

Multi-Task Gaussian Processes and Dilated Convolutional Networks for Reconstruction of Reproductive Hormonal Dynamics

Iñigo Urteaga*

*Department of Applied Physics and Applied Mathematics
Columbia University, New York City, NY, USA*

INIGO.URTEAGA@COLUMBIA.EDU

Tristan Bertin*

*Department of Biomedical Informatics
Columbia University, New York City, NY, USA*

TMB2183@COLUMBIA.EDU

Theresa M. Hardy

*Rory Meyers College of Nursing
New York University, New York City, NY, USA*

TH89@NYU.EDU

David J. Albers

*Division of Informatics, Department of Pediatrics
University of Colorado, Anschutz Medical Center, Denver, CO, USA*

DAVID.ALBERS@UCDENVER.EDU

Noémie Elhadad

*Department of Biomedical Informatics
Columbia University, New York City, NY, USA*

NOEMIE.ELHADAD@COLUMBIA.EDU

Abstract

We present an end-to-end statistical framework for personalized, accurate, and minimally invasive modeling of female reproductive hormonal patterns. Reconstructing and forecasting the evolution of hormonal dynamics is a challenging task, but a critical one to improve general understanding of the menstrual cycle and personalized detection of potential health issues. Our goal is to infer and forecast individual hormone daily levels over time, while accommodating pragmatic and minimally invasive measurement settings. To that end, our approach combines the power of probabilistic generative models (i.e., multi-task Gaussian processes) with the flexibility of neural networks (i.e., a dilated convolutional architecture) to learn complex temporal mappings. To attain accurate hormone level reconstruction with as little data as possible, we propose a sampling mechanism for optimal reconstruction accuracy with limited sampling budget. Our results show the validity of our proposed hormonal dynamic modeling framework, as it provides accurate predictive performance across different realistic sampling budgets and outperforms baselines methods.

1. Introduction

We propose and validate a statistical framework for personalized, accurate, and minimally invasive modeling of female reproductive hormonal patterns. Starting from sparse hormone

* Contributed equally

measurements at few and specific days, we aim to infer and predict daily hormone measurements over time. This is a challenging task, because there is much variation from one individual to another, both in the timing of the different phases of the menstrual cycle, and of the daily hormone levels during each phase (Treloar et al., 1967; Alliende, 2002). As such, models that can adjust to individuals’ patterns are needed; at the same time, population-level dynamics can also provide additional guidance towards our goal. Previous work was limited to modeling reproductive hormone levels of a given individual through independent Gaussian processes, and explored the impact of hormone measurement sampling frequency and sampling times on the ability to detect menstrual phases for that individual (Urteaga et al., 2017). Here, we present a novel framework which departs from previous work in several ways, provides further understanding of the dynamics of the menstrual cycle, and produces generalizable insights for reconstruction and prediction of continuous temporal signals.

1.1. Clinical Relevance

The menstrual cycle reflects underlying female reproductive hormonal function. It can be described according to its length: i.e., number of days from the first day of bleeding until the day before the next bleeding period (Hornsby et al., 2007; Small et al., 2007), its phases (e.g. follicular, luteal), and its daily reproductive hormonal dynamics, e.g. estrogen (E), progesterone (P), inhibin (Ih), follicle stimulating hormone (FSH), and luteinizing hormone (LH). The menstrual cycle is a powerful indicator of overall health in women (Small et al., 2007; Vassena et al., 2014). Beyond its relevance to reproductive health and fertility (Filiberto et al., 2012; Vassena et al., 2014), cycle characteristics can inform risk for chronic diseases such as cardiovascular disease, osteoporosis and cancer (Bedford et al., 2010; Zittermann et al., 2000; Solomon et al., 2002; Shuster et al., 2010; Mahoney, 2010). Further, knowledge of an individual’s full hormone dynamics can pinpoint hormonal irregularities and indicate potential diagnoses (Carmina and Lobo, 1999; Giudice, 2010; Tworoger et al., 2013). In fact, the medical field has suggested menstruation should be considered “the fifth vital sign” (ACOG, 2015; Bobel, 2019).

Despite this realization, reproductive female physiology and the exact relationship between hormone levels and hormonal dynamics are still enigmatic in many ways. Due to challenges associated with timed hormone sample collection and the cost of repeated sampling (Filiberto et al., 2012), continuous hormone measurement across the menstrual cycle has not been conducted on a large scale.

Currently, real-world population-level data is mostly available for cycle length as recorded through menstrual trackers. For a small subset of these populations, some phase information is also available through reports of ovulation test results. While these self-reported data have been validated (Hornsby et al., 2007; Small et al., 2007), there are no real-world, large-scale datasets of hormonal dynamics throughout the menstrual cycle. Rather, our current understanding of continuous reproductive hormone levels is based on small-scale studies, e.g. 30 women in McLachlan et al. (1990), and mechanistic models validated against these small datasets (Selgrade and Schlosser, 1999). How to infer the reproductive hormonal dynamics of a given individual in a minimally-invasive, low-cost fashion is an open research question.

In this work, we contribute to **knowledge of the menstrual cycle** through real-world, self-tracked data obtained through a popular menstrual tracking app for menstrual cycle length and ovulation. This mitigates the potential limitations of mechanistic models which have been validated in small-scale studies only. Further, we provide and validate a method to **reconstruct and predict an individual’s menstrual cycle** (i.e., daily hormonal measurements, and thus phases and length) based on a few hormone measurements only. This data-driven approach to reconstructing an individual’s cycle patterns can help with minimally-invasive, low-cost data collection at a massive scale.

1.2. Technical Significance

Mechanistic models of hormonal levels have been proposed in the literature (Selgrade and Schlosser, 1999; Clark et al., 2003; Selgrade et al., 2009). These constitute a system of non-linear differential equations that describe the dynamics of the cycle, their physiology, and the inter-relations of the different reproductive hormones, and can be tuned through a set of parameters to generate different types of hormonal cycles. However, there are a number of limitations with these models. While they can simulate realistic cycles, additional machinery is needed for the inverse problem, namely inferring the full hormonal curve from few measurements. Further, joint estimation of model states and parameters in a way that synchronizes the model to an individual, via techniques such as data assimilation (Albers et al., 2018) is challenging. Accurate inference and avoidance of identifiability problems require measuring all the hormones at multiple time points; prerequisites currently not possible on a large scale (Albers et al., 2018). Previous hormone reconstruction work, which opted for a physiology agnostic approach to inferring hormonal cycles, experimented with individual Gaussian processes per-hormone for a single individual’s data, where each hormone’s dynamics were independently predicted (Urteaga et al., 2017).

The technical significance of this work is two-fold, as we contribute both from a data and a methods perspective. From a data standpoint, **(1)** we propose a method to generate synthetic, yet realistic datasets of full female hormonal cycles. The characteristics of these cycles are diverse, but grounded on real-world evidence: we anchor the high level cycle characteristics (namely cycle length and ovulation day) to those of a large, real-world dataset of self-reports. From a modeling standpoint, our technical contributions are as follows. Since the overall aim is to devise models which necessitate as little data as possible, both in the number of individuals needed and the number of their hormonal measurements, we devise an end-to-end statistical framework for personalized modeling of female reproductive hormonal patterns under realistic sampling budgets: **(2)** We combine the power of generative processes with the ability of neural networks to learn complex temporal mappings. To this end, we first use Gaussian processes (GPs) —well adapted to the characteristics of our input, namely temporal signals with missing data, and whose outputs are distributions over time— to learn meaningful signals about individual hormonal dynamics. A dilated convolutional network is then trained to map from those uncertain, potentially noisy GP outputs, to target hormone levels through time. **(3)** To capture the interactions among the different hormones in a single individual through time (e.g., $hormone_i$ at day_t and $hormone_j$ at $day_{t'}$), we use multi-task or multi-output GPs, rather than univariate and independent GPs as experimented with in previous work. **(4)** We propose to use a non-causal dilated

convolutional neural network architecture to reconstruct personalized time-series, which yields reduced reconstruction error when compared to recurrent networks traditionally used for time series analysis. **(5)** Finally, we propose a Bayesian optimization-like approach to identify times at which to measure hormone levels that are most likely to yield accurate predictions of the proposed framework. In our set-up, where the hormonal dynamics operate at different time resolutions, yet the sampling occurs at specific times, the timing of the input samples plays a critical role on the pattern reconstruction. We describe, given a sampling budget, a way to identify the sampling times for best reconstruction of hormone levels at the population level, further understanding general hormone physiology.

A Python implementation of the proposed framework for personalized modeling of female reproductive hormonal patterns, along with an example dataset of 60 individuals, is publicly available in <https://github.com/iurteaga/hmc>.

