



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Input-Output Non-Linear Dynamical Systems applied to Physiological Condition Monitoring

Citation for published version:

Georgatzis, K, Williams, CKI & Hawthorne, C 2016, Input-Output Non-Linear Dynamical Systems applied to Physiological Condition Monitoring. in *Proceedings of the 1st Machine Learning for Healthcare Conference 2016*. Proceedings of Machine Learning Research, vol. 56, PMLR, Machine Learning for Healthcare 2016, Los Angeles, California, United States, 19/08/16. <<http://proceedings.mlr.press/v56/Georgatzis16.html>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Proceedings of the 1st Machine Learning for Healthcare Conference 2016

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Input-Output Non-Linear Dynamical Systems applied to Physiological Condition Monitoring

Konstantinos Georgatzis

*School of Informatics
University of Edinburgh
Edinburgh, United Kingdom*

K.GEORGATZIS@SMS.ED.AC.UK

Christopher K.I. Williams

*School of Informatics
University of Edinburgh
Edinburgh, United Kingdom
Alan Turing Institute
London, United Kingdom*

CKIW@INF.ED.AC.UK

Christopher Hawthorne

*Institute of Neurological Sciences
Queen Elizabeth University Hospital
Glasgow, United Kingdom*

CWHAWTHORNE@DOCTORS.ORG.UK

Abstract

We present a non-linear dynamical system for modelling the effect of drug infusions on the vital signs of patients admitted in Intensive Care Units (ICUs). More specifically we are interested in modelling the effect of a widely used anaesthetic drug (Propofol) on a patient's monitored depth of anaesthesia and haemodynamics. We compare our approach with one from the Pharmacokinetics/Pharmacodynamics (PK/PD) literature and show that we can provide significant improvements in performance without requiring the incorporation of expert physiological knowledge in our system.

1. INTRODUCTION

We are concerned with the problem of predicting the effect of infused drugs on the vital signs of a patient in intensive care. In the biomedical literature this problem is broken into pharmacokinetics (PK) and pharmacodynamics (PD). Pharmacokinetics concerns the absorption, distribution, metabolism, and elimination of drugs from the body, while pharmacodynamics addresses the biochemical and physiological effects of drugs on the body. PK models are typically expressed as sets of ordinary differential equations (ODEs), while PD models are typically nonlinear functions relating drug concentration to observed vital signs.

From a machine learning point of view, these models consist of a linear dynamical system with control inputs (the drug infusion rates), and a non-linear output model. The PK models are based on quite a number of assumptions that are arguably not highly accurate representations of what is going on in the body (see section 2.1 for more details). Thus our contribution is to take a more "data-driven" approach to the problem, fitting an input-output non-linear dynamical system (IO-

NLDS¹) to predict the vital signs based on input drug infusion rates. A notable difference to the PK/PD approach is that the latent process is not constrained to be mapped to any physiologically interpretable quantity and the model is free to learn any latent representation that might better explain the observed data. Our results show clear improvements in performance over the PK/PD approach.

2. MODEL DESCRIPTION

2.1 PK/PD model

We start by describing the standard PK approach which is called *compartmental modelling*. Compartmental models are an abstraction used to describe the rate of change of a drug’s concentration in a patient by accounting for the processes of absorption, distribution, metabolism and excretion of the drug in different parts (compartments) of the human body. This approach then involves building a system of ODEs that describe the evolution of drug concentrations at different compartments. The standard compartmental model for modelling the pharmacokinetical properties of the anaesthetic drug Propofol is comprised of three compartments and dates back to Gepts et al. (1987). Based on that work, a model widely used in practice is the one introduced by Marsh et al. (1991) which has been further improved upon by White et al. (2008). This line of work only addresses the PK aspect of the task. In order to quantify the effect of the drug on the observed vital signs, one needs to add an extra “effect” compartment for each observed vital sign, and link the concentration at the effect compartment with the observed physiology as done e.g. in Bailey and Haddad (2005).

A graphical representation of this overall PK/PD approach is shown in Figure 1 (left), where x_i is the concentration of drug in compartment i and k_{ij} ² is the drug’s transfer rate from compartment i to j . Thus, x_1 denotes the drug concentration in the central compartment, which is the site for drug administration and includes the intravascular blood volume and highly perfused³ organs (e.g. heart, brain). The highest fraction of the administered drug is assumed to reside in the central compartment. The remainder is diffused in two peripheral compartments which represent the body’s muscle and fat. The drug’s concentration in these two compartments is denoted as x_2 and x_3 respectively, while x_e refers to the drug concentration at the effect site. For example, if one needs to measure the effect of an anaesthetic drug on a patient’s consciousness level, then x_e would refer to the patient’s brain. Also, u denotes the drug infusion rate and k_{10} denotes the drug’s elimination rate from the central compartment. Finally, one needs to establish the functional relationship between x_e and the observed vital sign y . Since this relationship has a clear sigmoid shape, it is traditionally modelled via a generalised logistic function (also known as Richards’ curve; see Richards, 1959) of the form: $g(x_e) = m + (M - m)/(1 + e^{-\gamma x_e})^{(1/\nu)}$, where the parameters m , M govern the lower/upper asymptote respectively, γ controls the decrease rate and ν determines near to which asymptote maximum decrease occurs. This model can be described by a system of linear ODEs

1. Inspired by the IO-HMM of Bengio and Frasconi (1995).
2. Parameter k_{1e} is known as k_{e0} in the PK literature.
3. High ratio of blood flow to weight.

with the addition of a generalised logistic function as follows:

$$\begin{aligned}
dx_{1t}/dt &= -(k_{10} + k_{12} + k_{13})x_{1t} + k_{21}x_{2t} + k_{31}x_{3t} + u_t , \\
dx_{2t}/dt &= k_{12}x_{1t} - k_{21}x_{2t} , \\
dx_{3t}/dt &= k_{13}x_{1t} - k_{31}x_{3t} , \\
dx_{et}/dt &= k_{1e}(x_{1t} - x_{et}) , \\
y_t &= g(x_{et}) .
\end{aligned} \tag{1}$$

An important aspect of this model is that the parameters associated with the PK part (i.e. the set of k_{ij} s) are considered fixed and their values have been determined by Marsh et al. (1991) based on principles of human physiology. In White et al. (2008) it was established that this PK model could be improved by allowing k_{10} to vary according to a patient's age and gender. Hence, in this work, we follow the improved version of White et al. (2008).

