Published in final edited form as:

Biometrics. 2018 June ; 74(2): 714–724. doi:10.1111/biom.12790.

Modeling Associations between Latent Event Processes Governing Time Series of Pulsing Hormones

Huayu Liu¹, Nichole E. Carlson^{2,*}, Gary K. Grunwald², and Alex J. Polotsky³

¹Eli Lilly and Company, Indianapolis, Indiana 46285, U.S.A

²Department of Biostatistics and Informatics, University of Colorado Anschutz Medical Campus, Aurora, Colorado, U.S.A

³Department of Obstetrics and Gynecology, University of Colorado Anschutz Medical Campus, Aurora, Colorado, U.S.A

Summary

This work is motivated by a desire to quantify relationships between two time series of pulsing hormone concentrations. The locations of pulses are not directly observed and may be considered latent event processes. The latent event processes of pulsing hormones are often associated. It is this joint relationship we model. Current approaches to jointly modeling pulsing hormone data generally assume that a pulse in one hormone is coupled with a pulse in another hormone (one-to-one association). However, pulse coupling is often imperfect. Existing joint models are not flexible enough for imperfect systems. In this article, we develop a more flexible class of pulse association models that incorporate parameters quantifying imperfect pulse associations. We propose a novel use of the Cox process model as a model of how pulse events co-occur in time. We embed the Cox process model into a hormone concentration model. Hormone concentration is the observed data. Spatial birth and death Markov chain Monte Carlo is used for estimation. Simulations show the joint model works well for quantifying both perfect and imperfect associations and offers estimation improvements over single hormone analyses. We apply this model to luteinizing hormone (LH) and follicle stimulating hormone (FSH), two reproductive hormones. Use of our joint model results in an ability to investigate novel hypotheses regarding associations between LH and FSH secretion in obese and non-obese women.

Keywords

Bivariate point processes; Follicle stimulating hormone; Joint point process models; Luteinizing hormone; Pulsatile hormone; Reproductive hormones

*nichole.carlson@ucdenver.edu

7. Supplementary Materials

Web Appendix, Tables and Figures as referenced in Sections 1, 2, 3, 5.2 are available with this article at the *Biometrics* website on Wiley Online Library. Statistical code is publicly available at www.github.com/BayesPulse/jointpopmodel.

1. Introduction

A wide body of literature exists on methods for analyzing hormone curves, in particular female reproductive hormones. Two common study designs collect hormone values annually over multiple years (Jiang et al., 2015; Quintana et al., 2016) or daily for a menstrual cycle (Zhang et al., 1998, 2000) and use longitudinal analysis techniques. It is not well known that many reproductive hormones must be secreted in an intermittent fashion (e.g., every 60–120 minutes) for normal monthly cycling to occur. To study intermittent pulse secretion, investigators perform rapid sampling studies where blood is drawn every 5–10 minutes for a period of 6–24 hours (Web Figure S1). It is this type of design that is modeled in this article.

The signature pattern in these data is a set of rapid jumps in concentrations created by the pulsatile release of the hormone from a gland. Characterizing pulsatile secretion is challenging because the signal is not directly observed; at any time point, both secretion and elimination are simultaneously occurring. There are various non-linear hierarchical models that have been developed to model a single hormone (Veldhuis and Johnson, 1992; Keenan et al., 1998; Johnson, 2003). These models contain parameters with biologic interpretation such as pulse mass (size of the pulse), pulse width (duration of the release), basal (non-pulsatile) concentration, and hormone half-life.

Hormones in endocrine axes, like the reproductive axis, require well-orchestrated signaling between the hormones for effective functioning (Strauss and Barbieri, 2009). For example, the signal network of reproductive hormones is initiated by secretion of pulses of gonadotropin-releasing hormone (GnRH) from the hypothalamus, which triggers pulsatile release of luteinizing hormone (LH) and less perfectly triggers the pulsatile release of follicle-stimulating hormone (FSH) (Hall et al., 1990). This difference in pulse coupling occurs as a result of differential binding of the GnRH molecules to the pituitary GnRH receptors (Thompson and Kaiser, 2014). Alterations in the pulsatile release or their associations can change the entire system. An extreme result is reproductive dysfunction. Thus, there is interest in characterizing the associations between pulsing hormones.

The goal of this article is to develop a model of the associations between two pulsing hormones with well defined parameters that can be compared between groups. The types of associations seen in endocrine systems are trigger-response associations in that a pulse in one hormone is believed to trigger a pulse in another hormone (Greenspan and Gardner, 2004). Nearly all of the existing trigger-response models of pulsatile hormone associations (Guo and Brown, 2001; Keenan et al., 2008; Carlson et al., 2009) assume there is a one-to-one relationship between the pulsing of the trigger and response hormones. However, for many pulsing hormones, such as LH and FSH, the relationship is imperfect (a response pulse is not always triggered). In these instances, existing approaches fail because response pulses are estimated in places where they did not occur. The innovation of this article is a new flexible class of trigger-response models that does not require the strict one-to-one assumption.

We treat the pulse locations as latent point processes and propose a novel use of the Cox process model (Cox and Isham, 1984) to flexibly model how pulse events co-occur in time.

We embed the Cox model into an existing deconvolution model of hormone concentrations (Veldhuis and Johnson, 1992). The Cox process model is a parametric framework that is nicely defined by a set of biologically interpretable parameters. In particular, we can estimate directly: 1) the chance that a trigger pulse is accompanied by a response pulse and 2) the tightness in time of the response to a trigger. These are the quantities that our clinical collaborator hypothesized could differ between groups. Spatial birth-death Markov chain Monte–Carlo (SBDMMCMC) is used for estimation. We exhibit the strengths of this model on reproductive studies in women. Simulation is used to investigate model performance under different designs, pulse sizes and types of association models.

The article is organized as follows: Section 2 defines the joint deconvolution model and the new joint pulse location model. In Section 3, SBDMMCMC implementation details are provided. We apply our new model to two studies of LH and FSH in Section 4. Simulations are described in Section 5. Discussion can be found in Section 6. Before defining the model, we note that the trigger-response language implies causation; however the models developed are not casual. They characterize co-occurrence of events in time.