2. Dataset

Due to the lack of a publicly available gold-standard dataset of daily female reproductive hormonal measurements throughout a cycle, a synthetic dataset is for now the only alternative to design modeling techniques. Here we describe our approach to generate a dataset, which is both **diverse** —it provides heterogeneous hormone level patterns, consistent with different types of individuals in a population— and **realistic** —it is grounded in real-world, high-level characteristics of the menstrual cycle.

We leverage the previously clinically validated mechanistic model of [Clark et al. \(2003\)](#). The model accepts a large number of parameters, each with a range of admissible values that produce heterogeneous yet realistic hormone level patterns, resulting in different cycle characteristics (e.g. cycle and phase lengths). To ensure that the simulated signals reflect a range of realistic characteristics of menstrual cycles, we turned to a real-world dataset of women tracking their menstrual cycle ([Clue, 2019](#)). While menstrual trackers do not collect specific hormone levels, users can opt to track the results of their ovulation tests in addition to their cycle length. These two variables (cycle length and ovulation day) help in turn anchor different phases of the menstrual cycle.

We select a cohort of healthy women with natural cycles (i.e., not taking any contraceptive hormonal or non-hormonal therapy) in the ages of 22-30 (when menstrual cycles are more stable during the reproductive years ([Treloar et al., 1967](#))) and who tracked their ovulation test results. Figure 1 shows the empirical density of these two variables of interest across the resulting cohort, for a total of 50,000 cycles. While it confirms common clinically-validated knowledge about the menstrual cycle (most likely cycle length is 29 days, and most likely ovulation day is 15), it also reflects admissible variations in the population ([Treloar et al., 1967](#)). Figure 1 also shows a Gaussian approximation to the ob-

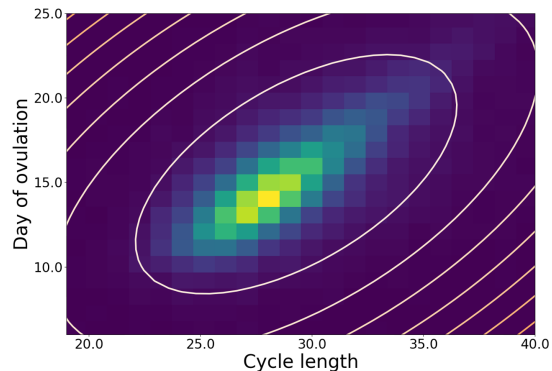


Figure 1: Density distribution of cycle length and ovulation day from a cohort of individuals aged 22-30 with natural cycles.

served ovulation-day/cycle-length empirical distribution, with mean $\mu = [15.5 \ 29.1]^T$ and covariance matrix $\Sigma = \begin{bmatrix} 25.5 & 8.0 \\ 8.0 & 12.6 \end{bmatrix}^T$.

To simulate a single individual’s hormonal dynamics, one can first sample an ovulation-day/cycle-length pair from the Gaussian approximation of the observed empirical distribution—ensuring that the underlying high-level characteristics of the simulated cycles are consistent with a real-world population—and then select a random synthetic set of simulated cycles that match these characteristics (i.e., an individual with realistic cycle features).

This proposed method can be used to simulate as large of a dataset as needed. Since our goal here is to identify a low-cost method to hormone level reconstruction, we aim for few individuals for which we would practically obtain gold-standard hormonal daily measurements. Therefore, in this work we simulate a small dataset of 60 diverse individuals.

3. Methods

The overall aim of our work is an end-to-end statistical framework for personalized, accurate, minimally invasive modeling of female reproductive hormonal dynamics. Specifically, our goals are to **(1)** infer and forecast individual hormone levels over time with high-fidelity; and **(2)** accommodate pragmatic and minimally invasive measurement settings. To attain accurate hormone level reconstruction with as little data as possible—both in the number of individuals needed for training, as well as in the number of hormonal measurements required—we devise a flexible statistical framework with two main components: a model for reconstruction of hormone levels from few measurements, and a sampling mechanism for optimal reconstruction accuracy with limited sampling budget (i.e., minimally invasive requirements).

3.1. Accurate reconstruction of hormone level dynamics

With our framework, we seek to reconstruct the daily evolution of hormone levels that regulate the menstrual cycle (i.e., *E*, *P*, *Ih*, *FSH*, and *LH*), from a few of their irregularly-observed measurements. We cast the problem of reconstructing these hormones over time as a **multivariate time-series regression task**. For our model to have clinical relevance, we must accommodate realistic hormone measurement practices: in essence, the less-invasive our assumptions are, the better. To that end, we aim at a very reduced and flexible measurement schedule, i.e., women are asked to provide hormone levels as scarcely as possible, and not necessarily subject to a regular schedule (e.g. once per week). Furthermore, all five hormone levels are attained simultaneously from the same serum-based measurement. These non-invasive, yet stringent, requirements raise several challenges for a reconstruction algorithm: available measurements will be sparse, uncertain and not uniformly sampled in time.

Figure 2 describes our end-to-end approach to reconstruct and predict daily hormone levels. For a single individual, a Multi-task Gaussian Process (MGP) learns, given a few non-uniformly sampled measurements, a personalized posterior distribution of hormone levels over time. These hormone level distributions are subsequently used as input, over equally spaced time instants, to a Dilated Convolutional Neural Network (DCNN) whose mission is to correct the potential mispredictions of each individual MGP by learning at the

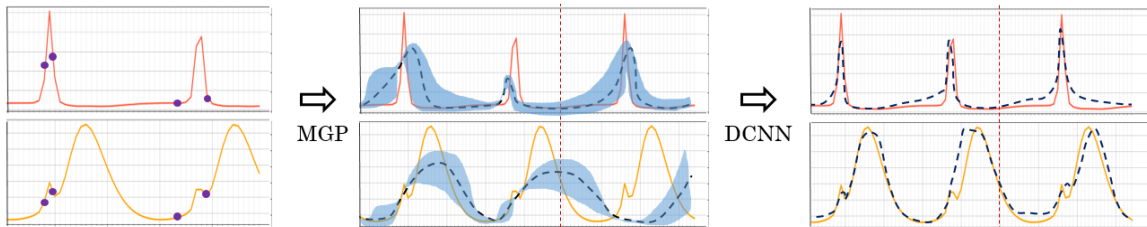


Figure 2: Reconstruction of hormone levels. A few measurements (purple dots) are sampled from the true hormonal levels (red and yellow solid lines) of the two first cycles and constitute the input of the MGP. From this subset of points, the MGP generates a probabilistic distribution (blue area) on an extended range of time (2 original cycles and one forecasted). Samples from the MGP posterior distributions for all individuals in the cohort are then fed into a population-level DCNN, whose goal is to reconstruct the hormonal levels (blue dotted line) over the three cycles.

population level. The DCNN incorporates general knowledge of hormonal dynamics, as it is trained over the set of MGP distributions for a cohort of individuals — which, to ensure our pragmatic goal of low-cost and limited measurements, is of reduced size.

3.1.1. MULTI-TASK GAUSSIAN PROCESSES (MGP) FOR INDIVIDUAL SPARSE DATA

Gaussian processes (Rasmussen and Williams, 2006) are useful for time-series modeling, because they can naturally accommodate irregular and uncertain observations, and provide probabilistic predictions. We leverage their flexible and probabilistic nature to model irregularly sampled individual sparse hormone levels. Here we extend on the previous success of GPs to model reproductive hormones (Urteaga et al., 2017) by considering multi-task Gaussian processes, which allow for an interpretable, yet more flexible, learning of the patterns across hormone levels: both in time and across hormones.

A GP is a stochastic process such that any finite collection of random variables has a multivariate Gaussian distribution. In particular, a collection of random variables $f(x) : x \in X$ is said to be drawn from a GP with mean function $m(\cdot)$ and covariance function $k(\cdot, \cdot)$ if for any finite set of elements $x_1, \dots, x_k \in X$, the associated finite set of random variables $f(x_1), \dots, f(x_k)$, follows $f(x) \sim GP(m(\cdot), k(\cdot, \cdot))$ where $m(x) = \mathbb{E}[f(x)]$ and $k(x, x') = \mathbb{E}[(f(x) - m(x))(f(x') - m(x'))]$. A GP can be seen as a probability distribution over arbitrary functions, with $m(\cdot)$ its mean function, and $k(\cdot, \cdot)$ the covariance kernel or function.

GPs have been successfully applied to the analysis of medical time-series data, e.g. (Stegle et al., 2008; Schulam and Saria, 2015). When analyzing multiple covariates over time, the MGP regression framework allows for joint modeling of multiple (irregularly sampled) time-series, as shown in several healthcare and other biomedical applications (Bonilla et al., 2008; Alvarez and Lawrence, 2009; Dürichen et al., 2014, 2015; Cheng et al., 2017). We here consider an MGP whose input is time, and the output a five dimensional vector (one dimension or task per hormone), using a multi-output kernel. Let’s consider the full set of H hormone levels over T time steps in matrix form $\mathbf{Y} \in \mathbb{R}^{H \times T}$, where $y_{h,t}$ is the hormone h level at time t , and $\mathbf{y} = \text{vec}\mathbf{Y}$ its vector counterpart. By extension of the Gaussian process principles to multivariate series, we can write the distribution of the observed hormone levels as $\mathbf{y} \sim \mathcal{N}(\mathbf{y}|\mu, \Sigma)$, with $\mu \in \mathbb{R}^{HT}$ and $\Sigma \in \mathbb{R}^{HT \times HT}$.

