Blending Physics and Data to Model Hemodynamic Effects Under General Anesthesia

Mattia Fregolent * **Michele Schiavo** * **Nicola Latronico** **,*** **Massimiliano Paltenghi** *** **Simone Del Favero** **** **Mirco Rampazzo** **** **Antonio Visioli** *

* *Dipartimento di Ingegneria Meccanica e Industriale, University of Brescia, Italy (e-mail: {mattia.fregolent, michele.schiavo, antonio.visioli}@unibs.it)*
** *Department of Medical and Surgical Specialties, Radiological Sciences and Public Health, University of Brescia, Brescia, Italy (e-mail: nicola.latronico@unibs.it)*
*** *Department of Anesthesiology, Critical Care and Emergency, Spedali Civili di Brescia, Brescia, Italy (e-mail: massimiliano.paltenghi@asst-spedalicivili.it)*
**** *Dipartimento di Ingegneria dell'Informazione, University of Padova, Italy (e-mail: {simone.delfavero, mirco.rampazzo}@unipd.it)*

Abstract:
General anesthesia, typically induced using a combination of hypnotic (propofol) and analgesic (remifentanil) drugs, is crucial for the success of surgical procedures, but it can cause dangerous cardiovascular side effects. In this context, models and simulations offer new opportunities to address the intrinsic complexity of the process, accelerating advances and innovation in the technology of anesthesia. This study aims to improve the modeling of hemodynamic effects under general anesthesia by expanding the applicability of a recent mechanistic model in combination with data-driven modules. In particular, we use a dataset related to plastic surgery for both model calibration and testing, preserving the physical interpretability of the mechanistic model while integrating it with data-driven components to enhance its predictive capabilities. The results demonstrate a significant improvement in the model ability to simulate hemodynamic variables under surgical conditions, offering potential applications for anesthesia monitoring and control systems design that consider the patient's cardiovascular safety. This enhanced hybrid model provides a more accurate representation of the complex interactions between anesthetic drugs and cardiovascular dynamics in real surgical settings.

Keywords: Control of Total Intravenous Anesthesia, first-principle model, data-driven model, hemodynamic effects.

1. INTRODUCTION

Surgical procedures are a routine part of medical practice, with millions of people worldwide undergoing surgery every day. One of the critical factors contributing to the success of these complex procedures is general anesthesia, which is accomplished through the combined administration of hypnotic and analgesic drugs, typically propofol and remifentanil (Smith et al., 2023). Over the past two decades, the integration of technology into various aspects of medicine has led to the adoption of automation in anesthesiology practice. This has been achieved through a range of solutions, including PID control and model predictive control (Pawlowski et al., 2023; Schiavo et al., 2023). These solutions are designed to optimize drug infusion, enhancing patient safety by preventing over- or under-dosing, detecting critical events, and reducing the overall workload of anesthesiologists as well as anesthesia-related costs (Ghita et al., 2020). However, closed-loop control systems often focus solely on utilizing an indicator of depth of hypnosis as the controlled variable, without considering the effects of the drugs on other vital variables critical to patient safety. For instance, propofol and remifentanil can cause cardiovascular side effects, particularly arterial hypotension, which significantly impacts surgical outcomes (Elliott et al., 2000; Sahinovic et al., 2018). Intraoperative hypotension has been associated with an increased risk of postoperative mortality, myocardial injury, myocardial infarction, cardiogenic shock, acute renal failure, delirium, and stroke (Guarracino and Bertini, 2022). In this context, predicting hemodynamic variables can enhance patient safety and stability by enabling precise, personalized drug dosing, particularly in automated anesthesia systems, to prevent complications and support faster recovery. However, existing hemodynamic models are limited, with many being empirical and lacking physi-

ological grounding, making them difficult to interpret and to accept clinically (Su et al., 2023). To the best of our knowledge, the mechanistic model proposed in (Su et al., 2023) is the first to describe both the functioning of the cardiovascular system and how the interaction between propofol and remifentanil affects it. The model is a human-scale extension of a previously identified hemodynamic model that was developed in rats (Snelder et al., 2014). Although the model is potentially clinically acceptable due to its physical nature, we hypothesize that it is not yet suitable for simulating hemodynamic variables in practical applications. Its development was based on data from a single study (Kuizenga et al., 2018) involving a healthy, homogeneous population without surgical stimulation, which oversimplifies the physiological responses occurring during invasive procedures in realistic settings.

The aim of this study is to to extend the applicability of the (Su et al., 2023) mechanistic model for forecasting hemodynamic data during surgeries. This is achieved by preserving as much of the original model structure as possible while balancing goodness-of-fit and interpretability. To this end, we propose a hybrid approach that combines the mechanistic model with black box models in a first-principle-data-driven fashion. In particular, the overall model is calibrated and tested exploiting a dataset which includes anesthesia-related observations from 48 plastic surgery procedures.

