-mail: rod@dcs.gla.ac.uk*

** *Departamento de Ingeniería Eléctrica, Universidad de Concepción,
Chile. E-mail: dsbarbar@die.udec.cl*

*** *Max Planck Institute, Tübingen, Germany. E-mail:
carl@tuebingen.mpg.de*

**** *Department of Computing Science, University of Glasgow,
Glasgow G12 8QQ, Scotland, UK. E-mail: agathe@dcs.gla.ac.uk*

Abstract: Nonparametric Gaussian Process models, a Bayesian statistics approach, are used to implement a nonlinear adaptive control law. Predictions, including propagation of the state uncertainty are made over a k -step horizon. The expected value of a quadratic cost function is minimised, over this prediction horizon, without ignoring the variance of the model predictions. The general method and its main features are illustrated on a simulation example. Copyright ©2003 IFAC

Keywords: Cautious control, Gaussian process priors, nonparametric models, nonlinear model-based predictive control, propagation of uncertainty.

1. INTRODUCTION

Gaussian process priors provide a flexible, nonparametric approach to modelling nonlinear systems. In Murray-Smith and Sbarbaro (2002), the use of Gaussian process priors allowed us to analytically obtain a control law which perfectly minimises the expected value of a quadratic cost function, without disregarding the variance of the model prediction as an element to be minimised. This led naturally and automatically to a regularising *caution* in control behaviour in following the reference trajectory, depending on model accuracy close to the current state. This paper expands on the previous work by making the cost function more flexible, using our recent results on inference with uncertain states Girard et al. (2003) and allowing control based on multistep predictions.

1.1 Background

Several authors have proposed the use of non-linear models as a base to build nonlinear adaptive controllers. Agarwal and Seborg (1987), for instance, have proposed the use of known nonlinearities, capturing the main characteristic of the process, to design a Generalized Minimum Variance type of self-tuning controller. In many applications, however, these nonlinearities are not known, and non-linear parameterisation must be used instead. A popular choice has been the use of Artificial Neural Networks for estimating the nonlinearities of the system Narendra and Parthasarathy (1990); Bittanti and Piroddi (1997). These researchers adopted the *certainty equivalence principle* for designing the controllers, where the model is used in the control law as if it were the true system.

In order to improve the performance of nonlinear adaptive controllers based on nonlinear models, the accuracy of the model predictions should also be taken into account. A common approach to consider the uncertainty in the parameters, is to add an extra term in the cost function of a Minimum Variance controller, which penalizes the uncertainty in the parameters of the nonlinear approximation Fabri and Kadiramanathan (1998). Another similar approach, based on the minimization of two separate cost functions, has been proposed in Filatov et al. (1997). The first one is used to improve the parameter estimation and the second one to drive the system output to follow a given reference signal. This approach is called *bicriterial* and it has also been extended to deal with nonlinear systems Sbarbaro et al. (1998).

The above ideas are closely related to the work done on dual adaptive control Fel'dbaum (1960), where the main effort has been concentrated on the analysis and design of adaptive controllers based on the use of the uncertainty associated with parameters of models with fixed structure. The dual control literature has struggled with the challenges of numerical integration over high-dimensional spaces. Most approaches to dual control have pursued approximations which lead to sub-optimal solutions Wittenmark (1995); Filatov and Unbehauen (2000).

1.2 Model structure – parametric vs. nonparametric?

Most control engineering applications are still based on parametric models, where the functional form is fully described by a finite number of parameters, often a linear function of the parameters. Even in the cases where flexible parametric models are used, such as neural networks, spline-based models, multiple models etc, the uncertainty is usually expressed as uncertainty of parameters (even though the parameters often have no physical interpretation), and do not take into account uncertainty about model structure, or distance of current prediction point from training data used to estimate parameters.

Non-parametric models retain the available data and perform inference conditional on the current state and local data (called ‘smoothing’ in some frameworks). As the data are used directly in prediction, unlike the parametric methods more commonly used in control contexts, non-parametric methods have advantages for off-equilibrium regions, since normally in these regions the amount of data available for identification is much smaller than that available in steady state. The uncertainty of model predictions can be made dependent on local data density, and the model complexity automatically related to the amount and distribution of available data (more complex models need more evidence to make them likely). Both aspects are very useful in sparsely-populated transient regimes.