We propose to learn an MGP for each individual i from the inferred hormone patterns $y_i(t_i)$, sampled at (few) time-instants, i.e., $t_i \subset \{1, \dots, T\}$. Given a set of observations, we can learn the parameters θ_h of the per-hormone time kernel $k^h(t, t')$ and the multi-output matrix $K(h, h')$ for hormones $h, h' \in 1, \dots, H$, that maximize the marginal likelihood of the data. The output of the learned MGP is a posterior distribution over the hormone levels $z_i(t)$ at the evenly-spaced time index $t = \{1, \dots, T\}$, i.e., $z_i(t) \sim \text{MGP}(m_i(\cdot), k_i(\cdot, \cdot) | \theta_i)$. We focus on a model that shares covariance function forms on the time input, and uses a free-form positive semi-definite kernel as the covariance function across hormones. This modeling choice is flexible and has been successful for capturing (interpretable) inter-task dependencies, while avoiding the need for large training data (Bonilla et al., 2008). Mathematically, the full covariance matrix of size $HT \times HT$ follows $\Sigma = K(h, h') \otimes k(t, t')$, where \otimes denotes the Kronecker product, $K(h, h')$ a positive semi-definite matrix that specifies the inter-hormone dependencies, and $k(t, t')$ the covariance matrix over time-inputs. An important property of this kernel choice is that, due to the not block-diagonal structure with regards to hormones, observations of one hormone can affect the MGP predictions of another hormone. If per-hormone i.i.d. noise is included in the model, then we simply extend the covariance structure to accommodate a $H \times H$ diagonal matrix \mathbf{D} of per-task noise variance, resulting in an updated covariance structure $\mathbf{y} \sim \mathcal{N}(\mathbf{y} | \mu, \Sigma)$ where $\Sigma = K(h, h') \otimes k(t, t') + D \otimes I$.

Leveraging the recommendations of Urteaga et al. (2017), we employ the exponential periodic kernel $k(t, t') = \exp\left(\left[2 \sin^2\left(\frac{\pi}{p}|t - t'|\right)\right] / l^2\right)$, with period length p and lengthscale parameter l to be learned; as well as a low-rank approximation to the tasks matrix through the across-hormone kernel $K(h, h') = VV^\top + \text{diag}(v)$. After learning all kernel parameters (we assume a zero-mean prior function), the MGP provides a posterior distribution over unobserved levels at any given time-instant for the hormone(s) of interest. We finally emphasize that, because of our modeling choices, the kernel function of the learned MGP provides an individualized and interpretable signature: for a given individual i , $k_i(t, t')$ reflects the time-varying pattern, while $K_i(h, h')$ captures the correlations between hormone levels.

For the remainder of this work, we focus on the $H = 5$ hormones of interest: E , P , Ih , FSH , and LH . Therefore, both the observed and the inferred variables are five dimensional hormone level vectors indexed by time, i.e., $y_i(t), z_i(t) \in \mathbb{R}^5$, with θ_i denoting the mean and kernel parameters learned for individual i . We will write y_i, z_i , when it is not necessary to specify the argument (i.e., the time dependence).

3.1.2. DILATED CONVOLUTIONAL NEURAL NETWORK (DCNN)

Conventional wisdom within the deep-learning community typically considers recurrent networks for time-series analysis (Goodfellow et al., 2016). However, several recent results indicate that convolutional architectures can outperform recurrent networks on several sequence to sequence tasks such as audio synthesis, word-level language modeling, and machine translation (van den Oord et al., 2016; Kalchbrenner et al., 2016; Dauphin et al., 2016; Gehring et al., 2016, 2017). Aligned with these in spirit, and based on a recent comparison benchmark by Bai et al. (2018), we hereby consider convolutional networks for hormonal time-series reconstruction.

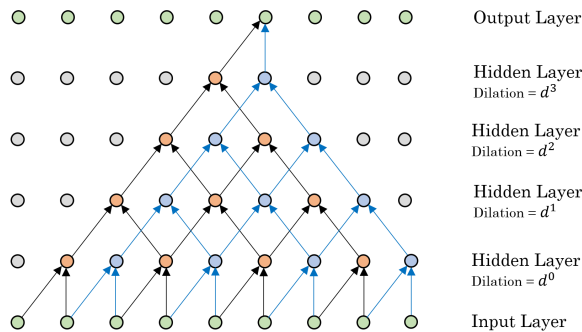


Figure 3: Dilated Convolutional Neural Network with 4 hidden layers, $d = 2$ and $K = 2$.

The key benefits on the use of convolutional structures for sequence modeling are **(1)** reduced gradient instability (i.e., exploding and vanishing gradients) than in recurrent architectures; **(2)** ability to capture long dependencies in the input sequence via dilated layers; and **(3)** a lower memory requirement in training (especially for long input sequences), since the convolution filters are shared across layers, and the backpropagation path depends only on the network depth.

Intuitively, the idea of applying CNNs to time-series analysis is to learn convolutional filters that capture the patterns in the data. A CNN consists of a sequence of convolutional layers, the output of which is connected only to local regions of the input. This is achieved by sliding a filter f — a weight matrix — over the input, and at each point computing the dot product between the two. A classic linear-causal convolution is only able to look back at a history proportional to the depth of the network, which limits its capability to accommodate long dependencies. To overcome this limitation, we leverage dilated convolutions and extend them to the non-causal setting: i.e., the filters compute the convolution over past, present and future data. Our proposed CNN architecture is based on recent convolutional networks for sequential data (van den Oord et al., 2016; Lea et al., 2016), and extends the dilated layers to the non-causal setting (Yu and Koltun, 2015; Bai et al., 2018): both past and future samples are used for the reconstruction of each time-instant. By employing non-causal dilated convolutions, we accommodate exponentially large receptive fields in our hormone level reconstruction. Note that the word *causal* here indicates, as in signal-processing, that the filter output depends only on past and present inputs. Our proposed filters, whose output also depends on future inputs, are *non-causal*.

Mathematically, for a sequence input X of size $t = \{1, \dots, T\}$ and a filter f of size K , the causal d -dilated convolution operation F on element $X[t]$ of the sequence is defined as $F[t] = (X *_d f)[t] = \sum_{k=0}^{K-1} f[k]X[t-dk]$. Dilation is equivalent to introducing a fixed step between adjacent filter taps —when $d = 1$, a dilated convolution reduces to a regular convolution. For a fixed dilation factor d , the outputs at the top level depend on an exponential range of inputs: specifically, the effective history for the hidden layer l is $d^l \times (K - 1)$. With this architecture, for a fixed number of hidden layers, there are two degrees of freedom to increase the receptive field of the network: the filter size K , and the dilation factor d . In this work, we extend the dilation architecture to be non-causal: i.e., for each output $F[t]$, the filter is applied both backward and forward in time, centered at t (see Figure 3 for an illustration).

In each hidden layer h_i of the proposed dilated, non-causal convolutional architecture, we use a ReLU non-linearity as our activation function, along with a residual connection, to further transform the output from the dilated convolution: i.e., $h_{i+1} = \text{ReLU}(h_i + F(h_i))$. The residual connections are implemented by skipping one or more layer(s), thus adding unmodified inputs from the skipped layers to the output activation layer.

3.1.3. TRAINING THE DCNN ACROSS A POPULATION OF MGP POSTERIOR DISTRIBUTIONS

The MGP is a full probabilistic model of the hormone levels of interest over time, which we leverage to train our proposed DCNN. We harness full potential of the generative model by drawing S hormone level sample streams $z_i^{(s)}, s = \{1, \dots, S\}$ from the learned MGP, instead of resorting to its expected value: i.e., we leverage the full posterior distribution, and not only its mean prediction. The advantage of this approach is two-pronged: the DCNN is trained with uncertain inputs (i.e., S drawn signals $z_i^{(s)}$ are mapped to one true output y_i), and further, the neural network is trained to learn these mappings at the population level. By mapping all these samples to the same target function—the true hormone measurements—we build a noise resistant neural network at the population level. Note that with such approach, training of a convolutional neural network is possible, even when data from few individuals is available.

