An Adaptive Brain-Machine Interface Algorithm for Control of Burst Suppression in Medical Coma

Yuxiao Yang[1] and Maryam M. Shanechi[2], *Member, IEEE*

***Abstract*— Burst suppression is an electroencephalogram (EEG) indicator of profound brain inactivation in which bursts of electrical activity alternate with periods of isoelectricity termed suppression. Specified time-varying levels of burst suppression are targeted in medical coma, a drug-induced brain state used for example to treat uncontrollable seizures. A brain-machine interface (BMI) that observes the EEG could automate the control of drug infusion rate to track a desired target burst suppression trajectory. Such a BMI needs to use models of drug dynamics and burst suppression observations, whose parameters could change with the burst suppression level and the environment over time. Currently, these parameters are fit prior to real-time control, requiring a separate system identification session. Moreover, this approach cannot track parameter variations over time. In addition, small variations in drug infusion rate may be desired at steady state. Here we develop a novel adaptive algorithm for robust control of medical coma in face of unknown and time-varying system parameters. We design an adaptive recursive Bayesian estimator to jointly estimate drug concentrations and system parameters in real time. We construct a controller using the linear-quadratic-regulator strategy that explicitly penalizes large infusion rate variations at steady state and uses the estimates as feedback to generate robust control. Using simulations, we show that the adaptive algorithm achieves precise control of time-varying target levels of burst suppression even when model parameters are initialized randomly, and reduces the infusion rate variation at steady state.**

I. INTRODUCTION

Medical coma is a drug-induced state of profound brain inactivation used after traumatic brain injuries and for treatment of uncontrollable seizures. The electroencephalogram (EEG) signal in medical coma, termed burst suppression, consists of bursts of electrical activity alternating with suppression periods. In current practice, medical coma is maintained by the intensive care unit (ICU) staff who target a specified suppression level by manually adjusting the anesthetic infusion rate based on the patient's EEG activity. Since the state of coma for such treatments is usually required for days, it is often infeasible for the ICU staff to continuously monitor the EEG and adjust the drug infusion rate to achieve tight control. To enable automatic control of burst suppression in medical coma based on the EEG, brain-machine-interface (BMI) systems can be developed [1], [2].

There has been considerable work on developing BMIs, often termed closed-loop anesthetic delivery (CLAD) systems, for control of sedation and general anesthesia (e.g., [3], [4]). For burst suppression, CLAD systems using non-model based control have been implemented in a rat model [5], [6]. However, these studies controlled a constant level of burst suppression rather than time-varying levels needed in medical coma, and reported average control results over animals. BMIs for management of medical coma only appeared recently [1], [2] and used the concept of burst suppression probability (BSP) [7] to quantify the burst suppression level. The BMI in [1] used a stochastic controller that enabled precise control of time-varying target BSP levels in individual rodents and allowed prompt transitions between levels without overshoot or undershoot. To build this BMI, simple models of drug dynamics and burst suppression observations are used. Currently, the parameters of these models are estimated in a system identification session prior to real-time control, and are assumed to be static over time. However, in the operating room or ICU, it is desirable to avoid a separate system identification session. Moreover, system parameters can change over time with changes in BSP level and environment. In addition, small variations in the drug infusion rate at steady state may be desired for robust control.

Here we develop an adaptive BMI algorithm, called online-identification (OnlineID), which jointly estimates the drug concentrations and system parameters in real time and then uses these estimates as feedback to generate robust control of burst suppression. The BMI algorithm is designed in a stochastic optimal control framework that has been used in BMIs for medical coma [1] and in motor BMIs [8], [9], [10]. We use BSP as the control signal and design a recursive Bayesian estimator by constructing and linearizing a nonlinear state-space model. We use a linear-quadratic-regulator (LQR) strategy to form a control law that takes both the estimated system parameters and estimated BSP as feedback to adjust the drug infusion rate and achieve a desired BSP level. Finally, to ensure small steady-state variations in the drug infusion rates, we add a novel penalty term to the controller objective function. Using numerical simulations, we show that the adaptive BMI algorithm achieves reliable control in face of both unknown and time-varying system parameters. Even with random initial parameters, the algorithm can achieve and maintain multiple target levels with low bias and error, can enable prompt transitions between target levels without overshoot and undershoot, and can ensure small drug infusion rate variations at fixed BSP levels.

