2020 American Control Conference
Denver, CO, USA, July 1-3, 2020

Virtual Patient Generation using Physiological Models through a Compressed Latent Parameterization

Ali Tivay, George C. Kramer, and Jin-Oh Hahn, *Senior Member, IEEE*

***Abstract*—This paper presents a data-driven approach to generating virtual patients using mathematical models of physiological processes. Such models often contain a large number of tunable parameters that must be calibrated to capture the observed characteristics of each real patient in a dataset. By sampling from this parameter space, potentially new virtual patients can be generated. However, it is often the case that the resulting set of virtual patients contains members that exhibit physiologically unrealistic behavior. In the present work, we employ a practically important case study on the modeling of cardiovascular responses to hemorrhage and fluid resuscitation in order to demonstrate that subject-specific characteristics observed in a dataset can be alternatively represented within a highly compressed latent parameter space without significant losses in calibration error for each real patient. Then, we show that by sampling from this latent parameter space, it is possible to generate new virtual patients that also exhibit physiologically realistic behavior.**

I. Introduction

The task of automating patient care using planning and control algorithms is worthy of extensive research attention due to its potential for achieving superiority in vigilant and precise performance of patient care routines, especially for critically ill patients. However, effective prototyping and testing of such algorithms are currently challenging due to the expensive nature and ethical limitations of conducting clinical trials on real patients.

Testing patient care algorithms based on populations of virtual patients is a promising direction that can potentially replace clinical trials in the early stages of algorithm and device development, increasing the maturity of the designs before they advance to more expensive stages of testing. For this purpose, using mathematical models of physiological processes as virtual patients has recently received notable attention in the research community [1]. To name a few: the diabetes simulator introduced in [2] is used to develop and test artificial pancreas control algorithms; the model of hemodynamic responses to hemorrhage presented in [3] has been used in a hardware-in-the-loop setup to test fluid resuscitation algorithms [4]; a synthetic virtual cohort of heart electro-grams has been used in [5] to run computer-aided clinical trials for implantable cardiac devices; and models of physiological responses to interacting drugs have been used in [6] to develop and test medication control algorithms. In addition, the U.S. Food and Drug Administration (FDA) has recently acknowledged the potential for computer simulations to complement regulatory submissions for new medical devices [7].

To reproduce the behavior of real patients in the form of virtual patients, a dynamic model of relevant physiological mechanisms is needed. Such models often consist of differential equations with possibly nonlinear elements and appropriately defined input-output signals, and a potentially large number of tunable parameters [8]–[10]. Generating virtual patients in this scenario consists of sampling from a distribution over the model parameters. A prevalent approach to conduct this sampling is to identify parameter values based on a dataset of subjects and use the identified models as virtual patients (e.g. as in [4], [11]). Alternatively, suitable physiological ranges for each parameter (calculated from the literature or based on the calibration of the model to a dataset) can be sampled to obtain potentially larger cohorts (e.g. as in [5], [12]). In some cases, additional information may be available about the distribution of each parameter and/or the relationship between the parameters, which can also be incorporated into the sampling procedure (e.g. as in [2], [13]). Furthermore, data "bootstrapping" can also be cited as a generation method, where subsets of data are randomly sampled and the corresponding maximum-likelihood parameter estimates are regarded as new virtual patients [14].

The possibility of generating virtual patients with realistic and reliable behavior through sampling from the parameter space of a model is limited by at least two important challenges: First, parameter values associated with a physiological system may be related in potentially unknown ways, and thus breaking relations by independently sampling from each model parameter could create virtual patients that would not have existed in reality. Second, a vast array of mathematical models proposed in biology and physics are known to exhibit the "sloppiness" property [15] (i.e., a lack of practical identifiability [16], [17]) in many directions in their parameter space, which is known to cause parameter estimates to drift out of proportion during model calibration, giving larger-than-reality values for parameter ranges, which can in turn result in unrealistic virtual patients when sampled. As an effective ad-hoc solution, objectively un-realistic simulations can be omitted from a virtual population after sampling from the parameters [18], however, a systematic way of generating virtual patients that takes into account both parameter interactions and parameter sloppiness is desirable.

In an attempt to address this challenge, we investigate the generation of virtual patients through sampling from a compressed latent parameter space for the model, where both parameter interaction and sloppiness are minimal. Focusing on a practically important case study on the physiological modeling of cardiovascular responses to hemorrhage and fluid resuscitation, a model structure is first presented that

*Research supported by National Science Foundation CAREER Award (Grant No. 1748762), and CDMRP (Grant No. W81XWH-19-1-0322).