Learning the DCNN given all the time-series $z_i^{(s)}, \forall(i, s)$, involves learning the parameters w of the DCNN that minimize the mean-squared error at the population level. We optimize the DCNN based on gradient descent, in which we incrementally update the weights based on the gradient of the error function: $w(\tau + 1) = w(\tau) - \eta \nabla \mathcal{L}_2(w(\tau))$ for $\tau \in \{1, \dots, \tau_{max}\}$ where τ_{max} is the maximum number of training iterations, and η is the learning rate. We compute the prediction error via the expected \mathcal{L}_2 norm: the distance between the DCNNs prediction $g(z_i)$ and the true hormone levels y_i , averaged over all individuals, i.e., $\forall i \in 1, \dots, I$. Further, since z_i is a random variable itself, we optimize the expectation of the loss function via its empirical average:

$$w^* = \underset{w}{\operatorname{argmin}} \sum_{i=1}^I \mathbb{E}_{z_i \sim \text{MGP}(\cdot|\theta_i)} \{ \mathcal{L}_2 [g(z_i, w), y_i] \} = \underset{w}{\operatorname{argmin}} \sum_{i=1}^I \sum_{s=1}^S \left[g(z_i^{(s)}, w) - y_i \right]^2 \quad (1)$$

3.2. Sampling for optimal hormone reconstruction via Expected Distance

As in any other signal processing problem, the reconstruction and prediction of hormone levels hinges upon the sampling strategy of the measurements. In the context of independent GP based hormone prediction task, [Urteaga et al. \(2017\)](#) showed that sampling times influence critically the ability to predict phases accurately, and suggest that knowledge of the timing of “peaks and valleys” of hormone levels is needed (e.g. day at which LH surge takes place for a given individual).

However, an open question remains: how to determine, given a realistic budget of measurements (i.e., a reasonable number of days for a women to collect serum levels), the sampling strategy that allows for most accurate reconstruction. In this section, we describe a method to devise, given a fixed amount of possible days, which specific days are best to collect the hormone measurements. Note that, consistent with our goal of minimally

invasive measurements, we assume that all five hormones are sampled simultaneously. As such, we do not look to optimize the best sampling scheme for each hormone, but rather consider optimal sampling for all hormones at once.

The notion of selecting a subset of points for regression was first suggested by [Silverman \(1985\)](#), and later applied via a greedy selection of points within the GP context by [Smola and Bartlett \(2001\)](#). Methods for determining accurate inducing points for sparse GP regression have been popular, e.g. pseudo-inputs in ([Snelson and Ghahramani, 2006](#)) and a more general view of probabilistic GP approximations in ([Bauer et al., 2016](#)).

However these strategies have been developed to approximate a full GP with as few datapoints as possible, so that the computational cost of the GP is reduced. Our goal differs as we are not working with big datasets nor do we want to reduce computational burden—in our setting, it is feasible to compute the full posterior of the observed hormone levels.

An alternative of interest to the task at hand is that of Bayesian optimization ([Frazier, 2018](#)), where one optimizes an objective function based on a surrogate which quantifies uncertainty via an stochastic process. By defining an appropriate acquisition function, one can then use the surrogate to decide where to sample towards the goal of maximizing the objective. Inspired by this framework in general, and by the *Expected Improvement* acquisition function in particular ([Frazier, 2018](#)), we hereby present our sampling approach for hormone reconstruction.

We propose to greedily select input points where the expected absolute difference between the true levels and the predictions of the MGP trained on the limited subset is maximized. Unlike with the expected improvement acquisition function, and more broadly any Bayesian optimization function, we are not trying to find the maximum (or minimum) of an unknown function. Instead, we greedily find the set of points that allow for a MGP with limited inputs to learn a good approximation to the true hormonal levels. That is, we find the time instant with the maximum **Expected Distance**: a function that measures the expected (over the underlying uncertainties) distance from the true hormone levels to the MGP predictions, trained over a subset of the points.

We define the distance function $\Psi_i^h(t|\mathcal{D}) = \mathbb{E} \{|y_i^h(t) - z_i^h(t|\mathcal{D})|\}$ for hormone h and individual i , where the expectation is computed over the randomness of the MGP output $z_i(t|\mathcal{D})$. We explicitly indicate that the MGP has been trained over a subset of the data $\mathcal{D} = \{y(t), t \subset \{1, \dots, T\}\}$. We find the time-instant that maximizes this expected distance with respect to the population distribution, which we approximate with the observed sample average

$$t^* = \operatorname{argmax}_t \mathbb{E} \left\{ \Psi_i^h(t|\mathcal{D}) \right\} = \operatorname{argmax}_t \sum_{i=1}^I \sum_{h=1}^H \Psi_i^h(t|\mathcal{D}), \quad (2)$$

where $\Psi_i^h(t|\mathcal{D})$ can be computed in closed form (see details below). Because at each time instant all hormones are measured simultaneously, and we consider their accuracy equally important, a third empirical average over hormones is computed. An illustrative example of the proposed optimal sampling on a single hormone is presented in [Figure 4](#).

We now derive the expectation over the MGP posterior, i.e., $\Psi_i(t) = \mathbb{E} \{|y_i(t) - z_i(t|\mathcal{D})|\}$. First, since the posterior of the MGP is parameterized by its mean $\mu_i(t|\mathcal{D})$ and covariance $k_i(t, t'|\mathcal{D})$ functions, we have for any given time instant t , $z_i(t|\mathcal{D}) = \mathcal{N}(\cdot|\mu_i(t|\mathcal{D}), \sigma_i^2(t|\mathcal{D})) =$

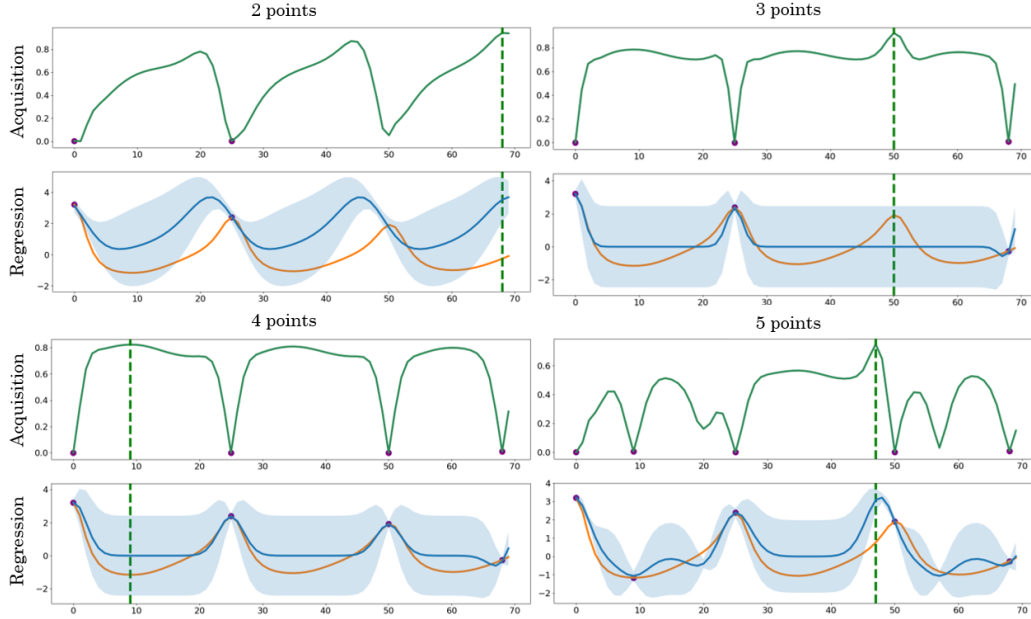


Figure 4: ED sampling with a budget of 5 days for a single hormone, starting from 2 samples: each additional sample (in purple) is identified in a greedy fashion based on the difference between, given the available sampled days, the learned posterior $z_i(\mathcal{D}_i)$ (in blue), and the true hormone dynamics (in orange), indicated by the green vertical line.

