Modeling Heart Rate Response to Exercise with Wearables Data

Achille Nazaret*
Columbia University†
aon2108@columbia.edu

Sana Tonekaboni*
University of Toronto†
stonekaboni@cs.toronto.edu

Gregory Darnell
Apple
gdarnell@apple.com

Shirley Ren
Apple
shirleyr@apple.com

Guillermo Sapiro
Apple
gsapiro@apple.com

Andrew C. Miller
Apple
acmiller@apple.com

Abstract

Heart rate (HR) dynamics in response to workout intensity measure key aspects of an individual's fitness and cardiorespiratory health. Models of exercise physiology have been used to characterize cardiorespiratory fitness in well-controlled laboratory settings, but face additional challenges when applied to wearables in noisy, real-world settings. Here, we introduce a hybrid machine learning model that combines a physiological model of HR during exercise with complex neural networks in order to learn user-specific fitness representations. We apply this model at scale to a large set of workout data collected with wearables and show that it can accurately predict HR response to exercise demand in new workouts. We further show that the learned embeddings correlate with traditional metrics of cardiorespiratory fitness. Lastly, we illustrate how our model naturally incorporates and learn the effects of environmental factors such as temperature and humidity.

1 Introduction

An increase in development of wearable technologies has given individuals the power to track and monitor their overall health and well-being through daily activities. This leads to significant potential for machine learning (ML) models to uncover correlates of human health from wearable signals [1]. Previous work has shown success in a variety of applications from clinical monitoring tools to fitness and activity planners [2]. In the clinical domain, modern ML methods show strong predictive power for cardiovascular events from wearable data [3–5], they are successfully used in population disease surveillance [6, 7] and monitoring [8]. Wearable measures can also be predictive of overall health conditions that can be measured by different lab test [9]. These works clearly demonstrates that wearable data is a rich source of information that can provide insight into individual's overall health.

In this work, we introduce a scalable model that learns representations of health of an individual from the workout profiles gathered using wearable devices — step count, speed, elevation change, and heart rate. These personalized representations capture the response of heart rate to activity and are therefore representative of an individual's fitness and cardio-respiratory health. We introduce a personalized physiological model of heart rate (HR) based on ordinary differential equations (ODEs) that uses neural networks and representation learning to estimate user-specific parameters. We show that our personalized representations and model can accurately estimate the heart rate profile in response to a workout, by simply using an individual's workout history. Unlike most existing work in the literature that perform short term HR prediction [10, 11], our method can predict the full

Workshop on Learning from Time Series for Health, 36th Conference on Neural Information Processing Systems (NeurIPS 2022).

HR trend of a workout of up to 2 hours, even for those workouts the subject has yet to experience. This can be used in a variety of applications for personalized workout planning, and estimating HR zones or calories burned during a workout; furthermore, it shows that the representations summarize meaningful health information that well characterizes the HR response to a workout. To further investigate this, we show that the learned representations correlate well with traditional metrics of cardiorespiratory fitness.

2 Method

Study design and participants Our study uses workout measurements contributed to the Apple Heart and Movement Study [12] (AHMS). We analyze $270,707$ outdoor runs from $7,465$ subjects enrolled between 2019 and 2022. The AHMS provides the heart rate during each workout as well as four measures of the exercise intensity: speed from the pedometer and from the GPS, step cadence, and elevation gain. The sensor measurements are interpolated on a 10 seconds grid to form, for each workout w, a heart rate time-series $HR^{(w)} \in \mathbb{R}^d$ and a multivariate time-series of exercise intensity $I^{(w)} \in \mathbb{R}^{4\times d}$; where d is the duration of the workout. The workouts that we analyze are between 15 and 120 minutes long and also contains the weather information $W^{(w)}$ at the time of each workout.