II. METHODS

The adaptive BMI algorithm, OnlineID, consists of an estimator that estimates the drug concentrations and system

[1,2]Yuxiao Yang and Maryam M. Shanechi are with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853 USA yy566@cornell.edu, shanechi@cornell.edu.

978-1-4244-7929-0/14/$26.00 ©2014 IEEE

parameters online, and a controller that takes the estimates as feedback to reliably track desired BSP levels and ensure small steady-state infusion rate variations. We now present the formulation, and the estimator and controller designs.

A. Problem formulation

We use BSP, denoted by p_t, as our measure of the burst suppression level. BSP indicates the brain's instantaneous probability of being suppressed and takes values in $[0, 1]$. BSP is computed by filtering and thresholding the EEG to identify the activity in each small interval as a burst or a suppression. This converts the EEG into a binary time-series, with 1 indicating a suppression and 0 indicating a burst [1]. This binary signal is the input to our adaptive algorithm. We relate BSP to the brain's anesthetic concentration, denoted by $x_e(t)$, via a hyperbolic transform [1], [7]

$$p_t = (1 - \mathrm{e}^{-x_e(t)})/(1 + \mathrm{e}^{-x_e(t)}). \tag{1}$$

The BMI aims to control the BSP or equivalently the brain's anesthetic concentration.

To develop the adaptive BMI, we build a state-space model using the two-compartment model in [1], [2], [11] to describe the anesthetic drug's dynamics in burst suppression. The two compartments represent the central plasma and the brain. We also take the unknown time-varying system parameters into consideration. We first rewrite the linear two-compartment state-space model as

$$\mathbf{x}(t) = \mathbf{A}(t-1)\mathbf{x}(t-1) + \mathbf{b}u(t-1), \tag{2}$$

with $\mathbf{A}(t) = \begin{bmatrix} 1 - \Delta(k_{ce}(t) + k_{c0}(t)) & \Delta k_{ec}(t) \\ \Delta k_{ce}(t) & 1 - \Delta k_{ec}(t) \end{bmatrix}$, and $\mathbf{x}(t) = [x_c(t), x_e(t)]^{'}$, where $x_c(t)$ denotes the central plasma drug concentration. Here $\mathbf{b} = [\Delta, 0]^{'}$, where Δ is the discretization time step. The control signal $u(t)$ is the instantaneous drug infusion rate. The system parameters, i.e., $k_{ce}(t)$, $k_{ec}(t)$ and $k_{c0}(t)$, represent the rate of the drug flowing from the central plasma into the brain, flowing from the brain back to the central plasma, and being eliminated from the central plasma, respectively. These parameters are assumed to be unknown and time-varying.

Our adaptive algorithm estimates both the time-varying system parameters and the anesthetic concentrations $x_c(t)$ and $x_e(t)$ online. To jointly estimate these variables, we incorporate them into an augmented five-dimensional state

$$\mathbf{y}(t) = [x_c(t), x_e(t), k_{ce}(t), k_{ec}(t), k_{c0}(t)]^{'}. \tag{3}$$

B. Estimator design

In this section, we derive a recursive Bayesian estimator for $\mathbf{y}(t)$ based on the EEG binary time-series. As all components of $\mathbf{y}(t)$ are positive, we estimate its logarithm

$$\mathbf{z}(t) = \log \mathbf{y}(t) = [z_c(t), z_e(t), z_{ce}(t), z_{ec}(t), z_{c0}(t)]^{'}. \tag{4}$$

After estimating $\mathbf{z}(t)$, we can find $\mathbf{y}(t)$ as $\mathbf{y}(t) = \mathrm{e}^{\mathbf{z}(t)}$.

