



Supervisory observer for parameter and state estimation of nonlinear systems using the DIRECT algorithm

Michelle Chong, Romain Postoyan, Sei Zhen Khong, Dragan Nesic

► To cite this version:

Michelle Chong, Romain Postoyan, Sei Zhen Khong, Dragan Nesic. Supervisory observer for parameter and state estimation of nonlinear systems using the DIRECT algorithm. 56th IEEE Conference on Decision and Control, CDC 2017, Dec 2017, Melbourne, Australia. <hal-01652497>

HAL Id: hal-01652497

<https://hal.science/hal-01652497v1>

Submitted on 25 Mar 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Supervisory observer for parameter and state estimation of nonlinear systems using the DIRECT algorithm

Michelle S. Chong, Romain Postoyan, Sei Zhen Khong and Dragan Nešić

Abstract—A supervisory observer is a multiple-model architecture, which estimates the parameters and the states of nonlinear systems. It consists of a bank of state observers, where each observer is designed for some nominal parameter values sampled in a known parameter set. A selection criterion is used to select a single observer at each time instant, which provides its state estimate and parameter value. The sampling of the parameter set plays a crucial role in this approach. Existing works require a sufficiently large number of parameter samples, but no explicit lower bound on this number is provided. The aim of this work is to overcome this limitation by sampling the parameter set automatically using an iterative global optimisation method, called Dividing RECTangles (DIRECT). Using this sampling policy, we start with $1 + 2n_p$ parameter samples where n_p is the dimension of the parameter set. Then, the algorithm iteratively adds samples to improve its estimation accuracy. Convergence guarantees are provided under the same assumptions as in previous works, which include a persistency of excitation condition. The efficacy of the supervisory observer with the DIRECT sampling policy is illustrated on a model of neural populations.

I. INTRODUCTION

Multiple-model approaches have traditionally been employed in the stochastic setting for estimation algorithms. Strong contenders include the Gaussian sum estimators [3, Section 8.5], the variable structure multiple-model framework [14], [17] and statistical filters such as particle filters [4]. In the deterministic domain, the focus has been mainly on improving robustness of adaptive schemes for control purposes [16], [2], [15, Chapter 6]. In the context of *parameter* and *state* estimation for dynamical systems, the multiple-model approach provides an alternative to other online estimation algorithms, such as adaptive observers [5, Chapter 7]. The idea is to design a bank of *state* observers (also known as a multi-observer), in which the state observer is designed for some sampled parameter values in a given set. Parameter and state estimates are then derived by combining the information from a subset of the state observers. This approach has the added benefit of modularity in the separation of the state and parameter estimation problems. This has been pursued

for continuous-time linear systems by [3], [14, Section 8.5] in the stochastic setting, and by [1], [16] in the deterministic setting, to name a few references. Most recently, we extended the deterministic setup to continuous-time nonlinear systems in [7] inspired by results in supervisory control [15, Chapter 6], which we called a *supervisory observer*.

In [7], the unknown system parameters are assumed to be constant and to belong to a known compact parameter set. Finite samples are drawn from this parameter set and a state observer is designed for each sample in a way that is robust against parameter mismatches. A selection criterion based on the mismatch between the estimated output of each observer and measured output from the plant then provides the state and parameter estimates at any given time. We showed in [7] that the estimates are guaranteed to converge to be within a required margin, as long as the parameter set is sufficiently densely sampled and a persistency of excitation condition holds. To potentially ease the need for a large number of samples, we also introduced a dynamic sampling scheme which updates the sampling of the parameter set by iteratively zooming in on the region of the parameter set where the true parameter is ‘most likely’ to reside. However, a major drawback is that the user is required to choose a pre-determined number of parameter samples at the start of the algorithm, which is hard to estimate.

In this paper, we aim to overcome this major drawback by automatically sampling the parameter set using an iterative global optimisation method for Lipschitz cost functions on a compact domain, called DIRECT (DIviding RECTangles) which was initially proposed in [12]. This sampling procedure starts from sampling the center point of the parameter set and additional given sample points in its neighbourhood. At subsequent iterations, DIRECT takes additional samples. This is achieved by dividing the domain into non-overlapping hyperrectangles and sampling the center of each one. A lower bound on the cost function in each hyperrectangle is obtained using the fact that the cost function is Lipschitz. The search at the local level is achieved by identifying potentially optimal hyperrectangles based on its lower bound of the cost function and its size, which are then further divided. In other words, the algorithm automatically generates additional samples in potentially optimal regions. The procedure is stopped once a pre-calculated number of iterations is reached.

