Correcting Hypothalamic-Pituitary-Adrenal Axis Dysfunction Using Observer-based Explicit Nonlinear Model Predictive Control

Ankush Chakrabarty[1], Gregery T. Buzzard[2], Martin J. Corless[3], Stanisław H. Żak[1] and Ann E. Rundell[4]

***Abstract*—The hypothalamic-pituitary-adrenal (HPA) axis is critical in maintaining homeostasis under physical and psychological stress by modulating cortisol levels in the body. Dysregulation of cortisol levels is linked to numerous stress-related disorders. In this paper, an automated treatment methodology is proposed, employing a variant of nonlinear model predictive control (NMPC), called explicit MPC (EMPC). The controller is informed by an unknown input observer (UIO), which estimates various hormonal levels in the HPA axis system in conjunction with the magnitude of the stress applied on the body, based on measured concentrations of adreno-corticotropic hormones (ACTH). The proposed closed-loop control strategy is tested on multiple *in silico* patients and the effectiveness of the controller performance is demonstrated.**

I. Introduction

The hypothalamic-pituitary-adrenal (HPA) axis is a self-regulating neuro-endocrine system that maintains homeostasis in response to physiological or psychological stress [1]. Malfunctions in cortisol regulation by the HPA-axis results in hypocortisolic conditions. This has been linked to severe stress-related disorders such as chronic fatigue syndrome [2] and post-traumatic stress disorder [3]. The HPA-axis operates as shown in Figure 1. The system is described in more detail in [4], [5]. A common approach to analyze a physiological

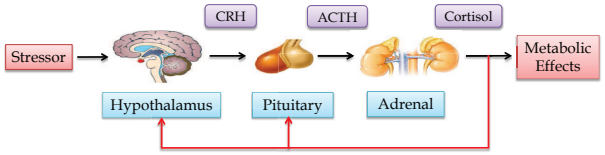

Fig. 1. Dynamics of the HPA Axis system. The red lines indicate negative feedback/inhibitory signals. CRH=Corticotropic Releasing Hormone, ACTH=Adreno-corticotropic Hormone.

system is to construct a mathematical model of the system dynamics. A model of the HPA-axis dynamics based on clinical patient data is proposed in [4]. This model contains two stable steady states/equilibrium points upon removal of stress. One steady-state corresponds to a *healthy equilibrium*. The other is a hypocortisolic equilibrium state, responsible for HPA axis-related disorders. The objective of this paper is to design a control strategy to drive the patient HPA-axis system to the healthy equilibrium.

Recently, numerous control strategies have been proposed to automate drug dosage for therapeutic applications. A neural-network MPC approach was proposed in [6] with reinforcement learning to provide optimal drug dosage values in renal anemia management. Furthermore, the authors in [7] provide clinical data to support the robustness and performance of the model predictive control paradigm in basal insulin control in a clinical setting. A previous study of the application of model predictive control on the HPA axis system is reported in [5] with steady-state analysis to derive control actions which correct HPA-axis dysfunction.

The key contributions of this paper are twofold. First, we employ an explicit model predictive controller (EMPC) to achieve our control objective. It is well-known that nonlinear model predictive control (NMPC) is widely used in the control of constrained nonlinear systems [8]. However, one of its drawbacks is the computational burden involved in computing optimal control actions iteratively. The EMPC addresses this issue by transforming the iterative optimization problem into a form that can be solved offline to create an *explicit* map of control actions as a function of the patient's current state. An additional advantage is that the constructed EMPC map can be inspected by medical professionals over numerous patient states and the safety of derived control actions can be judged before implementation. Second, we introduce an unknown input observer (UIO) to estimate the patient state and stress inputs. We consider circulatory ACTH to be the only state that can be measured as in [5], and all other concentrations of the hormones as well as the stress is estimated by the UIO.

