



Published in final edited form as:

Math Biosci. 2019 October ; 316: 108242. doi:10.1016/j.mbs.2019.108242.

The Parameter Houlihan: a solution to high-throughput identifiability indeterminacy for brutally ill-posed problems

DJ Albers^{a,b}, ME Levine^{a,c}, L Mamykina^a, G Hripcsak^a

^aDepartment of Biomedical Informatics, Columbia University, 622 West 168th Street, PH-20, New York, NY

^bDepartment of Pediatrics, Division of Informatics, University of Colorado Medicine, Mail: F443, 13199 E. Montview Blvd. Ste: 210-12 – Aurora, CO 80045

^cDepartment of computational and mathematical sciences, 1200 E California Blvd M/C 305-16 Pasadena, CA 91125

Abstract

One way to interject knowledge into clinically impactful forecasting is to use data assimilation, a nonlinear regression that projects data onto a mechanistic physiologic model, instead of a set of functions, such as neural networks. Such regressions have an advantage of being useful with particularly sparse, non-stationary clinical data. However, physiological models are often nonlinear and can have many parameters, leading to potential problems with parameter identifiability, or the ability to find a unique set of parameters that minimize forecasting error. The identifiability problems can be minimized or eliminated by reducing the number of parameters estimated, but reducing the number of estimated parameters also reduces the flexibility of the model and hence increases forecasting error. We propose a method, the parameter Houlihan, that combines traditional machine learning techniques with data assimilation, to select the right set of model parameters to minimize forecasting error while reducing identifiability problems. The method worked well: the data assimilation-based glucose forecasts and estimates for our cohort using the Houlihan-selected parameter sets generally also minimize forecasting errors compared to other parameter selection methods such as by-hand parameter selection. Nevertheless, the forecast with the lowest forecast error does not always accurately represent physiology, but further advancements of the algorithm provide a path for improving physiologic fidelity as well. Our hope is that this methodology represents a first step toward combining machine learning with data assimilation and provides a lower-threshold entry point for using data assimilation with clinical data by helping select the right parameters to estimate.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Conflict of interest

No author has a conflict of interest nor a competing interest to report.

Keywords

data assimilation; identifiability; machine learning; inverse problems; physiology; Markov Chain Monte Carlo

1. Introduction

We want to use data and our understanding of the world to better manage health — we want evidence and understanding to guide clinical and personal health-related decisions. Of course at a high level this is generally what medicine is about: interventions are undertaken only when they are understood or predicted to improve an individual’s health. However, traditionally this prediction is done in a non-personalized manner, meaning that interventions treat the “mean” person or patient. Personalized and precision medicine were conceptualized to relax this constraint by tailoring an intervention to a person. While genetics offers a path to personalizing treatment, we can also use data science machinery together with personal ([1]) and population-scale data to better personalize treatment ([2, 3, 4]). Specifically here, we want to leverage our knowledge encapsulated in mechanistic physiologic models and combine it with free living or clinical data to allow this knowledge and data to be used to make decisions related to health. In this context, computational problems related to personalized medicine can be broken into two broad categories: *forecasting*, where we make quantitative predictions about a patient’s future state that can be used by clinicians and patients to take corrective action, and *phenotyping* ([5, 6, 7, 8, 9]), where we identify properties of macroscopic observables that can be used to classify patients into subgroups that can give clinicians and researchers actionable insight into commonly occurring treatment outcomes and biological phenomena.

The idea of using mechanistic models and data assimilation in biomedicine or healthcare is old, but what is new is attempting to integrate models with variable complexity with sparse, noisy free-living and clinically collected data. Many mathematical biology models were designed to have variable degrees of biological fidelity, fidelity that we do not necessarily want to eliminate or reduce, but fidelity that we generally need to constrain in the usual case where we cannot estimate all the parameters because of data limitations that always exist in practice. This problem poses a significant barrier to using data assimilation—enough of a barrier that often data assimilation is not even attempted because the models, given data are hopelessly poorly resolved. This paper poses a machine learning solution to this problem—by using machine learning to identify and rank-order which model parameters are the most necessary to estimate.

Returning to the more practical contexts of phenotyping and forecasting, both applications impose particular demands on certain aspects of computational machinery used to model data. The properties we focus on here are the selection of the model parameters to estimate and the ensuing *identifiability* of a model, or ability to uniquely solve for parameters that yield optimal solutions ([10, 11, 12]). Our goal is to strike a balance a between identifiability and model fidelity [5] in situations where a model is not fully identifiable, where estimating all model parameters is hopeless, and where choosing model parameters to estimate is non-

trivial, *given the available data*. The method we develop here can facilitate both forecasting and phenotyping studies, and we evaluate this method in the context of modeling glucose dynamics using mechanistic models, machine learning and data assimilation.

The Houlihan, or the Houlihan throw, is a lasso throw used for roping livestock, e.g., a horse. It is used often under difficult circumstances such as picking out, from a substantial distance, a single horse from among a crowd of horses standing close together. It is a particularly flexible technique that can be used in a variety of circumstances. In this spirit we intuitively define the Houlihan method(s) as a collection of methods that can be used for selecting the most productive model parameters to estimate; specifically, the collection of methods uses machine learning techniques applied to simulated model output under parameter variation subject to a set of features, e.g., the mean of a state.

2. Background

The larger biomedical context of this work is the application of data science machinery used to personalize forecasts and phenotypes via a broadly defined regression. While there are many linear versions of regression that have been successfully applied to healthcare data ([13, 14, 15, 16]), here we focus on a specific type of nonlinear regression—data assimilation—in an effort to take advantage of potentially important nonlinearities present in most biological systems. Nonlinear regression approaches such as deep learning and related methods ([17, 18, 9, 19]) have seen some success in a number of biomedical applications thanks to their ability to approximate arbitrary, non-linear functions. While the flexibility of universal approximator approaches ([20, 21]) is particularly useful when little is known about the system and data are plentiful, this approach does not always work well when data are sparse and non-stationary, leading to problems such as poor generalization to new or unobserved individuals, problems with quickly changing health conditions, and difficulties with fast, accurate prediction with very few, e.g., 20, data points. Unfortunately, many health data and healthcare situations fit one or more of these data pathologies ([22, 23]).

In order to exploit the complex yet rich quantity of available health data, it is natural to consider ways of constraining the search space for machine learning methods. One way to do this is to constrain the model search space in accordance with as much expert knowledge as possible. To achieve this here we turn to mechanistic models developed by mathematical biologists and physiologists [24], which are typically formulated as dynamical systems ([25, 26, 27]), e.g., $x_{t+1} = f(x_t, \theta)$, or differential equations ([28, 29]), e.g., $\frac{dx}{dt} = f(x, \theta, t)$, where x are the time-varying states of the system and θ are the physiologic parameters that govern the process. For example, in the case of phenotyping type 2 diabetes one way of constraining the search space of a regression is to regress the data onto a nonlinear physiologic model [1, 30] instead of regressing the data onto a universal approximator [20, 21] function space such as neural networks. The way this is done is using data assimilation.

Data assimilation (DA) is a collection of methods ([31, 32, 33, 34, 35, 36, 37, 38, 39, 40]) concerned with performing the types of non-linear regressions we describe for dynamical systems, and centers itself around forecasting and inferring mechanistic states under available observations; it solves both forward and inverse problems ([41, 42, 43]). There

have been many successful applications of mechanistic modeling and data assimilation in biomedicine ([30, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60, 61, 62, 42, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 1, 74]). However, mechanistic models that are typically developed in biological laboratory settings are often not designed to interface with health data collected in the process of delivering care or in free-living situations—in particular, the physiologic models often model macroscopic states that *are* observable from routinely collected data but are governed by a composition of *unobservable* or partially mechanisms such as the health care process [22, 16, 23, 5]. While these models capture the dynamics we are interested in and constrain the regression to a smaller class of functions, their high-fidelity creates issues of identifiability and ill-posedness, problems for which this paper develops a practical, machine-learning-based work-around.

To understand how identifiability works for these machines, consider a trivial case of identifiability for the model $\frac{dx}{dt} = abx$. If we assume that a and b are unknown parameters, they cannot both be identifiable without another equation that could uniquely determine one of them. This topic and the associated methods for handling this situation are too old and wide ranging to give complete background ([11, 10, 75, 12]). We can, however, give a broad sketch of how identifiability has been traditionally approached. Identifiability analysis generally follows one of three pathways: *analytical methods*, e.g., showing algebraically that all parameters can be uniquely solved for ([76, 77, 78]); *numerical methods* ([12, 10, 79]); and *heuristic, knowledge-driven* sensitivity analyses where certain parameters are chosen based on computational experiments or knowledge of the system. In many complex, non-linear mechanistic physiologic models algebraic methods and linear computational methods are not tractable or applicable. In these situations nonlinear methods can be applied, but nonlinear methods usually have to be constructed for a particular situation ([80]), and, much like nonlinear optimization, generally do not have clearcut or simple resolutions ([80, 75]). *These problems pose a significant roadblock to parameter inference in the context of DA.* Nevertheless, there exist methods for working to remain within a traditional identifiability framework, e.g., [75] uses Bayesian inference to determine when parameters can be made identifiable.

The usual way of addressing these issues focuses on making sure the model is identifiable or finding ways of making it more identifiable ([75, 12, 10]). This work is often performed using substantial intuition about the important features encoded in the model, and parameterizing and grouping sub-processes. However, this creates silos of expertise and prevents wide-spread dissemination and evaluation of mechanistic models in potential application domains. Therefore, to progress toward understanding complex physiology via model refinement and selection, and to provide solutions in clinical situations that come with constraints of time-sensitive solutions, we must find a robust way of coping with brutally ill-posed problems and accept certain impurities and inaccuracies.

Here, we develop and evaluate a method—the Houlihan method—for rank-ordering mechanistic parameters based on their “influence” on important dynamical features, in order to improve forecast accuracy and help determine which models most faithfully

represent a given system. This provides a starting point from which to estimate parameters, prune the model, etc., that can be automated.

3. Conceptual construction of the Houlihan approach

3.1. Conventional operational use of data assimilation with ill-posed problems

The standard method of applying data assimilation (DA) or control in generic situations follows roughly the following steps: (i) select a model, (ii) work out identifiability, (iii) select a filter or inference method, (iv) find an optimal solution for states and parameters. This requires very careful experimental constructions, generally dense data streams, can be expensive, and requires relatively simple models, all situations that lie outside of what is possible in applications and even many basic science settings. The approach for applying DA in *operational*, complex, high-dimensional settings where accurate *real-time* forecasts are imperative is to: (i) select or develop a model, (ii) tune and fix parameters offline, often by hand or using a combination of by-hand and numeric tuning that allows the model to reconstruct or forecast states within some tolerance, (iii) select an inference scheme, and (iv) estimate states only and make a forecast. This is a tried and true method and is used in situations such as weather and climate forecasting ([81, 39, 40]). Neither of these approaches apply to biomedical situations that, by comparison, have a different set of constraints and problems, including: (i) the models are smaller, so they can be simulated faster and estimated faster, allowing for potentially many models to be used simultaneously; (ii) there are less data relative to the number of *unknown* parameters, so while parameter estimation is necessary [1] not all parameters can be estimated; (iii) models are not generated from first principles and their application to given individuals is potentially highly variable necessitating the use and potentially the averaging of many models; and (iv) tuning would have to be done for millions of people frequently, e.g., every patient in every ICU potentially every day, a process that is not likely to be practically possible. Because of these reasons, choosing which parameters to estimate is a significant barrier to the adoption and use of DA in biomedical situations.