2. The Model

The Bayesian deconvolution model is developed elsewhere (Johnson, 2003; Horton et al., 2017). Let $\{x_i(t_{ij}), y_i(t_{ij})\}$ be a pair of observed hormone concentrations for the trigger (x) and response (y) hormones at time t_{ij} for the j th observation on the i th subject, where $i = 1, 2, 3, \dots, m$, and $j = 1, 2, 3, \dots, n_i$. In our application, x is LH and y is FSH. Without loss of generality, we define the deconvolution model for the trigger hormone concentration and replace x with y for the model of the response hormone. The true hormone concentration, $C^x(t_{ij})$, is linked with the observed concentration through the following statistical model:

$\log \{x_i(t_{ij})\} = \log C^x(t_{ij}) + \varepsilon_{ij}^x$, where $C^x(t_{ij}) = \theta_{b,i}^x + \int_{-\infty}^{t_{ij}} S(z; \theta_{s,i}^x) E(t_{ij} - z; \theta_{h,i}^x) dz$. $\theta_{b,i}^x$ defines the non-pulsatile baseline concentration. $S(z; \theta_{s,i}^x)$ is the pulsatile secretion for subject i described further in Section 2.2. $E(z; \theta_{h,i}^x)$ is the elimination function and modeled as a single exponential decay, that is, $E(z; \theta_{h,i}^x) = \exp \left[- \left\{ \frac{\log 2}{\theta_{h,i}^x} \right\} z \right]$, where $\theta_{h,i}^x$ is the hormone half-life

for subject i . ε_{ij}^x represents model error. The errors of the trigger (x) and response (y)

hormones may be correlated. Thus, we can jointly model the errors:

$(\varepsilon_{ij}^x, \varepsilon_{ij}^y)' \sim MVN((0, 0)', \Sigma_{e,i})$. Thus, the likelihood $L(\mathbf{x}_i, \mathbf{y}_i)$ is a multivariate normal likelihood with a mean vector $(C^x(t_{ij}), C^y(t_{ij}))'$ and variance-covariance matrix $\Sigma_{e,i}$. All notation is further described in Web Table S1.

2.2. The Pulse Secretion Function, $S(z)$

The pulsatile secretion function of the trigger hormone for subject i is defined as the

superposition of all trigger pulse secretion events for subject i , $S(z; \theta_{s,i}^x) = \sum_{k=1}^{N_{s,i}^x} p(z; \theta_{s,ik}^x)$, where p is a non-negative user defined function that depends on pulse specific parameters $\theta_{s,ik}^x$, and $N_{s,i}^x$ is the unknown number of trigger pulses for subject i over the period of

observation. The k th trigger pulse event for subject i is assumed to have a Gaussian shape

$$(Veldhuis and Johnson, 1992), \text{ that is, } p(z; \theta_{s,ik}^x) = \frac{\alpha_{ik}^x}{\sqrt{2\pi\omega_{ik}^x}} \exp \left\{ -\frac{(z - \tau_{ik}^x)^2}{2\omega_{ik}^x} \right\}. \text{ Three}$$

parameters define each Gaussian shape pulse: pulse location, τ_{ik}^x , pulse mass, α_{ik}^x , the total secretion for the pulse, and pulse width, ω_{ik}^x , representing the duration of the pulse. These parameters together define the set of pulse specific parameters, $\theta_{s,ik}^x = (\tau_{ik}^x, \alpha_{ik}^x, \omega_{ik}^x)'$. The association model links these parameters between two hormones through their priors.

2.3. Specifications of the Priors

In the deconvolution model, there are pulse specific parameters (location, mass and width) and subject specific parameters (e.g., average pulse mass, baseline concentration, and half-life). The complete hierarchical specification is in Table 1. Boxed priors highlight the joint priors developed to link pulsing between the trigger and response hormones.

2.3.1. The joint model of pulse locations, Table 1(b)—The general framework is described first. A Cox cluster process (Cox and Isham, 1984) is used to jointly model how the pulses of two hormones co-occur in time. A Cox cluster process has two levels, offspring events (Level 1) and latent parent events (Level 2). Offspring events cluster in time close to parent events, with the expected number of offspring events and the temporal spread of the offspring defined by additional parameters. In our framework, the Level 2 events are trigger hormone (x) pulse locations, τ^x . The Level 1 events are the response hormone (y) pulse locations, τ^y .

Trigger pulse location model (Level 2): Let $\tau_i^x = \{\tau_{i1}^x, \dots, \tau_{iN_{s,i}^x}^x\}$ be a set of trigger pulse

locations for subject i from point process $\mathcal{T}_i^x$. We model $\mathcal{T}^x$ as a multilevel Strauss process (Penttinen, 1984). The multilevel Strauss process is a point-wise interacting point process that induces repulsion between events. The levels of repulsion in this model lessen or eliminate the chance that two trigger pulses will occur very close in time and induce some regularity in the trigger pulse locations, as is expected biologically. The density function,

$$\pi(\tau_i^x) \text{ is proportionate to } \beta^{N_i^x} \prod_{r=1}^2 \gamma_r^{s_r(\tau)} \text{ where } s_r = \sum_{\tau_{i,q}^x, \tau_{i,q'}^x \in \tau_i^x} I(\|\tau_{i,q}^x - \tau_{i,q'}^x\| < R_r). \text{ In other}$$

words, s_r counts the number of trigger pulse pairs that are less than R_r units apart in time. In this intensity function, β ($\beta > 0$) is the rate parameter. A larger β is associated with more pulses in a given time period. $\gamma_r \in [0, 1]$ is the strength of interaction of a pair of points within distance R_r ($R_r > 0$) for the r th level. Strict repulsion is when $\gamma_r = 0$, and $\gamma_r = 1$ indicates no repulsion (i.e., a Poisson process). Given that subjects are independent, the joint trigger location model for all subjects, $\pi(\tau^x)$, is the product of $\pi(\tau_i^x)$. In application, the user selects β , R_r and γ_r .

