ABSTRACT

Title of Dissertation: INFERENCE-BASED MODELING, MONITORING, AND CONTROL ALGORITHMS FOR AUTONOMOUS MEDICAL CARE

Ali Tivay, Doctor of Philosophy, 2022

Dissertation directed by: Dr. Jin-Oh Hahn, Associate Professor, Department of Mechanical Engineering

Autonomous medical care systems are relatively recent developments in biomedical research that aim to leverage the vigilance, precision, and processing power of computers to assist (or replace) humans in providing medical care to patients. Indeed, past research has demonstrated initial promise for autonomous medical care in applications related to anesthesia, hemodynamic management, and diabetes management, to name a few. However, many of these technologies yet do not exhibit the maturity necessary for widespread real-world adoption and regulatory approval. This can be attributed, in part, to several outstanding challenges associated with the design and development of algorithms that interact with physiological processes. Ideally, an autonomous medical care system should be equipped to exhibit (i) *transparent behavior,* where the system's perceptions, reasoning, and decisions are human-interpretable; (ii) *context-aware behavior,* where the system is capable of remaining mindful of contextual and peripheral information in addition to its primary goal; (iii) *coordinated behavior*, where the system can coordinate multiple actions in

synergistic ways to best achieve multiple objectives; (iv) *adaptable behavior,* where the system is equipped to identify and adapt to variabilities that exist within and across different patients; and (v) *uncertainty-aware behavior,* where the system can handle imperfect measurements, quantify the uncertainties that arise as a result, and incorporate them into its decisions. As these desires and challenges are specific to autonomous medical care applications and not fully explored in past research in this area, this dissertation presents a sequence of methodologies to model, monitor, and control a physiological process with special emphasis on addressing these challenges. For this purpose, first, a collective variational inference (C-VI) method is presented that facilitates the creation of personalized and generative physiological models from low-information and heterogeneous datasets. The generative physiological model is of special importance for the purposes of this work, as it encodes physiological knowledge by reproducing the patterned randomness that is observed in physiological datasets. Second, a population-informed particle filtering (PIPF) method is presented that fuses the information encoded in the generative model with real-time clinical data to form perceptions of a patient's states, characteristics, and events. Third, a population-informed variational control (PIVC) method is presented that leverages the generative model, the perceptions of the PIPF algorithm, and user-defined definitions of actions and rewards in order to search for optimal courses of treatment for a patient. These methods together form a physiological decision-support and closed-loop control (PCLC) framework that is intended to facilitate the desirable behaviors sought in the motivations of this work. The performance, merits, and limitations of this framework are analyzed and discussed based on clinically-important case studies on fluid resuscitation for hemodynamic management.

INFERENCE-BASED MODELING, MONITORING, AND CONTROL ALGORITHMS FOR AUTONOMOUS MEDICAL CARE

by

Ali Tivay

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2022

Advisory Committee:
Professor Jin-Oh Hahn, Chair
Professor Perinkulam S. Krishnaprasad
Professor Hosam Fathy
Professor Nikhil Chopra
Professor Yancy Diaz-Mercado

Dedication

To my beloved family: my mother Ashraf, my father Ali-Asghar, and my sister Shadi.

To Rumi, whose spiritual legacy guides me through these confusing times.

To the friendly One, who blesses me with inspiration when I need it the most.

Acknowledgements

I would like to express my warmest thanks to my doctoral advisor Prof. Jin-Oh Hahn for his invaluable guidance and support in bringing this work to fruition. This dissertation would not have been possible without his encouragement and insight, and I will always be striving to emulate his admirable work ethic, consistency, and intellectual curiosity in my future career. I would also like to express my gratitude to Prof. Perinkulam S. Krishnaprasad, Prof. Hosam Fathy, Prof. Nikhil Chopra, and Prof. Yancy Diaz-Mercado for serving on my dissertation committee and for their valuable inputs and discussions on this work.

I would like to thank Prof. George C. Kramer of the University of Texas Medical Branch for the provision of experimental data and physiological insight for the case studies presented in this work. I would also like to thank Dr. Christopher G. Scully and Dr. Ramin Bighamian of the U.S. Food and Drug Administration for their valuable inputs and discussions on this work.

I would like to express my appreciation to the U.S. National Science Foundation (Grant No. 1748762; Grant No. 1760817) and the U.S. Congressionally Directed Medical Research Programs (Grant No. W81XWH-19-1-0322) for providing research funding for projects that inspired and resulted in the contents of this dissertation.

During my years at UMD, I had the pleasure of knowing and working with the brilliant students and members of the Laboratory for Control and Information Systems, Alex Lo, Azin Sadat Mousavi, Bahram Parvinian, Drew Hohenhaus, Ghazal ArabiDarrehDor, Jesse Parreira, Mohammad Alsalti, Peyman Yousefian, Sungtae Shin,

Weidi Yin, Xin Jin, Yekanth Ram, Yi-Ming Kao, and Zahra Ghasemi. I will always be grateful for their collaboration, friendship, and support.

Studying in the United States as an international student came with its own set of excitements and challenges, and I am deeply grateful to my dear friends Amin Alimardani, Yazdan Movahedi, Peyman Yousefian, Nariman Niknejad, Mehdi Kohani, Fatholah Salehi, Samrad Babaee, Roohollah Heidary, Hossein Souri, Miead Nikfarjam, Amir Mortazavi, Edris Shirvani, and MirSaleh Bahavarnia for their friendship, encouragement, and support along the way. I wish them all the best in their personal and professional lives.

Finally, I would like to express my deepest gratitude to my mother, my father, and my sister for their unconditional love and support throughout the years. This dissertation would not have been possible without their continuous encouragement.

Table of Contents

List of Tables

Tables:

Algorithms:

List of Figures

Chapter 1: Introduction to Autonomous Medical Care Systems

1.1. Background

The human body is susceptible to a broad range of physical illnesses. In some instances, especially in case of critical illness (e.g., severe infections, physical trauma, and cardiovascular diseases), patients need to undergo continuous observation and treatment from trained clinicians. The clinician's toolbox includes a wide range of therapies that can be used to treat critical patients. For instance, in the context of hemodynamic management, which is an integral part of critical care, treatments involve the infusion of substances (e.g., fluids and drugs) into the bloodstream with the intention of stabilizing a patient's vitals and helping their recovery [1]–[4]. Fluids are usually infused with the purpose of increasing a patient's circulating blood volume (as well as blood pressure) and maintaining the perfusion of blood into their organs and tissues. The commonly used fluids in critical care include crystalloids (e.g., saline, Rigner's lactate), colloids (e.g., albumin, dextran), and blood products (e.g., plasma, red blood cells). Drugs are usually infused with the purpose of altering the patient's physiological state. The commonly used drugs in critical care include vasoactive medications (e.g., vasopressors and inotropic agents), anesthetic agents, and pain management medications (e.g., opioids) [5]–[8].

Naturally, the physiological state of a critically ill patient can vary significantly over time. As a result, administering life-critical therapies demands continuous vigilance from the clinicians. However, successful accomplishment of this task is

practically limited by several important factors. First, the time and resources at a clinician's disposal are inherently limited, and the clinician may not succeed in administering all therapies at the right time. Second, there is always potential for human error in the administration of the therapies. Third, different patients show different responses to the same therapy, which makes it difficult for a clinician to administer the optimal therapy for any given patient. Previously published research suggests that these challenges are indeed prevalent in practice, subjecting critical patients to unnecessary risk. For instance, maintaining adequate blood pressure is an integral part of care for a variety of illnesses including circulatory shock, sepsis, trauma, and the effects of surgery. However, prior studies have shown that such patients routinely suffer from off-target blood pressure values during the course of their care [9], [10], which may have an adverse effect on the patient's illness and cause subsequent complications [11]–[13]. As another relevant example, recovering blood volume and blood composition through fluid infusions is an important part of burn care. However, it is well-documented that burn patients are frequently over-infused with fluids, which can increase the risk of severe edema and organ failure [14]–[16].

Autonomous physiological monitoring and medical intervention can potentially provide substantial improvements to the safety and effectiveness of medical care by making recommendations and/or performing interventions in a continuous, precise, and personalized manner. A promising approach toward realizing this level of autonomy in medical care is to conceptualize physiological processes as dynamic systems. In this way, patients can be perceived as dynamic systems with input signals (e.g., infusion rates for fluids and drugs) and output signals (e.g., blood pressure measurements),

defined in relation to the physiological problem at hand. Therefore, in theory, the vast array of existing engineering methods and tools (including dynamic systems analysis, system identification, filtering and estimation, and control theory) can be leveraged toward the realization of autonomy in medical care systems. If successful, existing knowledge in dynamic systems and control engineering can be utilized to obtain mathematical models of physiological processes, estimate important physiological states/parameters in a patient, and propose decision-support and closed-loop control systems for autonomous medical intervention.

1.2. Autonomous Medical Care in Past Literature

In recent years, aspiration toward designing effective autonomous medical care systems has motivated a considerable body of research on the modeling of physiological processes, and the design and testing of decision-support and closed-loop control systems for autonomous medical intervention [17]–[22]. More specifically, automating therapies associated with the care of critically ill patients has been a significant area of interest, including research on autonomous administration of anesthesia, autonomous fluid therapy for burn and hemorrhage resuscitation, and autonomous administration of vasoactive medications for hemodynamic management [19], [20], [22].

Autonomous systems for anesthesia have the goal of inducing a loss of sensation and/or awareness in patients for medical purposes, using medications that induce hypnosis, analgesia, and neuromuscular block [23]. The physiological processes relevant to this application can be modeled mathematically by considering

the kinetics of the drug substance traveling to its effect site within the body (called pharmacokinetics, i.e., PK), and the mechanism of action for the drug at its effect site (called pharmacodynamics, i.e., PD). A PK model thus relates drug infusions to the effect-site concentration of the infused drug, while a PD model relates effect site concentrations to measurable drug effects [24]–[27]. Having such a framework for conceptualizing drug responses as a dynamic system, researchers have employed an array of control engineering techniques to realize autonomous medication infusion in the context of anesthesia. This includes "model-free" approaches such as PID, rule-based, and reinforcement learning controllers [28]–[30], as well as "model-based" approaches such as adaptive and model-predictive controllers [31]–[39]. In addition, researchers have devoted notable effort toward designing control systems that infuse multiple drugs (e.g., propofol and remifentanil) to induce simultaneous hypnosis and analgesia. This is achieved both through designing multiple control loops that run independently [40]–[42], and, more recently, through unified control loops that consider the interactions between multiple medications [39], [43]. The merits of autonomous anesthesia have been demonstrated in simulation studies, animal experiments, and most importantly, randomized controlled trials conducted on human subjects. The results suggest that these systems can deliver superior control of anesthesia (e.g., in terms of BIS) compared to manual administration, indicating the feasibility and the general future promise of autonomous medical care systems [22], [23], [44].

Autonomous systems for hemodynamic management are tasked with maintaining blood volume and blood pressure in the circulatory system such that all

vital organs are continuously supplied with adequate blood. Hemodynamic management is necessary in a wide range of critical care scenarios, including, but not limited to, physical trauma, hemorrhage, and burn injury [45], [46]. Similar to the anesthesia case, the patient's hemodynamic physiology could in theory be viewed as a dynamic system that receives fluid/drug infusions and responds with changes in blood volume, cardiac output, and blood pressure. In the context of hemorrhage resuscitation, researchers have used a variety of model-free control methods including PID, rule-based, fuzzy-logic, and black-box machine learning, in order to successfully control aspects of a patient's hemodynamics [47]–[51]. Recently, model-based control methods have also been developed for hemorrhage resuscitation [52]. In the context of burn resuscitation, closed-loop and decision-assist systems have been devised based on rules derived from expert knowledge, which has shown promising results in terms of maintaining urinary output in a target range (which is important for favorable outcomes in burn patients) [53]–[56]. In the context of vasopressor infusion for hypotensive patients, recent research has demonstrated the feasibility and the preliminary effectiveness of controllers from the PID family in maintaining the patient's mean arterial pressure within a target range [57]–[61]. Recently, hemodynamic management using multiple independent loops has received attention from the researchers, where the autonomous system is tasked with multiple goals (e.g., anesthesia and fluid management) to be performed using multiple inputs (e.g., fluids and drugs), and preliminary studies show that these tasks may indeed be feasible [61]–[63]. Overall, these results indicate that, within controlled settings, computers can provide superior

hemodynamic management compared to manual administration, which points to the future promise of autonomous systems for medical care [45], [46], [64].

1.3. Outstanding Challenges

As presented in the previous section, past research has demonstrated the promise of autonomous medical care especially in the administration of anesthesia and hemodynamic management in critical patients. However, such technologies still await widespread regulatory approval and ubiquitous real-world adoption in clinical settings [65], [66]. This can be attributed, in part, to several outstanding challenges associated with the design and development of algorithms that interact with physiological processes: Ideally, an autonomous medical care system should be equipped to exhibit (i) *transparent behavior,* where the system's perceptions, reasoning, and decisions are human-interpretable; (ii) *context-aware behavior,* where the system is capable of remaining mindful of contextual and peripheral information in addition to its primary goal; (iii) *coordinated behavior*, where the system can coordinate multiple actions in synergistic ways to best achieve multiple objectives; (iv) *adaptable behavior,* where the system is equipped to identify and adapt to variabilities that exist within and across different patients; and (v) *uncertainty-aware behavior,* where the system can handle low-information and heterogeneous measurements, quantify the uncertainties that arise as a result, and incorporate them into its decisions. Some additional aspects of these challenges are further discussed in this section.

1.3.1. Low-Information and Heterogeneous Measurements

Physiological datasets are essential for obtaining mathematical models of physiological processes, and real-time physiological measurements are essential for realizing autonomous systems for physiological monitoring and medical intervention. However, acquiring high-quality physiological measurements is in general a challenging task. For example, acquiring a dataset for the purpose of hemorrhage resuscitation modeling involves applying stimuli (e.g., hemorrhage and fluid infusions) to subjects in laboratory settings, while measuring multiple physiological variables such as hematocrit, cardiac output, and mean arterial pressure over time. Such experimental datasets present several important limitations. First, each experiment typically provides *low information*, where (i) stimuli only partially excite the underlying physiological dynamics, (ii) insufficient measurements are available in each experiment, and (iii) the measured signals are of relatively low quality (e.g., in terms of noise and sampling rate). Second, the experiments are typically *heterogeneous*, in the sense that there are variations in (i) experimental protocols (e.g., shape/timing of stimuli), (ii) measured variables, and (iii) subject characteristics, including the possibility of atypical responses to stimuli. These limitations become even more prominent in real-world clinical settings, where measurements tend to be of lower quality than laboratory research experiments. In addition, certain measurement techniques that are available in laboratory settings may be unavailable in clinical settings due to their invasive or expensive nature. If not treated carefully, a low-information and heterogeneous physiological dataset can severely limit the predictive value of the corresponding physiological models. Similarly, low-information and

heterogeneous feedback measurements can limit the applicability of existing engineering tools (e.g., filtering, estimation, and control algorithms) that rely heavily on high-fidelity models and/or high-quality feedback measurements to perform their operations. As these limitations are ubiquitous in autonomous medical care, a formal treatment of this problem is a necessary step toward successful physiological modeling, estimation, and control.

1.3.2. Transparency and Context-Awareness

In clinical settings, an autonomous medical care system would be tasked with making life-critical alterations to a patient's physiological state. As a result, real-world adoption and regulatory approval of an autonomous medical care system is contingent upon an adequate level of trust in its decisions. Therefore, *transparency* is a highly desirable feature for an autonomous system, which is achieved when a human observer (e.g., an engineer or a clinician) can interpret how and why the system makes its therapeutic decisions. Similarly, *context-awareness* is also a highly desirable feature, which is achieved when the autonomous system is mindful of the broader physiological and environmental context (e.g., patient and sensor/actuator status) in making its therapeutic decisions. However, previously proposed autonomous medical care systems (see Section 1.2), especially those built upon black-box models and model-free control algorithms, were not developed with transparency and context-awareness as a high priority. Rather, many previously proposed algorithms were developed to simply regulate a measured physiological variable without further regard to interpretation and context. For example, the overarching goal in hemodynamic

management is to stabilize and maintain a patient's homeostasis through fluid and drug infusions. However, control algorithms for this purpose are instead tasked with, for example, increasing the patient's mean arterial pressure toward a set-point, which is a detrimental reduction of the overarching goal: in case a patient's mean arterial pressure does not respond to therapy, or in case a sensor malfunctions, a lack of context-awareness would result in even more aggressive (and misguided) therapies from the algorithms. In addition, a lack of transparency would make it difficult to pinpoint the reason for any misguided decisions made by the algorithms. Not surprisingly, these limitations are a major obstacle for the widespread adoption and regulatory approval of these autonomous systems. Thus, it is essential to devote substantial effort to designing autonomous systems that operate based on human-interpretable models of physiological phenomena and make therapeutic decisions based on context-aware reasoning.

1.3.3. Coordination of Multiple Therapies

In the context of medical care, especially for critically ill patients, therapies are rarely performed in isolation. Rather, treating patients involves the administration of several therapies, many of which need to be performed simultaneously. For example, hemodynamic management may involve fluid (e.g., with colloids) and vasopressor infusions (e.g., with norepinephrine) under anesthesia (e.g., with propofol) and sedation (e.g., with remifentanil). Such therapies often interact with each other, creating synergistic or antagonistic effects with respect to the therapy goals. This is an area of great challenge and also opportunity for autonomous medical care systems. In terms

of challenges, a lack of coordination between multiple control loops can create conflicts between different therapies, which can in turn result in a sub-optimal and/or dangerous delivery of therapy to the patients. On the opportunities front, a coordinated system can in theory be designed to take advantage of the interconnected effects of different therapies to provide optimal care to the patients. However, past work on the autonomous administration of multiple therapies is scarce, and most of the proposed designs use independent control loops for this purpose. The problem of coordination between multiple therapies is rarely addressed in the existing literature. Given the importance of coordination for safe and effective medical care, a rigorous treatment of this problem may enable autonomous systems that avoid dangerous loop interactions and leverage multiple therapies to provide optimal care.

1.4. Problem Statement and Chapter Outline

As presented in Section 1.3, modeling, monitoring, and controlling a physiological process involves a set of opportunities and challenges that are unique and, for the most part, seldom explored in past research on autonomous medical care. The overarching goal of this dissertation is thus to develop methodologies for modeling, monitoring, and controlling a physiological process, with special emphasis on the capability of the proposed solutions to exhibit *transparent, context-aware, coordinated, adaptable* and *uncertainty-aware* behavior. This dissertation makes effort to present a concrete sequence of steps toward these objectives, which are briefly described below:

- *Chapter 2* will focus on building a collective variational inference (C-VI) method to facilitate the creation of personalized and generative physiological models from low-information and heterogeneous datasets. Generative physiological models are of special importance for the purposes of this work, as they encode knowledge about likely physiological behaviors by absorbing and reproducing the patterned randomness that is observed in physiological datasets.
- *Chapter 3* will focus on building a population-informed particle filtering (PIPF) method that fuses the knowledge encoded in the generative physiological model with real-time clinical data to form perceptions of a patient's states, characteristics, and events. This method is expected to have application both as a standalone solution useful for monitoring and diagnostic purposes, and as a foundation to provide perceptions to physiological control algorithms.
- *Chapter 4* will focus on building a population-informed variational control (PIVC) method that leverages the generative model, the perceptions of the PIPF algorithm, and user-defined definitions of actions and rewards in order to search for optimal courses of treatment for a patient.

The abovementioned methods together form a physiological decision-support and closed-loop control (PCLC) framework that is intended to facilitate the desirable behaviors sought in the motivations of this work. As a running case study throughout the dissertation, the performance, merits, and limitations of the proposed algorithms and frameworks are analyzed and discussed based on a clinically-important case study on fluid resuscitation for hemodynamic management.

Chapter 2: A Collective Variational Inference Method for Personalized and Generative Physiological Modeling

Individual physiological experiments typically provide useful but incomplete information about a studied physiological process. As a result, inferring the unknown parameters of a physiological model from experimental data is often challenging. The objective of this chapter is to propose and illustrate the efficacy of a collective variational inference (C-VI) method, intended to reconcile low-information and heterogeneous data from a collection of experiments to produce robust personalized and generative physiological models. To derive the C-VI method, we utilize a probabilistic graphical model to impose structure on the available physiological data, and algorithmically characterize the graphical model using variational Bayesian inference techniques. To illustrate the efficacy of the C-VI method, we apply it to a case study on the mathematical modeling of hemorrhage resuscitation. In the context of hemorrhage resuscitation modeling, we investigate whether the C-VI method can reconcile heterogeneous combinations of hematocrit, cardiac output, and blood pressure data across multiple experiments to obtain (i) robust personalized models along with associated measures of uncertainty and signal quality, and (ii) a generative model capable of reproducing the physiological behavior of the population [67][1].

---

[1] This chapter is adapted from a journal publication by the author of this dissertation. Tivay, Ali, George C. Kramer, and Jin-Oh Hahn. "Collective Variational Inference for Personalized and Generative Physiological Modeling: A Case Study on Hemorrhage Resuscitation." *IEEE Transactions on Biomedical Engineering* (2022).

2.1. Background and Literature Review

Autonomous physiological monitoring and medical intervention can potentially provide substantial improvements to the safety and effectiveness of medical care by making recommendations and/or performing interventions in a continuous, precise, and personalized manner [19]–[22]. In recent years, this potential has motivated a considerable body of research on the design and testing of decision-support and closed-loop control systems for medical intervention [17], [18], [23], [30], [42], [51], [52], [54], [56], [61], [62], [68], [69]. Yet, regulatory approval and widespread real-world adoption of these technologies necessitate further advancements in the state of the art in terms of patient safety, physiological interpretability, awareness of physiological context, and the ability to coordinate multiple therapeutic objectives associated with multiple physiological outputs [65], [66]. Arriving at an interpretable, context-aware, and coordinated autonomous medical care system is highly contingent upon representative mathematical models of the relevant physiological mechanisms [70]–[72]. These mathematical models, however, are usually only determined up to a set of latent (i.e., unknown) parameters that must be inferred from experimental data.

A versatile array of methods has been proposed in the literature that can be utilized to infer the unknown parameters of a mathematical model. Maximum-Likelihood Estimation (MLE) is a popular technique that can be used to find point estimates for model parameters by maximizing the likelihood of observed data [73], while Bayesian inference techniques can provide posterior beliefs about model parameters based on prior beliefs and observed data. These posterior beliefs can in turn be used to extract point estimates for model parameters and quantify parameter

uncertainties [74], [75]. Obtaining posterior beliefs for model parameters is in general a non-trivial mathematical problem. However, many effective numerical solutions have been proposed for this purpose: the Markov Chain Monte Carlo (MCMC) class of algorithms such as Metropolis-Hastings and Hamiltonian-Monte-Carlo can provide high-fidelity samples from the posterior [74], [76], [77], while approaches such as Approximate Bayesian Computation can provide approximate samples from the posterior [78], [79]. In recent years, statistics and machine learning research has shown notable advances in Variational Inference (VI) techniques, through which it is possible to obtain analytical approximations to the posterior using optimization [80], [81]. VI algorithms tend to require fewer computations than MCMC, while stochastic and amortized variants of VI provide the opportunity to handle larger datasets and more complex problem formulations [82], [83].

In addition to the method of inference, the formulation of the inference problem has notable effects on the fidelity of the resulting mathematical models. For instance, in case experimental data are obtained from multiple non-identical subjects, the problem formulation must account for the existence of inter-subject variability. The Empirical Bayes framework and its closely related counterpart in Nonlinear Mixed-Effects Modeling are effective problem formulations for this purpose, where subject-level models (e.g., personalized models) are augmented with a population-level model (e.g., a model of inter-subject variability, or equivalently, an empirical prior) to represent the conditions under which the data were obtained [84]–[87]. In recent years, several interesting variants of such a hierarchical problem formulation have been proposed, especially for dynamic systems modeling and machine learning applications

[83], [88]–[91]. More generally, the Probabilistic Graphical Modeling (PGM) framework provides a rich set of tools for formulating effective probabilistic dependence structures for a given problem class according to problem-specific challenges and objectives [92], [93].

Inferring the unknown parameters of a physiological model from experimental data presents a unique set of challenges that appear frequently in this area of research. Acquiring data for the purpose of hemorrhage resuscitation modeling, for example, involves applying stimuli (e.g., hemorrhage and fluid infusions) to subjects while measuring their relevant physiological variables such as hematocrit (HCT), cardiac output (CO), and mean arterial pressure (MAP) over time. Such physiological experiments tend to exhibit challenging characteristics: First, each experiment provides *low information*, in the sense that (i) stimuli only partially excite the underlying physiological dynamics, (ii) insufficient physiological measurements are available in each experiment, and (iii) the measured signals are of relatively low quality (e.g., in terms of noise and sampling rate). Second, the experiments are *heterogeneous*, in the sense that there exist (i) variations in experimental protocols (e.g., shape/timing of stimuli), (ii) variations in the availability of measured variables (e.g., HCT and CO may not be available in some subjects), and (iii) variations in subject characteristics, including the possibility of atypical responses to stimuli. These challenges, if not explicitly addressed in the course of parameter inference, tend to produce physiological models with unrealistic parameter values and/or limited predictive capability [94]–[97].

The objective of this chapter is to propose and illustrate the efficacy of a *collective variational inference* (C-VI) method, intended to reconcile low-information

and heterogeneous data from a collection of experiments to obtain robust personalized and generative physiological models. The *personalized* model aims to reproduce the physiological behavior of a specific subject, while the *generative* model aims to reproduce the physiological behavior of the population. To derive the C-VI method, we compose a PGM to structurally represent the scenario in which low-information and heterogeneous experiments are conducted on a collection of non-identical subjects. Given this problem formulation, obtaining personalized and generative models for a physiological process boils down to inferring the latent parameters of this PGM structure. For this purpose, we leverage recent advances in stochastic VI to obtain an algorithmic procedure that computes approximate posteriors for the PGM parameters through stochastic optimization. To illustrate the efficacy of the C-VI method, we apply it to a practically important case study on the mathematical modeling of hemodynamic responses to hemorrhage resuscitation, and compare the models produced by the C-VI method with those produced by a non-collective method based on MLE. In this context, we demonstrate that the C-VI method can reconcile heterogeneous combinations of HCT, CO, and MAP data across multiple experiments to obtain (i) robust personalized models along with associated measures of uncertainty and signal quality, and (ii) a generative model capable of reproducing the physiological behavior of the population. Finally, we discuss how the resulting models may provide basis for the development and testing of interpretable physiological monitoring, decision-support, and closed-loop control algorithms.

2.2. Collective Variational Inference (C-VI)

In this section, we present the C-VI methodology, which is intended to enable personalized and generative modeling of physiological processes using low-information and heterogeneous data. This methodology formulates the physiological modeling problem in such a way that multiple experiments can collectively provide information to characterize a physiological system, both at the level of each subject and the population encompassing all subjects.

2.2.1. Generative Modeling for Physiological Data

In the first step toward deriving the C-VI method, we aim to formulate a generative model of the processes underlying the acquisition of low-information and heterogeneous physiological data in multiple experiments. This generative model is schematically shown in Figure 2.1(a) for one representative experiment. This model is a hierarchical model consisting of three main levels. At the highest level, a subject generator model is tasked with generating virtual subjects with varying physiological characteristics, which can be formalized as:

$$\theta_i \sim \mathcal{G}(\phi) \tag{2.1}$$

where $\mathcal{G}$ is the subject generator model, ϕ is the vector of latent parameters for the subject generator model, and θ_i denotes a generated parameter vector representing the physiological characteristics of a virtual subject.

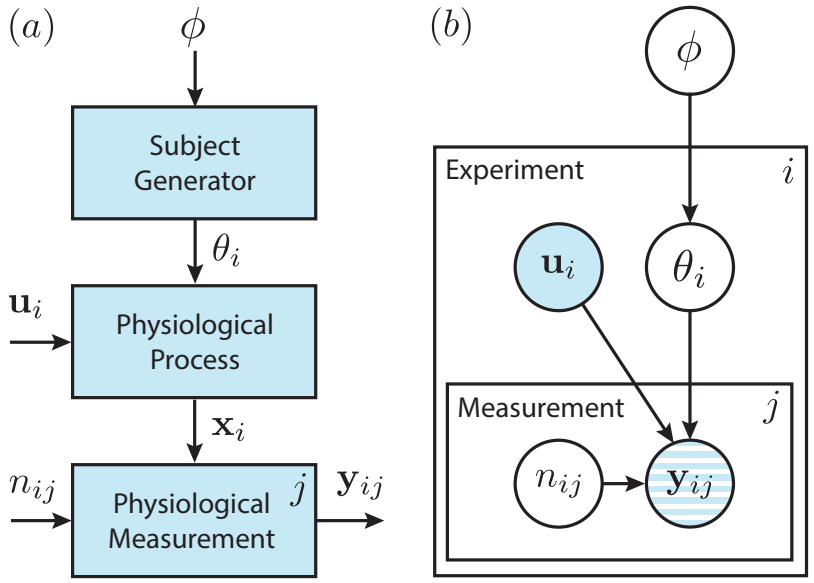

Figure 2.1. Formulation of the inference problem for personalized and generative physiological modeling. (a) Schematic view of the generative model structure for one experiment conducted on one subject sample, where several physiological variables are measured. (b) Probabilistic graphical representation of the dependencies between the latent parameters (white), the measured variables (blue), and the sometimes-measured variables (striped, blue) in the model structure.