2. GAUSSIAN PROCESS PRIORS

In a Bayesian framework the model is based on a prior distribution over the infinite-dimensional space of functions. As illustrated in O’Hagan (1978), such priors can be defined as Gaussian processes. These models have attracted a great deal of interest recently – see for example reviews such as Williams (1998). ? showed empirically that Gaussian processes were extremely competitive with leading nonlinear identification methods on a range of benchmark examples. The further advantage that they provide analytic predictions of model uncertainty makes them very interesting for control applications. Use of GPs in a control systems context is discussed in Murray-Smith et al. (1999); Leith et al. (2000). Integration of prior information in the form of state or control linearisations is presented in Solak et al. (2003). A simulation of Model Predictive Control with GPs is presented in Kocijan et al. (2003).

In the following, let \mathbf{x} and \mathbf{y} be the N input and output pairs used for identification and \mathbf{x}^* , \mathbf{y}^* the pair used for prediction. Instead of parameterising $y^i = f(\mathbf{x}^i)$ as a parametric model, we can place a prior directly on the space of functions where f is assumed to belong. A Gaussian process represents the simplest form of prior over functions – we assume that any p points have a p -dimensional multivariate Normal distribution. We will assume zero mean, so for the case with partitioned data \mathbf{y} and \mathbf{y}^* we will have the multivariate Normal distribution

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}^* \end{bmatrix} \sim \mathcal{N}(0, \Sigma_F), \quad \Sigma_F = \begin{bmatrix} \Sigma & \mathbf{k}(\mathbf{x}^*) \\ \mathbf{k}(\mathbf{x}^*)^T & C(\mathbf{x}^*, \mathbf{x}^*) \end{bmatrix}. \quad (1)$$

where Σ_F is the full covariance matrix. Like the Gaussian distribution, the Gaussian Process is fully specified by a mean and its covariance function, so we denote the distribution $GP(\mu, C)$. The covariance function $C(\mathbf{x}_i, \mathbf{x}_j)$ expresses the covariance between y_i and y_j .

2.1 The covariance function

The model’s prior expectations can be adapted to a given application by altering the covariance function. In this paper, we use a straightforward covariance function,

$$C(\mathbf{x}_i, \mathbf{x}_j; \Theta) = v_0 \rho(|\mathbf{x}_i - \mathbf{x}_j|, \alpha),$$

so that the hyperparameter vector is $\Theta = [v_0, \alpha_1, \dots, \alpha_p]^T$ and p is the dimension of vector \mathbf{x} . The function $\rho(d)$ is a distance measure, which should be one at $d = 0$ and which should be a monotonically decreasing function of d . The one used here was

$$\rho(|\mathbf{x}_i - \mathbf{x}_j|, \alpha) = e^{-\frac{1}{2} \sum_{k=1}^p \alpha_k (x_{ik} - x_{jk})^2}. \quad (2)$$

In most cases we will only have uncertain knowledge of Θ . With unknown hyperparameters, we can use




Fig. 1. Sampling from a Gaussian Process. Sample realisations are drawn from the Gaussian process prior (left hand figure), and then sample functions are drawn from prior conditioned on three training points (right hand figure).

maximum likelihood methods, with standard gradient-based optimisation tools to optimise the log-likelihood by adapting the hyperparameters. The hyperparameters are usually given a vague prior distribution, such as a gamma prior Neal (1997). Each hyperparameter of the covariance function can be given an independent prior distribution. The use of Gamma priors does not add significant complexity to the optimisation, and if used appropriately makes the model behaviour more robust with small training sets, leading to increased robustness and higher performance in the early stages of adaptation, but the relative advantage decreases with the amount of initial data available, as would be expected.

2.2 Prediction

As in the multinormal case, we can divide the joint probability into a marginal Gaussian process and a conditional Gaussian process. The marginal term gives us the probability of the training data,

$$P(\mathbf{y}|\mathbf{x}) = (2\pi)^{-\frac{N}{2}} |\Sigma|^{-\frac{1}{2}} e^{-\frac{1}{2}\mathbf{y}^T \Sigma^{-1} \mathbf{y}}. \quad (3)$$

The conditional part of the model, which best relates to a traditional regression model, is therefore the Gaussian process which gives us the output posterior density function conditional on the training data \mathbf{x} , \mathbf{y} and the test points \mathbf{x}^* .