Heart rate dynamics state equation. Several papers in sports physiology [13–16] have studied heart rate dynamics in response to exercise using ordinary differential equations (ODEs). These approaches translate the physical mechanisms of the human body into differential equations in order to incorporate domain knowledge in the modeling. This is an appealing method to build interpretable and ultimately trustworthy heart health models. The common approach introduces the body oxygen demand D as an intermediary quantity to link the exercise intensity I and the heart rate HR through the coupled ODEs:

$$\begin{cases} \dot{D}(t) &= B \cdot \left(f(I(t)) - D(t)\right), \\ \dot{HR}(t) &= A \cdot (HR(t) - HR_{min})^{\alpha} \cdot (HR_{max} - HR(t))^{\beta} \cdot (D(t) - HR(t)), \\ HR(0) &= HR_0 \quad \text{and} \quad D(0) = D_0. \end{cases}$$

In this dynamical system, f is a function translating the instantaneous activity intensity I into the necessary oxygen demand for I. The top equation attempts to match the current body oxygen demand D with the instantaneous demand $f(I)$. Parameter B controls how fast D adapts to $f(I)$. At the same time, the second equation drives the heart rate HR toward the pace required to deliver the demand D. Parameter A controls how fast the heart can adapt while the part with HR_{min}, HR_{max}, and α, β controls how difficult it is to reach the maximal heart rate or to rest down to the minimal heart rate.

To learn the parameters of this ODE, previous studies have limited their data collection to controlled laboratory environments on rarely more than ten individuals, and limited the form of $f(I)$ to simple functions. In this paper, we extend the ODE method to large scale uncontrolled environments and use it to model workout data from wearable devices.

Personalized large-scale heart rate model. A large scale study allows us to identify correlations between the ODE parameters across individuals and examine their evolution over time. Because these parameters capture the heart rate response to exercise, they are summarizing the fitness level and cardio-respiratory health of the subjects. In order to capture this information, we form a hierarchical model that relates the ODE parameters together.

We assume that the health state of individual i at date T can be represented by a low dimensional latent vector $z_{i,T} \in \mathbb{R}^{\ell}$. Then, we make each ODEs parameter a function of this representation. Each parameter's function, as well as function f, are learned as neural networks. Finally, we refine the physical model to incorporate the effect of weather $W =$ (temperature, humidity) into the demand equation, as well as the fatigue incurred over time t during workout. For instance, higher temperatures induce a higher oxygen demand [17]. We also parametrize these effects by neural networks $g(W)$ and $h(t)$. For health state z and intensity $t \mapsto I(t)$, the heart rate response in weather W is governed by:

$$\begin{cases} \dot{D}(t) &= B(z) \cdot (f(I(t), z) \cdot g(W) \cdot h(t) - D(t)), \\ \dot{HR}(t) &= A(z) \cdot (HR(t) - HR_{min}(z))^{\alpha(z)} \cdot (HR_{max}(z) - HR(t))^{\beta(z)} \cdot (D(t) - HR(t)), \\ HR(0) &= HR_0(z) \quad \text{and} \quad D(0) = D_0(z). \end{cases}$$

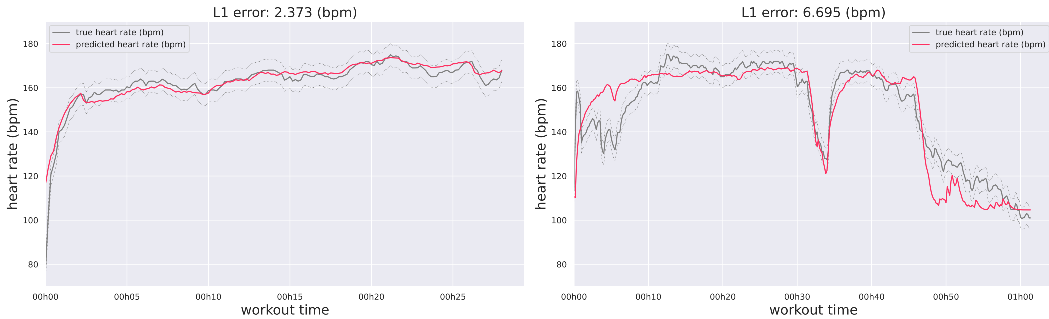

Figure 1: Examples of two HR predictions. The x-axis indicated time since the beginning of the workout and the Y-axis shows subject's instantaneous heart rate (beats per minute).