At the second level, a virtual physiological experiment is conducted on each virtual subject, whose response to the experiment is generated by a physiological process model $\mathcal{H}$:

$$x_i(t) = \mathcal{H}\left(\theta_i, u_i(t)\right) \tag{2.2}$$

where $u_i(t)$ represents the physiological stimuli associated with the experiment on subject i, and $x_i(t)$ represents the state evolution of subject i during the course of the experiment. At the third level, in each virtual experiment, one or several physiological variables (e.g., blood pressure) are measured from the state evolution $x_i(t)$ and recorded as virtual data. Thus, for each measured variable j, we consider a physiological measurement model $\mathcal{M}_j$, which can be formalized as:

$$y_{ij}^{m} = \mathcal{M}_{j}^{m}\left(\theta_i, x_i(t)\right) \tag{2.3}$$

$$y_{ij} \sim \mathcal{M}_{j}^{\sigma}\left(n_{ij}, y_{ij}^{m}\right) \tag{2.4}$$

where y_{ij}^{m} is a vector containing the model-generated outputs for the physiological variable j in subject i, y_{ij} is a vector containing the model-generated (and possibly noisy) virtual data for this physiological variable, and n_{ij} is a latent parameter modulating the signal quality of the virtual data. In summary, the generative model structure described by (2.1)-(2.4) is built to (i) generate a virtual subject cohort of arbitrary size, (ii) conduct virtual experiments on the generated cohort, and (iii) generate virtual datasets by compiling the results of the virtual experiments. Given this model structure, our aim is to infer the parameters ϕ, θ_i, n_{ij} using real physiological data, such that the n_{ij}'s capture the signal quality in each real experiment, the θ_i's capture the physiological characteristics of the real subjects, and the subject generator with ϕ produces virtual subjects that are representative of the population.

2.2.2. Collective Inference for Generative Modeling

The generative model presented in Section 2.2.1 imposes a hierarchical relationship between the variables of interest in the physiological modeling problem. A probabilistic graphical representation of this relationship is shown in Figure 2.1(b). In this representation, the generator model parameters ϕ act as global random variables that affect all subject characteristic vectors θ_i. In addition, each subject characteristic vector θ_i together with its corresponding stimuli u_i act as local random variables with respect to their own experiment. Finally, the virtual data y_{ij} and the signal quality parameters

n_{ij} act as local random variables with respect to their own experiment and measured variable. This hierarchical relationship can be formalized using the following joint density:

$$p(\phi, \boldsymbol{\theta}, n, u, y) = p(\phi)p(\boldsymbol{\theta}|\phi)p(\boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{n}, \boldsymbol{u})p(\boldsymbol{n})p(\boldsymbol{u}) \\ = p(\phi)\prod_i[p(\theta_i|\phi)p(\boldsymbol{u_i})]\prod_{i,j}\left[p\left(\boldsymbol{y_{ij}}\middle|\theta_i, n_{ij}, \boldsymbol{u_i}\right)p\left(n_{ij}\right)\right] \tag{2.5}$$

where the symbols $\boldsymbol{\theta}$, n, u, and y respectively denote the collection of all random variables corresponding to θ_i, n_{ij}, u_i, and y_{ij}. The physiological stimuli u and the virtual data y are observed random variables, while the parameters denoted by ϕ, $\boldsymbol{\theta}$, n are latent random variables that need to be inferred using their relationship to the observed random variables. This inference objective can be expressed in probabilistic terms as calculating the following conditional density:

$$p(\phi, \boldsymbol{\theta}, \boldsymbol{n}|\boldsymbol{u}, \boldsymbol{y}) = \frac{p(\phi, \boldsymbol{\theta}, \boldsymbol{n}, \boldsymbol{u}, \boldsymbol{y})}{p(\boldsymbol{u}, \boldsymbol{y})} \tag{2.6}$$

which is the *exact posterior density*, and represents the ultimate objective of inference for the purpose of personalized and generative physiological modeling in this work. Obtaining this exact posterior is tantamount to utilizing the available data in a collective manner to obtain personalized and generative physiological models along with associated measures of uncertainty and signal quality. However, computing this exact posterior is mathematically intractable and computationally expensive for many physiological applications, which motivates the derivation of an *approximate posterior* that can be computed with reasonable accuracy and computational efficiency.

2.2.3. Variational Inference for Generative Modeling

In this work, we employ a variational approach [80], [83] to finding analytical approximations to the exact posterior in (2.6). In this approach, a family of densities over the latent variables $q(\phi, \boldsymbol{\theta}, n|\boldsymbol{\nu})$ is formulated with each member (represented by the variational parameter $\boldsymbol{\nu}$) acting as a candidate for the best approximate posterior. To develop a procedure for finding the best approximate posterior, we start from the Kullback-Leibler (KL) divergence between the exact posterior and the candidate approximate posterior:

$$D_{KL}(\boldsymbol{\nu}) = \mathbb{E}_q[\log q\,(\phi, \boldsymbol{\theta}, \boldsymbol{n}|\boldsymbol{\nu}) - \log p\,(\phi, \boldsymbol{\theta}, \boldsymbol{n}|\boldsymbol{u}, \boldsymbol{y})] \tag{2.7}$$

where the operator $\mathbb{E}_q$ represents expectation with respect to samples from the approximate posterior. Substituting the exact posterior equation (2.6) and the joint density (2.5) into (2.7), and assuming that the stimuli u are observed accurately, we obtain the following equation for the KL-divergence:

$$\begin{aligned} D_{KL}(\boldsymbol{\nu}) = \mathbb{E}_q[\log q\,(\phi, \boldsymbol{\theta}, \boldsymbol{n}|\boldsymbol{\nu}) - \log p\,(\boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{n}, \boldsymbol{u}) \\ - \log p\,(\boldsymbol{\theta}|\phi) - \log p\,(\boldsymbol{n}) - \log p\,(\phi)] + \log p\,(\boldsymbol{y}) \end{aligned} \tag{2.8}$$

In this equation, two inherently hard-to-compute terms are present: (i) the $D_{KL}(\boldsymbol{\nu})$ term, i.e., the dissimilarity between the approximate and the exact posterior densities, which needs to be minimized, and (ii) the $p(y)$ term, i.e., the model evidence, which depends only on the model definition. Putting these two terms together, we obtain the following quantity:

$$L(\boldsymbol{\nu}) = \log p\,(y) - D_{KL}(\boldsymbol{\nu}) \tag{2.9}$$

which is the *evidence lower bound* and needs to be maximized. Combining equations (2.8) and (2.9), we obtain the following expression for this objective:

$$L(\boldsymbol{\nu}) = \mathbb{E}_q[\log p\,(\boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{n}, \boldsymbol{u}) + \log p\,(\boldsymbol{\theta}|\phi) + \log p\,(\boldsymbol{n}) + \log p\,(\phi) - \log q\,(\phi, \boldsymbol{\theta}, \boldsymbol{n}|\boldsymbol{\nu})] \tag{2.10}$$

which represents the main objective for the physiological modeling problem addressed in this work.

2.2.4. Interpretation of the Objective and Special Cases

The objective shown in (2.10) is based on an expectation with respect to samples from the approximate posterior $q(\phi, \boldsymbol{\theta}, n|\boldsymbol{\nu})$, the shape of which can be modulated through the variational parameters $\boldsymbol{\nu}$. The term $\log p\,(y|\boldsymbol{\theta}, n, u)$ is a log-likelihood term that promotes similarity between observed physiological data and the model-generated virtual data for every experiment. The term $\log p\,(\boldsymbol{\theta}|\phi)$ promotes personalized models that are likely under the subject generator, and a subject generator that is likely to generate the personalized models. The terms $\log p\,(\phi)$ and $\log p\,(n)$ represent prior densities that can encode prior knowledge about the subject generator and the physiological measurement model parameters. Finally, the term $\log q\,(\phi, \boldsymbol{\theta}, n|\boldsymbol{\nu})$ promotes a diffuse approximate posterior, acting as a mechanism for uncertainty quantification. As a result of this formulation, special cases of the inference problem can be obtained by removing a subset of terms from the objective. Removing $\log q\,(\phi, \boldsymbol{\theta}, n|\boldsymbol{\nu})$ promotes a concentrated approximate posterior, resulting in a point estimation problem over the unknown parameters ϕ, $\boldsymbol{\theta}$ and n. Removing the rest of the

terms except for $\log p\,(y|\boldsymbol{\theta}, n, u)$ results in a non-collective MLE problem over $\boldsymbol{\theta}$ and n, promoting separate estimation of model parameters for every subject.

2.2.5. Stochastic Optimization Algorithm

As presented in Section 2.2.3, performing inference for the purpose of personalized and generative physiological modeling boils down to maximizing the evidence lower bound objective shown in (2.10) over the variational parameters. In the absence of further algorithm engineering, this maximization can turn into a prohibitively expensive computational task due to the presence of the expectation operator $\mathbb{E}_q$. In this work, we employ an approach based on stochastic gradients of the objective [82], [83] to perform the desired maximization in a computationally feasible manner. For this purpose, we assume a function f_q that takes as its input the variational parameters ν, and a sample ϵ from the standard normal distribution (of appropriate dimension), and produces as its output samples from the approximate posterior $q(\phi, \boldsymbol{\theta}, n|\nu)$:

$$\left[\boldsymbol{z}_{\phi}; \boldsymbol{z}_{\boldsymbol{\theta}}; \boldsymbol{z}_{\boldsymbol{n}}\right] = f_q(\boldsymbol{\epsilon}, \boldsymbol{\nu}), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, I) \tag{2.11}$$

where z_{ϕ} is a sampled subject generator model parameter, $z_{\boldsymbol{\theta}}$ denotes sampled subject characteristic parameters, and z_n denotes sampled signal quality parameters. Substituting (2.11) into (2.10), and taking the gradient of both sides yields an equation of the form:

$$\nabla_{\boldsymbol{\nu}} L(\boldsymbol{\nu}) = E_{\boldsymbol{\epsilon}\sim\mathcal{N}(0,I)}[\nabla_{\boldsymbol{\nu}} \tilde{L}(\boldsymbol{\epsilon}, \boldsymbol{\nu})] \tag{2.12}$$

where the operator $E_{\boldsymbol{\epsilon}\sim\mathcal{N}(0,I)}$ denotes expectation with respect to samples from the standard normal distribution, and the gradient operator $\nabla_{\boldsymbol{\nu}}$ has been moved inside the

expectation since the expectation operator no longer depends on $\boldsymbol{\nu}$. The term inside the expectation $\nabla_{\boldsymbol{\nu}}\tilde{L}(\epsilon, \boldsymbol{\nu})$ is a stochastic gradient of the objective, which can be written in expanded form as:

$$\nabla_{\boldsymbol{\nu}}\tilde{L}(\boldsymbol{\epsilon}, \boldsymbol{\nu}) = \nabla_{\boldsymbol{\nu}}[\log p\,(\boldsymbol{y}|\boldsymbol{z_\theta}, \boldsymbol{z_n}, \boldsymbol{u}) + \log p\left(\boldsymbol{z_\theta}\middle|\boldsymbol{z_\phi}\right) + \log p\,(\boldsymbol{z_n}) + \log p\left(\boldsymbol{z_\phi}\right) - \log q\left(\boldsymbol{z_\phi}, \boldsymbol{z_\theta}, \boldsymbol{z_n}\middle|\boldsymbol{\nu}\right)] \tag{2.13}$$

According to (2.12), the stochastic gradient $\nabla_{\boldsymbol{\nu}}\tilde{L}(\epsilon, \boldsymbol{\nu})$ is an unbiased noisy sample from the actual gradient $\nabla_{\boldsymbol{\nu}}L(\boldsymbol{\nu})$. Therefore, this stochastic gradient can be used along with a stochastic optimization algorithm to maximize the objective $L(\boldsymbol{\nu})$. Stochastic optimization with unbiased gradients has been shown to exhibit favorable convergence properties in theory and practice [98]–[100]. In addition, although convergence results may apply only locally to non-convex objectives, the randomized nature of stochastic optimization has been shown to facilitate escapes from local extrema, resulting in state-of-the art solutions in many practical applications [101].

The stochastic optimization procedure used in this work is shown in Algorithm 2.1. This procedure operates by sampling the stochastic gradient of the objective in each iteration, and producing corresponding increments to the variational parameters through adaptive moment estimation [99]. Within each iteration, the stochastic objective is computed according to Algorithm 2.2 based on (2.11)-(2.13). This procedure computes the objective by considering a sample from the approximate posterior, and evaluating the consistency of the sample both with respect to the available data and the structure of the generative model. In this way, the proposed iterative procedure searches for a generative model that is consistent with the available

physiological data as a whole, and additionally produces signal quality estimates, personalized models, and uncertainty quantification as byproducts.

Algorithm 2.1. *Stochastic Optimization for C-VI*

Inputs: $\boldsymbol{\nu}_0, \beta_1, \beta_2, \alpha, \delta$ {initial guess and constants}

$t \leftarrow 0$ {initialize increments}

$\nu \leftarrow \nu_0, m_\nu \leftarrow 0, v_\nu \leftarrow 0$ {initialize optimization variables}

For l from 1 to l_{max}:

$t \leftarrow t + 1$

$\boldsymbol{\epsilon} \sim \mathcal{N}(0, I)$ {draw sample from standard normal distribution}

$g_{\boldsymbol{\nu}} \leftarrow \nabla_{\boldsymbol{\nu}} \tilde{L}(\boldsymbol{\epsilon}, \boldsymbol{\nu})$ {obtain (stochastic) gradient of the objective}

$m_{\boldsymbol{\nu}} \leftarrow [\beta_1 m_{\boldsymbol{\nu}} + (1 - \beta_1) g_{\boldsymbol{\nu}}]/(1 - \beta_1^l)$ {first moment}

$v_{\boldsymbol{\nu}} \leftarrow [\beta_2 v_{\boldsymbol{\nu}} + (1 - \beta_2) g_{\boldsymbol{\nu}}^2]/(1 - \beta_2^l)$ {second moment}

$\boldsymbol{\nu} \leftarrow \boldsymbol{\nu} + \alpha[m_{\boldsymbol{\nu}}/(\sqrt{v_{\boldsymbol{\nu}}} + \delta)]$ {ADAM update; see [99]}

End For.

Return: $\boldsymbol{\nu}$ {optimized variational parameters}

Algorithm 2.2. *Evaluation of Stochastic Objective for C-VI*

Inputs: $\boldsymbol{\epsilon}, \boldsymbol{\nu}, \boldsymbol{u}, \boldsymbol{y}^d$ {parameters, stimuli, and data}

$\boldsymbol{z}_\phi, \boldsymbol{z}_\theta, \boldsymbol{z}_n \leftarrow f_q(\boldsymbol{\epsilon}, \boldsymbol{\nu})$ {draw sample from posterior; see (2.11)}

$p_q \leftarrow \log q(\boldsymbol{z}_\phi, \boldsymbol{z}_\theta, \boldsymbol{z}_n | \boldsymbol{\nu})$ {evaluate approximate posterior density for sample}

For All i, j:

$\mathrm{x_i} \leftarrow \mathcal{H}(\mathrm{z}_{\theta_i}, \mathrm{u_i})$ {run physiological model for sample}

$y_{ij}^{m} \leftarrow \mathcal{M}_j^{m}(z_{\theta_i}, x_i)$ {get model outputs for sample}

End For.

$p_y \leftarrow \log p\,(y|z_{\boldsymbol{\theta}}, z_n, u)$ {evaluate likelihood function for sample}

$p_{\boldsymbol{\theta}} \leftarrow \log p\,(z_{\boldsymbol{\theta}}|z_{\phi})$ {evaluate subject characteristics density for sample}

$p_{\phi} \leftarrow \log p\,(z_{\phi})$ {evaluate generator parameter prior for sample}

$p_n \leftarrow \log p\,(z_n)$ {evaluate measurement prior for sample}

$\tilde{L} \leftarrow p_y + p_{\boldsymbol{\theta}} + p_{\phi} + p_n - p_q$

Return: $\tilde{L}$ {(stochastic) variational objective value}

2.3. The Hemorrhage Resuscitation Model

In this work, we use a hemorrhage resuscitation modeling case study to illustrate the efficacy of the C-VI method in the context of personalized and generative physiological modeling. In this modeling problem, the aim is to obtain interpretable mathematical models that can reproduce and predict the hemodynamic effects of hemorrhage and fluid resuscitation, both in specific subjects and in a given population. In addition to illustrating the efficacy of C-VI, the knowledge encoded in these models may be leveraged and extended to complement and further advance the state of the art in the development and testing of physiological monitoring [102], [103], decision support, and closed-loop control algorithms [52], [65] for hemorrhage resuscitation. In this

section, we build upon previous work [94], [104], [105] to derive a mathematical model of the main physiological phenomena in hemorrhage resuscitation, and specify the stimuli, states, and model parameters for this problem. A schematic representation of this model is shown in Figure 2.2. The model consists of four main components that are described in the following subsections.

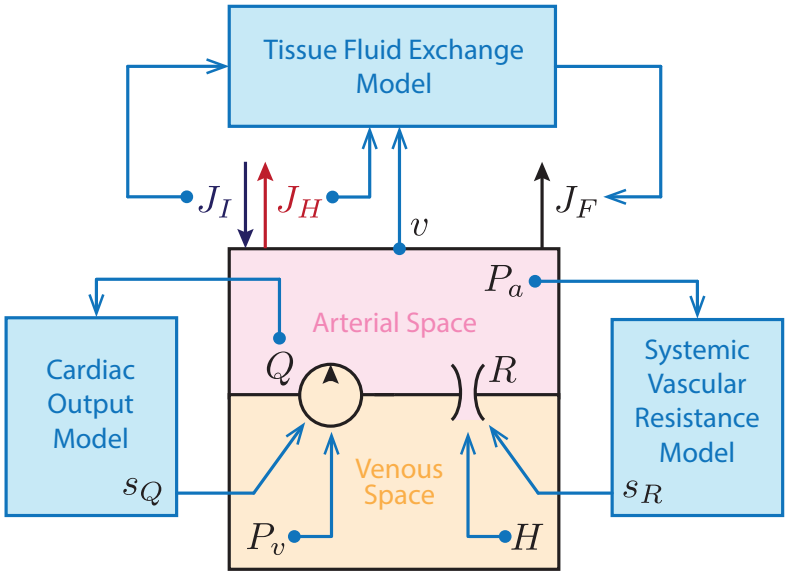

Figure 2.2. A schematic representation of the hemorrhage resuscitation model. The model consists of four main components: (i) the blood circulation model, (ii) the tissue fluid exchange model, (iii) the systemic vascular resistance model, and (iv) the cardiac output model. © 2022 IEEE.

2.3.1. Blood Circulation Model

In order to model the effects of hemorrhage and fluid resuscitation on the physiological state of a subject, a blood circulation model must be formulated that accounts for the volume and composition of blood in relevant circulatory spaces. A macro-state realization of such a model is described by the following differential equations:

$$\dot{v_a} = Q - (P_a - P_v)/R - J_H - J_F \tag{2.14}$$

$$\dot{v_v} = -Q + (P_a - P_v)/R + J_I \tag{2.15}$$

$$\dot{v_r} = -J_H H \tag{2.16}$$

where the states v_a and v_v respectively denote the arterial and venous blood volumes, and v_r denotes the total red blood cell volume. The changes in these volumes are driven by several flow-rate terms. The term Q denotes CO, which is the flowrate of blood pumped by the heart. The term $(P_a - P_v)/R$ is the flowrate of blood moving through the vascular system, where P_a is the MAP, P_v is the central venous pressure (CVP), and R is the systemic vascular resistance (SVR). The terms J_H and J_I respectively denote the flowrates of hemorrhage and fluid resuscitation, while J_F is the net rate of fluid exchange with the tissue space. Finally, the term $J_H H$ is the flowrate for red blood cell loss due to hemorrhage, where $H = v_r/(v_a + v_v)$ denotes the blood HCT. Prior to any perturbation, the system described by (2.14)-(2.16) is set up to be in equilibrium, with baseline arterial, venous, and red blood cell volumes at v_{a0}, v_{v0}, and v_{r0}, respectively, and $Q_0 = (P_{a0} - P_{v0})/R_0$, where Q_0 is the baseline CO, P_{a0} and P_{v0} denote baseline MAP and CVP, and R_0 is the baseline SVR. Changes in MAP and CVP are modeled to depend on changes in arterial and venous blood volume as follows:

$$P_a = P_{a0} + K_a(v_a - v_{a0}) \tag{2.17}$$

$$P_v = P_{v0} + K_v(v_v - v_{v0}) \tag{2.18}$$

where K_a is the arterial elastance, and K_v is the venous elastance. The model described by equations (2.14)-(2.18) is determined except for the responses of J_F, R, and Q, which are addressed in the the next three subsections.

2.3.2. Tissue Fluid Exchange Model

Circulating blood interacts with the fluid in the surrounding tissues (i.e., interstitial fluid) through lymphatic and micro-vascular exchange systems. By virtue of this interaction, the body can regulate the blood volume in the event of external perturbations such as hemorrhage and fluid resuscitation. This is achieved through shifting excess fluid from the blood to the tissue space, or compensating for a dearth of blood by drawing fluid from the tissue space into the blood [104], [106]. The net rate of fluid exchange with the tissue space J_F is thus an important quantity to model. For this purpose, we use a blood volume controller formulation of the form:

$$J_F = K_p(v - v_0 - r_F) \tag{2.19}$$

where $v = v_a + v_v$ is the total blood volume, $v_0 = v_{a0} + v_{v0}$ is the baseline total blood volume, K_p is the proportional gain of the controller, and r_F is the reference signal for the controller. The reference signal r_F determines the final value of blood volume change long after a blood volume perturbation (i.e., hemorrhage and/or fluid resuscitation), and is defined based on the history of blood volume perturbations as follows:

$$\dot{r_F} = \frac{1}{1 + \alpha_I} J_I - \frac{1}{1 + \alpha_H} J_H \tag{2.20}$$

where the parameter α_I determines the fraction of the resuscitation fluid that will remain in the blood after the exchange of fluid with the tissue space, and the parameter α_H determines the fraction of hemorrhaged blood that will remain uncompensated after

the exchange of fluid with the tissue space (see [104], [107] for more details on this modeling approach).

2.3.3. Systemic Vascular Resistance Model

The SVR (R) is the resistance of the vascular system to blood flow. In hemorrhage resuscitation scenarios, the SVR is affected by two main mechanisms: (i) the constriction and dilation of the blood vessels, and (ii) the viscosity of the blood moving through the blood vessels. The control mechanisms in the body (e.g., the baro-reflex mechanism) modulate SVR through vasoconstriction and vasodilation in order to maintain adequate MAP and tissue perfusion [108]. In addition, changes in blood HCT cause changes in blood viscosity, which in turn disturb the SVR [109]. To model these mechanisms, we formulate the SVR in the following form:

$$R = R_0 + K_h(H - H_0) + s_R \tag{2.21}$$

where $H_0 = v_{r0}/(v_{a0} + v_{v0})$ is the baseline blood HCT, K_h is a parameter representing the sensitivity of the SVR to changes in HCT, and s_R is a state representing the amount of SVR change prompted by the control mechanisms in the body. The state equation for s_R is therefore formulated as follows:

$$\dot{s_R} = -\frac{1}{\tau_R} s_R - \frac{K_R}{\tau_R}(P_a - P_{a0}) \tag{2.22}$$

where K_R is the controller gain, and τ_R is the time constant of the control system. Equations (2.21) and (2.22) together describe a control system whose objective is to maintain adequate MAP by changing SVR, while the SVR is disturbed by viscosity changes resulting from variations in blood HCT.

2.3.4. Cardiac Output Model

The CO (Q) is the flowrate of blood pumped into circulation by the heart. In hemorrhage resuscitation scenarios, the changes in CO stem from two main mechanisms. First, perturbations in CVP directly affect the right atrial pressure and subsequently the left ventricular preload. According to the Frank-Starling law, a higher preload results in higher cardiac muscle tension which in turn results in a more forceful stroke and higher CO. Second, the control mechanisms in the body modulate the heart rate and cardiac contractility in order to maintain adequate CO [108]. To model these mechanisms, we formulate the CO equation in the following form:

$$Q = Q_0 + \beta_v(P_v - P_{v0}) + s_Q \tag{2.23}$$

where β_v is a parameter representing the sensitivity of the CO to changes in CVP, and s_Q is a state representing the amount of CO change prompted by the control mechanisms in the body. The state equation for s_Q is formulated as follows:

$$\dot{s_Q} = -K_Q(Q - Q_0) \tag{2.24}$$

where K_Q is the controller gain. Equations (2.23) and (2.24) together describe a control system whose objective is to maintain adequate CO through changing s_Q, while the CO is disturbed by changes in CVP.

2.3.5. Physiological Model Summary

The presented hemorrhage resuscitation model corresponds to the physiological process model $\mathcal{H}$ defined in (2.2). For subject i, the physiological stimuli can be

summarized as $u_i(t) = \{J_I(t), J_H(t)\}_i$, and the physiological characteristics of the subject can be summarized as:

$$\theta_i = \left[v_0 \;\; H_0 \;\; Q_0 \;\; P_{a0} \;\; K_v \;\; K_a/K_v \;\; K_p \;\; \alpha_I \;\; \alpha_H \;\; K_h \;\; \tau_R \;\; K_R \;\; \beta_v \;\; K_Q \right]_i \qquad (2.25)$$

Given θ_i, the rest of the physiological parameters in the model are determined as follows: the baseline arterial and venous blood volumes are nominally set to $v_{a0} = 0.3v_0$, and $v_{v0} = 0.7v_0$; the baseline CVP is set to a nominal value P_{v0} from measured data; and the initial SVR is calculated from $R_0 = (P_{a0} - P_{v0})/Q_0$. Given these parameters, the state evolution of the physiological system $x_i(t)$ is obtained by numerically solving the differential equations described by (2.14)-(2.24).

2.4. Applying C-VI to Hemorrhage Resuscitation Modeling

In this section, we present the details of applying the C-VI method presented in Section 2.2 to the hemorrhage resuscitation modeling problem presented in Section 2.3. In addition, we present the details pertaining to the available physiological data and the methods used for data analysis to illustrate the efficacy of C-VI in the context of personalized and generative physiological modeling, especially in the presence of low-information and heterogeneous data. Further details follow.

2.4.1. The Approximate Posterior

As presented in Sections 2.2.3 and 2.2.5, variational inference is performed by leveraging a family of densities that act as candidates for the best approximate posterior. In this work, we employ a family of diagonal Gaussian densities for this purpose, which can be written in function form as:

$$\left[\boldsymbol{z}_{\phi}; \boldsymbol{z}_{\boldsymbol{\theta}}; \boldsymbol{z}_{\boldsymbol{n}}\right] = \boldsymbol{\nu}_{\boldsymbol{\mu}} + \mathrm{diag}(\boldsymbol{\nu}_{\boldsymbol{\sigma}})\boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, I) \tag{2.26}$$

where $\boldsymbol{\nu}_{\mu} = \left[\boldsymbol{\nu}_{\mu:\phi}; \boldsymbol{\nu}_{\mu:\boldsymbol{\theta}}; \boldsymbol{\nu}_{\mu:n}\right]$ is the mean vector of the approximate posterior, which represents most-likely values for the model parameters (ϕ, $\boldsymbol{\theta}$, and n), and $\boldsymbol{\nu}_{\sigma} = \left[\boldsymbol{\nu}_{\sigma:\phi}; \boldsymbol{\nu}_{\sigma:\boldsymbol{\theta}}; \boldsymbol{\nu}_{\sigma:n}\right]$ is the standard deviation vector of the approximate posterior, which represents the uncertainty associated with the model parameters. Thus, the variational parameters for this choice of approximate posterior can be summarized as $\boldsymbol{\nu} = \{\boldsymbol{\nu}_{\mu}, \boldsymbol{\nu}_{\sigma}\}$. Given this formulation, the approximate posterior density associated with a sample generated by (2.26) can be computed from:

$$\log q\left(\boldsymbol{z}_{\phi}, \boldsymbol{z}_{\boldsymbol{\theta}}, \boldsymbol{z}_{\boldsymbol{n}} \middle| \boldsymbol{\nu}\right) = \sum_k \left[-\frac{1}{2}\boldsymbol{\epsilon}_{\boldsymbol{k}}^{\boldsymbol{2}} - \log(\boldsymbol{\nu}_{\boldsymbol{\sigma}})_k - \frac{1}{2}\log(2\pi)\right] \tag{2.27}$$

where the sum Σ_k is computed over the elements of the vectors ϵ and $\boldsymbol{\nu}_{\boldsymbol{\sigma}}$. The density in (2.27) is used in Algorithm 2.2 as part of the stochastic objective.

2.4.2. The Subject Generator Model

As presented in Section 2.2.1, the proposed modeling scheme includes a subject generator model $\mathcal{G}(\phi)$. In this work, we employ a full-covariance Gaussian generator for this purpose, which can be written in function form as:

$$\theta_i = \phi_{\mu} + \phi_L \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, I) \tag{2.28}$$

where ϕ_{μ} is the mean vector of the subject generator, which represents the most typical subject, and ϕ_L is a lower-triangular matrix denoting the Cholesky decomposition of the covariance matrix of the subject generator. The covariance matrix itself can be computed from $\phi_{\Sigma} = \phi_L \phi_L^T$. Thus, the subject generator model parameters can be

summarized as $\phi = \{\phi_\mu, \phi_L\}$. Given this formulation, the subject generator density associated with a sample from the approximate posterior (as in (2.26)) can be computed from the following equation:

$$\log p\left(\boldsymbol{z_\theta} \middle| \boldsymbol{z}_\phi\right) = \sum_i \left[-\frac{1}{2}\left(\boldsymbol{z}_{\theta_i} - \boldsymbol{z}_{\phi_\mu}\right)^T \boldsymbol{z}_{\phi_\Sigma}^{-\mathbf{1}} \left(\boldsymbol{z}_{\theta_i} - \boldsymbol{z}_{\phi_\mu}\right) - \frac{1}{2}\log\left(\left|\boldsymbol{z}_{\phi_\Sigma}\right|\right) - \frac{d_\theta}{2}\log(2\pi) \right] \tag{2.29}$$

where z_{θ_i} is the sample associated with the physiological model parameters for subject i, z_{ϕ_μ} is the sample associated with the mean vector of the subject generator, z_{ϕ_Σ} is the sample associated with the covariance matrix of the subject generator, and d_θ is the dimension of physiological model parameters (and also the dimension of z_{θ_i} and z_{ϕ_μ}). Equation (2.29) is used in Algorithm 2.2 as part of the stochastic objective.

