



Published in final edited form as:

Proc Int World Wide Web Conf. 2019 May ; 2019: 2999–3005. doi:10.1145/3308558.3313512.

Predicting pregnancy using large-scale data from a women's health tracking mobile application

Bo Liu[†],

Dept. of Computer Science, Stanford

Shuyang Shi[†],

Dept. of Computer Science, Stanford

Yongshang Wu[†],

Dept. of Computer Science, Stanford

Daniel Thomas,

Clue by BioWink GmbH, Berlin

Laura Symul,

Dept. of General Surgery and Dept. of Statistics, Stanford

Emma Pierson^{*},

Dept. of Computer Science, Stanford

Jure Leskovec

Dept. of Computer Science, Stanford Chan-Zuckerberg Biohub

Abstract

Predicting pregnancy has been a fundamental problem in women's health for more than 50 years. Previous datasets have been collected via carefully curated medical studies, but the recent growth of women's health tracking mobile apps offers potential for reaching a much broader population. However, the feasibility of predicting pregnancy from mobile health tracking data is unclear. Here we develop four models – a logistic regression model, and 3 LSTM models – to predict a woman's probability of becoming pregnant using data from a women's health tracking app, Clue by BioWink GmbH. Evaluating our models on a dataset of 79 million logs from 65,276 women with ground truth pregnancy test data, we show that our predicted pregnancy probabilities meaningfully stratify women: women in the top 10% of predicted probabilities have a 89% chance of becoming pregnant over 6 menstrual cycles, as compared to a 27% chance for women in the bottom 10%. We develop a technique for extracting *interpretable* time trends from our deep learning models, and show these trends are consistent with previous fertility research. Our findings illustrate the potential that women's health tracking data offers for predicting pregnancy on a broader population; we conclude by discussing the steps needed to fulfill this potential.

^{*}To whom correspondence should be addressed: emmap1@cs.stanford.edu.

[†]indicates co-first authorship.

Keywords

Pregnancy prediction; mobile health tracking

1 INTRODUCTION

Predicting pregnancy is a fundamental problem in women's health [11, 13, 27, 41]. Modeling the probability of pregnancy, and identifying behaviors which affect it, allows for more efficient family planning. Identifying people with a very low chance of becoming pregnant is important for targeting and treating infertility, which can be a devastating experience [18]. Because of the importance of predicting and facilitating pregnancy, the fertility services market is predicted to grow to \$21 billion by 2020 [38].

Attempts to predict pregnancy go back decades, often using Bayesian methods [11, 13, 27, 41] to capture the complex dynamics which govern fertility. Many factors influence fertility, including age [12] and medical conditions [24, 30]. Fertility fluctuates over the course of the menstrual cycle (which averages 28-29 days in length, although this varies across women [8]), peaking during the "fertile window", around 14 days before the end of the cycle. For women with a regular cycle of 28-29 days, the fertile window occurs approximately midway through the cycle [3, 50]. Because the timing of the fertile window varies, and intercourse during the fertile window is much more likely to result in pregnancy, fertility prediction methods often focus on detecting and predicting the fertile window, using ovulation tests and measurements of cervical mucus and basal body temperature [5, 16].

Previous studies of fertility have relied on carefully curated data from medical studies [3, 11, 13, 27, 39-41, 43]: for example, participants are generally asked to regularly monitor critical features like basal body temperature and menstrual cycle starts and to take pregnancy tests after each menstrual cycle. Many people are unwilling or unable to engage in such careful monitoring, limiting the applicability of these methods to the general population.

In the last few years, the growing use of health tracking mobile apps offers a new potential data source for predicting pregnancy in broader populations. Health tracking apps have been recognized as delivering a "data bounty" [19] for healthcare [32, 49]. Within women's health specifically, health tracking apps are used by millions of women worldwide [7, 15, 37], and have already been used to study sexually transmitted infections [2], menstrual cycle fluctuations [35, 47], and menstrual cycle lengths [21]. These apps allow users to create profiles which include age and birth control information and keep daily logs of symptoms relevant to pregnancy, including basal body temperature, period starts, and ovulation tests. Crucially, they also allow women to log positive and negative pregnancy tests, providing a ground truth source of pregnancy data. Pregnancy data collected by women's health mobile apps differs from that collected in medical studies in three ways:

1. **Larger datasets.** The dataset we consider in this paper contains tens of millions of observations from tens of thousands of women, orders of magnitude more than previous datasets [14]. Complex models developed for these smaller datasets, which often rely on computationally intensive techniques like MCMC sampling,

will not scale to datasets from mobile health tracking apps, necessitating new methods.

2. **Broader populations.** While previous fertility studies have predominantly focused on filtered, mostly healthy or proven fertile populations that volunteer for research within a single country, health tracking apps can be used by anyone with a smartphone, and are used by women all over the world.
3. **Missing data.** Critical features like basal body temperature (BBT) and pregnancy tests are less reliably recorded in mobile health tracking datasets than in medical studies, making missing data a pressing concern.