A. Tivay and J.O. Hahn are with the Mechanical Engineering Department, University of Maryland, College Park, MD 20742 USA (phone: 301-405-7864; fax: 301-314-9477; e-mail: tivay@umd.edu).

G. Kramer is with the Anesthesiology Department, University of Texas Medical Branch, Galveston, TX 77555 USA.

can be used to simulate changes in blood hematocrit (HCT), cardiac output (CO), and mean arterial pressure (MAP) in response to hemorrhage and fluid resuscitation. Then, using patient-specific data available for HCT, CO, and MAP over time, the parameters of the model can be calibrated to match its behavior to each real patient. For this purpose, we introduce a compressed latent parameter space for the model where variabilities across different patients are represented by variations in a few latent directions, without significant losses in calibration error. Then, we demonstrate that it is possible to generate new virtual patients that exhibit realistic behavior by sampling from this latent parameter space.

II. Case Study: Modeling Cardiovascular Responses to Hemorrhage and Fluid Resuscitation

Hemorrhage (bleeding) is a serious event that can be incident in critical patients and patients subjected to trauma, the effects of which can be counteracted with appropriate fluid resuscitation. To generate virtual patients that can help design and test automated fluid resuscitation algorithms, a dynamic model is needed to represent the macroscopic responses of the cardiovascular system to both hemorrhage and fluid resuscitation. In this section, we present such a model structure.

Fig. 1 shows a schema of the proposed model structure for the macroscopic response of the cardiovascular system to fluid perturbation. We consider the following equations for the exchange and balance of fluid volume in the system:

$$\dot{v}_a = Q - (p_a - p_v)/R - J_h - J_f \quad (1)$$

$$\dot{v}_v = -Q + (p_a - p_v)/R + J_i \quad (2)$$

$$\dot{v}_r = -J_h H \quad (3)$$

$$H = v_r/(v_a + v_v) \quad (4)$$

$$V = v_a + v_v \quad (5)$$

where v_a and v_v denote arterial and venous blood volume, V is the total blood volume, v_r is the total volume of red blood cells, H represents hematocrit, p_a and p_v denote MAP and central venous pressure (CVP), R represents the systemic vascular resistance, Q is the cardiac output, J_h is the flow rate of hemorrhage, J_f is the flow rate of fluid exchange with the interstitial compartment, and J_i is the flow rate of fluid infusion into the bloodstream.

Changes in MAP and CVP are modeled to linearly depend on changes in arterial and venous blood volume through elastance parameters as follows:

$$\Delta p_a = K_a \Delta v_a \quad (6)$$

$$\Delta p_v = K_v \Delta v_v \quad (7)$$

where K_a and K_v represent the elastance of the arterial and venous volume compartments respectively, and $\Delta x = x\text{-}x_0$ for all quantities.

The blood volume compartment is known to be in relative equilibrium with the fluid in its surrounding tissue (called the interstitial compartment). A perturbation in blood volume is partially counteracted by a shift of fluid to/from this tissue compartment. The net rate of fluid shift J_f is thus

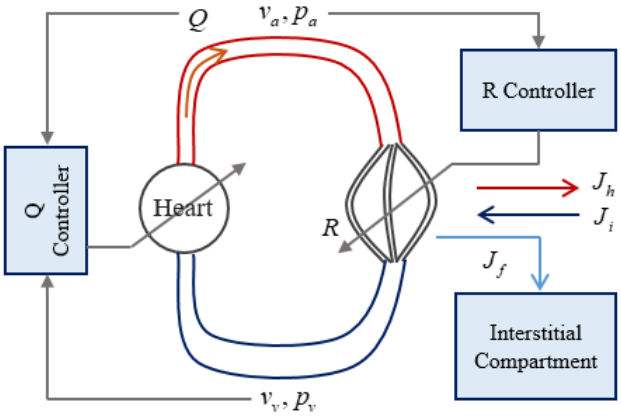

Figure 1. Schematic illustration of the proposed model of cardiovascular response to hemorrhage and fluid resuscitation

modeled as the control input of a hypothetical controller that has the goal of maintaining total blood volume as follows:

$$J_f = -K_p(r_V - \Delta V) \quad (8)$$

$$r_V = \frac{1}{1+\alpha_i}\int_0^t J_i(\tau)d\tau - \frac{1}{1+\alpha_h}\int_0^t J_h(\tau)d\tau \quad (9)$$

where K_p is a proportional gain for the controller, and r_v is the new value of blood volume after re-equilibrium. This value depends on the history of fluid perturbations, and parameters α_i and α_h determine the fraction of each perturbation that will be compensated for by a shift of fluid to/from the tissue compartment. Please refer to [19] for more information about this particular formulation.