We have $p(y^*|\mathbf{x}, \mathbf{y}, \mathbf{x}^*) = \frac{p(\mathbf{y}^*, \mathbf{y})}{p(\mathbf{y})}$, that is

$$p(y^*|\mathbf{x}, \mathbf{y}, \mathbf{x}^*) = \frac{1}{(2\pi)^{\frac{1}{2}} |\sigma^2(\mathbf{x}^*)|^{\frac{1}{2}}} e^{-\frac{1}{2} \frac{(y^* - \mu(\mathbf{x}^*))^2}{\sigma^2(\mathbf{x}^*)}}$$

where

$$\mu(\mathbf{x}^*) = \mathbf{k}(\mathbf{x}^*)^T \Sigma^{-1} \mathbf{y} \quad (4)$$

$$\sigma^2(\mathbf{x}^*) = C(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{k}(\mathbf{x}^*)^T \Sigma^{-1} \mathbf{k}(\mathbf{x}^*), \quad (5)$$

with $k^i(\mathbf{x}^*) = C(\mathbf{x}^*, \mathbf{x}^i)$. So we can use $\mu(\mathbf{x}^*)$ as the expected model output, with a variance of $\sigma^2(\mathbf{x}^*)$.

2.2.1. Prediction at a random input If we now assume that the test input \mathbf{x}^* has a Gaussian distribution, $\mathbf{x}^* \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{x}^*}, \boldsymbol{\Sigma}_{\mathbf{x}^*})$, the predictive distribution is now obtained by integrating over \mathbf{x}^* :

$$p(y^*|\boldsymbol{\mu}_{\mathbf{x}^*}, \boldsymbol{\Sigma}_{\mathbf{x}^*}) = \int p(y^*|\mathbf{x}, \mathbf{y}, \mathbf{x}^*) p(\mathbf{x}^*) d\mathbf{x}^*,$$

where $p(y^*|\mathbf{x}, \mathbf{y}, \mathbf{x}^*)$ is as specified by (4) and (5).

In Girard et al. (2002, 2003); Quinonero-Candela et al. (2003), we suggest an analytical Gaussian approximation to solve this integral, computing only the mean and variance of $p(y^*|\boldsymbol{\mu}_{\mathbf{x}^*}, \boldsymbol{\Sigma}_{\mathbf{x}^*})$. Assuming a covariance function such as (2), we arrive at

$$m(\boldsymbol{\mu}_{\mathbf{x}^*}, \boldsymbol{\Sigma}_{\mathbf{x}^*}) = v_0 \mathbf{q}^T \boldsymbol{\beta} \quad (6)$$

$$v(\boldsymbol{\mu}_{\mathbf{x}^*}, \boldsymbol{\Sigma}_{\mathbf{x}^*}) = C(\boldsymbol{\mu}_{\mathbf{x}^*}, \boldsymbol{\mu}_{\mathbf{x}^*}) + v_0^2 \text{Tr} [(\boldsymbol{\beta} \boldsymbol{\beta}^T - \boldsymbol{\Sigma}^{-1}) \mathbf{Q}] - m(\boldsymbol{\mu}_{\mathbf{x}^*}, \boldsymbol{\Sigma}_{\mathbf{x}^*})^2 \quad (7)$$

where $\boldsymbol{\beta} = \boldsymbol{\Sigma}^{-1} \mathbf{y}$, $\mathbf{W} = \text{diag}[\alpha_1, \dots, \alpha_p]$, I is the $p \times p$ identity matrix and

$$q_i = |\mathbf{W}^{-1} \boldsymbol{\Sigma}_{\mathbf{x}^*} + I|^{-\frac{1}{2}} e^{-\frac{1}{2} (\boldsymbol{\mu}_{\mathbf{x}^*} - \mathbf{x}_i)^T (\boldsymbol{\Sigma}_{\mathbf{x}^*} + \mathbf{W})^{-1} (\boldsymbol{\mu}_{\mathbf{x}^*} - \mathbf{x}_i)}$$