Response pulse location model (Level 1): Let τ_i^y be the set of response pulse locations for subject i from a Cox cluster process, $\mathcal{T}_i^y$, where $i = 1, \dots, m$. The response pulse location process, $\mathcal{T}_i^y$ is defined by its conditional intensity function, $\lambda_i^y(z | \tau_i^x)$, which depends on a realization of trigger pulse locations, τ_i^x . Since the Cox process is a type of Poisson process, the density is found by plugging the conditional intensity into the density equation for the Poisson process (Cox and Isham, 1984). The intensity function of our Cox process is as follows:

$$\lambda_i^y(z | \tau_i^x, \rho, v) = \sum_{k=1}^{N_i^x} \frac{\rho_{ik}}{\sqrt{2\pi v_{ik}^2}} \exp \left\{ -\frac{(z - \tau_{ik}^x)^2}{2v_{ik}^2} \right\}. \quad (1)$$

Except at the boundaries, ρ_{ik} is the expected number of response pulses for the k th trigger pulse in subject i and is typically less than one in the imperfect setting. Similarly, v_{ik} is the standard deviation (spread) of the Gaussian-shaped intensity function, which quantifies the mostly likely time region in which the response pulses will appear for the k th trigger pulse in subject i . We assume that all v_{ik} and ρ_{ik} are common across all subjects and triggers, that is, $\rho_{ik} = \rho$, $v_{ik} = v$. This restriction implies that the expected number of responses for each trigger is similar within the population. Importantly, this assumption does not restrict the number of responses to be identical for each trigger. Similarly, assuming that the spread of the response pulse distribution, v , is the same across individuals and triggers still allows for flexibility in the actual location of the response pulses for each trigger. It may seem biologically implausible to allow a response pulse to occur before a trigger. In practice, this can occur because the tightness of the association is often short compared to the sampling interval and the observed trigger may be a surrogate for an unmeasurable trigger, which will be the case for LH and FSH. Based on the superposition principle (Cox and Isham, 1984), the joint prior for the response pulse locations, $\pi(\tau^y)$, is still a Cox process with intensity function $\sum_i \lambda_i^y(z | \tau_i^x, \rho, v)$. To complete the prior specification, $\log \rho \sim N(m_\rho, v_\rho)$ and $\log v \sim N(m_v, v_v)$, where m_ρ , v_ρ , m_v and v_v are user specified.

2.3.2. Priors for pulse mass and width, Table 1(c)—The priors for the pulse specific masses and widths for the trigger and response hormones are multivariate truncated t-distributions with four degrees of freedom, $t_4^+(\mu_{s,i}^x, \Upsilon_s^x)$ and $t_4^+(\mu_{s,i}^y, \Upsilon_s^y)$, respectively. The distributions are truncated to be positive. These priors provide a robust alternative to the normal distribution and have been used in other pulse hormone models (Johnson, 2003; Carlson et al., 2009). The multivariate t-distribution with r degrees of freedom can be sampled using a conditional normal distribution with variance mixture scaled by a random variable $\kappa \sim \text{Gamma}(r/2, r/2)$ (Kleinman and Ibrahim, 1998). More precisely, if $(\alpha, \omega)' | \kappa \sim N(\mu_s, \kappa^{-1} \Upsilon)$, then the marginal distribution of (α, ω) is a multivariate t_r . These priors depend upon a set of population level hyper-parameters, $\mu_{s,i}^x, \mu_{s,i}^y, \Upsilon_s^x, \Upsilon_s^y$. The second component where we associate pulses between the trigger and response hormones is

between the subject level mean pulse parameters, $\mu_{s,i}^x$ and $\mu_{s,i}^y$. Specifically, these means are modeled as a multivariate truncated normal, truncated to be positive, $(\mu_{s,i}^x, \mu_{s,i}^y)' \sim MVN^+(\mu_s^{x,y}, \Sigma_s^{x,y})$. This imposes a linear association between the mean pulse mass and width of the trigger and response pulses for subject i . We complete the specification of the prior by assuming $\mu_s^{xy} \sim MVN^+(m_s^{xy}, V_s^{xy})$, where m_s^{xy} and V_s^{xy} are user specified. Further $\Upsilon_s^{xy} \sim IW_v(S_1)$ and $\Sigma_s^{xy} \sim IW_v(S_2)$, where IW is the Inverse-Wishart with v degrees of freedom.

2.3.3. Priors for baseline and half-life, Table 1(d,e)—As in Horton et al. (2017), $\theta_{b,i}^x \sim N(\mu_b^x, \sigma_{b,x}^2)$, $\mu_b^x \sim N(m_b^x, v_{b,x}^2)$ and $\sigma_{b,x}^{-2} \sim \text{Unif}(0, U_b^x)$ and $\theta_{h,i}^x \stackrel{iid}{\sim} N(\mu_h^x, \sigma_{h,x}^2)$, $\mu_h^x \sim N(m_h^x, v_{h,x}^2)$ and $\sigma_{h,x}^{-2} \sim \text{Unif}(0, U_{h,x})$. m , v , and U are user defined choices. The same priors are used for the response hormone with hormone specific choices for m , v , and U . These priors were independent across the two hormones because we were focused on modeling associations in the pulse parameters. It is certainly possible to generalize and have these parameters be associated also. Finally the prior on the model error is $IW_v(R_3)$.

3. Estimation and Implementation

The SBDMCMC algorithm used for simulating the posteriors is developed in Web Appendix A. Convergence of this algorithm is ensured by arguments in (Geyer and Møller, 1994; Stephens, 2000). The algorithm was custom coded in C and available at www.github.com/BayesPulse/jointpopmodel.

In the prior on the trigger locations, we set $\beta = 4$, $\gamma_1 = 0$, $R_1 = 0.6$ hours, $\gamma_2 = 0.1$ and $R_2 = 1.6$ hours. The parameters were chosen to reflect clinician input on how close pulses can be and how many pulses should occur on average over the period of observation. The results are not sensitive to a fairly wide selection for β . If the selection becomes too high, then pulses are modeled at exactly R_1 units apart trying to model a single pulse with multiple components. This scenario can be found by inspection of the posterior distribution of the trigger locations. The results are also not that sensitive to γ_2 and R_2 . Care does need to be taken in choosing R_1 because it firmly restricts the minimum distances two pulses can occur. We have found minimal changes in the results for R_1 values between 20 and 40 minutes. The priors on ρ and v were set at $\log \rho \sim N(-0.11, 1)$ and $\log v \sim N(1.5, 10)$, respectively.