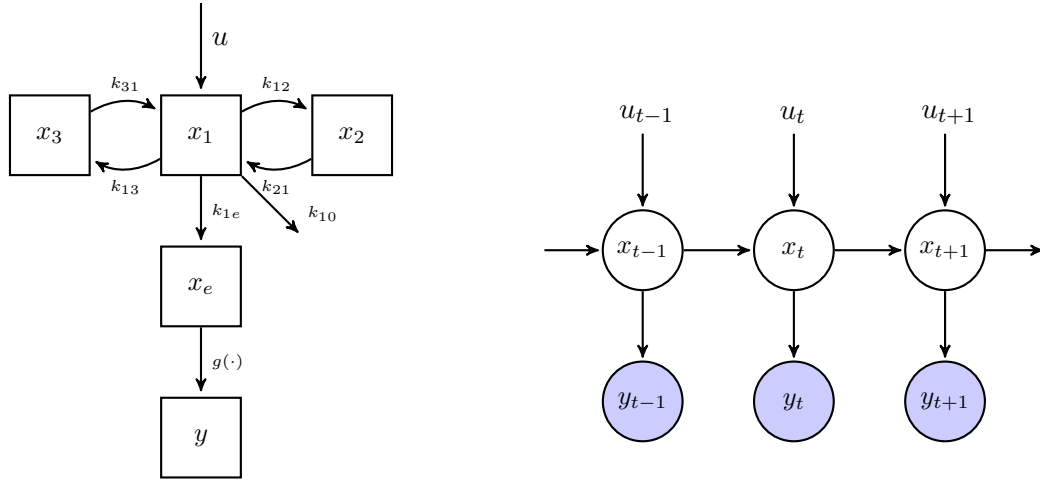


Figure 1: Left: a three compartment model for Propofol with one added effect site compartment. See text for details. Right: graphical model of IO-NLDS. The latent physiological state of a patient and the drug infusion rates at time t are denoted by \mathbf{x}_t and \mathbf{u}_t respectively. The shaded nodes correspond to the observed physiological values, \mathbf{y}_t .

2.2 IO-NLDS

In contrast to the PK/PD model we adopt an approach which makes no explicit use of expert physiological knowledge and is not restricted in associating the model's parameters with physiologically relevant quantities. Instead, the minimal assumptions made are that a latent temporal process with linear dynamics driven by control inputs (drug infusion), gives rise to observed quantities (vital signs) which are non-linearly dependent on the latent process. This gives rise to an input-output non-linear dynamical system (IO-NLDS) which is defined by the following joint distribution:

$$p(\mathbf{x}, \mathbf{y} | \mathbf{u}) = p(\mathbf{x}_1 | \mathbf{u}_1) p(\mathbf{y}_1 | \mathbf{x}_1) \prod_{t=2}^T p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_t) p(\mathbf{y}_t | \mathbf{x}_t) , \tag{2}$$

where $\mathbf{x}_t \in R^{d_x}$, $\mathbf{u}_t \in R^{d_u}$ and $\mathbf{y}_t \in R^{d_y}$ denote the latent states, control inputs and observed vital signs at time t , and T denotes the total length of the observed vital signs. We further assume that

the random variables are distributed according to:

$$\mathbf{x}_t \sim \mathcal{N}(\mathbf{A}\mathbf{x}_{t-1} + \mathbf{B}\mathbf{u}_t, \mathbf{Q}), \quad (3)$$

$$\mathbf{y}_t \sim \mathcal{N}(g(\mathbf{C}\mathbf{x}_t), \mathbf{R}). \quad (4)$$

The graphical model corresponding to these equations is shown in Figure 1 (right). Equation (4) encodes the assumption that the latent state is linearly projected onto the observation space via matrix \mathbf{C} , is subsequently non-linearly transformed via the generalised logistic function $g(\cdot)$, and corrupted by Gaussian observation noise with covariance \mathbf{R} to give rise to the observations. Function $g(\cdot)$ is parameterised as in the PK/PD approach. The latent state itself is following linear dynamics, governed by matrix \mathbf{A} and a linear transformation (via matrix \mathbf{B}) of control inputs, and additive Gaussian noise with covariance \mathbf{Q} . A notable difference to the PK/PD approach is that the latent process is not constrained to be mapped to any physiologically interpretable quantity and the model is free to learn any latent representation that might better explain the observed data. In contrast to the PK/PD model, the latent process exhibits a higher degree of flexibility, being unconstrained of any (simplifying) physiologically motivated assumptions and can thus be expected to model the patient-specific underlying dynamics in a more expressive way.

2.3 PK/PD model as NLDS

The PK/PD model as described in section 2.1 does not incorporate any uncertainty while the IO-NLDS is a probabilistic model. In order to compare the two models, we cast the PK/PD model into the IO-NLDS form as described in eqs. (3), (4). This involves two steps: a) the discretisation of the continuous time dynamics as described by the system of ODEs in the first four equations of eq. (1) and b) the addition of Gaussian noise on the discretised dynamics and on the non-linear output. We can do this by setting the following parameters (assuming here $d_y = 1$ for compactness) as $\mathbf{A} = \exp\{\mathbf{F}\Delta t\}$, $\mathbf{B} = [1 \ 0 \ 0 \ 0]^\top$, $\mathbf{C} = [0 \ 0 \ 0 \ 1]$, where \mathbf{F} is described in the supplementary material. The dynamics matrix \mathbf{A} is the discretised version of its continuous time counterpart \mathbf{F} , which involves computing the matrix exponential of \mathbf{F} times the discretisation step Δt as described e.g. in Aström and Murray (2010, sec. 5.3). We note that the PK model provided by White et al. (2008) provides estimates for the parameters involved in \mathbf{F} , except for k_{1e} . To fit an appropriate k_{1e} we perform a fine-grained one-dimensional grid search per observed channel around a clinically relevant value. The noise matrices \mathbf{Q} and \mathbf{R} have the same interpretation as in the case of the IO-NLDS and can be learned in the same way, as described in section 2.5. Under this form, the PK/PD model can be seen as an IO-NLDS with a constrained parametric form, where the parameters \mathbf{A} , \mathbf{B} , and \mathbf{C} are constrained in such a way as to capture the physiological processes involved with the infusion of an anaesthetic drug, as already described in section 2.1.