$$Q_{ij} = |2\mathbf{W}^{-1} \boldsymbol{\Sigma}_{\mathbf{x}^*} + I|^{-\frac{1}{2}} e^{-\frac{1}{2} (\mathbf{x}_i - \mathbf{x}_j)^T (2\mathbf{W})^{-1} (\mathbf{x}_i - \mathbf{x}_j)} e^{-\frac{1}{2} (\mathbf{x}_i \mathbf{b} - \boldsymbol{\mu}_{\mathbf{x}^*})^T (\frac{1}{2} \mathbf{W} + \boldsymbol{\Sigma}_{\mathbf{x}^*})^{-1} (\mathbf{x}_i \mathbf{b} - \boldsymbol{\mu}_{\mathbf{x}^*})},$$

where $\mathbf{x}_i \mathbf{b} = (\mathbf{x}_i + \mathbf{x}_j)/2$ (see Girard et al. (2002, 2003); Quinonero-Candela et al. (2003) for the detailed calculations).

Note that as $\boldsymbol{\Sigma}_{\mathbf{x}^*}$ tends to zero we have $m(\boldsymbol{\mu}_{\mathbf{x}^*}, \boldsymbol{\Sigma}_{\mathbf{x}^*}) \rightarrow \mu(\boldsymbol{\mu}_{\mathbf{x}^*})$ and $v(\boldsymbol{\mu}_{\mathbf{x}^*}, \boldsymbol{\Sigma}_{\mathbf{x}^*}) \rightarrow \sigma^2(\boldsymbol{\mu}_{\mathbf{x}^*})$, as we would expect.

3. DERIVATION OF CONTROL LAW

The objective of this paper is to control a multi-input, single-output, nonlinear system of the form,

$$y(t+1) = f(\mathbf{x}(t), u(t)) + e(t+1) \quad (8)$$

where $\mathbf{x}(t)$ is the state vector at time t , which in this paper will be defined as $\mathbf{x}(t) = [y(t), \dots, y(t-n), u(t-1), \dots, u(t-m), v_1(t), \dots, v_l(t)]$, $y(t+1)$ the output, $u(t)$ the current control vector, $v_i(t)$, $i = 1, \dots, l$ are external known signals, f is a smooth nonlinear function, bounded away from zero. For notational simplicity we consider single control input systems, but extending the presentation to vector $u(t)$ is trivial. The noise term $e(t)$ is assumed zero mean Gaussian, with unknown variance σ^2 .

The cost function proposed is:

$$J = E\{(y_d(t+1) - y(t+1))^2\}$$

where $y_d(t)$ is a bounded reference signal.¹

Using the fact that $\text{Var}\{y\} = E\{y^2\} - \mu_y^2$, where $\mu_y = E\{y\}$, the cost function can be written as:

$$J = (y_d(t+1) - E\{y(t+1)\})^2 + \text{Var}\{y(t+1)\}.$$

With most models, estimation of $\text{Var}\{y\}$, and (as we shall see below) $\frac{\partial \text{Var}\{y\}}{\partial u(t)}$ would be difficult, forcing the use of certainty equivalence assumptions, adding extra terms to the cost function, or pursuing other sub-optimal solutions Wittenmark (1995); Filatov and Unbehauen (2000). With the Gaussian process prior (assuming smooth, differentiable covariance functions – see O’Hagan (1992)) straightforward analytic solutions can be obtained.

3.1 Multistep optimisation

In Murray-Smith and Sbarbaro (2002) we presented the one-step ahead version of the controller, but at $T = 1$ we do not get any ‘probing’ or ‘active learning’ benefit from the one-step controller—it only minimises the loss at the next time-step, and we limited ourselves to non minimum-phase systems. We now wish to use multistep optimisation. The cost function proposed is:

$$J_T = \frac{1}{T} \sum_{k=1}^T E\{(y_d(t+k) - y(t+k))^2\} \quad (10)$$

3.1.1. Multistep prediction We use the framework described in section 2.2.1 to propagate model predictions and uncertainties as we predict ahead in time. That is, we compute $E\{y(t+k)\}$ and $\text{Var}\{y(t+k)\}$ using equations (6) and (7) respectively.² This enables us to incorporate the uncertainty about intermediate estimates $\hat{y}(t+k-1), \dots, \hat{y}(t+k-n)$, which affects both the expected value, and its variance at $t+k$. See Girard et al. (2002) for more details on the application of the prediction at a Gaussian random input to time-series iterative k -step ahead prediction.