The full-covariance subject generator in (2.28) is a relatively expressive model that may capture inter-subject variabilities in the form of a covariance matrix ϕ_Σ. However, effective characterization of this covariance matrix from data is contingent on the availability of a sufficient number of subjects in the dataset. In many physiological modeling applications, only a limited number of subjects are available for experimentation, which may in turn result in an "over-fitted" covariance matrix. To create a balance between generator model complexity and subject availability, we utilize regularization on the generator model parameters according to the following function:

$$\log p(\boldsymbol{z}_\phi) = -\lambda \left\| \boldsymbol{z}_{\phi_\Sigma} \right\|_* \tag{2.30}$$

where $||.||_*$ denotes the nuclear norm of the matrix, and λ is a scalar hyper-parameter. The nuclear norm $\left\|\boldsymbol{z}_{\phi_\Sigma}\right\|_*$ promotes a compressed subject generator in the sense of the sum of singular values for its covariance matrix, while the hyper-parameter λ can be used to modulate the rate of compression. Equation (2.30) is used in Algorithm 2.2 as part of the stochastic objective. The hyper-parameter λ is selected using a method conceptually similar to the L-curve approach [110]: λ is increased from zero while the likelihood in (2.34) (which represents the goodness of fit to the available data) is evaluated. λ is chosen as the value at which the likelihood starts to exhibit large deterioration. Such a choice of λ can achieve an adequate balance between generator complexity and subject availability [94], [110].

2.4.3. The Physiological Measurement Model

As presented in Section 2.2.1, the proposed modeling approach includes physiological measurement models $\mathcal{M}_j$ intended to represent the measurement processes used to obtain a real dataset. For the hemorrhage resuscitation problem, we consider three potential types of measurement: HCT, CO, and MAP. We model each of these measurements as the output of a process that observes the model outputs H, Q, and P_a, while the observations are corrupted by additive white Gaussian noise. This can be formulated as:

$$y_{i,H} = \{H(t) + n_{i,H}.\epsilon \mid t \in T_{i,H},\ \epsilon \sim \mathcal{N}(0,1)\} \tag{2.31}$$

$$y_{i,Q} = \{Q(t) + n_{i,Q}.\epsilon \mid t \in T_{i,Q},\ \epsilon \sim \mathcal{N}(0,1)\} \tag{2.32}$$

$$y_{i,P_a} = \{P_a(t) + n_{i,P_a}.\epsilon \mid t \in T_{i,P_a},\ \epsilon \sim \mathcal{N}(0,1)\} \tag{2.33}$$

where $y_{i,H}$, $y_{i,Q}$, and y_{i,P_a} respectively denote the model-generated virtual data for HCT, CO, and MAP in subject i. The sets $T_{i,H}$, $T_{i,Q}$, T_{i,P_a} contain the time points at which observations are made for subject i, and the latent parameters $n_{i,H}$, $n_{i,Q}$, n_{i,P_a} denote the standard deviations of the Gaussian noises corrupting the observations of HCT, CO, and MAP in subject i. For inference purposes, the likelihood associated with the physiological data can be obtained from:

$$\log p\,(y|z_{\boldsymbol{\theta}}, z_n, u) = \sum_i \sum_j \left[-\frac{1}{2z_{n_{ij}}^2}\left(y_{ij}^m - y_{ij}^d\right)^T\left(y_{ij}^m - y_{ij}^d\right) - d_{ij}\log\left(z_{n_{ij}}\right) - \frac{d_{ij}}{2}\log(2\pi) \right] \quad (2.34)$$

where $j \in \{H, Q, P_a\}$ is and index for the measured variables in the hemorrhage resuscitation problem, y_{ij}^m represents the (uncorrupted) physiological model outputs, y_{ij}^d denotes real data, $z_{n_{ij}}$ is an approximate posterior sample associated with the signal quality parameter n_{ij}, and d_{ij} denotes the length of the data vector y_{ij}^d. The likelihood (2.34) is substituted into its place in Algorithm 2.2 to complete the formulation of the stochastic objective. Furthermore, in this work, we do not assume further prior knowledge of the noise, which can be reflected in Algorithm 2.2 by setting $\log p\,(z_n)$ to zero.

2.4.4. The Physiological Data

The physiological data used in this work are derived from a series of hemorrhage resuscitation experiments conducted on sheep subjects ($N = 23$) in an array of previous work [48], [111], [112]. In each experiment, the animal is subjected to a large initial

hemorrhage and two subsequent smaller hemorrhages. To counter the physiological effects of hemorrhage over time, the subject is resuscitated using Ringer's Lactate infusions. The infusions are performed according to pre-determined closed-loop control laws designed to restore and regulate MAP. See Figure 2.3 for hemorrhage and resuscitation profiles received by two example subjects. During each experiment, the HCT, CO, and MAP responses of the subject are measured and recorded as data. These measurements are performed at $\sim$ 5-minute intervals over the course of 180 minutes. These animal experiments are useful for physiological modeling purposes, as (i) they can provide informative physiological measurements that are not commonly available in clinical settings, and (ii) the exact timing and amount of hemorrhages and fluid infusions applied to each subject are known [94], [104], [107].

2.4.5. Data Analysis

In the case study on hemorrhage resuscitation modeling, our aim is to illustrate the efficacy of the C-VI method in inferring personalized and generative physiological models from low-information and heterogeneous data. For this purpose, we consider four inference scenarios: (i) C-VI given full data, (ii) C-VI given partial data, (iii) non-collective MLE given full data, and (iv) non-collective MLE given partial data. The full data scenarios were created by presenting all available data to the methods, which allowed us to obtain highly informed physiological models. The partial data scenarios were created by random exclusion of measured variables from data, which allowed us to evaluate the methods against low-information and heterogeneous data. Further details follow.

C-VI versus Non-Collective MLE: To perform C-VI, we used the (partially or fully) available HCT, CO, and MAP data as inputs to the stochastic objective shown in Algorithm 2.2 and maximized this objective using Algorithm 2.1 to obtain the optimized variational parameters $\boldsymbol{\nu}$. According to the approximate posterior formulation in (2.26), the most-likely personalized physiological model parameters can be obtained from $\boldsymbol{\nu}_{\mu:\boldsymbol{\theta}}$, while their associated uncertainties can be obtained from $\boldsymbol{\nu}_{\sigma:\boldsymbol{\theta}}$. In addition, the most-likely subject generator model parameters can be obtained from $\boldsymbol{\nu}_{\mu:\phi}$, and additional virtual subjects can be generated using the generator in (2.28) with ϕ set to $\boldsymbol{\nu}_{\mu:\phi}$. To perform non-collective MLE, we modified the stochastic objective shown in Algorithm 2.2 by setting the non-likelihood terms (i.e., p_q, $p_{\boldsymbol{\theta}}$, p_{ϕ}, and p_n) to zero. We maximized this objective using Algorithm 2.1, and obtained the estimates for the personalized physiological model parameters from $\boldsymbol{\nu}_{\mu:\boldsymbol{\theta}}$. To generate virtual subjects in this case, we calculated the mean and standard deviation of the estimated personalized values for each physiological parameter, and generated additional virtual subjects by sampling from a Gaussian distribution with the same mean and standard deviation.

Full Data, Partial Data, and Method Evaluation: To study and compare the effects of information scarcity and data heterogeneity on the fidelity of the personalized models and the generated virtual subjects, we performed inference in two data availability cases. In the full data case, we used the available HCT, CO, and MAP data across all subjects to perform C-VI and non-collective MLE. In the partial data case, we randomly excluded two out of the three measured variables (HCT, CO, and MAP)

for each subject. The resulting partial dataset thus consisted of 8 subjects with only HCT measurements, 8 subjects with only CO measurements, and 7 subjects with only MAP measurements. We used this partial dataset to perform inference in both C-VI and non-collective MLE cases. To evaluate the fidelity of the personalized physiological models, we computed the mean-absolute error of the model predictions with respect to the excluded data for each subject. To determine significance in difference between C-VI and non-collective MLE cases, we utilized the Wilcoxon signed-rank test. To evaluate the fidelity of the subject generator models, we computed the likelihood of the excluded data in the generated virtual subjects, which can be written as:

$$p(y_E|n,u) = \mathbb{E}_{\theta_i \sim \mathcal{G}}[p(y_E|\theta_i,n,u)] \tag{2.35}$$

where y_E denotes the space of excluded data, $p(y_E|\theta_i,n,u)$ denotes the likelihood of the excluded data given a generated subject θ_i, and $p(y_E|n,u)$ denotes the likelihood of the excluded data under the subject generator $\mathcal{G}$.

2.5. Results and Discussion

In this section, we present the results of applying the C-VI method presented in Section 2.2 to the hemorrhage resuscitation modeling problem presented in Section 2.3, and discuss the effectiveness of C-VI in the context of personalized and generative physiological modeling. Further details follow.

2.5.1. Personalized Physiological Modeling

Figure 2.3 shows the personalized hemorrhage resuscitation model responses in two representative subjects. The model responses in the rest of the subjects are provided in the supplementary material. These models were obtained by presenting the full dataset to the C-VI method. For each subject in Figure 2.3, the most-likely physiological models (shown as solid blue lines) reproduced the trends in measured data for HCT, CO, and MAP. Among the N=23 subjects in the dataset, these physiological models reproduced the trends in measured data with an average mean-absolute error of 0.51% for HCT, 0.36 L/min for CO, and 7.77 mmHg for MAP. Furthermore, sampling the approximate posterior yielded other likely physiological models that also reproduced the trends in measured data. This is shown in Figure 2.3 using shaded blue areas representing the 2σ interval for the likely physiological model responses. In addition, for each subject, the measurement models captured the unexplained variations in measured data for HCT, CO, and MAP. This is shown in Figure 2.3 using dashed blue lines representing the 2σ interval for the measurement model responses. Overall, these results suggest that (i) the proposed hemorrhage resuscitation model can adequately reproduce the trends in the physiological data despite its relatively simple structure, and (ii) the C-VI method can infer personalized physiological models that capture the trends in measured data as well as personalized measurement models that capture the extent of unexplained variations in measured data.

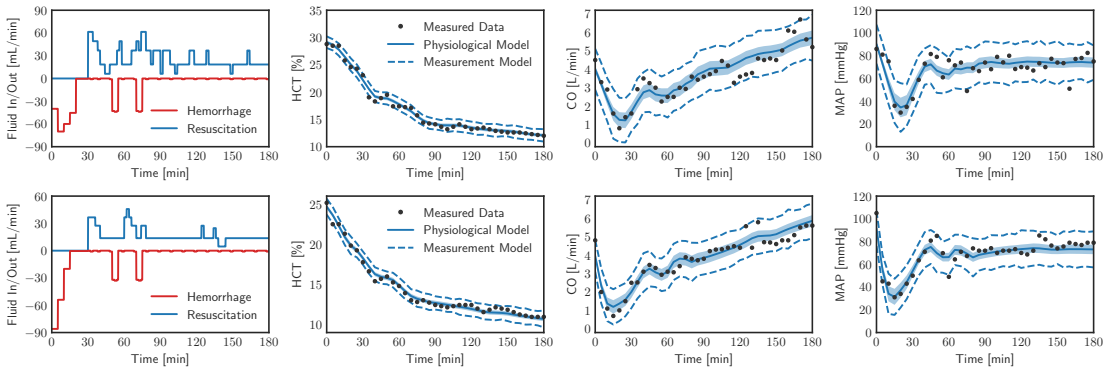

Figure 2.3. Personalized physiological model and measurement model responses to hemorrhage and fluid resuscitation for two example subjects. Each row corresponds to a subject. The first column shows the stimuli received by each subject, and the second to fourth columns respectively show hematocrit (HCT), cardiac output (CO), and mean arterial pressure (MAP) responses for each subject against measured physiological data. Shaded blue areas show the 2σ intervals associated with personalized physiological model responses, and dashed blue lines show the 2σ intervals associated with the measurements generated by the measurement model. © 2022 IEEE.

Figure 2.4. shows the 2σ intervals for the personalized physiological model parameters, which are plotted against the 2σ intervals for the parameters generated by the subject generator. The full set of parameter intervals is provided in the supplementary material. These intervals were obtained by presenting the full dataset to the C-VI method. Naturally, the size of the personalized intervals varied across different physiological model parameters. For some parameters (e.g., the pair shown on the right-hand side of Figure 2.4), the personalized intervals were smaller, and traveled further into the parameter space. For others (e.g., the pair shown on the left-hand side of Figure 2.4), the personalized intervals were larger, and resided closer to each other in the parameter space. This phenomenon is a consequence of the C-VI formulation. In case the data provides strong information about a parameter, the likelihood term $\log p\left(\boldsymbol{y}|\boldsymbol{z_\theta}, \boldsymbol{z_n}, \boldsymbol{u}\right)$ in (2.13) dominates the objective, causing the personalized intervals to shrink, move away toward their appropriate values, and alter

the shape of the subject generator so that it can reproduce these parameter variations. Conversely, if the data provides weak information about a parameter, the $\log p\left(\boldsymbol{z_\theta}|\boldsymbol{z_\phi}\right)$, $\log p\left(\boldsymbol{z_\phi}\right)$, and $\log q\left(\boldsymbol{z_\phi}, \boldsymbol{z_\theta}, \boldsymbol{z_n}|\boldsymbol{\nu}\right)$ terms in (2.13) dominate the objective, causing the personalized intervals to expand and the subject generator to shrink, gathering the personalized intervals together. Arguably, this provides opportunity for the personalized intervals to share their weak information in order to reach a consensus about appropriate values for the parameter. Overall, these results suggest that the C-VI method can utilize data across multiple experiments to (i) find personalized physiological model parameters for the subjects in a dataset and provide a measure of confidence about their values, and (ii) provide the opportunity to collectively determine plausible values for weakly-informed parameters by aggregating information across different experiments.

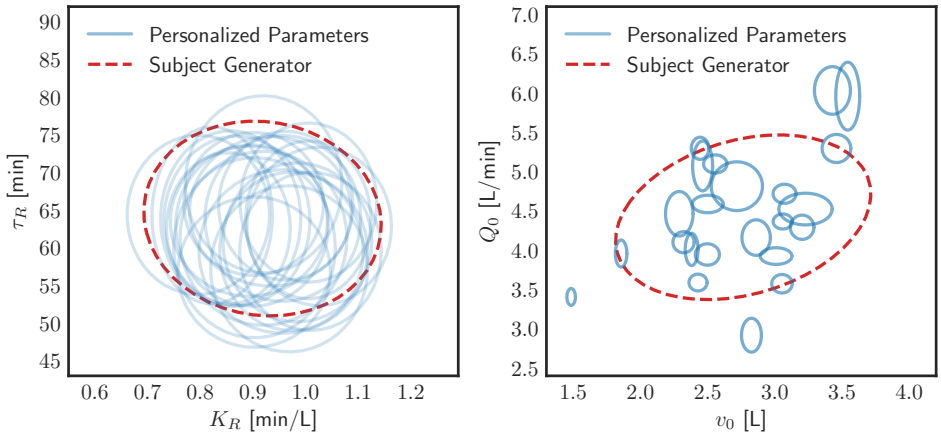

Figure 2.4. Personalized physiological model parameters versus the subject generator model for two example parameter pairs associated with weak information (left) and strong information (right) availability. Blue ellipses show 2σ intervals for personalized physiological model parameters, and the red ellipse shows the 2σ interval associated with the subjects generated by the subject generator model.

Table 2.1 shows the mean-absolute prediction errors for the personalized physiological models inferred from partial data. The partial dataset was constructed by randomly excluding two out of three measured variables (HCT, CO, and MAP) from each subject, simulating a case of low-information and heterogeneous data. The prediction errors (associated with excluded data) were computed against the excluded data for both C-VI and non-collective MLE methods (see Section 2.4.5 for details). As presented in Table 2.1, using C-VI on partial data resulted in lower HCT prediction error, and significantly lower CO and MAP prediction errors when compared to using the non-collective MLE method. This advantage can be attributed to the collective formulation of C-VI, in which the variables form an interconnected tree-like structure (as in Figure 2.1(b)). As a result, a loss of observation on some variables (in this case, a random subset of y_{ij}'s) is partially counteracted by the information coming from other observed variables (in this case, information from observed y_{ij}'s travels up to θ_i's, ϕ, and back down to unobserved y_{ij}'s), giving the personalized models superior prediction performance in the face of low-information and heterogeneous data. Overall, these results suggest that the C-VI method can reconcile data from a collection of experiments to produce superior (more predictive) personalized models when compared to a non-collective MLE method, especially when only low-information and heterogeneous data are available.

Table 2.1. Mean-absolute prediction errors for personalized models inferred from partial data [median (inter-quartile range); *denotes significance with respect to non-collective MLE ($p < 0.05$)].

	HCT [%]	CO [L/min]	MAP [mmHg]
Non-Collective MLE	5.10 (6.00)	1.88 (1.65)	35.2 (31.0)
C-VI	3.27 (3.93)	1.04 (1.04)*	9.98 (3.65)*

2.5.2. Generative Physiological Modeling

Figure 2.5 shows the HCT, CO, and MAP responses of the virtual subjects generated by two subject generators, obtained from applying C-VI to full and partial datasets. In the first row, we show the physiological responses of one hundred virtual subjects to an example hemorrhage resuscitation profile. In the second row, we show an instance of virtual data produced from a virtual subject in response to the same hemorrhage resuscitation profile. From visual inspection, the virtual subjects exhibited a reasonable range of behavior for HCT, CO and MAP, and did not show any objectively unrealistic behavior (e.g., responses that are out of a physiologically meaningful range). In addition, the generated virtual data appeared reasonable and visually similar to data acquired from a hemorrhage resuscitation experiment. In the following paragraph, we aim to analyze and discuss these observations more concretely.

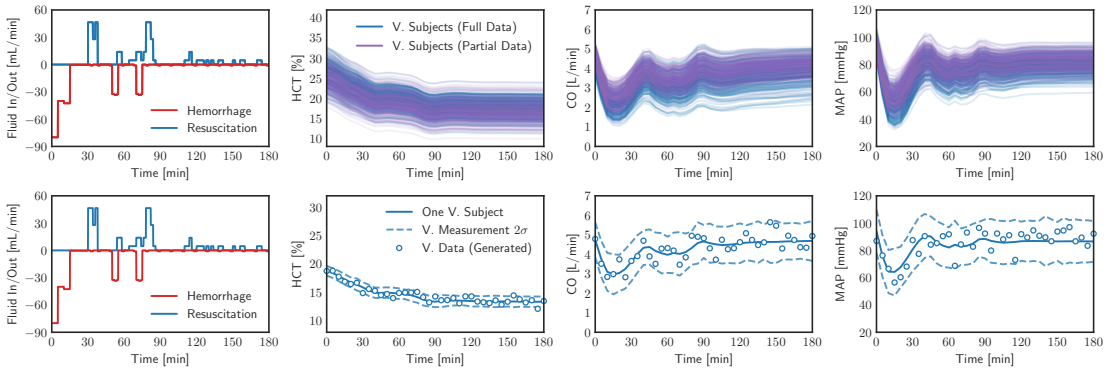

Figure 2.5. Generative physiological model responses to hemorrhage and fluid resuscitation: First row shows the physiological model responses of 100 virtual subjects generated by each subject generator model. Second row shows physiological and measurement responses of one example virtual subject. The first column shows the stimuli received by the subjects, and the second to fourth columns respectively show hematocrit (HCT), cardiac output (CO), and mean arterial pressure (MAP) responses.

Figure 2.6. shows histogram plots of post-hemorrhage (at 20 min in Figure 2.5 HCT, CO, and MAP responses in the generated virtual subjects. The first row in Figure 2.5 corresponds to virtual subject generation given full data, while the second row corresponds to virtual subject generation given partial data. Blue and purple histograms correspond to generation using the results of C-VI, while red histograms correspond to generation using the results of non-collective MLE (see Section 2.4.5 for details). With full data availability, the non-collective MLE approach resulted in scattered (and occasionally unrealistic) virtual subject responses especially for CO and MAP. With partial data availability, the non-collective MLE approach produced even more scattered responses for HCT, CO, and MAP, including many objectively unrealistic responses (e.g., MAP's above 100 mmHg, and CO's above 5 L/min post-hemorrhage). In contrast, the C-VI approach produced a more realistic range of post-hemorrhage behavior in the case of full data availability, and also retained its generation performance in the case of partial data availability. Furthermore, in the case of partial data availability, the negative log-likelihood of the excluded data (Table 2.2) was lower for the C-VI approach, indicating superior quality for the virtual subjects generated by this method. This can be attributed to the formulation of the C-VI method, which is built to simultaneously infer both the personalized parameters (θ_i's and n_{ij}'s) and the generator parameters (ϕ) in a consistent manner based on all available measured data. In contrast, in the non-collective MLE approach, we generated virtual subjects by mimicking parametric variations across separately inferred personalized parameters, which are sensitive to a lack of information in the data. Overall, these results suggest that the C-VI method can reconcile low-information and heterogeneous data from

multiple experiments to obtain a robust generative model of the studied physiological process. This generative model can in turn be used to create new virtual subjects not already in the population in the form of data, but distributed according to the population of real subjects in the dataset.

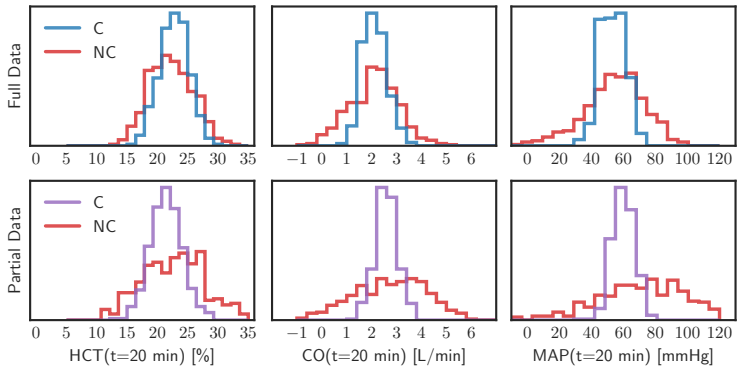

Figure 2.6. Histogram plots for post-hemorrhage (at 20 min in Figure 2.5) hematocrit (HCT), cardiac output (CO), and mean arterial pressure (MAP) in generated virtual subjects. First row corresponds to generation results given full data, while second row corresponds to generation results given partial data. Blue and purple histograms correspond to the collective (C) approach to generation while red histograms correspond to a non-collective (NC) approach to generation. © 2022 IEEE.

Table 2.2. Negative log-likelihood of the excluded data under subject generators inferred from partial data (lower is better). © 2022 IEEE.

	Negative Log-Likelihood $[-\log p(\boldsymbol{y}_E \mid \boldsymbol{n}, \boldsymbol{u})]$
Non-Collective MLE	2.09×10^{10}
C-VI	6.06×10^{2}

2.5.3. Potential Applications

As discussed, the C-VI method can be used to reconcile low-information and heterogeneous data from a collection of experiments to obtain robust *personalized* and *generative* physiological models. In the following paragraphs, we discuss several potential applications of these results.

In Silico Clinical Trials: In silico clinical trial is an emerging field of research concerned with the use of qualified mathematical/computational models and simulations to perform clinical trials [113]–[116]. As discussed in our case study, C-VI produces a generative physiological model that can generate cohorts of virtual subjects, compute their dynamic responses to given stimuli, and produce virtual data by mimicking specific measurement processes. In addition, C-VI produces personalized physiological models that can serve as "digital twins" of specific subjects in a dataset. Together, these physiological models provide a rich set of representations that may be used in conjunction with recent in silico clinical trial methodologies [113]–[115] to build credible systems and methodologies for systematic testing of physiological monitoring, decision-support, and closed-loop control algorithms.

Monitoring and Control with Limited Measurements: As presented in Section 2.1, a lack of sufficient information in experimental data is a long-standing challenge in physiological modeling. This challenge is expected to be even more pronounced in clinical settings, where invasive and/or expensive measurements (e.g., HCT and CO) are available only occasionally. This presents a potential use case for the C-VI method in clinical settings. Arguably, in a procedure similar to the partial data results presented in this work, the C-VI method can be used to reconcile limited clinical measurements with past information (e.g., measurements from past patients and/or from past lab experiments) to derive high-quality models of new patients. These results may in turn be leveraged to build model-based algorithms that can monitor and control the physiological states of the patient using limited clinical measurements.

2.5.4. Limitations

The presented study has limitations. First, the inference problem addressed is in general a non-convex problem. While stochastic optimization, as employed, is known to provide robust solutions for such problems in practice [101], theoretical guarantees are limited to local optimality [98]–[100]. Second, C-VI is built to produce a subject generator that follows the distribution of the real subjects in the dataset. Therefore, in case the dataset consists of homogeneous subjects and/or subjects biased toward a specific application, this tendency will likely be reflected in the generated virtual subjects. Thus, the resulting virtual subjects are most suited to contexts that are close to that of the original dataset. In our case study, the data were collected in the context of closed-loop control for hemorrhage resuscitation. Therefore, the resulting physiological models may be primarily suited to the design and in silico testing of hemorrhage resuscitation algorithms. In this regard, we expect that additional data gathering and physiological modeling efforts may be needed in order to extend the utility of the presented physiological models to applications beyond hemorrhage resuscitation.

2.6. Conclusions

In this chapter, we proposed the C-VI method to facilitate the personalized and generative modeling of physiological systems given low-information and heterogeneous data. To illustrate the effectiveness of the C-VI method, we applied it to a practically important case study on modeling the hemodynamic effects of hemorrhage and fluid resuscitation. In the context of this case study, we demonstrated that the C-

VI method can reconcile heterogeneous combinations of HCT, CO, and MAP data across multiple experiments to produce robust personalized and generative physiological models. In addition, we demonstrated that the C-VI method produces superior (more predictive) physiological models when compared to a non-collective MLE method, especially when only low-information and heterogeneous data are available. Over the subsequent chapters in this dissertation, effort is devoted to the study of approaches that incorporate the collective inference perspective into the design and testing of interpretable physiological monitoring, decision support, and closed-loop control algorithms [67][2].

[2] This chapter is adapted from a journal publication by the author of this dissertation. © 2022 IEEE. Reprinted, with permission, from: Tivay, Ali, George C. Kramer, and Jin-Oh Hahn. "Collective Variational Inference for Personalized and Generative Physiological Modeling: A Case Study on Hemorrhage Resuscitation." *IEEE Transactions on Biomedical Engineering* (2022).

Chapter 3: A Population-Informed Particle Filtering Method for Robust Physiological Monitoring Using Low-Information Time-Series Measurements

Forming reliable beliefs about a patient's physiological state is essential for algorithmic decision-making in medical care settings. However, conventional filtering algorithms may face challenges in forming beliefs when measurements are obtained intermittently or contain limited information. This chapter presents the population-informed particle filter (PIPF), a filtering approach that incorporates past experiences with patients into the filtering process to provide more reliable beliefs about a new patient's physiological state. To derive the PIPF, we formulate the filtering problem as recursive inference on a probabilistic graphical model, which includes representations for the pertinent physiological dynamics and the hierarchical relationship between past and present patient characteristics. Then, we provide an algorithmic solution to the filtering problem using Sequential Monte-Carlo techniques. To demonstrate the merits of the proposed approach, we apply it to a case study of physiological monitoring for hemodynamic management. In this context, we investigate whether the proposed approach can provide reliable beliefs about the likely values and uncertainties associated with a patient's unmeasured physiological variables (e.g., hematocrit levels), characteristics (e.g., tendency for atypical behavior), and events (e.g., hemorrhage) given low-information measurements. Further details follow.

3.1. Background and Literature Review

Physiological monitoring systems are foundational tools for patient care that provide continuous information about a patient's vital physiological variables, helping practitioners in deriving insight into the patient's condition and making informed therapeutic decisions [117]. The need for physiological monitoring also spans the non-clinical domain, where a wide range of monitoring products (e.g., wearables and consumer electronics) aim to provide users with insight into their health and bodily performance [118]. In addition, in recent years, there has been considerable research interest in the area of autonomous medical care systems [19], [21], [23], [119], where physiological decision-support and closed-loop control algorithms are built to assist (or replace) humans in making therapeutic or lifestyle decisions based on monitoring results, making it even more necessary to develop high-fidelity and reliable physiological monitoring systems.

As a fundamental challenge in physiological monitoring, the physiological variables that are relevant to decision-making are not always directly measurable in practice. As a result, to be useful, monitoring systems must have mechanisms to infer unmeasured physiological variables from measured ones. In the context of engineering systems, *filtering algorithms* are excellent candidates for such purposes [120], [121]. These algorithms typically leverage an underlying model of the studied system (which could be mechanistic or black box) to process measurement signals and infer unmeasured variables. The Kalman Filter (KF) is one of the most well-known and widely used filtering algorithms in the engineering domain, which can utilize a linear underlying model for estimation purposes [122], [123]. The Extended Kalman Filter

(EKF) and the Unscented Kalman Filter (UKF) are two extensions to the KF algorithm that use model approximation techniques to allow for a nonlinear underlying model to be used for estimation [124]–[126]. More generally, Bayesian filtering provides a concrete framework for reasoning about filtering problems, where beliefs about the possible values of a system's unmeasured variables are expressed using probability distributions [127]. In this framework, principled mathematical procedures exist to alter existing (i.e., prior) beliefs according to incoming data in order to turn them into updated (i.e., posterior) beliefs about unmeasured variables [128], [129]. However, these mathematical operations are often not analytically tractable, giving rise to a wide range of approximate approaches that provide solutions to such problems in practical applications. These approaches can be divided into two broad categories with respect to their approximation technique. Sequential Monte-Carlo (SMC) filtering approaches approximate belief distributions using (large) collections of weighted samples (i.e., particles), allowing for stochastic and/or nonlinear underlying models to be used in the filtering process with relatively high fidelity [130]–[134], while Variational Filtering (VF) approaches approximate belief distributions using tunable distribution models, allowing for the filtering problem to be converted into an optimization problem [135]–[138].