A recursive Bayesian estimator consists of a prior model on the states and an observation model that relates the EEG to these states. To make the adaptive algorithm robust under any system parameter dynamics, prior random-walk models are used for the system parameters $z_{ce}(t)$, $z_{ec}(t)$ and $z_{c0}(t)$. The two-compartment model (2) relates the drug concentration states to the system parameters. We also take model noise into consideration. We thus build the prior model as

$$\mathbf{z}(t) = f(\mathbf{z}(t-1), u(t-1)) + \mathbf{w}(t-1), \tag{5}$$

where $\mathbf{w}(t)$ is modeled as additive Gaussian noise with mean 0 and covariance matrix $\mathbf{W}$, and

$$f(\mathbf{z}, u) = [f_1(\mathbf{z}, u), f_2(\mathbf{z}), z_{ce}, z_{ec}, z_{c0}]^{'},$$

is a nonlinear function with

$$f_1(\mathbf{z}, u) = \log[(1 - \Delta \mathrm{e}^{z_{ce}} - \Delta \mathrm{e}^{z_{c0}})\mathrm{e}^{z_c} + \Delta \mathrm{e}^{z_{ec}} \mathrm{e}^{z_e} + \Delta u]$$
$$f_2(\mathbf{z}) = \log[\Delta \mathrm{e}^{z_{ce}} \mathrm{e}^{z_c} + (1 - \Delta \mathrm{e}^{z_{ec}})\mathrm{e}^{z_e}].$$

The observation in the estimator is the EEG binary time-series. The observation model assumes that the number of suppressions, N_t, in a time interval Δ with at most N suppressions, is binomially distributed with probability p_t (1). We estimate $\mathbf{z}(t)$ by finding the minimum mean-square error (MMSE) state estimate at each time, which is given by the mean of the posterior density $p(\mathbf{z}(t)|N_{1:t})$.

As the prior model (5) and binomial observation model are both nonlinear in the state $\mathbf{z}(t)$, we make a linear approximation to the prior model and a Gaussian approximation to the posterior model. Denoting the posterior mean $E(\mathbf{z}(t)|N_{1:t})$ by $\mathbf{z}_{t|t}$ and its covariance by $\mathbf{W}_{t|t}$, and the mean of the one step prediction density $p(\mathbf{z}(t)|N_{1:t-1})$ by $\mathbf{z}_{t|t-1}$ and its covariance by $\mathbf{W}_{t|t-1}$, estimator recursions are derived as:

$$\mathbf{z}_{t|t-1} = f(\mathbf{z}_{t-1|t-1}, u(t-1)) \tag{6}$$

$$\mathbf{W}_{t|t-1} = \mathbf{F}\mathbf{W}_{t-1|t-1}\mathbf{F}^{'} + \mathbf{W} \tag{7}$$

$$\mathbf{z}_{t|t} = \mathbf{z}_{t|t-1} + \mathbf{W}_{t|t}\left[0, \frac{c_t(N_t - Np_t)}{p_t(1-p_t)}, 0, 0, 0\right]^{'}_{\mathbf{z}_{t|t-1}} \tag{8}$$

$$\mathbf{W}_{t|t}^{-1} = \mathbf{W}_{t|t-1}^{-1} + \left[\begin{array}{cc} \begin{pmatrix} 0 & 0 \\ 0 & \gamma_t \end{pmatrix} & \mathbf{0}_{2\times 3} \\ \mathbf{0}_{3\times 2} & \mathbf{0}_{3\times 3} \end{array}\right]_{\mathbf{z}_{t|t-1}}, \tag{9}$$

where $[\cdot]_a$ indicates the evaluation of the inside expression at value a, $\mathbf{F} = [\frac{\partial f}{\partial \mathbf{z}}]_{(\mathbf{z}_{t-1|t-1}, u(t-1))}$, $\mathbf{0}_{p\times q}$ is a p by q zero matrix and

$$c_t = \frac{x_e(t)\mathrm{e}^{x_e(t)}}{1+\mathrm{e}^{x_e(t)}}(1 - p_t), \tag{10}$$

$$\gamma_t = \frac{Nc_t^2}{p_t(1-p_t)} - \frac{N_t - Np_t}{p_t(1-p_t)}\left[\frac{\partial^2 p_t}{\partial z_e^2(t)} - \frac{1-2p_t}{p_t(1-p_t)}c_t^2\right], \tag{11}$$

$$\frac{\partial^2 p_t}{\partial z_e^2(t)} = c_t[1 + x_e(t) - (1 - p_t)x_e(t)\mathrm{e}^{x_e(t)}]. \tag{12}$$