II. HPA model

The HPA-axis model dynamics that we use in the paper was proposed in [4], [5] and is reported to corroborate with clinical data. It has the form,

$$\begin{bmatrix}\dot{x}_1\\ \dot{x}_2\\ \dot{x}_3\\ \dot{x}_4\end{bmatrix} = \begin{bmatrix}\frac{1}{1+\frac{x_4}{k_{i1}}} - k_{cd}x_1\\ \frac{x_1}{1+\frac{x_3x_4}{k_{i2}}} - k_{ad}x_2\\ \frac{(x_3x_4)^2}{(x_3x_4)^2+k} + k_{cr} - k_{rd}x_3\\ x_2 - x_4\end{bmatrix} + \begin{bmatrix}\frac{1}{1+\frac{x_4}{k_{i1}}}\\ 0\\ 0\\ 0\end{bmatrix} d + \begin{bmatrix}0\\0\\0\\1\end{bmatrix} u.$$

[1] A. Chakrabarty (chakraa@purdue.edu) and S.H. Żak (zak@purdue.edu) are affiliated with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN.
[2] G.T. Buzzard (buzzard@purdue.edu) is affiliated with the Department of Mathematics, Purdue University, West Lafayette, IN.
[3] M.J. Corless (corless@purdue.edu) is affiliated with the School of Aeronautics and Astronautics, Purdue University, West Lafayette, IN.
[4] A.E. Rundell (rundell@purdue.edu) is affiliated with the Weldon School of Biomedical Engineering, Purdue University, West Lafayette, IN.

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

We represent this model compactly as,

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{B}_1(\boldsymbol{x})d + \boldsymbol{B}_2 u. \tag{1}$$

The state variables $x_1, \ldots, x_4$ correspond to normalized concentrations of CRH, ACTH, free glucocorticoid receptors (GR) and cortisol, respectively. The external stress input to the system is denoted by d, and the control action u represents the rate of addition or removal of cortisol from the peripheral blood by means of cortisol analogues. The model parameters $k_{i1}, k_{cd}, \ldots, k_{rd}$ are fixed at their nominal values as reported in [5]. A bifurcation diagram of the cortisol concentration with varying stress is shown in Figure 2. The intrinsic bistability of the system under absence of stress is illustrated, and the two equilibria are labeled.

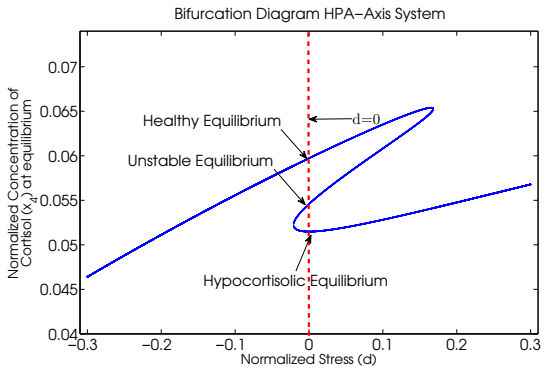

Fig. 2. Bifurcation diagram of normalized cortisol (x_4) with varying stress (d). The vertical red line indicates the absence of stress, that is, $d = 0$. The equilibrium corresponding to the lower steady-state cortisol value is the hypocortisolic state we wish to avoid. The healthy equilibrium to which the patient must be driven is also shown above.

III. Constructing the Controller to Achieve Healthy Equilibrium

The NMPC is an iterative methodology that involves the use of a mathematical model to predict the future states of a system based on the states at the current time. To alleviate the excessive computational burden of the NMPC, we construct an offline version called explicit model predictive control (EMPC).

We construct an EMPC on an admissible range of patient states, denoted $\mathcal{X}$ and a range of magnitudes of stress, $\mathcal{D}$. The EMPC in our construction is an interpolated nonlinear map which returns a control action given the patient's current state $\boldsymbol{x}$ and stress d, that is, $u : \mathcal{X} \times \mathcal{D} \to \mathcal{U}$. The construction of the interpolant $u(\boldsymbol{x}, d)$ can be divided into two steps: (1) the construction of a terminal region which guarantees asymptotic stability to the healthy equilibrium denoted $\boldsymbol{x}_h^*$, and (2) approximation of the controller surface using sparse-grid interpolation.

A. Finding the terminal region of guaranteed stability

To construct the EMPC, a terminal region $\mathcal{T} \subset \mathcal{X}$ is identified to which the patient's state can be driven such that the patient will attain the healthy equilibrium $\boldsymbol{x}_h^*$ upon removal of stress ($d = 0$). Such a region $\mathcal{T}$ is a domain of attraction of $\boldsymbol{x}_h^*$. Let $\Delta\boldsymbol{x} = \boldsymbol{x} - \boldsymbol{x}_h^*$. In this paper, we consider $\mathcal{T}$ to be an ellipsoidal terminal region of the form:

$$\mathcal{T} = \{\boldsymbol{x} : \Delta\boldsymbol{x}^\top \boldsymbol{P} \Delta\boldsymbol{x} \leq \varepsilon\}, \tag{2}$$

where $\boldsymbol{P}$ is a positive definite, real symmetric matrix that solves the continuous Lyapunov equation and $\varepsilon > 0$ is a scalar. We use the method described in [9] to compute $\mathcal{T}$.