3.1.2. Sensitivity equations for GPs Given the cost function (10), and observations to time t , if we wish to find the optimal $u(t)$, we need the derivative of J_T ,

¹ It is straightforward to generalise this to more practical cost function, such as $J = E\{(y_d(t+1) - y(t+1))^2\} + (R(q^{-1})u(t))^2$, where the polynomial $R(q^{-1})$ is defined as:

$$R(q^{-1}) = r_0 + r_1 q^{-1} + \dots + r_{n_r} q^{-n_r} \quad (9)$$

and its coefficients can be used as tuning parameters.

² Note that for $k = 1$, \mathbf{x}^* will be composed of *known* lagged outputs so that $\Sigma_{\mathbf{x}^*} = \mathbf{0}$ and we can use (4) and (5).

$$\frac{\partial J_T}{\partial u(j)} = \sum_{k=j+1}^{t+T} \frac{\partial J_k}{\partial u(j)} \quad (11)$$

$$\begin{aligned} \frac{\partial J_k}{\partial u(j)} = & -2(y_d(k+1) - \mu_y(k+1)) \frac{\partial \mu_y(k+1)}{\partial u(j)} \\ & + \frac{\partial \text{Var}\{y(k+1)\}}{\partial u(j)} \end{aligned} \quad (12)$$

where

$$\begin{aligned} \frac{\partial \mu_y(k+1)}{\partial u(j)} = & \frac{\partial \mu_y(k+1)}{\partial \mu_y(k)} \frac{\partial \mu_y(k)}{\partial u(j)} + \\ & \frac{\partial \mu_y(k+1)}{\partial \text{Var}\{y(k)\}} \frac{\partial \text{Var}\{y(k)\}}{\partial u(j)} \end{aligned} \quad (13)$$

$$\begin{aligned} \frac{\partial \text{Var}\{y(k+1)\}}{\partial u(j)} = & \frac{\partial \text{Var}\{y(k+1)\}}{\partial \mu_y(k)} \frac{\partial \mu_y(k)}{\partial u(j)} + \\ & \frac{\partial \text{Var}\{y(k+1)\}}{\partial \text{Var}\{y(k)\}} \frac{\partial \text{Var}\{y(k)\}}{\partial u(j)} \end{aligned} \quad (14)$$

The calculation of these terms requires use of the sensitivity equations for a change to $u(j)$, i.e. the derivatives of model predictions of mean and variance from $k = j, \dots, T$. This is relatively straightforward for GPs with differentiable covariance functions, with hyperparameters fixed during the T -step prediction. In other models the sensitivity of the variance term would be much more difficult to estimate than in the GP case.

Given the derivative $\partial J_T / \partial u(t)$ we can use standard algorithms to optimise the values of $u(t), \dots, u(t+T-1)$ to maximise J_T , i.e. we find

$$J_{opt} = \min_{u(t), \dots, u(t+T-1)} \frac{1}{T} \sum_{k=1}^T E\{(y_d(t+k) - y(t+k))^2\}$$

This can then be repeated in a step-wise fashion to produce a T -step-ahead model predictive controller. An obvious issue is the computational effort needed in each iteration of the optimisation, because of the need to recalculate the inverse of the covariance matrix of the training data, with each new training point.³

At $\frac{\partial J_T}{\partial u(t)} = 0$, the resulting optimal control signal is significantly affected by the derivative of the variance of the model over the prediction horizon. If we had not included the variance term in cost function (10), or if we were in a region of the state-space where the variance was zero, the optimal control law would be different, and designers typically have to tune control effort damping parameters. The regularisation inherent to the GP approach make a control effort penalty constant, or regulariser unnecessary in many applications.