In our supervisory observer, the cost function used by the DIRECT algorithm is an integral form of the output mismatch of each observer and the plant over a finite time interval, which we call the monitoring signal. The domain of the cost function is the compact parameter set to which the

M. Chong is with the Department of Automatic Control, Lund University, Sweden. Her work is supported by the LCCC Linnaeus Center and the ELLIIT Excellence Center at Lund University. michelle.chong@control.lth.se

R. Postoyan is with Université de Lorraine, CRAN, UMR 7039 and the CNRS, CRAN, UMR 7039, France. His work is partially supported by the ANR under the grant SEPICOT (ANR 12 JS03 004 01). romain.postoyan@univ-lorraine.fr

S. Khong is with the Institute for Mathematics and its Applications, The University of Minnesota, Minneapolis, MN 55455, USA. szkhong@umn.edu

D. Nešić is with the Department of Electrical and Electronic Engineering, University of Melbourne, Australia. dnesic@unimelb.edu.au

true parameter belongs. We wait a sufficiently long interval between each iteration to allow the transient effects of each state observer to decay in order to obtain an improved accuracy. More observers are added as DIRECT updates the sampling of the parameter set. Upon the termination of DIRECT, we implement the last chosen observer for the remainder of the algorithm's run-time. Contrary to the old dynamic sampling policy in [7], the DIRECT sampling policy provides the following improvements. First, DIRECT does not require the user to estimate the number of observers needed beforehand. Instead, we start with $1 + 2n_p$ observers, where n_p is the dimension of the parameter set. Then, DIRECT automatically takes samples until a given termination time, which is a clear advantage over the old sampling policy from the user's perspective. Second, the DIRECT policy potentially eases computational burden as we no longer need to implement all the required number of observers in parallel from initialisation. Third, after a pre-computed time, the supervisory observer implements only one observer for the rest of the run-time and thereby further reducing the computational resources required.

The DIRECT algorithm was employed in the context of extremum-seeking in [13] where the global minimum of the steady-state input-output map of a nonlinear time-invariant dynamical system is found using DIRECT without knowledge of the system model. Our setup is 'gray box' in nature, where the structure of the parameterised nonlinear dynamical system is known, but not the states and parameters. The supervisory observer we propose can be viewed as the problem of online extremisation in dynamical systems, where the estimation algorithm aims to provide estimates that minimises the estimation error based on the steady-state behaviour of each observer by waiting sufficiently long between sampling times.

The paper is structured as follows. We introduce the notation in Section II and formulate the problem in Section III. Section IV describes the supervisory observer with the DIRECT sampling policy in detail and we provide convergence guarantees in Section V. We illustrate the efficacy of the proposed algorithm in Section VI by revisiting a model of neuron populations considered in [7]. Due to space constraints, all proofs can be found in the full version of this paper [8].

II. PRELIMINARIES

Let $\mathbb{R} = (-\infty, \infty)$, $\mathbb{R}_{\geq 0} = [0, \infty)$, $\mathbb{R}_{> 0} = (0, \infty)$, $\mathbb{N} = \{0, 1, \dots\}$ and $\mathbb{N}_{\geq 1} = \{1, 2, \dots\}$. Let (u, v) where $u \in \mathbb{R}^{n_u}$ and $v \in \mathbb{R}^{n_v}$ denote the vector $(u^T, v^T)^T$. The smallest integer greater than $v \in \mathbb{R}$ is denoted by $\lceil v \rceil$. For a vector $x = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n$, the ∞ -norm of x is denoted $|x| := \max\{|x_1|, |x_2|, \dots, |x_n|\}$. Let $\mathcal{H}(p_c)$ denote the hypercube centered at $p_c \in \mathbb{R}^{n_p}$ where the distance of p_c to the edge is 1, i.e. $\mathcal{H}(p_c) := \{p \in \mathbb{R}^{n_p} : |p - p_c| \leq 1\}$. Hence, $\Delta\mathcal{H}(0)$ is the hypercube with center point at the origin and its distance to the edge is Δ , i.e. $\Delta\mathcal{H}(0) := \{p \in \mathbb{R}^{n_p} : |p| \leq \Delta\}$. For any $\Delta > 0$, the set of piecewise continuous functions from $\mathbb{R}_{\geq 0}$ to $\Delta\mathcal{H}(0)$ is denoted \mathcal{M}_Δ .

A continuous function $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a class \mathcal{K} function, if it is strictly increasing and $\alpha(0) = 0$; additionally, if $\alpha(r) \rightarrow \infty$ as $r \rightarrow \infty$, then α is a class \mathcal{K}_∞ function. A continuous function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a class \mathcal{KL} function, if: (i) $\beta(\cdot, s)$ is a class \mathcal{K} function for each $s \geq 0$; (ii) $\beta(r, \cdot)$ is non-increasing and (iii) $\beta(r, s) \rightarrow 0$ as $s \rightarrow \infty$ for each $r \geq 0$.

III. PROBLEM FORMULATION

Consider the following nonlinear system

$$\begin{aligned}\dot{x} &= f(x, p^*, u), \\ y &= h(x, p^*),\end{aligned}\tag{1}$$

where the state is $x \in \mathbb{R}^{n_x}$, the measured output is $y \in \mathbb{R}^{n_y}$, the measured input is $u \in \mathbb{R}^{n_u}$ and the *unknown* parameter vector $p^* \in \Theta \subset \mathbb{R}^{n_p}$ is constant. We assume that Θ is a known, normalised unit hypercube¹ $\mathcal{H}(p_c)$. The function $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_p} \rightarrow \mathbb{R}^{n_x}$ is locally Lipschitz and $h : \mathbb{R}^{n_x} \times \mathbb{R}^{n_p} \rightarrow \mathbb{R}^{n_y}$ is continuously differentiable.