Mobile tracking datasets thus present new challenges but also new opportunities for predicting pregnancy in a much broader population. To this end, a number of companies have announced initiatives to predict pregnancy or fertility using mobile tracking data [1, 28, 33, 42]. To date, however, it is unclear whether their fertility prediction algorithms are reliable, because they are proprietary (so they cannot be compared to or independently assessed) and their efficacy has been disputed [26, 31, 46].

This work.

Here we present what is to our knowledge the first study of pregnancy prediction using fully described methods applied to large-scale women's health tracking app data. Using a dataset from a women's health tracking app which includes ground truth data on pregnancy tests, we develop and assess four pregnancy prediction models — an interpretable logistic regression model, and three LSTM models — which integrate traditional fertility modeling methods and modern time series prediction methods. We first show that our models meaningfully stratify women by their probability of becoming pregnant: women in the top 10% of predicted probabilities have a 89% chance of becoming pregnant over 6 menstrual cycles, while women in the bottom 10% have only a 27% chance. We further provide an intuitive technique for extracting *interpretable* time trends from our LSTM models, overcoming a common failing of deep learning models, and show that these time trends are consistent with previous fertility research. Finally, we discuss the steps that should be taken to maximize the efficacy of pregnancy prediction from women's health tracking data.

2 PROBLEM SETUP

Our task is predicting, on the basis of the health tracking data a woman logs during a single menstrual cycle, whether she will log a positive pregnancy test at the end of that cycle. We first describe our dataset and then describe our task in more detail.

2.1 Dataset

We use data from the Clue women's health tracking app [17]. The app has been rated the most accurate menstrual cycle tracking app by gynecologists [29] and previously used in studies which show that it reliably replicates known women's health findings [2, 21, 35]. All data is de-identified and analysis was determined to be exempt from review by the Stanford Institutional Review Board.

Features.—Our dataset consists of two types of features:

- **User features.** Users can record age and birth control method (e.g., “None”, “Condoms”, “IUD”, or various types of birth control pill). We encode age using a two-element vector: a binary element which indicates whether age data is missing, and a continuous element which is the value of age if it is present, and 0 otherwise. We encode birth control using a vector with indicator variables for each type of birth control; the vector is all zeros if birth control data is missing.
- **Daily feature logs.** Each row in this dataset consists of one log of one feature for one user on one date. Users can log 110 *binary* features, including positive/negative pregnancy tests (which we use to define ground truth, and do not include in the inputs to our model), period bleeding, mood and behavior features (eg, happy mood or exercise:running), and taking daily birth control. Users can also log 3 *continuous* features: basal body temperature (BBT), resting heart rate, and body weight. We list all daily features in Table 1.

Data filtering.—To ensure that users are regularly using the app, we filter for users who enter at least 300 daily feature logs. We apply basic quality control filters to the three continuous features to remove unreliable values (for example, $BBT > 110^{\circ}F$). To minimize the chance that users are already pregnant, we filter out all cycles after a user has logged a positive pregnancy test.

Defining menstrual cycle starts.—Because fertility fluctuates over the course of the menstrual cycle, defining menstrual cycle starts – when the period begins – is necessary to predict pregnancy. We define a *cycle start* as a start of bleeding after the user has not recorded any bleeding for at least 7 days. So, for example, if the user records bleeding on May 1, 2, and 29, the cycle starts would be May 1 and May 29. A user’s *cycle day* is the number of days since their most recent cycle start, with 0 denoting the day of cycle start.

Encoding continuous features.—We encode continuous features (eg, weight) using a two-element vector: a binary element which indicates whether data is present or missing, and a continuous element which is the value for the feature if present, and 0 otherwise. For each cycle and each user, we subtract off the mean for each continuous feature (since the *change* in features like BBT is most indicative of cycle phase and therefore fertility [6]). We include each user’s average value of the three continuous features as a user-specific feature.

2.2 Prediction task

Our task is predicting from the first 24 days of a user’s cycle whether she will become pregnant in that cycle. We choose the 24-day interval because most women are very unlikely to become pregnant past this interval [50]. Our prediction task is binary, and each example is one cycle for one woman. We define positive examples as cycles in which the woman logs a positive pregnancy test after day 24 of the cycle and before day 24 of the next cycle (Figure 1); we define negative examples as cycles with a negative pregnancy test and no positive pregnancy test. Thus, our dataset consists only of cycles followed by positive or negative pregnancy tests¹. Importantly, predicting whether a test will be positive or negative is a

difficult and useful prediction task, because it occurs in settings where the woman herself is sufficiently uncertain about whether she is pregnant that she believes taking a test is worthwhile. Our dataset consists of 16,580 positive cycles, 88,685 negative cycles, 65,276 women, and 79,423,281 symptom logs. The proportion of cycles which are positive in our dataset (16%) is consistent with previous work [10]. We use data from 90% of users for training/validation and report results on a test set of the remaining 10% of users. We train on a dataset balanced for positive and negative examples by down-sampling each training batch (a standard technique for training on unbalanced data [23]). Following standard practice, we report all results on the unbalanced test dataset.

