

A simulation-based comparative analysis of PID and LQG control for closed-loop anesthesia delivery

Sourish Chakravarty^{1,*,**}, Ayan S. Waite^{1,*,**},
John H. Abel^{*,**}, Emery N. Brown^{*,**,***}

* *Picower Institute for Learning and Memory, Massachusetts Institute
of Technology (MIT), Cambridge, MA*

** *Department of Anesthesia, Critical Care and Pain Medicine at
Massachusetts General Hospital, Boston, MA*

*** *Institute for Medical Engineering and Science, MIT, Cambridge,
MA*

Abstract: Closed loop anesthesia delivery (CLAD) systems can help anesthesiologists efficiently achieve and maintain desired anesthetic depth over an extended period of time. A typical CLAD system would use an anesthetic marker, calculated from physiological signals, as real-time feedback to adjust anesthetic dosage towards achieving a desired set-point of the marker. Since control strategies for CLAD vary across the systems reported in recent literature, a comparative analysis of common control strategies can be useful. For a nonlinear plant model based on well-established models of compartmental pharmacokinetics and sigmoid-Emax pharmacodynamics, we numerically analyze the set-point tracking performance of three output-feedback linear control strategies: proportional-integral-derivative (PID) control, linear quadratic Gaussian (LQG) control, and an LQG with integral action (ILQG). Specifically, we numerically simulate multiple CLAD sessions for the scenario where the plant model parameters are unavailable for a patient and the controller is designed based on a nominal model and controller gains are held constant throughout a session. Based on the numerical analyses performed here, conditioned on our choice of model and controllers, we infer that in terms of accuracy and bias PID control performs better than ILQG which in turn performs better than LQG. In the case of noisy observations, ILQG can be tuned to provide a smoother infusion rate while achieving comparable steady state response with respect to PID. The numerical analysis framework and findings reported here can help CLAD developers in their choice of control strategies. This paper may also serve as a tutorial paper for teaching control theory for CLAD.

Keywords: Closed-loop control, physiological systems, LQG, PID, integral action, anti-windup, compartmental models, sigmoid-Emax, quantitative systems pharmacology

1. INTRODUCTION

More than 230 million patients undergo major surgeries world-wide under anesthesia (Weiser et al. (2008)). During major surgeries performed under general anesthesia (GA) (Brown et al. (2010)), conventionally, an anesthesiologist would monitor relevant physiological signals that are known to be associated with anesthetic depth, and manually adjust the drug dose. For example, changes in oscillatory patterns in the scalp electroencephalogram (EEG) has been shown to be correlated with behaviorally-defined changes in unconsciousness levels during GA (Purdon et al. (2013)). Such structured drug-dependent changes in the EEG signals have allowed developments of automated EEG-based closed-loop anesthesia delivery (CLAD) sys-

tems that can provide precise control of brain states under anesthesia (Bickford (1950); Schwilden et al. (1989); Absalom et al. (2002); Dumont et al. (2009); Shanechi et al. (2013); Puri et al. (2016); Yang et al. (2019)). Analogous systems are also available for blood pressure (BP)-based CLAD (Ngan Kee et al. (2007)).

A typical CLAD system works in a cyclic mode at a fast rate ($\sim 0.1 s^{-1}$) where in each cycle a computer calculates a scalar anesthetic marker from real-time physiological signals (say, EEG or BP) and uses a feedback control strategy to determine the adjustment of drug-dosage for the next cycle. As this iterative procedure continues over time, the anesthetic marker ideally would approach the user-prescribed set-point. The utility of such *autopilot* systems are in personalized medicine as they can aid anesthesiologists achieve patient-specific precise drug-dosing efficiently, even under resource constrained environments (Dumont and Ansermino (2013); Absalom et al. (2011)).

¹ S.C. and A.S.W. contributed equally.

² This work was partially supported by NIH Award P01-GM118629 (to E.N.B.), by funds from Massachusetts General Hospital (to E.N.B.), by funds from the Picower Institute for Learning and Memory (to E.N.B. and S.C.), and by NIH/NIA Award F32-AG064886 (to J.H.A.).

The CLAD systems reported in prior research, vary with respect to different features, e.g., regime of anesthesia, anesthetic drug, signal modality, marker definition, control strategy, and plant model, just to name a few. For example, the classical output-feedback proportional-integral-derivative (PID) control strategy and its variants have been found to be useful in CLAD applications in human studies (Dumont et al. (2009); Puri et al. (2016); Westover et al. (2015)). The more recent development of using an estimated-state feedback-based optimal control strategy via linear quadratic Gaussian (LQG) regulation have yielded promising results for maintenance of medically-induced coma in animal models (Shanechi et al. (2013)). The success of both PID and LQG, as general control paradigms for CLAD, motivated us to compare them through a numerical experiment in this work.