B. Approximating the EMPC using sparse-grid interpolation

Upon computing a terminal region, we then calculate an NMPC control action which drives the patient states to $\mathcal{T}$. The control action is a function of the initial condition $\boldsymbol{x}_0$, and is derived by solving a quasi-infinite horizon optimal control problem,

$$\begin{aligned} u(\boldsymbol{x}_0, d) = \arg\min_{u} \; & \Delta\boldsymbol{x}(T_f)^\top \boldsymbol{P} \boldsymbol{\Delta} \boldsymbol{x}(T_f) + \\ & \int_0^{T_f} \Delta\boldsymbol{x}^\top \boldsymbol{Q} \Delta\boldsymbol{x} + \boldsymbol{R} u^2 \; d\tau, \\ & \text{subject to (1)}, \; \boldsymbol{x}(0) = \boldsymbol{x}_0, \; \boldsymbol{x}(T_f) \in \mathcal{T} \\ & \boldsymbol{u}(t) \in \mathcal{U}, \boldsymbol{x}(t) \in \mathcal{X} \; \forall \, t \in [0, T_f], \end{aligned} \tag{3}$$

where $\boldsymbol{Q}, \boldsymbol{R}$ are positive definite, symmetric weighting matrices, $\boldsymbol{x}_0$ is an initial state of the patient, T_f is the specified prediction horizon within which time the state will enter the terminal region $\mathcal{T}$. As in conventional NMPC, only a small initial part of the control trajectory of $u(\boldsymbol{x}_0)$ is stored to construct the interpolated EMPC surface.

In order to construct an EMPC employing this NMPC formulation, we employ a sparse-grid framework as discussed in our previous work [10]. The sparse-grid architecture is chosen due to its tractability in construction of higher dimensional interpolants with high computational efficiency and reduced memory usage [11], [12]. We select N points (called samples/nodes) on $\mathcal{X} \times \mathcal{D}$. At the ith sample, $\left[\boldsymbol{x}^i, d^i\right]$, we solve (3) with $\boldsymbol{x}^i$ as the initial condition of the system, that is, $\boldsymbol{x}(0) = \boldsymbol{x}^i$ under a constant stress d^i.

This process is continued for $i = 1, 2, ..., N$. A sparse-grid interpolation algorithm is employed in order to approximate the EMPC control surface $u(\boldsymbol{x}, d)$ over $\mathcal{X} \times \mathcal{D}$ with information at the N nodes. Note that for the EMPC construction stage, we assume full knowledge of the states and stressor input values. In practice, however, such an implementation is unrealistic as only a subset of the states can be measured in real-time. This fact motivates the construction of the unknown input observer (UIO).

IV. Patient State and Stress Estimation Using an Observer

An observer is a deterministic dynamical system which can generate estimates of the patient's condition (states) given available measurements (outputs) and drug dosage history (control inputs). We represent the available measurements as $y = x_2 = \boldsymbol{C}\boldsymbol{x}$, where $\boldsymbol{C} = \begin{bmatrix} 0 & 1 & 0 & 0 \end{bmatrix}$ as we only measure x_2, the circulatory ACTH concentration. We denote the measurement error as $e_y = y - \hat{y}$, where hatted variables

