Benchmark Problem: A PK/PD Model and Safety Constraints for Anesthesia Delivery

Victor Gan, Guy A. Dumont and Ian M. Mitchell*

March 31, 2014

Abstract

Hypnosis, or depth of unconsciousness, is one of the goals of general anesthesia. In this brief paper we provide a differential equation model of how propofol, a commonly used intravenous anesthetic drug, leads to hypnosis. The model has two components: the pharmacokinetics (PK) describing how the drug is metabolized by the body, and the pharmacodynamics (PD) describing how the drug effects the depth of hypnosis. This standard PK/PD model is a compartmental model, and is linear except for an internal time delay and a nonlinear output mapping. We discuss how this model can be simplified and/or complexified so as to take best advantage of the capabilities of a particular analysis method. One goal of developing such a model is to ensure patient safety during surgery, so we describe an example safety verification problem. Finally, in order to demonstrate that this benchmark is operational, we provide code to simulate one version of the system.

Category: academic. **Difficulty:** medium

1 Introduction

General anesthesia is a broad term encompassing the use of drugs to induce and maintain three states during surgery: hypnosis (depth of unconsciousness), analgesia (absence of pain) and areflexia (lack of movement). While most surgeries will involve the use of multiple anesthetic drugs in order to manage all three states, in this brief paper we focus on the use of propofol [6], an intravenous anesthetic commonly used to induce and maintain hypnosis. Measuring the level of hypnosis is achieved by monitoring electroencephalography (EEG) signals to determine the depth of anesthesia [2]. The intensity of surgical stimulation during a surgery, haemodynamic effects of the drugs, underdosing due to equipment failure and overdosing due to inappropriate titration of the hypnotic components, among other factors, often result in imbalances between the amount of anesthetic required and the amount of anesthetic delivered [2, 4, 11]. Inappropriate anesthetic delivery can have severe consequences. If not enough anesthetic is delivered, the patient can remain conscious during a surgery which causes both trauma and anxiety [5, 12]. Too much anesthetic may have detrimental long term effects on the patient [10].

Delivery of anesthetics during surgery has traditionally been done manually, but the push to increase healthcare efficiency and reduce the number adverse incidents has driven improvements in sensing and modeling to the point where automated control of anesthetics is on the cusp of technological feasibility. That said, much work remains to be done in order to demonstrate the safety and efficacy of such systems. With safety analysis in mind, this paper and its accompanying code gathers together the following elements:
- A series of related models of the effect of propofol infusion on depth of hypnosis.
- A quantitative statement of the safety verification problem of ensuring suitable depth of hypnosis and drug concentration levels.
- MATLAB code to simulate trajectories for the core PK/PD model.

---

*VG & IMM are with the Department of Computer Science, University of British Columbia, {vhg,mitchell}@cs.ubc.ca. GAD is with the Department of Electrical and Computer Engineering, University of British Columbia, guyd@ece.ubc.ca
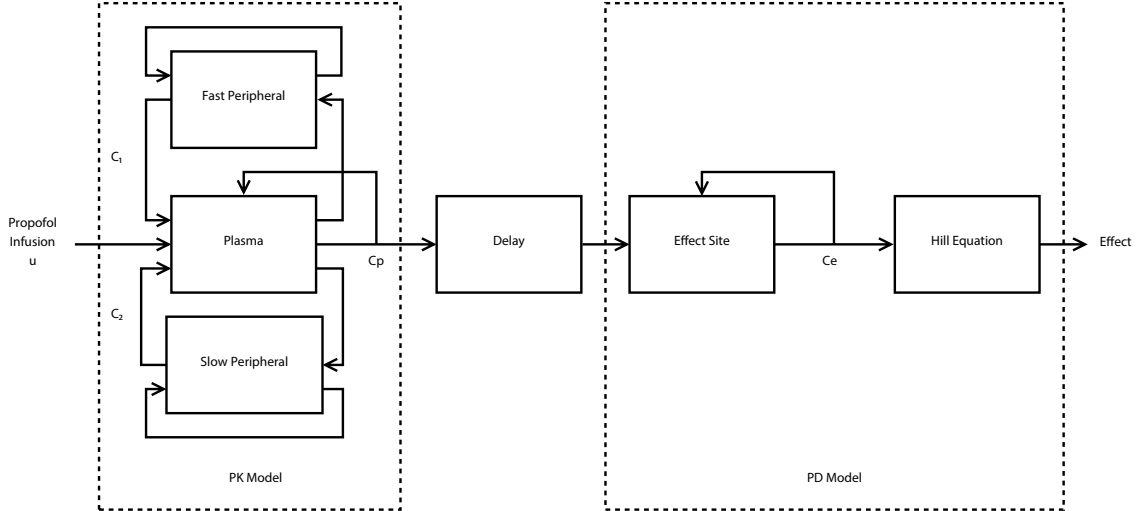
Figure 1: Block diagram of the full PK/PD model (1)– (3).