3 PREVIOUS WORK

Our predictive models rely on ideas both from the fertility modeling literature and from the deep learning time series prediction literature. We now summarize work in both areas.

3.1 Fertility modeling

We briefly summarize three main approaches in the extensive fertility modeling literature; [14] provides a lengthier review.

1. **Time-to-pregnancy (TTP) models** [20, 36,44,48] assume each individual menstrual cycle results in pregnancy with some probability μ which can vary across couples and over time. These models are useful for capturing the high level covariates which influence μ (eg, age) but they do not capture detailed daily dynamics within a single cycle (eg, sex during the fertile window is more likely to result in pregnancy). Therefore, they are less useful for modeling the health-tracking datasets we consider, whose strength is their detailed daily information.
2. **Barrett-Marshall and Schwartz (BMS) models** [3,11,13, 39, 43, 45] allow for more detailed modeling of daily activity within each cycle. They assume each act of sex contributes independently to the probability a cycle results in pregnancy. The probability that someone becomes pregnant in a cycle is thus

$$1 - \prod_d (1 - f_d)^{S_d},$$

where d is the cycle day, f_d is the probability of getting pregnant by having sex only on day d , and S_d is a binary variable indicating whether sex occurred on day d . The term inside the product is the probability that day d does not result in pregnancy; it is 1 if S_d is 0, indicating that no sex occurred, and $1 - f_d$ if S_d is 1. Schwartz [43] extends this model by assuming that the overall probability of pregnancy also depends on a couple-specific parameter which captures, for

¹We define positive and negative examples in this way to mitigate missing data concerns. The alternative (i.e. assuming that any cycle without a positive pregnancy test is a negative example) is problematic for two reasons. First, some users may not bother to log positive tests, so the negative label is a false negative. Second, some users may not bother to log birth control, concealing the fact that they have very little chance of getting pregnant—consistent with this, we find that the fraction of positive examples using this method, even for users who log no birth control, is far lower than the previous literature would imply.

example, the fact that some couples are infertile irrespective of sexual activity patterns. However, this additional parameter creates model identifiability problems.

3. **Extended Time-to-Pregnancy (ETTP) models** [9,40] provide an approximation to the BMS model by assuming that only sex on the most fertile day contributes to the probability of pregnancy, and sex on other days has no effect. As this model is only a pragmatic approximation to the true generative process, we instead develop a computationally efficient implementation of the BMS model.

All three of the above models were developed for datasets orders of magnitude smaller than the one we consider. Consequently, they perform parameter inference by estimating the posterior distribution, often using computationally intensive techniques like MCMC, and will not scale to our dataset. Therefore, in this work, we extend the BMS model using a scalable deep-learning-based model. Our extension scales to health-tracking datasets because it relies on more efficient backpropagation-based optimization.

3.2 Deep learning for time series prediction

Long Short Term Memory (LSTM) neural networks [22] have been successfully used to model medical time series data [4, 25, 34]. They maintain a hidden state at each timestep which can be used for prediction, and are a natural choice for prediction on time series with discrete timesteps (in our case, days of the menstrual cycle). Because they rely on more scalable backpropagation methods, they also scale to large health-tracking datasets.

4 MODELS

We develop four models. We select model hyperparameters – learning rate, hidden size, number of layers, batch size, dropout rate, and regularization strength – using grid search. Our models are publicly available at <https://github.com/AndyYSWoo/pregnancy-prediction>.

1. **Logistic regression.** As an interpretable baseline, we use logistic regression with coefficients for each feature and each cycle day (for example, having unprotected sex on cycle day 14), and coefficients for each user-specific feature. Our feature vector for each example consists of 2,771 features.
2. **LSTM.** We use an LSTM network with 24 timesteps (one for each cycle day). For each time-step, we feed as input the observed features for that cycle day (binary features plus continuous features, encoded as described in Section 2.1) along with the cycle averages for continuous features. The final hidden state of the LSTM is concatenated with the user-specific features and fed into a fully connected layer to produce the prediction.
3. **LSTM + BMS fertility model.** We develop an LSTM-based extension of the BMS model. As described in Section 3.1, the BMS model assumes each act of sex contributes independently to the probability a cycle results in pregnancy. Thus, we model a user's probability of becoming pregnant in a cycle as

$$\Pr(\text{pregnant}) = 1 - \prod_d \prod_t (1 - r_t \cdot f_d)^{s_{dt}}$$

where d is cycle day, t the type of sex (protected, unprotected, withdrawal, or none²), s_{dt} indicates if sex of type t was logged on day d , $f_d \in (0,1)$ is the contribution of sex on cycle day d towards pregnancy, and $r_t \in (0,1)$ is a learned parameter capturing the risk of a kind of sex (eg, $r_{\text{unprotected}} > r_{\text{protected}}$). The term inside the product is the probability that an act of sex of type t on cycle day d does *not* result in pregnancy. We model f_d as a function of the daily features x_d using an LSTM, as illustrated in Figure 2.