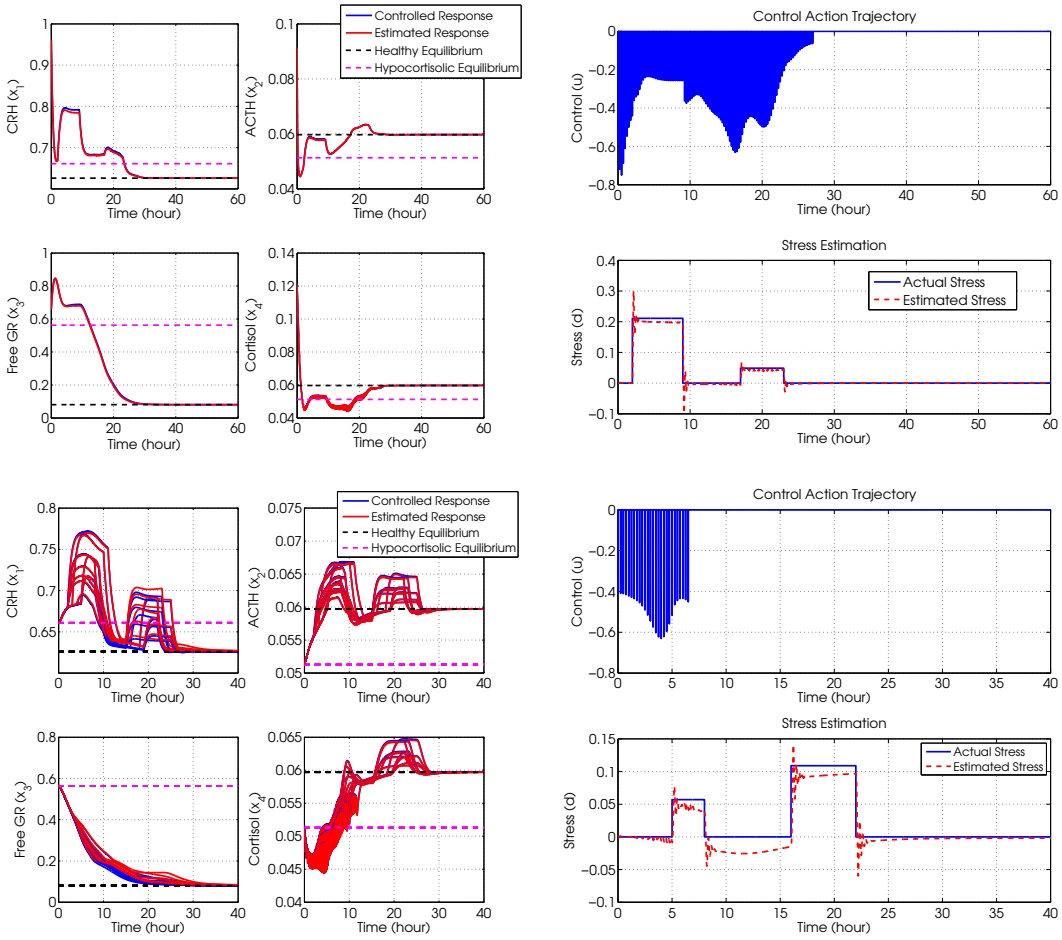

Fig. 3. (*Upper*) Controller performance for the nominal patient. The black dashed-line denotes the healthy equilibrium $\boldsymbol{x}_h^*$ whereas the pink dashed-line is the hypocortisolic steady-state that we wish to avoid. All concentrations of CRH, ACTH, free GR, and Cortisol, as well as the magnitudes of stress and control input are normalized. (*Lower*) Controller performance for ten model-mismatched patients and stress-estimation for one model-mismatched patient.

are the patient states estimated by the observer and non-hatted variables are actual patient states. Note that the input u is also known as it is computed from the EMPC interpolant. We use the observer proposed in [13],

$$\dot{\hat{\boldsymbol{x}}} = \boldsymbol{f}(\hat{\boldsymbol{x}}) + \boldsymbol{B}_1 u + \boldsymbol{B}_2(\hat{\boldsymbol{x}})L(e_y), \tag{4}$$

where $L(e_y) = Ke_y$ is a linear injection term with gain K and $\hat{\boldsymbol{x}}$ is the vector of state estimates. The system (4) is an observer of the system (1) if $\lim_{t\to\infty} \boldsymbol{x}(t) - \hat{\boldsymbol{x}}(t) = 0$ for a set of initial conditions $\boldsymbol{x}(0)$ and $\hat{\boldsymbol{x}}(0)$. Let the estimation error be denoted as $\boldsymbol{e} \triangleq \boldsymbol{x} - \hat{\boldsymbol{x}}$. The objective is to design an observer such that,

$$\lim_{t\to\infty} \boldsymbol{e}(t) = 0. \tag{5}$$

The dynamics of the estimation error system are given by,

$$\begin{aligned}\dot{\boldsymbol{e}} &= \boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{f}(\hat{\boldsymbol{x}}) + \boldsymbol{B}_2(\boldsymbol{x})d - \boldsymbol{B}_2(\hat{\boldsymbol{x}})L(e_y), \\ &= \boldsymbol{f}(\boldsymbol{e}+\boldsymbol{x}) - \boldsymbol{f}(\hat{\boldsymbol{x}}) + \boldsymbol{B}_2(\boldsymbol{e}+\hat{\boldsymbol{x}})d - \boldsymbol{B}_2(\hat{\boldsymbol{x}})L(e_y). \end{aligned} \tag{6}$$