We aim to estimate the parameter p^* and the state x of system (1) assuming that the output y and the input u are measured. To this end, we use the supervisory observer architecture we proposed in [7]. The idea is the following: the parameter set Θ is sampled and a state observer is designed for each parameter sample, forming a bank of observers. One observer is chosen to provide its state estimate and parameter value at any time instant based on a criterion. The results in [7] showed that both estimates converge to within a tunable margin of the corresponding true values provided that the number of samples is sufficiently large, under some conditions. The sampling policies in [7] require a sufficiently large number of samples and no explicit lower bound on this number was provided in [7]. The objective of this paper is to overcome these issues by sampling the parameter set Θ dynamically in a smarter manner using the DIRECT optimisation algorithm, such that the number of samples is automatically generated by the algorithm based on the desired estimation accuracy. In other words, the user no longer has to set the number of samples needed from the start, as the algorithm automatically achieves the required number of samples after some iterations and terminates at a pre-computed time.

We maintain all assumptions made in the original setup in [7]. In particular, we assume the following boundedness property for system (1).

Assumption 1: For any initial condition $x(0)$ and any piecewise-continuous input u , system (1) is forward complete and generates unique uniformly bounded solutions, i.e. for any $\Delta_x, \Delta_u \geq 0$, there exists a constant $K_x = K_x(\Delta_x, \Delta_u) > 0$ such that for all $x(0) \in \Delta_x\mathcal{H}(0)$ and $u \in \mathcal{M}_{\Delta_u}$, the corresponding unique solution to (1) satisfies $|x(t)| \leq K_x$, for all $t \geq 0$. \square

¹Any compact parameter set Θ can be embedded in a hyperrectangle and then normalised to be a unit hypercube. Hence, the assumption of Θ being a hypercube of edge 1 is made with no loss in generality as long as the parameter p^* lies in a known (potentially large) compact set.

IV. SUPERVISORY OBSERVER WITH THE DIRECT SAMPLING POLICY

In this section, we first recall the supervisory observer architecture proposed in [7]. We then explain how DIRECT is implemented to sample the parameter set Θ .

A. Architecture

The sampling of the parameter set Θ is carried out iteratively at each update time instant t_k , $k \in \mathbb{N}$ satisfying

$$t_{k+1} - t_k =: T_d, \quad (2)$$

where $T_d > 0$ is a design parameter.

Let $\hat{P}(k)$ denote the set of parameter sample points generated at time t_k , $k \in \mathbb{N}_{\geq 1}$; and the generation of these points will be described in Section IV-B. Also, let $\hat{\Theta}(k) := \hat{\Theta}(k-1) \cup \hat{P}(k)$ be the set of all parameter sample points from t_0 to t_k , where their corresponding cardinalities are $N_{\hat{P}}(k)$ and $N_{\hat{\Theta}}(k)$, respectively. Consequently, we have that $\hat{\Theta}(k-1)$ is a subset of $\hat{\Theta}(k)$. At the initial time t_0 , $\hat{P}(0)$ and $\hat{\Theta}(0)$ consist of the centre point of Θ and $2n_p$ additional points near it, i.e. $\hat{P}(0) = \hat{\Theta}(0) = \{p_c, p_c \pm \delta e_i\}$, where $\delta = 1/3$ and e_i is the i -th unit vector of \mathbb{R}^{n_p} . Hence, $N_{\hat{\Theta}}(0) = N_{\hat{P}}(0) = 1 + 2n_p$. In other words, the number of observers implemented over the first interval of time $[t_0, t_1)$ is $1 + 2n_p$. At subsequent update time instants t_k , $k \in \mathbb{N}$, new sampling points are added and an observer is designed for each of the newly generated samples $p_i^k \in \hat{P}(k)$, for $i \in \{1, \dots, N_{\hat{P}}(k)\}$, as follows

$$\begin{aligned} \dot{\hat{x}}_i &= \hat{f}(\hat{x}_i, p_i^k, u, y), & \forall t \in [t_k, t_{k+1}), \\ \hat{y}_i &= h(\hat{x}_i, p_i^k), \end{aligned} \quad (3)$$

where the function \hat{f} is continuously differentiable. At initialisation $t_0 = 0$, we set arbitrary initial conditions $\hat{x}_i(t_0) \in \mathbb{R}^{n_x}$ for $i \in \{1, \dots, 1 + 2n_p\}$. At subsequent update times t_k , $k \in \mathbb{N}_{\geq 1}$, each ‘old’ observer is kept running, i.e. the observers designed for each $p \in \hat{\Theta}(k-1)$, and each new observer is initialised as follows

$$\hat{x}_i(t_k^+) = \hat{x}_{\sigma(t_k)}(t_k), \quad (4)$$

where σ chooses one observer from the bank of observers and is defined below in (11).

