CAREER: Enabling “White-Box” Autonomy in Medical Cyber-Physical Systems

Jin-Oh Hahn (PI, University of Maryland), Ali Tivay (PhD Student, University of Maryland)
Christopher G. Scully (Co-I, U.S. Food and Drug Administration)

Interpretable Modeling of Cardiovascular Responses to Fluid Perturbation

1. Macroscopic View of the Circulatory System

- The circulatory nature of the cardio-vascular system is modeled as a loop.
- Through this loop, the heart pumps blood with rate Q (cardiac output).
- The resistance R (systemic vascular resistance) resists the flow of blood.
- Compartment fluid volumes determine mean arterial and venous pressures.
- Hemorrhage causes volume loss, while resuscitation causes volume expansion.

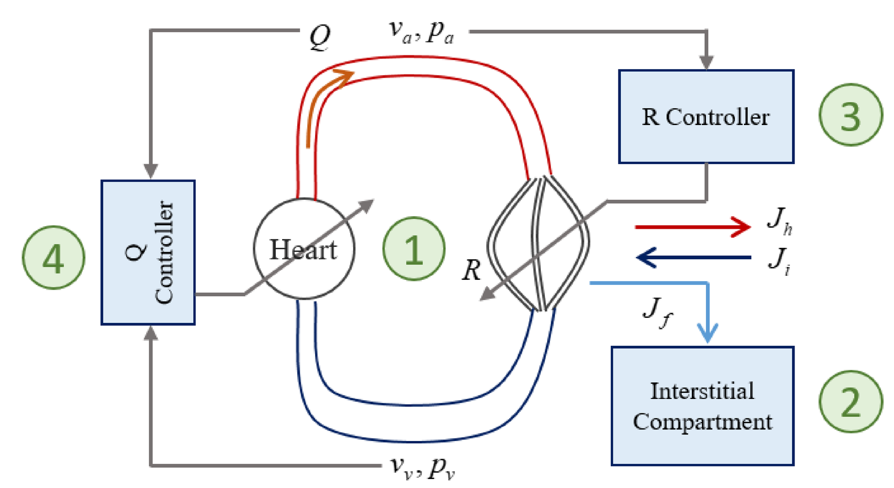

2. Fluid Shift (J_f) Controller

- Perturbation in blood volume is counteracted by an exchange of fluid with the interstitium.
- The net rate of fluid shift is modeled as a control input to regulate total blood volume.

3. Systemic Vascular Resistance (R) Controller

- The body can change the resistance R through vasoconstriction and vasodilation.
- The resistance *R* is modeled as a control input that has the goal of regulating arterial pressure.
- This control input is also disturbed by changes in the fraction of red blood cells in the blood.

4. Cardiac Output (Q) Controller

- The cardiac output *Q* is modeled as a controlled variable through heart rate and contractility.
- This controller is disturbed by changes in cardiac preload and loss of red blood cells.

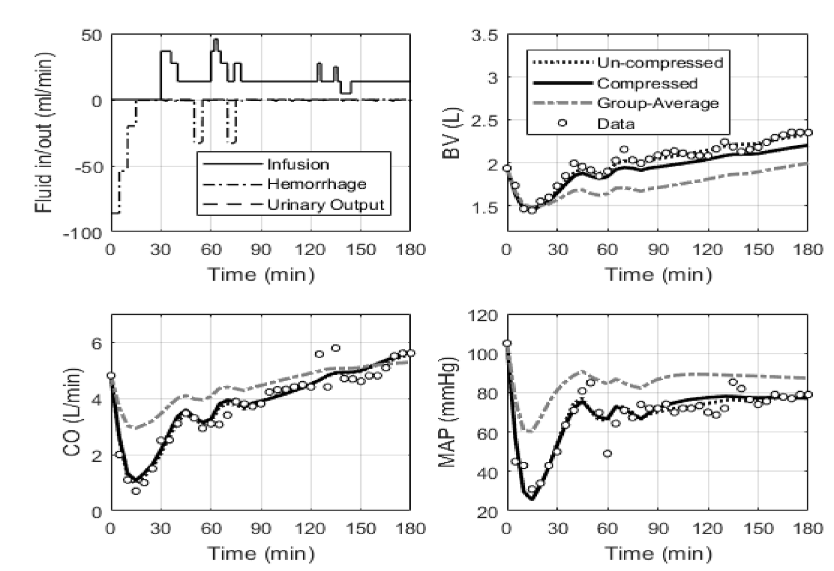

Compressed Patient-Specific Calibration of Physiological Models

1. Population-Average Model Calibration

- Diverse data from a population of subjects are used to calibrate the physiological model.
- The resulting model represents typical behavior in the population.