The construction of the observer thereby reduces to choosing the injection term to ensure that the system (6) is asymptotically stable at the origin, that is, (5) is satisfied. The estimated stress $\hat{d}$ is obtained by low-pass filtering the signal $L(e_y)$. A detailed analysis of the performance and implementation of the above UIO can be found in [13], [14].

V. *In-silico* TRIALS

In this section, we demonstrate the performance of the closed-loop EMPC control strategy on correction of HPA-axis dysfunction. We present our results into two parts. First, we apply the controller to the nominal patient and then test its performance under patient-model mismatch.

A. Nominal Patient

For the nominal *in-silico* patient, we assume there is no model-patient mismatch, that is, the patient dynamics are perfectly described by the model. The EMPC is constructed with a predictive horizon $T_f = 240$ min, and the terminal

region derived is described by (2), where

$$\boldsymbol{P} = \begin{bmatrix} 1.02 & 0.16 & 4.52 & 1.35 \\ 0.16 & 0.31 & 3.04 & 2.56 \\ 4.52 & 3.04 & 51.63 & 29.41 \\ 1.35 & 2.56 & 29.41 & 24.88 \end{bmatrix}, \ \varepsilon = 0.56.$$

The weighting matrices are taken to be $\boldsymbol{Q} = 10\boldsymbol{I}_4$ and $R = 0.1$, respectively. The admissible state-space is $\mathcal{X} := \{\boldsymbol{x} : \|\boldsymbol{x}\|_\infty \leq 1\}$. Furthermore, $\mathcal{D} := \{d : |d| \leq 0.25\}$ is the range of stressor magnitudes, and $\mathcal{U} := \{u : |u| \leq 0.8\}$ is the admissible control input space. The optimization problem (3) was solved in MATLAB 2013b, using `fmincon` and the sparse-grid sampling and interpolation were performed with the sparse-grid interpolation toolbox [15]. A linear injection term is used for the UIO with $L(e_y) = 900e_y$. A low-pass Butterworth filter is used to estimate the stress signal with cut-off frequency $F_c = 4$Hz, and sampling frequency $F_s = 200$Hz. MATLAB's `filtfilt` command is used for implementing this filter on the stress signal.

A random time-varying stress signal d is simulated as follows $d = \mu_1 s(t-t_1) - \mu_1 s(t-t_2) + \mu_2 s(t-t_3) - \mu_2 s(t-t_4)$, where $\mu_{1,2}$ are random numbers generated within the interval $[-0.25, 0.25]$, and $t_{1,3}$ are random integers denoting the activation times of the stress signal, $t_{2,4}$ denote deactivation times and $s(t)$ is the unit-step signal. The total time an *in-silico* patient is simulated is 60 hours and the rate of cortisol delivery is allowed to change every 15 minutes. The performance of the control strategy is shown in Figure 3. The results obtained here corroborate those in [5], as demonstrated by the negative control values. This implies that the derived treatment strategy is indeed a reduction of the rate of production of cortisol, which seems counter-intuitive to avoiding hypocortisolic steady-states. The UIO proves effective in estimating the states and the stress inputs without large errors throughout simulation time. Also the error between the actual states and the estimated states converge asympotically, as guaranteed in [13]. We notice that the patient enters the terminal region $\mathcal{T}$ after ≈ 28 hours of treatment, from where the patient state converges to the healthy equilibrium $\boldsymbol{x}_h^*$ with $u = d = 0$. Encouraged by the results obtained using the nominal patient model, we next test the control strategy in the presence of parameter mismatch between the patient and the model.