Despite the demonstrated success of filtering algorithms in estimating unmeasured variables in many engineering systems, unique filtering challenges remain in the physiological monitoring domain, especially when the available measurements contain limited information about the unmeasured variables. For instance, in the context of hemodynamic management for critical patients, beliefs about the patient's

blood volume/composition and cardiovascular function are expected to facilitate a practitioner's (or an algorithm's) decisions in administering therapy (e.g., fluid or drug infusions). However, clinically available measurements are limited to blood pressure and (in rare cases) intermittent cardiac output and hematocrit correlates, which are also typically affected by high levels of noise and artifacts. Moreover, a patient's physiological dynamics are typically not excited frequently in clinical settings, further limiting the information content of the available measurements. Such intermittent and low-information measurements limit the applicability of many established filtering algorithms (and underlying physiological models) to the task of characterizing patients and providing reliable estimates of their physiological state.

To address this challenge, in this chapter, we propose the population-informed particle filter (PIPF), a filtering approach that leverages a *generative physiological model* (as in Chapter 2) to incorporate past experiences with patients into the filtering process in order to provide more reliable beliefs about a new patient's physiological state. To derive the PIPF, we formulate the filtering problem as recursive Bayesian inference on a probabilistic graphical model, which includes representations for the pertinent physiological dynamics and the hierarchical relationship between past and present patient characteristics. Then, we provide an algorithmic solution to the filtering problem based on SMC (i.e., particle-based) techniques. To demonstrate the potential merits and limitations of the proposed approach, we apply it to a case study of physiological monitoring for hemodynamic management. In this context, we show that the proposed approach can provide reliable beliefs about the likely values and uncertainties associated with a patient's unmeasured physiological variables (e.g.,

hematocrit levels), characteristics (e.g., tendency for atypical behavior), and events (e.g., hemorrhage) given low-information measurements. Based on these results, we discuss whether this approach may be applied to a wider range of monitoring problems with intermittent and low-information measurements, and serve as a basis for designing interpretable and context-aware medical decision-support and closed-loop control algorithms.

3.2. Population-Informed Particle Filtering

In this section, we present the population-informed particle filter (PIPF). First, we provide a review of relevant concepts from generative physiological modeling. Based on these concepts, we present the population-informed filtering scheme, where a generative physiological model informs a recursive Bayesian filter. Then, we provide an algorithmic solution to this problem using SMC techniques and show how this algorithm and methodology may be leveraged to create a robust model-based physiological monitoring system. Further details follow.

3.2.1. Generative Physiological Modeling

Physiological models can serve as a concrete source of physiological knowledge in the design and development of patient monitoring algorithms. Generative modeling is a promising approach to physiological modeling, where the inherent variability and stochasticity of a physiological system is fully embraced with model components that behave in stochastic but patterned ways. The objective of the generative physiological model is therefore to reproduce and predict the patterned randomness that is often

observed in physiological datasets. As the proposed filtering approach in this work relies heavily on an underlying generative model, in this section, we present a general family of generative models for physiological systems and provide an overview of procedures used to characterize these models from data. We refer the readers to Chapter 2 further details on this topic.

The generative model considered in this work consists of a hierarchy of stochastic components that reflect the physiological data observed in a population of patients (see Figure 3.1(a)). At the highest level in the hierarchy, a patient generator model is tasked with generating variations in patient characteristics, which can be formalized as:

$$\theta_1 \sim \mathcal{G}_1(\phi) \tag{3.1}$$

$$\theta_k \sim \mathcal{G}_k(\theta_{k-1}, \phi) \tag{3.2}$$

where $\mathcal{G}_1$ is a component that instantiates virtual patients by generating patient characteristics, and $\mathcal{G}_k$ is a component that produces variations in a virtual patient's characteristics as time progresses. In this formulation, ϕ is the vector of parameters for the patient generator model, θ_k is the vector of patient characteristics at time k, and the symbol $\sim$ denotes sampling. At the second level, a physiological dynamics model is tasked with generating evolutions in the states of each virtual patient, which can be formalized as:

$$x_1 \sim \mathcal{H}_1(\theta_1) \tag{3.3}$$

$$x_k \sim \mathcal{H}_k(x_{k-1}, \theta_{k-1}, u_{k-1}) \tag{3.4}$$

where $\mathcal{H}$ is the physiological dynamics model, x_k represents the states of the virtual patient at time k, and u_k represents the known inputs/therapies given to the virtual patient at time k. Finally, at the third level, a physiological measurement model generates observations from the state and/or the characteristics of the patient, which can be formalized as:

$$y_k \sim \mathcal{M}(x_k, \theta_k, n) \tag{3.5}$$

where $\mathcal{M}$ is the physiological measurement model, which is parameterized by n, and y_k is the vector of virtual observations generated at time k.

Given a dataset containing physiological data from a cohort of patients, we are interested in inferring the unknown parameters of the generative model in (3.1)-(3.5) such that the model captures the characteristics of the dataset. We recently showed in [67] that an effective solution to this (often intractable) problem can be obtained using variational Bayesian inference methods [139], where the most-likely values and the uncertainties associated with the parameters of the generative model (i.e., ϕ, n) are computed through stochastic optimization. These inferred parameters can in turn be used with the generative model to generate virtual datasets with similar distribution to real data. In other words, the resulting generative model incorporates the knowledge needed to instantiate virtual patients, generate paths for patient characteristics, produce state evolutions in response to given stimuli, and generate realistic physiological measurements. In this work, we are interested in utilizing this encoded knowledge to inform a filtering algorithm's real-time perception of a patient, especially when only intermittent and low-information measurements are available from the patient.

Therefore, in the next section, we formulate a filtering problem where the filter is informed by a generative model.

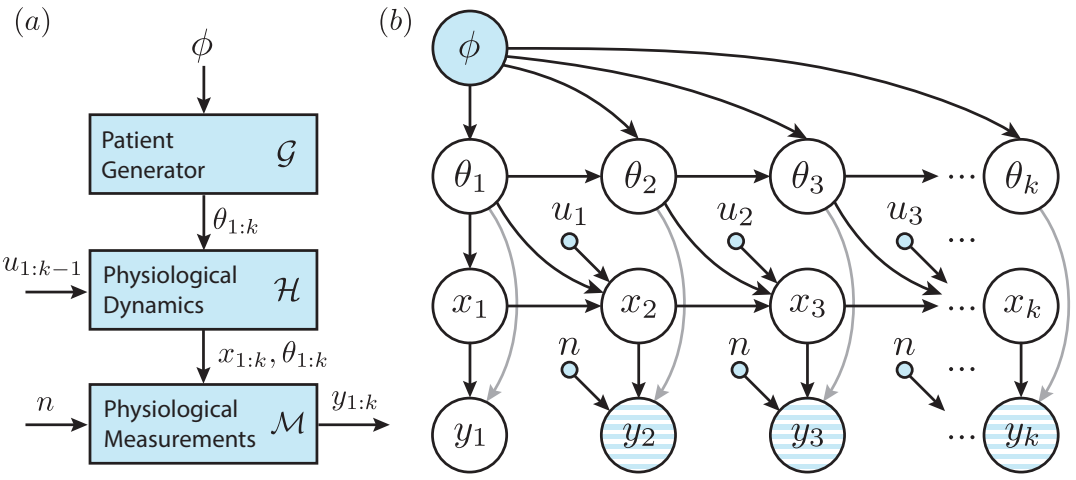

Figure 3.1. Schematic representations for (a) the hierarchical structure of the generative physiological model; and (b) the sequential relationship between the variables of interest in the population-informed filtering problem.

3.2.2. Population-Informed Filtering

Filtering (as referred to in this work) is the process of forming and updating beliefs about the unmeasured variables of a dynamic system using its measured variables. In the context of physiological monitoring, filtering may be used to form and update beliefs about the unmeasured states and characteristics of a patient using physiological measurements that arrive over time. Filtering processes typically operate based on an underlying model of the studied system. In this section, we formulate the problem of *population-informed filtering*, which is a filtering process that is informed by a generative physiological model.

To formulate the population-informed filtering problem, we adopt a Bayesian view of filtering, where beliefs are represented by probability density functions. In this

setting, the objective of filtering is to obtain the following *posterior probability density function* at each time k:

$$p(x_{1:k}, \theta_{1:k} | y_{1:k}, \phi, n) = \frac{p(x_{1:k}, \theta_{1:k}, y_{1:k} | \phi, n)}{p(y_{1:k} | \phi, n)} \tag{3.6}$$

This density represents our beliefs about the unmeasured states and characteristics of a patient up to time k (denoted by $x_{1:k}$ and $\theta_{1:k}$ respectively), given the patient's physiological measurements up to time k (denoted by $y_{1:k}$) and the information encoded in the parameters of the generative model ϕ, n. As shown on the right-hand-side of (3.6), it is conceptually possible to obtain the posterior density from the *joint density* shown in the numerator and the *marginal density* shown in the denominator. The joint density represents the relationship between the measured and unmeasured variables in the filtering problem, while the marginal density is obtained by integrating the joint density over the unmeasured variables.

Using the generative model in (3.1)-(3.5) to inform the filtering process imposes a specific structure on the joint density. This structure is shown in Figure 3.1(b) as a graphical representation. In this representation, the joint density has a hierarchical structure, where the patient generator parameters ϕ affect the patient's characteristics $\theta_{1:k}$, while the patient's characteristics affect the patient's state evolutions $x_{1:k}$. In turn, the patient's characteristics, state evolutions, and measurement parameters n affect the generated measurements $y_{1:k}$. The graphical representation also shows that the joint density consists of a sequential structure, where variables at time k are affected by variables at time $k-1$. This implies a recursive relationship between joint densities at times k and $k-1$, which can be formalized as:

$$p(x_{1:k}, \theta_{1:k}, y_{1:k}|\phi, n) = p(x_{1:k-1}, \theta_{1:k-1}, y_{1:k-1}|\phi, n).\, R_{k-1:k} \tag{3.7}$$

$$R_{k-1:k} = p(y_k|x_k, \theta_k, n)p(x_k|x_{k-1}, \theta_{k-1}, u_{k-1})p(\theta_k|\theta_{k-1}, \phi) \tag{3.8}$$

where $R_{k-1:k}$ is the ratio between joint densities at times k and $k-1$. The first multiplicative term in the ratio is the density associated with generating the measurement y_k given the current states x_k and characteristics θ_k; the second term is the density associated with generating a transition to the states x_k given the previous states x_{k-1}, characteristics θ_{k-1}, and inputs/therapies u_{k-1}; and the third term is the density associated with generating a transition to the characteristics θ_k given the previous characteristics θ_{k-1}.

The relationship described in (3.7)-(3.8) results in an important recursive relationship between posterior densities at times k and $k-1$. This relationship can be written as:

$$p(x_{1:k}, \theta_{1:k}|y_{1:k}, \phi, n) = p(x_{1:k-1}, \theta_{1:k-1}|y_{1:k-1}, \phi, n)\frac{R_{k-1:k}}{p(y_k|y_{1:k-1}, \phi, n)} \tag{3.9}$$

which suggests that it is conceptually possible to multiply the posterior density at time $k-1$ by an update term (i.e., the fraction on the right-hand side) to obtain the posterior density at time k. In other words, performing this update recursively would create a recursive filtering process, where previous beliefs about a patient's unmeasured states and characteristics can be updated according to the behavior of the generative model and any newly available measurements y_k. However, depending on the underlying generative model, this update procedure may often prove analytically intractable. In

the next section, we adopt a Sequential Monte-Carlo (i.e., particle-based) approach to derive a practical solution to this filtering problem.

3.2.3. A Sequential Monte-Carlo Solution

To derive a practical solution to the population-informed filtering problem presented in (3.9), we adopt a Sequential Monte-Carlo approach [133], [134], which is a sampling-based approach especially suited to solving inference problems with substantial nonlinearities, complex/multi-modal beliefs, and high levels of uncertainty, most of which are likely to arise in physiological monitoring applications. To perform this derivation, we first consider a *proposal density* of the following form:

$$q(x_{1:k}, \theta_{1:k}) = q(x_{1:k-1}, \theta_{1:k-1}) q(x_k, \theta_k | x_{k-1}, \theta_{k-1}) \tag{3.10}$$

which is designed to have a sequential structure, be easy to sample from, and have non-zero density wherever the posterior density (3.9) is expected to have non-zero density. The function of this proposal density is to propose beliefs about the unmeasured variables of the filtering problem. Multiplying and dividing (3.9) by (3.10) and rearranging the result yields the following variant of the recursive relationship between the posterior densities:

$$p(x_{1:k}, \theta_{1:k} | y_{1:k}, \phi, n) = p(x_{1:k-1}, \theta_{1:k-1} | y_{1:k-1}, \phi, n) \frac{\alpha(x_{k-1:k}, \theta_{k-1:k}) q(x_k, \theta_k | x_{k-1}, \theta_{k-1})}{p(y_k | y_{1:k-1}, \phi, n)} \tag{3.11}$$

where $\alpha(x_{k-1:k}, \theta_{k-1:k})$ is a *weighting function* defined as:

$$\alpha(x_{k-1:k}, \theta_{k-1:k}) = \frac{R_{k-1:k}}{q(x_k, \theta_k | x_{k-1}, \theta_{k-1})} \tag{3.12}$$

To perform the recursive filtering procedure described by (3.11)-(3.12), we start from the process of drawing samples from the proposal density in (3.10):

$$[x, \theta]_k^i \sim q\left(x_k, \theta_k | [x, \theta]_{k-1}^i\right) \tag{3.13}$$

where $[x, \theta]_k^i$ denotes a sample (at time k and indexed by i) from the proposal density, and it follows from the sequential structure of the proposal density that samples at time k can be generated recursively using samples at time $k-1$. Given N such samples, the proposal density itself can be expressed using the following equation:

$$\hat{q}\left(x_k, \theta_k | [x, \theta]_{k-1}^i\right) = \frac{1}{N} \sum_{i=1}^{N} \delta[x, \theta]_k^i \tag{3.14}$$

where $\delta[x, \theta]_k^i$ denotes the Dirac's delta function positioned at the sample $[x, \theta]_k^i$, and the equation indicates that the proposal density can be approximately represented by a collection of N samples $[x, \theta]_k^i$ drawn from it. Given the samples, the approximation in (3.14) can be substituted into (3.11) to yield the following sampling-based expression for the posterior density:

$$\hat{p}(x_{1:k}, \theta_{1:k} | y_{1:k}, \phi, n) = \sum_{i=1}^{N} w_k^i \delta[x, \theta]_{1:k}^i \tag{3.15}$$

where the posterior density is expressed in terms of weighted samples from the proposal density, and w_k^i denotes the weight for sample i at time k. Each sample path $[x, \theta]_{1:k}^i$ describes a proposed path for the patient's states and characteristics, and the weight w_k^i represents how probable the path is (relative to other sample paths) according to the generative model and the patient's physiological measurements up to time k. Each

weight at time k can be computed recursively from its counterpart at time $k-1$ using the following relationship:

$$w_k^i = \frac{w_{k-1}^i \alpha([x,\theta]_{k-1:k}^i)}{\sum_{i=1}^N w_{k-1}^i \alpha([x,\theta]_{k-1:k}^i)} \tag{3.16}$$

where w_{k-1}^i denotes the weight for sample i at time $k-1$, and $\alpha([x,\theta]_{k-1:k}^i)$ is the weighting function described in (3.12), evaluated at sample i for times k and $k-1$. Overall, the relationship described in (3.16) allows for the efficient calculation of the posterior density using recursive steps, namely: drawing N samples at time k from the proposal density (3.13) using the already available N samples from time $k-1$; evaluating the weighting function (3.12) using each of the N samples at times k and $k-1$; and calculating the new weights using (3.16), thereby obtaining the most up-to-date beliefs at time k. Realizing this procedure in practice, however, necessitates a few extra steps, which are presented in more detail in the next section.

3.2.4. The Population-Informed Particle Filtering (PIPF) Algorithm

As presented in Section 3.2.3, beliefs about a patient's unmeasured states and characteristics can be expressed and recursively updated using a collection of weighted samples (i.e., particles). In this section, we leverage this principle in conjunction with additional weight normalization and resampling techniques to build a practical procedure for population-informed particle filtering. Algorithm 3.1 shows an overview of this procedure. In this algorithm, we use the generative physiological model in (1)-(5) to generate proposal samples. As a result, the weighting function used in (16)

reduces to $\alpha\left([x,\theta]_{k-1:k}^{i}\right) \triangleq p\left(y_k \middle| x_k^i, \theta_k^i, n\right)$. The filtering procedure begins by using the generative model to generate (many) proposed samples for the patient's baseline states and characteristics. Then, at each time increment, the generative model proposes transitions to these states and characteristics, and the resulting samples are evaluated against data at the time to update the sample weights. The weight updates are performed in logarithmic scale and the weight normalizations are performed using the log-sum-exp trick [140] to achieve higher numerical accuracy and stability. In addition, weight variability is measured at each time increment, and if necessary, the samples are resampled using a systematic resampling [134] technique. This technique redraws each sample with a probability proportional to its weight, thereby allowing for a resetting of the weights. Overall, at each time k, this algorithm provides up-to-date beliefs about a patient's unmeasured variables in the form of a collection of weighted samples (i.e., particles).

Algorithm 3.1. *Population-Informed Particle Filtering (PIPF)*

$k \leftarrow 1$ {initialize time}

$\log w_1^i \leftarrow \log(1/N)$ {initialize N weights, indexed by i}

$\theta_1^i \sim \mathcal{G}_1(\phi),\ x_1^i \sim \mathcal{H}_1(\theta_1^i)$ {generate N samples, indexed by i; see (3.1), (3.3)}

Repeat:

$k \leftarrow k + 1$ {increment time}

$\theta_k^i \sim \mathcal{G}_k(\theta_{k-1}^i, \phi)$ {generate transition to characteristics; see (3.2)}

$x_k^i \sim \mathcal{H}_k(x_{k-1}^i, \theta_{k-1}^i, u_{k-1})$ {generate transition to states; see (3.4)}

$\log \alpha^i \leftarrow \log p(y_k|x_k^i, \theta_k^i, n)$ {get log-likelihood; see (3.12)}

$\log \overline{w}_k^i \leftarrow \log w_{k-1}^i + \log \alpha^i$ {update weights, un-normalized; see (3.16)}

$\log \overline{w}_k^* \leftarrow \max_i(\log \overline{w}_k^i)$ {get maximum weight}

$\log w_k^i \leftarrow \log \overline{w}_k^i - \log \overline{w}_k^* - \log \sum_i \exp(\log \overline{w}_k^i - \log \overline{w}_k^*)$ {normalize}

If $C\left(\log w_k^i\right) > T_w$: {check if weight variability is too high}

$[x, \theta]_k^i \leftarrow \mathcal{R}(\log w_k^i, [x, \theta]_k^i)$ {resample; see [134]}

$\log w_k^i \leftarrow \log(1/N)$ {reset weights}

Output: $\log w_k^i$, $[x, \theta]_k^i$ {provide up-to-date beliefs for time k}

3.3. Application to Hemodynamic Monitoring

Hemodynamic management is an essential aspect of care for critically ill patients, where the performance of a patient's cardiovascular system is monitored and, if necessary, therapies are administered to ensure adequate blood circulation [141], [142]. To demonstrate the merits and limitations of PIPF, we use this approach to build a monitoring algorithm for a typical critical care scenario where hypovolemia (i.e., low circulating blood volume) is treated with fluid infusions [143]. In this scenario, the monitoring algorithm receives a stream of blood pressure and fluid infusion data and processes this information to form real-time beliefs about a patient's unmeasured physiological variables (e.g., hematocrit levels), characteristics (e.g., the tendency for atypical behavior), and events (e.g., hemorrhage). These real-time beliefs can in turn inform a practitioner's decision to administer fluids, and furthermore serve as a

foundation for building automated decision-support and closed-loop control algorithms for hypovolemia treatment. This section presents the details of applying the PIPF approach to this problem and describes the methods and datasets used to evaluate the results. Further details follow.

3.3.1. Generative Modeling for Hemodynamic Management

As presented in Section 3.2, the PIPF approach relies on a generative physiological model to form and update beliefs about the unmeasured states and characteristics of a patient. In this section, we describe a generative model of the physiological phenomena relevant to the treatment of hypovolemia with fluid infusions. This model will be used in subsequent sections to build a monitoring algorithm for hypovolemia treatment.

Patient Generator Model: As presented in Section 3.2.1, the generative physiological model consists of a hierarchy of stochastic components that reflect the physiological behaviors observed in a population of patients. At the highest level in the hierarchy, a patient generator model $\mathcal{G}$ generates virtual patients with a variety of characteristics. In the context of hypovolemia treatment, we consider a patient generator model of the following form:

$$\theta_1 = \phi_\mu + \phi_L \epsilon, \quad \epsilon \sim \mathcal{N}(0, I) \tag{3.17}$$

$$\theta_k = (1 - b)\theta_{k-1} + b(\phi_\mu + \phi_L \epsilon), \quad b \sim \mathcal{B}(1, \gamma_\theta), \epsilon \sim \mathcal{N}(0, I) \tag{3.18}$$

In this model, first, the characteristics θ_1 of each virtual patient are instantiated in (17) by drawing samples from a full-covariance Gaussian distribution, where ϕ_μ determines the center of the distribution, ϕ_L is the Cholesky decomposition of the distribution's

covariance matrix (i.e., $\phi_\Sigma = \phi_L \phi_L^T$), and ϵ is a sample from the standard Gaussian distribution of appropriate dimension. Then, over time, each patient's characteristics θ_k follow the relationship in (18), where b is a sample from the Bernoulli distribution, which produces $b = 1$ with a probability of γ_θ, and $b = 0$ with a probability of $1 - \gamma_\theta$. The probability γ_θ (which is chosen to be small) acts as a forgetting factor for the characteristics of each instantiated virtual patient. In other words, each virtual patient will retain its characteristics over time, except for a small chance of transitioning to different characteristics drawn from the virtual patient generator in (17). For the purpose of hypovolemia treatment modeling, we define a vector of patient characteristics as follows:

$$\theta_k = \left[v_0 \;\; H_0 \;\; Q_0 \;\; P_{a0} \;\; K_v \;\; K_a/K_v \;\; K_p \;\; \alpha_I \;\; \alpha_H \;\; K_h \;\; \tau_R \;\; K_R \;\; \beta_v \;\; K_Q \right]_k \tag{3.19}$$

which contains $n_p = 14$ physiological parameters to be elaborated on later in this section. Overall, this patient generator model is built to inform a PIPF-based monitoring algorithm about the characteristics that are likely to occur in the patient population (along with their probability of occurrence) and furthermore allow for the algorithm to adapt to potential changes in a patient's characteristics over time.

Physiological Dynamics Model: At the second level in the hierarchical model in Section 3.2.1, a physiological dynamics model $\mathcal{H}$ generates state evolutions for each virtual patient. For this purpose, we utilize a dynamic model of the physiological phenomena relevant to hypovolemia treatment developed in Chapter 2. In this section, we present an overview of this model in discretized form. The model consists of macro-

state components that represent blood circulation and the mechanisms that affect blood circulation in the context of hypovolemia treatment.

To obtain a macro-state model of blood circulation, we consider equations of the following form:

$$[v_a]_k = [v_a]_{k-1} + \delta t\,[Q - (P_a - P_v)/R - J_H - J_F]_{k-1} \tag{3.20}$$

$$[v_v]_k = [v_v]_{k-1} + \delta t\,[-Q + (P_a - P_v)/R + J_I]_{k-1} \tag{3.21}$$

$$[v_r]_k = [v_r]_{k-1} + \delta t\,[-J_H H]_{k-1} \tag{3.22}$$

where v_a and v_v denote arterial and venous blood volumes, v_r is the total red blood cell volume, δt is the time increment between k and $k-1$ instances, Q is the cardiac output, P_a and P_v denote mean arterial and venous blood pressures, R is the systemic vascular resistance, $H = v_r/(v_a + v_v)$ is the hematocrit, J_F is the net rate of fluid exchange with the interstitial space, J_I is the rate of fluid infusion, and J_H is the rate of blood loss. In these equations, mean arterial pressure is related to arterial volume through $P_a = P_{a0} + K_a(v_a - v_{a0})$, and mean venous pressure is related to venous volume through $P_v = P_{v0} + K_v(v_v - v_{v0})$, where P_{a0}, P_{v0} are baseline (i.e., un-perturbed) arterial and venous blood pressures, v_{a0}, v_{v0} are baseline arterial and venous blood volumes, and K_a, K_v denote arterial and venous elastances in the patient.

To obtain a macro-state model of the mechanisms that affect blood circulation in the context of hypovolemia treatment, we consider: (i) fluid exchange with the interstitial space, (ii) changes in systemic vascular resistance (e.g., through vasoconstriction or vasodilation), and (iii) changes in cardiac output (e.g., through

changes to heart rate and contractility). Fluid exchange with the interstitial space is represented by:

$$J_F = K_p(v - v_0 - r_F) \tag{3.23}$$

$$[r_F]_k = [r_F]_{k-1} + \delta t\,[J_I/(1+\alpha_I) - J_H/(1+\alpha_H)]_{k-1} \tag{3.24}$$

where r_F is the steady-state change in blood volume after the exchange, $v = v_a + v_v$ is the total blood volume, $v_0 = v_{a0} + v_{v0}$ is the baseline total blood volume, K_p modulates the speed of exchange, and α_I, α_H determine the fraction of fluid infusion and blood loss that are compensated in the exchange. Changes in systemic vascular resistance are represented by:

$$R = R_0 + K_h(H - H_0) + s_R \tag{3.25}$$

$$[s_R]_k = [s_R]_{k-1} + \delta t\,[-s_R/\tau_R - K_R(P_a - P_{a0})/\tau_R]_{k-1} \tag{3.26}$$

where s_R is the change in resistance prompted by the body's compensatory mechanisms, K_R, τ_R modulate the characteristics of compensation, $R_0 = (P_{a0} - P_{v0})/Q_0$ is the baseline resistance, Q_0 is the baseline cardiac output, H_0 is the baseline hematocrit, and K_h modulates the effect of blood dilution on resistance. Changes in cardiac output are represented by:

$$Q = Q_0 + \beta_v(P_v - P_{v0}) + s_Q \tag{3.27}$$

$$[s_Q]_k = [s_Q]_{k-1} + \delta t\,[-K_Q(Q - Q_0)]_{k-1} \tag{3.28}$$

where s_Q is the change in cardiac output prompted by the body's compensatory mechanisms, K_Q modulates the characteristics of compensation, and β_v modulates the effect of variations in central venous pressure on cardiac output. The physiological

model described in (3.20)-(3.28) is intended to inform a PIPF-based monitoring algorithm about the physiological dynamics and behaviors that could occur in a patient, thereby giving the algorithm a basis to form and adjust its beliefs about the patient's state over time.

Hemorrhage Events Model: The model described above is built to represent the physiological response of a patient to fluid infusions (i.e., the rate J_I) and hemorrhage (i.e., the rate J_H). In the context of hypovolemia treatment, fluid infusions are typically known, as they would be administered to the patient by a caregiver (or an algorithm). In contrast, the presence of hemorrhage may be unknown in many cases. In such cases, a model should be devised to represent the possibility of unknown hemorrhage events. For this purpose, we consider a stochastic model of the following form:

$$[J_H]_1 = g, \quad g \sim GPD(0, \sigma_H, \xi_H) \tag{3.29}$$

$$[J_H]_k = (1-b)[J_H]_{k-1} + bg, \quad b \sim \mathcal{B}(1, \gamma_H), g \sim GPD(0, \sigma_H, \xi_H) \tag{3.30}$$

In this model, first, possible hemorrhage rates $[J_H]_1$ are instantiated in (3.29) by drawing samples g from a generalized Pareto distribution located at zero, where σ_H is the scale parameter and ξ_H is the shape parameter. This distribution represents a range of possible hemorrhage rates, where lower hemorrhages have a higher probability of occurrence. Over time, the rate of hemorrhage in a patient follows the relationship in (3.30), where b is a sample from the Bernoulli distribution, and the probability γ_H acts as a forgetting factor for the hemorrhage rate. According to this model, the rate of hemorrhage in a patient retains its value over short periods of time, except for the possibility of transitioning to a different rate drawn from the hemorrhage model in

(3.29). Overall, this model is built to inform a PIPF-based monitoring algorithm about the possibility of unknown hemorrhage events, and furthermore allow the algorithm to form beliefs about the presence of unknown hemorrhage in a patient over time.

Physiological Measurement Model: At the third level in the hierarchical model in Section 3.2.1, a physiological measurement model $\mathcal{M}$ generates observations from the characteristics and state evolutions generated by the first two levels (i.e., $\mathcal{G}$ and $\mathcal{H}$). In the context of hypovolemia treatment, mean arterial pressure measurements are typically available in clinical settings, while cardiac output and hematocrit measurements may be available only in experimental settings. To represent these potentially measured physiological variables, we consider the following model:

$$[y_{HCT}]_k = [H]_k + n_{HCT}\epsilon, \quad \epsilon \sim \mathcal{N}(0, I) \tag{3.31}$$

$$[y_{CO}]_k = [Q]_k + n_{CO}\epsilon, \quad \epsilon \sim \mathcal{N}(0, I) \tag{3.32}$$

$$[y_{MAP}]_k = [P_a]_k + n_{MAP}\epsilon, \quad \epsilon \sim \mathcal{N}(0, I) \tag{3.33}$$

where y_{HCT}, y_{CO}, and y_{MAP} denote the generated measurements for hematocrit, cardiac output, and mean arterial pressure respectively, and n_{HCT}, n_{CO}, and n_{MAP} denote the standard deviation of noise acting on these measurements. Overall, this measurement model is intended to inform a PIPF-based monitoring algorithm about the noises/artifacts that may corrupt each measured variable, thereby giving the algorithm a way to attribute such disturbances as they occur.