Inference of a large-scale heart rate model. In order to infer the user's health representations outside of the training set, and be able to study the evolution of a user's health over time to use the ODE model to predict future unobserved heart rates, we employ an amortized auto-encoder schema that concatenates user i's workout history up to T and encodes it into a health representation $z_{i,T}$.

$$z_{i,T} = e\left(HR^{(0)}, I^{(0)} ..., I^{(w^*)} HR^{(w^*)}\right), \quad \text{where } w^* \text{ is the last workout before date } T.$$

The encoder e is a convolutional neural network (CNN) with adaptive average pooling to accept variable input length. The ODE is solved using the Fourth Order Runge-Kutta method [18], such that it yields an HR solution that is differentiable against its input parameters. For training, we chain up the representation encoding with CNN followed by the heart rate ODE decoding and compute the Gaussian likelihood to form an objective function on which we run gradient descent. In practice, we subsample batches or portions of workouts and use stochastic gradient descent. The gradient updates simultaneously learn the representation encoder, and all the ODE internal neural networks parameters.

3 Results

Heart rate profile forecasting. The representation $z_{i,T}$ estimated using an individual's workout history can be used to predict the heart rate in future workouts. We measure the accuracy of heart rate prediction on workouts that were held-out for each subjects. Figure 1 shows two examples comparing the true heart rate to the heart rate estimated using our model. Note that for predicting HR for workout w happening at date T, our model only uses the workout intensity measures of that sample $I^{(w)}$ and the health representation $z_{i,T}$ – coming from encoding the previous workouts; i.e., no HR measurements $HR^{(w)}$ are ever observed by the model for making the predictions.

We compare the prediction performance of our model with a seq-to-seq deep model with similar modeling capacity and as a reference, a simple heuristic that predicts the HR as the mean of the subject-level historical HR observations. Table 1 shows the performance across these baselines and our model outperforms powerful deep recurrent models by a significant margin. We also measure the performance of our model in estimating the HR after the first 2 minutes of the workout. Indeed, it is difficult, if not impossible, for a model to predict the heart rate at the beginning of a workout. The initial heart rate depends on the user's activity prior to the workout, which is unobserved and unpredictable. Conversly, we hypothesize that the heart rate after two minutes can be explained by the user's activity in the first two minutes, that we observe.

Calories estimation. Estimating the calories burned during exercise can be accurately done using heart rate measurements during the workout with a linear formula [19]. This is useful for planning

Table 1: Heart rate reconstruction performance

Model	MAE (bpm)	Relative Error (%)
Our method	8.86 ± 4.54	5.0 ± 3.1
Our method (after 2 minutes)	6.88 ± 4.3	4.7 ± 3.2
Seq-to-Seq	12.81 ± 7.39	8.7 ± 5.0
Subject-level mean	16.38 ± 8.66	11.3 ± 7.8

Table 2: Predictive performance of heart rate zones

Population cohort	Full cohort	Female individuals	Male individuals
Accuracy	68.11 ± 21.73	67.96 ± 21.91	68.51 ± 21.07

Table 3: VO2Max prediction performance

	MSE	MAE	R2-score	Explained variance
Demographics only	30.53 ± 0.09	4.36 ± 0.006	0.36 ± 0.002	0.36 ± 0.002
ODE representations	16.2 ± 0.11	3.07 ± 0.01	0.66 ± 0.002	0.66 ± 0.002
Both	8.9 ± 0.10	2.16 ± 0.007	0.81 ± 0.003	0.81 ± 0.003

workouts based on calories burn goals and even more useful in cases where individuals are not wearing a wearable device that records heart rate. Our method can reliably estimate the amount of calories burned with 5% relative error (same relative error as the heart rate reconstruction), only using workout metrics that can be measured using a smart phone.

Heart rate zones prediction. Exercise heart rate zones are the percentage of an individual's maximum heart rate reached throughout the course of an exercise. Using our physiological model, we can predict heart rate zones using the workout data. This can help individuals plan personalized exercise routines to more effectively achieve their fitness goals. We define 6 zones (% intervals [0, 50, 60, 70, 80, 90, 100]) of maximum heart rate, and Table 2 shows the performance of our method on predicting the HR zone for the whole population, as well as different subgroups of the population.