The prior means for the population mean pulse mass were 3 and 1 for the LH and FSH, respectively. The population mean pulse width was 5 minutes for both hormones. The mean baseline was set at 4 units, and the mean half-life was 45 minutes and 200 minutes for LH and FSH, respectively. All had vague variances on these means of 100 units². These values were selected based on clinical input and prior studies (and formed the basis for our simulations). The prior means were set to the simulated values for the simulations with variances of 100 units². As in previous work (Carlson et al., 2009; Horton et al., 2017), the starting values were set to the medians of the posteriors. As long as the values are in a

biologically plausible range, the results are not sensitive to the starting values. In addition, each subject was initialized to have one randomly located trigger and response pulse.

In implementation, a single chain was run for 200,000 iterations (real data) and 100,000 iterations (simulations). As in previous work (Carlson et al., 2009; Horton et al., 2017), every 50th iteration was stored to save space. For the first 10,000 iterations, the proposal variances were adjusted every 500 iterations to achieve acceptance rates between 25 and 50%. The first 10,000 iterations were discarded as burn-in. Convergence was assessed graphically on the parameters that did not change dimension. One population of 10 subjects, each with 144 observations in 24 hours, took about 18 hours to run on a laptop with an Intel Core i5 2.40 GHz processor. Speed could be improved with parallelization of the subject level birth-death loops and subject level posterior draws.

4. Applications to LH and FSH Studies

4.1. Background

The motivating study was designed to investigate female gonadotrope secretion and regulation (LH and FSH) in obese women (Al-Safi et al., 2015). Because GnRH pulses are not directly measurable in human females, we use LH pulse as a surrogate trigger for FSH given its one-to-one association with GnRH (Strauss and Barbieri, 2009). This will allow us to indirectly estimate the association between GnRH and FSH along with potential differences in the relationship for obese compared to normal weight women. For 10 normal weight (NW) women and 12 obese women of age 18–42, blood was collected every 10 minutes for 6 hours in their early follicular phase. We were asked to characterize how LH and FSH pulses were associated and how those associations differ between ovulating obese and non-obese women. We fitted the obese and NW women separately and then the posterior differences of parameters calculated since the groups were independent. This design had a short duration and some of the parameter values in the analysis seemed less representative of a single imperfect trigger-response association. Thus, we analyzed a second natural history study with a more frequent sampling interval (7.5 minutes) and a longer study design (24 hours) (Pincus et al., 1998). Data were collected on 5 healthy women in the follicular phase. The data from one woman in her follicular phase was discarded for quality issues.

4.2. Characterizing Pulses and Pulse Associations in Obese and Non-Obese Women

Figure 1 (bottom panel) shows fits of the LH and FSH series and the posterior distribution of pulse location draws for one NW (Figure 1B) and one obese patient (Figure 1C). The fits for the 6 hour data were not as precise as the 24-hour data showing wider 95% posterior predictive intervals and more diffuse posterior distributions of pulse locations. The LH pulse frequency was similar between the groups [posterior probability that $NW > Obese$ (PP) = 0.23] (Table 2); however, the FSH pulse frequency was 2.6 pulses greater in the normal weight compared to the obese (PP = 0.99). Average LH pulse mass was greater in the NW compared to the obese (PP = 0.89). FSH pulse mass did not differ between the groups (PP = 0.47). The concordance of LH and FSH pulses (ρ) was stronger in the normal weight group compared to the obese group (PP = 0.99). The median average number of FSH pulses per LH pulse was 1.11 [95% CI: (0.78, 1.53)] in the NW group and only 0.59 [95% CI:

(0.37,0.86)] in the obese group. The spread of the FSH pulse locations relative to LH pulse locations was large at 30.9 minutes [95% CI: (18.5,71.2)] in NW women and 29.1 minutes [95% CI: (16.1,66.7)] in obese women but did not differ between groups ($PP = 0.56$).

4.3. 24-hour Study in Normal Weight Women

Figure 1 (top figure) shows posterior predictive fits of the LH and FSH obtained using the joint model for a randomly selected subject. The number of LH pulses was similar between the single hormone and joint model fits. However, the joint model detects approximately two more pulses over the 24-hour period (Table 2 bottom panel). The posterior median of the average number of FSH pulses for each LH pulse (ρ) was 0.67 [95% CI: (0.47,0.92)]. The spread of the timing of FSH response pulses was 6.5 minutes [95% CI: (6.1,8.8)], meaning that on average, when an FSH pulse was triggered by a GnRH pulse, the probability that the distance between a trigger and response is smaller than 12.9 minutes is 95%. These findings are more consistent with a single imperfect trigger-response model.

5. Simulation

5.1. Simulated Data

Various trigger-response scenarios were simulated to assess the performance of the Bayesian Cox cluster model. The parameters defining each scenario can be found in Table 3. Scenario 1: Imperfect association with large response pulse masses; Scenario 2: Imperfect association with small response pulse masses. This case closely resembles LH and FSH. Scenario 3: a perfect association model with large response pulse masses. In this scenario, a trigger pulse is always linked with a response pulse. In the imperfect associations, there is a response pulse for 70% of the trigger pulses. We also investigated the effects of study duration (6, 12, and 24 hour study durations) and sampling interval (5 vs. 10 min sampling intervals) for Scenario 2.

For each scenario, 100 sets of bivariate hormone concentrations, each with 10 subjects, were simulated. Typical mechanistic studies have 10–15 individuals in a group given the intense sampling protocol [see (Al-Safi et al., 2015), among many others]. For each simulation, joint hormone concentration series were simulated one subject at a time. We first drew the pulse locations of the trigger hormone from a multi-scale Strauss process (Penttinen, 1984) as described earlier. The parameter choices for level 1 were consistent with previous simulation studies (Mauger et al., 1995; Carlson et al., 2009; Horton et al., 2017). The parameter choices generated trigger pulse series similar to human female LH in the follicular phase. Given the simulated trigger locations, we simulated a response pulse as a Bernoulli trial with the probability of a response pulse occurring with probability ρ (Table 3). When a response pulse occurred, the location of the response pulse was simulated according to a normal distribution centered at the trigger pulse location with a standard deviation v (Table 3). Once the pulse locations were simulated, we completed simulating each individual hormone series as in Horton et al. (2017) by sampling individual level parameter values and then pulse specific parameter values. Given the pulse and individual specific parameter draws, we computed the true concentration curve according to the formula in Section 2.1, simulated model error and back transformed to the natural scale.