Remark 1: We only initialise the ‘new’ observers according to (4). Our results also apply when all the observers are reinitialised according to (4). However, no significant advantage is seen both in the analysis and when implemented in simulations in Section VI. \square

Denoting the state estimation error as $\tilde{x}_i := \hat{x}_i - x$, the output error as $\tilde{y}_i := \hat{y}_i - y$ and the parameter error as $\tilde{p}_i^k := p_i^k - p^*$, for all $i \in \{1, \dots, N_{\hat{\Theta}}(k)\}$, we obtain the following state estimation error systems for all $t \in [t_k, t_{k+1})$,

$$\begin{aligned} \dot{\tilde{x}}_i &= \hat{f}(\tilde{x}_i + x, \tilde{p}_i^k + p^*, u, y) - f(x, p^*, u) \\ &=: F(\tilde{x}_i, \tilde{p}_i^k, p^*, u, x), \\ \tilde{y}_i &= h(\tilde{x}_i + x, \tilde{p}_i^k + p^*) - h(x, p^*) =: H(\tilde{x}_i, \tilde{p}_i^k, p^*, x), \\ \tilde{x}_i(t_k^+) &= \begin{cases} \hat{x}_{\sigma(t_k)}(t_k) - x(t_k), & i \in \hat{P}(k), \\ \hat{x}_i(t_k) - x(t_k), & i \in \hat{\Theta}(k) \setminus \hat{P}(k). \end{cases} \end{aligned} \quad (5)$$

All the observers are designed such that the following property holds.

Assumption 2: Consider the state estimation error system (5) for $i \in \{1, \dots, N\}$, $N \in \mathbb{N}_{\geq 1}$. Let $\tilde{\Theta} := \{p - p^* : p \in \Theta \text{ and } p^* \in \Theta\}$. There exist scalars $a_1, a_2, \lambda_0 > 0$ and a continuous non-negative function $\tilde{\gamma} : \tilde{\Theta} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}_{\geq 0}$ where $\tilde{\gamma}(0, z, \bar{z}) = 0$ for all $z \in \mathbb{R}^{n_x}$, $\bar{z} \in \mathbb{R}^{n_u}$, such that for any $\tilde{p}_i \in \tilde{\Theta}$, there exists a continuously differentiable function $V_i : \mathbb{R}^{n_x} \rightarrow \mathbb{R}_{\geq 0}$ which satisfies the following for all $u \in \mathbb{R}^{n_u}$, $\tilde{x}_i \in \mathbb{R}^{n_x}$, $x \in \mathbb{R}^{n_x}$

$$a_1 |\tilde{x}_i|^2 \leq V_i(\tilde{x}_i) \leq a_2 |\tilde{x}_i|^2, \quad (6)$$

$$\frac{\partial V_i}{\partial \tilde{x}_i} F(\tilde{x}_i, \tilde{p}_i, p^*, u, x) \leq -\lambda_0 V_i(\tilde{x}_i) + \tilde{\gamma}(\tilde{p}_i, x, u). \quad (7)$$

\square

When there is no parameter mismatch $\tilde{\gamma}(0, x, u) = 0$, Assumption 2 implies that state estimates converge exponentially to the true state for all initial conditions. When there is a parameter mismatch $\tilde{p}_i \neq 0$, the state estimation error system satisfies an input-to-state exponential stability property with respect to \tilde{p}_i in view of Assumption 1. Examples of system (1) for which observers (3) can be designed are provided in Section VI of [7]. This includes linear systems and a class of nonlinear systems with monotone nonlinearities, such as the neural example considered later in simulations in Section VI. We assume that the output error of each of the observer \tilde{y}_i satisfies the following property.

Assumption 3: Consider the state estimation error system (5) for $i \in \{1, \dots, N\}$, $N \in \mathbb{N}_{\geq 1}$. For all $\Delta_{\tilde{x}}, \Delta_x, \Delta_u > 0$, there exist a constant $T_f = T_f(\Delta_{\tilde{x}}, \Delta_x, \Delta_u) > 0$ and a class \mathcal{K}_∞ function $\alpha_{\tilde{y}} = \alpha_{\tilde{y}}(\Delta_{\tilde{x}}, \Delta_x, \Delta_u)$ such that for all $\tilde{x}_i(0) \in \Delta_{\tilde{x}} \mathcal{H}(0)$, $x(0) \in \Delta_x \mathcal{H}(0)$, for any $u \in \mathcal{M}_{\Delta_u}$, and for all $\tilde{p}_i \in \tilde{\Theta}$, the corresponding solution to (5) satisfies

$$\int_{t-T_f}^t |\tilde{y}_i(\tau)|^2 d\tau \geq \alpha_{\tilde{y}}(|\tilde{p}_i|), \quad \forall t \geq T_f. \quad (8)$$

\square

The inequality (8) is a variant of the persistency of excitation (PE) condition found in many adaptive and identification schemes [6]. In [7, Proposition 1], we relate the PE-like condition (5) to the classical PE condition in the literature such that Assumption 3 can be guaranteed a priori for certain classes of systems.