³ An alternative to optimisation is to integrate over the variables of interest, by performing numerical integration of the joint density of future y and u using Monte Carlo techniques. Possibly the simplest strategy is the Metropolis method, which may be slow due to very small possible proposal widths. A more promising method is Hybrid Monte Carlo (HMC) Duane et al. (1987); MacKay (1998), which takes the derivative of the density into account and avoids slow random walk behaviour. HMC would be easily implemented since all the required derivatives are readily available, although the computational cost will still be significant.

4. SIMULATION

The example considers the following non-linear functions:

$$y(t+1) = \frac{y(t)y(t-1)y(t-2)u(t-1)(y(t-2)-1) + u}{1 + y(t-1)^2 + y(t-2)^2}$$

where $\mathbf{x} = [y(t) \ y(t-1) \ y(t-2) \ u(t-1)]^T$, from Narendra and Parthasarathy (1990). The observation noise has a variance $\sigma^2 = 0.001$, and we had 5 initial data points. We simulated a tracking experiment, where the model had to follow the reference trajectory (blue), while learning the behaviour of the system. We used a standard optimisation routine (MATLAB's `fminsearch`) to find the optimal control values at each point, and did not use explicit sensitivity equations to estimate the gradient. At each timestep, after $u(t)$ has been calculated, applied, and the output observed, we add the information $\mathbf{x}(t)$, $u(t)$, $y(t+1)$ to the training set, and the new Σ_1 increases in size to $N+1 \times N+1$. We then optimise the hyperparameters of the covariance function to further refine the model. The results are shown in Figure 2.



In simulation, including the variance term in the cost function led to more robust control and also led to a reduction in tracking error over the sample indicated in the figure, when compared to an identical model which had ignored the variance term in the cost function (the certainty equivalence approach).

T -step ahead predictions tend to show more caution for larger values of T . When the model state moves towards higher uncertainty regions (due to system complexity, or lack of training data), the control signal will be appropriately damped, and provides a smoother control signal. This leads to poorer immediate tracking performance, but greater robustness in the face of uncertainty. The choice of an appropriate prediction horizon T is obviously very important. A feature of the predictions is that as T increases, the mean prediction eventually tends towards zero (for a zero mean GP), and the variance of the prediction will also saturate at prior levels.


The hyperparameters in Figure 2(b) make few rapid changes, seeming well-behaved during learning.

5. CONCLUSIONS



Gaussian process priors can provide a flexible, general modelling tool for adaptive nonlinear control problems. The framework provides analytic estimates of model uncertainty and of derivatives of both model mean and uncertainty. In this paper we presented an approximate method for propagation of uncertainty for multistep predictions, which also allows you to calculate an analytic derivative of the cost function with respect to control actions. The experimental comparisons presented here used numerical optimisation over the prediction horizon, and found that the inclusion



(a) Simulation of nonlinear GP-based controller, showing it controlling the target system in an adaptive manner from an initial condition of 5 training points, prediction horizon $T = 1$



(b) Covariance function hyperparameters adapting for $T = 1$ case.



(c) Simulation of nonlinear GP-based controller, prediction horizon $T = 3$

Fig. 2. Simulation results for nonlinear system, showing modelling accuracy, control signals and development of the hyperparameters.

of the variance estimate, derived from estimates of propagated errors improve tracking performance, and robustness.

As a consequence of this multistep predictive controller, and the retention of the model variance in the cost function, we automatically get appropriate regularising behaviour (*caution*) which adapts to the uncertainty local to the current state.

ACKNOWLEDGEMENTS

RM-S & DS are grateful for support from FONDECYT Project 700397. RM-S, AG & CR gratefully acknowledge the support of the *Multi-Agent Control* Research Training Network (EC TMR grant HPRN-CT-1999-00107), and RM-S for the EPSRC grant *Modern statistical approaches to off-equilibrium modelling for nonlinear system control* GR/M76379/01.