The systemic vascular resistance (SVR) denoted by R represents the resistance to blood flow that is present throughout the vascular system. The body can change this resistance through vasoconstriction and vasodilation to restore a lower-than-normal MAP [20]. Also, a change in the fraction of red blood cells in the blood (H) directly affects blood viscosity, which in turn affects SVR [21]. The change in resistance (ΔR) is therefore modeled as a control input that has the goal of maintaining a normal MAP, and is also disturbed by changes in H as follows:

$$\Delta R = -\frac{K_r}{\tau_r s + 1}\Delta p_a + K_h \Delta H \quad (10)$$

where the parameters K_r and τ_r are the gain and time constant of the controller respectively, and K_h represents the sensitivity of SVR to changes in hematocrit.

The cardiac output denoted by Q is the flow rate of blood that is pumped by the heart. This flow rate can be affected by a few important mechanisms: The Frank-Starling mechanism is related to the inherent properties of cardiac muscles, where a higher preload (~proportional to p_v) results in a more forceful stroke in a single beat and thus higher Q. The cardiac contractility (force of contraction) and heart rate are controlled by the autonomic nervous system and the endocrine system in order to maintain a normal Q [20]. Overall, to obtain a minimal and lumped model of these

effects, Q is assumed to be regulated as a controlled variable by a controller that acts through manipulating a control input (which corresponds to heart rate and cardiac contractility), and is disturbed by changes in preload (Δp_v) as follows:

$$\Delta Q = \frac{s}{s + K_c}\left(\beta_v \Delta p_v\right) \tag{11}$$

where K_c is the controller gain, and β_v is the sensitivity of cardiac output to changes in preload (which corresponds to the slope of the cardiac function curve in the Frank-Starling mechanism). Note that (11) represents the closed-loop relationship between the disturbances and ΔQ.

To numerically simulate the model, the initial values for arterial and venous blood volume are set as $v_{a0}=0.3V_0$, and $v_{v0}=0.7V_0$, where V_0 is the initial total blood volume to be estimated; the values for initial arterial pressure p_{a0}, initial venous pressure p_{v0}, initial blood hematocrit H_0, and initial cardiac output Q_0 are set from baseline values in measured data; and the initial SVR is calculated from $R_0=(p_{a0}-p_{v0})/Q_0$.

In the presented model structure, each subject can be characterized by $n_p=11$ tunable parameters (denoted hereafter by the vector θ) as follows:

$$\theta = \left[\alpha_i \;\; \alpha_h \;\; K_p \;\; V_0 \;\; K_a \;\; K_v \;\; K_r \;\; \tau_r \;\; K_h \;\; \beta_v \;\; K_c\right] \tag{12}$$

The experimental data used in this work included HCT, CO, and MAP time-series measurements acquired from $N=23$ animal (sheep) experiments under hemorrhage and fluid resuscitation [22], [23]. The measurements were made at ~5 min intervals for 180 minutes.

III. Virtual Patient Generation using a Compressed Latent Parameterization

An important desirable when generating virtual patients is for the variations in the generated set to be representative of variations that are incident across real patients. Such a representation for inter-subject variability can be thought of as a joint distribution over the model parameters, where each sample represents a virtual patient. In the absence of additional assumptions, finding such a joint distribution is often infeasible given the amount of data that is available in physiological applications. In this section, we first argue that when certain conditions are met, the variations across real subjects can be alternatively represented in a compressed latent parameter space for the model. Then, we demonstrate that independently sampling from the dimensions of this latent space results in a set of generated patients that can represent the variations observed in the dataset and also exhibit realistic behavior.

A. Model Calibration in Compressed Latent Space

To find a compressed representation for the variabilities across patients (in case such a representation exists), we first consider a nominal model parameter vector $\bar{\theta}$, which represents a model of typical physiological behavior. One candidate for such a model is the "group-average" model, which is defined as the solution to the following optimization problem:

$$\bar{\theta} = \arg\min_{\theta} \left\| \mathbf{Y} - \hat{\mathbf{Y}}(\theta) \right\|_2^2 \tag{13}$$

where $\mathbf{Y}$ denotes the data from *all* patients and $\hat{\mathbf{Y}}(\theta)$ denotes the corresponding model predictions given the parameter θ. The solution $\bar{\theta}$ is a maximum-likelihood estimate using all available data, which can be interpreted as a group-average model that represents expected behavior in the population.