$\mu_i(t|\mathcal{D}) + \sigma_i(t|\mathcal{D})\epsilon$, where ϵ is a standard Gaussian variable $\epsilon \sim \mathcal{N}(0, 1)$. In the following derivation, and for the sake of clarity, we drop the dependency on t in all variables: $z_i = z_i(t|\mathcal{D})$, $\mu_i = \mu_i(t|\mathcal{D})$, $\sigma_i = \sigma_i(t|\mathcal{D})$, and denote with ϕ and Φ the probability density function and the cumulative density function of a standard Gaussian, respectively:

$$\begin{aligned}
\Psi_i &= \mathbb{E} \{|y_i - z_i|\} = \mathbb{E} \{|y_i - \mu_i - \sigma_i\epsilon|\} = \int_{-\infty}^{+\infty} |y_i - \mu_i - \sigma_i\epsilon| \phi(\epsilon) d\epsilon \\
&= \int_{-\infty}^{\frac{y_i - \mu_i}{\sigma_i}} (y_i - \mu_i - \sigma_i\epsilon) \phi(\epsilon) d\epsilon + \int_{\frac{y_i - \mu_i}{\sigma_i}}^{+\infty} (-y_i + \mu_i + \sigma_i\epsilon) \phi(\epsilon) d\epsilon \\
&= (y_i - \mu_i) \int_{-\infty}^{\frac{y_i - \mu_i}{\sigma_i}} \phi(\epsilon) d\epsilon - \sigma_i \int_{-\infty}^{\frac{y_i - \mu_i}{\sigma_i}} \epsilon \phi(\epsilon) d\epsilon + (-y_i + \mu_i) \int_{\frac{y_i - \mu_i}{\sigma_i}}^{+\infty} \phi(\epsilon) d\epsilon + \sigma_i \int_{\frac{y_i - \mu_i}{\sigma_i}}^{+\infty} \epsilon \phi(\epsilon) d\epsilon \\
&= (y_i - \mu_i) \Phi\left(\frac{y_i - \mu_i}{\sigma_i}\right) + \sigma_i \int_{-\infty}^{\frac{y_i - \mu_i}{\sigma_i}} \phi'(\epsilon) d\epsilon + (-y_i + \mu_i) [1 - \Phi\left(\frac{y_i - \mu_i}{\sigma_i}\right)] - \sigma_i \int_{\frac{y_i - \mu_i}{\sigma_i}}^{+\infty} \phi'(\epsilon) d\epsilon \\
&= (y_i - \mu_i) \left[2\Phi\left(\frac{y_i - \mu_i}{\sigma_i}\right) - 1 \right] + \sigma_i \left[\phi\left(\frac{y_i - \mu_i}{\sigma_i}\right) - \phi(-\infty) \right] - \sigma_i \left[\phi(+\infty) - \phi\left(\frac{y_i - \mu_i}{\sigma_i}\right) \right] \\
&= (y_i - \mu_i) \left[2\Phi\left(\frac{y_i - \mu_i}{\sigma_i}\right) - 1 \right] + 2\sigma_i \phi\left(\frac{y_i - \mu_i}{\sigma_i}\right)
\end{aligned}$$

$$\Psi_i(t|\mathcal{D}) = [y_i(t) - \mu_i(t|\mathcal{D})] \left[2\Phi \left(\frac{y_i(t) - \mu_i(t|\mathcal{D})}{\sigma_i(t|\mathcal{D})} \right) - 1 \right] + 2\sigma_i(t|\mathcal{D})\phi \left(\frac{y_i(t) - \mu_i(t|\mathcal{D})}{\sigma_i(t|\mathcal{D})} \right) \quad (3)$$

The classic tradeoff between exploration and exploitation is here easily interpretable. The first term counts for the exploitation, which is predominant when we are confident about our prediction (small σ_i), and the difference ($y_i - \mu_i$) determines our decision. On the contrary, the second term increases with the uncertainty of the MGP (high σ_i), therefore favoring exploration.

4. Evaluation

4.1. Experimental Setup

Train/test datasets: For evaluation purposes, we create a realistic synthetic dataset, that meets reasonable hormone measurement practices as well. As proposed in Section 2, we simulate a reduced set of individual hormone patterns with real-world cycle characteristics. Specifically, we sample $I = 60$ cycle lengths from the approximation to the empirical distribution in Figure 1, and select a random synthetic set of $H = 5$ hormone signals with matching features. We draw from the marginal Gaussian approximation to the cycle length $\mathcal{N}(\mu = 29.1, \sigma^2 = 12.6)$, and each associated hormone time series comprises $T = 105$ days (i.e., at least 3 full cycles are included for all individuals). Time-indexes are selected at random $t_i \subset \{1, T\}$ for each individual to simulate scarce, and non-necessarily regular measurement schedules, with different sampling budgets.

We randomly split the train/test population in $I_{train} = 50$ (with 10 individuals kept apart for validation purposes) and $I_{test} = 10$ individuals. To train the MGP, we only make use of subsampled data points in the first 2 cycles for each individual in the training set: i.e., $y_i(t_i), t_i \subset \{1, \dots, 70\}, i = \{1, \dots, I_{train}\}$. After learning each individual MGP, the DCNN is trained with the regularly spaced signals drawn from each individual MGP and the true hormone levels for the entire training population: i.e., $z_i(t), y_i(t), t = \{1, \dots, T\}, i = \{1, \dots, I_{train}\}$. For testing purposes, the same approach is replicated, where the MGP is trained over subsampled hormone levels of each individual in the test-set ($\mathcal{D}_i = \{y_i(t_i) | t_i \subset \{1, \dots, 70\}, i = \{1, \dots, I_{test}\}\}$), and the true hormone levels are only used to compute the mean squared error (MSE) across the test-population and across the five hormones: i.e., $MSE_{test} = \frac{1}{I_{test}} \sum_{i=1}^{I_{test}} (g(z_i(\mathcal{D}_i)) - y_i)^2$.

Subsampling schemes: In alignment with our aim of accommodating minimally invasive measurements (e.g. once or twice per week), we experiment with different sampling budgets of 10, 15, 25, 35 and 70 days across two cycles. Following the finding from [Urteaga et al. \(2017\)](#), we ensure that at least two LH peaks are sampled across the two cycles (corresponding to LH surge prior to ovulation). Pragmatically, these measurements are easily obtainable for women by using urine-based ovulation tests. To assess the value of our reconstruction framework, we experiment with two sampling strategies (that augment the number of measurements beyond these two pre-determined days): (i) random sampling, which selects with uniform probability a subset of days, and (ii) optimal sampling, which implements the Expected Distance (ED) based strategy described in Section 3.2.

Scaling: To facilitate the training of our MGP-DCNN model, we standardize each input-feature by removing its mean and scaling to unit variance. This scaling is applied based on the training set (i.e., the test set is scaled based on data seen only on the train set, to avoid any data leakage). This standardization of hormone levels helps the training process and allows for direct comparison of results across hormones.

Bayesian Optimization: To find the optimal configuration for the proposed DCNN, we tune its hyper parameters via Bayesian Optimization, using the Python [skopt module](#). The learning rate $\eta \in [0.5e^{-3}, 2e^{-2}]$, the number of hidden layers $L \in \{3, \dots, 6\}$, the dilation factor $d \in \{1, \dots, 3\}$, the kernel size $K \in \{2, \dots, 9\}$ (same for all hidden layers), and the number of filters (in range 5 to 12) per layer are optimized in training.

Baselines and Alternative Methods: To evaluate our proposed framework, we consider both simple baselines and alternative versions of the proposed end-to-end model:

Independent GPs This baseline replicates the work of [Urteaga et al. \(2017\)](#), where each hormone is analyzed with an independent GP learned with measurements within the first two cycles, and tested on the next cycle.

LSTM A neural network with one LSTM layer (with a five-dimensional hidden representation), followed by a fully connected layer which is applied to every time instant. Because recurrent networks do not accept missing data, we impute zero-values for the missing samples in the input time-series.

MGP This is the first component of our end-to-end model. An individual-level MGP is trained with subsampled hormone levels of the first two cycles. We learn the parameters for the across hormone kernel $K(h, h')$ and a single time correlation kernel $k(t, t')$ in the training set, and evaluate its mean prediction on the test set.

Blockwise-MGP We impose a blockwise structure of the kernel matrix of the MGP, allowing it to learn isolated MGPs on separable groups of hormones, each of them having its own parameters. Because LH and FSH exhibit dynamics different from the other three hormones (e.g. acute peaks compared to smoother peaks for the others), we train one block for *LH* and *FSH*, and another for *E*, *P*, and *Ih*.

Blockwise-MGP-DCNN This is our proposed end-to-end model, where the DCNN is trained with $S = 100$ samples from the blockwise MGP’s posterior distribution and the full set of training points.

4.2. Results

We present in [Figure 5](#) an illustrative example of the reconstruction and prediction of our end-to-end statistical framework: given 10 measurements within the first two cycles of an individual, the MGP-DCNN accurately reconstructs estrogen and luteinizing hormone levels across time. We further provide in [Table 1](#) results for the reconstruction and prediction accuracy of all the methods described above. Note that, because we consider all the hor-

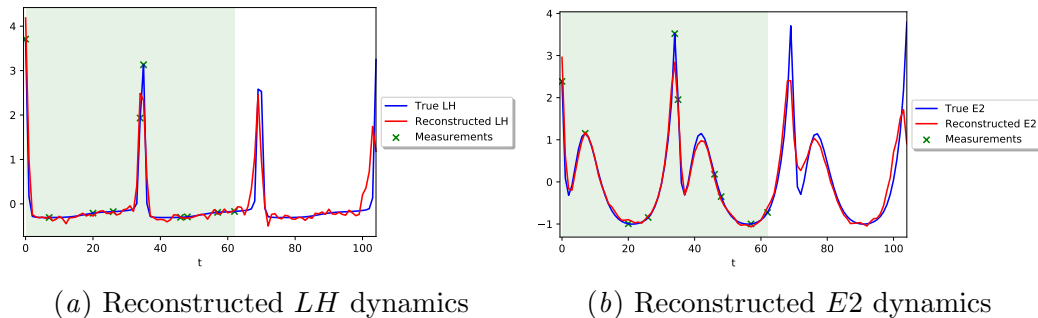


Figure 5: Example of a MGP-DCNN based reconstruction and prediction (in red), given 10 measurements (in green), of the true hormone levels (in blue). The shaded green area indicates reconstruction time (i.e., the training interval).