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

2. Compressed Patient-Specific Model Calibration

- Local parameter-output covariances (**C**) are computed around the population-average model.
- A singular value decomposition of **C** gives the directions of maximum parameter-output covariance.
- The model can be calibrated to each subject in a compressed way as follows:

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

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

The results suggest that:

- The model can be calibrated in a compressed way, to match the data from different patients.
- The variability across different subjects can be represented in the compressed latent space ϕ.

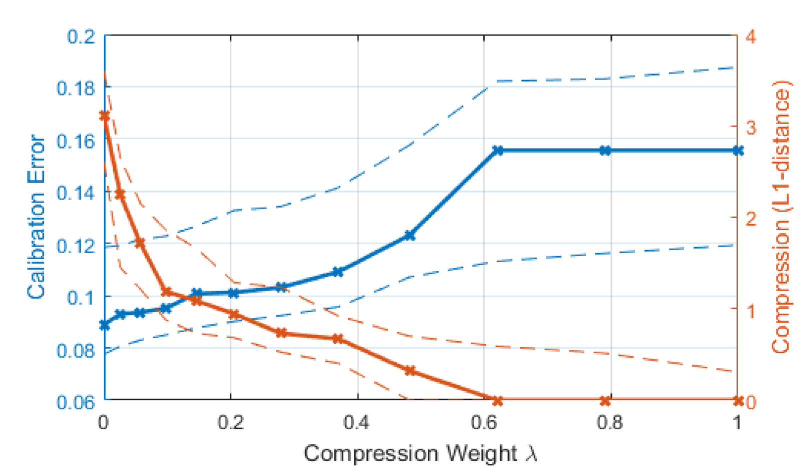

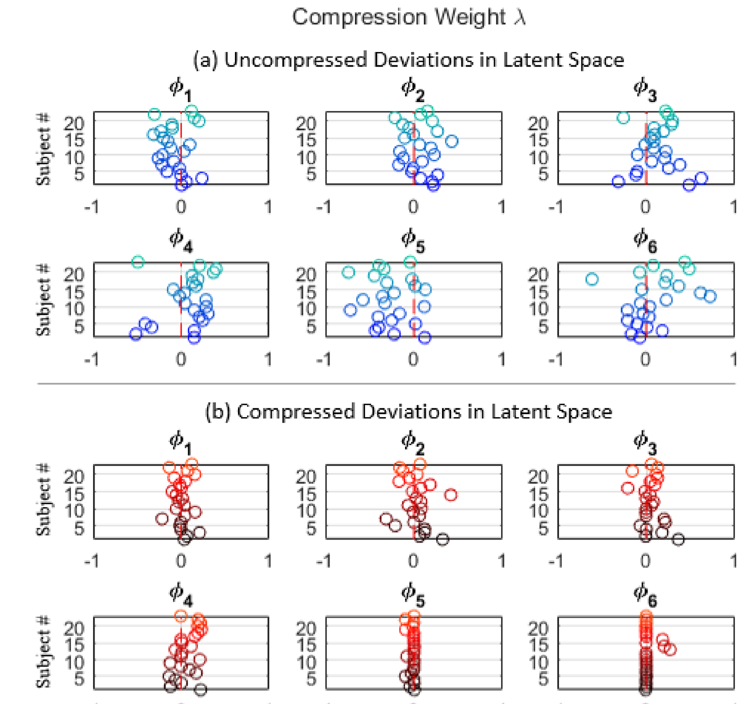

Data-Driven Generation of Virtual Patients using Physiological Models

A. Variational Approach to Virtual Patient Generation

1. Family of Distributions Over Latent Parameters ϕ

- A distribution from the mean-field family is selected in the compressed latent space.
- Each dimension $p_j(\phi_j)$ is characterized by a normal distribution.
- Parameters for $p_j(\phi_j)$ are estimated from compressed calibration results in each latent dimension.