Given the group-average model $\bar{\theta}$, variations across real patients in the dataset can be thought of as local deviations from $\bar{\theta}$. To find a potentially compressed representation for these deviations, we are interested in finding orthogonal directions in the vicinity of $\bar{\theta}$, sorted by the prominence of their effect on the predictions of the model. To find such directions, $k >> n_p$ random local *deviations* around $\bar{\theta}$ are obtained and stored in Θ ($n_p \times k$). The corresponding *changes* in model predictions are stored in $\hat{\mathbf{Y}}_\Theta$ ($n_d \times k$). Then, the following matrix can be constructed:

$$\mathbf{C} = \hat{\mathbf{Y}}_{\boldsymbol{\Theta}} \boldsymbol{\Theta}^{\mathbf{T}} = \mathbf{U}\mathbf{S}\mathbf{V}^{\mathbf{T}} \tag{14}$$

where the elements of $\mathbf{C}$ represent the (scaled) covariance between local parametric deviations from the group-average model and the corresponding changes in model predictions. The matrices $\mathbf{U}$, $\mathbf{S}$, and $\mathbf{V}$ are computed from the singular value decomposition of the covariance matrix. The columns of $\mathbf{V}$ constitute sorted orthogonal directions of maximum covariance in the parameter space, and the diagonal values of $\mathbf{S}$ represent the local sensitivity of model predictions to deviations along each of the columns of $\mathbf{V}$.

Depending on the structure of the physiological model, the matrix of (sorted) local sensitivities $\mathbf{S}$ can show interesting properties. For example, in the case that the proposed model exhibits the sloppiness property [15] (which is a prevalent property in a wide range of proposed models across many disciplines) the first few elements of $\mathbf{S}$ will be significantly larger than the rest. This means that locally deviating from the group-average model $\bar{\theta}$ in the first few directions in $\mathbf{V}$ will have a large effect on model outputs while deviating from $\bar{\theta}$ in the last few directions in $\mathbf{V}$ will have a small effect on model outputs and possibly only affect the internal behavior of the model.

Based on the observation above, we can define the following model calibration problem to find patient-specific models for each member of the dataset without unnecessary deviations from the group-average model:

$$\hat{\theta}_i = \arg\min_{\theta} \left\| \mathbf{Y_i} - \hat{\mathbf{Y}}(\theta) \right\|_2^2 + \lambda \left\| (\theta - \bar{\theta})^T \mathbf{V} \right\|_1 \tag{15}$$

where $\hat{\theta}_i$ denotes the vector of model parameters calibrated to match the behavior of real patient i. The second term in (15) measures the L_1-deviation of each patient from the group-average model in the latent space ($\phi^T = \theta^T\mathbf{V} - \bar{\theta}^T\mathbf{V}$). The well-known sparsity-promoting nature of the L_1-norm

induces compression in the latent space so that individual subjects deviate from $\bar{\theta}$ only in a few necessary latent directions. The rate of compression can be controlled through the choice of λ.

B. Generating Virtual Patients

Assuming that a compressed representation exists in the latent space ϕ for variations across different subjects, it is possible to generate new in-silico subjects by mimicking the observed variations in the latent space through sampling from the latent parameters. In this work, samples are drawn from a distribution of the mean-field variational family:

$$P(\phi) = \prod_{j=1}^{n_p} p(\phi_j) \tag{16}$$

where each dimension of the latent parameter space ϕ_j has its density $p(\phi_j)$. For each latent dimension, we use a uniform density with a range that is equal to that of the real patients represented in the compressed latent space. Sampling from $P(\phi)$ will generate virtual patients, and the corresponding model parameter values can be obtained from:

$$\theta = V\phi + \bar{\theta} \tag{17}$$

C. Comparing Generated Cohorts

To evaluate and compare different generated cohorts of virtual patients in terms of having realistic members and covering the range of variation observed in the real patient dataset, we first define the following evidence for the real patient i:

$$P(\mathbf{Y_i} \mid \mathbf{I_i}) = \int_{\Phi} P(\mathbf{Y_i} \mid \mathbf{I_i}, \phi) P(\phi) d\phi \tag{18}$$