2.4 Inference

If the observation model was linear, then the overall model would be a linear dynamical system (LDS) and exact filtering would be feasible via the well-known Kalman filter (Kalman, 1960). With the addition of the nonlinear observation function $g(\cdot)$, exact inference becomes intractable and one must resort to some form of approximation. We adopt a sigma-point filtering approach, and more specifically we use the unscented Kalman filter (UKF), as described in Särkkä (2013). The main component of UKF is the unscented transform (UT) as described in Julier and Uhlmann

(1996) (although the term “unscented” was introduced later). The idea behind it is that a non-linear transformation of a Gaussian distribution can be approximated by first deterministically selecting a fixed number of points (called sigma-points) from that distribution to capture its mean and covariance, then computing the exact non-linear transformation of these points, and subsequently fitting a Gaussian distribution to the non-linearly transformed points. Under this procedure, the UKF reduces to applying the UT twice at each time step t : a) once to compute the predictive density $p(\mathbf{x}_t|\mathbf{y}_{1:t-1}, \mathbf{u}_{1:t})$, by using the UT on the previous filtered density $p(\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}, \mathbf{u}_{1:t-1})$ and b) to compute the current filtered density $p(\mathbf{x}_t|\mathbf{y}_{1:t}, \mathbf{u}_{1:t})$ by using the UT on the predictive density to compute the current likelihood $p(\mathbf{y}_t|\mathbf{x}_t)$ and then using Bayes rule to obtain the required filtered (posterior) density. Since our model is governed by linear dynamics, the first step can be calculated by the standard Kalman filter equivalent step with both methods yielding the same results.

2.5 Learning

The parameters of the proposed model that need to be learned are $\boldsymbol{\theta} = \{\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1, \mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{Q}, \mathbf{R}, \boldsymbol{\eta}\}$, where we assume that $\mathbf{x}_1 \sim \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ and $\boldsymbol{\eta} = \{\mathbf{m}, \mathbf{M}, \boldsymbol{\gamma}, \boldsymbol{\nu}\}$ are the parameters of $g(\cdot)$. We learn maximum likelihood (ML) estimates of these parameters by using the expectation maximisation (EM) algorithm (Dempster et al., 1977). EM is an iterative algorithm that maximises the likelihood via maximising the following surrogate function:

$$\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{old}) = \mathbb{E}_{p(\mathbf{x}|\mathbf{y}, \mathbf{u}, \boldsymbol{\theta}^{old})}[\log p(\mathbf{x}, \mathbf{y}|\mathbf{u}, \boldsymbol{\theta})], \quad (5)$$

which, due to the Markov properties of the model, can be further decomposed as:

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{old}) &= \mathbb{E}_{p(\mathbf{x}_1|\mathbf{y}_{1:T}, \mathbf{u}_{1:T}, \boldsymbol{\theta}^{old})}[\log p(\mathbf{x}_1|\mathbf{u}_1, \boldsymbol{\theta})] + \mathbb{E}_{p(\mathbf{x}_t, \mathbf{x}_{t-1}|\mathbf{y}_{1:T}, \mathbf{u}_{1:T}, \boldsymbol{\theta}^{old})}[\log p(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{u}_t, \boldsymbol{\theta})] \\ &+ \mathbb{E}_{p(\mathbf{x}_t|\mathbf{y}_{1:T}, \mathbf{u}_{1:T}, \boldsymbol{\theta}^{old})}[\log p(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta})]. \end{aligned} \quad (6)$$

These terms involve expectations with respect to the smoothing distributions $p(\mathbf{x}_t|\mathbf{y}_{1:T}, \mathbf{u}_{1:T}, \boldsymbol{\theta}^{old})$ and pairwise joint smoothing distributions $p(\mathbf{x}_t, \mathbf{x}_{t-1}|\mathbf{y}_{1:T}, \mathbf{u}_{1:T}, \boldsymbol{\theta}^{old})$. These distributions can be computed in general via the unscented Rauch-Tung-Striebel (URTS) smoother as described in Särkkä (2008). In our case, since the system’s dynamics are linear, the backward smoothing step can be performed by the standard RTS smoother after obtaining the unscented filtered estimates during the forward filtering step via the UKF. Analogously to the UKF, where the UT is used to approximate the required expectations, the UT needs to be employed here to approximate the expectation appearing in the last term of the RHS of eq. (6). More details are given in Kokkala et al. (2014).

Computing the required distributions and expectations in order to calculate eq. (6) constitutes the E-step of EM. In the M-step the model parameters are set such that $\boldsymbol{\theta}^* \leftarrow \arg \max_{\boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{old})$. This maximisation step can be done analytically in the case of the LDS, but one needs to make use of numerical optimisation methods in the non-linear case. In our model, the subset of parameters $\boldsymbol{\theta}_L^* = \{\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1, \mathbf{A}, \mathbf{B}, \mathbf{Q}\}$ which correspond to the linear part of the model can be computed in closed form as shown in Cheng and Sables (2006), and \mathbf{R}^* can be computed similarly to the linear case as shown in Särkkä (2013, sec. 12.3.3). We provide those estimates in the supplementary material. The remaining set of parameters $\boldsymbol{\theta}_{NL}^* = \{\mathbf{C}, \boldsymbol{\eta}\}$ can be then optimised via numerical optimisation. In our experiments we use the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (see e.g. Fletcher, 2013). We note that EM is a natural choice in our case since our model involves a linear sub-component which can be exploited in the decomposition of eq. (6) to derive a subset

of parameters in closed form. In a fully non-linear case however one could use a numerical optimisation procedure to directly maximise the likelihood function instead of a surrogate function, as argued in Kokkala et al. (2015). To the best of our knowledge, this is the first time that the UKF and URTS are used within EM for ML estimation of parameters that include also control inputs.