3.2. Houlihan approach to ill-posedness

Here we are operating under a different situation from the more canonical DA application setting, one more heavily constrained by imperfections of data that will never disappear because the data are collected in the process of managing health instead of data collected in a controlled manner explicitly for the DA. In particular we assume: (a) we do not know the right model but we have some models we can try, (b) we do not know whether a given model is identifiable and that we do not have enough data to estimate all model parameters well anyway but that we have enough data to estimate at least one parameter, (c) for a given model, we don't know what parameters are the most useful to estimate, given that we cannot estimate all of them. Given this situation we develop a method for rank ordering which parameters to estimate, subject to features we want to capture, when we have no idea how to choose which parameters to estimate or when we must choose parameters in a more high-throughput setting where we are using many models at once.

This solution involves stacking machine learning on top of DA: machine learning is applied to *simulated* model output to select the important parameters to estimate to best synchronize the model with the data, and then we use DA restricted to estimation of the parameters chosen by machine learning. In this way, the method will scale to a high-throughput setting and can be applied to many different models with high dimensional parameter spaces more easily. And while we know that this method may not lead to a unique solution in function, parameter, or initial condition space, the set of solutions will be reduced to a workable set of solutions that allow forward progress to be made.

Conceptually, we are proposing to: (i) assume a model, (ii) simulate the model under discrete parameter variation creating a grid in parameter space for which at every point we have simulated data from the model (i.e., the instance of one attractor of the model for a given set of initial conditions at the parameter grid point), (iii) select features, e.g., the mean, of the attractors that are important for estimating the physiologic system, and then (iv) use a machine learning algorithm to identify the parameters that have the greatest impact on the features. While for the authors the geometric intuition of this method originates from bifurcation theory—we will discuss this in a later section [82]—one useful way to think about the problem is in the inverse problems context. *As was the case for the bifurcation theory context, this discussion is allegorical; we are not proposing a formal inverse problems regularization framework.* From a high level, given data, Y , and a model, \mathcal{F} with a state space x , the task is to find a *set* of parameter values, Θ_j , of which there may be many if the system is not identifiable, that minimize:

$$\|Y - \mathcal{F}(x, \theta)\|_Y^p \quad (1)$$

for some p , $p = 2$ being the commonly applied least squares minimization. The core of the identifiability issue is that, for complex models, and especially given sparse data, there may be many sets of parameters Θ_j that minimize the distance between the model and the data. In this case a goal might then be to balance the number of potential minimizing parameter sets, the number of Θ_j 's, against the distance between the model and the data via an optimization algorithm, e.g.,

$$\min_{\Theta} (w_1(\|Y - \mathcal{F}(x, \theta)\|_Y^p) + w_2(\# \{\Theta_j\})) \quad (2)$$

where the w_j 's are continuous functions. This framework, a formal regularization methodology, has many advantages, but can induce many complexities that increase rather than decrease the barrier to using data assimilation in more data-poor environments. Moreover, this relatively complex methodology may not be applicable in more high-throughput situations where, e.g., many models are used in a model averaging context. Therefore, motivated by the goal of an imperfect but practical solution, we postulate that if we carefully select the right parameters that maximize the parameter subspaces that can be explored relative to a set of desired features, we can often, effectively but imperfectly, solve the optimization problem. Effectively but not rigorously, we are *regularizing a priori*, by

selecting and reducing the parameter set to be estimated before we go about estimating the parameters given data. Given the framework above, such a solution may be well handled by a tool from sparse machine learning such as lasso [83] because it uniquely rank-orders parameters by their predictive power, but it is easy to imagine using other methods. But, it is important to be clear that we are hypothesizing that the parameter subspaces that allow maximal exploration of dynamics relative to a given feature, e.g., the mean, will contain sets of parameters, Θ_j that also find relatively good minima of Eq. 1. In our evaluation we will see cases where this hypothesis fails, but we will also see that this hypothesis generally holds true in our data set, and this conclusion is the point of the paper.

It is important to understand that whether or not the parameter vector(s) that globally minimize Eq. 1 are contained within the parameter subset identified by the Houlihan method is *unknown*, and relative to the work in this paper, is a hypothesis. There are reasons to believe that, given a feature metric related to the cost function, e.g., a feature metric of the mean and a cost function related to mean squared error, the Houlihan method will find productive or useful local minima, but even this remains a hypothesis. Formulating a precise theorem—e.g., that the parameter vectors that minimize Eq. 1 are included in, or excluded from, the set identified by the Houlihan (or other ML-based method)—would be a lovely result but is beyond the scope of the current paper. Results addressing whether a parameter subspace identified by a method like the Houlihan is the parameter subspace that minimizes Eq. 1 is akin to, but not as strong as, proving that a method is able to identify the global minimum of Eq. 1. As is the case with many nonlinear systems, providing convergence to a global minima is a difficult task. However, because the conceptual goal here is to reduce the parameter space that needs to be estimated to get a good solution, there is considerable flexibility for nice theoretical work. For example, one could imagine identifying the “first order” parameters or the “steepest gradient” parameters that lead to a global minimum. Or one could weaken the global minimum requirement and work towards convergence within some ϵ of the global minimum, or some minimum error related to Eq. 1; e.g., when using real data, the data themselves have error and therefore Eq. 1 could be recast in light of the error in data. Analysis along these lines is beyond the scope of the current paper, but we hope this formulation and these ideas will motivate work in the future.

In short, we are assuming a problem is ill posed and a system that is likely not identifiable, and given this situation, we are trying to cope. Therefore, we are not really solving an identifiability task because we are not trying to find *the* best or most representative model that admits *unique* parameter estimates; rather we are solving a problem more akin to, but not literally, a regularization task. We are starting from a point where the problem is both brutally ill-posed and likely non-identifiable, and where investigating identifiability using analytic methods, or even many numeric-by-hand methods are intractable. In this case we are assuming there will be a few different parameter combinations that represent reasonable parameter estimates. In this situation each combination of parameter represents a hypothesis for how the system works. More importantly, the method we present here is a flexible entry-point for using data assimilation with a complex nonlinear model and data collected in an uncontrolled environment rather than directly solving an identifiability problem.

4. Data cohort

We test and evaluate the Houlihan methodology in the context of modeling and forecasting blood glucose collected in a free-living setting — via a type 2 diabetes self-management mobile application. The blood glucose and nutrition data used here were collected retrospectively from four participants, two with type 2 diabetes and two without diabetes, using custom-designed mobile applications for capturing self-monitoring data ([84]). These data are summarized in Table 1. We acquired two types of data: 1) fingerstick blood glucose measurements taken at the discretion of each of the 4 participants (roughly 3-10 times per day) and 2) estimates of carbohydrate consumption over time (roughly 1-5 meals per day) determined by a certified dietitian's analysis of the daily meal logs (with photos and descriptions) reported by each participant. The data are documented more completely in ([1, 74]) and are available on PhysioNet upon request.

5. Methods

5.1. Glucose-insulin physiologic model

The Houlihan method was conceived in the context of DA with a mechanistic model, and while it could be used in any nonlinear regression context, this paper will be restricted to the setting where we begin by projecting data onto a mechanistic dynamical system and then work to decide which parameters of that dynamical system we should estimate to represent the data. The mechanistic model is more formally either a dynamical system when time is discrete or a system of ODEs when time is continuous. Explicit versions of such systems form parameterized families of functions that are physically meaningful but generally do not satisfy nice function space properties such as completeness and are not universal approximators. The more general theory of dynamical systems can be found in many books ([25, 26, 27]), but here we will restrict our use of these details to an absolute minimum. We assume that the systems—ODEs or dynamical systems—have, for a given set of parameters and initial conditions, and invariant measure in the SRB ([85, 86, 87, 88]) sense, meaning that the system has an invariant or physical measure with respect to the expanding and neutral subspaces. In this way each state can have a probability density function associated with it denoted Λ_{x_i} where x_i delineates the state. *This invariant measure can potentially depend on the parameters and the initial conditions for a set of parameters, as well as random perturbations of the orbit [89].*

As previously noted, we want to use DA to model the glucose-insulin system of a human being. We begin with a particular mechanistic glucose-insulin model, here the ultradian model that has been detailed in [90, 24, 2, 4, 1], and has 6 states and 21 parameters; its details can be found in the Appendix A. The model has unknown identifiability properties, especially when only glucose is measured, but we have strong evidence that at least some of the model parameters and states are not identifiable ([74]). The Houlihan method rests on quantifying how the invariant densities of the *synthetic data sets* and their properties vary as parameters of the mechanistic model(s) vary. Specifically, the Houlihan method decides which parameters to estimate by varying the parameters of the ultradian model, observing how the invariant densities and their properties vary, and then using this information to select parameters to estimate by rank-ordering their importance using statistical inference or

machine learning. The synthetic data used to select parameters to estimate will be generated by solving the ultradian model using an adaptive version of Runge-Kutta, ode23 in Matlab and will consist of 10^5 simulated data points.

5.2. Stochastic filtering and inverse problems methods

We use two previously documented data assimilation formulations, an unscented Kalman filter ([91, 92, 93, 94, 95, 96]) (UKF) whose details can be found in [1] and a Metropolis-within-Gibbs Markov Chain Monte Carlo (MCMC) method whose details can be found in [74, 97]. As previously mentioned, these DA methods are used with the ultradian model ([90]) for performing the DA tasks. We only use these methods over the course of evaluating the Houlihan methods; the exact implementation of the DA methods can be found in [1, 74].

5.3. Analytical construction and intuition for throwing the Houlihan around the right parameters

While the approach we are proposing is new, the *allegorical* geometric intuition motivating this approach comes from bifurcation theory and in particular the bifurcation sets defined in the 1970's ([98]) and the analytic geometry vision of bifurcation theory and singularities in parameter space [29]. Bifurcation sets are the low-dimensional sets or manifolds that denote transition/bifurcation surfaces between topologically equivalent invariant sets, *partitioning the parameter space into a set of equivalence classes*. It is this idea of partitioning the parameter space into equivalence classes that differently impact dynamical features we care about is the key motivational insight. In our context we want to partition the parameter space by influence on some feature or set of features, denoted the *feature-metric*, of the dynamics. Feature-metrics are calculated from the time-series of the simulated model (dynamical system), e.g., a mean. We do not want to be as rigid as requiring topological equivalence as was defined in the bifurcation sets framework, or necessarily strict classes, but we do want to partition the parameter space according to how parameters influence a dynamical feature we care about. The over-arching idea is that the subsets of parameter space that have the highest influence on the feature-metric are the parameters that will be the most useful to estimate to minimize Eq. 1. And, knowing the most useful parameters to estimate provides a systematic way of choosing the parameters to estimate until the system is either identifiable or identifiable enough to be serviceable; in practice serviceable might mean that the errors are within desired tolerances, that parameter estimates are unique, or that the parameter estimates have few enough equilibria or minima that they can be made useful. To make this more precise, begin with the following terms, which are functions of a parameter vector, p .