where $P(\mathbf{Y_i}|\mathbf{I_i})$ represents the probability that the generation method $P(\phi)$ produces patients that behave similarly to the measured data from the actual subject i, when given the input signals that subject i received in reality (I_i). In this setting, the similarity between a virtual patient and a real patient is measured by measuring the distance of their outputs through a Gaussian kernel:

$$P(\mathbf{Y_i} \mid \mathbf{I_i}, \phi) = \frac{\exp\left(-\frac{1}{2}(\mathbf{Y_i} - \hat{\mathbf{Y}}(\phi))^T \mathbf{\Sigma}^{-1} (\mathbf{Y_i} - \hat{\mathbf{Y}}(\phi))\right)}{\sqrt{(2\pi)^{n_d} |\mathbf{\Sigma}|}} \tag{19}$$

where the elements of the diagonal matrix $\mathbf{\Sigma}$ determine how close the simulation and the data should be to be considered similar. In the present work, diagonal elements of Σ were chosen according to the type of each data point (i.e., σ_{MAP} for MAP data, σ_{HCT} for HCT data, and σ_{CO} for CO data).

Having the probability in (18), which represents how likely is the incidence of patient i under the generation method $P(\phi)$, we can construct the following overall score:

$$S = -\log\left(\prod_{i=1}^{N} P(\mathbf{Y_i} \mid \mathbf{I_i})\right) = -\sum_{i=1}^{N} \log\left[P(\mathbf{Y_i} \mid \mathbf{I_i})\right] \tag{20}$$

which represents how likely is the incidence of *all* real patients under the generation method $P(\phi)$. The score S will be used in subsequent sections as one of the ways to compare different generated cohorts, and a lower score shows a better cohort in this sense.

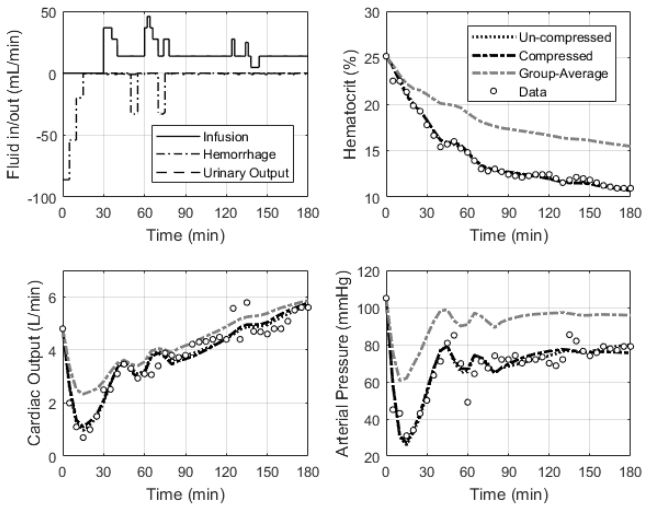

Figure 2. Simulation results of the compressed, un-compressed and group-average models versus the data for one representative subject.

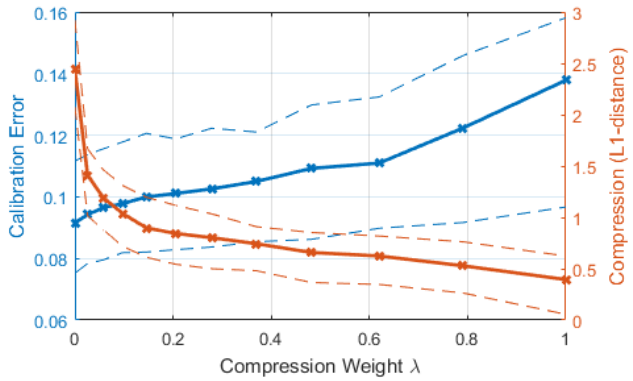

Figure 3. The effect of parameter space compression on calibration error (blue) and L_1-distance (red). Solid lines denote the mean value across all subjects, and dashed lines denote first and third quartile values.

IV. Results and Discussion

Fig. 2 shows the HCT, CO, and MAP responses of the uncompressed, compressed, and group-average calibrated models versus the data for a representative subject. The responses of the group-average model follow the overall trends observed in the data but do not exactly match the data, as this kind of model represents the expected behavior in the population. Uncompressed calibration (λ=0) of model parameters to the subject-specific data results in responses that match the data well. Alternatively, compressed calibration (λ=0.15) of the model to this data results in responses that match the data well and are close to the uncompressed case. This indicates that it is possible to calibrate the model to subject-specific data by limited deviations from the group-average model.