The paper is organized as follows: the surgery dataset is presented in Section 2.1. A brief overview of the (Su et al., 2023) mechanistic model is presented in Section 2.2. The proposed methodologies involving model calibration and hybrid modeling are described in Sections 2.3 and 2.4, respectively. Results are reported in Section 3. Finally, discussion of results and conclusions are given in Section 4.

2. METHODS

2.1 Data description and pre-processing

The dataset consists of anesthesia-related observations from 48 plastic surgery procedures performed at Brescia Hospital in Italy and it is limited to cases where only propofol and remifentanil were administered to minimize potential confounding effects from other drugs. Patient demographic data include age, sex, weight, height, and Body Mass Index (BMI). The cohort comprises 16 male and 32 female patients, with ages ranging from 27 to 82 years, thereby encompassing the entire adult age spectrum. BMI ranges from 19.2 to 36.4 $\mathrm{kg/m^2}$, representing a general population that includes both underweight and obese individuals. The considered variables are propofol infusion rate u_p [$\mu\mathrm{g{\cdot}s^{-1}}$], remifentanil infusion rate u_r [$\mathrm{ng{\cdot}s^{-1}}$], Hearth Rate (HR) [$\mathrm{min^{-1}}$], non-invasive Systolic blood Pressure (SP) [mmHg] and non-invasive Diastolic blood Pressure (DP) [mmHg]. Since specific and complex instrumentation is required, variables such as Total Peripheral Resistance (TPR) and Stroke Volume (SV) are not directly measured in practice. However, it is possible to obtain indirect measurements of these variables by exploiting some relationships that describe the physiology of the cardiovascular system. Starting from SV, assuming that the ability of the arterial system to expand and contract in response to changes in blood pressure remains constant, SP and DP can be used as a surrogate of SV as: SV = 1.5 (SP − DP) (Su et al., 2022). Mean Arterial Pressure (MAP) is calculated using MAP = (SP + 2DP) /3.
Propofol plasma concentrations Cp_p [μg/mL] are estimated by giving u_p in input to the pharmacokinetics (PK) model presented in (Eleveld et al., 2018). Remifentanil plasma concentrations Cp_r [ng/mL] are estimated by giving u_r as input to the PK model presented in (Eleveld et al., 2017).
We will refer to the considered patients group with $\mathbf{S}$. The following approaches are tuned on a group $\mathbf{S}_{\mathrm{TRAIN}} = \{s_1,\ s_2,\ \cdots,\ s_{40}\}$ of 40 subjects, while to analyse the generalizability of our methods, we use the remaining 8 subjects denoted as $\mathbf{S}_{\mathrm{TEST}} = \{s_{41},\ \cdots,\ s_{48}\}$.

2.2 Mechanistic model

The mechanistic model presented in (Su et al., 2023) distinguishes between two types of parameters: system-specific and drug-dependent. The former describes the cardiovascular system's physiological functioning, including the baroreceptor reflex and the inverse effect of HR on SV (Snelder et al., 2014). The latter details how the interaction between propofol and remifentanil affects the cardiovascular system itself. To have a deep understanding of each parameter, see (Su et al., 2023). Here, we propose to rewrite the model in its state space representation:

$$\begin{cases} \dot{\boldsymbol{x}}(t) = \boldsymbol{f}\left(\boldsymbol{x}(t)\right) + \boldsymbol{g}\left(\boldsymbol{x}(t)\right)\boldsymbol{l}\left(\boldsymbol{u}(t)\right), \\ y(t) = h\left(\boldsymbol{x}(t)\right), \end{cases} \tag{1}$$

where $\boldsymbol{x} \in \mathbb{R}^{3\times 1}$ is the state vector, $\boldsymbol{f} \in \mathbb{R}^{3\times 1}$, $\boldsymbol{g} \in \mathbb{R}^{3\times 5}$ and $h \in \mathbb{R}$ are nonlinear functions of $\boldsymbol{x}$, and $\boldsymbol{l} \in \mathbb{R}^{5\times 1}$ is a nonlinear function of the system inputs $\boldsymbol{u} \in \mathbb{R}^{2\times 1}$.
More specifically,

$$\boldsymbol{x} = \begin{bmatrix} x_1(t) \\ x_2(t) \\ x_3(t) \end{bmatrix} = \begin{bmatrix} \mathrm{TPR}(t) \\ \mathrm{HR}(t) \\ \mathrm{SV}^*(t) \end{bmatrix}, \tag{2}$$

$$\boldsymbol{f} = k\left(\frac{x_{10}x_{20}x_{30}}{1 - \mathrm{IE}\ln\dfrac{x_2(t)}{x_{20}}}\right)^{\mathrm{FB}} A(\boldsymbol{x}(t))\ \boldsymbol{x}(t)^{\circ(-\mathrm{FB})} - k\boldsymbol{x}(t), \tag{3}$$

$$\boldsymbol{g} = kB(\boldsymbol{x}(t)). \tag{4}$$

In (3), $^{\circ(-\mathrm{FB})}$ denotes the Hadamard power of the vector $\boldsymbol{x}$, obtained by raising each entry to the $-$FB power. Moreover, FB, IE, and k [$\mathrm{s^{-1}}$] are positive system-specific parameters (Su et al., 2023). The initial conditions of $\boldsymbol{x}(t)$, i.e., $\boldsymbol{x}(0)$ are x_{10}, x_{20}, and x_{30}, and they are set to the median values of the observed TPR, HR and SV before the start of drug administration. In (3) and (4), $A \in \mathbb{R}^{3\times 3}$ and $B \in \mathbb{R}^{3\times 5}$ are matrices that depend on the state vector and they are defined in (5) and (6).