We also assessed the performance of the joint model under misspecified trigger-response relationships and a case when the response hormone was not actually pulsatile, but assumed pulsatile. Scenario 4: The trigger and response pulse locations and masses were assumed independent (no pulse association). In this case, both the trigger and response locations were generated from the multiscale Strauss process. Scenarios 5 and 6: The response pulse was triggered by two hormones (one observed, one not observed). In these scenarios, additional response pulses were generated from a homogeneous Poisson process with rates shown in Table 3. Scenario 7: We simulated response hormone concentration profiles by drawing each log observed concentration from a Normal distribution $\mathcal{N}(\log 4, 0.005)$ and then transformed back to the natural scale. This scenario allowed us to assess the performance of the algorithm if the response hormone is mistakenly assumed to be pulsatile.

Results from the analyses were summarized by the posterior mean along with 95% equal tail area credible intervals. The mean of the posterior mean was computed and compared to truth, along with mean values of the boundaries of the 95% equal tails credible intervals across the simulations. False positive (FP) and false negative rates (FN) were computed by comparing whether there were estimated and true pulse location within 15 minutes of one another and averaged across all individuals and simulations for a scenario. All calculations were performed in R statistical software v2.15 (R Core Team, 2012).

5.2. Simulation Results

Figure 2 show a representative example series from simulation Scenario 1 (top panel) and Scenario 2 (bottom panel). The columns are median posterior predictive fits with corresponding 95% posterior predictive credible intervals from the new joint model (left) and the single hormone model (right). When the response hormone has large pulses (Scenario 1) the fits and the pulse locations are very similar between the two approaches. However, no measures of association of the trigger and response pulses are available from the single hormone fit. When the response hormone has a weaker signal (Scenario 2), the posterior estimates of the pulse locations are improved (tighter peaks in the posterior) and a greater number of the correct response pulses are detected when using the joint model (Web Table S2). As might be expected, the fitted curves do not differ substantially between the approaches because the pulses that do not get detected in a single hormone analysis are smaller and less impactful on the observed hormone curves.

With a 24 hour study duration, the parameters describing associations between trigger and response hormone concentrations were estimated with negligible bias using the joint model (Table 4) for Scenarios 1–3. The bias in the number of response pulses per trigger pulse, ρ , was less than 0.06 in the different simulation scenarios and coverage of the 95% credible intervals was $>85\%$. Bias in the spread of response pulses, ν , was less than 1.8 minutes (truth was 10 minutes) and the coverage was $>82\%$. The bias in the correlation of the mean pulse masses was also small at less than <0.06 with coverage rates of 100%. When the pulse association was perfect (Scenario 3) ρ is estimated close to one (1.02), showing that this model is also applicable to one-to-one associations. Web Table 2 shows FP and FN rates. The FN rates were lower using the joint model for Scenarios 1–3. The FP rate for the response hormone using the joint model was also improved in these scenarios; however, the

FP rate for the trigger hormone was slightly higher for the joint model, but still low ($<6.7\%$). Further investigation showed the FP pulses were being caused by pulses occurring at the boundaries.

Posterior predictive fits for the various study durations are provided in Web Figure S3. When the study duration was shorter, the bias increased (Table 4). The number of response pulses per trigger pulse was slightly underestimated (-0.16 and -0.13 for the 12 and 6 hour studies, respectively). The coverage rate of the 95% credible intervals was $>89\%$. The bias of the spread of the response pulse locations was largest (but still small) for the 6 hour design (2.8 minutes). The coverage was $>92\%$. Shorter study duration was associated with higher FN and FP rates in both hormones for both models (Web Table S2). When the study duration is 12 hours, the joint model outperformed the single hormone model. With only a 6 hour study duration the results were more mixed. The FN rate was higher for the joint model with 10 minute sampling largely due to mis-identification of pulses at the boundaries and the improvements in FP using the joint approach were smaller.

For brevity, the performance in estimating population-level parameters (population mean pulse mass and width and population mean baseline and half-life) are summarized in Web Table S3. Estimates pulse parameters were close to truth in the above scenarios.

When the model was misspecified, meaning the association model was incorrect, the parameter values for ρ and ν become biologically implausible. When the trigger and response pulses locations were not associated (Scenario 4), but an association was assumed, the posterior median for ρ was large at 6.15 response pulses per trigger pulse and ν was even more extreme at 1142 minutes (Table 4 bottom panel). These are non-biologically plausible estimates that essentially indicate a response pulse can occur at any time in the study period. When there was a second unobserved trigger for the response hormone (Scenario 5), a moderately enlarged estimation of both ρ (1.3 response per trigger) and ν (27.1 minutes) were observed. As the second trigger becomes more frequent (Scenario 6), the posterior medians were even larger. When there was no pulsing in the response, ν remains large, and the posterior distributions of the pulse locations are diffuse across the range of the study (data not shown).

6. Discussion

We developed a flexible framework for jointly modeling pulse locations and hormone concentrations for two pulsatile hormones using a Cox cluster process as a joint model of pulse location processes. This joint model of the pulse location has a small set of biologically interpretable parameters and is appropriate for both one-to-one and non-one-to-one associations between two pulsing hormones. This flexibility does not exist in most joint models of pulsing hormones. When pulses in both hormone series are well defined, the joint model generates similar fits compared to fitting each hormone series separately. The joint model has improved pulse quantification compared to separate hormone analyses for studies of appropriate duration.

The ability to estimate the pulse association parameters is reduced as the study duration shortens. One hypothesis is that there are just too few coupled events for accurate estimation of the association. This hypothesis is perhaps confirmed by our simulation with more rapid sampling (5 minutes between observations) as parameter estimation is not improved. Given the pulsing rates in this investigation (which mimic reproductive hormones), a 24 hour study design is recommended. Because rapid sampling studies are intensive and expensive, there has been some push to reduce the study time. If complex hypotheses such as associations between hormones is of interest, the work provides evidence that the traditional longer study duration is valuable.