The output error from each observer forms the monitoring signal used in the DIRECT sampling algorithm, defined as follows

$$\mu(p, t_k, t) := \int_{t_k}^t \exp(-\lambda(t-s)) |\tilde{y}(p, s)|^2 ds, \quad \forall t \geq t_k, \quad (9)$$

where $\lambda > 0$ is a design parameter and we use $\tilde{y}(p, t)$ in place of $\tilde{y}_i(t)$ to highlight its dependance on the parameter p . Our parameter and state estimates are chosen from the bank of observers to be, for any $t \geq 0$,

$$\hat{p}(t) := \hat{p}_{\sigma(t)}, \quad \hat{x}(t) := \hat{x}_{\sigma(t)}(t), \quad (10)$$

where σ is given by

$$\sigma(t) \in \arg \min_{i \in \{1, \dots, N_{\hat{\Theta}}(k)\}} \mu(p_i, t_k, t), \quad \forall t \in [t_k, t_{k+1}). \quad (11)$$

We will see in Section IV-C that the DIRECT sampling policy terminates according to a criterion described in Section IV-C at t_{k^*} , $k^* \in \mathbb{N}$, after which we only implement the observer for the last chosen sample at $p_{\sigma(t_{k^*})}$, i.e. the supervisory observer is reduced to one observer as follows, for all $t \in [t_{k^*}, \infty)$,

$$\begin{aligned} \dot{\hat{x}}_{\sigma(t_{k^*})} &= \hat{f}(\hat{x}_{\sigma(t_{k^*})}, p_{\sigma(t_{k^*})}, u, y) \\ \hat{y}_{\sigma(t_{k^*})} &= h(\hat{x}_{\sigma(t_{k^*})}, p_{\sigma(t_{k^*})}). \end{aligned} \quad (12)$$

B. The DIRECT sampling policy

We generate the sampling points of the parameter set Θ at every update time instant t_k , $k \in \mathbb{N}$ according to the DIRECT sampling policy as follows

- i. At the initial time $t_0 = 0$, set $k = 0$ (iteration counter).
- ii. At $t = t_1$, evaluate $\mu(p_c, t_0, t_1)$, where we recall that p_c is the center point of Θ and follow the procedure for dividing the hyperrectangle². Set $\hat{\mu}_1 = \mu(p_c, t_0, t_1)$. Increment the iteration counter to $k = 1$.
- iii. Identify the set $S(k)$ of the indices of potentially optimal hyperrectangles, i.e. for every $i \in \{1, \dots, N(k)\}$ samples for which there exists $L_k > 0$ such that,

$$\begin{aligned} \mu(p_i, t_{k-1}, t_k) - L_k d_i &\leq \mu(p_j, t_{k-1}, t_k) - L_k d_j, \\ j &\in \{1, \dots, N(k)\}, \\ \mu(p_i, t_{k-1}, t_k) - L_k d_i &\leq \hat{\mu}_{k-1} - \epsilon |\hat{\mu}_{k-1}|, \quad \epsilon > 0, \end{aligned} \quad (13)$$

where d_i denotes the distance from the centre point p_i to a vertex of the hyperrectangle i . Remarks 2 and 3 below discuss the search parameters L_k and ϵ in (13).

- iv. For each potentially optimal hyperrectangle indexed by $j \in S(k)$, subdivide the hyperrectangle indexed by j according to the procedure for dividing hyperrectangles².
- v. Increment the iteration counter, $k^+ = k + 1$ and set the estimate

$$\hat{\mu}_k := \min_{i \in \{1, \dots, N(k)\}} \mu(p_i, t_{k-1}, t_k). \quad (14)$$

- vi. Go to Step iii until k^* iterations have been reached. The selection of k^* is specified in Section IV-C.

Remark 2: The search parameter L_k in (13) can be thought of as a rate-of-change constant. In the case where the Lipschitz constant \tilde{L}_k of function μ is known, the user may restrict the search for L_k to $L_k \leq \tilde{L}_k$. However, the knowledge of Lipschitz constant is not a requirement for convergence to the global minimum and efficient algorithms can be used to find L_k such as one called Graham's scan, see [12]. \square

²Procedure for dividing a hyperrectangle: Given a hyperrectangle at time t_k for $k \in \mathbb{N}$, identify the dimensions $i \in I \subseteq \{1, \dots, n_p\}$ in which the hyperrectangle has the maximum edge length and let δ be a third of this value. Divide the hyperrectangle containing the sample point p_j into thirds according to the dimensions in $j \in I$, in ascending order of $\min\{\mu(p_j - \delta e_i, t_{k-1}, t_k), \mu(p_j + \delta e_i, t_{k-1}, t_k)\}$, where e_i is the i -th unit vector.

Remark 3: The search parameter $\epsilon > 0$ in (13) ensures that at the current iteration k , only hyperrectangles with cost μ that is much smaller than the minimum cost of the previous iteration $\hat{\mu}_{k-1}$ are identified as potentially optimal. Computational results in [12] show that DIRECT is insensitive to the choice of ϵ and a good value for ϵ ranges from 10^{-3} to 10^{-7} . \square

In practice, the DIRECT algorithm can be implemented with modifications of the code for DIRECT from [9]. This is due to the dynamic cost μ (c.f. (9)) in our setup as opposed to the static cost the code was originally written for.

C. Termination criterion of the DIRECT sampling policy

The final piece of the algorithm is the termination time of the DIRECT sampling policy. Before doing so, a critical sampling property required to show convergence is the fact that DIRECT will generate samples such that the distance between the true parameter p^* and the closest sample in $\hat{\Theta}(k)$ as defined below

$$d(p^*, \hat{\Theta}(k)) := \min_{p \in \hat{\Theta}(k)} |p^* - p| \quad (15)$$

tends to zero if $N_{\hat{\Theta}(k)}$ tends to infinity with increasing k . We formalize this in the next lemma.