$$B = \begin{bmatrix} x_{10}\left(\dfrac{x_{10}x_{20}x_{30}}{x_1(t)x_2(t)x_3(t)\left(1-\mathrm{IE}\ln\dfrac{x_2(t)}{x_{20}}\right)}\right)^{\mathrm{FB}} & 0 & x_1(t) & 0 & 0 \\ 0 & 0 & 0 & x_2(t) & 0 \\ 0 & x_{30}\left(\dfrac{x_{10}x_{20}x_{30}}{x_1(t)x_2(t)x_3(t)\left(1-\mathrm{IE}\ln\dfrac{x_2(t)}{x_{20}}\right)}\right)^{\mathrm{FB}} & 0 & 0 & x_3(t) \end{bmatrix} \tag{6}$$

$$A = \begin{bmatrix} \dfrac{x_{10}}{(x_2(t)x_3(t))^{\mathrm{FB}}} & 0 & 0 \\ 0 & \dfrac{x_{20}}{(x_1(t)x_3(t))^{\mathrm{FB}}} & 0 \\ 0 & 0 & \dfrac{x_{30}}{(x_1(t)x_2(t))^{\mathrm{FB}}} \end{bmatrix} \tag{5}$$

In (1), $\boldsymbol{l}$ is a function of $\boldsymbol{u}(t) = [Cp_p(t),\, Cp_r(t)]^\intercal$ defined as follows:

$$\boldsymbol{l} = \begin{bmatrix} Ep_{\mathrm{TPR}}(\boldsymbol{u}(t)) \\ Ep_{\mathrm{SV}}(\boldsymbol{u}(t)) \\ Er_{\mathrm{TPR}}(\boldsymbol{u}(t)) \\ Er_{\mathrm{HR}}(\boldsymbol{u}(t)) \\ Er_{\mathrm{SV}}(\boldsymbol{u}(t)) \end{bmatrix}, \tag{7}$$

where each row corresponds to a pharmacodynamic (PD) relation. In particular, $Ep_{\mathrm{TPR}}(\boldsymbol{u}(t))$ and $Ep_{\mathrm{SV}}(\boldsymbol{u}(t))$ describe propofol's effect, influenced by remifentanil, on TPR and SV, respectively. Similarly, $Er_{\mathrm{TPR}}(\boldsymbol{u}(t))$, $Er_{\mathrm{HR}}(\boldsymbol{u}(t))$ and $Er_{\mathrm{SV}}(\boldsymbol{u}(t))$ represent the effect of remifentanil, influenced by propofol, on TPR, HR and SV, respectively. These PD functions constitute the drug-specific part of the (Su et al., 2023) model.

Finally, MAP is the model output $y(t) \in \mathbb{R}$ defined by the h function in (1) as follows:

$$h = x_1(t)x_2(t)x_3(t)\left(1 - \mathrm{IE}\ln\frac{x_2(t)}{x_{20}}\right). \tag{8}$$

2.3 Model calibration

We recalibrate the mechanistic model to enhance its performance in predicting data specific to surgical procedures, as its original development dataset did not fully represent real invasive operations. Our approach seeks to optimize specific parameters to values closely aligned with those identified in (Su et al., 2023), thereby maintaining their biological significance.

First, we perform a Global Sensitivity Analysis (GSA) to gain a deep understanding of the mechanism-based model and to rank its parameters. Then, the most influential parameters are optimized to provide a calibrated version of the model. Thus, GSA avoids the problem of blindly selecting parameters for optimization. Our approach involves screening parameters using Morris indices evaluation (Morris, 1991), followed by ranking the significant parameters quantitatively by decomposing the variance of the model output into fractions attributable to parameters or a combination of parameters (Saltelli, 2008). High-ranked parameters are those that most significantly influence the variability of the model output. For the mechanistic model (1), we estimated the so called first-order and total-order Sobol indices (Saltelli, 2008). GSA was performed using the SAFE Toolbox (Pianosi et al., 2015) version for MATLAB, which is available at `https://safetoolbox.github.io`.

Formally, we denote by $\mathbf{p}$ the set of parameters of model (1). We choose to optimize the m most sensitive parameters, that are denoted by $\boldsymbol{\phi}$. The goal is to solve the following minimization problem:

$$\boldsymbol{\phi}^* = \arg\min_{\boldsymbol{\phi}\in\Phi}\{J(\boldsymbol{\phi})_{\mathbf{S}_{\mathrm{TRAIN}}}\} \tag{9}$$

where J is a cost function designed for quantifying the performance of (1), evaluated for a given $\boldsymbol{\phi}$, in predicting hemodynamics observations in $\mathbf{S}_{\mathrm{TRAIN}}$. The solution of (9) is the set of optimal parameters $\boldsymbol{\phi}^*$ that minimizes J, found within the m-dimensional parameters search space Φ. We denote by $\boldsymbol{p}^*$ the parameter set resulting from substituting the m optimized parameters $\boldsymbol{\phi}^*$ for their original values in $\boldsymbol{p}$.