It was reassuring to find that incorrectly assuming a single trigger and response association is easily identifiable in the posterior estimates of both association parameters. When the trigger and response pulses locations were not associated, the posterior values for ρ and ν were quite large and biologically implausible. When there was a second unobserved trigger for the response hormone, a moderately enlarged estimation of both ρ and ν occurred. ν was the more revealing parameter for misspecification. Interestingly, the simulation findings for the second trigger model mimic the findings in the obesity study, suggesting perhaps that FSH has more than one trigger. There is debate in the literature on whether this occurs in humans (Hall et al., 1990; Thompson and Kaiser, 2014). This work may offer weak evidence in favor of this hypothesis. In theory, it is possible to extend the Cox process model to incorporate a second trigger and investigate the two-trigger model more extensively. One approach would be to add a scatter process to the intensity function

$$(\lambda^y(z | \tau^x, \rho, \nu, \varepsilon) = \sum_{k=1}^{N_i^x} \frac{\rho_{ik}}{\sqrt{2\pi\nu_{ik}^2}} \exp \left[-\frac{(z - \tau_{ik}^x)^2}{2\nu_{ik}^2} \right] + \varepsilon) \text{ as has been done in other applications of}$$

the Cox process (Lawson and Denison, 2002). However, fitting this extension likely requires some additional identifiability constraints on the pulse masses for accurate estimation. How to do this is the subject of future work.

Although this model was developed in the context of reproductive hormones, stress hormones (ACTH and Cortisol) and growth hormones are also pulsatile hormones that exist in complex hormone networks with trigger-response relationships. The deconvolution model has been applied to these hormones (Johnson, 2007; van Esdonk et al., 2017). Thus, this work is more broadly applicable for investigating association hypotheses in these axes as well.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This article was completed as part of Huayu Liu's dissertation at the University of Colorado Anschutz Medical Campus. Supported in part by National Institutes of Health [NIH/NICHD R01HD081161]. Contents are the author's sole responsibility and do not necessarily represent official NIH views. The authors thank the associate editor and referee for their thoughtful comments that helped improve the manuscript.

References

- Al-Safi ZA, Liu H, Carlson NE, Chosich J, Lesh J, Robledo C, et al. Estradiol priming improves gonadotrope sensitivity and pro-inflammatory cytokines in obese women. *The Journal of Clinical Endocrinology and Metabolism*. 2015; 100:4372–4381. [PubMed: 26425884]
- Carlson NE, Johnson TD, Brown MB. A Bayesian approach to modeling associations between pulsatile hormones. *Biometrics*. 2009; 65:650–659. [PubMed: 18759850]
- Cox, DR., Isham, V. *Point Processes*. London: Chapman and Hall/CRC; 1984.
- Geyer CJ, Møller J. A new look at the statistical model identification. *Scandinavian Journal of Statistics*. 1994; 21:359–373.
- Greenspan, FS., Gardner, DG. *Basic and Clinical Endocrinology*, Seventh Edition. New York: McGraw-Hill; 2004.
- Guo W, Brown MB. Cross-related structural time series models. *Statistica Sinica*. 2001; 11:961–979.
- Hall JE, Whitcomb RW, Rivier JE, Vale WW, Crowley WF. Differential regulation of luteinizing hormone, follicle-stimulating hormone, and free α -subunit secretion from the gonadotrope by gonadotropin-releasing hormone (GnRH): Evidence from the use of two GnRH antagonists. *The Journal of Clinical Endocrinology and Metabolism*. 1990; 70:328–335. [PubMed: 2105329]
- Horton KW, Carlson NE, Grunwald GK, Mulvahill MJ, Polotsky AJ. A population-based approach to analyzing pulses in time series of hormone data. *Statistics in Medicine*. 2017; 36:2576–2589. [PubMed: 28393373]
- Jiang B, Wang N, Sammel MD, Elliott MR. Modelling short- and long-term characteristics of follicle stimulating hormone as predictors of severe hot flashes in the Penn Ovarian Aging Study. *Journal of the Royal Statistical Society C*. 2015; 64:731–753.
- Johnson TD. Bayesian deconvolution analysis of pulsatile hormone concentration profiles. *Biometrics*. 2003; 59:650–660. [PubMed: 14601766]
- Johnson TD. Analysis of pulsatile hormone concentration profiles with nonconstant basal concentration: A Bayesian approach. *Biometrics*. 2007; 63:1207–1217. [PubMed: 18078483]
- Keenan DM, Roelfsema F, Veldhuis JD. Endogenous ACTH concentration-dependent driver of pulsatile cortisol secretion in the human. *American Journal Physiology*. 2008; 287:E652–E661.
- Keenan DM, Veldhuis JD, Yang R. Joint recovery of pulsatile and basal hormone secretion by stochastic nonlinear random-effects analysis. *American Journal Physiology*. 1998; 275:R1939–R1949.
- Kleinman KP, Ibrahim JG. A semiparametric bayesian approach to the random effects model. *Biometrics*. 1998; 54:921–938. [PubMed: 9750242]
- Lawson, AB., Denison, DGT. *Spatial Cluster Modeling*. London: Chapman and Hall/CRC; 2002.
- Mauger DT, Brown MB, Kushler RH. A comparison of methods that characterize pulses in a time series. *Statistics in Medicine*. 1995; 14:311–325. [PubMed: 7724916]
- Penttinen, A. *Modelling Interactions in Spatial Point Patterns: Parameter Estimation by the Maximum Likelihood Method*. Jyväskylän yliopisto; 1984.
- Pincus S, Padmanabhan V, Lemon W, Randolph J, Rees MA. Follicle-stimulating hormone is secreted more irregularly than luteinizing hormone in both humans and sheep. *Journal of Clinical Investigation*. 1998; 101(6):1318–1324. [PubMed: 9502773]
- Quintana FA, Johnson WO, Waetjenc LE, Gold EB. Bayesian nonparametric longitudinal data analysis with embedded autoregressive structure: Application to hormone data. *Journal of the American Statistical Association*. 2016; 111:1168–1181. [PubMed: 28366967]
- R Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing; 2012.
- Stephens M. Bayesian analysis of mixture models with an unknown number of components: an alternative to reversible jump methods. *Annals of Statistics*. 2000; 28:40–74.
- Strauss, J., Barbieri, R. *Yen and Jaffe's Reproductive Endocrinology : Physiology, Pathophysiology, and Clinical Management*. 6. Philadelphia, PA: Saunders/Elsevier; 2009.
- Thompson IR, Kaiser UB. GnRH pulse frequency-dependent differential regulation of LH and FSH gene expression. *Molecular and Cellular Endocrinology*. 2014; 385:28–35. [PubMed: 24056171]