Test-set overall MSE	Subsampling budget				
Model	$ t_i = 10$	$ t_i = 15$	$ t_i = 25$	$ t_i = 35$	$ t_i = 70$
LSTM	0.358	0.247	0.203	0.186	0.168
Independent GPs	1.245	1.066	0.833	0.140	0.109
MGP	0.823	1.085	0.708	0.127	0.057
B-MGP	0.818	1.132	0.682	0.126	0.119
B-MGP (ED)	0.723	0.335	0.513	0.066	0.115
B-MGP-DCNN	0.302	0.120	0.189	0.041	0.071
B-MGP-DCNN (ED)	0.061	0.045	0.048	0.037	0.050

Table 1: Test-set overall average MSE for all hormones. Random sampling is used for all models, unless specifically (ED) indicated at different subsampling rates.

mones to be equally important, we compute MSE results with respect to the standardized hormone levels, so that the comparison across hormones is fair ¹.

As expected, the higher the available subsampling budget is, the better the results are. In general, note that the performance with once every other day sampling budgets ($|t_i| = 35$) is already satisfactory. At lower subsampling budgets (i.e., less than once every other day measurements), GP based reconstructions are less accurate and volatile, due to the uncertainty over unobserved data points, which hinder the learning of kernel hyperparameters — specially so if only one or two measurements are available per week. As more measurements are available, the GPs outperform the LSTM, which is otherwise quite consistent across sampling budgets.

We observe that learning different temporal patterns per hormone groups is beneficial for GP based models, as well as considering correlations across them. When measurements are scarce (i.e., $|t_i| = 10$), learning a correlation matrix across hormones helps improve reconstruction and prediction performance. As we learn separate kernel matrices with

1. All reconstruction and prediction MSE results across all hormones, as well as per-hormone plots and MSEs are provided in Appendix 5.

different time-correlations per-group, temporal dependencies do not need to be the same for all the hormones, thus matching different physiological patterns.

The proposed framework that combines MGPs and a DCNN consistently outperforms other baselines. The reconstruction and prediction is accurate across all subsampling budgets for the Blockwise-MGP-DCNN model and, most importantly so for the lowest subsampling budget of one measurement per week. We note that, due to the randomness on the training and optimization of MGP and DCNN parameters, performance results are volatile. In our experiments, we find that the more sample streams S drawn from the MGP for training the DCNN, the better and more stable the predictive accuracy of the framework is. We emphasize the two-pronged role of the MGP: (i) it allows training of neural nets, even with irregularly sampled time-series inputs, and (ii) it provides a full distribution over inputs, that allows for an uncertainty aware training of the DCNN. Further, these results support our hypothesis that the non-causal dilation allows for learning long time-horizon dependencies in our hormone level prediction task, reinforcing the claims of the literature (Bai et al., 2018) that convolutional architectures can be used for time-series analysis.

Finally, we note that our optimal sampling method (ED) has a strong impact on performance, both in the MGP alone and the combined MGP-DCNN setting as well. Our end-to-end framework, B-MGP-DCNN (ED), is the-best performing approach: note that optimal sampling with a budget of every other day ($|t_i| = 35$, 0.037 MSE) outperforms other methods with every day measurements ($|t_i| = 70$). More importantly, B-MGP-DCNN with optimal sampling and a budget of only $|t_i| = 10$ days within the first two cycles (i.e., a very realistic setting of measurements every 7 days) performs better than the alternatives at high-sampling rates.

5. Discussion and Related Work

The menstrual cycle is a complex physiological process, and with access to more and more self-tracked data, computational research has very recently started to look into its modeling. Most previous work, however, has focused on high-level characteristics of the cycle, such as predicting cycle length or identifying phases based on specific signals such as body basal temperature (Schervitzl et al., 2015; Sosnowski et al., 2018).

To our knowledge, the work of Urteaga et al. (2017) and this study are the only ones that attempt to reconstruct the full cycle, including hormone levels through time. In this work, we combine the power of generative processes (i.e., MGP) with the ability of convolutional neural networks (i.e., DCNN) to predict hormone levels through time, given a limited sampling budget.

Our proposed reconstruction model is inspired by and leverages several innovations in the machine learning community, yet is novel in both its goal—accurate hormone level inference and prediction with irregular and limited budget sampling—and its end-to-end framework—MGP-based probabilistic training of a non-causal dilated CNN. Combining Gaussian processes, which are well adapted to datasets with uncertain and missing measurements, with neural networks has been recently explored in the context of person-level classification, both in general tasks (Li and Marlin, 2016), as well as for sepsis prediction (Futoma et al., 2017; Moor et al., 2019). Our goal in this paper differs: we combine MGPs with neural nets to improved regression accuracy through time.

Furthermore, we align ourselves with a growing literature within the deep learning community (van den Oord et al., 2016; Kalchbrenner et al., 2016; Dauphin et al., 2016; Gehring et al., 2016, 2017; Bai et al., 2018) that advocates for the use of convolutional neural nets for time-series analysis. In particular, we leverage the exponentially growing receptive fields of dilated CNNs to capture long-horizon time-dependencies, and due to the extended view that the learned MGP provides, leverage non-causal filters to improve the prediction accuracy. To the best of our knowledge, this work is novel in the specific combination of techniques (individual MGPs and a population level non-causal dilated CNN) for prediction of hormone levels in a realistic setting of limited and irregular measurements.

Finally, our work highlights the importance of time-sampling planning in the reconstruction of hormonal levels. First, inclusion of the hormone level peaks is necessary, as suggested previously by Urteaga et al. (2017) and corroborated by our experiments. These peaks can be approximately derived from knowledge of a woman’s cycle length, as well as via ovulation tests that directly measure *LH* peaks, or basal-body temperature tests that do so indirectly. Second, our optimal-sampling strategy enables a minimal amount of measurements compared to random sampling. Our proposed Expected Distance function, which uses the probabilistic nature of GPs, selects the next optimal sampling time balancing the exploration-exploitation dilemma, and provides population-based insights which can be further applied in an individual-level reconstruction task. More precisely, the ED function is computed at the cohort level with respect to a normalized cycle length and, for each individual, the corresponding days of interest can be determined based on their specific cycle lengths. In practice, the optimal sampling requires pre-processing, as one would need to readjust the sampling schedule to each woman’s cycle length —this can be readily done for women with regular cycles.

In conclusion, we have proposed and validated a method to reconstruct and predict an individual’s daily hormonal levels throughout the menstrual cycle based on a few hormone measurements only. This data-driven approach to reconstructing an individual’s cycle patterns can help with minimally-invasive, low-cost data collection at a massive scale.

Acknowledgments

We thank the anonymous reviewers for their feedback and comments, as well as Clue (2019) by BioWink GmbH for the information on cycle length and ovulation day. This work is supported in part by awards from the National Science Foundation (#1344668) and the National Library of Medicine (R01 LM 912734).

References

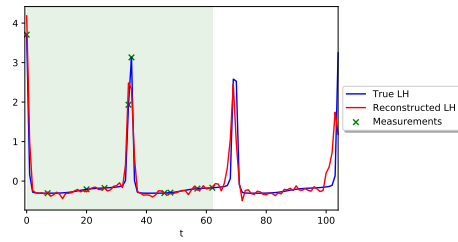
- ACOG. American College of Obstetricians and Gynecologists. Menstruation in girls and adolescents: Using the menstrual cycle as a vital sign. *Obstetrics & Gynecology*, 126(6):143–6, 2015.
- David J. Albers, Matthew E. Levine, Andrew Stuart, Mamykina Lena, Bruce Gluckman, and George Hripcsak. Mechanistic 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*, 25(10):1392–1401, 2018.
- María Elena Alliende. Mean versus individual hormonal profiles in the menstrual cycle. *Fertility and Sterility*, 78(1):90 – 95, 2002.
- Mauricio Alvarez and Neil D. Lawrence. Sparse Convolved Gaussian Processes for Multi-output Regression. In *Advances in Neural Information Processing Systems 21*, pages 57–64. 2009.
- Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *arXiv preprint arXiv:1803.01271*, 2018.
- Matthias Bauer, Mark van der Wilk, and Carl Edward Rasmussen. Understanding probabilistic sparse Gaussian process approximations. In *Advances in neural information processing systems*, pages 1533–1541, 2016.
- Jennifer L. Bedford, Jerilynn C. Prior, and Susan I. Barr. A prospective exploration of cognitive dietary restraint, subclinical ovulatory disturbances, cortisol, and change in bone density over two years in healthy young women. *The Journal of Clinical Endocrinology & Metabolism*, 95(7):3291–3299, 2010.
- Chris Bobel. Beyond the Managed Body: Putting Menstrual Literacy at the Center. In *The Managed Body: Developing Girls and Menstrual Health in the Global South*, pages 281–321. Springer International Publishing, 2019.
- Edwin V. Bonilla, Kian Ming A. Chai, and Christopher Williams. Multi-task Gaussian Process Prediction. In *Advances in Neural Information Processing Systems 20*, pages 153–160. 2008.
- Enrico Carmina and Rogerio A. Lobo. Polycystic Ovary Syndrome (PCOS): Arguably the Most Common Endocrinopathy Is Associated with Significant Morbidity in Women. *The Journal of Clinical Endocrinology & Metabolism*, 84(6):1897–1899, 1999.
- Li-Fang Cheng, Gregory Darnell, Corey Chivers, Michael E Draugelis, Kai Li, and Barbara E Engelhardt. Sparse Multi-Output Gaussian Processes for Medical Time Series Prediction. *ArXiv e-prints*, 2017.