In practice, to solve (9), we employ the Bayesian Optimization (BO) methodology (Shahriari et al., 2015), following an implementation protocol inspired by (Villaverde et al., 2022), which is summarized in Alg. (1).

Algorithm 1 Parameters Optimization

Initialize: $\boldsymbol{\phi}$ ▷ Select m most sensitive parameters
Initialize: k ▷ Number of folds for cross-validation
Initialize: $\{\mathbf{S}_1,\ \mathbf{S}_2,\ \cdots,\ \mathbf{S}_k\} = \mathbf{S}_{\mathrm{TRAIN}}$ ▷ Disjointed sub-groubs of training subjects
for $i = 1,\ \cdots,\ k$ **do**
  $\mathbf{S}^{(i)}_{validation} \leftarrow \mathbf{S}_i$
  $\mathbf{S}^{(i)}_{train} \leftarrow \mathbf{S}_{\sim i}$ ▷ All sub-groups other than $\mathbf{S}_i$
  $\boldsymbol{\phi}^{*(i)} = \arg\min_{\boldsymbol{\phi}\in\Phi}\left\{J(\boldsymbol{\phi})_{\mathbf{S}^{(i)}_{train}}\right\}$ ▷ Optimized parameters for current training group using BO. J is calculated through (10).
  $\Delta_i\left(\boldsymbol{\phi}^*\right) = \left|\frac{J(\mathbf{p}) - J(\boldsymbol{\phi}^{*(i)})_{\mathbf{S}^{(i)}_{validation}}}{J(\mathbf{p})}\right|$ ▷ Relative J drop on validation group
end for
$\boldsymbol{\phi}^* = \arg\max_{\boldsymbol{\phi}^*\in\{\boldsymbol{\phi}^{*(1)},\ \boldsymbol{\phi}^{*(2)},\ \cdots,\ \boldsymbol{\phi}^{*(k)}\}}\{\Delta_{1:k}(\boldsymbol{\phi}^*)\}$

The J we design depends on both y and the states x_2 and x_3, as we are able to compare them with the available data. Specifically, considering a generic subject $s \in \mathbf{S}_{\mathrm{TRAIN}}$, we denote the MAP observations over time with the vector

$$\boldsymbol{y}^{(s)} = \left[y^{(s)}(0),\ y^{(s)}(t_1),\ \cdots,\ y^{(s)}(t_{ns})\right]^\intercal \in \mathbb{R}^{(ns+1)\times 1}$$

and HR and SV observations with vectors $\boldsymbol{x}_2^{(s)}$ and $\boldsymbol{x}_3^{(s)}$. Similarly, we build the $(ns+1)$-dimension vectors $\boldsymbol{Cp}_p^{(s)}$

and $\boldsymbol{Cp}_r^{(s)}$ that compose $\boldsymbol{u}^{(s)}$ (see Sec. 2.2) to be given as input to the model (1) in order to obtain predictions $\hat{\boldsymbol{y}}^{(s)}$, $\hat{\boldsymbol{x}}_2^{(s)}$ and $\hat{\boldsymbol{x}}_3^{(s)}$. Then, we evaluate the Root Mean Squared Error (RMSE) for each prediction within s as an indicator of prediction quality. For $\hat{\boldsymbol{y}}^{(s)}$ it is calculated as

$$\mathrm{RMSE}_y^{(s)}(\boldsymbol{\phi}) = \sqrt{\frac{\left[\boldsymbol{y}^{(s)} - \hat{\boldsymbol{y}}^{(s)}\right]^\intercal \left[\boldsymbol{y}^{(s)} - \hat{\boldsymbol{y}}^{(s)}\right]}{n_s}}$$

and the same can be done to calculate $\mathrm{RMSE}_{x_2}^{(s)}(\boldsymbol{\phi})$ and $\mathrm{RMSE}_{x_3}^{(s)}(\boldsymbol{\phi})$. Finally, to build J with the required specifications to solve the minimization problem (9), we define:

$$J(\boldsymbol{\phi})_{\mathbf{S}} = \sum_{i \in \{y, x_2, x_3\}} w_i \left(\frac{\sum_{s \in \mathbf{S}} \mathrm{RMSE}_i^{(s)}(\boldsymbol{\phi})}{\max\{\mathrm{RMSE}_i^{(s)}(\boldsymbol{\phi})\}_{s \in \mathbf{S}}} \right). \quad (10)$$

The normalization in (10) allows us to compare the RMSE of different quantities. The weights w_y, w_{x_2} and w_{x_3} are positive user-defined values such that $w_y + w_{x_2} + w_{x_3} = 1$. $\mathbf{S}$ represents the general group of subjects for which we aim to globally assess the model performance.