3.3.2. The PIPF-Based Monitoring Algorithm

As suggested in Section 3.2, given a generative model of the relevant physiological phenomena (such as the model introduced in Section 3.3.1), Algorithm 3.1 can be used to create a monitoring system for hypovolemia treatment. The notable steps of this procedure are the following:

- To perform $\theta_1^i \sim \mathcal{G}_1(\phi)$, we use (3.17) to generate N proposed samples θ_1^i representing a patient's possible baseline characteristics.
- To perform $x_1^i \sim \mathcal{H}_1(\theta_1^i)$, we use the generated θ_1^i's to initialize the states of the physiological model in (3.20)-(3.22), (3.24), (3.26), (3.28) as follows: v_a is set to $v_{a0} = 0.3v_0$, v_v is set to $v_{v0} = 0.7v_0$, v_r is set to $v_{r0} = H_0 v_0$, while r_F, s_R, and s_Q are set to zero. Also, the possible hemorrhage rates are initialized according to (3.29).
- To perform $\theta_k^i \sim \mathcal{G}_k(\theta_{k-1}^i, \phi)$, we use the relationship in (3.18).
- To perform $x_k^i \sim \mathcal{H}_k(x_{k-1}^i, \theta_{k-1}^i, u_{k-1})$, we use the relationships in (3.20)-(3.28), (3.30).
- To perform $\log \alpha^i \leftarrow \log p(y_k | x_k^i, \theta_k^i, n)$, we consider the measurement model in (3.31)-(3.33), which implies a likelihood function of the following form:

$$\log p\left(y_k \middle| x_k^i, \theta_k^i, n\right) = \sum_m \left[-\frac{1}{2n_m^2} \left([y_m]_k - [y_m]_k^d \right)^2 - \log(n_m) - \frac{1}{2}\log(2\pi) \right] \quad (3.34)$$

where the index m enumerates the available measured variables at time k (e.g., if mean arterial pressure and hematocrit measurements are available at time k, then $m \in \{MAP, HCT\}$), and $[y_m]_k^d$ denotes the measured value. Executing Algorithm 3.1 as

outlined above results in a monitoring algorithm that takes in measurements as they become available over time, forming running beliefs about the physiological states (i.e., v_a, v_v, v_r, r_F, s_R, and s_Q), characteristics (i.e., shown in the θ_k vector in (3.19)) and events (i.e., J_H) in a patient.

3.3.3. Physiological Data

To demonstrate the potential merits and limitations of the PIPF-based monitoring algorithm, we utilize a physiological dataset from previous work [48], [144], [145], which contains time-series measurements pertaining to hypovolemia treatment in 23 animal (sheep) experiments. In each experiment, the animal is subjected to controlled hemorrhage and subsequently resuscitated over time using fluid (crystalloid) infusions, which are administered according to the recommendations of a control algorithm. Each experiment spans a course of 180 minutes, where the subject's hematocrit, cardiac output, and mean arterial pressure are measured at a target period of ~5 minutes. This dataset is especially suited to the purpose of analyzing the PIPF-based monitoring algorithm, as it can be used to assess the algorithm's ability to form reliable beliefs about variables that are often unmeasured or unknown in clinical settings (e.g., hematocrit, cardiac output, and hemorrhage rates) by processing measurements that are typically available in clinical settings (e.g., mean arterial pressure and infusion rates). In the next section, we provide an overview of our data analysis procedures for this purpose.

3.3.4. Data Analysis

In order to assess the merits and limitations of the PIPF-based monitoring algorithm, we consider its application to a realistic scenario for hypovolemia treatment, and quantify its performance based on the physiological dataset described in Section 3.3.3. In this section, we describe the problem setting and the details pertaining to the evaluation of the algorithm.

Problem Setting: To evaluate the algorithm in a scenario that resembles real-world hypovolemia treatment, we consider the task of monitoring each subject in our dataset as it undergoes controlled hemorrhage and fluid resuscitation. In this scenario, the monitoring algorithm receives a stream of *mean arterial pressure* and *infusion rate* measurements (which are typically available in clinical settings) and is tasked with forming and updating beliefs about the subject's states, characteristics, and hemorrhage events over time. These beliefs can in turn be evaluated against available measurements from the subject not shown to the monitoring algorithm (i.e., hematocrit, cardiac output, and hemorrhage rates) in order to assess its performance. It is important to note that this problem setting is a purposefully challenging one, where many unmeasured variables are inferred from few measured variables. This setting enables us to assess the merits and limitations of the PIPF-based monitoring algorithm especially when the clinically available data is limited.

Algorithm Evaluation: As presented in Section 3.2.1, obtaining a PIPF-based monitoring algorithm involves (i) training a generative physiological model on a physiological dataset obtained from the population, and (ii) providing this model to

Algorithm 3.1. To evaluate the performance of the monitoring algorithm, we follow a leave-one-out cross-validation procedure. For each studied subject, we exclude the subject from the population dataset used to train the generative physiological model and use the resulting model in Algorithm 3.1 to obtain a PIPF-based monitoring algorithm. Then, we test the resulting algorithm on the excluded subject by providing the subject's mean arterial pressure and infusion rate data to the algorithm as a stream of measurements and assessing the performance of the algorithm based on the adequacy of its beliefs about the subject's unseen measurements (i.e., hematocrit, cardiac output, and hemorrhage rates). To quantify the adequacy of the beliefs provided by the algorithm, we utilize the mean continuous ranked probability score (MCRPS), which is a proper metric suitable for comparing beliefs to measured values [146]. A particle-based form of this metric can be written as follows:

$$\mathrm{MCRPS}(F_m, y_m) = \mathbb{E}_k\left[\mathbb{E}_{z_1 \sim [F_m]_k}|z_1 - [y_m]_k| - \frac{1}{2}\mathbb{E}_{z_1, z_2 \sim [F_m]_k}|z_1 - z_2|\right] \quad (3.35)$$

where $[F_m]_k$ denotes the belief provided by the monitoring algorithm about variable m at time k, which is expressed as a collection of weighted points, z_1 and z_2 are samples from this belief, and $[y_m]_k$ denotes the measured value for variable m at time k. Intuitively, this scoring function maximally promotes predictions that are sharp (i.e., certain) and accurate, and maximally discourages predictions that are sharp and inaccurate. In the case of deterministic predictions, the MCRPS metric reduces to mean absolute error.

PIPF versus Particle Filtering (PF): To highlight the potential advantages of the PIPF-based monitoring algorithm with respect to a more established approach that

does not consider population-level information, we build an alternative monitoring algorithm that operates based on the traditional particle filtering (PF) scheme [127], [134]. This PF-based monitoring algorithm differs from the PIPF-based one in that it does not include a patient generator model in its underlying model. According to PF convention, the unmeasured variables of the problem (i.e., the subject's states, characteristics, and events) are instead initialized by drawing a uniform sample of particles over a plausible range in the space of unmeasured variables. In addition, to enable the PF-based monitoring algorithm to adapt to the subject's characteristics, we consider a random walk model for the parameters over time $\theta_k = \theta_{k-1} + \sigma_{PF}\epsilon$ where $\epsilon \sim \mathcal{N}(0, I)$. To compare the PIPF-based and PF-based monitoring algorithms, we use the MCRPS scoring procedure described above. To test the significance of the score differences between the two approaches, we use the Wilcoxon signed-rank test. Overall, this comparison study is designed primarily to highlight the potential merits and limitations of incorporating population-level information into the filtering process.

3.4. Results and Discussion

Forming reliable beliefs about a patient's state is essential for algorithmic decision-making in medical care settings. This task is especially challenging when the physiological measurements available from the patient are intermittent or contain limited information. To address this challenge, we proposed the PIPF scheme, where the information encoded in a generative physiological model is leveraged to form more robust beliefs about a patient's state. This section presents the results of applying the

PIPF scheme to the problem of monitoring for hypovolemia treatment and discusses the merits and limitations of the PIPF scheme in this context.

3.4.1. Beliefs about Unmeasured Physiological Variables

Figure 3.2 shows the beliefs formed by the PIPF-based monitoring algorithm about hematocrit (HCT), cardiac output (CO), and mean arterial pressure (MAP) in two representative subjects (marked by Subject A and Subject B). These beliefs were formed when providing the algorithm with a stream of MAP (see in Figure 3.2; right column) and infusion rate (see Figure 3.3; blue line) measurements. Figure 3.5 shows the beliefs formed about the internal states and characteristics of Subject A in this scenario. As it is expected from the formulation of the PIPF algorithm, whenever a new measurement becomes available, the algorithm adjusts its beliefs about the subject's physiological states and characteristics. As a result, beliefs about the subject's MAP closely follow the MAP data, and the algorithm can provide beliefs about the subject's HCT and CO (see Figure 3.2) as well as its internal states and characteristics (see Figure 3.5). These beliefs appear consistent with the HCT and CO data unseen by the algorithm.

These representative results also highlight two notable behavior patterns in the beliefs generated by the PIPF algorithm in relation to the informativeness of the data. First, CO is an example of a variable that could be *strongly informed* by MAP data because of its close physiological relationship with MAP. Namely, the pressure in the arterial space, which has a relatively low compliance, is strongly affected by the flow rate of blood coming into the space, especially within short timeframes. As a result,

the monitoring algorithm can initially infer CO in large part from the information contained in MAP, giving relatively sharp beliefs about CO that are also consistent with the CO data (see Figure 3.2; center column). Second, baseline HCT is an example of a variable that is *weakly informed* by MAP data because of its distant physiological relationship with MAP. Namely, MAP is primarily affected by blood volume kinetics (i.e., the volume of blood that ends up in the arterial space at each time) while baseline HCT is a measure of blood composition. In this scenario, the monitoring algorithm forms its beliefs about baseline HCT in large part from the population-level information encoded in the generative physiological model. As a result, the beliefs about baseline HCT appear to span the range of plausible HCT values in the population (see Figure 3.2; left column). Overall, these results suggest that the PIPF-based monitoring algorithm shows promise in forming beliefs about the unmeasured variables of the physiological system by combining the information contained in subject-specific measurements with the population-level information encoded in the generative physiological model.

3.4.2. Beliefs about Subject Characteristics

Figure 3.5 shows (in the top panel) the posterior beliefs about the characteristics of Subject A. These beliefs were formed when providing the algorithm with a stream of MAP (see in Figure 3.2; top-right column) and infusion rate (see Figure 3.3; top-panel; blue line) measurements. For most characteristics, the beliefs maintain a high entropy (i.e., large spread) throughout the filtering process, and resemble the characteristics generated by the generative model. However, in the initial phases of the experiment

(e.g., first 60 minutes), where the subject undergoes significant perturbations (i.e., hemorrhage and fluid infusions), the beliefs show some deviations from the characteristics generated by the generative model. The high entropy of the beliefs indicates that the characteristics of the subject are, for the most part, only weakly identifiable from a stream of MAP and infusion rate measurements, especially when the perturbations do not sufficiently excite the subject's physiology. By design, the PIPF algorithm reverts to the characteristics generated by the generative model in such conditions. In other words, in the absence of informative measurements or excitations, the algorithm opts for considering all possibilities with regards to subject characteristics as informed by the generative physiological model.

Despite the high entropy, the beliefs about subject characteristics may be used to calculate useful summarized information about the subject. The *subject atypicality index* [147] is a notable example of this, which measures the presence of atypical characteristics in a subject by comparing the subject's characteristics to those of the population. Figure 3.4 shows the posterior beliefs about atypicality in two representative subjects. The index values lie between zero and one, with a value of one indicating a highly atypical subject. For Subject A, the beliefs about atypicality show a high level of entropy and uniformly span the range between zero and one, which indicates that Subject A's atypicality cannot be established or rejected based on the given stream of MAP and infusion rate data. For Subject B, the beliefs about atypicality shift toward the higher end of the spectrum at two times during the experiment (starting at ~45-minute and ~90-minute marks). Inspecting Subject B's MAP data at those times reveals two instances of rapid rise in pressure to values even higher than the subject's

baseline (pre-hemorrhage) pressure, which is not generally expected from a typical subject undergoing hemorrhage and crystalloid resuscitation. Overall, these results suggest that, despite the high entropy, beliefs about a subject's characteristics may be summarized to obtain potentially useful information about the subject.

3.4.3. Beliefs about Physiological Events

Figure 3.3 shows beliefs about hemorrhage rate in two representative subjects when providing the monitoring algorithm with a stream of MAP and infusion rate measurements. In Subject A, the algorithm attributes the subject's sudden initial drop in MAP (see Figure 3.2; top-right panel) to the presence of hemorrhage, with beliefs about hemorrhage rate that are consistent with hemorrhage data unseen by the algorithm. Subsequently, the algorithm infers a cessation of hemorrhage (albeit with delay) and provides reasonable beliefs about the later instances of small hemorrhage. In Subject B, the algorithm detects the initial hemorrhage, but with some delay, which can be attributed to the first few MAP measurements remaining high despite the presence of large hemorrhage (see Figure 3.2; bottom-right panel). Subsequently, the algorithm detects one of the two smaller hemorrhages, depending on whether they affect the MAP measurements. Overall, these results suggest that the PIPF-based monitoring algorithm may be utilized to form useful beliefs about physiological events in a subject (such as hemorrhage), if the occurrence of the event leaves detectable effects on the measurements available from the subject.

3.4.4. The Effect of Population-Level Information

Table 3.1 compares cross-validation scores (based on MCRPS in N=23 subjects) between the PIPF algorithm and a conventional particle filtering (PF) algorithm described in Section 3.3.4. Comparing the scores between the two cases reveals that the PIPF approach provides superior beliefs about hemorrhage rate and HCT when compared to the conventional PF algorithm, while beliefs about CO are comparable between the two cases. This result may be explained based on whether the studied physiological variables are strongly informed or weakly informed by the available measurements. As presented in Section 3.4.1, CO is a variable that is strongly informed by MAP data due to its close physiological relationship with MAP. As a result, both algorithms can infer CO in large part from the information contained in MAP, giving relatively sharp beliefs about CO that are also consistent with the CO data. This results in comparable MCRPS scores for the two algorithms. In contrast, hemorrhage rate and HCT are variables that are weakly informed by MAP data because of their more distant physiological relationship with MAP. As a result, the PIPF-based monitoring algorithm forms its beliefs about hemorrhage rate and HCT in large part from the population-level information encoded in the generative physiological model, while the PF-based monitoring algorithm does not have access to the population-level information. This results in superior MCRPS scores for the PIPF-based monitoring algorithm. Overall, these results suggest that incorporating population-level information into the filtering process (as it is done in the PIPF algorithm) results in the formation of superior beliefs, especially for physiological variables that are weakly informed by the available subject-specific data.

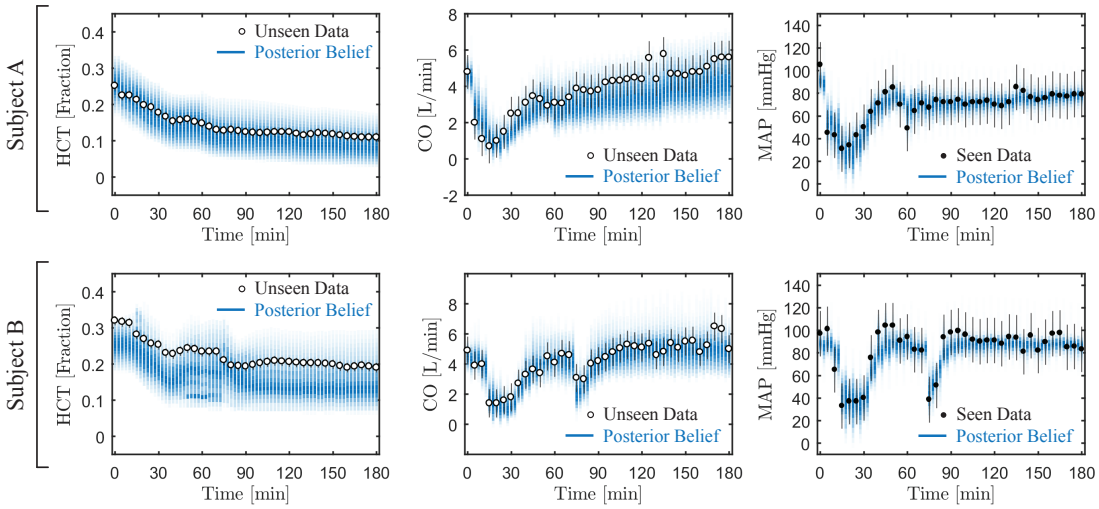

Figure 3.2. Marginalized posterior beliefs about cardiac output (CO) and hematocrit (HCT) in two representative subjects when presenting the PIPF algorithm with a sequence of mean arterial pressure (MAP) measurements. Bolder colors represent higher belief density. Black dots show data presented to PIPF, while white dots show data never presented to PIPF.

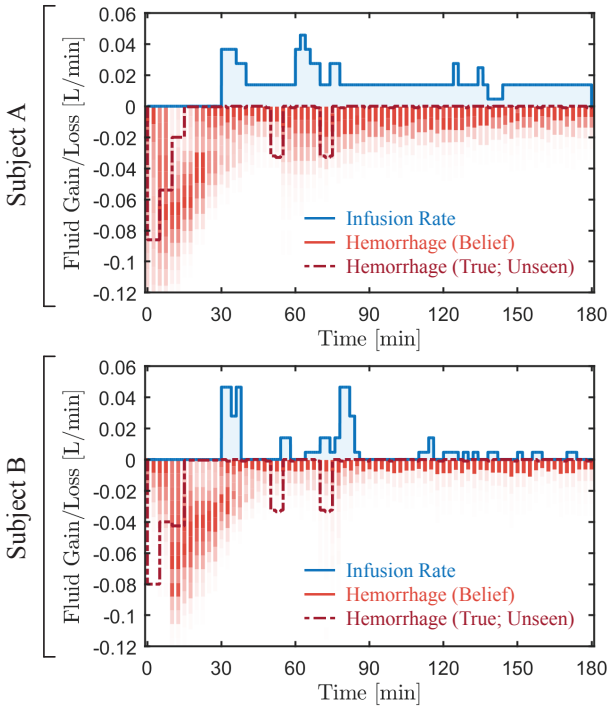

Figure 3.3. Marginalized posterior beliefs about hemorrhage rate in two representative subjects when presenting the PIPF algorithm with a sequence of mean arterial pressure (MAP) measurements. Bolder colors represent higher belief density. True hemorrhage rates (unknown to the algorithm) and infusion rates (known to the algorithm) are provided for reference.

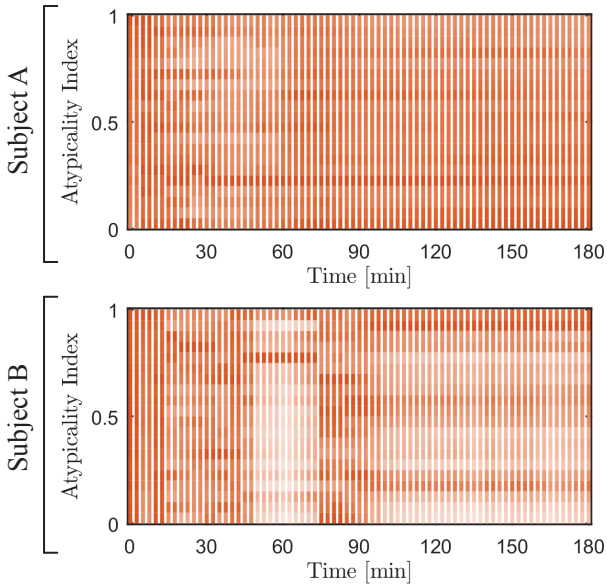

Figure 3.4. Marginalized posterior beliefs about atypicality in two representative subjects when presenting the PIPF algorithm with a sequence of mean arterial pressure (MAP) measurements. Bolder colors represent higher belief density.

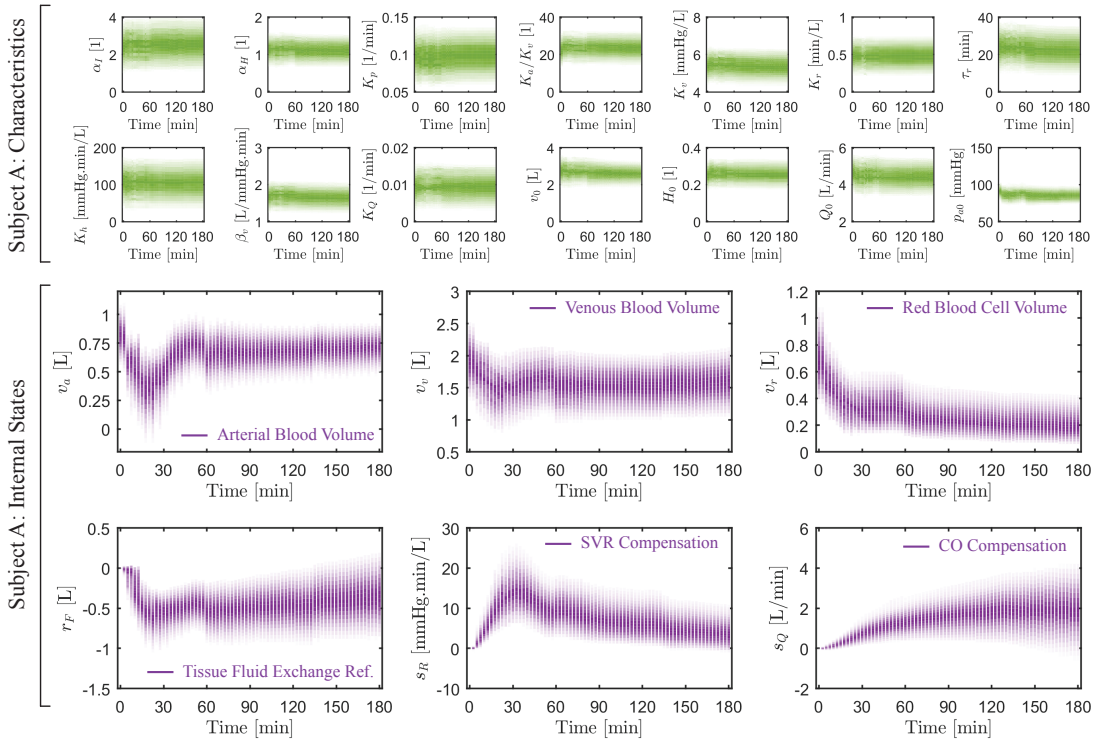

Figure 3.5. Marginalized posterior beliefs about a representative subject's states and characteristics when presenting the PIPF algorithm with a sequence of mean arterial pressure (MAP) measurements. Bolder colors represent higher belief density.

Table 3.1. Comparing cross-validation scores (based on MCRPS) between the PIPF algorithm and a conventional particle filtering (PF) approach, when presenting the algorithms with a sequence of mean arterial pressure (MAP) measurements [Median (Q1, Q3); N=23; Lower is better]

Method	Hemorrhage Rate [mL/min]	Hematocrit [%]	Cardiac Output [L/min]	Mean Arterial Pressure [mmHg]
PIPF	**10.03*** (8.29, 11.10)	**2.15*** (1.25, 3.91)	**0.52** (0.35, 0.89)	**3.85** (3.70, 4.11)
PF	**20.29** (17.17, 22.37)	**4.97** (4.18, 6.54)	**0.53** (0.42, 0.97)	**3.81** (3.67, 3.89)

3.5. Conclusions

In this chapter, we proposed the population-informed particle filter (PIPF), a Bayesian filtering approach that leverages a generative physiological model to provide beliefs about a patient's states, characteristics, and events in the context of physiological monitoring. Using a case study on monitoring for hemodynamic management, we showed that the PIPF approach can provide reasonable beliefs (as compared to excluded data) about the likely values and uncertainties associated with a patient's physiological variables (e.g., hematocrit levels), characteristics (e.g., tendency for atypical behavior), and events (e.g., hemorrhage). In addition, we demonstrated that incorporating population-level information into the filtering process (as is done in the PIPF algorithm) results in the formation of beliefs that are superior to those provided by a traditional particle filtering approach, especially for physiological variables that are weakly informed by the available patient-specific measurements. These results imply that PIPF is a promising candidate for use in physiological monitoring systems that process low-information and intermittent physiological measurements. Therefore, future efforts should be devoted to applying and evaluating the PIPF approach in a

wider range of physiological monitoring applications. In the next chapter, we investigate principled ways to leverage the beliefs provided by PIPF to design physiological decision-support and closed-loop control algorithms.

Chapter 4: An Inference-Based Framework Toward Transparent, Context-Aware, and Coordinated Physiological Closed-Loop Control

Physiological closed-loop control (PCLC) systems are recent developments in biomedical research that promise to leverage the vigilance, precision, and processing power of computers to replace humans in providing medical care to patients. Despite the demonstrated success of these systems in research settings, controlled experiments, and limited real-world applications, notable challenges remain to be addressed to enable their widespread adoption in real-world scenarios. In Chapter 1, we introduced five categories of outstanding desires and challenges related to the operation of such systems, which include considerations related to transparency, context-awareness, coordination, adaptation, and the handling of uncertainty. To consider these aspects in PCLC design, in this chapter we propose an inference-based framework for PCLC, where the process of perception and decision-making is formulated as recursive Bayesian inference on a probabilistic graphical model, and subsequently solved using a principled combination of sequential Monte-Carlo (SMC) and variational inference (VI) techniques. To demonstrate the merits and limitations of the proposed PCLC approach, we apply it to an in-silico case study on the autonomous resuscitation of hypovolemia (i.e., low circulating blood volume) using fluid and blood infusions. In this context, we investigate whether the proposed approach can successfully control a subject's blood pressure and blood composition through fluid and blood infusions, while exhibiting the desirable behaviors sought in the motivations of this work.

4.1. Background and Literature Review

Autonomous medical care systems [19], [23], [119], [148] are a relatively recent development in biomedical research that promise to leverage the continuous vigilance, precision, and processing power of computers to assist (or replace) humans in providing medical care to patients and/or general consumers. Active fields of research and product development in this area include autonomous anesthesia, hemodynamic management, and diabetes management, to name a few [20], [21], [23]. In general terms, an autonomous medical care system should be equipped to (i) measure relevant aspects of a patient's physiology, (ii) process those measurements to "understand" the patient's state, and (iii) leverage this understanding to make informed therapeutic decisions and apply them to the patient. As a result, decision-making algorithms are expected to be a core constituent of most autonomous medical care systems.

In the context of engineering systems, a wide array of methods and frameworks exist that may be utilized to create algorithms for decision-making. The field of *control theory* provides an especially wide range of methods for this purpose, where the overarching idea is to connect a judiciously-designed controller (e.g., a computational object containing static and/or dynamic rules) to the controlled system (e.g., the human physiology) to form a closed loop, and positively alter its behavior. Indeed, a large body of research has been devoted to designing and evaluating controllers for specific physiological control applications. This includes (but is not limited to) the use of PID, rule-based, fuzzy-logic, adaptive, and model-predictive controllers to perform anesthesia [28]–[30], [32], [34]–[36], [39], hemorrhage resuscitation [47]–[52], burn resuscitation [53]–[56], and vasopressor infusion [57]–[61], to name a few. As another

category of methods, the field of *probabilistic planning* provides a mathematical framework for decision-making, where model structures such as Markov Decision Processes (MDPs) and Partially Observable MDPs (POMDPs) are utilized to represent real-world decision-making scenarios and search for optimal strategies [149]–[151]. Similarly, the field of *reinforcement learning* provides principled ways to reason about decision problems, where the overarching idea is to consider decision-making algorithms as "agents" that interact with their environment to achieve desirable outcomes [152]. This point of view gives rise to algorithms that can reason about actions and rewards, make value judgements, and effectively search for and follow strategies to reach desirable outcomes [152]–[155]. These approaches have found application in a solving a wide range of real-world problems, including mastery of complex games, robot navigation and control, autonomous systems, and healthcare applications, to name a few [156]–[159].

Despite the demonstrated success of control theory, probabilistic planning, and reinforcement learning in solving decision problems in a variety of applications, notable outstanding desires and challenges remain to be addressed in the context of decision-making for autonomous medical care. To facilitate discussion, in this work, we introduce five categories of outstanding challenges specific to such applications. Ideally, autonomous medical care systems should be equipped to exhibit (i) *transparent behavior,* where the system's perceptions, reasoning, and decisions are human-interpretable; (ii) *context-aware behavior,* where the system is capable of remaining mindful of contextual and peripheral information in addition to its primary goal; (iii) *coordinated behavior*, where the system can coordinate multiple actions in synergistic

ways to best achieve multiple objectives; (iv) *adaptable behavior,* where the system is equipped to identify and adapt to variabilities that exist within and across different patients; and (v) *uncertainty-aware behavior,* where the system can handle imperfect and/or low-information measurements, quantify uncertainties that arise as a result, and incorporate them into its decisions. This combination of domain-specific challenges limits the direct applicability of most established decision-making frameworks to autonomous medical care problems, making it necessary to devise more specialized methods and frameworks tailored to such applications.

In this work, we present an *inference-based* physiological decision-support and closed-loop control (PCLC) framework intended to facilitate a solution toward the challenges summarized above. At its core, this framework is built to leverage the information encoded in a *generative physiological model* (as in Chapter 2; representing physiological behaviors that are likely to occur in the population and the patterned randomness that is typical of such behavior) to form perceptions and make decisions in the context of autonomous medical care. To address the perception problem, we use an inference-based filtering approach introduced in Chapter 3, which forms beliefs about the physiological states and characteristics of a patient by fusing the generative model's behaviors with real-time data. To address the decision problem, in this chapter, we formulate an inference-based control procedure that uses the generative model along with filter-provided beliefs, action definitions, and reward definitions to form beliefs about future actions that are likely to result in desirable outcomes. To demonstrate the merits and limitations of the proposed PCLC framework, we apply it to an in-silico case study on autonomous fluid resuscitation for hemodynamic

management. Based on this case study, we discuss the suitability of the proposed PCLC framework for application in decision-making problems for autonomous medical care, and the potential for this framework to serve as a basis for designing next-generation medical decision-support and closed-loop control systems.

4.2. Inference-Based Physiological Closed-Loop Control

This section presents an inference-based approach intended to facilitate the process of PCLC design with special emphasis on the motivations presented in Section 4.1. For this purpose, first, we present a general-purpose framework for reasoning about PCLC problems and their relevant considerations. Based on this framework, we demonstrate that the PCLC problem can be expressed as a Bayesian inference problem over a corresponding probabilistic graphical model, and provide an algorithmic solution to the problem using a principled combination of Sequential Monte-Carlo (SMC) and Variational Inference (VI) techniques. Finally, we show how the resulting algorithm may be leveraged to create an inference-based PCLC system. Further details follow.