4. **LSTM + user embeddings.** Using only the user's current cycle does not account for their full history: eg, a user who has frequently logged unprotected sex for a long time but has not yet become pregnant may be less fertile. To incorporate information prior to the current cycle, we use a "user history LSTM" to encode the 180 days of user history prior to the user's current cycle, then use this LSTM's final state as a user embedding vector which we concatenate onto the other features and feed into a second LSTM as before. We jointly train both LSTMs. Figure 3 illustrates the model architecture.

5 RESULTS

5.1 Predictive performance

Metrics.—We evaluate model performance using AUC. Importantly, becoming pregnant is inherently a somewhat random process: it is impossible to *guarantee* a woman will become pregnant in a particular cycle. Consequently, achieving a very high AUC on our task is unlikely to be feasible, and we would not expect AUC to be very high. Nonetheless, we report AUC because it is a standard metric. As a second metric, we compare the probability p_{preg} that a user will log a positive pregnancy test after a cycle for the users in the top 10% of predicted pregnancy probabilities and users in the bottom 10% of predicted probabilities. (Note that p_{preg} is the true, not the predicted, probability of a positive pregnancy test.) However, the goal of fertility counseling is not to guarantee a woman will get pregnant in a particular cycle, but rather that over some reasonable time period (eg, 6 cycles) she has a good chance of becoming pregnant if she continues her current pattern of behavior. Thus, we also compute the 6-cycle probability of becoming pregnant, assuming that each cycle contributes independently to the probability of pregnancy and that woman continues her current pattern of behavior. For example, if a woman's single-cycle p_{preg} is 0.2, her six-cycle p_{preg} is $1 - (1 - 0.2)^6$.

Results.—The LSTM with user embeddings has the highest AUC (0.67) (Table 2). This model is able to meaningfully stratify users by their probability of getting pregnant: for example, the top 10% of users have an 89% chance of pregnancy over 6 cycles, whereas the

²We allow the model to learn a non-zero probability of getting pregnant even if no sex is logged to account for missing data, where users neglect to log sex but still get pregnant; this significantly improves the BMS model's predictive accuracy.

bottom 10% have only a 27% chance. While the other three models have somewhat lower AUCs, all four models usefully stratify users by pregnancy probability. The BMS model slightly worsens LSTM performance, possibly because of the more restrictive assumptions of its probability model – that is, that the overall probability of not getting pregnant in a cycle is the product of the probabilities of not getting pregnant on each cycle day. This demonstrates that traditional fertility models may not yield optimal predictive performance on mobile health datasets.

5.2 Interpretability

We next assess whether our models learn *interpretable* time trends which are consistent with prior fertility research. Because sexual activity is the feature most fundamentally associated with pregnancy, we examine time trends for the three types of sex logged in Clue data: unprotected, protected, and withdrawal sex.

Interpreting the logistic regression model.—Extracting time trends from the logistic regression model is straightforward, because the model is specifically developed to be interpretable. To understand how a feature contributes to the probability of pregnancy, we simply plot the coefficients for logging the feature on each day (Figure 4, left). The daily feature which is most strongly associated with positive pregnancy tests (averaging weights across all cycle days) is unprotected sex; the feature which is most strongly associated with negative pregnancy tests is protected sex.

Interpreting the LSTM models.—Interpreting deep learning models is much less straightforward, since they are high-dimensional and nonlinear and their coefficients do not have clear meanings. We use the following technique to quantify how logging binary feature b on day d influences the model's inferred probability of pregnancy:

1. For each example in our test set, we set binary feature b to 1 on day d , and compute the modeled probability of pregnancy for the example, $\Pr(\text{preg}|x_{bd}=1)$.
2. For each example in our test set, we set binary feature b to 0 on day d , and compute the modeled probability of pregnancy for the example, $\Pr(\text{preg}|x_{bd}=0)$.
3. We compute $\Pr(\text{preg}|x_{bd}=1) - \Pr(\text{preg}|x_{bd}=0)$, averaging across examples. This corresponds to the average difference in modeled probabilities when the user does and does not log feature b on day d .

Figure 4 (right) shows results for the LSTM model with the best predictive performance (LSTM + user embeddings); results for the other two LSTM models are qualitatively similar. Like the logistic regression model, the LSTM model learns that unprotected sex (blue) is more positively associated with positive tests than protected sex (green) or withdrawal sex (orange). Unprotected sex during the middle of the cycle shows the strongest positive association.