2.4 Hybrid modeling

Techniques aimed to identify unknown or partially-known process mechanisms constrained by already defined first-principles equations are referred to as *hybrid modeling* (Sohlberg and Jacobsen, 2008; Czop et al., 2011). In this context, parallel and series hybrid options integrate first-principle (white-box) and data-driven (black-box) models. The difference lies in how these modules interact to generate the final prediction. In hybrid-parallel architectures, the two modules are both fed with the same data simultaneously: while the mechanistic module estimates the system behavior based on first principles, the data-driven component aims to predict the residuals that need to be added to the mechanistic model predictions to obtain the observed experimental data (Duarte et al., 2004). The main challenge is to train, using $\mathbf{S}_{\mathrm{TRAIN}}$, a generalizable black-box model $z_{par}\left(\mathbf{u},\ \hat{\boldsymbol{y}},\ \hat{\boldsymbol{x}}_2,\ \hat{\boldsymbol{x}}_3,\ \boldsymbol{\theta}_{par}\right)$ able to give a prediction of $\boldsymbol{r}_y = \boldsymbol{y} - \hat{\boldsymbol{y}}$, $\boldsymbol{r}_{x_2} = \boldsymbol{x}_2 - \hat{\boldsymbol{x}}_2$ and $\boldsymbol{r}_{x_3} = \boldsymbol{x}_3 - \hat{\boldsymbol{x}}_3$, where $\boldsymbol{y}$ is a vector built by concatenating $\boldsymbol{y}^{(s)}\ \forall s \in \mathbf{S}_{\mathrm{TRAIN}}$, formally:

$$\boldsymbol{y} = \left[\boldsymbol{y}^{(s_1)^\intercal},\ \boldsymbol{y}^{(s_2)^\intercal},\ \cdots,\ \boldsymbol{y}^{(s_{40})^\intercal}\right]^\intercal \in \mathbb{R}^{\left(\sum_{i=1}^{40}(n_{s_i}+1)\right)\times 1}. \quad (11)$$

The calculation in (11) is similarly performed to obtain $\hat{\boldsymbol{y}}$, $\boldsymbol{x}_2$, $\hat{\boldsymbol{x}}_2$, $\boldsymbol{x}_3$, $\hat{\boldsymbol{x}}_3$, $\boldsymbol{Cp}_p$ and $\boldsymbol{Cp}_r$. Then, we can form $\mathbf{u} = [\boldsymbol{Cp}_{\boldsymbol{p}},\ \boldsymbol{Cp}_{\boldsymbol{r}}]^\intercal$. $\boldsymbol{\theta}_{par}$ is the set of hyper-parameters of the black-box model. Once trained, the black-box model z_{par} can be used in the forecasting phase to predict residuals that are then added to the mechanistic model outputs in order to obtain the final parallel hybrid model predictions. For hybrid models developed using the series approach, the two modules are fed sequentially (Duarte et al., 2004; Czop et al., 2011). Here, we propose a sequential approach in which the black-box model $z_{ser}(\mathbf{u},\ \boldsymbol{\theta}_{ser})$ predicts for a given subject s, a subset of m' optimized parameters $\boldsymbol{\phi}^{*(s)}$ that update $\boldsymbol{p}^{*(s)}$ of the mechanistic model (1), which is used to model the system. This series approach can be seen as a personalized auto-calibration of some parameters during the simulation. The choice of the m' parameters to optimize is guided by GSA. $\boldsymbol{\theta}_{ser}$ denotes the hyperparameters of the model.
To train z_{ser}, we must first derive $\boldsymbol{\phi}^{*(s)}\ \forall s \in \mathbf{S}_{\mathrm{TRAIN}}$ by solving the minimization problem (9) separately $\forall s$. For this training procedure, we can still exploit Alg. (1) by performing the k-fold cross-validation on each subject sample points. Finally, the training is performed by providing $\mathbf{u}$ as input to z_{ser} and the concatenation of $\boldsymbol{\phi}^{*(s)} \in \mathbb{R}^{1\times m'}\ \forall s \in \mathbf{S}_{\mathrm{TRAIN}}$ as the output to predict.
As regards the last hybrid modeling technique we propose, it is a combination of the two previous ones. That is, a further data-driven model

$$z_{par'}\left(\mathbf{u},\ \hat{\boldsymbol{y}}_{ser},\ \hat{\boldsymbol{x}}_{2_{ser}},\ \hat{\boldsymbol{x}}_{3_{ser}},\ \boldsymbol{\theta}_{par'}\right),$$

is added in parallel to the previous described hybrid model $z_{ser}(\mathbf{u},\ \boldsymbol{\theta}_{ser})$ with the task of correcting its predictions. This involves an increase in the structural complexity of the overall model, as well as a decrease in its interpretability: in total 3 models are involved, 2 of these are black-box models.