4.2.1. A General-Purpose Framework for PCLC

As a framework for reasoning about PCLC problems, we consider the system shown in Figure 4.1. This system consists of a main decision/control loop, where, in each cycle, a patient's physiological measurements are processed by a *filtering algorithm* to form real-time beliefs about the patient's state and characteristics. These beliefs are then provided to a *control algorithm* tasked with finding and applying the best course of action for the patient. The overall loop is informed by additional contextual

information, including a *generative physiological model*, and definitions pertaining to available *actions* (e.g., specific therapies) and *rewards* (e.g., desirable therapeutic outcomes). In the following paragraphs, we present a more detailed overview of these components and their role in creating a PCLC system toward the objectives of this work.

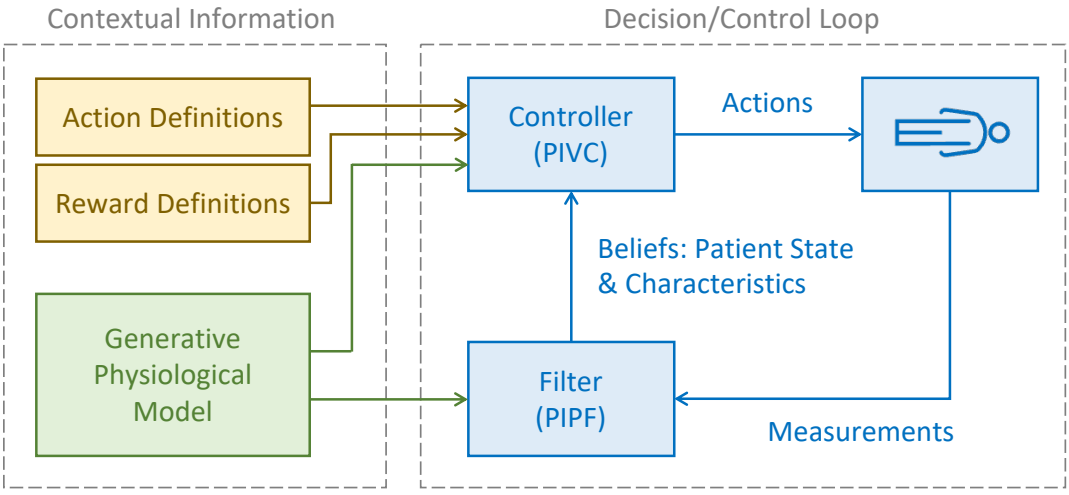

Figure 4.1. A general-purpose framework for physiological closed-loop control: A filter receives a stream of measurements from the patient and forms beliefs about the patient's state and characteristics. The controller plans an optimal sequence of actions based on these beliefs. The filter-controller pair is informed by a generative physiological model, along with contextual information pertaining to available actions and desirable outcomes (i.e., rewards).

The Generative Physiological Model: The generative physiological model is intended to serve as a foundational source of physiological knowledge in the PCLC system. In this type of model, the patterns and variabilities of physiological behavior, both within a patient and across the patient population, are represented by model components that behave in stochastic (but patterned) ways. For this purpose, we consider a generative modeling scheme similar to the one proposed in Chapter 2 and Chapter 3, which includes three main sub-models summarized as follows:

$$\theta_k \sim \mathcal{G}(\theta_{k-1}, \phi) \tag{4.1}$$

$$x_k \sim \mathcal{H}(x_{k-1}, \theta_{k-1}, u_{k-1}) \tag{4.2}$$

$$y_k \sim \mathcal{M}(x_k, \theta_k, n) \tag{4.3}$$

In this model, $\mathcal{G}$ is a patient generator model that instantiates "virtual patients" by generating vectors of patient characteristics and subsequently produces variations in those characteristics; $\mathcal{H}$ is a physiological dynamics model that generates evolutions in the states of each virtual patient; and $\mathcal{M}$ is a physiological measurement model that represents measurement processes along with their imperfections. In this formulation, k denotes the time index, θ_k is the vector of patient characteristics, x_k is the vector of patient states, u_k is the known inputs/therapies given to the patient, and y_k is the vector of generated measurements. Furthermore, ϕ denotes the tunable parameters of $\mathcal{G}$, and n denotes the tunable parameters of $\mathcal{M}$. Given physiological data from a cohort of patients, the unknown parameters of the generative model can be inferred using the method presented in [67], resulting in a model that can generate virtual datasets with similar distribution to real data. In this work, we are primarily interested in utilizing this encoded physiological knowledge to inform the PCLC system's filtering and control algorithms.

The Filter: In the context of PCLC, the filtering algorithm is tasked with forming and updating beliefs about the unmeasured states and characteristics of the patient using physiological measurements that arrive over time. For this purpose, we consider the filtering algorithm proposed in Chapter 3 named PIPF, which is a filter that is informed by the generative physiological model in (4.1)-(4.3). This filter

operates based on a Bayesian view of filtering, where beliefs are represented by probability densities, and the goal of filtering is to obtain the following posterior density at each time k:

$$p(x_{1:k}, \theta_{1:k} | y_{1:k}, \phi, n) = \frac{p(x_{1:k}, \theta_{1:k}, y_{1:k} | \phi, n)}{p(y_{1:k} | \phi, n)} \tag{4.4}$$

This density represents beliefs about the unmeasured states and characteristics of a patient up to time k (denoted by $x_{1:k}$ and $\theta_{1:k}$ respectively), given the patient's physiological measurements up to time k (denoted by $y_{1:k}$) and the information encoded in the parameters of the generative model ϕ, n. To compute this density in practice, we adopt an SMC (i.e., particle-based) approach, where beliefs are expressed as a (large) collection of weighted sample paths $\{w_k^i, [x, \theta]_{1:k}^i\}$, resulting in the following alternative expression for the posterior density:

$$\hat{p}(x_{1:k}, \theta_{1:k} | y_{1:k}, \phi, n) = \sum_{i=1}^{N} w_k^i \delta[x, \theta]_{1:k}^i \tag{4.5}$$

where each sample path $[x, \theta]_{1:k}^i$ describes a proposed potential path for the patient's states and characteristics, and the weight w_k^i represents how (relatively) likely the path is according to the generative physiological model and the patient's physiological measurements up to time k. In the PIPF algorithm, these sampling-based beliefs at time k can be computed recursively from their counterpart at time $k-1$, giving a recursive filtering algorithm that can be utilized in the PCLC system to form real-time beliefs about the patient's states and characteristics.

The Controller: In the proposed PCLC system, at each time k, the control algorithm receives (i) beliefs about the patient's states and characteristics (as generated

by the filtering algorithm); (ii) information about likely physiological behaviors (as encoded in the generative physiological model); and (iii) definitions pertaining to available actions (e.g., therapies) and rewards (e.g., desirable therapeutic outcomes). Given this information, the control algorithm is tasked with finding and applying a sequence of actions that would maximally benefit the patient. In this work, we present an inference-based approach to this planning/control problem, which relies on a look-ahead procedure to infer future actions that could result in maximal rewards. This approach is described in the following paragraphs.

To reason about actions in this control problem, we define an *action vector* of the following form:

$$a_t = [\, a_1 \;\; a_2 \;\; a_3 \;\; \ldots \,]_t \tag{4.6}$$

which is a vector of random variables representing actions (i.e., control inputs) that are available to the controller, with the index t denoting *look-ahead time*, which counts the number of time-steps into the future starting from the present time k. The vector a_t therefore refers to the set of actions taken at t steps into the future. To reason about physiological behaviors and rewards in the control problem, we consider a *look-ahead model* of the following form (see Figure 4.2(a)):

$$\theta_t \sim \mathcal{G}(\theta_{t-1}, \phi) \tag{4.7}$$

$$x_t \sim \mathcal{H}(x_{t-1}, \theta_{t-1}, a_{t-1}) \tag{4.8}$$

$$r_t \sim \mathcal{R}(x_t, \theta_t) \tag{4.9}$$

where $\mathcal{G}$ is the look-ahead patient generator model whose role is to generate variations in patient characteristics into the future, $\mathcal{H}$ is the look-ahead physiological dynamics

model whose role is to generate state evolutions in response to a given sequence of future actions, and $\mathcal{R}$ is the look-ahead rewards model whose role is to evaluate the look-ahead responses and return a vector of reward values r_t measuring their desirability.

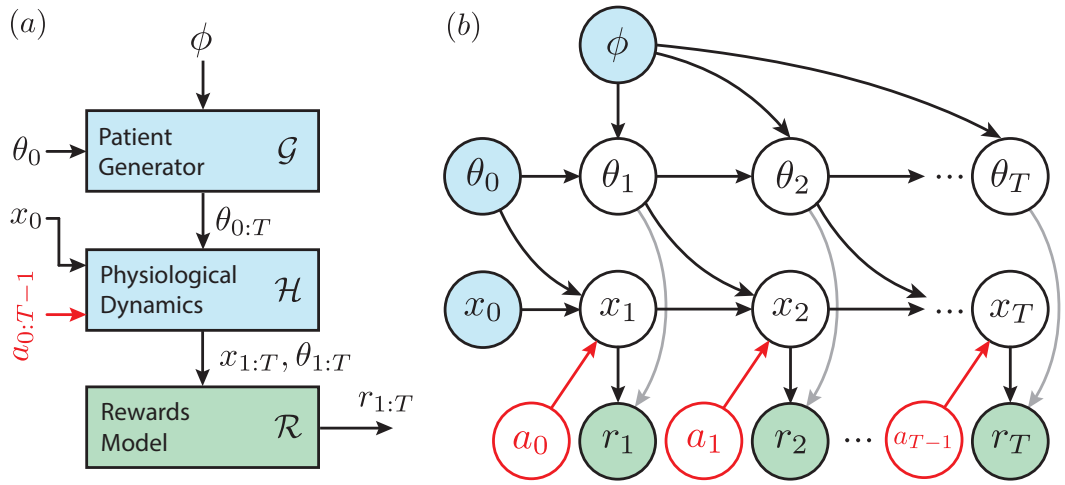

Figure 4.2. Schematic representation of the relationship between the variables of interest during look-ahead in the population-informed control problem: (a) block diagram of the generative physiological model used for look-ahead; (b) probabilistic graphical representation of the hierarchical and sequential relationship among the look-ahead variables.

Given the definitions above, the joint relationship between the variables of interest in the control problem during look-ahead is shown in Figure 4.2(b), which implies the following joint density:

$$p(a_{0:T-1}, x_{0:T}, \theta_{0:T}, r_{1:T}|\phi) = p(x_0, \theta_0) \prod_{t=1}^{T} [p(r_t|x_t, \theta_t) p(x_t|x_{t-1}, \theta_{t-1}, a_{t-1}) p(\theta_t|\theta_{t-1}, \phi) p(a_{t-1})] \quad (4.10)$$

This joint relationship has a hierarchical and sequential structure. In this structure, $p(x_0, \theta_0)$ denotes the (marginalized) beliefs about the patient's current states and characteristics (as provided by the filtering algorithm), T denotes the look-ahead time

horizon, $a_{0:T-1}$ is the look-ahead action sequence, $\theta_{0:T}$ is the look-ahead path for the patient's characteristics (which is affected by the parameters of the patient generator model ϕ), $x_{0:T}$ is the look-ahead path for the patient's state evolutions (which is affected by the action sequence as well as the patient's characteristics), and $r_{1:T}$ denotes look-ahead rewards.

Given the joint density in (4.10), the ultimate objective of control in this work can be expressed in more concrete terms using the following posterior density:

$$p(a_{0:T-1}|r_{1:T},\phi) = \frac{p(a_{0:T-1},r_{1:T}|\phi)}{p(r_{1:T}|\phi)} \tag{4.11}$$

which represents beliefs about the look-ahead action sequence $a_{0:T-1}$ conditioned on the attainment of future rewards $r_{1:T}$ and the information encoded in the parameters of the generative physiological model ϕ. As shown on the right-hand-side of this equation, it is conceptually possible to compute the posterior density from the joint density introduced in (4.10) through the following integration procedures:

$$p(a_{0:T-1},r_{1:T}|\phi) = \int_{x_{0:T},\theta_{0:T}} p(a_{0:T-1},x_{0:T},\theta_{0:T},r_{1:T}|\phi) \tag{4.12}$$

$$p(r_{1:T}|\phi) = \int_{a_{0:T-1}} p(a_{0:T-1},r_{1:T}|\phi) \tag{4.13}$$

Overall, obtaining the posterior density in (4.11) is equivalent to obtaining all future action sequences that are likely to result in desirable outcomes, which is the ultimate goal of PCLC in this work. However, in most practical problems, the procedures in (4.11)-(4.13) are expected to be analytically intractable due to complexities in the joint distribution and nonlinearities in the underlying model. In order to address this

challenge, in the next section, we derive a practical solution to this problem based on variational inference techniques.

4.2.2. A Variational Solution

This section presents a variational solution to the control problem in (4.11). Variational methods operate based on searching in a pre-defined family of distributions in order to find the ones that best approximate the desired posterior density [80], [81], [83], [153]. For this purpose, we first define the following family of *approximate posterior densities* for the control problem:

$$q_A(a_{0:T-1}|\nu) = \prod_{t=1}^{T}[q_A(a_{t-1}|\nu)] \tag{4.14}$$

which is a parametric density function defined over the space of action sequences $a_{0:T-1}$, the position/shape of which can be modulated through ν. The role of this density is to approximate the posterior density in (4.11), thus holding a solution to the control problem. In addition to this density, we define the following *integration density*:

$$\begin{gathered} p_I(x_{0:T}, \theta_{0:T}|a_{0:T-1}, \phi) = \\ p(x_0, \theta_0) \prod_{t=1}^{T}[p(x_t|x_{t-1}, \theta_{t-1}, a_{t-1}) p(\theta_t|\theta_{t-1}, \phi)] \end{gathered} \tag{4.15}$$

which represents the joint density between the look-ahead states and characteristics of the patient, given a sequence of actions and the information encoded in the generative model. Sample paths from the integration density can be drawn by giving an action sequence to the look-ahead generative model described in (4.7)-(4.8). The role of this

density is to assist in performing the integrations necessary to obtain a variational solution to the control problem in (4.11).

To derive a procedure for finding the best approximate posterior (through modulating ν in (4.14)), we start from the logarithm of the integral in (4.13) as follows:

$$\log p(r_{1:T}|\phi) = \log \int_{a_{0:T-1}} \int_{x_{0:T},\theta_{0:T}} p(a_{0:T-1}, x_{0:T}, \theta_{0:T}, r_{1:T}|\phi) \tag{4.16}$$

Multiplying and dividing the integrand in (4.16) by q_A and p_I yields:

$$\log p(r_{1:T}|\phi) = \log \int_{a_{0:T-1}} \int_{x_{0:T},\theta_{0:T}} p(a_{0:T-1}, x_{0:T}, \theta_{0:T}, r_{1:T}|\phi) \frac{q_A(a_{0:T-1}|\nu) p_I(x_{0:T},\theta_{0:T}|a_{0:T-1},\phi)}{q_A(a_{0:T-1}|\nu) p_I(x_{0:T},\theta_{0:T}|a_{0:T-1},\phi)} \tag{4.17}$$

After substituting (4.10), (4.14), and (4.15) into (4.17) and cancelling/rearranging the terms we obtain:

$$\log p(r_{1:T}|\phi) = \log \mathbb{E}_{a_{0:T-1}\sim q_A} \mathbb{E}_{x_{0:T},\theta_{0:T}\sim p_I} \left[\frac{\prod_{t=1}^{T} [p(r_t|x_t,\theta_t) p(a_{t-1})]}{\prod_{t=1}^{T} [q_A(a_{t-1}|\nu)]} \right] \tag{4.18}$$

Finally, applying Jensen's inequality to (4.18) yields:

$$\begin{aligned} \log p(r_{1:T}|\phi) &\geq L(\nu) \\ &= \mathbb{E}_{a_{0:T-1}\sim q_A} \mathbb{E}_{x_{0:T},\theta_{0:T}\sim p_I} \sum_{t=1}^{T} [\log p(r_t|x_t,\theta_t) + \log p(a_{t-1}) - \\ &\qquad \log q_A(a_{t-1}|\nu)] \end{aligned} \tag{4.19}$$

which constitutes the *variational objective*, to be maximized over ν, for the purpose of solving the control problem considered in this work. This objective function is based on two expectation operations, one with respect to $a_{0:T-1}\sim q_A$, which samples action sequences from the approximate posterior in (4.14), and another with respect to

$x_{0:T}, \theta_{0:T} \sim p_I$, which samples paths for the subject's states and characteristics from the integration density in (4.15) using the generative model described in (4.7)-(4.8). Inside the expectation operator, there is a sum over the look-ahead time horizon, containing the following terms: (i) the term $\log p(r_t|x_t, \theta_t)$ is the *rewards likelihood*, which promotes the desirability of the patient's states and characteristics at look-ahead time t; (ii) the term $\log p(a_{t-1})$ is the *action prior*, which promotes prior preferences and constraints on actions; and (iii) the term $-\log q_A(a_{t-1}|\nu)$ promotes a diffuse approximate posterior, acting as a means for uncertainty quantification around actions.

4.2.3. The Population-Informed Variational Control (PIVC) Algorithm

As presented in Section 4.2.2, a solution to the control problem considered in this work (seeking action sequences that are likely to result in desirable outcomes for the patient) can be obtained by maximizing the variational objective in (4.19) over the variational parameters ν. Due to the presence of expectation operators in (4.19), this task could become prohibitively expensive if the objective is evaluated in each iteration of an optimization algorithm. To address this challenge, we present an optimization procedure that operates based on stochastic gradients of the variational objective [82], [160], thereby enabling a practical way to maximize (4.19). For this purpose, we consider the sampling process within the expectation operator in (4.19). This process can be expressed conceptually as a transformation:

$$[a_{0:T-1}, x_{0:T}, \theta_{0:T}] = F(\epsilon, \nu), \quad \epsilon \sim \mathcal{S} \tag{4.20}$$

where all sources of randomness in sampling actions from (4.14), sampling beliefs from (5), and sampling behavior paths from (4.7)-(4.9), and (4.15) have been extracted and

denoted by ϵ. In other words, the transformation F receives a vector of random numbers (of appropriate dimension) along with the vector of variational parameters (which are to be optimized), and produces, deterministically, sample paths for action and behavior. As a result, the gradient of the variational objective in (4.19) can be expressed as:

$$\nabla_v L(v) = \mathbb{E}_{\epsilon\sim\mathcal{S}} \nabla_v \sum_{t=1}^{T} [\log p(r_t|x_t, \theta_t) + \log p(a_{t-1}) - \log q_A(a_{t-1}|v)] \quad (4.21)$$

which indicates that an unbiased but noisy version of the gradient can be obtained efficiently by drawing $\epsilon \sim \mathcal{S}$ and calculating the gradient inside the expectation based on the drawn ϵ. This noisy gradient can in turn be used with a stochastic optimization algorithm [98], [99], [101] to optimize the variational parameters v. This optimization procedure is outlined in Algorithm 4.1, while the procedure for evaluating the stochastic objective in each iteration is outlined in Algorithm 4.2. Recalling the PCLC framework in Figure 4.1, these procedures can be leveraged directly to realize the controller block, whose role is to recommend/apply optimal courses of action for the patient based on the defined therapeutic goals.

Algorithm 4.1. Stochastic Optimization of the Variational Objective for PIVC

Inputs: $\{w_k^i, [x, \theta]_k^i\}$ {current beliefs about the patient in particle form; see (4.5)}

Inputs: $v_0, \beta_1, \beta_2, \alpha, \delta$ {optimization parameters}

$v \leftarrow v_0, m_v \leftarrow 0, v_v \leftarrow 0$ {initialize variables}

For l from 1 to l_{max}:

$\epsilon_a, \epsilon_{\mathcal{G}}, \epsilon_{\mathcal{H}}, \epsilon_{\mathcal{R}} \sim \mathcal{S}$ {draw and store random samples}

$x_0, \theta_0 \sim \{w_k^i, [x, \theta]_k^i\}$ {draw and store sample from current beliefs}

$g_\nu \leftarrow \nabla_\nu \tilde{L}(\epsilon_a, \epsilon_{\mathcal{G}}, \epsilon_{\mathcal{H}}, \epsilon_{\mathcal{R}}, x_0, \theta_0, \nu)$ {obtain (stochastic) gradient}

$m_\nu \leftarrow [\beta_1 m_\nu + (1-\beta_1) g_\nu]/(1-\beta_1^l)$ {estimate first moment}

$v_\nu \leftarrow [\beta_2 v_\nu + (1-\beta_2) g_\nu^2]/(1-\beta_2^l)$ {estimate second moment}

$\nu \leftarrow \nu + \alpha[m_\nu/(\sqrt{v_\nu} + \delta)]$ {update variational parameters; see [99]}

End For.

Return: ν {optimized variational parameters}

Algorithm 4.2. Evaluating the Stochastic Objective for PIVC

Inputs: $\epsilon_a, \epsilon_{\mathcal{G}}, \epsilon_{\mathcal{H}}, \epsilon_{\mathcal{R}}, x_0, \theta_0, \nu$ {stochastic samples; parameters}

$p_r \leftarrow 0, p_a \leftarrow 0$ {initialize objective terms}

$a_{0:T-1} \leftarrow F_a(\nu, \epsilon_a)$ {sample action sequence; F_a is a sample generator for (4.14)}

For t from 1 to T:

$\theta_t \leftarrow F_{\mathcal{G}}(\theta_{t-1}, \phi, \epsilon_{\mathcal{G}})$ {sample characteristics using (4.7)}

$x_t \leftarrow F_{\mathcal{H}}(x_{t-1}, \theta_{t-1}, a_{t-1}, \epsilon_{\mathcal{H}})$ {sample states using (4.8)}

$r_t \leftarrow F_{\mathcal{R}}(x_t, \theta_t, \epsilon_{\mathcal{R}})$ {sample reward using (4.9)}

$p_r \leftarrow p_r + \log p(r_t | x_t, \theta_t)$ {reward terms for the objective}

$p_a \leftarrow p_a + \log p(a_{t-1}) + \log q_A(a_{t-1} | \nu)$ {action terms for the objective}

End For.

Return: $\tilde{L} = p_a + p_r$ {return stochastic sample from the objective; see (4.19)}

4.3. Application to Autonomous Hemodynamic Management

Hemodynamic management is an essential aspect of medical care, especially for critical patients, where the goal is to monitor the performance of a patient's cardiovascular system and administer therapies to ensure adequate blood circulation and tissue oxygenation for the patient [141], [142]. To demonstrate the merits of the proposed PCLC framework, we follow this framework to build a PCLC algorithm for a typical hemodynamic management scenario where hypovolemia (i.e., low circulating blood volume) is treated with colloid and/or blood infusions [143]. In this scenario, the control algorithm is tasked with (i) receiving a stream of blood pressure data from the patient (along with information about infusions and possible hemorrhage), (ii) forming a real-time perception of the patient's condition, and (iii) applying an optimized course of treatment (i.e., fluid and/or blood infusions) to maximally benefit the patient. In this section, we present the details of this PCLC algorithm and describe the in-silico case studies used to demonstrate its effectiveness.

4.3.1. Generative Model, Filter, and Controller

As presented in Section 4.2.1, building a PCLC algorithm based on the proposed framework involves, as its basis, a generative physiological model. For the hypovolemia treatment case study, we use a generative model structure from our previous work [67], which was designed to represent and reproduce the bodily responses to hemorrhage and fluid resuscitation in a population of subjects. This model thus includes components that can instantiate cohorts of "virtual subjects", generate

paths for subject characteristics, produce state evolutions in response to given stimuli, and generate realistic physiological measurements pertaining to hypovolemia treatment. The information encoded in this model is intended to inform the perceptions and decisions of the PCLC algorithm (as presented in Section 4.2). Given this generative model, the next steps are to determine the filtering and control algorithms. For the filtering algorithm, we utilize an algorithm proposed in Chapter 3 and summarized in Section 4.2.1, named the population-informed particle filter (PIPF), which is a filter that is built to be informed by a generative physiological model. For the control algorithm, we apply the PIVC algorithm introduced in Section 4.2.2 and Section 4.2.3. The details of this application are presented over the next few sections.

4.3.2. Defining Actions

As presented in Section 4.2.3, the PIVC algorithm operates based on action definitions that represent the types of actions that would be available to the controller. To formalize these actions, we consider a Gaussian formulation for the approximate posterior in (4.14) as follows:

$$q_A(a_{0:T-1}|\nu) = \prod_{t=1}^{T}\left[\mathcal{N}(a_{t-1}|\nu_{t-1}^{\mu}, \nu_{t-1}^{\sigma})\right] \tag{4.22}$$

which indicates that the approximate posterior at each look-ahead timestep $t-1$ is a Gaussian distribution $\mathcal{N}(a_{t-1}|\nu_{t-1}^{\mu}, \nu_{t-1}^{\sigma})$ defined over the action vector a_{t-1} and parameterized by the distribution's center ν_{t-1}^{μ} and standard deviation ν_{t-1}^{σ}. In case n_a types of actions are available, the action vector a_{t-1} would have n_a dimensions. For instance, in the hypovolemia treatment case study, where the available types of action

include colloid and blood infusions, the action vector is of the form $a_{t-1} = \left[a_{t-1,C} \;\; a_{t-1,B}\right]^T$. In this vector, $a_{t-1,C}$ denotes the fluid infusion rate and $a_{t-1,B}$ is the blood infusion rate. Given this formulation, the variational parameters for PIVC can be summarized as $\nu = \{\nu^{\mu}_{0:T-1}, \nu^{\sigma}_{0:T-1}\}$. In this sense, $\nu^{\mu}_{0:T-1}$ represents the sequence of actions that are most likely to result in maximum rewards over the time horizon T, while $\nu^{\sigma}_{0:T-1}$ quantifies the uncertainty around those actions.

The PIVC algorithm also needs a sample generator function $F_a(\nu, \epsilon_a)$ to sample action sequences from the approximate posterior (see Algorithm 4.2). For the Gaussian formulation in (4.22), this sample generator can be written as follows:

$$a_{0:T-1} = \nu^{\mu}_{0:T-1} + \nu^{\sigma}_{0:T-1} \odot \epsilon_a, \quad \epsilon_{a,ij} \sim \mathcal{N}(0,1) \tag{4.23}$$

where $a_{0:T-1}$ denotes a sampled action sequence, ϵ_a is a $n_a \times T$ matrix of samples from the standard Gaussian distribution, i denotes the row index pointing to action type, j is the column index pointing to look-ahead time, $\nu^{\mu}_{0:T-1}$, $\nu^{\sigma}_{0:T-1}$ are $n_a \times T$ matrices containing the means and standard deviations of the action posterior over the time horizon, and $\odot$ denotes element-wise multiplication. This sample generator function can be used in Algorithm 4.2 to generate action samples in the PIVC algorithm when necessary. Given this sample generator, the approximate posterior term in (4.21) can be computed as follows:

$$\log q_A(a_{t-1}|\nu) = \sum_i \left[-\tfrac{1}{2}\epsilon_{a,i} - \log(\nu^{\sigma}_{t-1,i}) - \tfrac{1}{2}\log(2\pi)\right] \tag{4.24}$$

The variational control objective in (4.21) also includes an action prior term, which can be leveraged to incorporate additional constraints and/or prior knowledge about actions

into the control problem. In the hypovolemia treatment case study, we consider prior limits on maximum infusion rate (due to rate limitations in infusion pumps) and blood infusions (due to the preference to preserve blood resources if possible). These action priors can be formulated as follows:

$$\log p(a_{t-1}) = D\sum_i\left[\min\left(a_{t-1,i}, 0\right) - \max\left(a_{t-1,i} - a_{max}, 0\right)\right] - D\left|a_{t-1,B}\right| + \mathcal{C} \tag{4.25}$$

where the first term discourages out-of-range infusion rates for colloid and blood, the second term discourages unnecessary blood infusions, D is a (large) constant, and $\mathcal{C}$ is a constant containing the normalizing terms for the action prior distribution and does not affect the variational objective.

4.3.3. Defining Rewards

As presented in Section 4.2.3, the PIVC algorithm operates based on reward definitions that represent desirable patient states and/or outcomes and determine the objectives of the control problem. In the hypovolemia treatment example, the overarching objective is to maintain adequate blood circulation and ensure that the organs and tissues receive adequate perfusion and oxygenation. To make this objective more concrete, we consider two main reward definitions for the PCLC framework, which are described below.

As a means to promote adequate blood circulation, the first reward definition promotes mean arterial pressures (MAP) that are in the vicinity of a desirable target blood pressure, which is formalized using the following reward density function:

$$\log p\left(r_{t,MAP}\middle|x_t,\theta_t\right) = \\ -\frac{1}{2\sigma_{r,MAP}^2}\left(y_{t,MAP}-\mu_{r,MAP}\right)^2 - \log\left(\sigma_{r,MAP}\right) - \frac{1}{2}\log(2\pi) \tag{4.26}$$

where $r_{t,MAP}$ denotes the MAP reward, $y_{t,MAP}$ denotes the MAP sample generated by the PIVC algorithm, $\mu_{r,MAP}$ is the mean of the reward density, and $\sigma_{r,MAP}$ is the standard deviation of the reward density. The mean $\mu_{r,MAP}$ determines the MAP with maximum reward density while the standard deviation $\sigma_{r,MAP}$ determines the spread rewards around the mean.

As a means to prevent blood over-dilution (which is likely to limit organ and tissue oxygenation, endangering the subject), the second reward definition promotes blood hematocrit (HCT) values that are above a minimum threshold, which is formalized as follows:

$$\log p\left(r_{t,HCT}\middle|x_t,\theta_t\right) = D[\min(y_{t,HCT}-h_{min},0)] + \mathcal{C} \tag{4.27}$$

where $r_{t,HCT}$ denotes the HCT reward, $y_{t,HCT}$ is the HCT sample generated by the PIVC algorithm using the generative model, h_{min} is the minimum threshold for HCT, D is a (large) constant, and $\mathcal{C}$ is a constant containing the normalizing terms for the reward density. Given the definitions above, the PIVC algorithm described in Algorithm 4.1 and Algorithm 4.2 can be leveraged to search for optimal sequences of action that result in maximal rewards.