$$P(\phi) = \prod_{j=1}^{n_p} p_j(\phi_j)$$

2. Sample from P(ϕ) to Generate Virtual Patients

- Each sample $\phi^{(i)}$ can be converted to model parameter $\theta^{(i)}$ to simulate the virtual patient.

$$\theta^{(i)} = V\phi^{(i)} + \bar{\theta}$$

The results suggest that:

- Estimating P(ϕ) *without compression* results in generation of *many* un-realistic virtual patients.
- Estimating P(ϕ) *with compression* omits the un-realistic virtual patients.
- With compression, the range of observed variations in virtual patient behavior are more realistic.

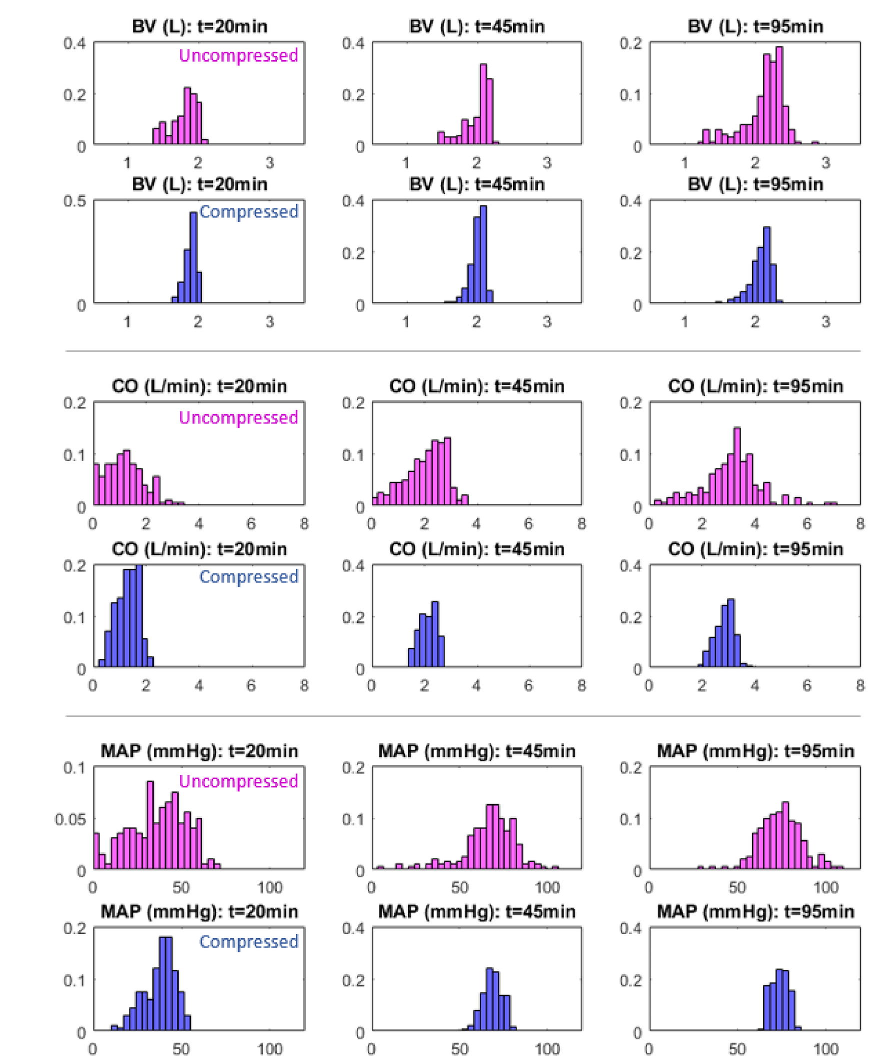

B. Markov-Chain Monte-Carlo (MCMC) Approach to Virtual Patient Generation

1. The Virtual Patient Generation Problem

- To successfully use MCMC methods for virtual patient generation, two important elements must be designed:
- A *likelihood function*, that represents how likely it is for a proposed virtual patient to behave like a real subject from the dataset.
- A *prior distribution*, that encodes prior knowledge about “reasonable” parameter samples.