Feature metric: the feature of the dynamical system we wish to influence, denoted $g(p)$; feature metrics are estimated from the time-series of the simulated model output and vary with parameter variation.

Influence: the amount that a parameter influences the feature-metric, denoted $F_i(g(p))$ for the i -th parameter.

Influence equivalence: a rule that defines equivalence of influence, e.g, all parameters i such that $a_j - F_i(g(p)) < a_{j+1}$. This allows for us to introduce a partition over influence, $\{a\}_j$

J called an *impact set*, which represents the transitions or boundaries between influence equivalence classes.

Parameter influence sets: the sets of parameters with equivalent influence according to the influence equivalence rule.

Demonstrative example: Begin by defining the dynamical system f with state variables x_j and parameters p_j assuming at least four parameters. Next define the *feature-metric* as the mean of a single state variable x_* , μ_{x_*} (i.e. we are interested in how each parameter “influences” the state’s mean). Set the *influence function* to be the absolute linear correlation, $|\beta_i|$, between the feature-metric, μ_{x_*} , and values of the parameter p_i . In this example, the influence function is a vector-valued function, with a scalar metric (linear correlation between parameter and the state’s mean it induces) corresponding to each parameter. The influence per parameter defines a probability mass function (PMF) with support $[0, 1]$ with values $\frac{|\beta_i|}{\sum_j |\beta_j|}$. Finally, we define *influence equivalence* as membership in a given quartile of the PMF defined by the influence function. Note that the *impact set* is defined by the PMF quartiles, and the *influence sets* are the parameters in respective quartiles of the PMF. Depending on the separation observed in the impact sets, we could ultimately choose to estimate parameters only from the upper equivalence class(es); i.e. the set of parameters with $|\beta|$ in the upper quartile. \square

This example takes a narrow interpretation of the flexible construct we develop for identifying equivalence classes of parameter influence. However, even the above example allows for wild topological variation within a given equivalence class. For example, within a given equivalence class one would easily imagine there being many topologically distinct invariant sets due to *both* parameter variation and initial condition variation. Presumably there are other similar equivalence class violations such as ergodicity properties ([99, 100]), k – *LCE* stability ([101, 102]), etc. These issues can all be addressed by defining the various properties, e.g., the influence function, differently, or more restrictively such that we end up with increasingly more restrictive constructions such as the original notion of bifurcation sets. This flexibility in equivalencies is the point of this construction: we can, depending on our goals, data, etc., have substantial flexibility in how we set up how to choose what parameters to estimate all while explicitly acknowledging *what we know we do not know we are preserving*. For example, if we define the feature-metric to be the mean, we know we are allowing the system to explore or have many different coexisting invariant densities as long as they have a mean that lies within a given equivalence class.

Visual example: Figure 1 shows three potential equivalence class outcomes of the Houlihan analysis. The left-most plot in fig. 1 shows the case where the rank-ordering of influence is calculated on a by-coordinate basis; meaning, the equivalence classes include collections of entire coordinates. Here in the left-most plot in fig. 1 there are two equivalence classes each with a single member corresponding to the horizontal and vertical coordinates. The middle plot in fig. 1 shows a case where the influence can consist of subsets or portions of coordinates. In this example there are thresholds or points above and below which the equivalence classes are defined. In this case the definition of the

equivalence class for each coordinate does not depend on the other coordinates; meaning that there is no co-coordinate-dependence of equivalence class definitions. The right-most plot in fig. 1 demonstrates an example where the influence equivalence includes joint coordinate relationships. In this paper we will only address the first of these cases, leaving the more complex situations for later work.

5.4. Computational moving parts for throwing the Houlihan around the right parameters

The computational task of selecting parameters to estimate involves defining the equivalence-like classes, finding their boundaries, rank ordering the parameters by importance and has, broadly, five moving parts. *First*, select the feature-metric(s), $g(p)$, e.g., mean. *Second*, formulate the representation of the space of parameters and their variation, including (i) parameter grid resolution, (ii) parameter perturbation range, (iii) parameter variation type, e.g., joint versus individual by-parameter parameter variation. *Third*, choose an influence function that defines how to model the parameterized variation of the feature-metric variation with parameter variation. *Fourth*, choose a method for rank ordering these parameterizations by influence. Sometimes steps three and four can be done using a single method, e.g., linear regression with a L_1 regularization or by using lasso with cross validation, and sometimes it is done in two steps, e.g., linear regression with a threshold on the β 's, partitioning the β 's into equivalence classes. And *fifth*, decide which parameters to keep or which equivalence classes, or which impact sets are important.

Feature metrics.—We use two feature metrics, mean and standard deviation of the invariant density generated by mechanistic model with set parameter values and initial conditions.

Parameter grid.—We begin with the nominal parameters ([90, 24, 1]), and then vary them in intervals of \log_2 over 10 decades in both directions. For example, for parameter i the parameter grid point for the k^{th} decade was set as $p_i(\text{nominal})2^k$. We did not consider joint-variation of parameters, but varied parameters independently while holding all other parameters fixed at their nominal values.

5.4.1. Parameter selection methods: Influence functions, impact sets, and ranking—Given a feature metric as a covariate or input vector, e.g., the means of attractor densities for a set of parameter values, we use several methods for selecting the best set of parameters to estimate in a DA context. Some of these methods are stock—linear regression with lasso—some are standard practice—parameter selection using knowledge of the model—and some are modifications of existing methods—see PCA-lariat below. We will see that the method for selecting the parameters matters, although not as much as the feature metric, and it is clear that sophisticated machine learning methods could be useful in this context.

Covariates or input vectors.: All of the methods below take a covariate matrix as input. The covariates correspond to vectors: one dimension of the covariate matrix corresponds to a feature metric calculated at every point along the parameter variation, e.g., the mean of a simulated attractor at every point along a one-dimensional parameter curve.

By hand selection parameter selection — parameter selection using knowledge. In our previous work we selected parameters to estimate by hand as they were tied to certain dynamical features, physiologic knowledge want to fit something in particular to solve a problem, e.g., phenotyping. We selected E and V_p because they seemed to have an impact on the mean ([4]) and t_p because it was related to liver function; the results can be found in [1].

Automatic parameter selection using linear regression. A basic method for determining influence is the linear dependence between the feature metric and parameter variation. In this setting we perform a linear regression between the feature metric and the parameters and we keep all β s for which $\beta_i > (\beta_1)(\kappa_{LR})$. Here we set $\kappa_{LR} = 20\%$ or 0.2, meaning that we keep all the parameters that have a regression coefficient that explains at least 20% of the regression coefficient with the highest influence.

Automatic parameter selection using Lasso and cross validation. A natural way of reducing the number of parameters in a model is to select parameters that have a lot of power explaining the feature metric while simultaneously being non-redundant. One way of achieving this is to use lasso, or L_1 regularization to enforce a sparse representation of the parameter system ([83, 103, 104, 105]). We use the standard lasso formulation ([83]) with cross validation to determine the rank-ordering of parameters; the optimal value of λ , or the optimal number of parameters, is set using a cutoff of one standard error. Lasso automatically and uniquely rank orders parameters. We keep the parameters within one standard error of the minimum mean squared error (MSE) ensuring a sparse representation of the model.

Automatic parameter selection using elastic net approximation of ridge regression. In addition to lasso regularization, we also use ridge regression, or L_2 regularization ([83, 105, 106]). We compute the ridge regression selected parameters using an elastic nets formulation with α set to 0.0001 where elastic nets formulation approaches L_2 regularization, and select the number of parameters using cross validation in the same way as is done in the lasso setting. We keep the parameters within one standard error of the minimum mean squared error (MSE) ensuring a sparse representation of the model.

Automatic parameter selection using PCA-lariat with a single metric. To add diversity to the set of methods for selecting parameters beyond linear regression-based methods, we devised a principle component analysis (PCA) ([107, 108, 109]) based algorithm for computing an influence function, then implement a rank-ordering scheme for defining influence equivalence. The method we develop, *PCA-lariat*, follows seven steps. *First*, estimate the PCs for the feature-metric, $g(p)$, taking care to de-trend the summary. *Second*, estimate the percentage of the variance captured by the i -th PC, $\sigma_{PC}(i)$. *Third*, identify the *important PCs*, or the PCs that explain variance above a threshold, κ_{PC} ; we use 5%. *Fourth*, for each important PC, rank-order the contribution of each parameter or coordinate to the PC. *Fifth*, collect all the coordinates for all the important PCs that contribute proportionally to a given PC above a set threshold, κ_C , $PC(i) > \kappa_C$; we use 10%. *Sixth*, for the important parameters for the important PCs, estimate the contribution per parameter:

$$PCR(i) = \sum_j \sigma_{PC}(j) * PC_j(i). \quad (3)$$

And *seventh*, rank order the important parameters by *PCR* and select the parameters above a given threshold, κ_j ; we use 0.1, or 10%.

Multi-directional parameter wrangling.: Combining models, or model averaging can be very useful for improving results ([110, 111, 112, 1]), especially when you either know you want to adjust to multiple feature-metrics, or you do not know what feature metrics are important. Here, we only consider using set operations over methods, and consider three cases. First, we take the union of: (number of rank-ordered parameters, feature-metric, influence function) using one parameter per influence function, two feature metrics, mean and standard deviation. Second and third, we take the union of: (number of rank ordered parameters, feature-metric, influence function) using one parameter per influence function and one feature metric, either mean or standard deviation.

5.5. Evaluation scheme

The evaluation of the Houlihan methods is done in four steps. *First*, we apply the Houlihan methods to the ultradian model to select parameters to estimate and compare the parameter selections as the method is perturbed. *Second*, we use both the UKF and the MCMC DA methods to estimate these Houlihan-selected parameters for the four people in our cohort and calculate the mean squared error (MSE) between the data and the model state estimates (MCMC methods) and forecasts (UKF methods). *Third*, we use both the UKF and the MCMC DA methods to estimate parameters for *both* parameters that were previously chosen by hand in previously published work ([1]) and parameters that the Houlihan methods determined were low-influence parameters and again calculate the MSE between the data and model state estimates and forecasts. *Fourth*, we compare the MSE for the variously selected parameter sets.

6. Results

The results come in two stages. *First*, we present the rank-ordered parameters selected by different methods in order to demonstrate: (i) which parameters the methods selected, (ii) that the methods selected some but not all parameters, (iii) how the parameter selection varied across methods, and (iv) the rank-ordering of parameters by method. *Second*, we evaluate the methods by using the parameters selected in each method to forecast glucose with the UKF and smooth glucose with MCMC; methods are compared via the MSE between measurements and predictions.