- Citations to the source literature for the models and parameters as well as some previously completed verifications.

2 The Models

Most surgeries involve two phases of anesthesia delivery: induction and maintenance. The induction phase begins with a bolus of anesthetic—a large dose delivered quickly—designed to achieve sufficient depth of hypnosis. As the bolus is metabolized, a much smaller and time-varying infusion of anesthetic is then used to maintain the hypnosis. Although models of the types discussed below are sometimes used in both phases, here we will focus on the maintenance phase and assume that the patient begins near the desired equilibrium of depth of hypnosis.

2.1 A PK/PD Model

Like many drugs, the pharmacologic effects of anesthesia are traditionally modelled in two stages. The pharmacokinectics (PK) track how the concentration of the drug changes as the body metabolizes it; in other words, what the body does to the drug. The pharmacodynamics (PD) track how the concentration of the drug leads to desired (or undesired) effects; in other words, what the drug does to the body. Figure 1 shows the components of the PK/PD model described in more detail below.

We start with the PK model, which is traditionally a linear multi-compartment model that describes drug concentrations between representative compartments of a patient. Anywhere from one to four compartments have been used, but we choose a three compartment version [7] here in order to push toward the more computationally challenging end of that spectrum and its suitability for describing propofol [13]. The state variables are the drug concentrations within each compartment; in this case the plasma c_p, the fast peripheral c_1 and the slow peripheral c_2 compartments, all measured in $[\mu g \ / \ \ell]$. The state space model is:

$$\begin{bmatrix} \dot{c}_p(t) \\ \dot{c}_1(t) \\ \dot{c}_2(t) \end{bmatrix} = \begin{bmatrix} -(k_{10} + k_{12} + k_{13}) & k_{12} & k_{13} \\ k_{21} & -k_{21} & 0 \\ k_{31} & 0 & -k_{31} \end{bmatrix} \begin{bmatrix} c_p(t) \\ c_1(t) \\ c_2(t) \end{bmatrix} + \begin{bmatrix} 1/V_1 \\ 0 \\ 0 \end{bmatrix} u(t) \tag{1}$$

The input to the PK model is the infusion rate of propofol $u(t)$ $[\mu g \ / \ min]$, while k_{10}, k_{12}, k_{13}, k_{21} and k_{31} are first order rate constants between the compartments. Between one and twelve years of age, these

2

parameters are estimated as $V_1 = 458.4*\text{weight}$, $k_{10} = 0.1527*\text{weight}^{-0.3}$, $k_{12} = 0.114$, $k_{13} = 0.0419$, $k_{21} = 0.055$, $k_{31} = 0.0033$ [1], where weight is measured in [kg].

The PD model can take many forms [3], but here we choose a linear single compartment version (the "effect" compartment) with input delay and a nonlinear output mapping. The state variable c_e measures the effect site concentration and its evolution is governed by

$$\dot{c}_e(t) = k_d c_p(t - T_d) - k_d c_e(t) \tag{2}$$

where $c_p(t)$ is the plasma concentration from the PK model, T_d is a delay to account for the lag associated with transport of the drug from the plasma to the effect site (the brain) and k_d is the rate constant between plasma and effect site. We use $T_d = 20s$ and $k_d = 40$. The actual anesthetic effect is described by mapping through the saturating Hill equation

$$E = E(c_e) = \frac{c_e^\gamma}{EC_{50}^\gamma + c_e^\gamma} \tag{3}$$

where EC_{50} is the 50% effect concentration and γ is the cooperativity coefficient. This normalized measure of effect E then ranges from $E = 1$ (awake) to $E = 0$ (deep coma). Values of $EC_{50} = 3$ and $\gamma = 2$ are used. Note that (3) may be nonlinear, but it is a (strictly monotonic) rational mapping and not a differential equation.