In this numerical study of CLAD, we consider a specific nonlinear state-space model as the plant model. Here the state dynamics, driven by the drug infusion, is described by a four-compartment linear pharmacokinetic (PK) model. The observation dynamics, relating the anesthetic marker and the dynamic state in the effect compartment, is assumed to be described by a *sigmoid-Emax* pharmacodynamic (PD) model (Meibohm and Derendorf (1997)). We choose realistic values of model parameters based on a pair of well-established models in literature for propofol-based PK and corresponding EEG-based PD (Schneider et al. (1998, 1999)). We linearize this nonlinear plant model around a steady state value and use the resulting linear state-space model to determine control parameters. Since LQG is known to have robustness issues (Doyle (1978)) and model-misspecification can lead to steady state errors during set-point tracking, we also consider LQG with an additional integral action (ILQG) on the output error. To address practical constraints, we impose actuator saturation and zero-order holds explicitly. We consider different operating conditions characterized by observation noise and pump-update rates, as well as control parameters. For fixed target trajectory and a given operating condition, we simulate CLAD for multiple subjects to mimic a practical scenario when only a nominal plant model, instead of the true one, is available for controller design.

In the ensuing discussion, Sec. 2 presents the core equations for the plant and the control strategies. Sec. 3 presents our numerical simulation design, results and associated discussion. Finally, Sec. 4 summarizes our work.

2. METHODS

2.1 Nonlinear plant model

We model the pharmacokinetics (PK) using a first-order ordinary differential equation as,

$$\dot{x}(t) = Ax(t) + Bu(t) \quad (1)$$

where, \dot{x} indicates dx/dt . Transition matrix, A , and input scaling matrix, B , are prescribed based on a four compartment mammillary model assumption as proposed by Shafer and Gregg (1992), where the effect site is augmented to a three-compartment model as an additional fourth compartment (much smaller in volume relative to the central compartment). The mass of drug in all the compartments at any instant of time, t , is indicated by

$x(t) = [x_1(t), x_2(t), x_3(t), x_4(t)]^T$, $u(t)$ denotes scalar infusion rate function, $B = [1, 0, 0, 0]^T$ and transition matrix is given by,

$$A = \begin{bmatrix} -(k_{10} + k_{12} + k_{13} + k_{14}) & k_{21} & k_{31} & k_{41} \\ k_{12} & -k_{21} & 0 & 0 \\ k_{13} & 0 & -k_{31} & 0 \\ k_{14} & 0 & 0 & -k_{41} \end{bmatrix} \quad (2)$$

Note, that $u(t) \geq 0$ and $x_i(t) \geq 0 \forall i \in \{1, 2, 3, 4\}$ and $\forall t \geq 0$. The PK model described above is illustrated in Fig. 1. We model the pharmacodynamic (PD)³ relationship between a scalar-valued dose-dependent function, say $y(t)$, of the physiological signal of interest and the effect site amount, $x_4(t)$, using the following sigmoid-Emax model,

$$y(t) = E_0 + E_{max} \left(\frac{(x_4(t)/x_{50})^\gamma}{1 + (x_4(t)/x_{50})^\gamma} \right) \quad (3)$$

where, $E_0, E_{max} > 0$, $x_{50} > 0$, and $\gamma > 0$ are scalar parameters. In the ensuing discussion, the symbol ' (t) ', in the notations indicating dynamic variables, will be omitted for brevity. In the computational simulation studies that




Fig. 1. Schematic of a 4-compartment PK model

we perform here, the core plant model taking input $u(t)$ and generating output $y(t)$ is given by Eqs. (1) and (3). The nonlinear plant model is characterized by the subject-specific PK-PD parameters $\Theta_P = \{A, x_{50}, \gamma, E_{max}, E_0\}$.

2.2 Linear control strategy with output feedback

Steady state relationships: For linear control design for set-point tracking, we determine relationships among y , x and u for a steady state condition. Imposing the steady state condition, $\dot{x} = 0$ in Eq. (1), and denoting the constant values at steady state as y_{ss} , x_{ss} and u_{ss} , we obtain

$$x_{ss} = -A^{-1}Bu_{ss}, \quad (4)$$

$$x_{4,ss} = [0, 0, 0, 1]x_{ss}. \quad (5)$$

y_{ss} can be calculated by substituting $x_{4,ss}$ in Eq. (3). When either y_{ss} or u_{ss} is specified, one can use the aforementioned algebraic relationships to calculate the rest of the steady state quantities.

Linearizing about a steady state: The relationship between the y and u , via Eqs. (1) and (3), is nonlinear. To incorporate a linear feedback control framework to regulate y around a steady state value, characterized by

³ Conventionally, k_{14} and k_{41} are treated as PD parameters (Schneider et al. (1999)), but here, for convenience, we consider them as PK model parameters so as to maintain separate parameter sets for the state and the observation equations.

the tuple (y_{ss}, x_{ss}, u_{ss}) , we linearize Eq. (3) about this steady state,

$$y - y_{ss} = \left(\frac{dy}{dx_4} \right)_{x_{4,ss}} (x_4 - x_{4,ss}) + \text{higher order terms} \quad (6)$$