C. Controller design

We have three goals for the controller, which can be quantified in a quadratic cost function. The main goal is to take the brain concentration $x_e(t)$ close to the target concentration level x^*, which is specified by the desired BSP level p^* through $x^* = \log((1 + p^*)/(1 - p^*))$ (1). This is done by minimizing $\sum_t (x_e(t) - x^*)^2$. We also want to use as little drug as possible, so we include $\sum_t u(t)^2$ in the cost function. In addition, we enforce small drug infusion rate

variations by penalizing the difference between the infusion rate of the current time step and that of the previous time step. Hence, we add a new term $\sum_t (u(t) - u(t-1))^2$ to the cost function. In summary, we form the cost function as

$$J = \sum_t (x_e(t) - x^*)^2 + w_r u(t)^2 + w_s v(t)^2, \quad (13)$$

where $v(t) = u(t) - u(t-1)$, and w_r and w_s are positive quantities chosen depending on the desired system response.

At each time, the controller chooses $u(t)$ that minimizes this cost function. To solve the minimization problem, we first define a new state $\tilde{\mathbf{x}}(t) = [x_c(t), x_e(t), u(t-1)]'$ and use $v(t)$ as the control variable. Then, we rewrite the cost as $J = \sum_t \tilde{\mathbf{x}}'(t)\tilde{\mathbf{Q}}\tilde{\mathbf{x}}(t) + w_s v(t)^2$, where $\tilde{\mathbf{Q}}$ is a 3×3 diagonal matrix with diagonal elements $[0, 1, w_r]$. Given parameter estimates k_{ce}^{est}, k_{ec}^{est} and k_{c0}^{est} at time t, we use the linear model in (2) and after simple manipulations, we derive the state-space model for $\tilde{\mathbf{x}}(t)$ as $\tilde{\mathbf{x}}(t) = \tilde{\mathbf{A}}^{est}\tilde{\mathbf{x}}(t-1) + \tilde{\mathbf{b}}v(t-1)$, with $\tilde{\mathbf{A}}^{est} = \begin{bmatrix} \mathbf{A}^{est} & \mathbf{b} \\ \mathbf{0}_{1\times 2} & 1 \end{bmatrix}$ and $\tilde{\mathbf{b}} = [\Delta, 0, 1]'$, where $\mathbf{A}^{est}$ is the system matrix in (2) evaluated at the estimated parameters. Hence by using the standard LQR solution and by transforming the origin of the state-space to x^*, we find $v(t) = \mathbf{L}^{est}(\tilde{\mathbf{x}}(t) - \tilde{\mathbf{x}}^{*est})$, where $\tilde{\mathbf{x}}^{*est} = [\frac{k_{ec}^{est}}{k_{ce}^{est}}x^*, x^*, \frac{k_{ec}^{est}k_{c0}^{est}}{k_{ce}^{est}}x^*]'$, and $\mathbf{L}^{est}$ is the Riccatti equation solution [12]. Finally, the control $u(t)$ is given by

$$u(t) = \mathbf{L}^{est}(\tilde{\mathbf{x}}^{est}(t) - \tilde{\mathbf{x}}^{*est}) + u(t-1), \quad (14)$$

where $\tilde{\mathbf{x}}^{est}(t) = [x_c^{est}(t), x_e^{est}(t), u(t-1)]'$ is given by the estimates and the known previous drug infusion rate. Note that $u(t-1) = 0$ for $t = 1$. We impose any rate constraints by bounding the unconstrained optimized rate (14).

III. OnlineID algorithm at steady state

In this section, we find sufficient conditions for the algorithm to achieve desired reliable control at steady state.