2.2 Simplifications

The PK/PD model described above is almost a four dimensional linear differential equation, except for the delay between the PK and PD compartments and the nonlinear output mapping. The latter is the easiest to deal with: Because (3) is strictly monotonic, we can map any constraints on the output E through the inverse of (3)

$$c_e(E) = \left(\frac{E}{1 - E}\right)^{\frac{1}{\gamma}} EC_{50} \tag{4}$$

to arrive at constraints on the state variable c_e. In fact, in the maintenance phase the system should be operating near $c_e = EC_{50}$ where (3) is essentially linear; consequently, if the safety constraints are close enough to that operating point then the output mapping can be further simplified to be linear:

$$E(c_e) \approx \frac{\gamma}{4EC_{50}} c_e$$

The time delay in (2) is a little more delicate. One option is to introduce additional state variables to model the delay; for example, additional compartments or a Padé approximation. The PK/PD model described above is suitable only for relatively low frequency input anyway, so these approximations are not unreasonable. A complementary approach builds on the fact that the PK concentrations are not observable in real-time and are independent of the PD state c_e, so we can shift the origin of time for the PK state variables by T_d and effectively shift the time delay all the way backward to the PK input u:

$$\begin{bmatrix} \dot{\tilde{c}}_p(t) \\ \dot{\tilde{c}}_1(t) \\ \dot{\tilde{c}}_2(t) \\ \dot{c}_e(t) \end{bmatrix} = \begin{bmatrix} -(k_{10} + k_{12} + k_{13}) & k_{12} & k_{13} & 0 \\ k_{21} & -k_{21} & 0 & 0 \\ k_{31} & 0 & -k_{31} & 0 \\ k_d & 0 & 0 & -k_d \end{bmatrix} \begin{bmatrix} \tilde{c}_p(t) \\ \tilde{c}_1(t) \\ \tilde{c}_2(t) \\ c_e(t) \end{bmatrix} + \begin{bmatrix} 1/V_1 \\ 0 \\ 0 \\ 0 \end{bmatrix} u(t - T_d) \tag{5}$$

where $\tilde{c}_p(t) = c_p(t + T_d)$, $\tilde{c}_1(t) = c_1(t + T_d)$, and $\tilde{c}_2(t) = c_2(t + T_d)$ from (1). Whether this time shifted model is useful will depend on whether the proposed analysis procedure has some way of dealing with input delay; for example, it is much easier in MATLAB to simulate with an input delay (5) than with a delay in internal state feedback (1)–(2).

2.3 Complexifications

Traditional PK/PD modeling of anesthesia would stop at this point; however, if closed-loop control is desired then a model of the sensor and actuator which close the loop may be required. As one example, a second order linear model of the WAV_{CNS} Monitor is derived in [14].

In a surgical setting, it may not be feasible to estimate the state or adjust the input in continuous time, so it is natural to consider a discrete time model. Such a model is easily derived from the continuous time models given above, but ignores the evolution (and hence safety) of the system between sampling times. A more rigorous safety analysis can be accomplished through a sampled data formulation, in which the continuous time evolution of the state is tracked but the feedback loop is only closed at discrete times (see [9] and the citations therein).

The deterministic model discussed above ignores several sources of variability which could be included. Depending on the degree of analgesia that has been achieved, surgical stimuli may perturb the state evolution. Because they are generated by the surgeon's actions, these effects are probably best modeled as an additional input to the PD component: either as a known function of time (for predictable stimuli) or as a disturbance (treated in a worst case fashion). More generally, it is expected that there will be significant variability of the model parameters between patients and even over time for a single patient. Ideally we would identify these parameters for each patient at the beginning of the surgery, but for medical reasons it is unlikely that a "model identification" phase will be permitted between induction and maintenance anytime soon. An obvious approach to handle this uncertainty is then to include stochastic parameter values or disturbance inputs in the model.