- Leona Harris Clark, Paul M. Schlosser, and James F. Selgrade. Multiple stable periodic solutions in a model for hormonal control of the menstrual cycle. *Bulletin of Mathematical Biology*, 65(1):157–173, Jan 2003.
- Clue. Clue by BioWink GmbH, Adalbertstraße 7-8, 10999 Berlin, Germany. <https://helloclue.com/>, 2019.
- Yann N. Dauphin, Angela Fan, Michael Auli, and David Grangier. Language Modeling with Gated Convolutional Networks. *arXiv preprint arXiv:1612.08083*, 2016.
- Robert Dürichen, Marco. A. F. Pimentel, Lei Clifton, Achim Schweikard, and David A. Clifton. Multi-task Gaussian process models for biomedical applications. In *IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*, pages 492–495, June 2014.
- Robert Dürichen, Marco. A. F. Pimentel, Lei Clifton, Achim Schweikard, and David A. Clifton. Multitask Gaussian Processes for Multivariate Physiological Time-Series Analysis. *IEEE Transactions on Biomedical Engineering*, 62(1):314–322, Jan 2015.
- Amanda C. Filiberto, Anna Z. Pollack, Donald R. Mattison, Enrique F. Schisterman, Neil J. Perkins, Paul S. Albert, Sunni L. Mumford, Anne Z. Steiner, and Jean Wactawski-Wende. The Utility of Menstrual Cycle Length as an Indicator of Cumulative Hormonal Exposure. *The Journal of Clinical Endocrinology & Metabolism*, 97(10):E1871–E1879, 10 2012.
- Peter I. Frazier. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.
- Joseph Futoma, Sanjay Hariharan, and Katherine Heller. Learning to Detect Sepsis with a Multitask Gaussian Process RNN Classifier. *arXiv preprint arXiv:1708.05894*, 2017.
- Jonas Gehring, Michael Auli, David Grangier, and Yann N. Dauphin. A Convolutional Encoder Model for Neural Machine Translation. *arXiv preprint arXiv:1611.02344*, 2016.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. Convolutional Sequence to Sequence Learning. *arXiv preprint arXiv:1705.03122*, 2017.
- Linda C. Giudice. Endometriosis. *NEJM*, 362(25):2389 – 2398, 2010.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- Paige Hornsby, Allen J. Wilcox, Anne Marie Zaura Jukic, Donna D. Baird, Clarice R. Weinberg, and D. Robert McConaughy. Accuracy of Reporting of Menstrual Cycle Length. *American Journal of Epidemiology*, 167(1):25–33, 11 2007.
- Nal Kalchbrenner, Lasse Espeholt, Karen Simonyan, Aäron van den Oord, Alex Graves, and Koray Kavukcuoglu. Neural Machine Translation in Linear Time. *arXiv preprint arXiv:1610.10099*, 2016.
- Colin Lea, Michael D. Flynn, René Vidal, Austin Reiter, and Gregory D. Hager. Temporal Convolutional Networks for Action Segmentation and Detection. *arXiv preprint arXiv:1611.05267*, 2016.

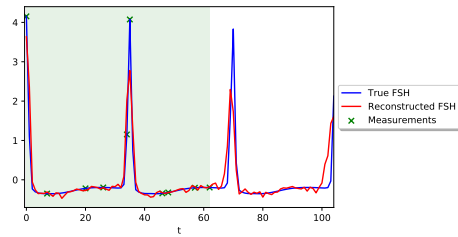
- Steven Cheng-Xian Li and Benjamin Marlin. A Scalable End-to-end Gaussian Process Adapter for Irregularly Sampled Time Series Classification. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 1812–1820, 2016. ISBN 978-1-5108-3881-9.
- Megan M. Mahoney. Shift work, jet lag, and female reproduction. *International Journal of Endocrinology*, 2010, 2010.
- Robert I. McLachlan, Nancy L. Cohen, Kristine D. Dahl, William J. Bremner, and Michael R. Soules. Serum inhibin levels during the periovulatory interval in normal women: Relationships with sex steroid and gonadotrophin levels. *Clinical Endocrinology*, 32(1):39–48, 1990.
- Michael Moor, Max Horn, Bastian Rieck, Damian Roqueiro, and Karsten M. Borgwardt. Temporal Convolutional Networks and Dynamic Time Warping can Drastically Improve the Early Prediction of Sepsis. *arXiv preprint arXiv:1902.01659*, 2019.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- Elina Berglund Scherwitzl, Angelica Lindén Hirschbergand, and Raoul Scherwitzl. Identification and prediction of the fertile window using naturalcycles. *The European Journal of Contraception & Reproductive Health Care*, 20(5):403–408, 2015.
- Peter Schulam and Suchi Saria. A Framework for Individualizing Predictions of Disease Trajectories by Exploiting Multi-Resolution Structure. In *Advances in Neural Information Processing Systems*, pages 748–756. 2015.
- James F. Selgrade and Paul M. Schlosser. A model for the production of ovarian hormones during the menstrual cycle. *Fields Institute Communications*, pages 429 – 446, 1999.
- James F. Selgrade, Leona A. Harris, and R.Drew. Pasteur. A model for hormonal control of the menstrual cycle: Structural consistency but sensitivity with regard to data. *Journal of Theoretical Biology*, 260(4):572 – 580, 2009.
- Lynne T. Shuster, Deborah J. Rhodes, Bobbie S. Gostout, Brandon R. Grossardt, and Walter A. Rocca. Premature menopause or early menopause: long-term health consequences. *Maturitas*, 65(2):161–166, February 2010.
- Bernhard W. Silverman. Some aspects of the spline smoothing approach to non-parametric regression curve fitting. *Journal of the Royal Statistical Society: Series B (Methodological)*, 47(1):1–21, 1985.
- skopt module. <https://scikit-optimize.github.io/>.
- Chanley M. Small, Amita K. Manatunga, and Michele Marcus. Validity of Self-Reported Menstrual Cycle Length. *Annals of Epidemiology*, 17(3):163 – 170, 2007.
- Alex J. Smola and Peter L. Bartlett. Sparse Greedy Gaussian Process Regression. In *Advances in Neural Information Processing Systems 13*, pages 619–625. MIT Press, 2001.