Quantifying the impact of the weather on heart rate. Leveraging the interpretability of our ODE model, we analyze the learned neural network g and quantify the relative effect of weather on the body oxygen demand, this constitutes one of the largest study of this kind (over 270,000 workouts). Figure 2b shows an increase in body oxygen demand by up to 10% in high temperatures and humidity.

Learning about cardiorespiratory health. To show that our representations summarize information about cardiorespiratory health, we use a summary of cardio fitness, VO_2max, estimated by wearable devices. VO_2max is the maximum amount of oxygen the body can consume during exercise, normalized to body mass. This value is estimated using the heart and motion sensors on wearable devices. Using the health representations $z_{i,t}$, we aim to predict the estimated VO_2max with a simple linear regression model, and achieve an accuracy of ± 3 mL/(kg $\cdot$ min). Table 3 reports the performance of a linear regression model on the ODE representations only, on demographics only, or on both. Figure 2a shows a 2D projection of the health representation for different workouts where we can see the separation of higher and lower values of VO_2max.

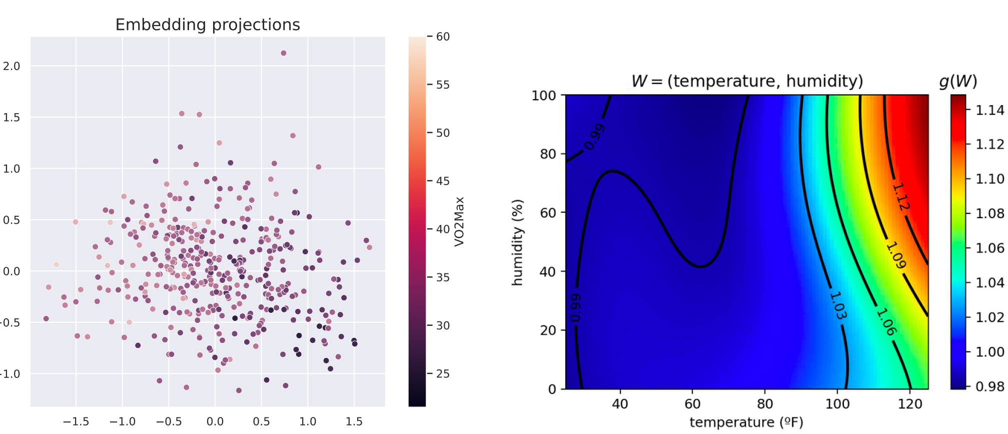

(a) PCA projection of health representations z (b) Impact $g(\cdot)$ of weather on the oxygen demand

4 Discussion

The increased availability of wearable devices empowers individuals to track their health. Our goal is to quantify this measure through modelling the heart rate response to workout. We learn representations that summarize the dynamics of the HR response, and serve as a measure of cardiorespiratory fitness. This can help track fitness levels over time and aid personalized workout planning; future work will investigate how such measures can predict changes in cardiovascular health.

References

[1] Ty Ferguson, Timothy Olds, Rachel Curtis, Henry Blake, Alyson J Crozier, Kylie Dankiw, Dorothea Dumuid, Daiki Kasai, Edward O'Connor, Rosa Virgara, et al. Effectiveness of wearable activity trackers to increase physical activity and improve health: a systematic review of systematic reviews and meta-analyses. *The Lancet Digital Health*, 4(8):e615–e626, 2022.

[2] Matilda Swee Sun Tang, Katherine Moore, Andrew McGavigan, Robyn A Clark, and Anand N Ganesan. Effectiveness of wearable trackers on physical activity in healthy adults: systematic review and meta-analysis of randomized controlled trials. *JMIR mHealth and uHealth*, 8(7): e15576, 2020.

[3] B Ballinger, J Hsieh, A Singh, N Sohoni, J Wang, GH Tison, GM Marcus, JM Sanchez, C Maguire, JE Olgin, et al. Deepheart: Semi-supervised sequence learning for cardiovascular risk prediction. thirty-second aaai conference on artificial intelligence. thirty-second aaai conference on artificial intelligence, 2018.