3 Safety Verification Challenge

The goal of verification in this anesthesia model is relatively straightforward: under constraints on the input, compute the maximal controlled invariant set (or viability kernel) which lies within a "safe" portion of the state space (the constraint set); in other words, the subset of states in the constraint set for which there exist input signals which lead to trajectories which remain inside the constraint set over the time horizon. Note that the maximal controlled invariant set is complementary to the reachable set, although in many cases it is possible to use tools for one to compute the other (for example, [8]). Ideally this set would be computed over an infinite horizon, although in practical terms the length of a surgery is finite and can be estimated in advance.

For the PK/PD model discussed above, a typical set of constraints would be

$$c_p \in [1, 6]$$
$$c_1 \in [0, 10]$$
$$c_2 \in [0, 10]$$
$$c_e \in [1, 8]$$

all measured in $[\mu g \ / \ \ell]$, and the input constraint is $u \in [0, 200] \ [\mu g \ / \ kg \ / \ min]$.

As an example, an underapproximation of the viability kernel for the PK version of this model (1) in discrete time with a three dimensional Padé approximation of the input delay was calculated in [8] for a 90 minute surgery on an eleven year-old child weighing 35kg. The constraint set was approximated by an ellipsoid contained within the box constraints described above, and the underapproximation represented with support vectors. Code implementing this calculation can be found at `http://www.ece.ubc.ca/~jmaidens/ viability_supplement.zip`.

For low dimensional models with only a single control input, the controlled invariant set can be estimated relatively easily through simulation: for every state in a sampling of the constraint set, search for an input signal (eg: through model predictive control) to see if one exists which maintains safety. This brute force approach becomes less tractable as the state space dimension rises; for example, the model analyzed in [8] was six dimensional once the Padé approximation was included, and a sampling of a six dimensional space is coarse, large, or both. It is even less feasible if disturbance inputs are included in the model, in which case
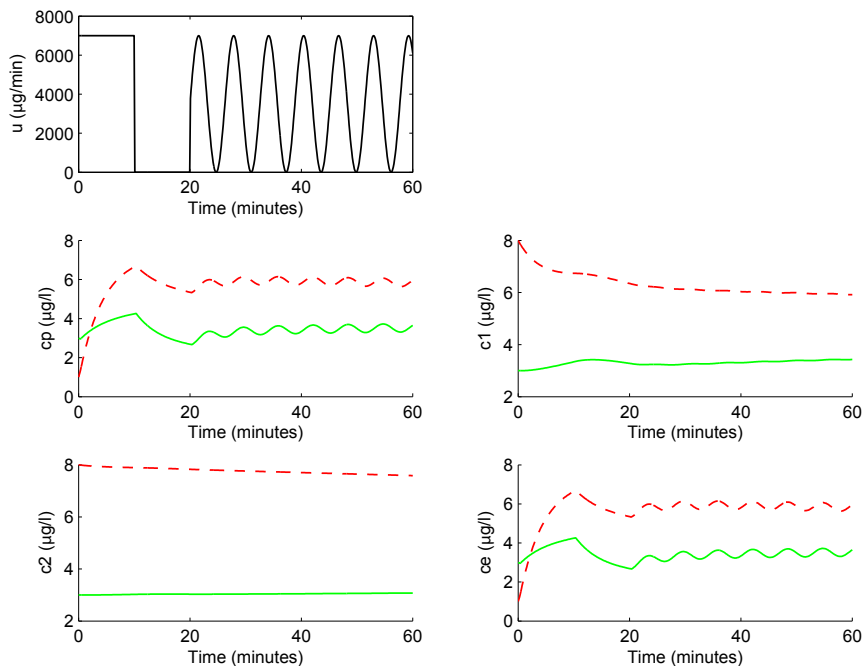
Figure 2: Simulated trajectories of (5) for the input u shown at the top. Dashed red trajectories exit the safety constraints defined in section 3, while solid green trajectories satisfy the constraints.

the goal is computation of the maximal robust controlled invariant set (or discriminating kernel): the set of states for which there exist input signals that for all disturbances lead to trajectories which remain inside the constraint set over the time horizon. In general, the combination of control and disturbance inputs leads to a nonconvex optimization problem even for linear models, and hence the search for suitable control signals for even a single initial state becomes computationally challenging.