2.6 Related work

The PK/PD model presented in section 2.1 is a standard approach in the PK/PD literature that stems from a long line of research. Starting with the PK component, in Benet (1972), the Laplace transform method is presented as a way of analysing multi-compartmental PK models. Work conducted by Gepts et al. (1987) is focussed on the PK properties of the anaesthetic drug Propofol under a continuous drug infusion, as opposed to bolus injections which used to be standard practice. In that work, it was concluded that a three-compartmental PK model best fits the observed data, and physiologically relevant parameters for this model were established. In Marsh et al. (1991), the well-known (in PK literature) *Marsh* model was presented, which investigated the previous model and suggested a new set of parameters that better reflected the distribution and elimination of Propofol so as to include children. In White et al. (2008), it was shown that the *Marsh* model should be modified to allow for the elimination rate (k_{10}) of Propofol from the central compartment to vary depending on a patient's age and gender.

The PD component has a similarly long history. In Colburn (1981), one, two and three compartmental PK models with the addition of an effect site compartment to model PD effects were investigated. Similarly, in Fuseau and Sheiner (1984), a three compartmental PK model with an additional effect site compartment was studied. An approach which is closer to our proposed PK/PD model is presented in Bailey and Haddad (2005), where a three-compartmental PK model with an effect site compartment is cast as a state space model and a sigmoid function links the drug concentration at the effect site to the patient's observed level of consciousness. However, that model is non-stochastic, and is evaluated on simulated data with parameters fixed on a priori known values.

In terms of research which is closer to the machine learning community, there has been work which tackles various tasks of interest with respect to physiological monitoring in ICUs. In Quinn et al. (2009), a factorial switching linear dynamical system (FSLDS) was used to infer artifactual and physiological processes of interest, whereas in Georgatzis and Williams (2015) a discriminative SLDS is used for the same purposes. In Lehman et al. (2015), a switching vector autoregressive process was used to extract features from vital signs which were used as input in a logistic regression classifier to predict patient outcome, while work by Nemati et al. (2013) was focussed on discriminatively training a SLDS for learning dynamics associated with patient outcome. These methods are primarily concerned with the classification of events of interest ignoring any administered drugs. Additionally, all of these works use linear models as their core elements, while our model is non-linear. More similar work to ours is presented by Enright et al. (2011) and Enright et al. (2013), where a non-linear dynamical system is developed to model the glucose levels of patients under the intravenous administration of glucose and insulin. However, their approach relies solely on converting existing systems of ODEs into a probabilistic model, which in our case corresponds only to the PK/PD model and is clearly outperformed by the fully data-driven approach.

In terms of methodology, the task of inference in non-linear dynamical models has been thoroughly explored and methods such as the extended Kalman filter (see e.g. Särkkä, 2013, sec. 5.2), the UKF (Wan and Van Der Merwe, 2000) and the particle filter (Gordon et al., 1993) have been

proposed. The UKF is shown empirically to outperform the extended KF (see e.g. Haykin, 2001, Ch. 7), and was chosen over the particle filter (PF) because the PF can require “orders of magnitude” (see Wan and Van Der Merwe, 2000) more sample points compared to the UKF’s fixed, small number of sigma-points to achieve high accuracy, rendering it computationally prohibitive for a real-time application such as ours. Learning is not as thoroughly explored as inference in NLDSs but nonetheless there has been considerable work recently. In Schön et al. (2011) the EM algorithm is used in conjunction with PF for parameter estimation in NLDSs, while in Gašperin and Juričić (2011) and Kokkala et al. (2014) the EM is used in conjunction with the UKF for the same task. However, they do not include control inputs in their formulation as we do.

3. EXPERIMENTS

In this section we describe experiments that were conducted in order to establish if the newly proposed approach can model accurately the effect of drug infusions on the observed physiology of patients in ICUs. To this end, we compare the IO-NLDS and the PK/PD model for the task of predicting the effect of Propofol on patients’ vital signs.

3.1 Data Description

The dataset comprises of 40 Caucasian patients admitted in the neuro ICU of the Golden Jubilee National Hospital in Glasgow, Scotland. All patients were spontaneously breathing with the maximum airway intervention being an oropharyngeal airway, and none was classified higher than Class 2 with respect to illness severity according to the American Society of Anesthesiologists. Two controlled Propofol infusion protocols were investigated on these patients pre-operatively during a period of approximately 45 minutes. This investigation was part of an independent clinical study and the data have been anonymised. Each patient was randomly assigned to one of the two protocols. The first protocol involved a target Propofol concentration of $2 \mu\text{g/ml}$ for the first 15 minutes followed by a target concentration of $5 \mu\text{g/ml}$ for the next 15 minutes and $2 \mu\text{g/ml}$ for the last 15 minutes. The second protocol was the inverse of the first with a 5-2-5 $\mu\text{g/ml}$ target sequence. The drug pumps which automatically administer Propofol calculate internally the desired infusion rates and these rates were used as control inputs in our model. Propofol doses are at peak during the early phase of the protocol making it very likely that the maximum haemodynamic effects will be seen during the study period and not after discontinuation of the drug. Also, during the study very few interventions were required in the case of hypotension (treated by small incremental doses of either ephedrine or Metaraminol) and bradycardia (treated by Glycopyrolate). The frequency and duration of these interventions were deemed sufficiently low so as to treat the overall dataset as unaffected by them. During the duration of the protocol, six vital signs were recorded, namely systolic, mean and diastolic blood pressure (BP_{sys}, BP_{mean}, BP_{dia}), heart rate (HR), respiratory rate (RR) and the bispectral index (BIS). BIS was recorded only in 27 of the 40 patients. Of those channels, the first five were measured at a time interval $\Delta t = 15$ seconds and BIS was measured every 5 seconds and was subsequently subsampled to 15 seconds. BIS is a standard index (see e.g. Rampil, 1998) that is based on features extracted from the spectrum of a patient’s electroencephalogram (EEG) and measures depth of anaesthesia, with values ranging between 0 – 100 (lower values representing deeper anaesthesia). From these values, BP_{sys}, BP_{dia} and BIS were expected to be affected by Propofol administration and the signals were of adequate quality so that they could be further analysed.