[4] Shamim Nemati, Mohammad M Ghassemi, Vaidehi Ambai, Nino Isakadze, Oleksiy Levantsevych, Amit Shah, and Gari D Clifford. Monitoring and detecting atrial fibrillation using wearable technology. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 3394–3397. IEEE, 2016.

[5] Takuji Suzuki, Ken-ichi Kameyama, and Toshiyo Tamura. Development of the irregular pulse detection method in daily life using wearable photoplethysmographic sensor. In *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 6080–6083. IEEE, 2009.

[6] Jennifer M Radin, Nathan E Wineinger, Eric J Topol, and Steven R Steinhubl. Harnessing wearable device data to improve state-level real-time surveillance of influenza-like illness in the usa: a population-based study. *The Lancet Digital Health*, 2(2):e85–e93, 2020.

[7] H Ceren Ates, Ali K Yetisen, Firat Güder, and Can Dincer. Wearable devices for the detection of covid-19. *Nature Electronics*, 4(1):13–14, 2021.

[8] Robert Avram, Geoffrey Tison, Peter Kuhar, Gregory Marcus, Mark Pletcher, Jeffrey E Olgin, and Kirstin Aschbacher. Predicting diabetes from photoplethysmography using deep learning. *Journal of the American College of Cardiology*, 73(9S2):16–16, 2019.

[9] Jessilyn Dunn, Lukasz Kidzinski, Ryan Runge, Daniel Witt, Jennifer L Hicks, Sophia Miryam Schüssler-Fiorenza Rose, Xiao Li, Amir Bahmani, Scott L Delp, Trevor Hastie, et al. Wearable sensors enable personalized predictions of clinical laboratory measurements. *Nature medicine*, 27(6):1105–1112, 2021.

[10] Melanie Ludwig, Katrin Hoffmann, Stefan Endler, Alexander Asteroth, and Josef Wiemeyer. Measurement, prediction, and control of individual heart rate responses to exercisebasics and options for wearable devices. *Frontiers in physiology*, 9:778, 2018.

[11] Jianmo Ni, Larry Muhlstein, and Julian McAuley. Modeling heart rate and activity data for personalized fitness recommendation. In *The World Wide Web Conference*, pages 1343–1353, 2019.

[12] Apple. Apple heart & movement study. `https://clinicaltrials.gov/ct2/show/NCT04198194`, 2019. ClinicalTrials.gov Identifier: NCT04198194.

[13] James Robert Stirling, Maria Zakynthinaki, Ignacio Refoyo, and Javier Sampedro. A model of heart rate kinetics in response to exercise. *Journal of Nonlinear Mathematical Physics*, 15 (sup3):426–436, 2008.

[14] Maria S Zakynthinaki. Modelling heart rate kinetics. *PloS one*, 10(4):e0118263, 2015.

[15] Michael J Mazzoleni, Claudio L Battaglini, Kerry J Martin, Erin M Coffman, and Brian P Mann. Modeling and predicting heart rate dynamics across a broad range of transient exercise intensities during cycling. *Sports Engineering*, 19(2):117–127, 2016.

[16] Michael J Mazzoleni, Claudio L Battaglini, Kerry J Martin, Erin M Coffman, Jordan A Ekaidat, William A Wood, and Brian P Mann. A dynamical systems approach for the submaximal prediction of maximum heart rate and maximal oxygen uptake. *Sports Engineering*, 21(1): 31–41, 2018.

[17] SD Galloway and Ronald J Maughan. Effects of ambient temperature on the capacity to perform prolonged cycle exercise in man. *Medicine and science in sports and exercise*, 29(9): 1240–1249, 1997.

[18] Ricky T. Q. Chen. torchdiffeq, 2018. URL `https://github.com/rtqichen/torchdiffeq`.

[19] LR Keytel, JH Goedecke, Timothy D Noakes, H Hiiloskorpi, Raija Laukkanen, Lize van der Merwe, and EV Lambert. Prediction of energy expenditure from heart rate monitoring during submaximal exercise. *Journal of sports sciences*, 23(3):289–297, 2005.