References

- M. Agarwal and D. E. Seborg. Self-tuning controllers for nonlinear systems. *Automatica*, (2):209–214, 1987.
- S. Bittanti and L. Piroddi. Neural implementation of GMV control schemes based on affine input/output models. *Proc. IEE Control Theory Appl.*, (6):521–530, 1997.
- S. Duane, A. D. Kennedy, and D. Roweth. Hybrid Monte Carlo. *Phys. Lett. B*, 195:216–222, 1987.
- S. Fabri and V. Kadirkamanathan. Dual adaptive control of nonlinear stochastic systems using neural networks. *Automatica*, 34(2):245–253, june 1998.
- A. A. Fel'dbaum. Dual-control theory I-IV. *Avtomatika i Telemekhanika*, 21 and 22(9,11 (vol. 21), 1, 2 (vol. 22)):1240–1249, 1453–1464, 3–16,129–142, 1960.
- N.M. Filatov and H. Unbehauen. Survey of adaptive dual control methods. *Proc. IEE Control Theory Appl.*, (1):119–128, 2000.
- N.M. Filatov, H. Unbehauen, and U. Keuchel. Dual pole placement controller with direct adaptation. *Automatica*, 33(1):113–117, 1997.
- A. Girard, C. E. Rasmussen, and R. Murray-Smith. Gaussian process priors with uncertain inputs: multiple-step ahead prediction. Technical Report TR-2002-119, Department of Computing Science, University of Glasgow, 2002.
- A. Girard, C. E. Rasmussen, J. Quinero-Candela, and R. Murray-Smith. Gaussian process priors with uncertain inputs – application to multiple-step ahead time series forecasting. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information processing Systems 15*, 2003.
- J. Kocijan, R. Murray-Smith, C.E. Rasmussen, and B. Likar. Predictive control and Gaussian process models. Submitted for publication, 2003.
- D. J. Leith, R. Murray-Smith, and W. E. Leithead. Nonlinear structure identification: A Gaussian Process prior/Velocity-based approach. In *Control 2000, Cambridge*, 2000.
- D.J.C. MacKay. Introduction to Monte Carlo methods. In M. I. Jordan, editor, *Learning and Inference in Graphical Models*, pages 175–204. Kluwer, 1998.
- R. Murray-Smith, T. A. Johansen, and R. Shorten. On transient dynamics, off-equilibrium behaviour and identification in blended multiple model structures. In *European Control Conference, Karlsruhe, 1999*, pages BA–14, 1999.
- R. Murray-Smith and D. Sbarbaro. Nonlinear adaptive control using non-parametric Gaussian process prior models. In *15th IFAC World Congress on Automatic Control, Barcelona*, 2002.
- K. Narendra and K. Parthasarathy. Identification and Control of Dynamical Systems Using Neural Networks. *IEEE Trans. on Neural Networks*, 1(1):4–27, January/February 1990.
- R. M. Neal. Monte Carlo implementation of Gaussian process models for Bayesian regression and classification. Technical Report 9702, Department of Statistics, University of Toronto, 1997.
- A. O'Hagan. On curve fitting and optimal design for regression (with discussion). *Journal of the Royal Statistical Society B*, 40:1–42, 1978.
- A. O'Hagan. Some Bayesian numerical analysis. In J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, editors, *Bayesian Statistics 4*, pages 345–363. Oxford University Press, 1992.
- J. Quinero-Candela, A. Girard, J. Larsen, and C. E. Rasmussen. Propagation of uncertainty in Bayesian kernels models – application to multiple-step ahead forecasting. In *ICASSP, Hong Kong*, 2003.
- D. Sbarbaro, N.M. Filatov, and H. Unbehauen. Adaptive dual controller for a class of nonlinear systems. In *Proceedings of the IFAC Workshop on Adaptive systems in Control and Signal Processing*, pages 28–33, Glasgow, U.K., 1998.
- E. Solak, R. Murray-Smith, W. E. Leithead, D. J. Leith, and C. E. Rasmussen. Derivative observations in Gaussian process models of dynamic systems. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information processing Systems 15*, 2003.
- C. K. I. Williams. Prediction with Gaussian processes: From linear regression to linear prediction and beyond. In M. I. Jordan, editor, *Learning and Inference in Graphical Models*, pages 599–621. Kluwer, 1998.
- B. Wittenmark. Adaptive dual control methods: An overview. In *Preprints of 5th IFAC symposium on Adaptive Systems in Control and Signal processing, Hungary*, pages 67–72, 1995.