- Edward Snelson and Zoubin Ghahramani. Sparse Gaussian Processes using Pseudo-inputs. In *Advances in Neural Information Processing Systems*, pages 1257–1264. 2006.
- Caren G. Solomon, Frank B. Hu, Andrea Dunaif, Janet E. Rich-Edwards, Meir J. Stampfer, Walter C. Willett, Frank E. Speizer, and Joann E. Manson. Menstrual cycle irregularity and risk for future cardiovascular disease. *The Journal of Clinical Endocrinology and Metabolism*, 87(5):2013–2017, May 2002.
- Lukasz Sosnowski, Wojciech Chaber, Lukasz Miłobedzki, Tomasz Penza, Jadwiga Sosnowska, Karol Zaleski, Joanna Fedorowicz, Iwona Szymusik, and Dominik Ślęzak. Similarity-based Detection of Fertile Days at OvuFriend. In *IEEE International Conference on Big Data*, pages 2663–2668, 2018.
- Oliver Stegle, Sebastian V. Fallert, David J. C. MacKay, and Soren Brage. Gaussian Process Robust Regression for Noisy Heart Rate Data. *IEEE Transactions on Biomedical Engineering*, 55(9):2143–2151, Sep. 2008.
- Alan E. Treloar, Ruth E. Boynton, Borghild G. Behn, and Byron W. Brown. Variation of the human menstrual cycle through reproductive life. *International journal of fertility*, 12(1 Pt 2):77–126, 1967.
- Shelley S. Tworoger, A. Heather Eliassen, Xuehong Zhang, Jing Qian, Patrick M. Sluss, Bernard A. Rosner, and Susan E. Hankinson. A 20-year prospective study of plasma prolactin as a risk marker of breast cancer development. *Cancer research*, 73(15):4810–4819, 2013.
- Iñigo Urteaga, David J. Albers, Marija Vlajic Wheeler, Anna Druet, Hans Raffauf, and Noémie Elhadad. Towards Personalized Modeling of the Female Hormonal Cycle: Experiments with Mechanistic Models and Gaussian Processes. *NIPS ML4H workshop*, 2017.
- Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. WaveNet: A Generative Model for Raw Audio. *arXiv preprint arXiv:1609.03499*, 2016.
- Rita Vassena, Ricard Vidal, Oriol Coll, and Valrie Vernaève. Menstrual cycle length in reproductive age women is an indicator of oocyte quality and a candidate marker of ovarian reserve. *European Journal of Obstetrics & Gynecology and Reproductive Biology*, 177:130 – 134, 2014.
- Fisher Yu and Vladlen Koltun. Multi-Scale Context Aggregation by Dilated Convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
- Armin Zittermann, I. Schwarz, K. Scheld, T. Sudhop, H. K. Berthold, K. von Bergmann, H. van der Ven, and P. Stehle. Physiologic fluctuations of serum estradiol levels influence biochemical markers of bone resorption in young women. *The Journal of Clinical Endocrinology and Metabolism*, 85(1):95–101, 2000.

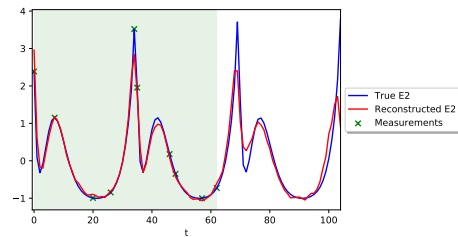
Evaluation



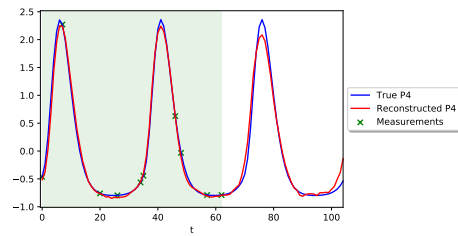
(a) Reconstructed LH dynamics



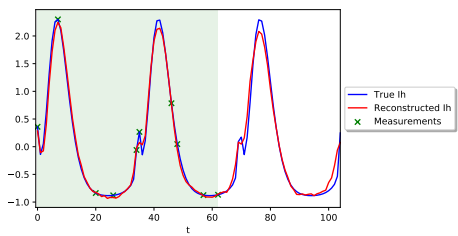
(b) Reconstructed FSH dynamics



(c) Reconstructed E dynamics



(d) Reconstructed P dynamics



(e) Reconstructed Ih dynamics

Figure 6: Example of a MGP-DCNN based reconstruction and prediction (in red), given 10 measurements (in green), of the true hormone levels (in blue).

Test-set reconstruction MSE	Subsampling budget				
Model	$ t_i = 10$	$ t_i = 15$	$ t_i = 25$	$ t_i = 35$	$ t_i = 70$
LSTM	0.336	0.240	0.176	0.173	0.131
Independent GPs	1.198	0.600	0.721	0.140	0.927
MGP	0.735	1.042	0.563	0.119	0.045
B-MGP	0.729	1.134	0.524	0.123	0.100
B-MGP (ED)	0.557	0.239	0.242	0.064	0.108
B-MGP-DCNN	0.227	0.096	0.118	0.035	0.054
B-MGP-DCNN (ED)	0.048	0.042	0.035	0.034	0.045

Table 2: Test-set reconstruction average MSE for all hormones. Random sampling is used for all models, unless specifically (ED) indicated at different subsampling rates.

Test-set prediction MSE	Subsampling budget				
Model	$ t_i = 10$	$ t_i = 15$	$ t_i = 25$	$ t_i = 35$	$ t_i = 70$
LSTM	0.402	0.261	0.258	0.213	0.241
Independent GPs	1.338	0.659	1.057	0.141	0.875
MGP	0.998	1.171	0.998	0.143	0.081
B-MGP	0.998	1.126	0.998	0.131	0.125
B-MGP (ED)	1.055	0.250	1.055	0.070	0.129
B-MGP-DCNN	0.453	0.134	0.329	0.049	0.106
B-MGP-DCNN (ED)	0.088	0.050	0.074	0.043	0.062

Table 3: Test-set prediction average MSE for all hormones. Random sampling is used for all models, unless specifically (ED) indicated at different subsampling rates.

Test-set overall MSE for LH	Subsampling budget				
Model	$ t_i = 10$	$ t_i = 15$	$ t_i = 25$	$ t_i = 35$	$ t_i = 70$
LSTM	0.703	0.475	0.406	0.389	0.348
Independent GPs	1.671	1.220	0.930	0.382	1.002
MGP	0.671	2.845	0.560	0.282	0.135
B-MGP	0.674	1.775	0.548	0.295	0.261
B-MGP (ED)	0.563	0.485	0.419	0.186	0.282
B-MGP-DCNN	0.542	0.196	0.336	0.071	0.132
B-MGP-DCNN (ED)	0.105	0.072	0.086	0.046	0.078

Table 4: Test-set overall average MSE for LH . Random sampling is used for all models, unless specifically (ED) indicated at different subsampling rates.

Test-set overall MSE for FSH	Subsampling budget				
Model	$ t_i = 10$	$ t_i = 15$	$ t_i = 25$	$ t_i = 35$	$ t_i = 70$
LSTM	0.583	0.431	0.324	0.289	0.272
Independent GPs	1.564	0.900	0.933	0.228	0.965
MGP	0.772	1.587	0.699	0.205	0.085
B-MGP	0.769	2.189	0.671	0.260	0.188
B-MGP (ED)	0.569	0.356	0.426	0.107	0.209
B-MGP-DCNN	0.494	0.164	0.297	0.055	0.107
B-MGP-DCNN (ED)	0.084	0.051	0.066	0.037	0.060

Table 5: Test-set overall average MSE for FSH . Random sampling is used for all models, unless specifically (ED) indicated at different subsampling rates.

Test-set overall MSE for E	Subsampling budget				
Model	$ t_i = 10$	$ t_i = 15$	$ t_i = 25$	$ t_i = 35$	$ t_i = 70$
LSTM	0.319	0.205	0.160	0.132	0.105
Independent GPs	1.136	0.731	0.779	0.077	0.927
MGP	0.786	0.638	0.706	0.099	0.049
B-MGP	0.809	0.860	0.720	0.060	0.074
B-MGP (ED)	0.809	0.170	0.646	0.031	0.069
B-MGP-DCNN	0.295	0.115	0.188	0.041	0.073
B-MGP-DCNN (ED)	0.073	0.049	0.056	0.051	0.064

Table 6: Test-set overall average MSE for E . Random sampling is used for all models, unless specifically (ED) indicated at different subsampling rates.

Test-set overall MSE for P	Subsampling budget				
Model	$ t_i = 10$	$ t_i = 15$	$ t_i = 25$	$ t_i = 35$	$ t_i = 70$
LSTM	0.098	0.053	0.048	0.060	0.053
Independent GPs	0.918	0.091	0.757	0.002	0.800
MGP	0.938	0.146	0.794	0.020	0.009
B-MGP	0.920	0.395	0.736	0.003	0.009
B-MGP (ED)	0.831	0.085	0.528	0.001	0.007
B-MGP-DCNN	0.089	0.033	0.058	0.014	0.020
B-MGP-DCNN (ED)	0.020	0.024	0.014	0.025	0.023

Table 7: Test-set overall average MSE for P . Random sampling is used for all models, unless specifically (ED) indicated at different subsampling rates.

Test-set overall MSE for Ih	Subsampling budget				
Model	$ t_i = 10$	$ t_i = 15$	$ t_i = 25$	$ t_i = 35$	$ t_i = 70$
LSTM	0.086	0.073	0.079	0.062	0.061
Independent GPs	0.935	0.156	0.767	0.013	0.853
MGP	0.948	0.211	0.781	0.028	0.009
B-MGP	0.920	0.439	0.736	0.010	0.012
B-MGP (ED)	0.843	0.117	0.546	0.004	0.010
B-MGP-DCNN	0.091	0.038	0.063	0.016	0.023
B-MGP-DCNN (ED)	0.023	0.027	0.017	0.028	0.027

Table 8: Test-set overall average MSE for Ih . Random sampling is used for all models, unless specifically (ED) indicated at different subsampling rates.