Ignoring the higher-order terms, we obtain a new set of equations in terms of dynamic quantities that denote perturbations with respect to respective steady state values,

$$\dot{\tilde{x}} = A\tilde{x} + B\tilde{u} \quad (7)$$

$$\tilde{y} = C\tilde{x} \quad (8)$$

where, $\tilde{y} \equiv y - y_{ss}$, $\tilde{x} \equiv x - x_{ss}$, $\tilde{u} \equiv u - u_{ss}$, and

$$C = \left[0 \ 0 \ 0 \ \left(\frac{dy}{dx_4} \right)_{x_{4,ss}} \right] \quad (9)$$

$$\text{where, } \left(\frac{dy}{dx_4} \right)_{x_{4,ss}} = \frac{E_{max}}{x_{50}} \frac{\gamma(x_{4,ss}/x_{50})^{\gamma-1}}{(1 + (x_{4,ss}/x_{50})^\gamma)^2} \quad (10)$$

PID control: To achieve set-point tracking of y about a constant set-point y_{sp} , we consider linear system of equations in Eqs. (7) and (8) where linearization is achieved around a user-prescribed state (y_o, x_o, u_o) . Using PID control gains, K_P , K_I and K_D , tuned based on this linear system, we implement the following control strategy (Aström and Murray (2010)),

$$u_{pid} = u_o + K_P(y - y_{sp}) + K_I \int (y - y_{sp}) dt + K_D \dot{y} \quad (11)$$

where, y_{sp} denotes the user-prescribed set-point which may or may not be equal to y_o , but should be in the neighborhood around y_o where the linear dynamics regime can still be assumed. The parameter set for the PID is given by $\Theta_{PID} = \{K_P, K_I, K_D, A, x_{50}, \gamma, E_{max}, E_0, u_o\}$. This parameter set includes controller gains (K_P, K_I, K_D) as well as the parameters of the nominal linear model that will depend both on the nonlinear plant parameters ($A, x_{50}, \gamma, E_{max}, E_0$) as well as the state (u_o) at which the linearization occurs.

LQG control: To achieve constant set-point tracking when using LQG strategy, we reformulate the problem in the form of,

$$\dot{x} = A(x - x_{sp}) + B(u - u_{sp}) + w \quad (12)$$

$$y - y_{sp} = C_o(x - x_{sp}) + v \quad (13)$$

where, w and v are continuous-time zero-mean Gaussian white noise processes characterized by positive-definite covariance matrix W and positive scalar V , respectively (Crassidis and Junkins (2011)). The C_o is determined from Eq. (10) using the linearization state (y_o, x_o, u_o) prescribed in terms of either y_o or u_o (see Sec. 2.2.1). We assume that during a session the controller parameters that depend on (y_o, x_o, u_o) will be kept constant, but the set-point y_{sp} may vary around y_o in a session. With the problem setup in Eqs. (12) and (13), we can pose the LQG optimal control problem as,

$$\{u^*(t)\} = \operatorname{argmin} E[J(\{u(t)\})] \quad (14)$$

$$J \equiv \int_0^\infty (x - x_{sp})^T Q (x - x_{sp}) + (u - u_{sp})^T R (u - u_{sp}) dt \quad (15)$$




Fig. 2. Schematic block diagram of CLAD simulation setup

such that Eq. (12) holds at every time-point, and $E[\cdot]$ indicates an expectation operation. Assuming the system (A, B) is stabilizable and (A, C_o) is detectable and matrices Q and R are positive semi-definite and positive definite, respectively, the solution of the aforementioned optimal control problem leads to a steady state linear quadratic regulator with control action u_{lqq} given by,

$$u_{lqq} = u_{sp} - K_C(\hat{x} - x_{sp}) \quad (16)$$

where, \hat{x} is the output of a steady state Kalman filter,

$$\dot{\hat{x}} = (A - K_E C_o)(\hat{x} - x_{sp}) + B(u - u_{sp}) + K_E(y - y_{sp}) \quad (17)$$

Controller gain matrix K_C depends on A, B, Q and R . Estimator gain matrix K_E depends on A, C_o, W and V . The formulae to calculate K_C and K_E can be found in literature (Athans (1971); Crassidis and Junkins (2011)). The parameter set for the PID is given by $\Theta_{LQG} = \{Q, R, W, V, A, x_{50}, \gamma, E_{max}, E_0, u_o\}$.

ILQG control: To minimize any steady state bias that can ensue due to model mismatch between the plant and the controller, we add an integral action to minimize the steady state errors. To do this we consider an additional state z driven by the tracking error,

$$\dot{z} = y_{sp} - y, \text{ such that,} \quad (18)$$

$$u_{ilqq} = u_{lqq} + K_I z \quad (19)$$

where, $K_I > 0$ is an additional user-prescribed gain. The parameter set for the ILQG is given by $\Theta_{ILQG} = \{K_I, Q, R, W, V, A, x_{50}, \gamma, E_{max}, E_0, u_o\}$. Since we are interested in how the LQG performance changes with regard to the current CLAD problem in the presence and absence of an integral action, we explicitly specify K_I . This is in the same spirit as the LQG servocontroller designs that can be simulated using the *lqgtrack* function in MATLAB (2019). Alternatively, one can also design K_I implicitly within the LQG framework (Grimble (1979)).