Lemma 1: The DIRECT method of sampling the parameter set Θ results in the following property,

$$N(k) \rightarrow \infty \implies d(p^*, \hat{\Theta}(k)) \rightarrow 0, \quad (16)$$

where we recall that $\hat{\Theta}(k)$ is the set of all sample points at time $t \in [t_k, t_{k+1})$. \square

Lemma 1 can be used to derive a termination time for the DIRECT sampling policy. This is the purpose of the next lemma which follows from Theorem 4.2 in [10] and Lemma 1.

Lemma 2: Given any $d^* > 0$, let $k^* := 3^{n_p-1} \left(\frac{3^{n_p(i+1)} - 1}{3^{n_p} - 1} \right)$, where i satisfies

$$\frac{(n_p 3^{-2i})^{1/2}}{2} \leq d^*. \quad (17)$$

Then, the DIRECT algorithm described in Section IV-B samples the parameter set $\Theta \subset \mathbb{R}^{n_p}$ such that

$$d(p^*, \hat{\Theta}(k)) \leq d^*, \quad \forall k \geq k^*. \quad (18)$$

\square

Hence, given a desired bound $d^* > 0$ on the distance between the true parameter p^* and the closest sampling point, the algorithm can be terminated once the number of iterations reaches k^* as defined in Lemma 2. In practice, once we have decided on d^* , we know that after $k^* T_d$ units of time, a single observer can be run as described at the end of Section IV-A.

Remark 4: The estimation of the number of iterations k^* to achieve the desired resolution (18) in Lemma 2 is calculated based on the assumption that only one hyperrectangle gets divided at each iteration. In reality, the number of hyperractangles identified to be potentially optimal according to (13) can be more than one (see Section VI) and hence, the calculations in Lemma 2 is an over-approximation. The

calculation is tight only when the cost function $\mu(p, t, \tau)$ is constant for all $p \in \Theta$ and $t, \tau > 0$. \square

V. CONVERGENCE GUARANTEES

We provide the following convergence guarantees and its proof is provided in [8].

Theorem 1: Consider system (1), the multi-observer (3)-(4) and (12), the monitoring signals (9), the selection criterion (11) and the estimates (10), under Assumptions 1-3 and the DIRECT sampling policy. Given any $\Delta_x, \Delta_{\tilde{x}}, \Delta_u > 0$, $d^* > 0$ and $\eta > 0$, there exist a class \mathcal{K}_∞ function $\gamma_{\tilde{x}}$, a sufficiently large $T > 0$ such that for any sampling interval $T_d \geq T$, a class \mathcal{K}_∞ function $\nu_{\tilde{p}}$ and a constant $T^* > 0$ such that the following holds

$$\begin{aligned} |\tilde{p}_{\sigma(t)}(t)| &\leq \nu_{\tilde{p}}(d^*) + \eta, \quad \forall t \geq T^*, \\ \limsup_{t \rightarrow \infty} |\tilde{x}_{\sigma(t)}(t)| &\leq \gamma_{\tilde{x}}(d^*) + \eta, \end{aligned} \quad (19)$$

for all $(x(0), \tilde{x}(0)) \in \Delta_x \mathcal{H}(0) \times \Delta_{\tilde{x}} \mathcal{H}(0)$ and for any $u \in \mathcal{M}_{\Delta_u}$. \square

The convergence guarantees (19) show that the upper bound of the estimation error on the parameter and states decreases with d^* , after a sufficiently long time. Additionally, note that η can be taken to be as small as desired by increasing T . Hence, the estimation accuracy can be tuned by adjusting d^* and by ensuring a sufficiently large sampling time T_d .

VI. NUMERICAL SIMULATIONS: A NEURAL MASS MODEL

We apply the setup on a neural mass model in [11]. This model describes the dynamics of the mean membrane potential (states) and the synaptic gain (parameters) of neuron populations, which is used to capture the electroencephalogram (measurement) patterns related to various brain activities. By taking the states to be $x = (x_{11}, x_{12}, x_{21}, x_{22}, x_{31}, x_{32}) \in \mathbb{R}^6$ and the *unknown* parameter vector to be $p^* = (p_1^*, p_2^*) \in \mathbb{R}^2$, which is known to reside in a compact set $\Theta := [2, 8] \times [22, 28]$ which we normalise to a unit hypercube. Its dynamics is given by

$$\dot{x} = Ax + G(p^*)\gamma(Hx) + B(p^*)\phi(u, y), \quad y = Cx, \quad (20)$$

where $A = \text{diag}(A_a, A_a, A_b)$, where $A_a = \begin{bmatrix} 0 & 1 \\ -a^2 & -2a \end{bmatrix}$

$$\text{and } A_b = \begin{bmatrix} 0 & 1 \\ -b^2 & -2b \end{bmatrix}, \quad G(p) = \begin{bmatrix} 0_{3 \times 1} & 0_{3 \times 1} \\ p_1 a c_2 & 0 \\ 0 & 0 \\ 0 & p_2 b c_4 \end{bmatrix},$$

$$B(p) = \begin{bmatrix} 0 & 0 \\ p_1 a & 0 \\ 0 & 0 \\ 0 & p_1 a \\ 0_{2 \times 1} & 0_{2 \times 1} \end{bmatrix}, \quad H = \begin{bmatrix} c_1 & 0_{1 \times 5} \\ c_3 & 0_{1 \times 5} \end{bmatrix} \text{ and } C =$$