- van Esdonk MJ, Burggraaf J, van der Graaf PH, Stevens J. A two-step deconvolution-analysis-informed population pharmacodynamic modeling approach for drugs targeting pulsatile endogenous compounds. *Journal of Pharmacokinetics and Pharmacodynamics*. 2017:1–12.
- Veldhuis JD, Johnson ML. Deconvolution analysis of hormone data. *Methods in Enzymology*. 1992; 210:539–575. [PubMed: 1584051]
- Zhang D, Lin X, Raz J, Sowers M. Semiparametric stochastic mixed models for longitudinal data. *Journal of the American Statistical Association*. 1998; 93:710–719.
- Zhang D, Lin X, Sowers M. Semiparametric regression for periodic longitudinal hormone data from multiple menstrual cycles. *Biometrics*. 2000; 56:31–39. [PubMed: 10783774]

Figure 1.

Observed (dots) and median posterior predictive fits and 95% posterior predictive intervals for LH (black) and FSH (grey). The histograms are posterior distributions of the pulse locations for LH (top histograms) and FSH (bottom histograms). The top panel (A) shows hormone levels for one subject in Study 2. The normal weight woman was sampled during her follicular phase for 24-hours every 7.5 minutes. The bottom panel are a NW (B) and an Obese (C) woman from Study 1. Patients were sampled for 6 hours at 10 minute intervals.

Figure 2.

Observed and fitted response hormone series in one randomly selected simulated subject from simulation scenario 1 (top panel) and scenario 2 (bottom panel). The left column are fits from the new joint model and the right column fits from the single hormone analyses. In each figure, the solid dots are observed hormone concentrations, the grey solid line is the median of the posterior predictive distribution of the concentration, and the grey dashed lines are the 95% posterior predictive intervals. The marks at the bottom of each figure are the true pulse locations for the trigger hormone (circle) and the response hormone (cross). The histograms represent the posterior distributions for the pulse locations.

Table 1

The unnormalized joint posterior distribution of the parameters given the data, $\mathbf{x}_i$ and $\mathbf{y}_i$. Following work by Horton et al. (2017), we incorporate population level parameters for pulse mass, width, baseline and half-life as we are jointly modeling a population of subjects. $\mathbf{x}_i$ and $\mathbf{y}_i$ are the vector of hormone concentration values for subject i .

Model component	Hierarchical Specification
(a) Likelihood	$\prod_{i=1}^m L(\mathbf{x}_i, \mathbf{y}_i \mid \boldsymbol{\theta}_{s,i}^x, \theta_{b,i}^x, \theta_{h,i}^x, \boldsymbol{\theta}_{s,i}^y, \theta_{b,i}^y, \theta_{h,i}^y, \boldsymbol{\Sigma}_{e,i}) \times$
(b) Joint prior pulse locations	$\prod_{i=1}^m \left[\pi(\boldsymbol{\tau}_i^y \mid \boldsymbol{\tau}_i^x, \rho, \nu) \pi(\boldsymbol{\tau}_i^x \mid \beta, R, \gamma) \right] \pi(\rho) \pi(\nu) \times$
(c) Joint prior pulse Mass & width	$\prod_{i=1}^m \left\{ \prod_{k=1}^{N_{s,i}^x} \pi((\boldsymbol{\alpha}_{s,ik}^x, \omega_{s,ik}^x)' \mid N_{s,i}^x, \boldsymbol{\mu}_{s,i}^x, \boldsymbol{\Upsilon}_s^x) \prod_{l=1}^{N_{s,i}^y} \pi((\boldsymbol{\alpha}_{s,il}^y, \omega_{s,il}^y)' \mid N_{s,i}^y, \boldsymbol{\mu}_{s,i}^y, \boldsymbol{\Upsilon}_s^y) \right\} \times$ $\prod_{i=1}^m \left[\pi((\boldsymbol{\mu}_{s,i}^x, \boldsymbol{\mu}_{s,i}^y)' \mid \boldsymbol{\mu}_s^{x,y}, \boldsymbol{\Sigma}_s^{x,y}) \right] \pi(\boldsymbol{\mu}_s^{x,y}) \pi(\boldsymbol{\Sigma}_s^{x,y}) \pi(\boldsymbol{\Upsilon}_s^x) \pi(\boldsymbol{\Upsilon}_s^y) \times$
(d) Priors baseline	$\prod_{i=1}^m \pi(\theta_{b,i}^x \mid \mu_{b,i}^x, \sigma_{x,b}^2) \pi(\mu_{b,i}^x) \pi(\sigma_{x,b}^2) \prod_{i=1}^N \pi(\theta_{b,i}^y \mid \mu_{b,i}^y, \sigma_{y,b}^2) \pi(\mu_{b,i}^y) \pi(\sigma_{y,b}^2) \times$
(e) Priors half-life	$\prod_{i=1}^m \pi(\theta_{h,i}^x \mid \mu_{h,i}^x, \sigma_{x,h}^2) \pi(\mu_{h,i}^x) \pi(\sigma_{x,h}^2) \prod_{i=1}^N \pi(\theta_{h,i}^y \mid \mu_{h,i}^y, \sigma_{y,h}^2) \pi(\mu_{h,i}^y) \pi(\sigma_{y,h}^2) \times$
(f) Priors model error	$\prod_{i=1}^m \pi(\boldsymbol{\Sigma}_{e,i})$

Top panel: Posterior distribution summaries of the pulse parameters in Study 1 with comparisons between normal weight (NW) and obese (OB) women.
Bottom panel: Posterior distribution summaries of the pulse parameters in Study 2 with comparisons between the joint and single hormone models. PM: Posterior median, 95% CI: 95% equal tailed credible interval, $\cdot$: The difference in normal weight minus obese, PP: Posterior probability.