3.2 Model fitting

For both models we use EM as described in section 2.5 to learn the parameters on each patient separately and then obtain predicted values using those fitted parameters. For both models, 100 iterations of EM were used and the BFGS algorithm was run with 1000 function evaluations during the first 10 iterations and 100 evaluations during the remaining iterations. An important advantage of our model’s linear dynamics is that we can enforce stability constraints which are of paramount importance in a real-world setting. If the systems dynamics matrix, \mathbf{A} , becomes unstable (i.e if the modulus of its largest eigenvalue is greater than one) we project this matrix back to the space of stable matrices such that it is also closer (in a least-squares sense) to the originally learned matrix. We follow the approach proposed by Siddiqi et al. (2007) to achieve this, which involves solving for a quadratic program inside EM. This process is very fast and is performed only if an instability is observed. We also constrain the noise covariance matrices \mathbf{Q} and \mathbf{R} to be diagonal. Furthermore, we set $d_x = 4$ for the IO-NLDS. We fix it to this value because it directly mirrors the four compartments of the PK/PD model, and thus allows for a more direct and fair comparison between the two models. Under these assumptions, the IO-NLDS has 65 parameters compared to 50 in the case of the PK/PD model. The higher number of parameters in the case of the IO-NLDS reflects its greater expressive power since it is not constrained by physiological assumptions. We note, however, that in the general case the IO-NLDS’ latent space dimensionality does not have to be restricted in such a way. In contrast with the PK/PD model, one has the flexibility to decide on alternative latent space dimensionalities; e.g. by making use of an information criterion to determine d_x under a more formal measure of optimality, turning the whole process into a standard model order selection procedure. In the PK/PD case, a higher latent space dimension would correspond to an increased number of compartments, which could require years of additional research to validate physiologically appropriate extensions of the already existing literature. However, to perform a more fair model comparison between the two models, we also compute Bayesian Information Criterion (BIC) scores to account for the increased number of parameters in the case of the IO-NLDS. Finally, the UT involves determining three parameters α, β, κ (see supplementary material for details). Following Särkkä (2013, sec. 5.5), we set $\alpha = 1, \beta = 0$, and $\kappa = 3 - d_x$.

3.3 Results

We use the standardised mean squared error (SMSE)⁴ between predictions and actual observations as our evaluation metric, and also provide representative examples of curve fits on the observed data produced by the two models. The SMSE takes into account the variance of the observed data and provides a natural baseline with $SMSE = 1$ denoting the mean prediction. For both models we compare the SMSE between model predictions and measured outputs of BPsys, BPdia and BIS across a number of different prediction horizons; namely $h = 1, 10$ and 20-step ahead predictions, (corresponding to 15 seconds, 2.5 minutes and 5 minutes respectively). We also evaluate the SMSE for the “free-running” case, which corresponds to predictions for the whole duration of the protocol. These time intervals, apart from the 1-step interval, were decided as clinically relevant. We include results on the 1-step ahead prediction task, since this is a standard evaluation task when assessing predictive performance of dynamical models.

In Figure 2, three representative examples of observed vital signs and the fitted traces produced by the two models are presented. In the left panel, an example on the 10-step ahead prediction task

4. $SMSE = MSE/var_y$, where MSE is the mean squared error and var_y is the variance of the observed vital sign.

for the BPdia channel is given. Both models manage to capture the observed temporal structure with satisfactory accuracy. In the middle panel, an example of the 20-step ahead task on BIS is given. Here the difference in performance is clearer, with the IO-NLDS tracking the temporal evolution of the observed signal more accurately, especially during the time that BIS decreases more rapidly which corresponds to a rapid increase in the patient’s depth of anaesthesia. Finally, in the right panel, the “free-running” predictions of the two models are shown along with the observed BPsys channel. The IO-NLDS manages to stay very close to the actual signal, while on the other hand the PK/PD model showcases a considerably higher degree of error in its predictions, as it has not managed to learn an appropriate rate of decrease on its predictions. In all cases, the IO-NLDS manages to stay closer to the observed signal especially at the beginning of the protocol during the steepest decrease of the observed signal. This constitutes the most critical phase of the protocol during which the uncertainty of an anaesthetist about the patient’s reaction to the drug is considerable and the risk of an undesirable episode (e.g. hypotension) is at its highest. Therefore, accurate predictions during that phase are much more critical compared to the rest of the signal.

The results for the three channels and the four prediction horizons are shown as boxplots in Figure 3, where the central mark denotes the median, the edges of the box are the lower and upper quartiles and the whiskers extend to outliers with extreme outliers being denoted separately by a red cross. For all three channels, the IO-NLDS’ predictions are consistently more accurate than the PK/PD model’s across all four prediction horizons. Moreover, the difference in favour of the IO-NLDS becomes more obvious as the prediction horizon increases. Also, in almost all cases the predictive errors for both models are increased as the prediction horizon increases, as expected.

A summary of the mean SMSEs of the two models per prediction horizon and per channel is shown in Table 1. The IO-NLDS’ prediction errors are consistently lower than the PK/PD model’s and by a large margin in most cases. Both models’ errors are low at the 1-step ahead prediction level. In the same table, we also provide BIC scores per prediction horizon for both models to account for IO-NLDS’ higher number of parameters. The BIC is defined as: $BIC = -2 \ln(L) + b \ln(N)$, where L is the data likelihood under the model, b is the number of free parameters and N is the number of data points. The IO-NLDS achieves a lower (better) score in all cases and thus demonstrates that the increase in likelihood is not just an effect of higher statistical capacity but rather that the IO-NLDS manages to model the observed temporal structure better than the PK/PD model.

A more fine-grained comparison shows that in 92% of all investigated cases the accuracy of the IO-NLDS predictions is better than the PK/PD model’s (see Figure S1 in the supplementary material). Finally, we perform twelve paired, right-tailed t-tests (one per channel and per prediction horizon) on the differences of the obtained SMSEs ($SMSE_{PKPD} - SMSE_{NLDS}$) across the 40 patients. The null hypothesis that the mean of this difference is zero, is rejected in 10 out of 12 cases with p -values ranging from 1.6×10^{-2} to 4.8×10^{-17} . In the case of the 1-step ahead predictions on BPdia and BIS, p -values of 0.14 and 0.37 respectively were obtained and thus the null hypothesis could not be rejected.

4. DISCUSSION

We first describe a use case scenario for the IO-NLDS and then provide suggestions for future work.

The current practice in an ICU is that an anaesthetist will start drug administration on a patient with an initial target dosage and then readjust it according to the evolution of the patient’s vital signs. This implies that the anaesthetist needs to be at bedside continuously and for the whole duration of the drug administration, and by definition can only be reactive to potential vital sign deteriorations.