6.1. Parameter selections by method

Table 2 shows the rank-ordered parameters selected by each parameter selection method. The methods were sensitive to the feature metric; the mean and standard deviation-based methods did not select the same parameters as important.

For a given feature metric, all selection methods identified the same top two parameters — all methods ranked a_1 and C_1 as the top influencers of the mean, and ranked R_g and C_3 as the top influencers of the standard deviation. However, the entire influence sets differed substantially. This indicates that influence set structure, as defined (upper quartile of influence), is sensitive to choices of influence functions and influence equivalence definitions.

Interestingly, the equivalence classes of high and low parameter influence are preserved under perturbations to the influence function. Fig. 2 shows how the I_1 , I_2 and PCA-based methods rank-order parameters according to how they influence the mean. While lasso is expected to preserve the ordering with different λ (it fits one-at-a-time), ridge regression also remains robust to variations in the regularization term, λ , adding parameters one at a time.

Most methods find only 5 – 6 influential parameters out of 21, greatly reducing the dimension of the parameter space. In all cases, the methods gave an entry point for which parameters to begin estimating; the next question, then, is whether using the Houlihan approach helps to reduce forecasting errors and improve convergence of parameter estimates.

Redundancy and influence.—Our goal is to select parameters to estimate during forecasting and smoothing tasks. We aim to facilitate this goal by identifying small parameter sets that have significant, minimally redundant influence over important dynamical features. Accomplishing this can minimize problems in identifiability, multiple coexisting invariant sets, etc. Fig. 3 visualizes variation of the feature-metric, mean and standard deviation of the invariant density with parameter variation, as well as how the methods partitioned parameters into a high and low-influence equivalence class. It is clear that some variations in some parameters create large shifts in the mean and variance (e.g. a_1), whereas the mean and variance features are far less sensitive to other parameters, like E and t_d .

While the mean and standard deviation are not always influenced by the same parameters, the methods select parameters that have both high influence and relatively orthogonal influence; e.g., in the case of the mean the methods generally select a_1 and C_1 , both parameters managing insulin secretion whereas the standard deviation-based methods often select R_g that controls insulin *dependent* glucose use. The low influence parameters, by comparison, are not able to move the mean or standard deviation appreciably and are therefore not able to fully explore the space. Similarly, the low influence parameters are relatively redundant. Following this logic one might predict that estimating alpha and C_2 would lead to the most accurate model estimates while estimating E and t_d would lead to the least accurate model estimates.

Comparison with by-hand selection.—In our previous work ([1]) we selected parameters to estimate by hand based on our desire to estimate certain parameters related to physiologic function, e.g., t_p , and because of their obvious influence on parameters, e.g., V_p as could be deduced from other previous work ([4]) to influence the mean state. The

automated methods selected V_p and t_p as high influence parameters, but not E , a parameter the methods determined was a low-influence parameter.

6.2. Parameter selection method evaluation

To evaluate the effectiveness of the machine-selected parameters compared to low-influence parameters as characterized by their β_i 's, and the by-hand-selected parameters we used in our previous work, we compare the mean squared error (MSE) between the data and the forecasts for the various parameter combinations as shown in table 3. Fig 4 provides a visual summary of the results in table 3—the plots are calculated directly from table 3—for the MCMC smoothing setting, and demonstrates that all Houlihan-based parameter sets (of any size) noticeably out-performed the by-hand and low-influence parameter sets. Moreover, we see that most Houlihan-based methods achieve similar overall accuracy for parameter sets of cardinality ≤ 3 . In addition, Houlihan-based methods that selected parameter sets with 4 or more parameters achieved the best performance, and there is a general trend of improved fit with more parameters—this contrasts sharply with the by-hand parameter selections, whose performance tapered with more than 3 parameters (probably due to unforeseen issues of identifiability).

In particular, lasso chose parameters with the lowest MSE between forecasts and measurements in 7 of 8 cases. In one case, taking the union over methods shared the same MSE with lasso. And, in one case, the lowest MSE was observed with a pair of low-influence parameters. In this case it was the parameter-pair combination, α with E , that mattered. This result implies that generally low influence parameters may, for some people, be physiologically important and explore particular pathophysiology necessary to synchronize to the individual. We also know that as the number of parameters increased to 3, some of the MCMC parameter estimates with the lowest MSE found multiple, competing equilibria, were not unique, and sometimes did not fully converge. For example, Fig. 5 shows parameter estimates of two different parameters—one that converges and one that does not—for two parameter sets for P1 with standard deviation as the feature-metric. When lasso-selected parameters are restricted to two parameters for P1, then both parameters, R_g and C_3 converge producing a MSE of 600; C_3 is shown in Fig. 5. In contrast, lasso restricted to the one standard error minimum selects eight parameters, has a lower MSE of 375 but at least one of the parameters, t_p , does not converge well as shown in Fig. 5. This means that as we increased the flexibility, we lowered the MSEs but possibly came at the expense of physiology or convergent parameter estimates.

7. Discussion

Summary

Our most broad conclusion is that the machine-selected parameters work better than hand-selected parameters and that the Houlihan methods are a scalable method for selecting which parameters of a mechanistic model to estimate using DA methods. This means that stacking machine learning techniques on top of, or together with, DA is a helpful strategy, especially when models are complex and data are sparse, as in our glucose modeling example.

Houlihan methods

We intuitively define the Houlihan method(s) as a collection of methods for selecting the most productive model parameters to estimate with machine learning techniques applied to simulated model output under parameter variation subject to a set of features, e.g., the mean of a state.

Feature metric selection matters:

For all methods, the feature metric (mean or standard deviation) was the first-order driver of differences in parameter rank orderings. This implies two key questions, one applied and one abstract. The applied question involves how to make an appropriate choice of feature-metric for a given application, problem or to capture a given desired dynamical feature. The more abstract question addresses how to quantify the impact, variability, or dependence of a given feature-metric selection on the parameters the Houlihan method identifies as important—in other words, variance or invariance of the feature-metric relative to the Houlihan-selected parameters.

Regarding the applied problem, choosing a metric is highly problem-dependent. In some biomedical applications, sensitivity of the mean to parameter perturbation is not especially important for a good fit; e.g., there are physiologic systems where variation in the mean across people is small, but excursions, peaks, number of peaks, location of peaks, dynamical features related to invariant measures, etc., may be a more important types of features to capture. The way we constructed the Houlihan method allows for a great deal of flexibility in the metric choice for this reason. One can imagine developing feature metrics based on many features of the model such as spectral properties of the system, e.g., a feature metric based on the L^2 (Euclidean) distance from a baseline power spectrum, the proportion of power found at certain frequencies, Lyapunov exponent magnitudes or numbers, types or features of the invariant structures, e.g., fixed points, as well as broad ergodic type averages, etc. Given the wide possible choices, anchoring the metric to explicit clinical decision-making or to a scientific hypothesis and its evaluation will help guide the feature metric choice.

The diversity available for choosing a feature-metric motivates some very natural theoretical/mathematical questions related to sensitivity to the metric choice and necessary conditions for global and/or unique or even satisfactory convergence. For example, we might want to quantify how sensitive—here we did this by comparing the mean and standard deviation as feature-metrics—a Houlihan identified parameter set is to a given feature-metric, whether there is there an equivalence between classes of or individual feature-metrics, and a more ambitious goal of identifying whether there are certain classes of feature-metrics might be universal or yield the same results via some sort of ergodic-like hypothesis, etc. Similarly, one could imagine placing constraints on the model spaces; e.g., compactness, such that all metric of a type converge to the same answer. These types of questions follow in the lineage of dynamical systems and mathematics in general.

The cutoff matters:

The cutoff for influence has a substantial impact on the ability to estimate parameters. For example, lasso-selected parameters usually minimized MSE, but the induced MSE and MCMC convergence were both sensitive to the influence cutoff. All the methods had this sensitivity, and estimating optimal cutoffs automatically would be beneficial.

The selection method sometimes matters:

For the high-ranked parameter choices, the feature-metric was the primary difference between selected parameter sets. However, as the number of parameters included was increased, the methods diverged. We suspect that as the complexity of feature metrics and ranking methods increases, e.g., using nonlinear regressions, there will be more sensitivity of the parameter selections to the methods.

Physiology matters:

We know from carefully considering the convergence properties of the MCMC chains that some of the lowest MSEs for the runs with three or more parameters didn't converge well. Meaning, as we increased the flexibility, we lowered the MSE but possibly at the expense of physiologic fidelity or convergent parameter estimates. For pure forecasting applications this may or may not matter, but when we want the parameters to be meaningful, we need the parameter estimates to converge, not necessarily to a unique set of parameters, but to distinctly different parameter estimates that can be treated as hypotheses. Another problem that can arise because of physiology is that different people with different physiology can be sensitive to different parameters. For example, the physiological feature that is important to personalize the model for a particular person may not be related to the properties captured by the feature metric, e.g., the mean, and in this circumstance parameters identified as low influence relative to the feature-metric will not be estimated. A potential example of this is P1, for whom estimating α and E achieved the lowest MSE despite E and α being low-influence parameters relative to both the mean and standard deviation.

Effective parameter space exploration:

Abstractly, a mechanistic model is a parameterized family of functions whose parameters, depending on the model, have varying degrees of independence. From this perspective, the goal of the Houlihan methodology is to find a way to explore the maximal amount of the parameter space while minimizing the redundancy between parameters. The feature-metrics and the influence functions define which subsets of the parameter space are most useful to open for exploration, which in turn defines which dynamics can be explored. For example, focusing on variations of the mean may close off other dynamical features such as amplitude variations or any feature that is not uniquely defined by the variation of the mean. We do not yet have a good method for understanding how a feature-metric may influence other, potentially valuable explorations. We acknowledge that understanding and quantifying how limited feature-metrics influence the effective parameter space of a model is an important, unexplored problem.

Incorporating knowledge into the Houlihan method:

There are several ways—both systematic and ad hoc—that knowledge can be incorporated into the Houlihan framework; here we will discuss four. First, the selection of the *feature metric* or the feature by which rank-ordering of parameters is calculated can include a great deal of knowledge. The feature metric is somewhat akin to the creation of a cost function, although it differs from a cost function in that it is not explicitly optimized; this is a systematic method for incorporating knowledge. A second systematic way to incorporate external knowledge into the Houlihan method is to incorporate constraints into the ML machinery, e.g., by limiting the space over which parameters can vary or how parameters must co-vary according to a functional relationship. A third and more ad hoc but powerful way—a way we often use-of incorporating external knowledge is to apply the Houlihan rank ordering to select parameters and then include or keep any additional parameters known to be important; e.g., parameters that are dynamically important such as bifurcation parameters or parameters that are known to be physiologically meaningful for a given application. And a fourth ad hoc way of incorporating knowledge into the Houlihan method is to apply the Houlihan method to a limited set of parameters selected because they are useful to estimate for some reason, e.g., because external knowledge has determined that the other parameters can be held fixed because they cannot change in a meaningful way. From these examples it is clear that there are likely other ways of incorporating external knowledge into the Houlihan workflow.