Additional constraints for CLAD simulation: In the context of CLAD, practical considerations require $u \geq u_{min} = 0$ and $u \leq u_{max}$; the latter constraint can be due to constraints on the physiological system and/or the mechanical actuation system. To mimic this in our simulation, we allow output of the controller, say u_c (such that $c \in \{pid, lqq, ilqq\}$), goes through a saturation block prescribed as,

$$\begin{aligned} u &= u_c, \text{ if } 0 < u_c < u_{max}, \\ u &= 0, \text{ if } u_c \leq 0, \\ u &= u_{max}, \text{ if } u_c \geq u_{max}, \end{aligned} \quad (20)$$

to yield the control signal u that is input to the drug delivery system. Within the drug delivery system, the modulation of infusion rates for anesthetic fluids is a mechanical phenomenon which can have a short transient before the infusion rate stabilizes around the desired flow rate. To simulate this condition without modeling the drug delivery dynamics explicitly we impose a *zero-order hold* (zoh) condition on u , such that the flow rate u is updated

at an interval T_u , and held constant until the next update. A similar zoh condition is applied on the output side assuming that it is sampled at rate T_y which may or may not be equal to T_u . To mimic the noise in the observations, we also add Gaussian white noise with variance V and sampled at every interval T_y , to the observation. If $T_y = T_u$ holds, then one can alternatively choose to model plant and controller dynamics in discrete-time governed by incremental dynamic equations. However, this would require the increment time-interval to enter the system matrices that are used to calculate controller gains. Furthermore, modeling the plant in continuous time is conceptually closer to the physiological phenomena governing the true PK-PD phenomena. Therefore, in this work, we try to keep the plant and control blocks independent of any prescribed time-interval, but add the aforementioned constraints, via user-prescribed values of u_{max} , T_u and T_y , externally in the complete simulation block diagram (Fig. 2).

Table 1. Performance metrics for PID (Fig.3(a-d)), LQG with $\rho = 10^{-13}$ (Fig.4(a-d)), LQG with $\rho = 10^{-6}$ (Fig.4(e-h)), ILQG with $K_I = 10$, $\rho = 10^{-13}$ (Fig.5(a-d)), LQG with $K_I = 1$, $\rho = 10^{-6}$ (Fig.5(e-h)). Median (min, max) values of performance metrics from 44 CLAD simulations per case.

| Fig. | $60T_u$ | V | ϵ_I (%) | ϵ_B (%) | u_{avg} |
|------|---------|-----------|------------------|------------------|--------------|
| 3(a) | 5s | 10^{-6} | 0.4(0.4,0.4) | 0.1(0.1,0.1) | 5.6(4.0,7.3) |
| 3(b) | 5s | 10^{-4} | 3.2(3.2,3.3) | 0.2(0.1,0.2) | 5.6(4.0,7.3) |
| 3(c) | 10s | 10^{-6} | 0.3(0.3,0.4) | 0.1(0.1,0.1) | 5.7(4.1,7.3) |
| 3(d) | 10s | 10^{-4} | 2.4(2.3,2.4) | 0.2(0.2,0.3) | 5.6(4.0,7.3) |
| 4(a) | 5s | 10^{-6} | 2.6(1.2,6.7) | -2.5(-6.6,1.6) | 5.4(4.2,6.6) |
| 4(b) | 5s | 10^{-4} | 4.2(3.7,7.0) | -2.7(-6.8,1.7) | 5.4(4.2,6.6) |
| 4(c) | 10s | 10^{-6} | 2.6(1.5,6.5) | -2.4(-6.4,1.4) | 5.4(4.2,6.6) |
| 4(d) | 10s | 10^{-4} | 3.6(3.0,6.8) | -2.6(-6.6,1.9) | 5.4(4.2,6.6) |
| 4(e) | 5s | 10^{-6} | 16.9(3.4,47.2) | -16.9(-47.2,4.7) | 4.6(4.6,4.6) |
| 4(f) | 5s | 10^{-4} | 14.8(5.5,42.8) | -14.8(-42.8,7.0) | 4.6(4.6,4.6) |
| 4(g) | 10s | 10^{-6} | 17.2(3.4,47.5) | -17.0(-47.5,4.3) | 4.6(4.6,4.6) |
| 4(h) | 10s | 10^{-4} | 15.2(4.7,44.7) | -15.2(-44.7,7.0) | 4.6(4.6,4.6) |
| 5(a) | 5s | 10^{-6} | 1.0(0.9,1.2) | 0.5(0.2,0.9) | 5.6(4.1,7.2) |
| 5(b) | 5s | 10^{-4} | 3.6(3.5,3.7) | 0.3(-0.3,1.2) | 5.6(4.1,7.2) |
| 5(c) | 10s | 10^{-6} | 1.0(1.0,1.2) | 0.5(0.1,1.0) | 5.7(4.1,7.2) |
| 5(d) | 10s | 10^{-4} | 2.9 (2.8,3.0) | 0.5(-0.1,1.3) | 5.7(4.1,7.2) |
| 5(e) | 5s | 10^{-6} | 1.2(0.9,1.6) | 0.7(0.5,0.9) | 5.6(4.1,7.3) |
| 5(f) | 5s | 10^{-4} | 3.7(3.5,3.9) | 0.6(0.3,0.9) | 5.6(4.1,7.3) |
| 5(g) | 10s | 10^{-6} | 1.5(1.0,1.9) | 0.8(0.6,1.0) | 5.7(4.1,7.3) |
| 5(h) | 10s | 10^{-4} | 2.8(2.6,3.1) | 0.8(0.5,1.1) | 5.7(4.1,7.3) |