Fig. 3 shows changes in calibration error and deviation distance (in the L_1 sense, from the group-average model)

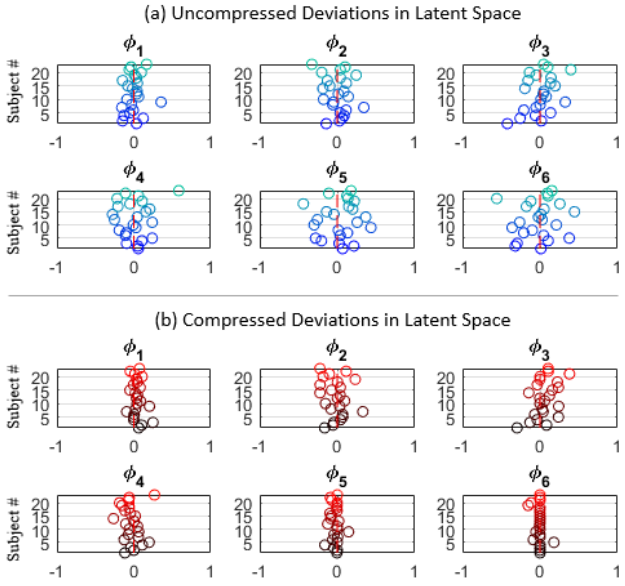

Figure 4. Comparison of the first six latent parameter values identified for the dataset of 23 subjects: (a) latent parameter values in the uncompressed case where λ=0, and (b) in the compressed case where λ=0.15.

with respect to the compression weight λ for the subjects in the dataset. In the uncompressed case (λ=0), the calibrated parameter values tend to have very high L_1-distances from the group-average model, which indicates that they are highly dispersed in the parameter space. As λ is increased, deviation distances drop with a steep slope, while at the same time calibration errors only mildly increase. Furthermore, in the case that λ is increased to very high values, calibration errors will increase significantly and tend to the calibration error for the group-average model. This indicates that for a middle-ground value of the compression weight (e.g. λ=0.15) the variations across different subjects can be represented in a compressed way without noticeable losses in calibration error.

Fig. 4 shows values for the first six latent parameters in case of uncompressed and compressed calibration. In Fig. 4(a), which represents the uncompressed calibration case, variations are visible in all dimensions of the latent space. However, in the compressed calibration case shown in Fig. 4(b), most of the variations are represented by the first four dimensions of the latent space, while from the fifth dimension onward, the latent parameters are nearly zero for most of the real subjects. This result suggests that the variations across different subjects have a compressed representation in the latent space.

Fig. 5 compares the outcomes of virtual subject generation when the compressed versus uncompressed latent parameter values are used in (16) to generate subjects. These outcomes have been shown as histograms of model output values at three points in time (20min, 45min, and 95min) in response to a typical hemorrhage and resuscitation profile. In the uncompressed case, objectively un-realistic subjects are incident (and even common) in the set: (i) many virtual

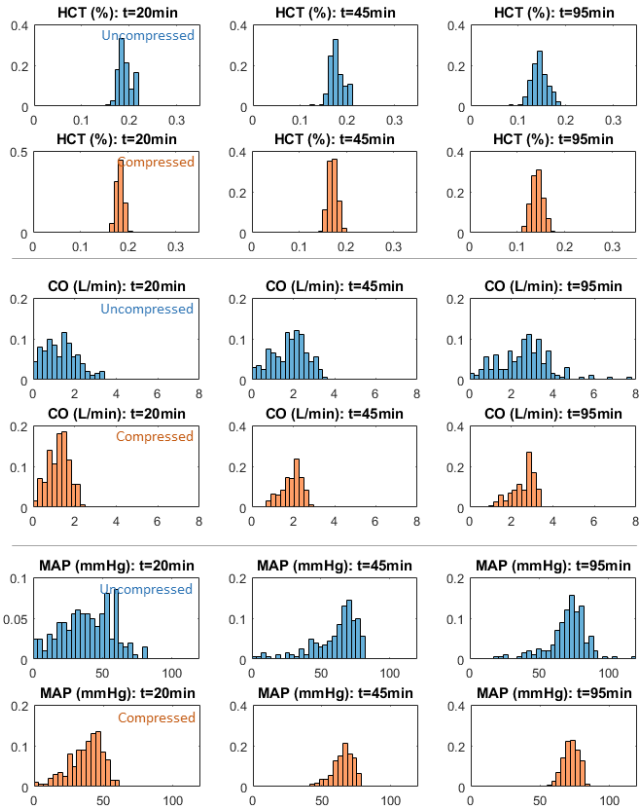

Figure 5. Representative histograms of the HCT, CO, and MAP responses from the virtual subjects: the comparison is made between subjects that were generated using compressed vs uncompressed parameter values.