Computational complexity:

We consider only the case here where we vary any one single parameter while leaving all other parameters fixed at their nominal values; this means that the dimension of the input for regression used to select the most useful parameters scales linearly in the number of parameters. If we were to co-vary parameters, meaning if varied all parameters at once, depending on how one choose to partition the parameter space, the computational complexity would explode. In this way, the framework we present here does not solve the computational complexity problem of exploring parameter space. Instead, the results in this paper show that even by only considering feature-metric variation along one-dimensional subspaces of the full parameter space we can gain substantial insight into which parameters have the most impact on the features we are interested in approximating. Moreover, we can also see the limitations of this approach — we do observe synergy between parameters where combinations of some low-influence parameters for some people can end up having a high influence on the model fit.

Obvious extensions:

In this paper, we stack machine learning on top of DA, which has many potential extensions; here we identify six directions we find particularly compelling in order of increasing difficulty. *First*, feature-metrics could be generalized to be multi-dimensional both over states and over types of feature-metrics. *Second*, feature selection methods could be developed or employed to select feature metrics from among many. *Third*, estimates of influence could be calculated to include jointly varying parameters—this would be computationally expensive and would require computational innovation in high-dimensional

settings, e.g., the computational complexity discussion above. Moreover, this problem is not necessarily a simple extension because the parameter spaces of mechanistic models are not likely to form a basis for the model space, in contrast to the parameters of the space of polynomials which do form a basis. Of course this lack of a basis structure is part of the problem—parameters of mechanistic models and likely the physiology they represent are redundant, likely for biological reasons such as robustness. *Fourth*, we use linear regression and PCA-based machine learning methods; it is likely that more sophisticated machine learning methods e.g., full elastic nets, support vector machines which are a direct generalization of both lasso and ridge regularizations [106], deep learning, sparse machine learning (compressed sensing), Bayesian methods, model averaging and ensemble learning could all be used and would likely improve the parameter selections. However, it does seem that methods that preform both regularization and parameter selection are particularly helpful, e.g., lasso regression. *Fifth*, further stacking of machine learning techniques on top of the Houlihan methods would likely be productive. For example, greedy, Gibbs-sampling-like rotation between sets of parameters that are identifiable to explore different subsets of the parameter space could minimize both model errors and identifiability issues. And *sixth*, feature-metrics could be made substantially more sophisticated, insightful and tailored to circumstance or physiologic knowledge, such as preserving power in certain frequency bands. More sophisticated feature metrics could also be used to gain insight into potentially meaningful constraints on parameters for use in operational DA.

8. Conclusion

We devised a methodology for rank-ordering parameters of a mechanistic model and using this rank-ordering to select an effective subset of parameters to estimate when projecting biomedical data onto the model via data assimilation. This methodology specifically targets parameter sets that avoid issues of model identifiability and parameter-estimation convergence problems, improving forecasting and phenotyping performance of data assimilation methods that use mechanistic biological models. Using machine learning to select parameters to estimate worked: the machine-chosen parameters reduced the mean-squared error between estimates and forecasts and data in nearly all cases by factors as large as three. These results imply that combining mechanistic and non-mechanistic machine learning could be a particularly productive direction of future research and could greatly aid in our ability to use computational machinery to both help deepen our physiologic understanding and help clinicians achieve more positive outcomes in clinical settings.

Acknowledgements

We would like to acknowledge the helpful comments from three anonymous reviewers.

We acknowledge financial support from NIH RO1 LM012734 “Mechanistic machine learning” and LM006910 “Discovering and applying knowledge in clinical databases.”

Appendix A.: Ultradian model

The model is comprised of a set of six ordinary differential equations; the model is non-autonomous because it has an external, time-dependent driver, consumed nutrition. The six

dimensional state space made up of three physiologic variables and a three stage filter. The physiologic state variables are the glucose concentration G , the plasma insulin concentration I_p and the interstitial insulin concentration I_i . The three stage filter (h_1, h_2, h_3) which reflects the response of the plasma insulin to glucose levels [90]. The model was designed to capture ultradian oscillations missing in previous models. The ordinary differential equations that define the model are [24]:

$$\frac{dI_p}{dt} = f_1(G) - E \left(\frac{I_p}{V_p} - \frac{I_i}{V_i} \right) - \frac{I_p}{t_p} \quad (\text{A.1})$$

$$\frac{dI_i}{dt} = E \left(\frac{I_p}{V_p} - \frac{I_i}{V_i} \right) - \frac{I_i}{t_i} \quad (\text{A.2})$$

$$\frac{dG}{dt} = f_4(h_3) + I_G(t) - f_2(G) - f_3(I_i)G \quad (\text{A.3})$$

$$\frac{dh_1}{dt} = \frac{1}{t_d} (I_p - h_1) \quad (\text{A.4})$$

$$\frac{dh_2}{dt} = \frac{1}{t_d} (h_1 - h_2) \quad (\text{A.5})$$

$$\frac{dh_3}{dt} = \frac{1}{t_d} (h_2 - h_3) \quad (\text{A.6})$$

The state variables include physiologic processes that have been parameterized, including: $f_1(G)$ represents the rate of insulin production; $f_2(G)$ represents insulin-independent glucose utilization; $f_3(I_i)G$ represents insulin-dependent glucose utilization; $f_4(h_3)$ represents delayed insulin-dependent glucose utilization. These functions are defined by:+

$$f_1(G) = \frac{R_m}{1 + \exp\left(\frac{-G}{v_g c_1} + a_1\right)} \quad (\text{A.7})$$

$$f_2(G) = U_b(1 - \exp(\frac{-G}{C_2 V_g})) \quad (\text{A.8})$$

$$f_3(I_i) = \frac{1}{C_3 V_g} (U_0 + \frac{U_m - U_0}{1 + (\kappa I_i)^{-\beta}}) \quad (\text{A.9})$$

$$f_4(h_3) = \frac{R_g}{1 + \exp(\alpha(\frac{h_3}{C_5 V_p} - 1))} \quad (\text{A.10})$$

$$\kappa = \frac{1}{C_4} (\frac{1}{V_i} - \frac{1}{E t_i}) \quad (\text{A.11})$$

The nutritional driver of the model $I_G(t)$ is defined over N discrete nutrition events [4], where k is the decay constant and event j occurs at time t_j with carbohydrate quantity m_j

$$I_G(t) = \sum_{j=1}^N \frac{m_j k}{60} \exp(k(t_j - t)) \mathbb{1}_{\{t_j < t\}} \quad (\text{A.12})$$

Table A.4.

Full list of parameters for the ultradian glucose-insulin model [24]. Note that IIGU and IDGU denote insulin-independent glucose utilization and insulin-dependent glucose utilization, respectively.

Ultradian model parameters		
Name	Nominal Value	Meaning
V_p	3 l	plasma volume
V_i	11 l	interstitial volume
V_g	10 l	glucose space
E	0.2 l min ⁻¹	exchange rate for insulin between remote and plasma compartments
t_p	6 min	time constant for plasma insulin degradation (via kidney and liver filtering)
t_i	100 min	time constant for remote insulin degradation (via muscle and adipose tissue)
t_d	12 min	delay between plasma insulin and glucose production
k	0.5 min ⁻¹	rate of decayed appearance of ingested glucose
R_m	209 mU min ⁻¹	linear constant affecting insulin secretion

Ultradian model parameters		
Name	Nominal Value	Meaning
a_1	6.6	exponential constant affecting insulin secretion
C_1	300 mg l ⁻¹	exponential constant affecting insulin secretion
C_2	144 mg l ⁻¹	exponential constant affecting IIGU
C_3	100 mg l ⁻¹	linear constant affecting IDGU
C_4	80 mU l ⁻¹	factor affecting IDG
C_5	26 mU l ⁻¹	exponential constant affecting IDGU
U_b	72 mg min ⁻¹	linear constant affecting IIGU
U_0	4 mg min ⁻¹	linear constant affecting IDGU
U_m	94 mg min ⁻¹	linear constant affecting IDGU
R_g	180 mg min ⁻¹	linear constant affecting IDGU
α	7.5	exponential constant affecting IDGU
β	1.772	exponent affecting IDGU

References

- [1]. Albers D, Levine M, Gluckman B, Ginsberg H, Hripcsak G, Mamykina L, Personalized, nutrition-based glucose forecasting using data assimilation, *PLoS Comp Bio* 13 (4) (2017) e1005232.
- [2]. Albers D, Hripcsak G, Schmidt M, Population physiology: leveraging electronic health record data to understand human endocrine dynamics, *PLoS One* 7 (2012) e480058.
- [3]. Xu Y, Xu Y, Saria S, A non-parametric bayesian approach for estimating treatment-response curves from sparse time series, in: Doshi-Velez F, Fackler J, Kale D, Wallace B, Weins J (Eds.), *Proceedings of the 1st Machine Learning for Healthcare Conference*, Vol. 56 of *Proceedings of Machine Learning Research*, PMLR, Northeastern University, Boston, MA, USA, 2016, pp. 282–300. URL <http://proceedings.mlr.press/v56/Xu16.html>
- [4]. Albers D, Elhadad N, Tabak E, Perotte A, Hripcsak G, Dynamical phenotyping: Using temporal analysis of clinically collected physiologic data to stratify populations, *PLoS One* 6 (2014) e96443.
- [5]. Hripcsak G, Albers DJ, High-fidelity phenotyping: richness and freedom from bias, *Journal of the American Medical Informatics Association* (2017) ocx110.
- [6]. Pathak J, Kho A, Denny J, Electronic health records-driven phenotyping: challenges, recent advances, and perspectives, *J Am Med Inform Assoc.* 20 (2013) e206–11. [PubMed: 24302669]
- [7]. Pivovarov R, Perotte A, Grave E, Angiolillo J, Wiggins C, Elhadad N, Learning probabilistic phenotypes from heterogeneous ehr data, *Journal of Biomedical Informatics*.
- [8]. Halpern Y, Choi Y, Horng S, Sontag D, Electronic medical record phenotyping using the anchor and learn framework, *JAMIA*.
- [9]. Rajkomar A, Oren E, Chen K, Dai AM, Hajaj N, Liu PJ, Liu X, Sun M, Sundberg P, Yee H, Zhang K, Duggan GE, Flores G, Hardt M, Irvine J, Le QV, Litsch K, Marcus J, Mossin A, Tansuwan J, Wang D, Wexler J, Wilson J, Ludwig D, Volchenboum SL, Chou K, Pearson M, Madabushi S, Shah NH, Butte AJ, Howell M, Cui C, Corrado G, Dean J, Scalable and accurate deep learning for electronic health records, *CoRR abs/1801.07860*. arXiv:1801.07860. URL <http://arxiv.org/abs/1801.07860>
- [10]. Westwick DT, Kearney RE, Identification of nonlinear physiological systems, *IEEE Engineering in Medicine and Biology*, 2003.
- [11]. Ljung L, *System Identification*, Prentice Hall, 1987.
- [12]. Schoukens J, Vaes M, Pintelon R, Linear system identification in a nonlinear setting, *IEEE Control Systems* (2016) 38–69.