Consistency with prior research.—These modeled time trends are consistent with previous fertility research, which finds that unprotected sex during the “fertile window” [50] is most likely to result in pregnancy. The modeled increase in pregnancy probability due to

unprotected sex on a single day is fairly small (<2% for the LSTM model); a plausible explanation for this is that many users are taking birth control. The negative weights for protected and withdrawal sex are of interest because they indicate that the associations learned by the model do not necessarily have causal interpretations. Obviously, having protected sex never reduces the chance of pregnancy in a *causal* sense; rather, it indicates that the woman is trying to avoid becoming pregnant. Overall, this analysis demonstrates that we can extract interpretable, biologically plausible time trends for sexual activity, the feature most fundamentally associated with pregnancy, for both the simple logistic regression model and the predictively superior LSTM models.

6 DISCUSSION

We develop four models for pregnancy prediction, combining ideas from both classical fertility modeling and modern deep learning techniques, and assess them using women's health tracking mobile data. We show we can stratify women by their probability of becoming pregnant and learn interpretable time trends consistent with prior fertility research. We develop a simple, intuitive technique to extract time trends from our LSTM models which is more broadly applicable to other time series datasets.

Based on our study, we recommend two steps to enable women's health tracking apps to reach their full potential for pregnancy prediction. First, companies should fully describe the algorithms used for pregnancy prediction; this will facilitate constructive criticism and increase confidence in the algorithms. Second, incentivize users to provide more complete data. In the Clue dataset, most users do not log pregnancy tests or predictively useful features like BBT; encouraging users to provide this data will improve predictions.

These two improvements can facilitate prediction tasks which build on the one described here. For example, early identification of people who are very unlikely to become pregnant can allow faster referrals to infertility treatment. Modeling the behaviors which *causally* contribute to pregnancy can be used to recommend behavior: for example, having sex during a person-specific fertile window. Progress on these prediction tasks will improve well-being for the millions of people pursuing pregnancy.

Acknowledgments.

We thank Paula Hillard, Pang Wei Koh, Chris Olah, Nat Roth, and Camilo Ruiz for helpful comments; Amanda Shea at Clue for data assistance; all Clue users who volunteered their data for research; and the Hertz, NDSEG, and Swiss National Science Foundations for funding. This research was supported in part by the Stanford Data Science Initiative and the Chan Zuckerberg Biohub.

REFERENCES

- [1]. 2017. Glow and National Institutes of Health Collaborate to Advance Fertility Model. PR Newswire (2017).
- [2]. Alvergne Alexandra, Marija Vljajic Wheeler, and Vedrana Höggqvist Tabor. 2018. Do sexually transmitted infections exacerbate negative premenstrual symptoms? Insights from digital health. *Evolution, Medicine, and Public Health* (2018).
- [3]. Barrett John C and Marshall John. 1969. The risk of conception on different days of the menstrual cycle. *Population Studies* 23, 3 (1969), 455–461. [PubMed: 22073960]