The developed data-driven models were simple Feedforward Neural Networks (FNNs), chosen based on prior research highlighting their extensive use in hybrid modeling approaches (Duarte et al., 2004; Sohlberg and Jacobsen, 2008). we considered FNNs with tanh as the activation function for a single hidden layer composed of $\mathscr{L}$ neurons, where $\mathscr{L}$ is a hyperparameter to be tuned. In general, the nets identification process involves dividing $\mathbf{S}_{\mathrm{TRAIN}}$ data into training data (85%) and validation data (15%). Then, to tune $\mathscr{L}$ we manually train the network with an increasing $\mathscr{L}$ from 2 to 20. The final $\mathscr{L}$ value we choose is the one that lead the network to have the lowest J calculated for validation subjects as in (10). Additionally, the inclusion of demographic variables among the network inputs is assessed. Demographic variables are evaluated individually for inclusion in the model. A variable is retained if its addition reduces prediction error on the validation data; otherwise, it is excluded if it does not improve or worsens prediction quality. The implementation was performed using the Deep Learning Toolbox in MATLAB.

3. RESULTS

The mechanistic model presented in (Su et al., 2023), its calibrated version, and the three implemented hybrid models are compared by performing predictions on subjects not used during the development of our approaches. We employ the cost function $J_{\mathbf{S}_{\mathrm{TEST}}}$ as in (10) to evaluate models performance of the whole test group. We observe that our calibration results in a 23% reduction in $J_{\mathbf{S}_{\mathrm{TEST}}}$ compared to the original model (see Fig. 1). Furthermore, all hybrid models performed better than the calibrated model, particularly due to the inclusion of FNNs in parallel for residual estimation. The best prediction quality was achieved by the hybrid series-parallel model, which shows a 34% reduction in $J_{\mathbf{S}_{\mathrm{TEST}}}$ compared to the original model. We also compare the performance of the original model and the best hybrid model for each test subject individually. As shown in Fig. 2, the hybrid model outperforms the original model in all test subjects except one. To achieve these results, Alg. (1) is iterated for increasing values of m. We found that using $m = 3$ parameters to optimize yields the best result in terms of cost function reduction during validation. The $m = 3$

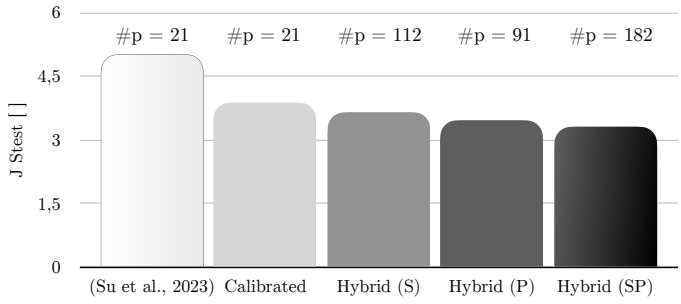

Fig. 1. Models performance comparison in terms of the RMSE-based cost function (10) evaluated for the whole group of test subjects. The number of parameters #p is also reported.

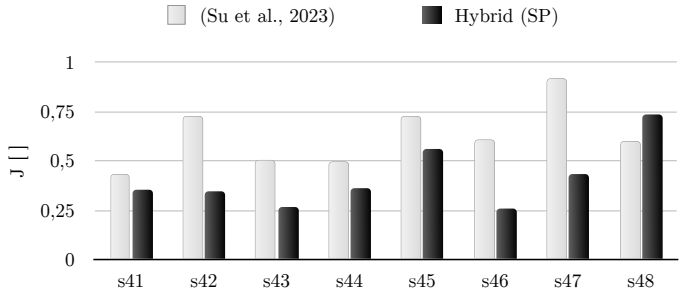

Fig. 2. Performance comparison between the (Su et al., 2023) mechanistic model and the hybrid series-parallel hybrid model in terms of the RMSE-based cost function (10) evaluated for each test subject separately.

most sensitive mechanistic model parameters according to GSA are the system-specific parameter FB and the drugs-specific parameters $\mathrm{INT_{HR}}$ and $\mathrm{Emax}_{p\_\mathrm{TPR}}$. Once optimized, these parameters experienced the following absolute changes from their nominal values: -5%, +29% and -30%, respectively. As regards the design of $J_{\mathbf{S}}$ in (10), we assign a greater weight w_y with respect to w_{x_2} and w_{x_3}, because MAP is the most critical hemodynamic variable to evaluate for patient safety among the three. Considering that we aim to employ the model to impose hemodynamics constraints during automatic anesthesia control, an accurate MAP prediction would be preferred. Then, an intermediate value is assigned to w_{x_2}. As for w_{x_3}, a lower weight is assigned since it is an indirect measure that holds true under certain assumptions (Su et al., 2022). Following these considerations the weights $w_y = 0.4$, $w_{x_2} = 0.35$ and $w_{x_3} = 0.25$ were set. As for the final structure of the FNNs adopted for the hybrid models, z_{par} has a number $\mathscr{L} = 4$ of hidden neurons, 9 input and 3 output neurons. The patient's demographics BMI, weight and sex were added to the net inputs. The FNN for z_{ser} has a number $\mathscr{L} = 5$ of hidden neurons, 6 input and 8 output neurons. The inclusion of patient's demographics BMI, weight and sex proved to be influential in achieving a better fit for this network as well. Note that for this model the 8 outputs are the 8 most sensitive parameters according to GSA, as we found that this number of parameters gave better results in training phase. Specifically, these are FB, $\mathrm{INT_{HR}}$, $\mathrm{Emax}_{p\_\mathrm{TPR}}$, $\mathrm{INT_{SV}}$, $\mathrm{INT_{TPR}}$, $\mathrm{Emax}_{p\_\mathrm{SV0}}$ and $\mathrm{Emax}_{r\_\mathrm{TPR}}$ (see (Su et al., 2023) for description). Finally, the FNN adopted for the hybrid model $z_{par'}$ has a number $\mathscr{L} = 6$ of hidden neurons, 7 input and 3 output neurons. The patient's demographics BMI and sex were added to the net inputs.