subjects have a CO at or near zero at t=20min, and similarly, (ii) many have a MAP near zero at t=20min, and (iii) some virtual subjects have a cardiac output greater than 5 at t=95min, which is also considered un-realistic. Interestingly, in the compressed case, most of these un-realistic subjects vanish from the generated population. The potential reason behind this advantage is the following: In the uncompressed calibration case, the sloppiness property [15] of the model structure with respect to available data causes unnecessary drifts in some parametric directions (visible in Fig. 4(a) and also Fig. 3 at λ=0). A sampling method that uses these values would also sample in-between the drifted values and thus create un-realistic virtual subjects. In contrast, latent space compression prevents any such unnecessary deviations from the group-average model, resulting in a lower number of un-realistic virtual subjects.

Fig. 6 compares the quality of the generated cohort for different compression weights, using the score introduced in equation (20). At λ=0, where there is no compression, the score shows a poor value, which indicates that the probability of generating realistic (as compared to the data) patients in the uncompressed case is relatively low. As λ increases above zero, the score improves, which corresponds to using the compressed method for virtual patient generation. Finally, for large values of λ, most generated virtual patients will become too similar to the group-average

model, which again results in a poor score. As a result, it is beneficial to pick a moderate compression weight according to Fig. 3 for the purpose of model calibration and virtual subject generation.

V. Conclusion

In this paper, we investigated the data-driven generation of virtual patients using physiological models. For this purpose, a parameter space compression method and a virtual patient generation method were proposed and applied to a practically important case study on the physiological modeling of cardiovascular responses to hemorrhage (bleeding) and fluid resuscitation. The results suggested the validity of the proposed approach: A set of virtual patients generated using the proposed compressed sampling method showed higher similarity to a real dataset when compared to the uncompressed sampling case. Furthermore, unlike the uncompressed sampling case, the compressed sampling method generated fewer virtual patients with unrealistic behavior. Future effort should be devoted to the investigation of the advantages and limitations of the proposed method for a wider range of physiological modeling applications, and to the possibility of utilizing more advanced calibration and compression techniques to further improve the quality of the generated virtual patients.

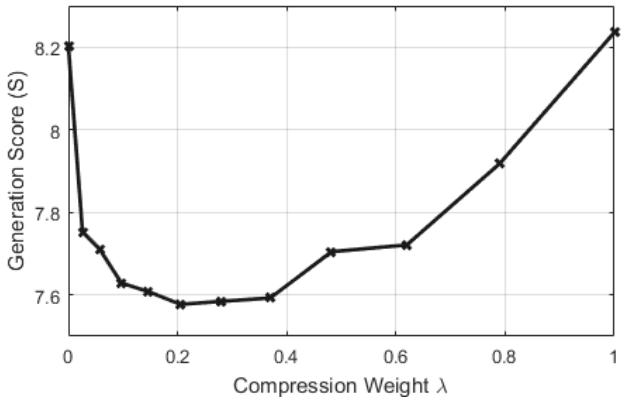

Figure 6. The quality of virtual subjects generated by sampling from a compressed latent space with compression weight λ (lower is better).

References

[1] F. Pappalardo, G. Russo, F. M. Tshinanu, and M. Viceconti, "In silico clinical trials: concepts and early adoptions," *Brief. Bioinform.*, Jun. 2018.

[2] C. D. Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli, "The UVA/PADOVA Type 1 Diabetes Simulator: New Features.," *J. Diabetes Sci. Technol.*, vol. 8, no. 1, pp. 26–34, Jan. 2014.

[3] R. Bighamian, B. Parvinian, C. G. Scully, G. Kramer, and J.-O. Hahn, "Control-oriented physiological modeling of hemodynamic responses to blood volume perturbation.," *Control Eng. Pract.*, vol. 73, no. April 2018, pp. 149–160, Apr. 2018.

[4] H. Mirinejad *et al.*, "Evaluation of Fluid Resuscitation Control Algorithms via a Hardware-in-the-Loop Test Bed," *IEEE Trans. Biomed. Eng.*, pp. 1–1, 2019.

[5] K. Jang *et al.*, "Computer Aided Clinical Trials for Implantaule Cardiac Devices," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2018, pp. 1–4.