4.3.4. The Physiological Data

As presented in Section 4.3.1, the PCLC framework works based on a generative physiological model. To infer the latent parameters of this generative model, we utilize a physiological dataset from previous work [48], [144], [145], which contains time-series measurements pertaining to hypovolemia treatment with colloids in 5 animal (sheep) experiments. In each experiment, the animal is subjected to controlled hemorrhage and subsequently resuscitated over time using colloid infusions, which are administered according to the recommendations of a control algorithm. Each experiment spans a course of 180 minutes, where the subject's hematocrit, cardiac output, and mean arterial pressure are measured at a target period of ~5 minutes. Given this dataset, we leverage the inference algorithm introduced in Chapter 2 to infer the latent parameters of the generative model, resulting in a model that can be used as a basis for the case studies in this work. These case studies are designed to showcase the merits and limitations of the proposed PCLC framework, and are described in the following section.

4.3.5. Problem Setting and Case Studies

To discuss the merits and limitations of the proposed PIVC controller, and more generally, the proposed PCLC framework, we consider the in-silico application of this framework in three scenarios pertaining to hypovolemia treatment. In Scenarios A, B and C, the objective of control is to maintain the subject's MAP according to a Gaussian reward definition (centered at 90 mmHg with standard deviation of 10 mmHg; shown using green lines in Figure 4.3). In Scenarios B and C, the objective of control also

includes avoiding the over-dilution of the subject's blood below an HCT of 10% (shown using orange lines in Figure 4.3). In Scenario B, the controller's actions are limited to colloid infusions with a maximum rate of 0.04 L/min (see Figure 4.4), while in Scenario C, the controller can infuse both colloids and blood, albeit with preference toward colloid infusions. During its operation, the PCLC system receives a stream of MAP measurements from an in-silico subject (along with information about infusion rates and hemorrhage) and forms beliefs about the subject's states and characteristics (e.g., see posterior beliefs about HCT and CO in Figure 4.3, where bolder colors denote higher belief density). Based on these beliefs, the PCLC system performs a look-ahead procedure (see Section 4.2.2) to obtain action sequences to apply to the subject. It is important to note that this problem setting is a purposefully challenging one, where many unmeasured variables are inferred from few measured variables and used as basis to make therapeutic decisions. This setting enables us to assess the merits and limitations of the PCLC framework especially when the clinically available data is limited, and the beliefs have high entropy (pointing to high uncertainty about the patient's states and characteristics).

4.4. Results and Discussion

Developing decision-support and closed-loop control algorithms for the purpose of autonomous medical care typically involves numerous considerations and challenges related to transparency, context-awareness, coordination, adaptation, uncertainty, and the appropriate handling of low-information measurements. In Section 4.2, we proposed a framework that leverages concepts from probabilistic inference to facilitate

the incorporation of these considerations into PCLC algorithms. In this section, we present the results of applying this framework to build an algorithm for autonomous hypovolemia treatment, and use this example to discuss the capabilities, merits, and limitations of the PCLC framework in solving such problems. Further details follow.

4.4.1. Controller Behavior and Potential for Context-Awareness

Figure 4.3 and Figure 4.4 show representative examples of control behavior in two in-silico hypovolemia treatment scenarios. In Scenario A and Scenario B (see Section 4.3.5 for scenario descriptions), the controller successfully responds to the initial hemorrhage by infusing colloids at maximum rate until the hemorrhage stops and MAP recovers back to 90 mmHg. Then, the controller infuses smaller amounts of colloid to maintain MAP at 90 mmHg. In Scenario A, this behavior pattern continues in response to subsequent hemorrhages. However, in Scenario B, the controller, being aware of the possibility of over-diluting the subject's blood (see posterior beliefs about HCT in the middle-left panel of Figure 4.3), infuses smaller amounts of colloid while settling for MAP values that are moderately smaller than 90 mmHg, thereby limiting the possibility of over-dilution. In Scenario C, where the option to infuse blood exists, the controller similarly responds to the initial hemorrhage by infusing colloids to bring the subject's MAP back to 90 mmHg. Interestingly, in response to subsequent hemorrhages, the algorithm performs limited blood infusions, thereby mitigating the risk of over-dilution (see the bottom-left panel in Figure 4.3) and enabling further colloid infusions. Following this strategy, the controller is able to bring the subject's MAP back to 90 mmHg after the cessation of hemorrhage. These example scenarios suggest that the

proposed PCLC framework may be utilized to build a PCLC system that can form perceptions of a subject's condition from a stream of measurements and act on those perceptions to achieve pre-defined therapeutic goals.

The example results above also highlight a notable aspect of the proposed PCLC framework aimed at facilitating the realization of context-awareness in PCLC: A recurring challenge in the design and adoption of PCLC systems is the legitimate fear that computers, when put into uncertain and evolving situations, tend to make naïve decisions that are blind to the broader context. Ideally, a PCLC system designed for use in real-world scenarios should be equipped with comprehensive contextual knowledge of the task at hand and the ability to effectively use this knowledge to make informed therapeutic decisions. Toward this objective, the proposed PCLC framework is equipped with a generative physiological model to carry contextual information about physiological behaviors that are likely to occur, as well as action and reward definitions to carry contextual information about available therapies and therapeutic goals. This setup results in a control system that is, for example, aware of blood over-dilution and its undesirable nature, and takes steps to minimize (see Scenario B in Figure 4.3) or mitigate (see Scenario C in Figure 4.3) this risk, thereby avoiding a potentially dangerous outcome that could have occurred with a traditional context-blind approach to controlling the subject's MAP.

4.4.2. Handling of Low-Information Measurements and Uncertainty

A recurring challenge in the design and development of algorithms for physiological decision-support and closed-loop control is the existence of variabilities in

physiological behavior, both within individuals and across the patient population. This challenge is further complicated by the often intermittent and low-information nature of real-world physiological measurements, adding further uncertainty around a patient's condition. Ideally, a PCLC system should be equipped to handle such variabilities and uncertainties, and incorporate them into its decisions. The PCLC framework proposed in this work (i) utilizes a generative physiological model to represent and reproduce real-world variabilities in physiological behavior, and (ii) adopts a Bayesian formulation to reason about uncertainties that arise during the algorithm's operation. The advantages of this PCLC scheme become especially apparent when available measurements lack sufficient information about important aspects of a patient's physiology. For instance, in the hypovolemia treatment example shown in Figure 4.3 and Figure 4.4, where the available MAP measurements are not sufficiently informative to result in accurate estimation of the subject's HCT and CO, the algorithm leverages the behaviors and variabilities that are generated by the generative physiological model and systematically combines them with measurements (i.e., MAP, infusions, and hemorrhage) to form real-time beliefs (see blue lines in Figure 4.3) about HCT and CO. In this way, the algorithm produces a real-time collection of "potential explanations" for the observed data in the form of posterior probability densities, laying a foundation to reason about best courses of action in the face of uncertainty. As a result of this setup, in the hypovolemia treatment example, the algorithm is able to minimizes the risk of blood over-dilution (HCTs below 10%) while maintaining the subject's MAP around 90 mmHg. Overall, these results suggest that the proposed PCLC framework may be utilized to create control algorithms that

are aware of the presence of uncertainty and can use this knowledge of uncertainty to make informed therapeutic decisions.

4.4.3. Transparency of Perceptions and Decisions

Figure 4.5 and Figure 4.6 show a cross-section of the look-ahead procedure for control, taken at the 72-minute mark in Scenario C (shown using a dotted line in Figure 4.3 and Figure 4.4), where Figure 4.6 shows the controller's action plan (i.e., colloid and blood infusions) while Figure 4.5 shows the plan for the subject's internal states and physiological variables (e.g., HCT, CO, and MAP). At this particular cross section in time, the algorithm believes that the subject's blood could be on the brink of over-dilution (see dotted line in bottom-left panel of Figure 4.3) and that the subject's MAP is below its most desirable value of 90 mmHg (see dotted line in bottom-right panel of Figure 4.3). In this scenario, infusing more colloid is expected to increase the subject's circulating blood volume and MAP (which is desirable) while creating a risk of blood over-dilution. On the other hand, infusing blood is expected to increase HCT and MAP (both of which are desirable), albeit at the expense of tapping into valuable blood resources. In this scenario, the algorithm's proposed action plan (shown in Figure 4.6) is to first infuse a moderate amount of blood to the subject, which increases MAP and prevents potential blood over-dilution (see bottom-right and bottom-left panels in Figure 4.5). Subsequently, the algorithm plans to leverage the newfound HCT margin to infuse further colloid and increase the subject's MAP without significant risk of over-dilution.

The example above highlights the potential for the proposed PCLC framework to facilitate increased transparency around the perceptions and actions of the controller. On the perceptions front, the filtering algorithm in the PCLC framework is built to operate based on a generative model that encodes a joint relationship between the variables of interest in the PCLC problem (e.g., HCT, CO, and MAP), thereby allowing for a human observer to inspect the algorithm's perception of these variables in real-time (as in Figure 4.3). Furthermore, in case the generative model is built with a physical/mechanistic structure, there is additional opportunity for transparency through inspecting the states of the mechanistic model (e.g., the top two rows in Figure 4.5). On the actions front, the control algorithm in the PCLC framework is built to leverage the same generative model to reason about future actions and the desirability of their effects, which in turn allows for a human observer to inspect the controller's action plans (e.g., Figure 4.6) as well as its rationale for planning those actions (e.g., Figure 4.5).

4.4.4. Coordinating Multiple Actions and Goals

An important real-world challenge in the context of autonomous medical care is the observation that patient care is rarely an isolated single-input single-output endeavor. Rather, providing adequate care to a patient often involves reasoning about a multitude of physiological variables and coordinating multiple possible actions to maximally benefit the patient. For instance, in the context of care for critically-ill patients, there may be a need for fluid infusions and/or blood transfusions to ensure adequate tissue perfusion and oxygenation for the patient, while simultaneously, there may be a need

for anesthesia and/or sedation through appropriate drug infusions. Ideally, an autonomous medical care system should be equipped to handle such multivariable scenarios and (even more importantly) leverage the multivariable nature of the problem to provide superior care to patients through judicious coordination of its actions. The proposed PCLC framework is highly amenable to such scenarios due to a structure that is, for the most part, agnostic to the number of variables, actions, and goals. For instance, in Scenario C of the hypovolemia treatment example (shown in bottom rows of Figure 4.3 and Figure 4.4; also Figure 4.5 and Figure 4.6) the decisions are made based on two physiological variables (HCT and MAP), two available actions (colloid infusions and blood infusions), and three main goals (maintaining MAP close to 90 mmHg, preventing blood over-dilution below 10% HCT, and avoiding infusion rates above 0.04 L/min). Arguably, it would be possible to involve even more variables, actions, and goals without significant changes to the PCLC setup. As a result of this feature, the fundamental step in PCLC design using this framework would be to construct a generative model of the relevant physiological phenomena, and infer its latent parameters based on physiological data. Given such a model, the proposed PCLC framework can be utilized to form perceptions of the patient's condition from multivariable measurements, and search for multivariable action sequences that best achieve the control problem's multiple goals.

4.5. Conclusions

In this chapter, we proposed an *inference-based* framework intended to facilitate the principled design and development of closed-loop control algorithms for autonomous

medical care. Using an in-silico case study on autonomous fluid resuscitation for hemodynamic management, we demonstrated that the control system arising from the proposed framework shows promise in exhibiting (i) *transparent behavior,* where the system's perceptions, reasoning, and decisions are human-interpretable; (ii) *context-aware behavior,* where the system is capable of remaining mindful of contextual and peripheral information in addition to its primary goal; (iii) *coordinated behavior*, where the system can coordinate multiple actions in synergistic ways to best achieve multiple objectives; (iv) *adaptable behavior,* where the system is equipped to identify and adapt to variabilities that exist within and across different patients; and (v) *uncertainty-aware behavior,* where the system can handle imperfect and/or low-information measurements, quantify uncertainties that arise as a result, and incorporate them into its decisions. Given the demonstrated promise of the proposed framework, future efforts should be devoted to (i) applying the framework to develop decision-support and closed-loop control algorithms for a wider range of autonomous medical care applications, and (ii) systematically evaluating the resulting algorithms in real-world in-vivo scenarios.

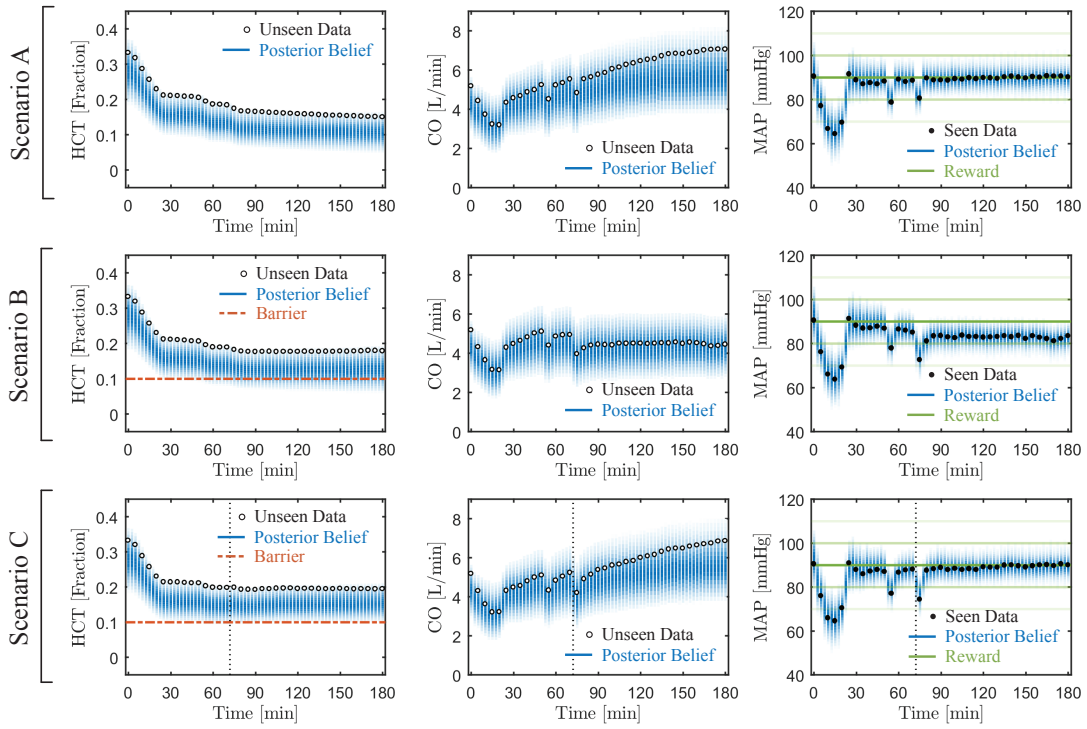

Figure 4.3. Representative closed-loop control behaviors for three in-silico scenarios pertaining to hypovolemia treatment: (A) the controller uses colloid infusions to maintain mean arterial pressure (MAP) around 90 mmHg; (B) the controller uses colloid infusions to maintain MAP around 90 mmHg while avoiding hematocrit (HCT) values lower than 0.1; (C) the controller uses both colloid and blood infusions (with colloids preferred) toward the same objective as B.

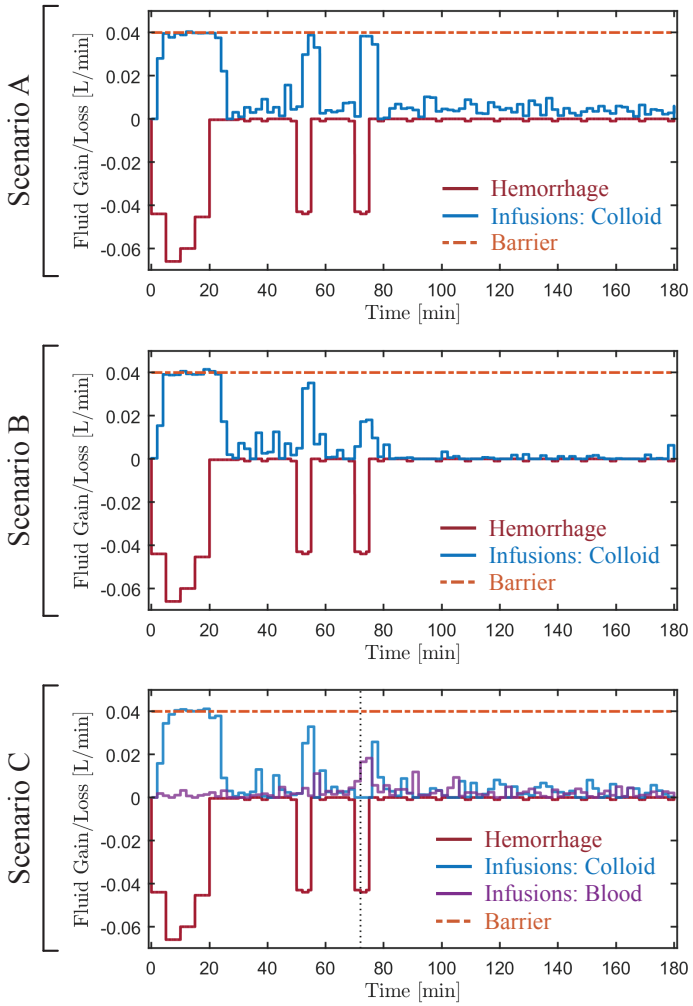

Figure 4.4. Representative control input behaviors for three in-silico scenarios pertaining to hypovolemia treatment: (A) the controller uses colloid infusions to maintain mean arterial pressure (MAP) around 90 mmHg; (B) the controller uses colloid infusions to maintain MAP around 90 mmHg while avoiding hematocrit (HCT) values lower than 0.1; (C) the controller uses both colloid and blood infusions (with colloids preferred) toward the same objective as B.

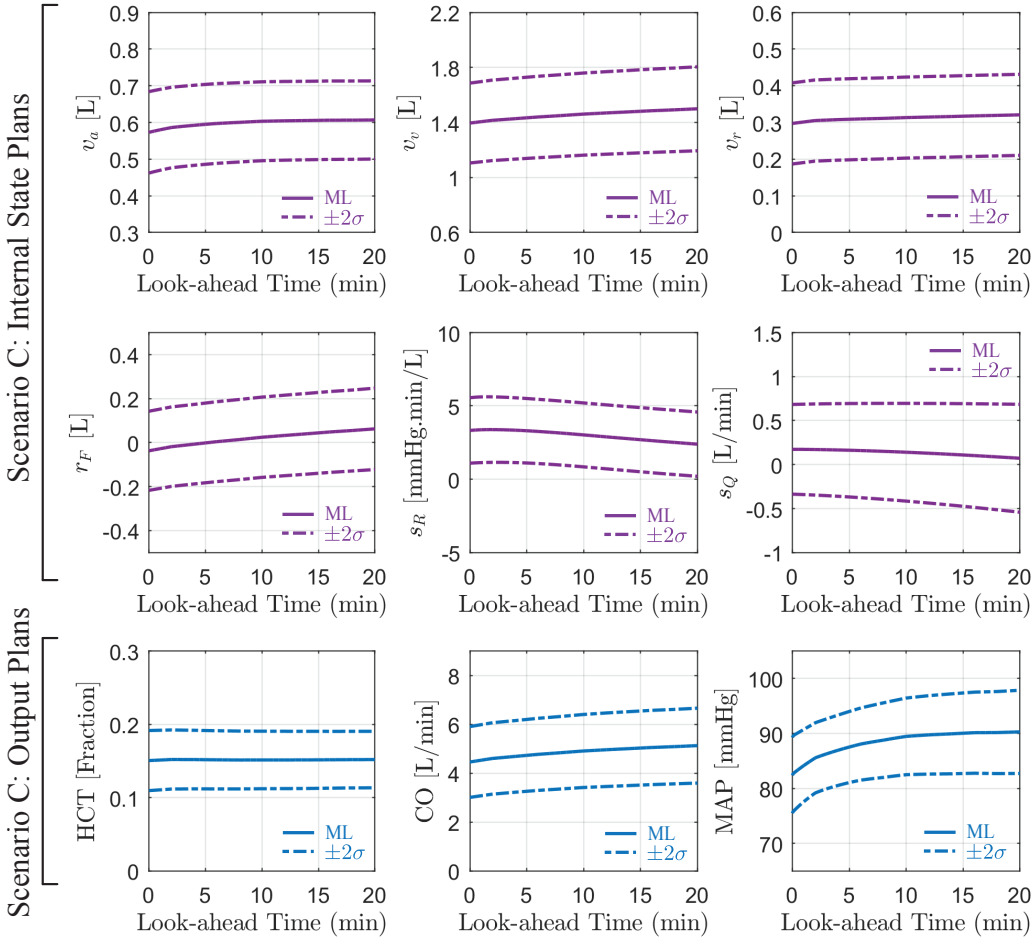

Figure 4.5. Representative look-ahead state and output plans (starting at time 72 min in Scenario C; see dotted lines in Figure 3 and Figure 4), where the controller plans to use both colloid and blood infusions (with colloids preferred) to regulate mean arterial pressure (MAP) around 90 mmHg, while avoiding hematocrit (HCT) values lower than 0.1.

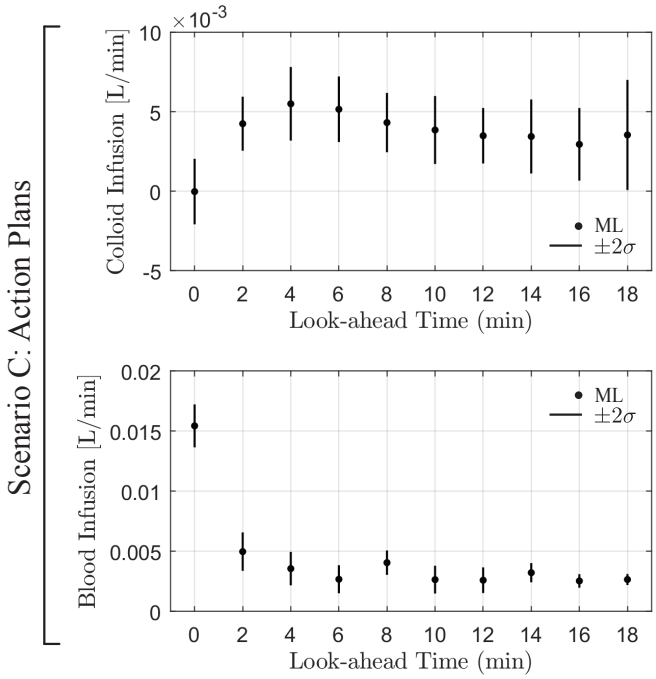

Figure 4.6. Representative look-ahead action plans (starting at time 72 min in Scenario C; see dotted lines in Figure 3 and Figure 4), where the controller plans to use both colloid and blood infusions (with colloids preferred) to regulate mean arterial pressure (MAP) around 90 mmHg, while avoiding hematocrit (HCT) values lower than 0.1.

Summary and Future Directions

In this dissertation, three inference-based methodologies were presented to facilitate the modeling, monitoring, and control of physiological processes, with special emphasis on enabling and supporting transparency, context-awareness, coordination, adaptation, and the appropriate handling of uncertainties in autonomous medical care systems.

First, we presented a collective variational inference (C-VI) method to facilitate the creation of personalized and generative physiological models from low-information and heterogeneous datasets. To illustrate the effectiveness of the C-VI method, we applied it to a practically important case study on modeling the hemodynamic effects of hemorrhage and fluid resuscitation. In the context of this case study, we demonstrated that the C-VI method can reconcile heterogeneous combinations of HCT, CO, and MAP data across multiple experiments to produce robust personalized and generative physiological models. In addition, we demonstrated that the C-VI method produces superior (more predictive) physiological models when compared to a non-collective maximum-likelihood estimation method, especially when only low-information and heterogeneous data are available.

Second, we presented a population-informed particle filtering (PIPF) method that fuses the information encoded in a generative physiological model with real-time clinical data to form perceptions of a patient's states, characteristics, and events. Using a case study on monitoring for hemodynamic management, we showed that the PIPF approach can provide reasonable beliefs (as compared to excluded data) about the

likely values and uncertainties associated with a patient's physiological variables (e.g., hematocrit levels), characteristics (e.g., tendency for atypical behavior), and events (e.g., hemorrhage). In addition, we demonstrated that incorporating population-level information into the filtering process (as is done in the PIPF algorithm) results in the formation of beliefs that are superior to those provided by a traditional particle filtering approach, especially for physiological variables that are weakly informed by the available patient-specific measurements. These results implied that PIPF is a promising candidate for use in physiological monitoring systems that process low-information and intermittent physiological measurements.

Third, we presented a population-informed variational control (PIVC) method and an associated physiological decision-support and closed-loop control (PCLC) framework that leverages the generative model, the perceptions of the PIPF algorithm, and user-defined definitions of actions and rewards in order to search for optimal courses of treatment for a patient. Using an in-silico case study on autonomous fluid resuscitation for hemodynamic management, we demonstrated that the control system arising from the proposed framework shows promise in exhibiting (i) *transparent behavior,* where the system's perceptions, reasoning, and decisions are human-interpretable; (ii) *context-aware behavior,* where the system is capable of remaining mindful of contextual and peripheral information in addition to its primary goal; (iii) *coordinated behavior*, where the system can coordinate multiple actions in synergistic ways to best achieve multiple objectives; (iv) *adaptable behavior,* where the system is equipped to identify and adapt to variabilities that exist within and across different patients; and (v) *uncertainty-aware behavior,* where the system can handle imperfect

and/or low-information measurements, quantify uncertainties that arise as a result, and incorporate them into its decisions. The results and analysis presented suggested that the proposed PCLC framework facilitates the desirable behaviors sought in the motivations of this work.

In terms of methodology, several future directions are conceivable. First, the C-VI method could be further extended to (i) better accommodate large and/or computationally expensive dynamic models through alternative gradient computation processes and/or gradient-free formulations; and (ii) form more expressive approximate posterior beliefs toward improved accuracy in the inference procedure, which could prove useful in applications with highly nonlinear generative models. Second, a variational variant of the proposed PIPF method could be devised with the capability to perform fast (but approximate) population-informed filtering using optimization, which may prove useful in scenarios with very large underlying models and/or large patient data streams. Third, a particle-based variants of the proposed PIVC method could be devised, with higher capability to reason about look-ahead beliefs with complex shapes (albeit at the cost of high computational burden), which could prove useful for applications with very nonlinear look-ahead procedures.

In terms of applications, the proposed methods and frameworks are expected to find use in future work on PCLC systems, especially in the context of critical care, including (i) simultaneous anesthesia and fluid resuscitation; (ii) hemodynamic management with fluid and drug combinations; and (iii) the resuscitation of burn injury.

Bibliography

[1] R. P. Dellinger *et al.*, "Surviving Sepsis Campaign: international guidelines for management of severe sepsis and septic shock, 2012," *Intensive Care Med*, vol. 39, no. 2, pp. 165–228, 2013.

[2] J. C. Hemphill III *et al.*, "Guidelines for the management of spontaneous intracerebral hemorrhage: a guideline for healthcare professionals from the American Heart Association/American Stroke Association," *Stroke*, vol. 46, no. 7, pp. 2032–2060, 2015.

[3] P. C. Wing, "Early acute management in adults with spinal cord injury: a clinical practice guideline for health-care providers. Who should read it?," *J Spinal Cord Med*, vol. 31, no. 4, p. 360, 2008.

[4] Y. Yoshino *et al.*, "The wound/burn guidelines--6: Guidelines for the management of burns," *The Journal of Dermatology*, vol. 43, no. 9, pp. 989–1010, 2016.

[5] P. Kaur, S. Basu, G. Kaur, and R. Kaur, "Transfusion protocol in trauma," *Journal of Emergencies, Trauma and Shock*, vol. 4, no. 1, p. 103, 2011.

[6] M. N. Bangash, M.-L. Kong, and R. M. Pearse, "Use of inotropes and vasopressor agents in critically ill patients," *Br J Pharmacol*, vol. 165, no. 7, pp. 2015–2033, 2012.

[7] A. Senz and L. Nunnink, "Inotrope and vasopressor use in the emergency department," *Emergency Medicine Australasia*, vol. 21, no. 5, pp. 342–351, 2009.

[8] B. E. Cooper, "Review and update on inotropes and vasopressors," *AACN Advanced critical care*, vol. 19, no. 1, pp. 5–13, 2008.

[9] R. Bighamian, C. Rubbo, J. E. Thorsen, J.-O. Hahn, and A. T. Reisner, "Is there opportunity for automated decision-support and closed-loop control in ICU patients receiving vasopressor infusion?," in *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2014, pp. 1949–1952.

[10] M. Saeed *et al.*, "Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): a public-access intensive care unit database," *Crit Care Med*, vol. 39, no. 5, p. 952, 2011.

[11] C. S. Anderson *et al.*, "Rapid blood-pressure lowering in patients with acute intracerebral hemorrhage," *N Engl j Med*, vol. 368, pp. 2355–2365, 2013.

[12] G. Hawryluk *et al.*, "Mean arterial blood pressure correlates with neurological recovery after human spinal cord injury: analysis of high frequency physiologic data," *J Neurotrauma*, vol. 32, no. 24, pp. 1958–1967, 2015.

[13] B. A. Cotton, J. S. Guy, J. A. Morris Jr, and N. N. Abumrad, “The cellular, metabolic, and systemic consequences of aggressive fluid resuscitation strategies,” *Shock*, vol. 26, no. 2, pp. 115–121, 2006.

[14] F. K. Bacomo and K. K. Chung, “A primer on burn resuscitation,” *Journal of Emergencies, Trauma and Shock*, vol. 4, no. 1, p. 109, 2011.