Let's consider the case where target brain drug concentration is fixed at a specific level x^* and assume we have an oracle algorithm that knows the underlying true system parameters k_{ce}^{true}, k_{ec}^{true} and k_{c0}^{true}. The oracle steady-state control, which is what we want to achieve, can be calculated as

$$u_{oracle} = \frac{k_{ec}^{true}k_{c0}^{true}}{k_{ce}^{true}}x^*. \quad (15)$$

Then, sufficient conditions for $u(t) = u_{oracle}$ are found as

$$\frac{k_{ec}^{est}k_{c0}^{est}}{k_{ce}^{est}} = \frac{k_{ec}^{true}k_{c0}^{true}}{k_{ce}^{true}}, \quad (16)$$

$$x_e^{est}(t) = x^*, \quad (17)$$

$$x_c^{est}(t) = \frac{k_{ec}^{est}}{k_{ce}^{est}}x^*. \quad (18)$$

The sufficient conditions are found as follows. First, these conditions indicate that the estimator converges. Then, by plugging (17) and (18) into (14), one can show that the controller also converges to a steady-state control $u_s = \frac{k_{ec}^{est}k_{c0}^{est}}{k_{ce}^{est}}x^*$. This combined with (15) and (16), indicates that the steady-state control u_s is identical to u_{oracle}.

IV. Results

Numerical experiments are conducted to evaluate the proposed adaptive BMI algorithm. For illustration purposes, we assume that the system parameters $k_{ce}(t)$, $k_{ec}(t)$ and $k_{c0}(t)$ are changing linearly with both the brain and central drug concentrations, $x_e(t)$ and $x_c(t)$, at each time step as

$$k_{ce}(t) = (1 + \alpha_{ce}x_e(t) + \beta_{ce}x_c(t))k_{ce}^0,$$
$$k_{ec}(t) = (1 + \alpha_{ec}x_e(t) + \beta_{ec}x_c(t))k_{ec}^0,$$
$$k_{c0}(t) = (1 + \alpha_{c0}x_e(t) + \beta_{c0}x_c(t))k_{c0}^0,$$

where k_{ce}^0, k_{ec}^0 and k_{c0}^0 are the initial parameters at the beginning of the experiment, and α_{ce}, α_{ec}, α_{c0}, β_{ce}, β_{ec} and β_{c0} are noisy model parameters. But note that our algorithm is a general adaptive approach that can apply to any other parameter dynamics. This linear model, which also incorporates adequate uncertainties, merely serves to test the algorithm. In the experiments, we choose k_{ce}^0, k_{ec}^0 and k_{c0}^0 randomly from independent uniform distributions. The noisy model parameters are set as $\alpha_{ce} = 4 + w_{ce}^e$, $\alpha_{ec} = 4 + w_{ec}^e$, $\alpha_{c0} = 4 + w_{c0}^e$, $\beta_{ce} = 0.004 + w_{ce}^c$, $\beta_{ec} = 0.004 + w_{ec}^c$ and $\beta_{c0} = 0.004 + w_{c0}^c$, where w_{ce}^e, w_{ec}^e and w_{c0}^e are independent model noises drawn from uniform distributions on $[-2, 2]$, and w_{ce}^c, w_{ec}^c and w_{c0}^c are independent model noises drawn from uniform distributions on $[-0.002, 0.002]$.

To assess the performance of OnlineID, we introduce an oracle (Oracle) algorithm that knows the true underlying parameters exactly, and a static-parameter (StaticP) algorithm that is given the true parameters at the first time step (this can be thought of as system identification) but never updates the parameters. Thus, StaticP is actually the nonadaptive algorithm introduced in [1]. Our OnlineID algorithm is only given a set of random starting parameters. To assess the effect of OnlineID in reducing steady-state drug variation, we set $w_s = 0.025$ in OnlineID and introduce an elementary-OnlineID (E-OnlineID) algorithm that is identical to OnlineID except that it uses $w_s = 0$. This means that E-OnlineID doesn't incorporate the novel cost term in the controller. Oracle and StaticP also use $w_s = 0$. We set $w_r = 0.0005$ in all of the above four algorithms.