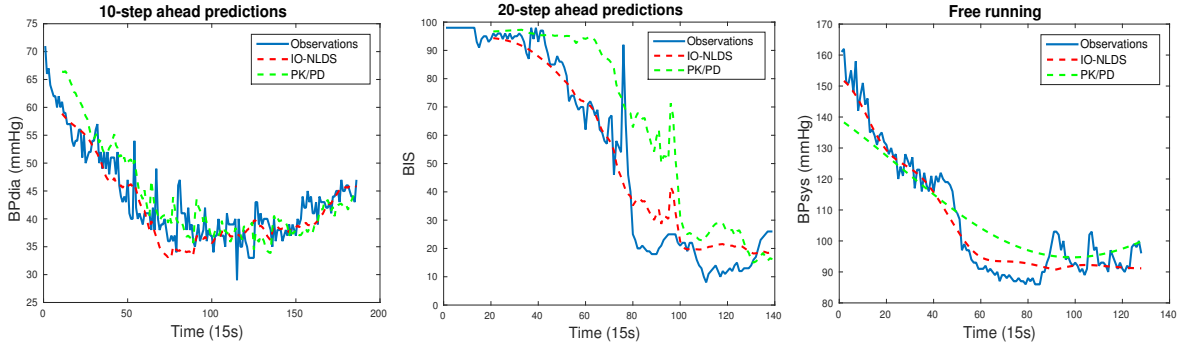


Figure 2: Examples of observed vital sign traces and the predictions made by the IO-NLDS (red dashed line) and the PK/PD model (green dashed line) for different prediction horizons.

Table 1: Comparison of IO-NLDS and PK/PD model with respect to mean SMSE per channel and prediction horizon and mean BIC per prediction horizon. Lower numbers are better.

SMSE/BIC	1-step				10-step				20-step				free-running			
	BP _s	BP _d	BIS	BIC	BP _s	BP _d	BIS	BIC	BP _s	BP _d	BIS	BIC	BP _s	BP _d	BIS	BIC
IO-NLDS	0.13	0.20	0.08	2592	0.19	0.27	0.21	4458	0.20	0.27	0.25	6656	0.25	0.36	0.33	16588
PK/PD	0.30	0.23	0.10	2709	0.52	0.53	0.41	5314	0.73	0.75	0.74	7919	0.84	0.82	0.89	19008

Also, this monitoring is not always straightforward. For example, during drug induction when the risk of hypotension is high, the anaesthetist’s attention is divided as they will be also managing the airway (e.g. by inserting an endotracheal tube). Therefore a use case for the system is to take a clinician-specified drug infusion protocol, and make predictions for the time-course of the vital signs. The clinician could then inspect these forecasts, and either continue with the original protocol or adjust it accordingly. (In that case the model could also automatically notify the clinician, e.g. if a threshold is crossed by the predicted values). This extra functionality would help the anaesthetist to be proactive and to avert undesirable episodes for the patient (e.g. an episode of hypotension). The PK/PD literature has become very focussed on compartmental models, but more data-driven models may give rise to better predictions which could thus translate to better clinical outcomes.

We have presented a purely data-driven approach for the important application of modelling drug effects on the physiology of patients in ICUs. We show that our new approach outperforms the previous expert-knowledge based one. To the best of our knowledge, this is the first time that a complex, real-world model has been fully learned via a combination of unscented filters and EM, making this a promising direction for related tasks. As our model is drug-agnostic, we plan to apply it to other drugs in future work. Finally, in our experiments a separate NLDS model was fitted per patient. However, the model would be more powerful if its parameters could be predicted from patient covariates (e.g. age, gender etc.) and this is the target of our future work.

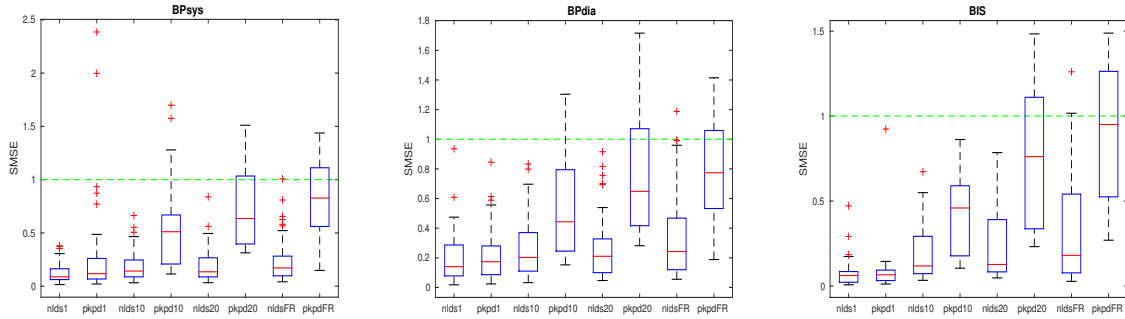


Figure 3: Summarised SMSE results for IO-NLDS compared to PK/PD for BPsys, BPdia and BIS across the four prediction horizons for all patients. The green dashed line corresponds to the mean predictions. The x-axis is labelled according to the model name followed by the prediction horizon.

ACKNOWLEDGEMENTS

We extend our thanks to Martin Shaw for preprocessing code and valuable discussions. Author KG was funded by the Scottish Informatics and Computer Science Alliance. The work of CW is supported in part by EPSRC grant EP/N510129/1.

References

- Karl Johan Aström and Richard M Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, 2010.
- James M Bailey and Wassim M Haddad. Drug Dosing Control in Clinical Pharmacology. *Control Systems, IEEE*, 25(2):35–51, 2005.
- Leslie Z Benet. General Treatment of Linear Mammillary Models with Elimination from Any Compartment as Used in Pharmacokinetics. *Journal of Pharmaceutical Sciences*, 61(4):536–541, 1972.
- Yoshua Bengio and Paolo Frasconi. An Input Output HMM Architecture. In G. Tesauro, D. S. Touretzky, and T. K. Leen, editors, *Advances in Neural Information Processing Systems 7*, pages 427–434. MIT Press, 1995.
- Sen Cheng and Philip N Sabes. Modeling Sensorimotor Learning with Linear Dynamical systems. *Neural Computation*, 18(4):760–793, 2006.
- Wayne A Colburn. Simultaneous Pharmacokinetic and Pharmacodynamic Modeling. *Journal of Pharmacokinetics and Biopharmaceutics*, 9(3):367–388, 1981.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.
- Catherine G Enright, Michael G Madden, Niall Madden, and John G Laffey. Clinical time series data analysis using mathematical models and dbns. In *Artificial Intelligence in Medicine*, pages 159–168. Springer, 2011.