2. The Likelihood Function

- S is the set of real subjects in data, and S_i (i=1,…,m) random subsets of S with size K (m is large).
- We accept a proposed virtual patient as likely if its behavior resembles any K patients in the dataset.

$$L(\theta) = \frac{1}{m} \sum_{i=1}^{m} \prod_{j \in S_i} P(\mathbf{Y}_j \mid \mathbf{I}_j, \theta)$$

3. The Prior Distribution

- The regularizer used in compressed model calibration can be thought of as a Laplace prior.
- Thus, a Laplace prior (in latent space) is used, which is centered at the population average model.

The results suggest that:

- MCMC generates virtual patients that behave realistically and cover the range of observed behavior.
- The MCMC-sampled parameter values lie within their reasonable physiological ranges.

Observed Range of Behavior in Generated Patients (red) VS Identified Patients (blue)

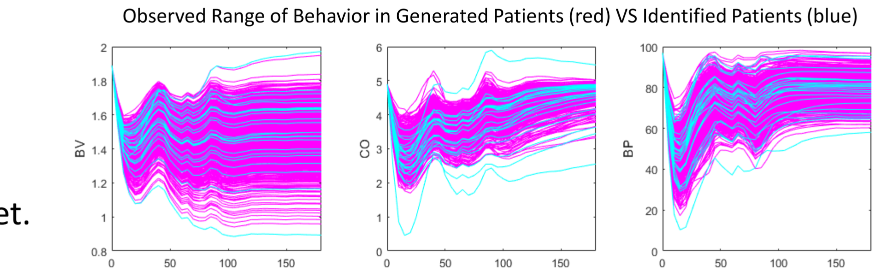

Histograms of Marginal Distributions for MCMC-Sampled Model Parameters

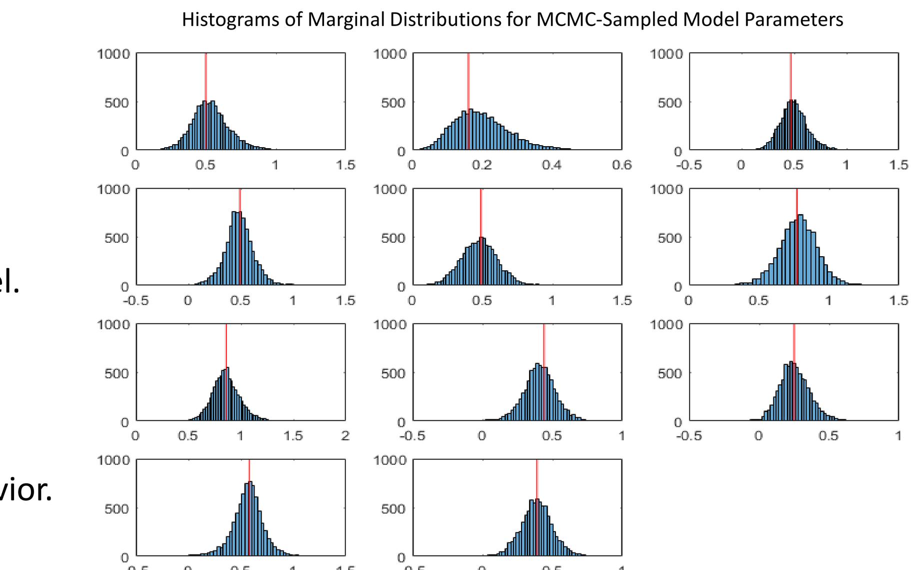

2019 NSF Cyber-Physical Systems Principal Investigators' Meeting
November 21-22, 2019 | Crystal City, Virginia

Award ID#: CNS-1748762