4. DISCUSSION AND CONCLUSIONS

The only published mechanistic model for predicting hemodynamic variables during anesthesia with propofol and remifentanil (Su et al., 2023) was utilized in this study. However, we observed that its predictions did not adequately capture anesthesia-induced hemodynamic variations in a surgical setting. This limitation may arise because the cardiovascular system is highly complex and, during clinical interventions, is significantly affected by adverse drug reactions, complications, surgical incisions, noxious stimuli, fluid loss, and other factors that are likely not fully captured by the mechanistic model's structure. Moreover, the data used for model identification originate from a study design that does not fully account for the system's complexity under surgical conditions.

We aimed to improve mechanistic model adherence to experimental data by preserving its structure in order to maintain its biological interpretability. To achieve this, we calibrated the model by optimizing its most sensitive parameters. This adjustment improved the model's fit to test subjects (Fig. 1); however, it did not fully capture the higher frequency dynamics. Notably, parameter optimization resulted in a 5% decrease in FB, indicating a reduced magnitude of the baroreceptor reflex. This is likely due to the greater heterogeneity of subjects in our dataset or the higher drug concentrations (primarily remifentanil), which may contribute to baroreflex inhibition. The substantial increase in the absolute value of $\mathrm{INT_{HR}}$ appears to mitigate the HR-raising effect of remifentanil, suggesting that at high concentrations, remifentanil may not exhibit the same HR-raising effect reported in (Su et al., 2023). It is crucial to emphasize that these considerations hold true only when these parameters are uniquely identifiable; otherwise, their biological significance becomes ambiguous or even lost. However, conducting an identifiability analysis for a highly nonlinear model is challenging. As suggested by (Dobre et al., 2010), GSA can serve as a valuable tool for gaining insights into parameter identifiability. In this study, we leveraged GSA to establish a ranking of parameters, providing an initial step toward understanding their relative importance. In future work, we intend to further explore GSA techniques specifically for identifiability assessment, along with a potential re-parameterization of the model.

Our results confirm that incorporating data-driven models significantly improves the goodness of fit. In particular, estimating residuals proves highly effective, whether applied in parallel with the calibrated mechanistic model or within the hybrid series model. This effectiveness arises not only from the increased model complexity and parameter count but also from the ability of the black-box model to capture unexplained dynamics in residuals during training. Notably, the hybrid approach outperformed the calibrated mechanistic model in 7 out of 8 test subjects (Fig. 2). The slightly poorer performance observed in subject s48 may be attributed to peculiar surgical events not represented within the training group. Overall, only 48 patients were included in this study, and developing more generalizable black-box models may require a significantly larger

dataset, potentially incorporating data from multiple centers and hospitals. However, the analysis of residuals from the best-performing hybrid model indicates that not all system dynamics are captured. In a future work we will employ advanced neural networks to achieve white noise residuals.

Another limitation of our study is that integrating black-box models with the mechanistic model negatively impacts its interpretability, a crucial factor for clinical approval. To address this, methods such as SHAP analysis (Ali et al., 2023) will be implemented to enhance the explainability of black-box model predictions. Additionally, we excluded intraoperative boluses of drugs to avoid confounding effects on hemodynamics. While propofol and remifentanil are the primary agents used for hypnosis and analgesia during general anesthesia, boluses of other drugs are frequently administered and contribute to variations in hemodynamic variables. Future work should incorporate these additional drugs into the model to improve its physiological accuracy.

In conclusion, we proposed a methodology for adapting a mechanistic model predictions to hemodynamic observations in a heterogeneous patient population undergoing real surgical procedures by optimizing most sensitive parameters. Additionally, we developed hybrid approaches that balance interpretability and goodness-of-fit by integrating black-box models with the calibrated mechanistic model. With access to a larger, multi-center dataset and more sophisticated neural network approaches, alongside with their explainability analyses, our approach could be further refined and potentially integrated into automated anesthesia control strategies, ensuring that drug infusion rates maintain the patient's hemodynamic variables within a safe range.