[15] J. R. Saffle, “The phenomenon of ‘fluid creep’ in acute burn resuscitation,” *Journal of Burn Care & Research*, vol. 28, no. 3, pp. 382–395, 2007.

[16] J. R. Saffle, “Fluid creep and over-resuscitation,” *Crit Care Clin*, vol. 32, no. 4, pp. 587–598, 2016.

[17] F. J. Doyle, L. M. Huyett, J. B. Lee, H. C. Zisser, and E. Dassau, “Closed-loop artificial pancreas systems: engineering the algorithms,” *Diabetes Care*, vol. 37, no. 5, pp. 1191–1197, 2014.

[18] C. K. Boughton and R. Hovorka, “Advances in artificial pancreas systems,” *Sci Transl Med*, vol. 11, no. 484, 2019.

[19] E. Brogi, S. Cyr, R. Kazan, F. Giunta, and T. M. Hemmerling, “Clinical Performance and Safety of Closed-Loop Systems: A Systematic Review and Meta-analysis of Randomized Controlled Trials,” *Anesth Analg*, vol. 124, no. 2, pp. 446–455, Feb. 2017.

[20] S. Coeckelenbergh *et al.*, “Automated systems for perioperative goal-directed hemodynamic therapy,” *J Anesth*, vol. 34, no. 1, pp. 104–114, Feb. 2020.

[21] A. Weisman, J. W. Bai, M. Cardinez, C. K. Kramer, and B. A. Perkins, “Effect of artificial pancreas systems on glycaemic control in patients with type 1 diabetes: a systematic review and meta-analysis of outpatient randomised controlled trials,” *Lancet Diabetes Endocrinol*, vol. 5, no. 7, pp. 501–512, Jul. 2017.

[22] L. Pasin *et al.*, “Closed-Loop Delivery Systems Versus Manually Controlled Administration of Total IV Anesthesia: A Meta-analysis of Randomized Clinical Trials,” *Anesthesia & Analgesia*, vol. 124, no. 2, pp. 456–464, Feb. 2017.

[23] C. Zaouter, A. Joosten, J. Rinehart, M. M. R. F. Struys, and T. M. Hemmerling, “Autonomous Systems in Anesthesia: Where Do We Stand in 2020? A Narrative Review,” *Anesth Analg*, vol. 130, no. 5, pp. 1120–1132, 2020.

[24] J. Schüttler and H. Ihmsen, “Population pharmacokinetics of propofola multicenter study,” *Anesthesiology: The Journal of the American Society of Anesthesiologists*, vol. 92, no. 3, pp. 727–738, 2000.

[25] C. Jeleazcov *et al.*, “Pharmacodynamic modelling of the bispectral index response to propofol-based anaesthesia during general surgery in children,” *Br J Anaesth*, vol. 100, no. 4, pp. 509–516, 2008.

[26] J. F. Standing, G. B. Hammer, W. J. Sam, and D. R. Drover, "Pharmacokinetic--pharmacodynamic modeling of the hypotensive effect of remifentanil in infants undergoing cranioplasty," *Pediatric Anesthesia*, vol. 20, no. 1, pp. 7–18, 2010.

[27] J.-O. Hahn, G. A. Dumont, and J. M. Ansermino, "A direct dynamic dose-response model of propofol for individualized anesthesia care," *IEEE transactions on biomedical engineering*, vol. 59, no. 2, pp. 571–578, 2011.

[28] J.-S. Shieh, M. F. Abbod, C. Y. Hsu, S.-J. Huang, Y.-Y. Han, and S.-Z. Fan, "Monitoring and control of anesthesia using multivariable self-organizing fuzzy logic structure," in *Fuzzy Systems in Bioinformatics and Computational Biology*, Springer, 2009, pp. 273–295.

[29] K. van Heusden *et al.*, "Design and clinical evaluation of robust PID control of propofol anesthesia in children," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 2, pp. 491–501, 2013.

[30] M. Yousefi, K. van Heusden, I. M. Mitchell, J. M. Ansermino, and G. A. Dumont, "Falsified model-invariant safety-preserving control with application to closed-loop anesthesia," *IEEE Transactions on Control Systems Technology*, 2018.

[31] A. Gentilini *et al.*, "Modeling and closed-loop control of hypnosis by means of bispectral index (BIS) with isoflurane," *IEEE transactions on biomedical engineering*, vol. 48, no. 8, pp. 874–889, 2001.

[32] Y. Sawaguchi, E. Furutani, G. Shirakami, M. Araki, and K. Fukuda, "A model-predictive hypnosis control system under total intravenous anesthesia," *IEEE transactions on biomedical engineering*, vol. 55, no. 3, pp. 874–887, 2008.

[33] T. J. Sieber, C. W. Frei, M. Derighetti, P. Feigenwinter, D. Leibundgut, and A. M. Zbinden, "Model-based automatic feedback control versus human control of end-tidal isoflurane concentration using low-flow anaesthesia," *Br J Anaesth*, vol. 85, no. 6, pp. 818–825, 2000.

[34] J. Nino, R. de Keyser, S. Syafiie, C. Ionescu, and M. Struys, "EPSAC-controlled anesthesia with online gain adaptation," *Int J Adapt Control Signal Process*, vol. 23, no. 5, pp. 455–471, 2009.

[35] C. S. Nunes, T. Mendonça, J. M. Lemos, and P. Amorim, "Feedforward adaptive control of the Bispectral Index of the EEG using the intravenous anaesthetic drug propofol," *Int J Adapt Control Signal Process*, vol. 23, no. 5, pp. 485–503, 2009.

[36] T. Mendonça, J. M. Lemos, H. Magalhaes, P. Rocha, and S. Esteves, "Drug delivery for neuromuscular blockade with supervised multimodel adaptive control," *IEEE Transactions on Control systems technology*, vol. 17, no. 6, pp. 1237–1244, 2009.

[37] T. de Smet, M. M. R. F. Struys, S. Greenwald, E. P. Mortier, and S. L. Shafer, "Estimation of optimal modeling weights for a Bayesian-based closed-loop system for propofol administration using the bispectral index as a controlled

variable: a simulation study," *Anesthesia & Analgesia*, vol. 105, no. 6, pp. 1629–1638, 2007.

[38] T. de Smet, M. M. R. F. Struys, M. M. Neckebroek, K. den Hauwe, S. Bonte, and E. P. Mortier, "The accuracy and clinical feasibility of a new bayesian-based closed-loop control system for propofol administration using the bispectral index as a controlled variable," *Anesthesia & Analgesia*, vol. 107, no. 4, pp. 1200–1210, 2008.

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

[40] N. Liu *et al.*, "Closed-loop coadministration of propofol and remifentanil guided by bispectral index: a randomized multicenter study," *Anesthesia \& Analgesia*, vol. 112, no. 3, pp. 546–557, 2011.

[41] O. Simanski, M. Janda, A. Schubert, J. Bajorat, R. Hofmockel, and B. Lampe, "Progress of automatic drug delivery in anaesthesia—The 'Rostock assistant system for anaesthesia control (RAN),'" *International Journal of Adaptive Control and Signal Processing*, vol. 23, no. 5, pp. 504–521, 2009.

[42] K. van Heusden, J. M. Ansermino, and G. A. Dumont, "Robust MISO control of propofol-remifentanil anesthesia guided by the NeuroSENSE monitor," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 5, pp. 1758–1770, 2017.

[43] X. Jin, C.-S. Kim, S. T. Shipley, G. A. Dumont, and J.-O. Hahn, "Coordinated semi-adaptive closed-loop control for infusion of two interacting medications," *International Journal of Adaptive Control and Signal Processing*, vol. 32, no. 1, pp. 134–146, 2018.

[44] J. Rinehart and C. Canales, "Closed-loop pharmacology in anesthesia and critical care: Benefits and limitations," *International Anesthesiology Clinics*, vol. 53, no. 2, pp. 91–101, 2015.

[45] F. Michard, "Decision support for hemodynamic management: from graphical displays to closed loop systems," *Anesthesia & Analgesia*, vol. 117, no. 4, pp. 876–882, 2013.

[46] J. Rinehart, N. Liu, B. Alexander, and M. Cannesson, "Closed-loop systems in anesthesia: Is there a potential for closed-loop fluid management and hemodynamic optimization?," *Anesthesia & Analgesia*, vol. 114, no. 1, pp. 130–143, 2012.

[47] N. F. Chaisson, R. A. Kirschner, D. J. Deyo, J. A. Lopez, D. S. Prough, and G. C. Kramer, "Near-infrared spectroscopy-guided closed-loop resuscitation of hemorrhage," *Journal of Trauma and Acute Care Surgery*, vol. 54, no. 5, pp. S183--S192, 2003.

[48] N. R. Marques *et al.*, "Automated closed-loop resuscitation of multiple hemorrhages: a comparison between fuzzy logic and decision table controllers in a sheep model," *Disaster Mil Med*, vol. 3, no. 1, Dec. 2017.

[49] G. C. Kramer *et al.*, "Closed-loop control of fluid therapy for treatment of hypovolemia," *Journal of Trauma and Acute Care Surgery*, vol. 64, no. 4, pp. S333--S341, 2008.

[50] J. Rinehart, E. Chung, C. Canales, and M. Cannesson, "Intraoperative stroke volume optimization using stroke volume, arterial pressure, and heart rate: closed-loop (learning intravenous resuscitator) versus anesthesiologists," *J Cardiothorac Vasc Anesth*, vol. 26, no. 5, pp. 933–939, 2012.

[51] G. Hundeshagen *et al.*, "Closed-Loop– and Decision-Assist–Guided Fluid Therapy of Human Hemorrhage," *Critical Care Medicine*, vol. 45, no. 10, pp. e1068–e1074, Oct. 2017.

[52] X. Jin, R. Bighamian, and J.-O. Hahn, "Development and in Silico Evaluation of a Model-Based Closed-Loop Fluid Resuscitation Control Algorithm," *IEEE Transactions on Biomedical Engineering*, 2018.

[53] J. Salinas *et al.*, "Closed-loop and decision-assist resuscitation of burn patients," *Journal of Trauma and Acute Care Surgery*, vol. 64, no. 4, pp. S321--S332, 2008.

[54] J. Salinas *et al.*, "Computerized decision support system improves fluid resuscitation following severe burns: An original study," *Critical Care Medicine*, vol. 39, no. 9, pp. 2031–2038, Sep. 2011.

[55] G. Arturson, T. Groth, A. Hedlund, and B. Zaar, "Computer simulation of fluid resuscitation in trauma: first pragmatic validation in thermal injury," *J Burn Care Rehabil*, vol. 10, no. 4, pp. 292–299, 1989.

[56] R. Cartotto, D. G. Greenhalgh, and C. Cancio, "Burn state of the science: fluid resuscitation," *Journal of Burn Care & Research*, vol. 38, no. 3, pp. e596--e604, 2017.

[57] T. Wassar *et al.*, "Automatic control of arterial pressure for hypotensive patients using phenylephrine," *International Journal of Modelling and Simulation*, vol. 34, no. 4, pp. 187–198, 2014.

[58] M. Merouani *et al.*, "Norepinephrine weaning in septic shock patients by closed loop control based on fuzzy logic," *Critical Care*, vol. 12, no. 6, p. R155, 2008.

[59] J. Rinehart, M. Ma, M.-D. Calderon, and M. Cannesson, "Feasibility of automated titration of vasopressor infusions using a novel closed-loop controller," *J Clin Monit Comput*, vol. 32, no. 1, pp. 5–11, 2018.

[60] A. Joosten *et al.*, "Feasibility of closed-loop titration of norepinephrine infusion in patients undergoing moderate-and high-risk surgery," *Br J Anaesth*, vol. 123, no. 4, pp. 430–438, 2019.

[61] A. Joosten *et al.*, "Feasibility of fully automated hypnosis, analgesia, and fluid management using 2 independent closed-loop systems during major vascular surgery: a pilot study," *Anesthesia & Analgesia*, vol. 128, no. 6, pp. e88--e92, 2019.

[62] N. Libert *et al.*, "Performance of closed-loop resuscitation of haemorrhagic shock with fluid alone or in combination with norepinephrine: an experimental study," *Ann Intensive Care*, vol. 8, no. 1, p. 89, 2018.

[63] K. Kashihara, "Automatic regulation of hemodynamic variables in acute heart failure by a multiple adaptive predictive controller based on neural networks," *Ann Biomed Eng*, vol. 34, no. 12, pp. 1846–1869, 2006.

[64] J. Rinehart, C. Lee, C. Canales, A. Kong, Z. Kain, and M. Cannesson, "Closed-loop fluid administration compared to anesthesiologist management for hemodynamic optimization and resuscitation during surgery: an in vivo study," *Anesthesia & Analgesia*, vol. 117, no. 5, pp. 1119–1129, 2013.

[65] B. Parvinian, C. Scully, H. Wiyor, A. Kumar, and S. Weininger, "Regulatory considerations for physiological closed-loop controlled medical devices used for automated critical care: food and drug administration workshop discussion topics," *Anesth Analg*, vol. 126, no. 6, p. 1916, 2018.

[66] B. Parvinian *et al.*, "Credibility evidence for computational patient models used in the development of physiological closed-loop controlled devices for critical care medicine," *Frontiers in Physiology*, vol. 10, p. 220, 2019.

[67] A. Tivay, G. C. Kramer, and J. O. Hahn, "Collective Variational Inference for Personalized and Generative Physiological Modeling: A Case Study on Hemorrhage Resuscitation," *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 2, pp. 666–677, Feb. 2022.

[68] A. Joosten *et al.*, "Automated Titration of Vasopressor Infusion Using a Closed-loop Controller: In Vivo Feasibility Study Using a Swine Model," *Anesthesiology: The Journal of the American Society of Anesthesiologists*, vol. 130, no. 3, pp. 394–403, 2019.

[69] F. H. El-Khatib, S. J. Russell, D. M. Nathan, R. G. Sutherlin, and E. R. Damiano, "A bihormonal closed-loop artificial pancreas for type 1 diabetes," *Science Translational Medicine*, vol. 2, no. 27, pp. 27ra27--27ra27, 2010.

[70] O. Wolkenhauer, "Why model?," *Front Physiol*, vol. 5, p. 21, 2014.

[71] R. E. Baker, J.-M. Peña, J. Jayamohan, and A. Jérusalem, "Mechanistic models versus machine learning, a fight worth fighting for the biological community?," *Biology Letters*, vol. 14, no. 5, p. 20170660, 2018.

[72] I. Tavassoly, J. Goldfarb, and R. Iyengar, "Systems biology primer: the basic methods and approaches," *Essays in Biochemistry*, vol. 62, no. 4, pp. 487–500, 2018.

[73] R. J. Rossi, *Mathematical statistics: an introduction to likelihood based inference*. John Wiley & Sons, 2018.

[74] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian data analysis*. CRC press, 2013.

[75] D. J. Albers, M. E. Levine, A. Stuart, L. Mamykina, B. Gluckman, and G. Hripcsak, "Mechanistic machine learning: how data assimilation leverages physiologic knowledge using Bayesian inference to forecast the future, infer the present, and phenotype," *Journal of the American Medical Informatics Association*, vol. 25, no. 10, pp. 1392–1401, 2018.

[76] S. Chib and E. Greenberg, "Understanding the metropolis-hastings algorithm," *The American Statistician*, vol. 49, no. 4, pp. 327–335, 1995.

[77] M. Girolami and B. Calderhead, "Riemann manifold langevin and hamiltonian monte carlo methods," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 73, no. 2, pp. 123–214, 2011.

[78] T. Toni, D. Welch, N. Strelkowa, A. Ipsen, and M. P. H. Stumpf, "Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems," *Journal of The Royal Society Interface*, vol. 6, no. 31, pp. 187–202, Feb. 2009.

[79] M. A. Beaumont, "Approximate bayesian computation," *Annual Review of Statistics and Its Application*, vol. 6, pp. 379–403, 2019.

[80] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational inference: A review for statisticians," *Journal of the American statistical Association*, vol. 112, no. 518, pp. 859–877, 2017.

[81] C. Zhang, J. Bütepage, H. Kjellström, and S. Mandt, "Advances in variational inference," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 8, pp. 2008–2026, 2018.

[82] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, "Stochastic variational inference," *The Journal of Machine Learning Research*, vol. 14, no. 1, pp. 1303–1347, 2013.

[83] D. P. Kingma, M. Welling, and others, "An Introduction to Variational Autoencoders," *Foundations and Trends® in Machine Learning*, vol. 12, no. 4, pp. 307–392, 2019.

[84] C.-H. Zhang, "Compound decision theory and empirical Bayes methods," *The Annals of Statistics*, vol. 31, no. 2, pp. 379–390, 2003.

[85] B. Efron and others, "Bayes, oracle Bayes and empirical Bayes," *Statistical Science*, vol. 34, no. 2, pp. 177–201, 2019.

[86] M. Davidian and D. M. Giltinan, "Nonlinear models for repeated measurement data: An overview and update," *Journal of Agricultural, Biological, and Environmental Statistics*, vol. 8, no. 4, pp. 387–419, Dec. 2003.

[87] Y. Wang, “Derivation of various NONMEM estimation methods,” *Journal of Pharmacokinetics and pharmacodynamics*, vol. 34, no. 5, pp. 575–593, 2007.

[88] W. Pan, Y. Yuan, L. Ljung, J. Gonçalves, and G.-B. Stan, “Identification of nonlinear state-space systems from heterogeneous datasets,” *IEEE Transactions on Control of Network Systems*, vol. 5, no. 2, pp. 737–747, 2017.

[89] W. R. Jacobs, T. Baldacchino, T. Dodd, and S. R. Anderson, “Sparse Bayesian nonlinear system identification using variational inference,” *IEEE Transactions on Automatic Control*, vol. 63, no. 12, pp. 4172–4187, 2018.

[90] G. Roeder, P. Grant, A. Phillips, N. Dalchau, and T. Meeds, “Efficient Amortised Bayesian Inference for Hierarchical and Nonlinear Dynamical Systems,” in *International Conference on Machine Learning*, 2019, pp. 4445–4455.

[91] A. Nazabal, P. M. Olmos, Z. Ghahramani, and I. Valera, “Handling incomplete heterogeneous data using vaes,” *Pattern Recognition*, p. 107501, 2020.

[92] D. Koller and N. Friedman, *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

[93] M. J. Johnson, D. Duvenaud, A. B. Wiltschko, S. R. Datta, and R. P. Adams, “Structured {VAEs}: Composing probabilistic graphical models and variational autoencoders,” *arXiv preprint arXiv:1603.06277*, vol. 2, p. 2016, 2016.

[94] A. Tivay, X. Jin, A. K.-Y. Lo, C. G. Scully, and J.-O. Hahn, “Practical Use of Regularization in Individualizing a Mathematical Model of Cardiovascular Hemodynamics Using Scarce Data,” *Frontiers in Physiology*, vol. 11, 2020.

[95] M. K. Transtrum, B. B. Machta, and J. P. Sethna, “Why are nonlinear fits to data so challenging?,” *Phys Rev Lett*, vol. 104, no. 6, p. 60201, 2010.

[96] 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. 07B201_1, 2015.

[97] H. H. Mattingly, M. K. Transtrum, M. C. Abbott, and B. B. Machta, “Maximizing the information learned from finite data selects a simple model,” *Proceedings of the National Academy of Sciences*, vol. 115, no. 8, pp. 1760–1765, 2018.

[98] J. Duchi, E. Hazan, and Y. Singer, “Adaptive subgradient methods for online learning and stochastic optimization.,” *Journal of machine learning research*, vol. 12, no. 7, 2011.

[99] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2014.

[100] A. Défossez, L. Bottou, F. Bach, and N. Usunier, “A Simple Convergence Proof of Adam and Adagrad,” *arXiv preprint arXiv:2003.02395*, 2020.

[101] B. Kleinberg, Y. Li, and Y. Yuan, “An alternative view: When does SGD escape local minima?,” in *International Conference on Machine Learning*, 2018, pp. 2698–2707.

[102] V. A. Convertino and N. J. Koons, “The compensatory reserve: potential for accurate individualized goal-directed whole blood resuscitation,” *Transfusion (Paris)*, vol. 60, pp. S150--S157, 2020.

[103] V. A. Convertino *et al.*, “Validating clinical threshold values for a dashboard view of the compensatory reserve measurement for hemorrhage detection,” *Journal of Trauma and Acute Care Surgery*, vol. 89, no. 2S, pp. S169--S174, 2020.

[104] 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 Engineering Practice*, vol. 73, no. April 2018, pp. 149–160, Apr. 2018.

[105] A. Tivay, G. C. Kramer, and J. O. Hahn, “Virtual Patient Generation using Physiological Models through a Compressed Latent Parameterization,” *Proceedings of the American Control Conference*, vol. 2020-July, pp. 1335–1340, Jul. 2020.

[106] R. G. Hahn, “Volume kinetics for infusion fluids,” *Anesthesiology*, vol. 113, no. 2. Lippincott Williams and Wilkins, pp. 470–481, Aug. 2010.

[107] R. Bighamian, A. T. Reisner, and J.-O. Hahn, “A Lumped-Parameter Subject-Specific Model of Blood Volume Response to Fluid Infusion,” *Frontiers in Physiology*, vol. 7, no. AUG, p. 390, Aug. 2016.

[108] J. E. Hall, *Guyton and Hall textbook of medical physiology*.

[109] Y. Çınar, G. Demir, M. Paç, and A. B. Çınar, “Effect of hematocrit on blood pressure via hyperviscosity,” *American Journal of Hypertension*, vol. 12, no. 7, pp. 739–743, Jul. 1999.

[110] P. C. Hansen and D. P. O’Leary, “The Use of the L-Curve in the Regularization of Discrete Ill-Posed Problems,” *SIAM Journal on Scientific Computing*, vol. 14, no. 6, pp. 1487–1503, Nov. 1993.

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

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

[113] M. Viceconti, M. A. Juárez, C. Curreli, M. Pennisi, G. Russo, and F. Pappalardo, “Credibility of In Silico Trial Technologies—A Theoretical Framing,” *IEEE Journal of Biomedical and Health Informatics*, vol. 24, no. 1, pp. 4–13, 2019.

[114] M. Viceconti, F. Pappalardo, B. Rodriguez, M. Horner, J. Bischoff, and F. M. Tshinanu, "In silico trials: Verification, validation and uncertainty quantification of predictive models used in the regulatory evaluation of biomedical products," *Methods*, 2020.

[115] O. T. Inan *et al.*, "Digitizing clinical trials," *npj Digital Medicine*, vol. 3, no. 1, pp. 1–7, 2020.

[116] S. A. Niederer *et al.*, "Creation and application of virtual patient cohorts of heart models," *Philosophical Transactions of the Royal Society A*, vol. 378, no. 2173, p. 20190558, 2020.

[117] J. F. Dyro, *Clinical engineering handbook*. Elsevier Science Ltd., 2004.

[118] M. N. Sawka and K. E. Friedl, "Emerging Wearable Physiological Monitoring Technologies and Decision Aids for Health and Performance," *Journal of Applied Physiology*, vol. 124, no. 2, pp. 430–431, 2018.

[119] J.-O. Hahn and O. T. Inan, "Physiological closed-loop control in critical care: Opportunities for innovations," *Progress in Biomedical Engineering*, May 2022.

[120] A. Gelb and others, *Applied optimal estimation*. MIT press, 1974.

[121] B. D. O. Anderson and J. B. Moore, *Optimal filtering*. Courier Corporation, 2012.

[122] B. D. O. Anderson and J. B. Moore, "Kalman Filtering: Whence, What and Whither?," *Mathematical System Theory*, pp. 41–54, 1991.

[123] J. Humpherys, P. Redd, and J. West, "A fresh look at the kalman filter," *SIAM Review*, vol. 54, no. 4, pp. 801–823, 2012.

[124] E. A. Wan and R. van der Merwe, "The unscented Kalman filter for nonlinear estimation," *IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium, AS-SPCC 2000*, pp. 153–158, 2000.

[125] H. M. T. Menegaz, J. Y. Ishihara, G. A. Borges, and A. N. Vargas, "A Systematization of the Unscented Kalman Filter Theory," *IEEE Transactions on Automatic Control*, vol. 60, no. 10, pp. 2583–2598, Oct. 2015.

[126] G. A. Einicke and L. B. White, "Robust extended Kalman filtering," *IEEE Transactions on Signal Processing*, vol. 47, no. 9, pp. 2596–2599, 1999.

[127] S. Särkkä, *Bayesian filtering and smoothing*, no. 3. Cambridge university press, 2013.

[128] R. van de Schoot *et al.*, "Bayesian statistics and modelling," *Nature Reviews Methods Primers 2021 1:1*, vol. 1, no. 1, pp. 1–26, Jan. 2021.

[129] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis*. Chapman and Hall/CRC, 1995.

[130] O. Cappe, S. J. Godsill, and E. Moulines, “An overview of existing methods and recent advances in sequential Monte Carlo,” *Proceedings of the IEEE*, vol. 95, no. 5, pp. 899–924, 2007.

[131] H. R. Künsch, “Particle filters,” *Bernoulli*, vol. 19, no. 4, pp. 1391–1403, Sep. 2013.

[132] N. Kantas, A. Doucet, S. S. Singh, J. Maciejowski, and N. Chopin, “On particle methods for parameter estimation in state-space models,” *Statistical science*, vol. 30, no. 3, pp. 328–351, 2015.

[133] C. A. Naesseth, F. Lindsten, and T. B. Schön, “Elements of Sequential Monte Carlo,” *Foundations and Trends® in Machine Learning*, vol. 12, no. 3, pp. 307–392, Nov. 2019.

[134] A. Doucet and A. M. Johansen, “A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later,” *Oxford Handbook of Nonlinear Filtering*, vol. 12, pp. 656–704, 2011.

[135] V. Šmídl and A. Quinn, “Variational Bayesian filtering,” *IEEE Transactions on Signal Processing*, vol. 56, no. 10 II, pp. 5020–5030, 2008.

[136] K. J. Friston, “Variational filtering,” *Neuroimage*, vol. 41, no. 3, pp. 747–766, Jul. 2008.

[137] J. Marino, M. Cvitkovic, and Y. Yue, “A General Method for Amortizing Variational Filtering,” in *Advances in Neural Information Processing Systems*, 2018, vol. 31.

[138] C. Naesseth, S. Linderman, R. Ranganath, and D. Blei, “Variational Sequential Monte Carlo,” in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, May 2018, vol. 84, pp. 968–977.

[139] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational Inference: A Review for Statisticians,” *Journal of the American Statistical Association*, vol. 112, no. 518. American Statistical Association, pp. 859–877, Apr. 03, 2017.

[140] P. Blanchard, D. J. Higham, and N. J. Higham, “Accurately computing the log-sum-exp and softmax functions,” *IMA Journal of Numerical Analysis*, vol. 41, no. 4, pp. 2311–2330, Oct. 2021.

[141] A. E. Laher *et al.*, “A review of hemodynamic monitoring techniques, methods and devices for the emergency physician,” *The American Journal of Emergency Medicine*, vol. 35, no. 9, pp. 1335–1347, Sep. 2017.

[142] B. Saugel, J. L. Vincent, and J. Y. Wagner, “Personalized hemodynamic management,” *Current Opinion in Critical Care*, vol. 23, no. 4, pp. 334–341, Aug. 2017.

[143] A. Bougié, A. Harrois, and J. Duranteau, “Resuscitative strategies in traumatic hemorrhagic shock,” *Annals of Intensive Care*, vol. 3, no. 1, pp. 1–9, Jan. 2013.

[144] 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.

[145] A. D. Rafie *et al.*, "Hypotensive resuscitation of multiple hemorrhages using crystalloid and colloids," *Shock*, vol. 22, no. 3, pp. 262–269, Sep. 2004.

[146] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *Journal of the American statistical Association*, vol. 102, no. 477, pp. 359–378, 2007.

[147] A. Tivay, G. C. Kramer, and J. O. Hahn, "Inference-based subject atypicality and signal quality indicators for physiological data," *MCPS 2021 - Proceedings of the 2021 Medical Cyber Physical Systems and Internet of Medical Things*, pp. 7–11, May 2021.

[148] A. E. W. Johnson, M. M. Ghassemi, S. Nemati, K. E. Niehaus, D. Clifton, and G. D. Clifford, "Machine Learning and Decision Support in Critical Care," *Proceedings of the IEEE*, vol. 104, no. 2, pp. 444–466, Feb. 2016.

[149] O. Sigaud and O. Buffet, *Markov decision processes in artificial intelligence*. John Wiley & Sons, 2013.

[150] V. Krishnamurthy, *Partially observed Markov decision processes*. Cambridge university press, 2016.

[151] A. Piché, V. Thomas, C. Ibrahim, Y. Bengio, and C. Pal, "Probabilistic planning with sequential monte carlo methods," 2018.

[152] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[153] S. Levine, "Reinforcement learning and control as probabilistic inference: Tutorial and review," *arXiv preprint arXiv:1805.00909*, 2018.

[154] S. Kamthe and M. Deisenroth, "Data-efficient reinforcement learning with probabilistic model predictive control," in *International Conference on Artificial Intelligence and Statistics*, 2018, pp. 1701–1710.

[155] J. Marino, A. Piche, A. D. Ialongo, and Y. Yue, "Iterative Amortized Policy Optimization," in *Advances in Neural Information Processing Systems*, 2021, vol. 34, pp. 15667–15681.

[156] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT press, 2005.

[157] D. Silver *et al.*, "Mastering the game of Go with deep neural networks and tree search," *Nature 2016 529:7587*, vol. 529, no. 7587, pp. 484–489, Jan. 2016.

[158] D. Silver *et al.*, "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play," *Science (1979)*, vol. 362, no. 6419, pp. 1140–1144, Dec. 2018.

[159] C. Yu, J. Liu, S. Nemati, and G. Yin, “Reinforcement Learning in Healthcare: A Survey,” *ACM Computing Surveys (CSUR)*, vol. 55, no. 1, pp. 1–36, Nov. 2021.

[160] R. Ranganath, S. Gerrish, and D. Blei, “Black box variational inference,” in *Artificial Intelligence and Statistics*, 2014, pp. 814–822.