4 Simulation Code

We provide MATLAB code that simulates (for the eleven year old child weighing 35kg) trajectories of either (1) or (5) for sample initial states and a specific open-loop propofol administration rate $u(t)$. Trajectories from each initial state are numerically integrated up to a fixed time horizon. Figure 2 shows the concentration response of (5) to a time-varying input $u(t)$ for initial states $\begin{bmatrix} \tilde{c}_p(t) & \tilde{c}_1(t) & \tilde{c}_2(t) & c_e(t) \end{bmatrix}^T = \begin{bmatrix} 1 & 8 & 8 & 1 \end{bmatrix}^T$ in dashed red and $\begin{bmatrix} 3 & 3 & 3 & 3 \end{bmatrix}^T$ in solid green. The code also produces a three dimensional phase space plot (not shown). Trajectories are automatically displayed with dashed red or solid green curves depending on whether they exit or stay inside the constraint set respectively.

5 Conclusions

We provided a four dimensional PK/PD model of the effect of propofol infusion on depth of hypnosis during the anesthesia maintenance phase of a surgery, as well as discussing various ways of simplifying or complexifying the model. The safety verification problem is to determine the maximal controlled invariant subset of a box-like constraint set. Citations for the model and an approximation of the maximal controlled

invariant subset were given.

We are continuing to investigate new models and safety verification algorithms with the goal of developing and proving the safety of a closed-loop anesthesia delivery system.

6   Acknowledgements

References

[1] A. Absalom and G. Kenny. Paedfusor pharmacokinetic data set. *British J. Anaesthesia*, 95(1):110–110, 2005.

[2] J. Bruhn, P.S. Myles, R. Sneyd, and M.M.R.F. Struys. Depth of anaesthesia monitoring: What's available, what's validated and what's next? *British J. Anaesthesia*, 97(1):85–94, 2006.

[3] Hartmut Derendorf and Bernd Meibohm. Modeling of pharmacokinetic/pharmacodynamic (PK/PD) relationships: Concepts and perspectives. *Pharmaceutical Research*, 16(2):176–185, 1999.

[4] Karen B. Domino, Karen L. Posner, Robert A. Caplan, and Frederick W. Cheney. Awareness during anesthesia: A closed claims analysis. *Anesthesiology*, 90(4):1053–1061, 1999.

[5] Michael P. Dosch. On being aware: Patient recall of intraoperative events. *AANA Journal*, 56:238–245, 1988.

[6] Bret Fulton and Eugene M. Sorkin. Propofol. *Drugs*, 50(4):636–457, 1995.

[7] Jin-Oh Hahn, Guy A. Dumont, and J. Mark Ansermino. Closed-loop anesthetic drug concentration estimation using clinical-effect feedback. *IEEE Trans. Biomedical Engineering*, 58(1):3–6, 2011.

[8] John N. Maidens, Shahab Kaynama, Ian M. Mitchell, Meeko M. K. Oishi, and Guy A. Dumont. Lagrangian methods for approximating the viability kernel in high-dimensional systems. *Automatica*, 49(7):2017–2029, July 2013.

[9] Ian M. Mitchell, Shahab Kaynama, Mo Chen, and Meeko Oishi. Safety preserving control synthesis for sampled data systems. *Nonlinear Analysis: Hybrid Systems*, 10:63–82, 2013.

[10] Terri G. Monk, Vikas Saini, B. Craig Weldon, and Jeffrey C. Sigl. Anesthetic management and one-year mortality after noncardiac surgery. *Anesthesia & Analgesia*, 100(1):4–10, 2005.

[11] G.A. Osborne, A.K. Bacon, William Ben Runciman, and Stephen Clayton Helps. Crisis management during anaesthesia: awareness and anaesthesia. *Quality and Safety in Health Care*, 14(3):e16–e16, 2005.

[12] R. Sandin and O. Nordström. Awareness during total IV anaesthesia. *British J. Anaesthesia*, 71(6):782–787, 1993.

[13] Jürgen Schüttler and Harald Ihmsen. Population pharmacokinetics of propofol: a multicenter study. *Anesthesiology*, 92(3):727–738, 2000.

[14] Tatjana Zikov, Stéphane Bibian, Guy A. Dumont, Mihai Huzmezan, and Craig R. Ries. Quantifying cortical activity during general anesthesia using wavelet analysis. *IEEE Trans. Biomedical Engineering*, 53(4):617–632, 2006.