B. Model-Mismatched Patient

We allow the parameters k_{i1}, the CRH inhibitor constant, and k_{cr}, the GR synthesis constant, to vary $\pm 10\%$ about their nominal values under a uniformly distributed random variable for 10 *in-silico* patients. This emulates the condition when the patient dynamics are not perfectly captured by the model. Other design parameters are maintained constant and the nominal controller designed in Section V-A is applied to ten *in-silico* patients for a simulation time of 40 hours. The simulation results are shown in Figure 3. We observe that for all the patients, the healthy equilibrium is attained in spite of model mismatch. The model uncertainty causes a degradation in the stress estimation quality, but due to the inherent robustness of the EMPC the controller performs satisfactorily.

VI. Conclusions

In this paper, we propose an observer-based EMPC controller to correct HPA axis dysfunction. The controller designed on the nominal system (with no uncertainties) is tested on ten simulated patients, demonstrating the effectiveness of the proposed control scheme. Some limitations, however, remain to be addressed. First, we rely on the inherent robustness of the EMPC strategy in the parameter uncertain case without explicitly taking into account the magnitude of parametric uncertainty. Second, we assume the measurement of ACTH possible in real-time, which may be prohibitive in practice.

References

[1] L. Jacobson. Hypothalamic–pituitary–adrenocortical axis regulation. *Endocrinology and metabolism clinics of North America*, 34(2):271–292, 2005.

[2] L.J. Crofford, E.A. Young, N.C. Engleberg, A. Korszun, C.B. Brucksch, L.A. McClure, M.B. Brown, and M.A. Demitrack. Basal circadian and pulsatile acth and cortisol secretion in patients with fibromyalgia and/or chronic fatigue syndrome. *Brain, behavior, and immunity*, 18(4):314–325, 2004.

[3] N. Rohleder, L. Joksimovic, J.M. Wolf, and C. Kirschbaum. Hypocortisolism and increased glucocorticoid sensitivity of pro-inflammatory cytokine production in bosnian war refugees with posttraumatic stress disorder. *Biological psychiatry*, 55(7):745–751, 2004.

[4] S. Gupta, E. Aslakson, B.M. Gurbaxani, and S.D. Vernon. Inclusion of the glucocorticoid receptor in a hypothalamic pituitary adrenal axis model reveals bistability. *Theoretical Biology and Medical Modelling*, 4:8, 2007.

[5] A. Ben-Zvi, S. D. Vernon, and G. Broderick. Model-based therapeutic correction of hypothalamic-pituitary-adrenal axis dysfunction. *PLoS computational biology*, 5(1):e1000273, 2009.

[6] A.E. Gaweda, M.K. Muezzinoglu, A.A. Jacobs, G.R. Aronoff, and M.E. Brier. Model predictive control with reinforcement learning for drug delivery in renal anemia management. In *Engineering in Medicine and Biology Society, 2006. EMBS '06. 28th Annual International Conference of the IEEE*, pages 5177–5180, Aug 2006.

[7] N.N. Razak, J.G. Chase, F.M. Suhaimi, G.M. Shaw, and U. Jamaluddin. A model-based control protocol for transition from ICU to HDU: Robustness analysis. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, volume 2013, pages 205–208, 2013.

[8] J. M. Maciejowski. *Predictive control: with constraints*. Pearson education, 2002.

[9] H. Chen and F. Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Automatica*, 34(10):1205–1217, 1998.

[10] A. Chakrabarty, V. Dinh, G.T. Buzzard, S.H. Żak, and A.E. Rundell. Robust explicit nonlinear model predictive control with integral sliding mode. In *American Control Conference (ACC), 2014*. IEEE, 2014.

[11] V. Barthelmann, E. Novak, and K. Ritter. High dimensional polynomial interpolation on sparse grids. *Advances in Computational Mathematics*, 12(4):273–288, 2000.

[12] G.T. Buzzard. Efficient basis change for sparse-grid interpolating polynomials with application to T-cell sensitivity analysis. *Computational Biology Journal*, 2013.

[13] S. Hui and S.H. Żak. Stress estimation using unknown input observer. In *American Control Conference (ACC), 2013*, pages 259–264. IEEE, 2013.

[14] S. H. Żak. Stress estimation using an hypothalamic-pituitary-adrenal axis model. In M. Busłowicz and K. Malinowski, editors, *Advances in Control Theory and Automation*, pages 283–294. Białystok: Printing House of Białystok University of Technology, 2012.

[15] A. Klimke. Sparse-grid interpolation toolbox. *Universität Stuttgart*, 2008.