- [4]. Inci M Baytas, Cao Xiao, Zhang Xi, Wang Fei, K Jain Anil, and Zhou Jiayu. 2017 Patient subtyping via time-aware LSTM networks. In Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining ACM, 65–74.
- [5]. Bigelow Jamie L, Dunson David B, Stanford Joseph B, Ecochard René, Gnath Christian, and Colombo Bernardo. 2004. Mucus observations in the fertile window: a better predictor of conception than timing of intercourse. *Human Reproduction* 19, 4 (2004), 889–892. [PubMed: 14990542]
- [6]. Buxton Charles L and Atkinson William B. 1948. Hormonal factors involved in the regulation of basal body temperature during the menstrual cycle and pregnancy. *The Journal of Clinical Endocrinology* 8, 7 (1948), 544–549.
- [7]. Chakradhar Shraddha. 2018. Discovery cycle. *Nature Medicine* (2018).
- [8]. Chiazzè Leonard, Brayer Franklin T, Macisco John J, Parker Margaret P, and Duffy Benedict J. 1968. The length and variability of the human menstrual cycle. *JAMA* 203, 6 (1968), 377–380. [PubMed: 5694118]
- [9]. Clayton David and Ecochard René. 1997 Artificial insemination by donor: discrete time survival data with crossed and nested random effects. In Proceedings of the First Seattle Symposium in Biostatistics Springer, 99–122.
- [10]. Colombo Bernardo, Masarotto Guido, and Menstrual Cycle Fecundability Study Group. 2000. Daily fecundability: first results from a new data base. *Demographic Research* 3 (2000).
- [11]. Dunson David B. 2001. Bayesian modeling of the level and duration of fertility in the menstrual cycle. *Biometrics* 57, 4 (2001), 1067–1073. [PubMed: 11764245]
- [12]. Dunson David B., Baird Donna D., and Colombo Bernardo. 2004. Increased Infertility With Age in Men and Women. *Obstetrics & Gynecology* (2004).
- [13]. Dunson David B and Stanford Joseph B. 2005. Bayesian inferences on predictors of conception probabilities. *Biometrics* 61, 1 (2005), 126–133. [PubMed: 15737085]
- [14]. Ecochard René. 2006. Heterogeneity in fecundability studies: issues and modelling. *Statistical Methods in Medical Research* 15, 2 (2006), 141–160. [PubMed: 16615654]
- [15]. A Epstein Daniel, Lee Nicole B, H Kang Jennifer, Agapie Elena, Schroeder Jessica, Pina Laura R, Fogarty James, Kientz Julie A, and Munson Sean. 2017 Examining menstrual tracking to inform the design of personal informatics tools. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems ACM, 6876–6888.
- [16]. Fehring Richard J. 2002. Accuracy of the peak day of cervical mucus as a biological marker of fertility. *Contraception* 66, 4 (2002), 231–235. [PubMed: 12413617]
- [17]. GmbH Biowink. 2017 Clue by Biowink GmbH. helloclue.com
- [18]. Greil Arthur L. 1997. Infertility and psychological distress: a critical review of the literature. *Social Science & Medicine* 45, 11 (1997), 1679–1704. [PubMed: 9428088]
- [19]. Hayden Erika Check. 2016. Mobile-phone health apps deliver data bounty. *Nature* 531, 7595 (2016), 422–423. [PubMed: 27008946]
- [20]. Heckman James J, Robb Richard, and Walker James R. 1990. Testing the mixture of exponentials hypothesis and estimating the mixing distribution by the method of moments. *J. Amer. Statist. Assoc.* 85, 410 (1990), 582–589.
- [21]. Adams Hillard Paula J and Wheeler Marija Vljajic. 2017. Data from a Menstrual Cycle Tracking App Informs our Knowledge of the Menstrual Cycle in Adolescents and Young Adults. *Journal of Pediatric and Adolescent Gynecology* 30, 2 (2017), 269–270.
- [22]. Hochreiter Sepp and Schmidhuber Jürgen. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780. [PubMed: 9377276]
- [23]. Japkowicz Nathalie. 2000 The class imbalance problem: Significance and strategies. In Proceedings of the International Conference on Artificial Intelligence.
- [24]. Kamel Remah M. 2010. Management of the infertile couple: an evidence- based protocol. *Reproductive Biology and Endocrinology* (2010).
- [25]. Lipton Zachary C, Kale David C, Elkan Charles, and Wetzel Randall. 2015. Learning to diagnose with LSTM recurrent neural networks. *arXiv preprint arXiv:1511.03677* (2015).

- [26]. Lomas Natasha. 2018. Clinic reports Natural Cycles app for 37 unwanted pregnancies since September. TechCrunch (2018).
- [27]. Lum Kirsten J, Sundaram Rajeshwari, Buck Louis Germaine M, and Louis Thomas A. 2016. A Bayesian joint model of menstrual cycle length and fecundity. *Biometrics* 72, 1 (2016), 193–203. [PubMed: 26295923]
- [28]. Miller Courtney. 2013. How the Period and Fertility Prediction Feature Works. Kindara Blog (2013).
- [29]. Moglia Michelle L, Nguyen Henry V, Chyjek Kathy, Chen Katherine T, and Castaño Paula M. 2016. Evaluation of smartphone menstrual cycle tracking applications using an adapted APPLICATIONS scoring system. *Obstetrics & Gynecology* 127, 6 (2016), 1153–1160. [PubMed: 27159760]
- [30]. Moos Merry-K., Dunlop Anne L., Jack Brian W., Nelson Lauren, Coonrod Dean V., Long Richard, Boggess Kim, and Gardiner Paula M. 2008. Healthier women, healthier reproductive outcomes: recommendations for the routine care of all women of reproductive age. *American Journal of Obstetrics & Gynecology* (2008).
- [31]. Morley Katie. 2018. Advert for “Natural Cycles” contraceptive app banned for exaggerating effectiveness. *The Telegraph* (2018).
- [32]. O’shea Catherine J, McGavigan Andrew D, Clark Robyn A, Chew Derek PB, and Ganesan Anand. 2017. Mobile health: an emerging technology with implications for global internal medicine. *Internal medicine journal* 47, 6 (2017), 616–619. [PubMed: 28580749]
- [33]. Pardes Arielle. 2018. In Contraceptive Tech, the App’s Guess Is as Good as Yours. *Wired* (2018).
- [34]. Pham Trang, Tran Truyen, Phung Dinh, and Venkatesh Svetha. 2016. Deepcare: A deep dynamic memory model for predictive medicine. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining Springer*, 30–41.
- [35]. Pierson Emma, Althoff Tim, and Leskovec Jure. 2018. Modeling Individual Cyclic Variation in Human Behavior. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web International World Wide Web Conferences Steering Committee*, 107–116.
- [36]. Potter Robert G. 1960. Length of the observation period as a factor affecting the contraceptive failure rate. *The Milbank Memorial Fund Quarterly* 38, 2 (1960), 140–152. [PubMed: 14488300]
- [37]. Rabin Roni Caryn. 2015. How Period Trackers Have Changed Girl Culture. *New York Times* (2015).
- [38]. Raphael Rina. 2018. Can Silicon Valley Get You Pregnant? *Fast Company* (2018).
- [39]. Royston J Patrick. 1982. Basal body temperature, ovulation and the risk of conception, with special reference to the lifetimes of sperm and egg. *Biometrics* (1982).
- [40]. Royston Patrick and Ferreira Alberto. 1999. A new approach to modeling daily probabilities of conception. *Biometrics* 55, 4 (1999), 1005–1013. [PubMed: 11315041]
- [41]. Scarpa Bruno and Dunson David B. 2007. Bayesian methods for searching for optimal rules for timing intercourse to achieve pregnancy. *Statistics in Medicine* 26, 9 (2007), 1920–1936. [PubMed: 17328097]
- [42]. Scherwitzl E Berglund, Lundberg O, Kallner H Kopp, Danielsson K Gemzell, Trussell J, and Scherwitzl R. 2017. Perfect-use and typical-use Pearl Index of a contraceptive mobile app. *Contraception* 96, 6 (2017), 420–425. [PubMed: 28882680]
- [43]. Schwartz D, MacDonald PDM, and Heuchel V. 1980. Fecundability, coital frequency and the viability of ova. *Population Studies* 34, 2 (1980), 397–400. [PubMed: 22077133]
- [44]. Sheps Mindel C. 1964. On the time required for conception. *Population Studies* 18, 1 (1964), 85–97.
- [45]. Stanford Joseph B and Dunson David B. 2007. Effects of sexual intercourse patterns in time to pregnancy studies. *American Journal of Epidemiology* 165, 9 (2007), 1088–1095. [PubMed: 17289774]
- [46]. Sudjic Olivia. 2018. “I felt colossally naive”: the backlash against the birth control app. *The Guardian* (2018).
- [47]. Symul Laura, Wac Katarzyna, Hillard Paula, and Salathe Marcel. 2018. Assessment of Menstrual Health Status and Evolution through Mobile Apps for Fertility Awareness. *bioRxiv* (2018).