[6] X. Jin and J.-O. Hahn, "Semi-adaptive switching control for infusion of two interacting medications," *Biomed. Signal Process. Control*, vol. 43, pp. 183–195, May 2018.

[7] O. Faris and J. Shuren, "An FDA Viewpoint on Unique Considerations for Medical-Device Clinical Trials," *N. Engl. J. Med.*, vol. 376, no. 14, pp. 1350–1357, Apr. 2017.

[8] R. Moss, T. Grosse, I. Marchant, N. Lassau, F. Gueyffier, and S. R. Thomas, "Virtual patients and sensitivity analysis of the Guyton model of blood pressure regulation: Towards individualized models of whole-body physiology," *PLoS Comput. Biol.*, vol. 8, no. 6, Jun. 2012.

[9] A. Mousavi, A. Tivay, B. Finegan, M. S. McMurtry, R. Mukkamala, and J.-O. Hahn, "Tapered vs. Uniform Tube-Load Modeling of Blood Pressure Wave Propagation in Human Aorta," *Front. Physiol.*, vol. 10, no. JUL, p. 974, Aug. 2019.

[10] M. Doosthosseini, E. Hansen, and H. K. Fathy, "On the accuracy of drug-resistant cell population estimation from total cancer size measurements," in *2019 18th European Control Conference, ECC 2019*, 2019, pp. 343–350.

[11] M. Davidian and D. M. Giltinan, "Nonlinear models for repeated measurement data: An overview and update," *J. Agric. Biol. Environ. Stat.*, vol. 8, no. 4, pp. 387–419, Dec. 2003.

[12] D. Brown *et al.*, "Trauma in silico: Individual-specific mathematical models and virtual clinical populations," *Sci. Transl. Med.*, vol. 7, no. 285, p. 285ra61-285ra61, 2015.

[13] A. Haidar, M. E. Wilinska, J. A. Graveston, and R. Hovorka, "Stochastic Virtual Population of Subjects With Type 1 Diabetes for the Assessment of Closed-Loop Glucose Controllers," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 12, pp. 3524–3533, Dec. 2013.

[14] S. Barish, M. F. Ochs, E. D. Sontag, and J. L. Gevertz, "Evaluating optimal therapy robustness by virtual expansion of a sample population, with a case study in cancer immunotherapy," *Proc. Natl. Acad. Sci.*, vol. 114, no. 31, pp. E6277–E6286, Aug. 2017.

[15] M. K. Transtrum, B. B. Machta, K. S. Brown, B. C. Daniels, C. R. Myers, and J. P. Sethna, "Perspective: Sloppiness and emergent theories in physics, biology, and beyond," *J. Chem. Phys.*, vol. 143, no. 1, p. 010901, Jul. 2015.

[16] A. Raue *et al.*, "Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood," *Bioinformatics*, vol. 25, no. 15, pp. 1923–1929, Aug. 2009.

[17] A. Tivay, G. Arabi Darreh Dor, R. Bighamian, G. C. Kramer, and H. Jin-Oh, "A Regularized System Identification Approach to Subject-Specific Physiological Modeling with Limited Data," in *2019 American Control Conference (ACC)*, 2019, pp. 3468–3473.

[18] P. Pathmanathan, J. M. Cordeiro, and R. A. Gray, "Comprehensive Uncertainty Quantification and Sensitivity Analysis for Cardiac Action Potential Models.," *Front. Physiol.*, vol. 10, p. 721, 2019.

[19] R. Bighamian, A. T. Reisner, and J. Hahn, "A Lumped-Parameter Subject-Specific Model of Blood Volume Response to Fluid Infusion," *Front. Physiol.*, vol. 7, p. 390, 2016.

[20] J. E. (John E. Hall, *Guyton and Hall textbook of medical physiology*. .

[21] Y. Çınar, G. Demir, M. Paç, and A. B. Çınar, "Effect of hematocrit on blood pressure via hyperviscosity," *Am. J. Hypertens.*, vol. 12, no. 7, pp. 739–743, Jul. 1999.

[22] A. D. Rafie *et al.*, "Hypotensive Resuscitation of Multiple Hemorrhages Using Crystalloid and Colloids," *Shock*, vol. 22, no. 3, pp. 262–269, Sep. 2004.

[23] S. U. Vaid *et al.*, "Normotensive and hypotensive closed-loop resuscitation using 3.0% NaCl to treat multiple hemorrhages in sheep," *Crit. Care Med.*, vol. 34, no. 4, pp. 1185–1192, Apr. 2006.