- Catherine G Enright, Michael G Madden, and Niall Madden. Bayesian networks for mathematical models: techniques for automatic construction and efficient inference. *International Journal of Approximate Reasoning*, 54(2):323–342, 2013.
- Roger Fletcher. *Practical Methods of Optimization*. John Wiley & Sons, 2013.
- Eliane Fuseau and Lewis B Sheiner. Simultaneous modeling of pharmacokinetics and pharmacodynamics with a nonparametric pharmacodynamic model. *Clinical Pharmacology & Therapeutics*, 35(6):733–741, 1984.
- Matej Gašperin and Dani Juričić. Application of Unscented Transformation in Nonlinear System Identification. In *Proceedings of the 18th IFAC World Congress*, volume 18, pages 4428–4433, 2011.
- Konstantinos Georgatzis and Christopher K I Williams. Discriminative Switching Linear Dynamical Systems applied to Physiological Condition Monitoring. In *Proceedings of the Thirty-first Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 306–315, 2015.
- E Gepts, F Camu, ID Cockshott, and EJ Douglas. Disposition of Propofol Administered as Constant Rate Intravenous Infusions in Humans. *Anesthesia & Analgesia*, 66(12):1256–1263, 1987.
- Neil J Gordon, David J Salmond, and Adrian FM Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *Radar and Signal Processing, IEE Proceedings of*, volume 140, pages 107–113. IET, 1993.
- Simon S Haykin. *Kalman Filtering and Neural Networks*. Wiley Online Library, 2001.
- Simon J Julier and Jeffrey K Uhlmann. A General Method for Approximating Nonlinear Transformations of Probability Distributions. Technical report, Robotics Research Group, Department of Engineering Science, University of Oxford, 1996.
- Rudolph Emil Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.
- Juho Kokkala, Arno Solin, and Simo Särkkä. Expectation Maximization Based Parameter Estimation by Sigma-Point and Particle Smoothing. In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–8. IEEE, 2014.
- Juho Kokkala, Arno Solin, and Simo Sarkka. Sigma-Point Filtering and Smoothing Based Parameter Estimation in Nonlinear Dynamic Systems. *arXiv:stat.ME/1504.06173v2*, 2015.
- Li-wei H Lehman, Ryan P Adams, Louis Mayaud, George B Moody, Atul Malhotra, Roger G Mark, and Shamim Nemati. A Physiological Time Series Dynamics-Based Approach to Patient Monitoring and Outcome Prediction. *Biomedical and Health Informatics, IEEE Journal of*, 19(3):1068–1076, 2015.
- B Marsh, M White, N Morton, and GNC Kenny. Pharmacokinetic Model Driven Infusion of Propofol in Children. *British Journal of Anaesthesia*, 67(1):41–48, 1991.
- Kevin P Murphy. *Machine Learning: A Probabilistic Perspective*. MIT Press, 2012.

- Shamim Nemati, L-WH Lehman, and Ryan P Adams. Learning outcome-discriminative dynamics in multivariate physiological cohort time series. In *Engineering in Medicine and Biology Society (EMBS), 2013 35th Annual International Conference of the IEEE*, pages 7104–7107. IEEE, 2013.
- John A Quinn, Christopher K I Williams, and Neil McIntosh. Factorial Switching Linear Dynamical Systems applied to Physiological Condition Monitoring. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 31(9):1537–1551, 2009.
- Ira J Rampil. A Primer for EEG Signal Processing in Anesthesia. *The Journal of the American Society of Anesthesiologists*, 89(4):980–1002, 1998.
- FJ Richards. A Flexible Growth Function for Empirical Use. *Journal of Experimental Botany*, 10(2):290–301, 1959.
- Simo Särkkä. Unscented Rauch–Tung–Striebel Smoother. *Automatic Control, IEEE Transactions on*, 53(3):845–849, 2008.
- Simo Särkkä. *Bayesian Filtering and Smoothing*. Cambridge University Press, 2013.
- Thomas B Schön, Adrian Wills, and Brett Ninness. System identification of nonlinear state-space models. *Automatica*, 47(1):39–49, 2011.
- Sajid M Siddiqi, Byron Boots, and Geoffrey J Gordon. A constraint generation approach to learning stable linear dynamical systems. In *Advances in Neural Information Processing Systems*, pages 1329–1336, 2007.
- Eric A Wan and Ronell Van Der Merwe. The Unscented Kalman Filter for Nonlinear Estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, pages 153–158. Ieee, 2000.
- Martin White, Gavin NC Kenny, and Stefan Schraag. Use of Target Controlled Infusion to Derive Age and Gender Covariates for Propofol Clearance. *Clinical Pharmacokinetics*, 47(2):119–127, 2008.

SUPPLEMENTARY MATERIAL

We provide details about the PK/PD model’s continuous time dynamics matrix \mathbf{F} , the UT and the EM algorithm.

S.1 PK/PD continuous time dynamics matrix

$$\mathbf{F} = \begin{bmatrix} -(k_{10} + k_{12} + k_{13}) & k_{21} & k_{31} & 0 \\ k_{12} & -k_{21} & 0 & 0 \\ k_{13} & 0 & -k_{31} & 0 \\ k_{1e} & 0 & 0 & -k_{1e} \end{bmatrix}.$$