We characterize the performance of the algorithms at steady state by using the median prediction error (MDPE) given by MDPE $= \text{median}(e(t))/p^*(t) \times 100\%$, and the median absolute performance error (MDAPE) given by MDAPE $= \text{median}(|e(t)|)/p^*(t) \times 100\%$, where $e(t) = p^*(t) - p(t)$ is the error between the desired BSP level $p^*(t)$ and the controlled BSP level $p(t)$. The MDPE is a measure of bias and the MDAPE is a measure of normalized error. We measure the variation of the infusion rate by calculating the normalized absolute error (NMAE) given by NMAE $= \text{median}(|d(t)|)/\text{mean}(u(t)) \times 100\%$, where $d(t) = u(t) - u(t-1)$ is the difference between adjacent controls.

We compare the performance of the algorithms by applying them to 6 experiments, each consisting of an initial period with a low target BSP level of 0.1 followed by a random permutation of three target BSP levels (high 0.8, medium 0.5 and low 0.3). The initial low target level allows time for the

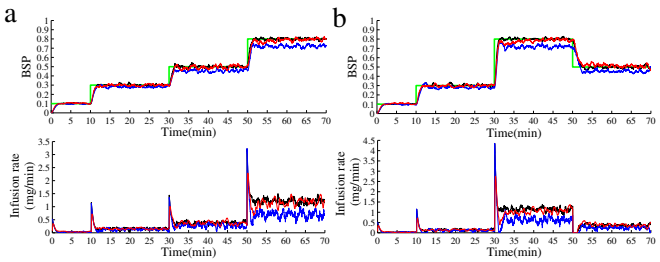

Fig. 1. Control of burst suppression in two simulation experiments. In each subfigure, the top panel shows the controlled BSP trace (black for Oracle, red for OnlineID and blue for StaticP) and the desired time-varying target level (green), and the bottom panel shows the corresponding drug infusion rate administered by the BMI.

TABLE I

COMPARISON OF MDPE AND MDAPE OF DIFFERENT ALGORITHMS

	Oracle	OnlineID	E-OnlineID	StaticP
MDPE	0.1935%	-0.5081%	-0.5079%	9.4067%
MDAPE	1.5022%	2.0709%	1.9389%	9.4067%

TABLE II

COMPARISON OF NMAE OF ONLINEID AND E-ONLINEID ALGORITHM

	OnlineID	E-OnlineID
NMAE	1.1695%	5.0146%

parameters to converge in OnlineID before starting control at higher levels. For each experiment, we simulate 100 trials of the closed-loop system using different random starting parameters. The results of 2 experiments are shown in Fig. 1. We find the average MDPE and MDAPE for the algorithms over all trials (Table I) and then compute the NMAE of the infusion rate differences for both OnlineID and E-OnlineID (Table II). Table I indicates that OnlineID has a bias that is close to 0 and is less than $1/18$ of the bias of StaticP. The MDAPE of OnlineID is comparable with Oracle, and is much smaller than StaticP. This shows that OnlineID can achieve both low bias and error at steady state, even when its parameters are initialized randomly. In Table II, the NMAE of infusion rate variation for OnlineID is less than $1/4$ of E-OnlineID and from Table I, OnlineID and E-OnlineID have comparable MDPE and MDAPE. This shows that the novel cost term can effectively reduce infusion rate variations at steady state at nearly no expense of bias or error. We also find that for each experiment, OnlineID indeed satisfies the three sufficient conditions (16), (17) and (18) asymptotically (Fig. 2). It thus ensures that the generated control converges to the oracle control (see bottom panels in Fig. 1).

V. CONCLUSIONS

We develop an adaptive BMI algorithm, termed OnlineID, for control of burst suppression in medical coma in face of unknown and time-varying system parameters. By using this algorithm, we can achieve reliable control without a system identification session. The controlled BSP has both low bias and error at steady state and makes reliable changes between different target BSP levels. In addition, the infusion rates