- [48]. Tietze Sara L and Lincoln Richard. 1987 Differential Fecundity and Effectiveness of Contraception In Fertility Regulation and the Public Health. Springer, 129–134.
- [49]. Wilbanks John T and Topol Eric J. 2016. Stop the privatization of health data. Nature News 535, 7612 (2016), 345.
- [50]. Wilcox Allen J, Dunson David, and Baird Donna Day. 2000. The timing of the “fertile window” in the menstrual cycle: day specific estimates from a prospective study. British Medical Journal 321, 7271 (2000), 1259–1262. [PubMed: 11082086]


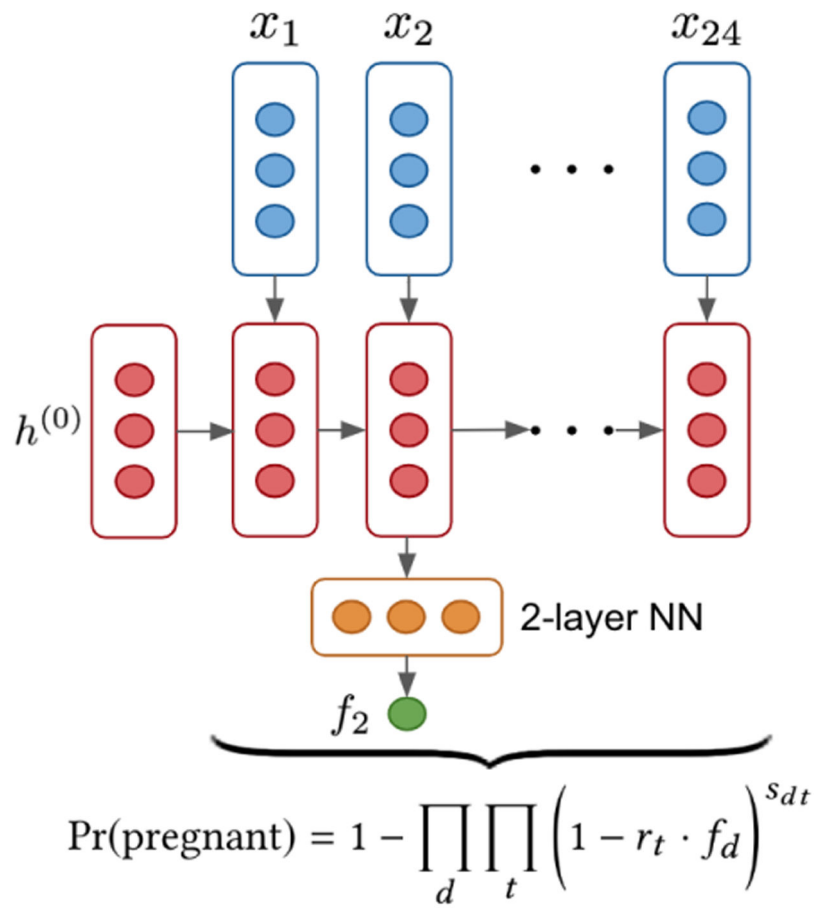



Figure 1:

Prediction task. The model makes predictions using logs from the first 24 days of a cycle (green interval), and the cycle is labeled using pregnancy tests taken after day 24 of the cycle and before day 24 of the next cycle (red interval). The vast majority of pregnancy tests in our dataset are taken near when the user's cycle is supposed to start, consistent with proper use of pregnancy tests, so any positive pregnancy tests likely result from activity during the green interval, which will be included in the feature vector.

**Figure 2:**

LSTM + BMS fertility model: we feed the daily features x (blue) into an LSTM model to obtain the hidden state h (red). f_d is then a function of the hidden state h parameterized by a neural network. The probability of becoming pregnant in a cycle is the function shown at bottom.

**Figure 3:**

LSTM + user embedding: we use the history features, x_h , in the previous $H = 180$ days for the user and feed them into the “user history LSTM” (shown at left). The final hidden state becomes the user embedding vector e , which is then concatenated with the daily features from the current cycle and fed into the pregnancy prediction LSTM (right).




Figure 4:

Model-learned time trends are interpretable for both the simple logistic regression model (left plot) and the best-performing LSTM + user embedding model (right plot). The horizontal axis is the cycle day. The vertical axis in the left plot is the logistic regression weight for logging a feature on that cycle day. The vertical axis in the right plot is how much logging a feature on a cycle day affects the LSTM-inferred probability of pregnancy. In both plots, positive y-values indicate associations with positive pregnancy tests, and negative y-values indicate associations with negative pregnancy tests. Both models learn that protected sex (green line) is negatively associated with pregnancy, while unprotected sex (blue line) is positively associated, particularly during the fertile window, and withdrawal sex (orange line) is intermediate.

All daily features which appear in our data. Each feature has both a category and a type. Features are binary except for those in the “continuous” category. While pregnancy tests are included in the binary features, we do not use them as predictive features. HBC=hormonal birth control; TNP=type not provided.

Table 1:

Category	Type
Ailment	Allergy, Cold/Flu Ailment, Fever, Injury
Appointment	Date, Doctor, Ob Gyn, Vacation
Collection Method	Menstrual Cup, Pad, Panty Liner, Tampon
Continuous	BBT, Resting Heart Rate, Weight
Craving	Carbs, Chocolate, Salty, Sweet
Digestion	Bloated, Gassy, Great Digestion, Nauseated
Emotion	Happy, PMS, Sad, Sensitive
Energy	Energized, Exhausted, High Energy, Low Energy
Exercise	Biking, Running, Swimming, Yoga
Fluid	Atypical, Creamy, Egg White, Sticky
Hair	Bad, Dry, Good, Oily
Injection HBC	Administered, Type Not Provided
IUD	Inserted, Removed, Thread Checked, TNP
Medication	Antibiotic, Antihistamine, Cold/Flu Medication, Pain
Mental	Calm, Distracted, Focused, Stressed
Motivation	Motivated, Productive, Unmotivated, Unproductive
Pain	Cramps, Headache, Ovulation Pain, Tender Breasts
Party	Big Night Party, Cigarettes, Drinks Party, Hangover
Patch HBC	Removed, Removed Late, Replaced, Replaced Late, TNP
Period	Heavy, Light, Medium, Spotting
Pill HBC	Double, Late, Missed, Taken, TNP
Poop	Constipated, Diarrhea, Great, Normal
Ring HBC	Removed, Removed Late, Replaced, Replaced Late, TNP
Sex	High Sex Drive, Protected, Unprotected, Withdrawal
Skin	Acne, Dry, Good, Oily
Sleep	0-3 Hrs, 3-6 Hrs, 6-9 Hrs, 9 Hrs, TNP

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Category	Type
Social	Conflict, Sociable, Supportive, Withdrawn
Test	Ovulation Neg, Ovulation Pos, Pregnancy Neg, Pregnancy Pos

Table 2:

Predictive performance of all models. The third column provides the probability a pregnancy test is positive for users in the top 10% vs bottom 10% of pregnancy risk; the fourth column provides the probability over six cycles.

Model	AUC	Single cycle p_{preg}	Six cycle p_{preg}
Logistic regression	0.63	28% vs 5%	86% vs 26%
LSTM	0.65	30% vs 4%	88% vs 23%
LSTM + BMS	0.64	29% vs 4%	87% vs 22%
LSTM + user embeddings	0.67	30% vs 5%	89% vs 27%