$\begin{bmatrix} 0 & 0 & 1 & 0 & -1 & 0 \end{bmatrix}$. We have used the notation $0_{r \times s}$ to denote an r by s matrix with all 0 entries. The known parameters are $a = 100$, $b = 50$, $c_1 = 135$, $c_2 = 108$, $c_3 = 33.75$, $c_4 = 33.75$ are assumed to be known. The nonlinear terms are $\gamma = (S, S)$ and $\phi(u, y) = (S(y), u)$,

where $S(v) := 2e_0/(1 + e^{\tau(v_0 - v)})$, for $v \in \mathbb{R}$, with known constants $e_0 = 2.5$, $v_0 = 6$, $r = 0.56$.

This model satisfies Assumption 1 because the matrix A is Hurwitz and the nonlinear terms only contain a bounded input $u \in \mathcal{M}_{\Delta_u}$ and the function S which is bounded. By also noting that the S is slope-restricted, we can employ Proposition 4 in [7] to design our multi-observer (3) to satisfy Assumption 2. See [7, Section VI-B] for details.

With $d^* = 0.8$, we calculate according to Lemma 2 that the termination iteration of the DIRECT algorithm is $k^* = 6$. Other parameters for the algorithms are $\lambda = 0.05$ in (9), sampling interval $T_d = 10$ s and DIRECT search parameter $\epsilon = 10^{-5}$. The results obtained are shown in Figures 1 and 2.

To compare DIRECT with the previous dynamic sampling scheme in [7], we ran both algorithms for $t_f = 100$ s with update instants at t_k , $k \in \{0, 1, \dots, 9\}$ and sampling interval $T_d = 10$ s. We compare the convergence time for a desired margin of the parameter estimation error, the average number of observers used during simulation time $[0, 100]$ s, the parameter estimation error at the end of the run-time t_f as well as the corresponding normalised state estimation error. We summarise this comparison in Table I, which shows that DIRECT achieves a significantly better parameter estimation accuracy. Namely, given the desired accuracy, DIRECT requires a fewer number of observers on average compared to the dynamic sampling policy in [7]. The average number of observers used for DIRECT decreases as the run-time t_f is increased since only one observer is used from $t = k^* T_d$ onwards, while the average number of observers used remains constant for the dynamic sampling policy in [7]. However, this trend does not necessarily translate to the state estimation error, due to the fact that the parameter mismatch gain function $\bar{\gamma}_{\tilde{x}}$ (c.f. Lemma 3 in [8]) can differ between observers.

REFERENCES

- [1] A P. Aguiar, V. Hassani, A.M. Pascoal, and M. Athans. Identification and convergence analysis of a class of continuous-time multiple-model adaptive estimators. *IFAC Proceedings Volumes*, 41(2):8605–8610, 2008.
- [2] B.D.O. Anderson, T.S. Brinsmead, F. De Bruyne, J. Hespanha, D. Liberzon, and A.S. Morse. Multiple model adaptive control. part 1: Finite controller coverings. *International Journal of Robust and Nonlinear Control*, 10(11-12):909–929, 2000.
- [3] B.D.O. Anderson and J.B. Moore. *Optimal filtering*. Prentice-Hall Englewood Cliffs, NJ, 1979.
- [4] M.S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002.
- [5] G. Besançon. *Nonlinear observers and applications*. Lecture notes in control and information sciences. Springer, 2007.
- [6] R. Bitmead. Persistence of excitation conditions and the convergence of adaptive schemes. *IEEE Transactions on Information Theory*, 30(2):183–191, 1984.
- [7] M. S. Chong, D. Nešić, R. Postoyan, and L. Kuhlmann. Parameter and state estimation of nonlinear systems using a multi-observer under the supervisory framework. *IEEE Transactions on Automatic Control*, 60(9):2336–2349, 2015.
- [8] M.S. Chong, R. Postoyan, S. Khong, and Nesic D. Supervisory observer for parameter and state estimation of nonlinear systems using the DIRECT algorithm. 2017. Preprint, <https://arxiv.org/abs/1709.01272>.