- [13]. Levine M, Albers D, Hripcsak G, Comparing lagged linear correlation, lagged regression, granger causality, and vector autoregression for uncovering associations in ehr data, in: Annual Symposium Proceedings, AMIA, 2016.
- [14]. Levine M, Albers D, Hripcsak G, Methodological variations in lagged regression for detecting physiologic drug effects in ehr data, arXiv:1801.08929 (2018).
- [15]. Hripcsak G, Albers D, Perotte A, Exploiting time in electronic health record correlations, JAMIA 18 (2011) 109–115.
- [16]. Hripcsak G, Albers D, Correlating electronic health record concepts with healthcare process events, JAMIA 0 (2013) 1–8.
- [17]. Perotte A, Ranganath R, Hirsch J, Blei D, Elhadad N, Risk prediction for chronic kidney disease progression using heterogeneous electronic health record data and time series analysis, JAMIA 22 (4) (2015) 872–880. [PubMed: 25896647]
- [18]. Ranganath R, Perotte A, Elhadad N, Blei D, Deep survival analysis, in: Proceedings of Machine Learning for Healthcare, Vol. 56, 2016.
- [19]. Lasko T, Denny J, Levy M, Computational phenotype discovery using unsupervised feature learning over noisy, sparse, and irregular clinical data, PLOS One.
- [20]. Hornik K, Stinchcombe M, White H, “Multilayer Feedforward Networks are Universal Approximators”, Neural Networks 2 (1989) 359–366.
- [21]. Hornik K, Stinchcombe M, White H, “Universal Approximation of an Unknown Mapping and its Derivatives Using Multilayer Feedforward Networks”, Neural Networks 3 (1990)551.
- [22]. Hripcsak G, Albers D, Next-generation phenotyping of electronic health records, JAMIA 10 (2012) 1–5.
- [23]. Pivovarov R, Albers D, Sepulveda J, Elhadad N, Identifying and mitigating biases in ehr laboratory tests, Journal of Biomedical Informatics.
- [24]. Keener J, Sneyd J, Mathematical physiology II: Systems physiology, Springer, 2008.
- [25]. Brin M, Stuck G, Introduction to Dynamical Systems, Cambridge University Press, 2004.
- [26]. Guckenheimer J, Holmes P, Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields, Springer-Verlag, New York, 1983.
- [27]. Arrowsmith DK, Place CM, An introduction to dynamical systems, Cambridge University Press, 1990.
- [28]. Arnold VI, Ordinary differential equations, Springer-Verlag, 1992.
- [29]. Arnold V, Geometric methods in the theory of ordinary differential equations, 2nd Edition, Grundlehren de mathematischen Wissenschaften, Springer-Verlag, 1983.
- [30]. Albers D, Levine L, Gluckman B, Hripcsak G, Mamykina L, Stuart A, Mechanistic machine learning: how data assimilation leverages physiologic knowledge using bayesian inference to forecast the future, infer the past, and phenotype, in revision (2018).
- [31]. Jazwinski A, Stochastic processes and Filtering Theory, Dover, 1998.
- [32]. Lorenc A, Analysis methods for numerical weather prediction, Q. J. R. Meteorol. Soc 112 (1988) 1177–1194.
- [33]. Law K, Stuart A, Zygalakis K, Data assimilation, Springer, 2015.
- [34]. Ash M, Bocquet M, Nodet M, Data assimilation: methods, algorithms and applications, SIAM, 2016.
- [35]. Reich S, Cotter C, Probabilistic forecasting and Bayesian data assimilation, Cambridge University Press, 2015.
- [36]. Haug A, Bayesian estimation and tracking, Wiley, 2012.
- [37]. Ristic B, Arulampalam S, Gordon N, Beyond the Kalman filter: particle filters for tracking and applications, Artech house, 2004.
- [38]. Candy J, Bayesian signal processing: classical, modern, and particle filtering methods, Wiley, 2009.
- [39]. Evensen G, Data assimilation, the ensemble kalman filter, Springer, 2009
- [40]. Evensen G, The ensemble kalman filter: theoretical formulation and practical implementation, Ocean Dynamics.

- [41]. Stuart A, Inverse problems: A bayesian perspective, *Acta Numerica* 19 (2010) 451–559.
- [42]. Zenker S, Rubin J, Clermont G, From inverse problems in mathematical physiology to quantitative differential diagnoses, *PLoS Comput Biol* 3 (11).
- [43]. Banks HT, Hu S, , Thompson WC, Modeling and inverse problems in the presence of uncertainty, CRC Press, 2014.
- [44]. Hirata Y, Bruchovsky N, Aihara K, Development of a mathematical model that predicts the outcome of hormone therapy for prostate cancer, *Journal of Theoretical Biology*, 264 (2) (2010) 517–527. [PubMed: 20176032]
- [45]. Hirata Y, de Bernardo M, Bruchovsky N, Aihara K, Hybrid optimal scheduling for intermittent androgen suppression of prostate cancer, *CHAOS* 20 (2010) 0451251.
- [46]. Schiff S, Neural control engineering: The emerging intersection between control theory and neuroscience, MIT Press, 2011.
- [47]. Dukic V, Lopes H, Polson N, Tracking epidemics with google flu trends data and a state-space seir model, *Journal of the American Statistical Association* 107 (2012) 1410–1426.
- [48]. Miao H, Xia X, Perelson A, Wu H, On identifiability of nonlinear ode models and applications in viral dynamics, *SIAM Review* 53 (2011) 3–39. [PubMed: 21785515]
- [49]. Chee F, Fernando T, Closed-loop control of blood glucose, Springer, 2007.
- [50]. P PB, Orsini M, Benedetti MM, The artificial pancreas, *Artif Cells Blood Substit Immobil Biotechnol.* 2 (2003) 127–38.
- [51]. Fabietti P, Canonico V, Orsini-Federici M, Sarti E, Massi-Benedetti M, Clinical validation of a new control-oriented model of insulin and glucose dynamics in subjects with type 1 diabetes, *Diabetes Technol Ther.* 4 (2007) 327–338.
- [52]. Kovatchev B, Breton M, Man C, Cobelli C, In silico preclinical trials: a proof of concept in closed-loop control of type 1 diabetes, *J Diabetes Sci Technol* 3 (2009) 44–55. [PubMed: 19444330]
- [53]. Coutinho MAFAB, Azevedo RS, Burattini MN, Lopenz LF, Massad E, Vaccination against rubella: analysis of the temporal evolution of the age-dependent force of infection and the effects of different contact patterns, *Phys. Rev. E* 67 (2003) 051907.
- [54]. Mirowski M, Reid P, Mower M, Watkins L, Gott V, Schauble J, Langer A, Heilman M, Kolenik S, Fischell R, Weisfeldt M, Termination of malignant ventricular arrhythmias with an implanted automatic defibrillator in human beings, *New England Journal of Medicine* 303 (1980)322–324. [PubMed: 6991948]
- [55]. Thabit H, Tauschman M, Allen J, Leelarathna L, Hartnell S, Wilinska M, Acerini C, Dellweg S, Benesch C, Heinemann L, Made J, Holzer M, Kojzar H, Exall J, Yong J, Pichierri J, Barnard K, Kollman C, Cheng P, Hindmarsh P, Campbell F, Arnolds S, Pieber T, Evans M, Dunger D, Hovorka R., Home use of an artificial beta cell in type 1 diabetes, *New England Journal of Medicine*.
- [56]. Leelarathna L, English SW, Thabit H, Caldwell K, Allen JM, Kumareswaran K, Wilinska ME, Nodale M, Mangat J, Evans ML, Burnstein R, Hovorka R, Feasibility of fully automated closed-loop glucose control using continuous subcutaneous glucose measurements in critical illness: a randomized controlled trial, *Crit Care* 17 (4) (2013) R159. [PubMed: 23883613]
- [57]. Cobelli C, Man C, Sparacino G, Magni L, De Nicolao G, Kovatchev B, Diabetes: models, signals, and control, *IEEE Reviews in Biomedical Engineering* 2 (2009) 54–96. [PubMed: 20936056]
- [58]. Thabit H, Tauschman M, Allen JM, Leelarathna L, Hartnell S, Wilinska ME, Acerini CL, Dellweg S, Benesch C, Heinemann L, Made JK, Holzer M, Kojzar H, Exall J, Yong J, Pichierri J, Barnard KD, Kollman C, Cheng P, Hindmarsh PC, Campbell FM, Arnolds S, Pieber TR, Evans ML, Dunger DB, Hovorka R, Home use of an artificial beta cell in type 1 diabetes, *N. Engl. J. Med.*
- [59]. Glass L, Courtemanche M, Cardiac Arrhythmias and Device Therapy: Results and Perspectives for the New Century, *Futura*, 2000, Ch. Control of atrial fibrillation: A theoretical perspective, pp. 87–94.
- [60]. Christini D, Glass L, Mapping and control of complex cardiac arrhythmias, *Chaos* 12 (2002) 732–739. [PubMed: 12779601]