3. RESULTS

3.1 CLAD simulation setup

Plant model parameters: As candidate plant model parameters we used the weight, height, age and sex information available for $N_p = 44$ patients reported in prior work (Dumont et al., 2009, Table A1). For the PK-PD model, we used the model parameters reported for 3-compartment propofol pharmacokinetics (Schnider et al., 1998, Table 2) and corresponding EEG-based pharmacodynamics (Schnider et al., 1999, Table 3), where for the latter we use only the age-adjusted parameters for the ascending sigmoidal function in the bi-phasic response (Schnider et al., 1999, Fig. 1). We set $E_0 = 0$ and $E_{max} = 1$ since they reflect a constant shift and scaling. Per suggestion by Shafer and Gregg (1992), we use $k_{41} = 10^{-5}k_{14}$.

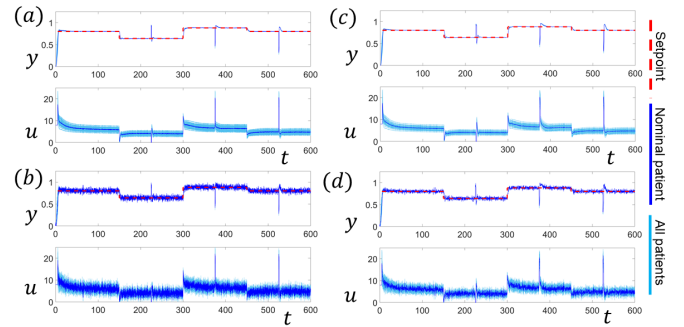


Fig. 3. PID: Simulated trajectories for 45 patients (including the nominal subject, highlighted in solid blue) with 4 sets of $(60T_u, V)$: (a) $(5s, 10^{-6})$, (b) $(5s, 10^{-4})$, (c) $(10s, 10^{-6})$, (d) $(10s, 10^{-4})$. Also see Table 1.

Simulation parameters: Each simulation mimics a CLAD session of duration, $T_{max} = 600min$. In all simulations, we maintain the maximum allowable infusion rate for a patient weighing W (kg) as $u_{max} = 1mg/kg/min \times W$. We set $T_y = T_u$ for simplicity. To mimic a practical situation where the CLAD system is activated after a patient is anesthetized, we run the system in an open-loop mode for $T_{ind} = 5min$ with a fixed infusion rate of $u = u_{max}/10$ prior to activating the closed-loop mode. For the closed-loop mode, the following target profile is maintained: $y_{sp} = 0.8$ for $t \in (T_{ind}, T_{max}/4]$, then $y_{sp} = 0.64$ for $t \in (T_{max}(1/4), T_{max}(2/4)]$, followed by $y_{sp} = 0.88$ for $t \in (T_{max}(2/4), T_{max}(3/4)]$, and finally back to $y_{sp} = 0.9$ for $t \in (T_{max}(3/4), T_{max}]$. We also simulate disturbance profiles at $t = 225min$, $375min$ and $575min$ where the effect-site dynamic state x_4 is scaled by a factor 2, 0.5 and 0.5, respectively, for a brief period ($= 10T_u$).

Controller parameters: For any given simulation run we linearize the PK-PD model, available to the controller, around a steady state value of $u_o = u_{max}/10$. For LQG and ILQG, we set Q to be a diagonal matrix such that $diag(Q) = [1/(x_{1,max} - x_o)^2, 1/(x_{2,max} - x_o)^2, 1/(x_{3,max} - x_o)^2, \rho/(x_{4,max} - x_o)^2]$, where x_{max} is calculated using Eq. (4) and u_{max} . We set scalar $R = 1/(u_{max} - u_o)^2$. For the process noise variance W we choose a diagonal matrix such that $diag(W) = 10^{-6}V[1, 1, 1, 1]$. For the PID controller, we determine K_P , K_I and K_D for linear system parameters A , B and C_o by invoking the Matlab function, *pidtune* with its default setting of balance between performance and robustness. Since we will be simulating with noisy observations, we incorporate a derivative filter such that transfer function contribution to the derivative component of the PID transfer function is $K_D s / (T_f s + 1)$ where we set $T_f = 5T_u$. Furthermore, since we have actuation saturation, to avoid integral error build-up, we employ the following anti-windup strategy (Aström and Murray (2010)): when the output of the controller, u_c , falls outside the range $[u_{min}, u_{max}]$ we enforce a zero input into the integrator blocks of both the PID and the ILQG.