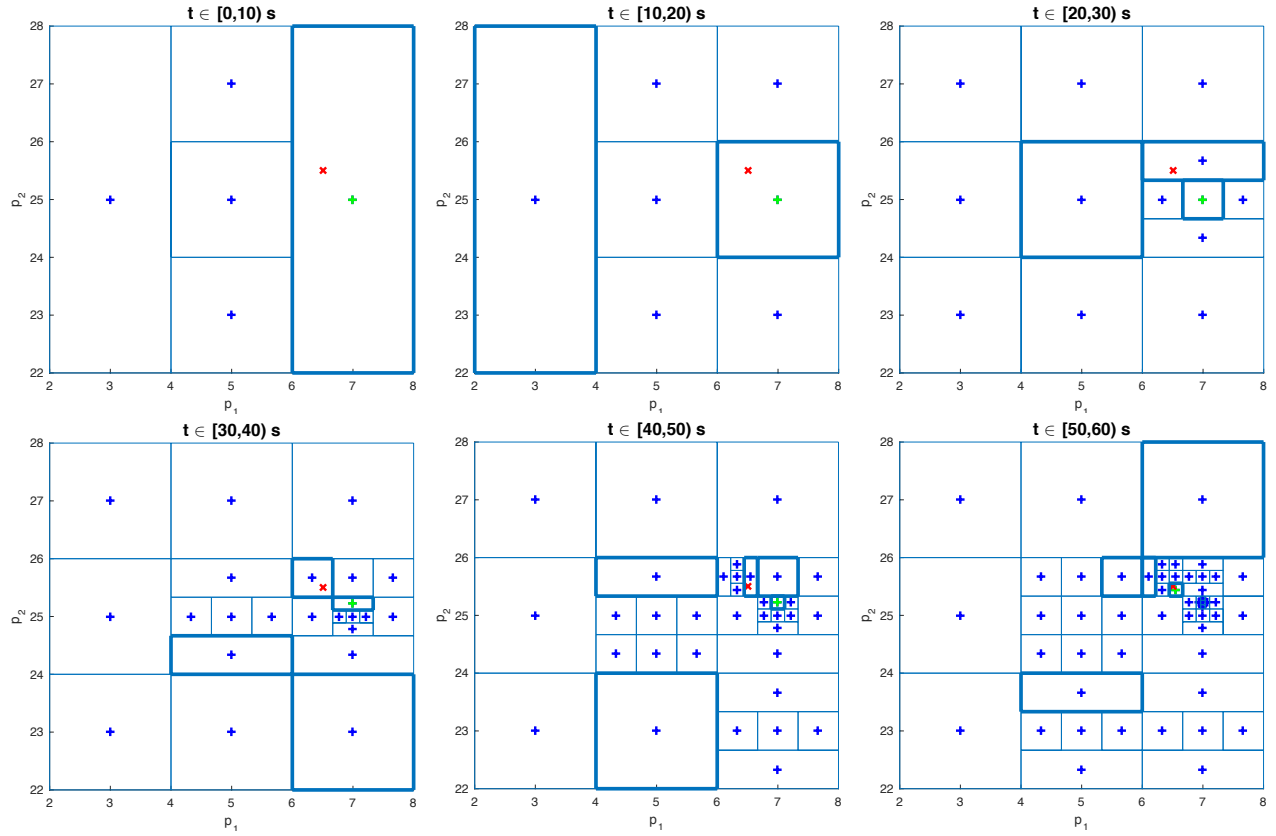


Fig. 1. Legend: \times is the true parameter p^* ; $+$ are the sampled parameters; $+$ is the final selected parameter in the interval $t \in [t_k, t_{k+1})$; the potentially optimal hyperrectangles at each iteration is highlighted with thicker lines.




Fig. 2. Parameter error $|\hat{p}(t) - p^*|$ and normalised state estimation error $\frac{|\hat{x}_{\sigma(t)}(t)|}{\max_t |x(t)| - \min |x(t)|}$.

	DIRECT	Dynamic policy in [7]
Convergence time T^* such that $\hat{p}_{\sigma(t)}(t) \leq 0.72, \forall t \geq T^*$.	45s	90s
Average number of observers	13.8	16
$\frac{\sum_{k=0}^9 N(t_k)}{\lceil t_f/T_d \rceil}$		
where $N(t_k)$ is the number of observers used at $t \in [t_k, t_{k+1})$, $k \in \{0, \dots, 9\}$.		
Parameter estimation error	0.03	1.04
Normalised state estimation error	0.0043	0.0025
$\frac{ \hat{x}_{\sigma(t_f)}(t_f) }{\max_{t \in [0, t_f]} x(t) - \min_{t \in [0, t_f]} x(t) }$		

TABLE I

NUMERICAL RESULTS COMPARING THE SUPERVISORY OBSERVER WITH DIRECT VS. THE PREVIOUS DYNAMIC POLICY IN [7] WITH $t_f = 100$ s.

- [9] D.E. Finkel. DIRECT: Research and codes, January 2004. http://www4.ncsu.edu/~ctk/Finkel_Direct/.
- [10] J.M. Gablonsky. *Modifications of the DIRECT algorithm*. PhD thesis, North Carolina State University, 2001.
- [11] B.H. Jansen and V.G. Rit. Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biological cybernetics*, 73(4):357–366, 1995.
- [12] D.R. Jones, C.D. Perttunen, and B.E. Stuckman. Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1):157–181, 1993.
- [13] S.Z. Khong, D. Nešić, C. Manzie, and Y. Tan. Multidimensional global extremum seeking via the DIRECT optimisation algorithm. *Automatica*, 49(7):1970–1978, 2013.
- [14] X. Li and Y. Bar-Shalom. Multiple-model estimation with variable structure. *IEEE Transactions on Automatic Control*, 41(4):478–493, 1996.
- [15] D. Liberzon. *Switching in systems and control*. Springer, 2003.
- [16] K.S. Narendra and Z. Han. The changing face of adaptive control: the use of multiple models. *Annual Reviews in Control*, 35(1):1–12, 2011.

- [17] S. N. Sheldon and P. S. Maybeck. An optimizing design strategy for multiple model adaptive estimation and control. *IEEE Transactions on Automatic Control*, 38(4):651–654, April 1993.