- [61]. Hall K, Christini DJ, Tremblay M, Collins JJ, Glass L, Billette J, Dynamic control of cardiac alternans, *Phys. Rev. Lett* 78 (1997) 4938.
- [62]. Mackey M, Glass L, Oscillation and chaos in physiological control systems, *Science* 197 (1977) 287–289. [PubMed: 267326]
- [63]. Parker RS, III FJD, Ward JF, Peppas NA, Robust h-infinity glucose control in diabetes using a physiological mode, *AIChE Journal* 46 (12)(20xx) 2537–2549.
- [64]. Parker RS, III FJD, , N. A. Peppas, The intravenous route to blood glucose control. a review of control algorithms for noninvasive monitoring and regulation in type i diabetic patients, *IEEE Engineering in Medicine and Biology* 20 (1) (2001) 65–73.
- [65]. Donnet S, Samson A, A review on estimation of stochastic differential equations for pharmacokinetic/pharmacodynamic models, *Advanced Drug Delivery Reviews*.
- [66]. Bonate P, Recommended reading in population pharmacokinetic pharmacodynamics, *AAPS*.
- [67]. Sadean M, Glass P, Pharmacokinetic-pharmacodynamic modeling of anesthesia, intensive care, and pain medicine, *Curr Opin Anaesthesiol*.
- [68]. Kristensen NR, Madsen H, Ingwersen SH, Using stochastic differential equations for pk/pd model development, *XXXX* 32(1).
- [69]. Sedigh-Sarvestan M, Albers D, Gluckman B, Data assimilation of glucose dynamics for use in the intensive care unit, in: 34th Annual International IEEE EMBS Conference, 2012.
- [70]. Selgrade JF, Harris LA, Pasteur RD, A model for hormonal control of the menstrual cycle: Structural consistency but sensitivity with regard to data, *Journal of Theoretical Biology* 260 (2009) 572–580. [PubMed: 19560471]
- [71]. Sedigh-Sarvestani M, Schiff SJ, Gluckman BJ, Reconstructing mammalian sleep dynamics with data assimilation, *PLoS Comput Biol* 8(11).
- [72]. Lin J, Chase JG, Shaw GM, Doran CV, Hann CE, Robertson MB, Browne PM, Lotz T, Wake GC, , Broughton B, Adaptive bolus-based set-point regulation of hyperglycemia in critical care, in: 34th Annual International IEEE EMBS Conference, 2004.
- [73]. Lin J, Razak NN, Pretty CG, Compte AL, Docherty P, Parente JD, Shaw GM, Hann CE, , Chase G, A physiological intensive control insulin-nutrition-glucose (icing) model validated in critically ill patients, *Computer Methods and Programs in Biomedicine* 102 (2011) 192–205. [PubMed: 21288592]
- [74]. Levine M, Hripesak G, Mamykina L, Stuart A, Albers D, Offline and online data assimilation for real-time blood glucose forecasting in type 2 diabetes, *arXiv:1709.00163* (2017).
- [75]. Hines E, Middendorf T, Aldrich R, Determination of parameter identifiability in nonlinear biophysical models: A bayesian approach, *J Gen. Physiol* 143 (3) (2014) 401–416. [PubMed: 24516188]
- [76]. Simens J, Cree-Green M, Bergman B, Nadeau K, Behn CD, Structural identifiability analysis of a labeled oral minimal model for quantifying hepatic insulin resistance, in: *Association of Women in Mathematics Annual Symposium*, 2017.
- [77]. Eisenberg M, Hayashi M, Determining identifiable parameter combinations using subset profiling, *Math Biosciences* 256 (2014) 116–125.
- [78]. Walch O, Eisenberg M, Parameter identifiability and identifiable combinations in generalized hodgkin-huxley models, *Neurocomputing* 199 (2016) 137–143.
- [79]. Brouwer A, Meza R, Eisenberg M, A systematic approach to determining the identifiability of multistage carcinogenesis models, *Risk analysis*.
- [80]. Chis O-T, Banga J, Balsa-Canto E, Structural identifiability of systems biology models: a critical comparison of methods, *PLoS ONE* 6 (11)(2011)e27755. [PubMed: 22132135]
- [81]. Hamill TM, Ensemble-based atmospheric data assimilation, *XXXX*.
- [82]. Kuznetsov Y, *Bifurcation Theory*, 2nd Edition, Springer-Verlag, 1998.
- [83]. Hastie T, Tibshirani R, Wainwright M, *Statistical learning with sparsity*, CRC, 2015.
- [84]. Mamykina L, Levine M, Davidson P, Smaldone A, Elhadad N, Albers D, Reasoning about personally generated nutritional data in diabetes with information technologies, *J Am Med Inform Asso*.

- [85]. Young L-S, What are SRB measures, and which dynamical systems have them?, *J. Stat. Phys* 108 (2002) 733–754.
- [86]. Ruelle D, A measure associated with axiom A attractors, *Amer. J. Math* 98 (1976) 619–654.
- [87]. Sinai Y, Gibbs measures in ergodic theory, *Russian Math. Surveys* 27 (1972) 21–69.
- [88]. Bowen R, Equilibrium states and the ergodic theory of Anosov diffeomorphisms, Vol. 470 of *Lect. Notes in Math.*, Springer-Verlag, Berlin, 1975.
- [89]. Lin K, Young L, Shear-induced chaos, *Nonlinearity* 21 (2008) 899–922.
- [90]. Sturis J, Polonsky KS, Mosekilde E, Cauter EV, Computer model for mechanisms underlying ultradian oscillations of insulin and glucose, *Am J Physiol Endocrinol Metab* 260 (1991) E801–E809.
- [91]. Julier S, Uhlmann J, Unscented filtering and nonlinear estimation, *Proc. IEEE* 92 (2004) 401–22.
- [92]. Julier S, Uhlmann J, Durrant-Whyte H, A new approach for filtering nonlinear systems, in: *American Control Conference, IEEE, IEEE, 1995.*
- [93]. Wan E, Merwe R, The unscented kalman filter for nonlinear estimation, in: *Adaptive Systems for Signal Processing, Communications, and Control Symposium, IEEE, Wiley, 2000, pp. 153–158.*
- [94]. Gove J, Hollinger D, Application of a dual unscented Kalman for simultaneous state and parameter estimation problems of surface-atmosphere exchange, *J Geophys Res* 111 (2006) D08S07.
- [95]. Wan E, Merwe R, Nelson A, Dual estimation and the unscented transformation, in: *NEURAL INFORMATION PROCESSING SYSTEMS, MIT Press, 2000, pp. 666–672.*
- [96]. Wan EA, Merwe RVD, The unscented kalman filter, in: *Kalman Filtering and Neural Networks, Wiley, 2001, pp. 221–280.*
- [97]. Cotter SL, Roberts GO, Stuart AM, White D, Mcmc methods for functions: Modifying old algorithms to make them faster, *Statist. Sci* 28 (3) (2013)424–446.
- [98]. Sotomayor J, Structural stability and bifurcation theory, in: *Dynamical Systems, 1973, pp. 549–560.*
- [99]. Pugh C, Shub M, Ergodic attractors, *Trans. Amer. Math. Soc* 312 (1989) 1–54.
- [100]. Burns K, Wilkinson A, On the ergodicity of partially hyperbolic systems, *Annals of Mathematics.*
- [101]. Albers DJ, Sprott JC, Structural stability and hyperbolicity violation in high-dimensional dynamical systems, *Nonlinearity* 19 (2006) 1801–1847.
- [102]. Albers DJ, Sprott JC, Crutchfield JP, Persistent chaos in high dimensions, *Phys. Rev. E* 74 (2006) 057201.
- [103]. Tibshirani R, Regression shrinkage and selection via lasso, *Journal of the Royal Statistical Society, Series B* 58 (1) (1996) 267–288.
- [104]. Breiman L, Better subset regression using the nonnegative garrote, *Technometrics* 37 (4)(1995) 373–84.
- [105]. Zou H, Hastie T, Regularization and variable selection via the elastic net, *Journal of the Royal Statistical Society, Series B* (2005) 301–320.
- [106]. Jaggi M, *An Equivalence between the Lasso and Support Vector Machines.*, Chapman and Hall/CRC, 2014.
- [107]. Pearson K, On lines and planes of closest fit to systems of points in space, *Philosophical Magazine* 2 (1901) 559–572.
- [108]. Hotelling H, The most predictable criterion, *J. Ed. Psych*
- [109]. Jolliffe I, *Principal Component Analysis*, Springer, 2010.
- [110]. Hoeting J, Madigan D, Raftery A, Volinsky C, Bayesian model averaging: A tutorial, *Statistical Science* 14 (1999) 382–417.
- [111]. Raftery A, Gneiting T, Balabdaoui F, Polakowski M, Using bayesian model averaging to calibrate forecast ensembles, *Monthly Weather Review* 133 (2005) 1155–1174.
- [112]. Claeskens G, Hjort N, *Model selection and model averaging*, Cambridge University Press, 2008.
- [113]. Hastie T, Tibshirani R, *Statistical learning with sparsity: the lasso and generalizations*, Chapman and Hall, 2015.

Highlights

- Data assimilation using sparse data and complex models with many parameters can lead to non-unique or non-convergent parameter estimates.
- When identifiability failure arises it can be difficult to decide which parameters to estimate from among the 10s to 100s of potential parameters.
- The parameter Houlihan is a framework for selecting which parameters to estimate using data assimilation in the context of sparse data and identifiability failure to minimize non-uniqueness of parameter estimates and error.




Figure 1. Shown are three different Houlihan constructions: left shows equivalence class by coordinate —this is the construction we use in this paper; middle shows equivalence by subsets of coordinates but retains the non-joint parameter dependency assumption; right shows a fully joint equivalence where combinations of parameters can generate influence when individual parameters do not, similar to the notion of bifurcation sets.

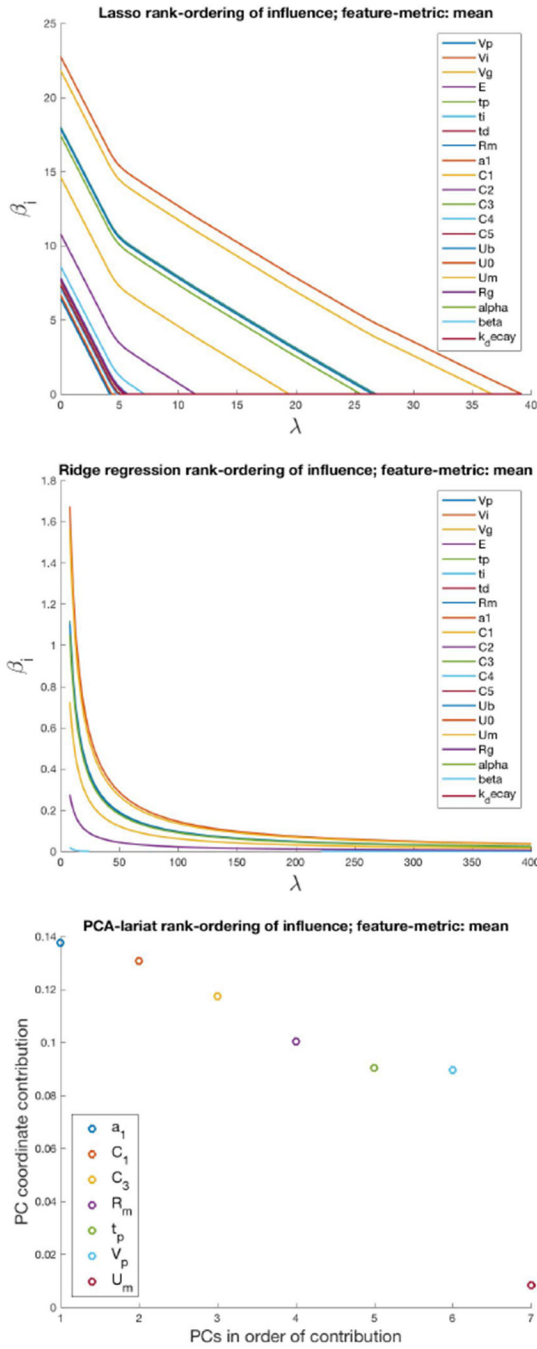


Figure 2. The rank-ordered influence function with a feature-metric set to the mean for lasso, ridge regression, and PCA-lariat methods. Note that β_i delineates the i^{th} parameter of the lasso/ridge regression and λ is the Lagrange multiplier associated with the lasso/ridge regularization where $\lambda = 0$ is the un-regularized regression. As λ increases the regularization is increased, decreasing the magnitude of all β_i for $i > 1$ while increasing the influence of β_0 [113].