3.2 CLAD simulations

For each control strategy we simulate for every patient setpoint tracking with $T_u \in \{5/60min, 10/60min\}$, and $V = \{10^{-6}, 10^{-4}\}$ under conditions described in the following subsection. For each combination of (T_u, V) , we simulate:

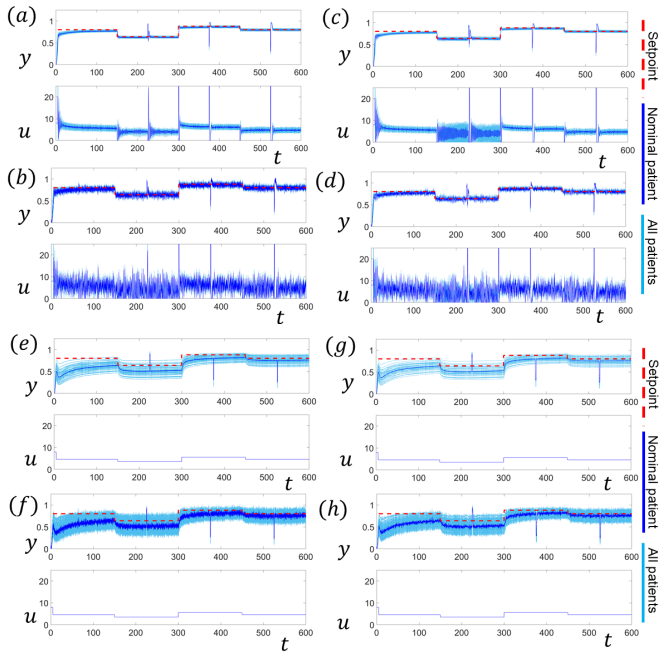


Fig. 4. **LQG**: Simulated trajectories for 45 patients (including the nominal subject, highlighted in solid blue) with 4 sets of $(60T_u, V)$: (a,e) $(5 s, 10^{-6})$, (b,f) $(5 s, 10^{-4})$, (c,g) $(10 s, 10^{-6})$, (d,h) $(10 s, 10^{-4})$. (a-d) $\rho = 10^{13}$, $K_I = 0$, (e-h) $\rho = 10^6$, $K_I = 0$. See Table 1.

a PID-based CLAD (Fig. 3(a-d)), an ILQG-based CLAD with two control designs [$\rho = 10^{-13}$, $K_I = 10$] (Fig. 5(a-d)) and [$\rho = 10^{-6}$, $K_I = 1$] (Fig. 5(e-h)), and an LQG-based CLAD simulation under same conditions, but with $K_I = 0$ (Fig. 4(a-d) and Fig. 4(e-h)). For each of these operating conditions, we simulate CLAD sessions for N_p subjects such that the nominal plant and true plant are different. In these simulations, every subject is characterized by a respective Θ_P value, whereas controller parameters are characterized by a fixed set of parameters Θ_C determined based on a nominal model of the plant, $\Theta_{P_{nom}}$. To calculate the $\Theta_{P_{nom}}$, we consider a nominal subject (age = 35.5 years, height = 177.5 cm, weight = 79.5 kg, male) using the median of each of these features. The parameters, for this nominal patient are: $k_{10} = 0.4282 \text{ min}^{-1}$, $k_{12} = 0.4005 \text{ min}^{-1}$, $k_{13} = 0.1958 \text{ min}^{-1}$, $k_{21} = 0.0664 \text{ min}^{-1}$, $k_{31} = 0.0035 \text{ min}^{-1}$, $k_{41} = 0.4560 \text{ min}^{-1}$, $x_{50} = 6.8 \times 10^{-5} \text{ mg}$, $\gamma = 3.0965$. All simulations are performed using MATLAB (2019)/Simulink. To assess the performance across controller designs and across multiple simulation conditions, we use the following inaccuracy and bias metrics (Dumont et al. (2009)),

$$\text{Inaccuracy, } \epsilon_I = \text{median}\{100|(y_k - y_{sp})/y_{sp}|\} \quad (21)$$

$$\text{Bias, } \epsilon_B = \text{median}\{100(y_k - y_{sp})/y_{sp}\} \quad (22)$$

where, y_k denotes the observation at the k -th time-point. We also report the average infusion, u_{avg} , calculated across each simulated session.

3.3 Discussion

Our final assessment of controller performance across the simulated CLAD sessions is based on both the numerical outputs of ϵ_I , ϵ_B , u_{avg} (Table 1) and visual inspection of output and control trajectory plots (Figs. 3, 4, 5). The inaccuracy and bias metrics for the PID are found to be lower