Table 2

Study 1: Normal weight vs. obese women							
Parameters	Normal ($n = 10$)		Obese ($n = 12$)		(NW-OB)		NW>OB
	PM	95%CI	PM	95%CI	PM	95%CI	PP
LH Pulse # (6h)	6.4	(5.8, 7.0)	6.6	(6.1, 7.1)	-0.2	(-1.0, 0.6)	0.23
LH Pulse mass(IU/L)	2.86	(0.64, 4.01)	1.22	(0.27, 1.71)	1.68	(-0.66, 3.56)	0.89
FSH Pulse # (6h)	6.2	(4.4, 8.0)	4.0	(2.9, 5.3)	2.6	(0.6, 4.7)	0.99
FSH Pulse mass(IU/L)	1.63	(1.14, 2.36)	1.64	(0.94, 3.53)	-0.04	(-1.86, 1.06)	0.47
# of FSH per LH (ρ)	1.11	(0.78, 1.53)	0.58	(0.37, 0.86)	0.52	(0.085, 1.00)	0.99
Spread ($min; v$)	30.86	(18.54, 71.20)	29.13	(16.10, 66.74)	2.18	(-38.06, 41.08)	0.56

Study 2: A 24-hour natural history in normal weight women				
Parameters	Joint		Single	
	PM	95%CI	PM	95%CI
LH Pulse # (24 h)	17.9	(11.5, 20.0)	17.0	(9.1, 22.2)
LH Pulse mass(ng/ml)	2.6	(1.6, 3.3)	2.9	(1.8, 3.6)
FSH Pulse # (24 h)	11.9	(8.1, 12.3)	9.8	(7.3, 14.4)
FSH Pulse mass(ng/ml)	0.9	(0.7, 1.1)	0.9	(0.2, 1.5)
# of FSH per LH (ρ)	0.67	(0.47, 0.92)		
Spread ($min; v$)	5.0	(1.9, 8.2)		

Table 3

Table of the parameters combinations used for the simulation scenarios. The parameters listed are on the population level unless otherwise specified.

Parameter	Units (all)	Trigger (x)	Scenario						
			1	2	3	4	5	6	7
			Response (y)						
Pulse mass									
Pop. mean (μ_a)	ng/ml	3	3	1	3	3	1	1	
Pop. std. (σ_a)	ng/ml	0.5	0.5	0.5	0.5	0.5	0.5		
Subj. std. (v_a)	ng/ml	0.5	0.5	0.5	0.5	0.5	0.5		
Correlation ($\mu_{a,t}^x$ & $\mu_{a,t}^y$)		0.8	0.8	0.8	0.8	0.8	0.8	0.8	
Pulse width									
Pop. mean (μ_ω)	$min.^2$	25	25	25	25	25	25	25	
Pop. std. (σ_ω)	$min.^2$	2	2	2	2	2	2	2	
Subj. std. (v_ω)	$min.^2$	0.5	0.5	0.5	0.5	0.5	0.5	0.5	
Baseline									
Pop. mean (μ_b)	ng/ml	4	4	4	4	4	4	4	4
Pop. std. (σ_b)	ng/ml	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4
Half-life									
Pop. mean (μ_h)	min.	45	45	200	45	45	200	200	
Pop. std. (σ_h)	min.	2	2	2	2	2	2	2	
Model error (log-scale)									
σ_ε^2	log conc. ²	0.005	0.005	0.005	0.005	0.005	0.005	0.005	0.005
Trigger pulse strauss model									
Rate(β)		4							

Scenario							
	1	2	3	4	5	6	7
Parameter	Units (all)			Trigger (x)	Response (y)		
Repulsion Level 1 (γ_1)				0			
Repulsion Level 2 (γ_2)				0.1			
Repulsion dist. Level 1 (R_1)	min.			36			
Repulsion dist. Level 2 (R_2)	min.			96			
Response pulse model							
Cox model							
Cluster size (ρ)	0.7	1			0.7	0.7	
Cluster width (v)	10	10			10	10	
2nd Trigger Pulse rate					0.10	0.30	
Strauss model							
Rate(β)					4		
Repulsion Level 1 (γ_1)					0		
Repulsion Level 2 (γ_2)					0.1		
Repulsion dist. Level 1 (R_1)					36		
Repulsion dist. Level 2 (R_2)					96		

Table 4

Posterior distributions of association parameters under different scenarios. Unless specified, the sampling interval was 10 minutes. MPM: Mean of the posterior mean across all simulations. 95% CI: Mean of the boundaries of the 95% equal tailed credible intervals across the simulations.

	# of Res. per trigger, ρ			Spread, $\sigma(\text{min})$			Mass correlation, ρ_a		
	Truth	MPM	95% I	Truth	MPM	95% I	Truth	MPM	95% I
24 hour study									
Scenario 1 (Imperfect, large resp.)	0.7	0.74	(0.61,0.83)	10	10.7	(8.8,12.8)	0.7	0.75	(0.57,0.89)
Scenario 2 (Imperfect, small resp.)	0.7	0.76	(0.65,0.97)	10	11.8	(9.3,17.0)	0.7	0.70	(0.61,0.95)
Scenario 3 (Perfect, large resp.)	1.0	1.02	(0.93,1.06)	10	10.6	(9.8,11.2)	0.7	0.68	(0.57,0.84)
12 hour study									
Scenario 2	0.7	0.54	(0.37,0.76)	10	10.8	(7.1,16.0)	0.7	0.62	(0.05,0.98)
6 hour study									
Scenario 2	0.7	0.57	(0.33,0.89)	10	12.8	(6.5,23.7)	0.7	0.50	(0.02,0.95)
Scenario 2 (5 min. sampling)	0.7	0.79	(0.48,1.21)	10	14.5	(8.3,25.5)	0.7	0.64	(0.03,0.98)
Association model misspecified									
Scenario 4 (No association)	–	6.15	(0.18,11.38)	–	1142	(370,2350)	0	–0.20	(–0.98,0.95)
Scenario 5 (2 triggers, average of 2–3 extra resp. pulses)	0.7	1.30	(0.70,1.69)	10	27.1	(20.2,35.1)	0.7	0.25	(–0.73,0.96)
Scenario 6 (2 triggers, average 6–7 extra resp. pulses)	0.7	4.75	(0.16,11.43)	–	142	(8.2,459)	0.7	–0.02	(–0.91,0.92)
Scenario 7 (Resp. not pulsing)	–	0.24	(0.16,0.43)	–	74.0	(8.2,459)	0	0.12	(–0.80,0.89)