S.2 Unscented transform

The UKF and EM involve the calculation of integrals of the form $\mathbb{E}_{p(\mathbf{x})}[g(\mathbf{x})] = \int g(\mathbf{x})\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})d\mathbf{x}$. Since $g(\cdot)$ is non-linear, these integrals are approximated via Gaussian cubature (i.e. multidimensional Gaussian quadrature) as: $\mathbb{E}_{p(\mathbf{x})}[g(\mathbf{x})] \approx \sum_i w_i g(\mathbf{x}_i)$. Following similar notation to Murphy (2012, sec. 18.5.2.1), we first define a set of $2d + 1$ sigma points \mathbf{x}_i :

$$\mathbf{x} = \{\boldsymbol{\mu}, \{\boldsymbol{\mu} + (\sqrt{(d + \lambda)\boldsymbol{\Sigma}})_i\}, \{\boldsymbol{\mu} - (\sqrt{(d + \lambda)\boldsymbol{\Sigma}})_i\}\},$$

where $\boldsymbol{\Sigma}_i$ denotes the i 'th column of matrix $\boldsymbol{\Sigma}$ and $i = 1, \dots, d$. This set of sigma points is propagated via $g(\cdot)$, producing a new set of transformed points \mathbf{y}_i with mean and covariance:

$$\begin{aligned} \boldsymbol{\mu}_y &= \sum_{i=0}^{2d} w_m^i \mathbf{y}_i, \\ \boldsymbol{\Sigma}_y &= \sum_{i=0}^{2d} w_c^i (\mathbf{y}_i - \boldsymbol{\mu}_y)(\mathbf{y}_i - \boldsymbol{\mu}_y)^\top, \end{aligned}$$

with weights w 's defined as:

$$\begin{aligned} w_m^0 &= \frac{d}{d + \lambda}, \\ w_c^0 &= \frac{d}{d + \lambda} + (1 - \alpha^2 + \beta), \\ w_m^i, w_c^i &= \frac{1}{2(d + \lambda)}, \end{aligned}$$

where $\lambda = \alpha^2(d + \kappa) - d$ and α, β , and κ are method-specific parameters.

S.3 EM

We provide ML estimates for the parameters involved in eq. (6) for the i 'th iteration of EM that can be derived in closed form. We first define the following:

$$\begin{aligned}
S_{x-x-} &= \sum_{t=2}^T \mathbf{P}_{t-1|T} + \boldsymbol{\mu}_{t-1|T} \boldsymbol{\mu}_{t-1|T}^\top, \\
S_{xx-} &= \sum_{t=2}^T \mathbf{P}_{t,t-1|T} + \boldsymbol{\mu}_{t|T} \boldsymbol{\mu}_{t-1|T}^\top, \\
S_{xx} &= \sum_{t=2}^T \mathbf{P}_{t|T} + \boldsymbol{\mu}_{t|T} \boldsymbol{\mu}_{t|T}^\top, \\
S_{x-u} &= \sum_{t=2}^T \boldsymbol{\mu}_{t-1|T} \mathbf{u}_t^\top, \\
S_{xu} &= \sum_{t=2}^T \boldsymbol{\mu}_{t|T} \mathbf{u}_t^\top, \\
S_{uu} &= \sum_{t=2}^T \mathbf{u}_t \mathbf{u}_t^\top,
\end{aligned}$$

where $\boldsymbol{\mu}_{t|T}$, $\mathbf{P}_{t|T}$ denote the smoothed mean and covariance estimates at time t and $\mathbf{P}_{t,t-1|T}$ denotes the smoothed pairwise covariance estimates at times $t-1, t$, as outputted by the URTS smoother. We then have:

$$\begin{aligned}
\boldsymbol{\mu}_1^i &= \boldsymbol{\mu}_{1|T}, \\
\boldsymbol{\Sigma}_1^i &= \mathbf{P}_{1|T}, \\
[\mathbf{A}^i \ \mathbf{B}^i] &= [S_{xx-} \ S_{xu}] \begin{bmatrix} S_{x-x-} & S_{x-u} \\ S_{x-u}^\top & S_{uu} \end{bmatrix}^{-1}, \\
\mathbf{Q}^i &= \frac{1}{T-1} (S_{xx} - \mathbf{A}^i S_{xx-}^\top - \mathbf{B}^i S_{xu}^\top).
\end{aligned}$$

Finally, we have:

$$\mathbf{R}^i = \frac{1}{T} \sum_{t=2}^T \mathbb{E}_{p(\mathbf{x}_t | \mathbf{y}_{1:T}, \mathbf{u}_{1:T}, \boldsymbol{\theta}^{i-1})} [(\mathbf{y}_t - g(\mathbf{x}_t, \boldsymbol{\theta}_{NL}))(\mathbf{y}_t - g(\mathbf{x}_t, \boldsymbol{\theta}_{NL}))^\top],$$

where the expectation is another Gaussian integral that can be computed via the unscented approximation as described in section 4.

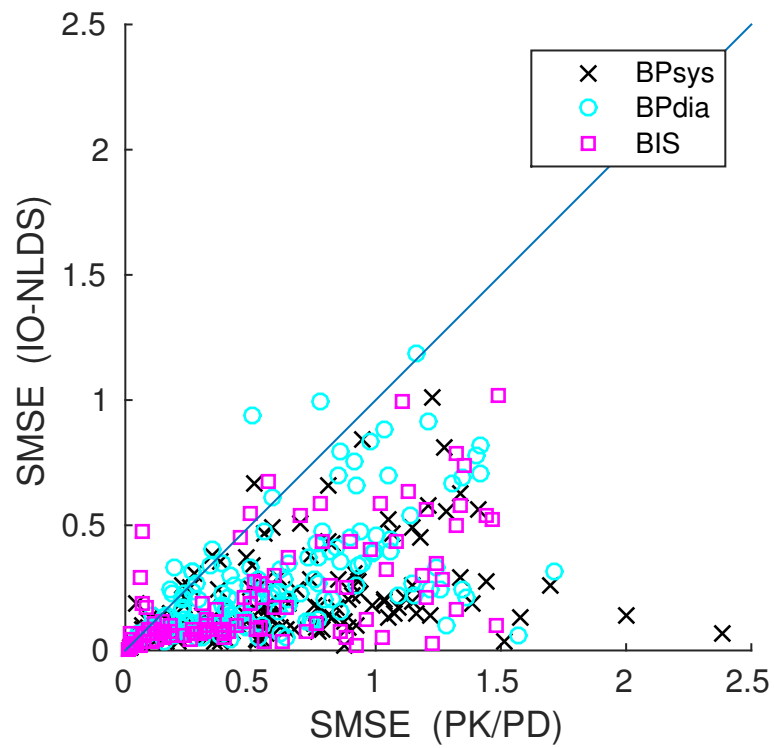


Figure S1: SMSE results for all patients across all four prediction horizons and all three measured channels. Points below the 45 degree line correspond to lower errors for the IO-NLDS on the same task and points above that line correspond to lower errors for the PK/PD model. The diagonal line represents equal performance between the two models.