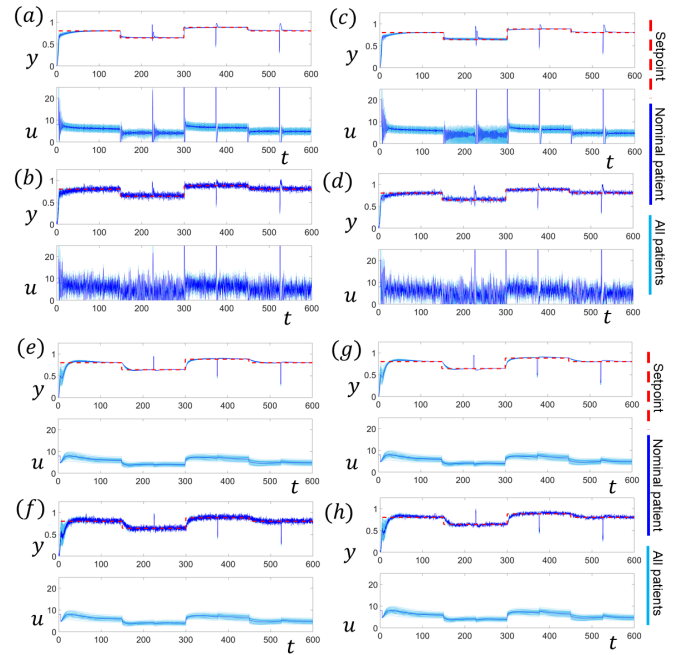


Fig. 5. **ILQG**: Simulated trajectories for 45 patients (including the nominal subject, highlighted in solid blue) with 4 sets of $(60T_u, V)$: (a,e) $(5 s, 10^{-6})$, (b,f) $(5 s, 10^{-4})$, (c,g) $(10 s, 10^{-6})$, (d,h) $(10 s, 10^{-4})$. (a-d) $\rho = 10^{13}$, $K_I = 10$, (e-h) $\rho = 10^6$, $K_I = 1$. See Table 1.

than both ILQG and LQG. The ILQG simulations show better set-point tracking performance relative to LQG, e.g. compare Fig.4(a-d) vs. Fig. 5(a-d) and Fig.4(e-h) vs. Fig. 5(e-h). Since the model mis-specification (difference between true and nominal models) has a significantly adverse effect on the performance metrics for Fig. 4(e-h) this particular LQG design is unacceptable for any practical CLAD implementation. Also, from visual inspection we can infer that PID control leads to faster convergence to set-point relative to LQG or ILQG, but all three control strategies are able to provide stable regulation. The plots in Figs. 3(a-d), 5(a-d), and 5(e-f) indicate that by lowering ρ (the relative weighting of state cost to input cost in the LQG quadratic objective function) we can tune the ILQG to achieve a control signal that is not as sensitive to the observation noise and leads to tracking performance comparable to PID. In all the simulations except for Fig. 4(e-h), the median u_{avg} is found to be almost similar (at least to the first significant digit). A plausible reason for this observation is that in all cases, except Fig. 4(e-h), the control strategies are able to converge on the same average infusion rate that would achieve and track the same y_{sp} for most part of the simulated sessions (as indicated by the median $\epsilon_I \leq 4.2\%$ across all these sessions, which is markedly lower than the same metric in Fig. 4(e-h)). Furthermore, the influence of the transient dynamics of the control signal due to set-point changes and noise become less pronounced in the total infused volume estimate (used to calculate u_{avg}) over such long sessions. In the summary our aforementioned observations reflect the salient differences among the three control strategies. The PID controller is designed to have fast, stable dynamics on the output feedback error. On the other hand, the LQG controller is designed to minimize the quadratic cost

Eq. (15) and a higher penalty on the deviation of state trajectory from x_{sp} (relative to the deviation of control from u_{sp}) can lead to better set-point tracking of y_{sp} by y . For ILQG, the addition of integral action leads to a sub-optimal performance relative to the LQG cost function but with improved steady state set-point tracking relative to LQG, and comparable to PID.

4. CONCLUSION

In this work, we compared three control strategies: PID, LQG, and ILQG for a CLAD application. For the CLAD problem of set-point tracking of output from a single-input single-output system described by Eqs. (1) and (3), we draw the following inference from our analyses. PID can help with faster convergence to set-point when a set-point change occurs. Once a set-point is achieved, PID, ILQG and LQG can each provide stable set-point tracking. LQG requires a very high penalty on state cost relative to actuation cost to provide acceptable set-point tracking performance, but adding an integral action via an ILQG can lead to acceptable tracking performance with lower relative cost. In the presence of observation noise, ILQG can be tuned to achieve temporally smoother actuation with comparable set-point tracking performance relative to a PID controller. Our results indicate that an integral action, incorporated within a linear feedback control strategy, can improve set-point tracking of pharmacodynamic response in such a CLAD setting.

We envisage that this work will be useful to guide controller choices in future CLAD designs, and to numerically assess their performance prior to their usage in animal models or humans. Since the framework involving PK-PD model considered in this study is quite general (Meibohm and Derendorf (1997)), our work can be relevant to CLAD applications beyond the specific anesthetic drug and physiologic signal definitions of the assumed PK-PD model (Schneider et al. (1998, 1999)). Furthermore, this paper may also serve as a tutorial introduction to applications of control theory in CLAD systems.

ACKNOWLEDGEMENTS

We thank anonymous reviewers for helpful feedback and members of the Neuroscience Statistics Research Laboratory at MIT for helpful discussions.