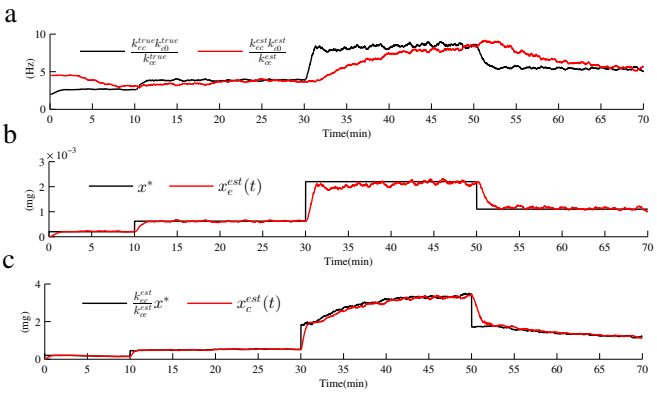

Fig. 2. Sufficient conditions for reliable control. The figure shows that OnlineID algorithm tracks the three sufficient conditions described in (16) (subfigure a), (17) (subfigure b) and (18) (subfigure c) in the experiment corresponding to Fig. 1b. Red and black lines represent the left and right hand side expressions in the sufficient conditions, respectively.

have small steady-state variations. We also derive sufficient conditions for reliable control at steady state and verify that the algorithm satisfies these sufficient conditions asymptotically. Our results suggest that the adaptive algorithm has the potential to achieve robust control in experimental setups.

REFERENCES

[1] M. M. Shanechi, J. J. Chemali, M. Liberman, K. Solt, and E. N. Brown, "A brain-machine interface for control of medically-induced coma," *PLoS Comput. Biol.*, vol. 9, no. 10, p. e1003284, 2013.

[2] S. Ching, M. Y. Liberman, J. J. Chemali, M. B. Westover, J. Kenny, K. Solt, P. L. Purdon, and E. N. Brown, "Real-time closed-loop control in a rodent model of medically induced coma using burst suppression," *Anesthesiology*, vol. 119, no. 4, pp. 848–860, Oct. 2013.

[3] J. Rinehart, N. Liu, B. Alexander, and M. Cannesson, "Review article: closed-loop systems in anesthesia: is there a potential for closed-loop fluid management and hemodynamic optimization," *Anesth Analg*, vol. 114, pp. 130–143, 2012.

[4] P. S. A. Glass, "Automated control of anesthesia ten years later: futuristic novelty or present day reality," *Can J Anesth*, vol. 57, pp. 715–719, 2010.

[5] P. C. M. Vijn and J. R. Sneyd, "I.v. anesthesia and EEG burst suppression in rats: bolus injections and closed-loop infusions," *Br. J. Anaesth.*, vol. 81, no. 3, pp. 415–421, Sept. 1998.

[6] J. F. Cotten, R. L. Ge, N. Banacos, E. Pejo, S. S. Husain, J. H. Williams, and D. E. Raines, "Closed-loop continuous infusions of etomidate and etomidate analogs in rats: A comparative study of dosing and the impact on adrenocortical function," *Anesthesiology*, vol. 115, no. 4, pp. 764–773, 2011.

[7] J. Chemali, S. Ching, P. L. Purdon, K. Solt, and E. N. Brown, "Burst suppression probability algorithms: state-space methods for tracking EEG burst suppression," *J. Neural Eng.*, vol. 10, no. 5, 2013.

[8] M. M. Shanechi, G. W. Wornell, Z. M. Williams, and E. N. Brown, "Feedback-controlled parallel point process filter for estimation of goal-directed movements from neural signals," *IEEE Trans. Neural Syst. Rehabil. Eng.*, doi: 10.1109/TNSRE.2012.2221743, Oct. 2012.

[9] M. M. Shanechi, Z. M. Williams, G. W. Wornell, R. Hu, M. Powers, and E. N. Brown, "A real-time brain-machine interface combining motor target and trajectory intent using an optimal feedback control design," *PLOS ONE*, vol. 8, no. 4, p. e59049, Apr. 2013.

[10] M. M. Shanechi and J. M. Carmena, "Optimal feedback-controlled point process decoder for adaptation and assisted training in brain-machine interfaces," in *Proc. IEEE EMBS conference on neural engineering (NER)*, San Diego, CA, Nov. 2013, pp. 653–656.

[11] M. Y. Liberman, S. Ching, J. Chemali, and E. N. Brown, "A closed-loop anesthetic delivery system for real-time control of burst suppression," *J. Neural Eng.*, vol. 10, no. 4, p. 046004, 2013.

[12] K. Kwakernaak and R. Sivan, *Linear Optimal Control Systems*. New York, NY: Wiley-Interscience, 1972.