REFERENCES

Ali, S., Abuhmed, T., El-Sappagh, S., Muhammad, K., Alonso-Moral, J.M., Confalonieri, R., Guidotti, R., Del Ser, J., Díaz-Rodríguez, N., and Herrera, F. (2023). Explainable artificial intelligence (xai): What we know and what is left to attain trustworthy artificial intelligence. *Information fusion*, 99, 101805.

Czop, P., Kost, G., Sławik, D., and Wszołek, G. (2011). Formulation and identification of first-principle data-driven models. *Journal of Achievements in materials and manufacturing Engineering*, 44(2), 179–186.

Dobre, S., Bastogne, T., and Richard, A. (2010). Global sensitivity and identifiability implications in systems biology. *IFAC Proceedings Volumes*, 43(6), 54–59.

Duarte, B., Saraiva, P., and Pantelides, C. (2004). Combined mechanistic and empirical modelling. *International Journal of Chemical Reactor Engineering*, 2(1).

Eleveld, D., Colin, P., Absalom, A., and Struys, M. (2018). Pharmacokinetic–pharmacodynamic model for propofol for broad application in anaesthesia and sedation. *British journal of anaesthesia*, 120(5), 942–959.

Eleveld, D.J., Proost, J.H., Vereecke, H., Absalom, A.R., Olofsen, E., Vuyk, J., and Struys, M.M. (2017). An allometric model of remifentanil pharmacokinetics and pharmacodynamics. *Anesthesiology*, 126(6), 1005–1018.

Elliott, P., O'Hare, R., Bill, K.M., Phillips, A.S., Gibson, F.M., and Mirakhur, R.K. (2000). Severe cardiovascular depression with remifentanil. *Anesthesia & Analgesia*, 91(1), 58–61.

Ghita, M., Neckebroek, M., Muresan, C., and Copot, D. (2020). Closed-loop control of anesthesia: Survey on actual trends, challenges and perspectives. *IEEE Access*, 8, 206264–206279.

Guarracino, F. and Bertini, P. (2022). Perioperative hypotension: causes and remedies. *Journal of Anesthesia, Analgesia and Critical Care*, 2(1), 17.

Kuizenga, M.H., Colin, P.J., Reyntjens, K.M., Touw, D.J., Nalbat, H., Knotnerus, F.H., Vereecke, H.E., and Struys, M.M. (2018). Test of neural inertia in humans during general anaesthesia. *British Journal of Anaesthesia*, 120(3), 525–536.

Morris, M.D. (1991). Factorial sampling plans for preliminary computational experiments. *Technometrics*, 33(2), 161–174.

Pawlowski, A., Schiavo, M., Latronico, N., Paltenghi, M., and Visioli, A. (2023). Experimental results of an mpc strategy for total intravenous anesthesia. *IEEE Access*, 11, 32743–32751.

Pianosi, F., Sarrazin, F., and Wagener, T. (2015). A matlab toolbox for global sensitivity analysis. *Environmental Modelling & Software*, 70, 80–85.

Sahinovic, M.M., Struys, M.M., and Absalom, A.R. (2018). Clinical pharmacokinetics and pharmacodynamics of propofol. *Clinical pharmacokinetics*, 57(12), 1539–1558.

Saltelli, A. (2008). *Global Sensitivity Analysis: the Primer*. John Wiley & Sons.

Schiavo, M., Padula, F., Latronico, N., Paltenghi, M., and Visioli, A. (2023). Experimental results of an event-based pid control system for propofol and remifentanil coadministration. *Control Engineering Practice*, 131, 105384.

Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., and De Freitas, N. (2015). Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1), 148–175.

Smith, G., D'Cruz, J.R., Rondeau, B., and Goldman, J. (2023). General anesthesia for surgeons. In *StatPearls [Internet]*. StatPearls Publishing.

Snelder, N., Ploeger, B., Luttringer, O., Rigel, D., Fu, F., Beil, M., Stanski, D., and Danhof, M. (2014). Drug effects on the cvs in conscious rats: separating cardiac output into heart rate and stroke volume using pkpd modelling. *British Journal of Pharmacology*, 171(22), 5076–5092.

Sohlberg, B. and Jacobsen, E.W. (2008). Grey box modelling–branches and experiences. *IFAC Proceedings Volumes*, 41(2), 11415–11420.

Su, H., Eleveld, D.J., Struys, M.M., and Colin, P.J. (2022). Mechanism-based pharmacodynamic model for propofol haemodynamic effects in healthy volunteers. *British Journal of Anaesthesia*, 128(5), 806–816.

Su, H., Koomen, J.V., Eleveld, D.J., Struys, M.M., and Colin, P.J. (2023). Pharmacodynamic mechanism-based interaction model for the haemodynamic effects of remifentanil and propofol in healthy volunteers. *British Journal of Anaesthesia*, 131(2), 222–233.

Villaverde, A.F., Pathirana, D., Fröhlich, F., Hasenauer, J., and Banga, J.R. (2022). A protocol for dynamic model calibration. *Briefings in bioinformatics*, 23(1), bbab387.