REFERENCES

- Absalom, A.R., De Keyser, R., and Struys, M.M. (2011). Closed loop anesthesia: are we getting close to finding the holy grail? *Anesthesia & Analgesia*, 112(3), 516–518.
- Absalom, A.R., Sutcliffe, N., and Kenny, G.N. (2002). Closed-loop control of anesthesia using bispectral index performance assessment in patients undergoing major orthopedic surgery under combined general and regional anesthesia. *Anesthesiology: The Journal of the American Society of Anesthesiologists*, 96(1), 67–73.
- Åström, K.J. and Murray, R.M. (2010). *Feedback systems: an introduction for scientists and engineers*. Princeton university press.
- Athans, M. (1971). The role and use of the stochastic linear-quadratic-gaussian problem in control system design. *IEEE transactions on automatic control*, 16(6), 529–552.
- Bickford, R.G. (1950). Automatic electroencephalographic control of general anesthesia. *Electroencephalography and Clinical Neurophysiology*, 2(1-4), 93–96.
- Brown, E.N., Lydic, R., and Schiff, N.D. (2010). General anesthesia, sleep, and coma. *New England Journal of Medicine*, 363(27), 2638–2650.
- Crassidis, J.L. and Junkins, J.L. (2011). *Optimal estimation of dynamic systems*. Chapman and Hall/CRC.
- Doyle, J.C. (1978). Guaranteed margins for lqg regulators. *IEEE Transactions on automatic control*, 23(4), 756–757.
- Dumont, G.A. and Ansermino, J.M. (2013). Closed-loop control of anesthesia: a primer for anesthesiologists. *Anesthesia & Analgesia*, 117(5), 1130–1138.
- Dumont, G.A., Martinez, A., and Ansermino, J.M. (2009). Robust control of depth of anesthesia. *International Journal of Adaptive Control and Signal Processing*, 23(5), 435–454.
- Grimble, M. (1979). Design of optimal stochastic regulating systems including integral action. In *Proceedings of the Institution of Electrical Engineers*, volume 126, 841–848. IET.
- MATLAB (2019). *Version: 9.7.0.1216025 (R2019b)*. The MathWorks Inc., Natick, Massachusetts.
- Meibohm, B. and Derendorf, H. (1997). Basic concepts of pharmacokinetic/pharmacodynamic (pk/pd) modelling. *International journal of clinical pharmacology and therapeutics*, 35(10), 401–413.
- Ngan Kee, W., Tam, Y., Khaw, K., Ng, F., Critchley, L., and Karmakar, M. (2007). Closed-loop feedback computer-controlled infusion of phenylephrine for maintaining blood pressure during spinal anaesthesia for caesarean section: a preliminary descriptive study. *Anaesthesia*, 62(12), 1251–1256.
- Purdon, P.L., Pierce, E.T., Mukamel, E.A., Prerau, M.J., Walsh, J.L., Wong, K.F.K., Salazar-Gomez, A.F., Harrell, P.G., Sampson, A.L., Cimenser, A., et al. (2013). Electroencephalogram signatures of loss and recovery of consciousness from propofol. *Proceedings of the National Academy of Sciences*, 110(12), E1142–E1151.
- Puri, G.D., Mathew, P.J., Biswas, I., Dutta, A., Sood, J., Gombar, S., Palta, S., Tsering, M., Gautam, P., Jayant, A., et al. (2016). A multicenter evaluation of a closed-loop anesthesia delivery system: a randomized controlled trial. *Anesthesia & Analgesia*, 122(1), 106–114.
- Schneider, T.W., Minto, C.F., Gambus, P.L., Andresen, C., Goodale, D.B., Shafer, S.L., and Youngs, E.J. (1998). The influence of method of administration and covariates on the pharmacokinetics of propofol in adult volunteers. *Anesthesiology: The Journal of the American Society of Anesthesiologists*, 88(5), 1170–1182.
- Schneider, T.W., Minto, C.F., Shafer, S.L., Gambus, P.L., Andresen, C., Goodale, D.B., and Youngs, E.J. (1999). The influence of age on propofol pharmacodynamics. *Anesthesiology: The Journal of the American Society of Anesthesiologists*, 90(6), 1502–1516.
- Schwilden, H., Stoeckel, H., and Schüttler, J. (1989). Closed-loop feedback control of propofol anaesthesia by quantitative EEG analysis in humans. *BJA: British Journal of Anaesthesia*, 62(3), 290–296.
- Shafer, S.L. and Gregg, K.M. (1992). Algorithms to rapidly achieve and maintain stable drug concentrations at the site of drug effect with a computer-controlled infusion pump. *Journal of pharmacokinetics and biopharmaceutics*, 20(2), 147–169.
- Shanechi, M.M., Chemali, J.J., Liberman, M., Solt, K., and Brown, E.N. (2013). A brain-machine interface for control of medically-induced coma. *PLoS computational biology*, 9(10), e1003284.
- Weiser, T.G., Regenbogen, S.E., Thompson, K.D., Haynes, A.B., Lipsitz, S.R., Berry, W.R., and Gawande, A.A. (2008). An estimation of the global volume of surgery: a modelling strategy based on available data. *The Lancet*, 372(9633), 139–144.
- Westover, M.B., Kim, S.E., Ching, S., Purdon, P.L., and Brown, E.N. (2015). Robust control of burst suppression for medical coma. *Journal of neural engineering*, 12(4), 046004.
- Yang, Y., Lee, J.T., Guidera, J.A., Vlasov, K.Y., Pei, J., Brown, E.N., Solt, K., and Shanechi, M.M. (2019). Developing a personalized closed-loop controller of medically-induced coma in a rodent model. *Journal of neural engineering*, 16(3), 036022.