Figure 3. The influence for two feature metrics, mean and standard deviation, versus parameter variation for high impact and low impact parameters.




Figure 4. The overall performance of each method in the smoothing setting. The vertical axis indicates the %-optimal MSE for a given method, averaged over the four patient data sets. Note that methods are labeled as blue to red, where the minimally-performing methods are blue and the maximally-performing methods are red. The plots are estimated directly from the information in Table 3.

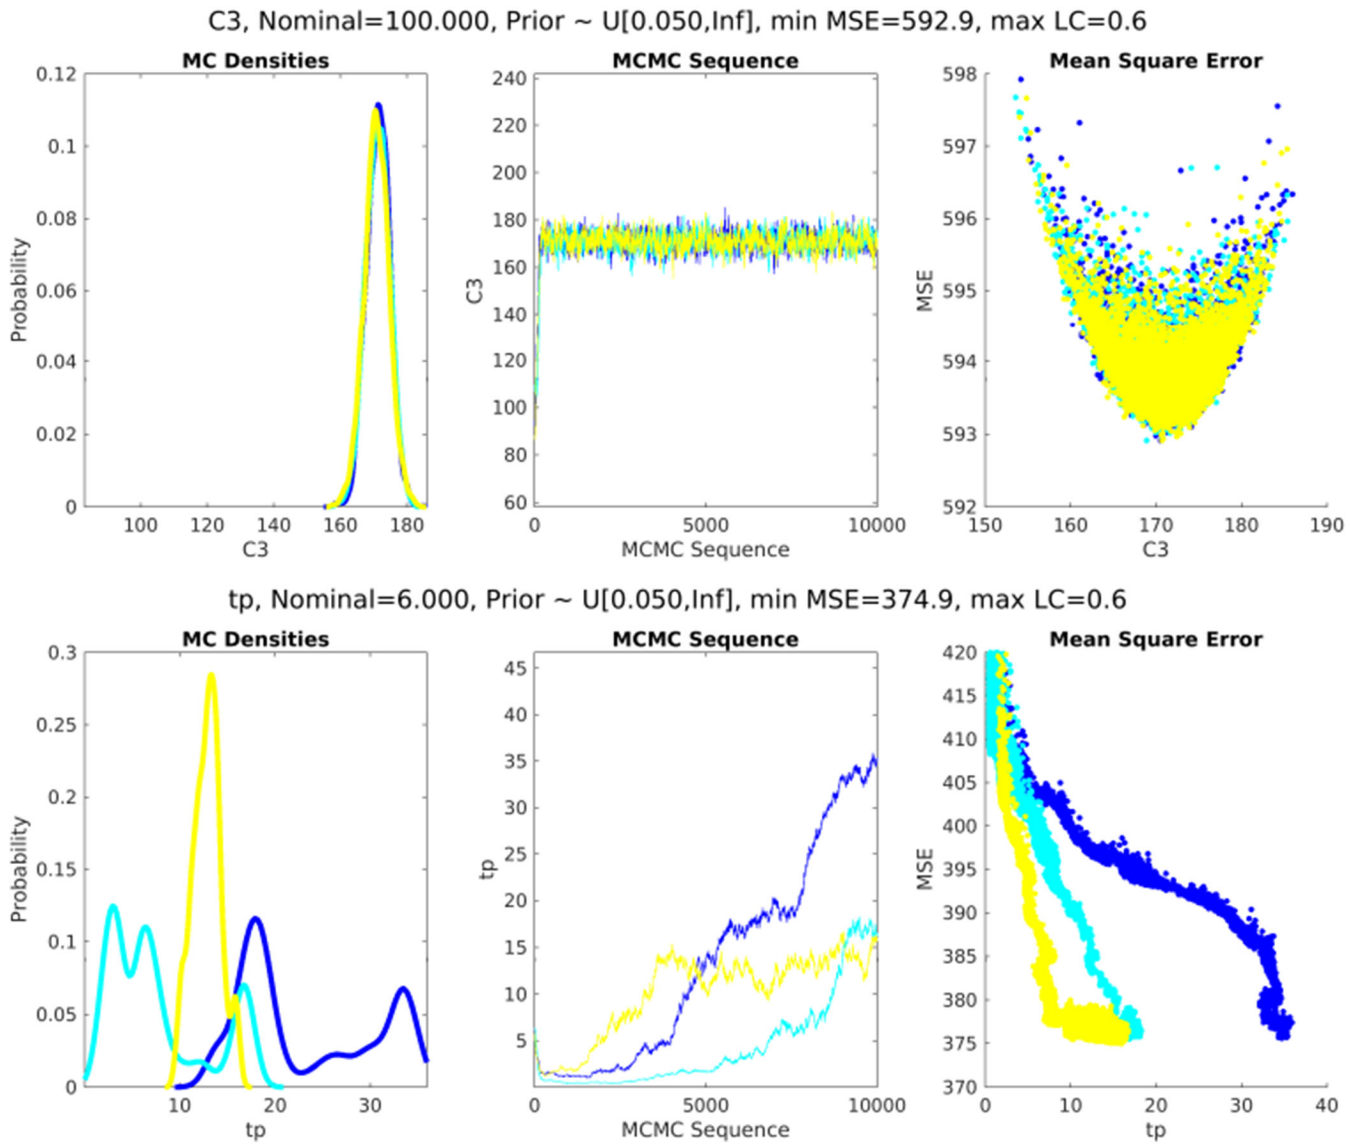


Figure 5. The posterior densities, Markov chains, and MSE surfaces, for two parameters taken from two sets of parameters. The top set of plots shows C_3 estimates for P1 where lasso is allowed to select two parameters with standard deviation set as the feature metric; C_3 converges well. The bottom set of plots show t_p estimates for P1 for lasso-selected parameters at one standard error minimum—eight parameters are selected in this case—with standard deviation set as the feature metric; t_p does not converge to a unique minimum but has a lower MSE than cases where the parameters are uniquely identified.

Table 1.

Demographic information and summary statistics are reported for the four participants whose retrospectively collected data are included in the study.

Data Summary				
Participant ID	P1	P2	P4	P5
Age	40 – 50	40 – 50	40 – 50	40 – 50
Disease Status	T2D	T2D	No Diabetes	No Diabetes
Medications	metformin	metformin	—	—
Total # glucose measurements	304	211	520	322
Total # meals recorded	124	76	370	184
Total # days measured	16	16	53	52
Mean measured glucose	113 ± 25	127 ± 32	92 ± 17	101 ± 16

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2.

The rank ordering choice of the four parameter selection methods for the feature-metrics mean, μ , and standard deviation, σ .

Rank-ordered parameters per selection method out of 21 possible parameters											
method	1	2	3	4	5	6	7	8	9	10	11
LASSO μ	a_1	C_1	V_p	t_p	R_m	C_3	—	—	—	—	—
LASSO σ	R_g	C_3	U_m	a_1	C_1	t_p	R_m	V_p	—	—	—
Linear regression μ	a_1	C_1	C_3	R_m	t_p	V_p	U_m	R_g	C_4	U_b	U_0
Linear regression σ	R_g	C_3	U_m	a_1	C_1	R_m	V_p	t_p	k_{decay}	—	—
Ridge regression μ	a_1	—	—	—	—	—	—	—	—	—	—
Ridget regression σ	R_g	—	—	—	—	—	—	—	—	—	—
PCA μ	a_1	C_1	C_3	R_m	t_p	V_p	U_m	—	—	—	—
PCA σ	R_g	C_3	U_m	a_1	C_1	—	—	—	—	—	—

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 3.

The mean squared error (MSE) between forecast/smoothed and measured glucose. The machine-based methods, almost always selected the parameter set that achieved the MSE minimum, but for some individuals, certain hand-chosen parameters matter.

Rank-ordered parameters per selection method									
parameter	MSE for MCMC				MSE for UKF				method-feature-metric pairs
	P1	P2	P4	P5	P1	P2	P4	P5	
a_1	822	1140	338	296	809	1270	304	356	LASSO(μ), LR(μ), PCA(μ)
R_g	655	1180	475	288	672	1490	470	401	LASSO(σ), LR(σ), ridge(σ), PCA(σ)
t_p	807	1020	448	349	788	1050	407	420	by-hand, high-influence
V_p	820	1120	332	320	805	1300	313	362	by-hand, high-influence
E	681	1250	655	500	721	1380	704	724	by-hand, low-influence
α	501	1250	526	346	526	1580	528	394	low-influence
t_d	530	1080	730	674	NaN	1260	NaN	480	low-influence
Rank-ordered parameter pairs per selection method									
(a_1, C_1)	570	1080	285	276	698	1290	258	330	LASSO(μ), LR(μ), PCA(μ)
(R_g, C_3)	593	923	210	297	613	1260	215	385	LASSO(σ), LR(σ), ridge(σ), PCA(σ)
(a_1, R_g)	578	1130	292	296	614	1400	269	343	Union of rank 1 over methods
(α, E)	454	1174	518	345	483	1310	535	520	low-influence
(α, t_d)	432	993	525	347	NaN	1120	NaN	NaN	low-influence
(E, t_d)	462	1030	592	487	643	1190	NaN	490	low-influence
Rank-ordered parameter 3-tuple per selection method									
(a_1, C_1, V_p)	569	1060	284	276	663	1310	260	329	LASSO(μ) (1 st)
(a_1, C_1, t_p)	518	864	247	275	NaN	NaN	234	294	LASSO(μ) (2 nd)
(R_g, C_3, U_m)	590	922	190	294	618	1140	228	391	LASSO(σ), LR(σ), PCA(σ)
(a_1, C_1, C_3)	431	1020	261	274	1330	1110	251	340	LR(μ) PCA(μ)
(C_2, E, α)	442	1020	518	346	479	1250	535	515	low-influence
(t_d, C_2, α)	432	894	525	347	NaN	1190	NaN	NaN	low-influence
(t_d, E, α)	398	956	479	343	NaN	1120	NaN	520	low-influence
(t_d, E, C_2)	464	941	592	489	630	1190	NaN	NaN	low-influence
Rank-ordered parameter 4-tuple per selection method									
(a_1, R_g, C_1, C_3)	398	864	182	288	649	985	265	324	Union of rank 2 over methods
Full Houlihan for μ and σ									
$(a_1, C_1, V_p, t_p, R_m, C_3)$	414	862	217	229	661	NaN	236	291	Lasso(μ)

Rank-ordered parameters per selection method									
	MSE for MCMC				MSE for UKF				
parameter	P1	P2	P4	P5	P1	P2	P4	P5	method-feature-metric pairs
$(R_g, C_3, U_m, a_1, C_1, t_p, R_m, V_p)$	375	863	182	231	632	942	224	289	Lasso(σ)
Method with the lowest MSE									
	Lasso	Lasso	Lasso/Union	Lasso	low-influence	Lasso	Lasso	Lasso	

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript