# A Study of Machine Learning Models for Personalized Heart Rate Forecasting in Mountain Biking

Xiaoxing Qiu[1], Jules White[1] and Douglas C. Schmidt[1]

[1]*Department of Computer Science, Vanderbilt University, Nashville, TN 37212 USA*
*{xiaoxing.qiu, jules.white, d.schmidt}@vanderbilt.edu*

Abstract:     Heart rate forecasting in cycling is most effective when it is personalized and course-specific to account for the influence of individual and terrain factors. This paper empirically assesses various personalized and course-specific heart rate forecasting models based on four machine learning models, including random forest, feed forward neural networks (FFNNs), recurrent neural networks (RNNs), and long short term memory (LSTM). The mean square error (MSE) is selected as the metric for model comparison. The results of our experiments show that despite the severely overfitted random forest models the LSTM models have the lowest MSE in the heart rate forecasting on our test dataset.

## 1 INTRODUCTION

### 1.1 Background

Monitoring exercise intensity during cycling is important. While insufficient exercise may slow improvement in an athlete's performance, excess exertion can result in over-training or even muscle damage. Over-training in cycling can result in a debilitating syndrome that degrades the performance of cyclists for several months and may ultimately result in failure to meet competition goals (Gleeson, 2002). Moreover, over-trained athletes—especially those involved in endurance sports like cycling—are susceptible to infections and require significantly longer time for recovery than non-athletes (Gleeson, 2002; Nieman, 1994).

Several metrics are available to estimate exercise intensity, including heart rate, power output, and $\dot{V}O_2max$. According to Jeukendrup et al's study (Jeukendrup and Diemen, 1998), heart rate is a reliable indicator of cycling exercise intensity since it is largely independent of a specific course. However, heart rate can be affected by terrain factors, such as grade of the hill, which varies on different riding courses.

Predicting heart rate at different points in time on a cycling course is hard since heart rate is affected by course-specific features, as well as other personalized physiological factors. Research has shown that a cyclist's heart rate drifts upwards after exercising for 20 to 60 minutes despite unchanged work loads (Le et al., 2009), which is a condition called "cardiac drift."

Cardiac drift is associated with an increase of core body temperature during exercise (Dawson et al., 2005; Collinson et al., 2001), which may cause athletes to lower their speed to maintain their heart rate. On the other hand, higher speed can yield a higher heart rate in a given condition (Le et al., 2009). Moreover, heart rate responses vary with a cyclist's individual factors, such as gender and age (Le et al., 2009), as well as their cadence on different parts of a given course.

It is hard to predict heart rate on a specific point in a course at a target speed. Training plans therefore focus on setting target heart rate or power goals (e.g., power applied to the pedals by the rider) and riding as fast as possible without exceeding those targets. Ideally, training plans could be developed where heart rate at a given speed and point in a course is predicted and riders are given a set of target speeds to ride instead. Achieving this ideal, however, requires building heart rate forecasting models that are personalized and course-specific, which yields the following benefits:

- **Achieve speed goals without over-training.** A personalized performance forecasting model for specific courses is vital for cyclists to establish achievable speed goals at different courses in advance and maximize training effects without over-training.

- **Improve course-specific estimation accuracy.** Riders can accurately estimate how much faster they could ride in different sections of the course, yet still remain within heart rate targets. For example, a rider who has only ridden a course at 70-80% of max heart rate could estimate how fast they would ride at a race intensity of up to 90% of max heart rate. This estimate could give them a benchmark of what they could currently achieve without actually riding the course and potentially over-exerting themselves.

A biker may need to climb up on a steep uphill trail on uneven terrain with 6-8in rocks, which requires significantly slower speeds and greater intensity than a similar climb on smooth ground. Similarly, a high-speed downhill ride on smooth terrain will place less stress on a rider's core muscle groups compared to the same downhill speed across tree roots. The individual terrain features, turns, gradients, and other aspects of a course significantly impact a rider's speed and intensity, but current models are mainly course-independent.

To maximize their efficacy, personalized and course-specific heart rate models should be easily trainable from limited data—ideally a single test ride of a course by an athlete. If an athlete rides a course multiple times at multiple target heart rates, they are less likely to need a predictive model since they already know the course well. Therefore, new courses (i.e., where the athlete has limited knowledge) are those where personalized and course-specific predictive models are most valuable. Ideally, a cyclist should be able to ride a course once and then predict how different target speeds would impact heart rate at different points on the course.

## 1.2 Research Question

This paper presents our research on comparing personalized multivariate models to forecast the heart rate of a cyclist on a specific course using data from a single ride. These models consider course-specific factors at each part of the course (such as the grade of road and the altitude), as well as current rider details (such as the cadence), and then forecast the heart rate of the athlete based on them. We compare the results of heart rate forecast by the following machine learning models from a single ride and report which models perform best:

- The first type of model experimented with is *random forest*, which is a traditional machine learning model. Forecasting results show that random forest models have severe overfitting and there-

fore cannot be utilized in heart rate forecasting from a single ride.
- To mitigate the random forest overfitting problem, we implement *feed forward neural networks* (FFNNs). FFNNs do not exhibit overfitting problems on the dataset, though their forecasting accuracy is lower. In particular, FFNNs do not consider historical information when forecasting heart rate.
- To address the limitations with FFNNs, we also implemented two other types of neural networks: *simple recurrent neural networks* (RNNs) and *long short term memory* (LSTM) networks.

We empirically compare all these models via experiments in our testbed. The results of these experiments indicate that the LSTM models have the lowest mean square error among machine learning models despite severely overfitted random forest models and produce predictions that closely match real-world heart rate sequences.

## 1.3 Paper Organization

The remainder of this paper is organized as follows: Section 2 describes prior work on heart rate forecasting; Section 3 reviews the data processing methods and all the models applied in this paper, including basic concepts of each model and their pros and cons; Section 4 describes how we built random forest, feed forward neural network (FFNN), recursive neural network (RNN), and long short term memory (LSTM) models to forecast the heart rate of athletes on a specific route and then compares the best models among these four types with two course-independent models (FitRec (Ni et al., 2019) and Minmin's LSTM-based model (Luo and Wu, 2020)); Section 5 presents concluding remarks and outlines future work.

## 2 RELATED WORK

Researchers have built various models to predict the performance of elite cyclists. Le et al. (Le et al., 2009) proposed a mathematical model to evaluate athletes' heart rate response under moderate exercise intensity based on physical and physiological principles. Lucia et al. (Lucía et al., 2001) analyzed the preferred cadence of elite cyclists and found that on flat stages they tend to adopt higher cadences (around 90 rpm) while on mountain ascents cadences are around 70 rpm. However, their models focus on laboratory conditions. Course-specific factors, such as the slope of the road in real courses, are not consider, though these

factors significantly influence a cyclist's heart rate response.

Due to the advent of wearable devices, large amounts of data can be collected and processed via mobile devices, which offers an opportunity to build personalized performance forecasting models. Xiao and Ming et al. (Ming and Jun, 2008; Xiao et al., 2010) used an FFNN to investigate the relationship between heart rate and physical activity in daily life with the help of a wearable physical activity recorder that monitors the 3-D accelerations of the body.

RNNs exhibit sequential correlation and can seamlessly model problems with multiple inputs. These models are therefore widely used in natural language processing and time series prediction (Cho et al., 2014). In athletic performance forecasting, RNN based models can take personalized factors (such as blood pressure and running speed during exercise) to make a heart rate predictions. Ni et al. (Ni et al., 2019) proposed an LSTM-based model to learn a user's heart rate profile during exercise and offer workout route recommendation and short term heart rate prediction. Luo et al. (Luo and Wu, 2020) also proposed an LSTM-based model to predict heat rate based on heart rate signal, gender, age, accelerations and mental state.

In summary, although there are other performance forecasting models, most studies focus on the heart rate profile collected either during daily activities or under laboratory conditions. There are few models that are course-specific and personalized to forecast a cyclist's heart rate. However, such models can be beneficial to both cyclists and coaches.

For example, a cyclist needs a model to establish various speed goals at different parts of a course before a competition or predict their heart rate based on given speed goals. A coach can use such a model to evaluate the heart rate of a given athlete on a given course at a given speed to ensure the exercise intensity and avoid over-training. Likewise, when given specific heart rate goals, these models can be used to predict how fast a cyclist can/should ride at different parts of a course.

## 3 METHODOLOGY

To address the limited understanding of personalized and course-specific heart rate forecasting from a single ride, we evaluated the performance of prior course/cyclist-independent models, course/cyclist-specific traditional machine learning models, and course/cyclist-specific neural networks on forecasting rider heart rate on a single ride of a course. These comparisons allowed us to investigate a number of important research questions and collect important lessons learned to guide future research, as discussed in this section.

### 3.1 Key Research Questions

The key research questions we investigated in our study include the following:

- Are current models that do not consider course-specific features as good as models that consider specific features, such as location?

- Which machine learning approaches perform best for personalized course-specific heart rate prediction from a single ride?

- For traditional machine learning models, which features are most salient for learning?

### 3.2 Experimentation Approach

We began our study by surveying prior work on heart rate forecasting. We then selected and applied both traditional machine learning models and neural networks to a cycling dataset that we collected. Important cyclist and course-specific factors must be considered for cycling performance forecasting. It is natural to consider multivariate models for performance forecasting, e.g., random forest and neural networks are popular machine learning algorithms because they work for both regression and classification and can handle multiple inputs.

According to Leijnen et al. (Leijnen and van Veen, 2020), there are 13 major neural network architectures used by researchers. Among all these neural networks, three of them are widely used in performance forecasting, including feed forward neural networks (FFNN) (Ming and Jun, 2008; Xiao et al., 2010), basic recurrent neural networks (RNN) (Chowdhury et al., 2019), and long short term memory (LSTM) (Bian et al., 2019; Ni et al., 2019) models. We therefore selected these three neural networks and random forest and built a personalized model to predict the performance of a cyclist on a specific routes. LSTM is a special type of RNN, so to distinguish these two models we call basic RNN models "simple RNN" models in this paper.

Pruning irrelevant features to an athlete's performance can reduce model overfitting and improve forecasting accuracy. Traditional machine learning models such asrandom forest models are sensitive to data variation, so small differences in the dataset can cause a large variance in the prediction. Neural networks have built-in mechanisms to mitigate overfitting and

perform feature selection by assigning significant features larger weights. This process, however, consumes a large amount of time and requires a large amount of data. Removing less important factors can accelerate the training process of traditional machine learning models and neural networks and reduce the amount of data needed. This paper therefore uses feature selection for all four types of models.

In statistics, a correlation coefficient is used to characterize how strong a relationship is between two variables. Two common correlation coefficients are widely used: the Pearson correlation coefficient and the Spearman's rank correlation coefficient. According to Bishara et al. (Bishara and Hittner, 2017), calculating Spearman's correlation coefficient for non-normal data may be an optimal strategy when the data size is larger than 20. Therefore, for heart rate forecasting, Spearman's rank correlation coefficient is more suitable for calculating the correlation coefficient, so we therefore select this approach to filter the personalized and terrain factors.

## 3.3 Overview of Our Dataset

For heart rate forecasting, the dataset used in this paper contains the grade of course, speed, heart rate, altitude, cadence, and distance at each second. We are interested in understanding how course-specific features impacted forecasting performance. We therefore use mountain biking data from trails in the Nashville, Tennessee, USA region.

Our dataset was collected on a Ripmo AF mountain bike instrumented with a Garmin 830 biking computer connected to an accelerometer-based speed sensor mounted to the front hub of the bike. The bike's crank arms included a Quarq Sram XX1 Eagle Dub power meter that used embedded strain gauges to measure the power applied to the pedals within +/- 2%. The Quarq power meter directly measured the rider's pedaling cadence from accelerometers embedded in the crank arms. The Garmin 830 included GPS positioning and improved location tracking using a fusion of wheel rotation, speed, and GPS fix data. Finally, a Wahoo Tickr elctrode-based chest strap was used to measure heart rate and communicate the data to the Garmin 830.

For our study, we selected 8.71 miles of riding on a 10-mile courses in the Nashville area. The rider was a 40-year old male weighing approximately 210 pounds.[1] The dataset and all the source code we

---

[1]Our current work focuses on a single rider to maximize understanding of how course-specific features could be learned since individual riders have been more thoroughly studied in prior work.

used to evaluate the machine learning and neural network models discussed in this paper is available from github.com/EricXQiu/SportDataProcessing.

## 3.4 Feature Selection

The first step in our dataset processing involved selecting features for model training. Before selecting these features their significance must be determined. Spearman's correlation coefficients ($\rho$) for each sequence are listed in Table 1. This table shows that the

Table 1: Spearman Correlation Coefficient Between Factors and Heart Rate

| Features | Spearman correlation coefficient |
|---|---|
| speed | 0.1826 |
| grade of course | 0.2524 |
| cadence | 0.2389 |
| distance | 0.0466 |
| altitude | 0.2586 |

grade of course, cadence, speed, and altitude significantly influence heart rate more than the other factors. We therefore selected these four factors as the features for our heart rate forecasting model.

We split the dataset between a training set and test set. A window average is carried out to metigate the mesurement error of heart rate. By convention, we used an 80% vs 20% train-test split ratio to split data into the training set and test set. For the LSTM-based model, the training set was the first 80% of the data and the test set was the remaining 20% rather than a random selection to account for the order dependence in heart rate data. To compare the results with LSTM-based models, the same train-test split process was also performed for the random forest model and neural networks.

## 4 HEART RATE FORECASTING MODEL COMPARISON

This section describes how we applied four machine learning models (random forest, feed forward neural networks, simple recurrent neural networks, and long short term memory) to build heart rate forecasting models that predict an cyclist's heart rate on a given course. The results of applying these heart rate forecasting models are then analyzed. We also compare the mean squared error (MSE) of all four models to glean insights into which models perform best and whether they have severe overfitting. Finally, we compare our models with other models (i.e., Ni's

model (Ni et al., 2019) and Luo's model (Luo and Wu, 2020)) that are *not* course-specific, but which we trained on our dataset described in Section 3.3.

## 4.1 Random Forest Heart Rate Model

We used the scikit-learn library (Pedregosa et al., 2011) to construct the heart rate forecasting models based on random forest. The criterion is MSE and the number of decision tree regressors can be modified. The number of decision trees regressors of four random forest models are 5, 10, 15, and 20.

Each decision tree regressor uses all factors in regression. Among them, model 1 with 5 decision tree regressors shows the best performance. Its MSEs are 0.455 $bpm^2$ and 125.15 $bpm^2$ on the training set and test set respectively. The maximum depth of tree depth is 31.

For all four models, the MSE on the test set is hundreds of times larger than on the training set, which shows that these random forest models have severe overfitting. The depth of decision tree regressors is approximately six times larger than the number of factors. This result indicates some factors are split multiple times, which increases the complexity of the decision tree, thereby yielding severe overfitting of the random forest model.

The heart rate forecasting result of model 1 is chosen and shown in Figure 1. The predicted value follows the trend of the heart rate and shows the model learned some patterns in the heart rate sequence due to the course. However, the predicted heart rate shows a large error around time 8,900s. Moreover, the predicted heart rate remains stable for the time range from 8,600s to 8,700s, whereas the real heart rate shows a sharp decreasing trend. The error in percentage between predicted heart rate and real heart rate shows a similar trend as the forecasting result.
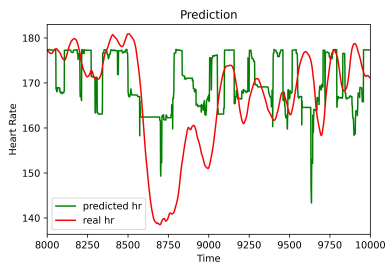


Figure 1: Heart Rate Forcasting of Random Forest Model

To investigate the structure of random forest model, the Gini importance[2] of each factor is calcu-

---

[2]The Gini importance is the average decrease of variance, which indicates the probability of whether there is a

lated. A factor with large Gini importance means that more nodes are split by this factor in the decision tree and therefore this feature is considered significant.

The average Gini importance of factors (in the order of time, grade, speed, cadence, altitude) of model 1 is 0.294, 0.0376, 0.531, 0.028, 0.108 respectively. It shows that speed, time, and altitude are three major factors used to split nodes in the first several layers of decision tree. Likewise, the table also shows the influence of cadence and grade are less important in decision tree construction.

We hypothesize that the source of the error may occur for the following reasons:

- The dataset only indirectly measures course features via speed and position, so it does not effectively learn course-specific influences on heart rate.

- There may be some internal rider conditions, such as the mental activities of the cyclist over time. For example, a cyclist may be anxious due to the terrain difficulty, thereby increasing their heart rate. All these features can influence heart rate, though they are hard to measure and express quantitatively.

- The random forest models may overfit on the training set and provide erroneous predictions.

## 4.2 The FFNN Heart Rate Model

Due to the overfitting problems of the random forest models, we built feed forward neural network (FFNN) models with different hyper-parameters. Two of them have two dense layers and the other two have three dense layers.

An important phase of training a neural network involves finding a proper learning rate. To search the learning rate systematically, we employed Smith's method (Smith, 2017) and utilized learning rate finder implemented by Pavel Surmenok et. al (Surmenok and Mackenzie, 2017). By running at each learning rate for 10 epochs, the optimal start learning rate for these four FFNN models is $3 \times 10^{-2}$.

Figure 2 shows the predicted heart rate and the real heart rate of the model applied to the same sample as the random forest.

The sample shows that the FFNN model can follow the trend of the cyclist's heart rate on the specific course. The FFNN model shows large errors from 8,500s to 9,000s. To investigate the error source, The MSEs of the training set and test set are calculated. The MSE on the training set and test set are 262 $bpm^2$ and 364 $bpm^2$.

---

split on this feature (Menze et al., 2009).
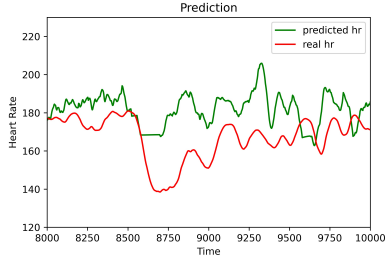
Figure 2: Heart Rate Forecasting of FFNN Model



Figure 3: Heart Rate Forecasting of Simple RNN Models

Compared with the random forest model, the variances of the FFNN model on the training set and test set are much closer, indicating less overfitting severity. We therefore expect the FFNN model to generalize better than the random forest model because they do not incur such severe overfitting.

## 4.3 The Simple RNN Model

Unlike FFNNs, recurrent neural networks (RNNs) can utilize heart rate data in the past to forecast the heart rate at the current moment. The architecture of a simple RNN model for heart rate forecasting consists two parts: 1. one or two Simple RNN layers 2. several dense layers for output.

Simple RNN Models with different hyper parameters were built and four groups of hyper-parameters are experimented. The models with one layer of Simple RNN layer perform dramatically better than models with two Simple RNN layers. The optimal start learning rates were identified via Smith's method and an exponential decay learning rate function was utilized for learning rate searching. With Keras learning rate finder (Surmenok and Mackenzie, 2017), the optimal start learning rate was selected as $2 \times 10^{-1}$.

The result of heart rate forecasting model with the highest accuracy is predicted in Figure 3. The MSEs on the training set and test set are 275.11 and 248.83 respectively. Compared with FFNN models, simple RNN Models exhibit fewer errors and the forecasting results are closer to the real heart rate sequence. The reason of limited improvement may be ascribed the gradient vanishing and gradient exploding problem in Simple RNN.

## 4.4 The LSTM Model

Simple RNN models can suffer from vanishing gradient problems, which limits their application when the sequence of input data items is very long. To address this issue, we implemented LSTM models. Four sets of hyper-parameters are experi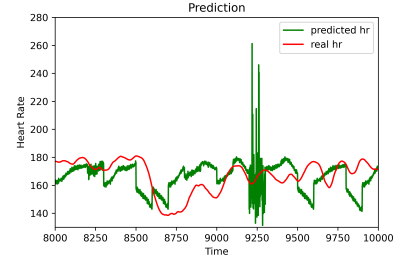mented and models with one LSTM layer model performs dramatically better than models with two LSTM layers. This may also indicate that the simple one-dimensional signals of speed, position, and cadence are insufficient to learn course-specific heart rate variation.

With the help of keras learning rate finder (Surmenok and Mackenzie, 2017), the optimal start learning rate for models with one LSTM layer (i.e., Models 1 and 3) is selected to be $1 \times 10^{-1}$.

The models with two LSTM layers show large MSE on the test set and they do not learned the general pattern of the heart rate compared with one LSTMN layer (Model 1 and Model 3). The results for these two one-layer LSTM models are shown in Figure 4.
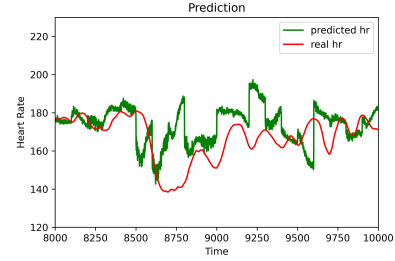


(a) Model 1
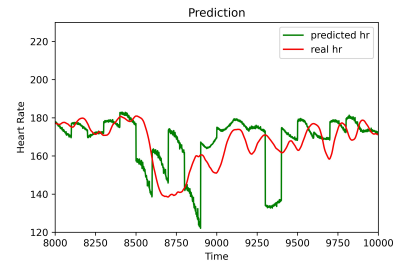


(b) Model 3

Figure 4: Heart Rate Forecasting of LSTM Models

The MSEs on the training set and test set of LSTM Model 1 is 141.17 $bpm^2$ and 200.50 $bpm^2$. The MSEs on the training set and test set of LSTM Model 3 is 62.47 $bpm^2$ and 196.61 $bpm^2$.

## 4.5 Comparison with Course Independent Models

Based on the results presented above, it appears that two LSTM models with only one LSTM layer perform better than the other three types of model. We therefore also compared these LSTM models with Luo's LSTM model (Luo and Wu, 2020) and Ni's LSTM-based model (Ni et al., 2019), which are course-independent and rely only on personalized factors and contextual factors. The input features to these latter two models were cadence, speed, altitude, and time.

Since Luo's and Ni's models are not course-specific, the grade of the biking course is excluded in the input factors. In particular, only the structure of their models are utilized and some layers (such as the encoding layers in Ni's model) are removed since heart rate forecasting is the main focus. The structure of Ni's model consists an LSTM layer, a dense layer followed by a dropout layer while Luo's model includes two LSTM layers and two dropout layers.

These two models were first trained on the same training set as our heart rate forecasting models with course-specific factors excluded. They were then tested on the same test set. The MSE of all four of these models is shown in Table 2.

Table 2: MSE of LSTM models

| Models | MSE on Training Set | MSE on Test Set |
|---|---|---|
| Ni's model | 1822.23 | 407 |
| Luo's model | 239.89 | 157.40 |
| LSTM model 1 | 141.17 | 200.50 |
| LSTM model 3 | 62.47 | 196.61 |

The comparison of forecasted heart rate is shown in Figure 5. While Ni et al.'s model showed less error
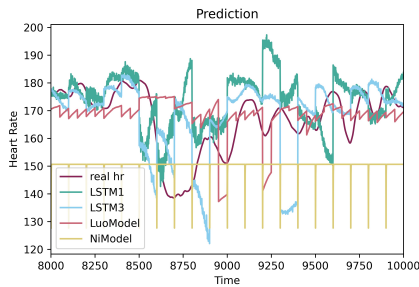


Figure 5: Comparison of Different Heart Rate Forecasting Models

on both the training set and test set, the heart rate it forecasted is close to constant, except for several drop

outs. Luo et al.'s model captures the heart rate pattern at around 170 beats per minute, but cannot follow the heart rate pattern overall.

## 5 CONCLUDING REMARKS

This paper presents an empirical analysis of personalized and course-specific models to forecast heart rates for cyclists. In particular, we explored the performance and feasibility of learning a cyclist's course-specific heart rate model from a single ride on a given course. We implemented models using long short term memory (LSTM), recursive neural network (RNN), random forest, feed forward neural network (FFNN), and previously published architectures and compared their performance learning a cyclist. We also compared our model with Ni et al's (Ni et al., 2019) and Luo et al.'s (Luo and Wu, 2020) models.

The heart rate forecasting results in Section 4 showed that the accuracy of two LSTM models outperformed the other two types of neural network models and did not have as severe overfitting as the random forest models. In Section 4.5, the LSTM course-specific models also performed better than the two LSTM-based course-independent models.

For research question 1 in Section 3.1 both course-independent models do not capture the heart rate trend of the cyclist well from 8,000s to 10,000s, as discussed in Section 4.5. In particular, the LSTM model proposed by Ni et al. (Ni et al., 2019) mainly outputs a rectangular heart rate sequence probably due to the lack of course-specific factors, such as grade. Ni et al. (Ni et al., 2019) state that their model focuses primarily on short-term predictions (typically in a window of 10 seconds), so a 2,000-second sequence may be too long for their model. Luo's model forecasted a somewhat continuous heart rate around 170 beats per minute with small dips around 9000s and 9250s. In general, our results show that these two models are not as accurate as our LSTM models since they do not consider course-specific factors.

For research question 2 in Section 3.1, random forest models exhibit lower MSE on the dataset. However, their substantial difference in MSE on the training set and test set indicates that they all incur severe overfitting. Moreover, as the data set grows larger, the number of erroneous readings will unavoidably increase, thereby increasing the probability of erroneous prediction due to the accumulated influence of error. As a result, conventional machine learning models have difficulty on course-specific heart rate forecasting.

For research question 3 in Section 3.1 speed has

the highest Gini importance, which aligns with prior work on heart rates for cyclists (Le et al., 2009) and is utilized as the splitting factor for the first layer. The three major factors are speed, time, and altitude. In contrast, the influence of cadence and grade are considered less important in decision tree construction.

Our future work focuses on scaling up our validation on a larger body of cyclists to determine whether these results hold true across a range of riders. We are also exploring how imagery of the course can aid in understanding complicated course features, such as terrain roughness. For personalized factors, we are evaluating the extent to which learned course-specific models transfer to other riders of the same gender and age, as well as bike types. We are also considering dynamic personalized factors, such as breathing rate.

For course-specific factors, image data and videos are being collected and analyzed via neural networks. We are analyzing roughness and course conditions at different parts of the course from these images and applying them in our heart rate forecasting model.

# REFERENCES

Bian, M., Peng, B., Wang, W., and Dong, J. (2019). An accurate LSTM based video heart rate estimation method. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 11859 LNCS, pages 409–417. Springer International Publishing.

Bishara, A. J. and Hittner, J. B. (2017). Confidence intervals for correlations when data are not normal. *Behavior Research Methods*, 49(1):294–309.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, pages 1724–1734.

Chowdhury, S. S., Hasan, M. S., and Sharmin, R. (2019). Robust Heart Rate Estimation from PPG Signals with Intense Motion Artifacts using Cascade of Adaptive Filter and Recurrent Neural Network. In *IEEE Region 10 Annual International Conference, Proceedings/TENCON*, volume 2019-October, pages 1952–1957.

Collinson, P. O., Boa, F. G., and Gaze, D. C. (2001). Measurement of Cardiac Troponins. *Annals of Clinical Biochemistry: International Journal of Laboratory Medicine*, 38(5):423–449.

Dawson, E. A., Shave, R., George, K., Whyte, G., Ball, D., Gaze, D., and Collinson, P. (2005). Cardiac drift during prolonged exercise with echocardiographic evidence of reduced diastolic function of the heart. *European Journal of Applied Physiology*, 94(3):305–309.

Gleeson, M. (2002). Biochemical and immunological markers of overtraining. *Journal of Sports Science and Medicine*, 1(2):31–41.

Jeukendrup, A. and Diemen, A. V. (1998). Journal of Sports Sciences Heart rate monitoring during training and competition in cyclists Heart rate monitoring during training and competition in cyclists. *Journal of Sports Sciences*, 16(1):91–99.

Le, A., Jaitner, T., Tobias, F., and Litz, L. (2009). A Dynamic Heart Rate Prediction Model for Training Optimization in Cycling (P83). *The Engineering of Sport 7*, pages 425–433.

Leijnen, S. and van Veen, F. (2020). The Neural Network Zoo. *Proceedings*, 47(1):9.

Lucía, A., Hoyos, J., and Chicharro, J. L. (2001). Preferred pedalling cadence in professional cycling. *Medicine and Science in Sports and Exercise*, 33(8):1361–1366.

Luo, M. and Wu, K. (2020). Heart rate prediction model based on neural network. *IOP Conference Series: Materials Science and Engineering*, 715(1):12060.

Menze, B. H., Kelm, B. M., Masuch, R., Himmelreich, U., Bachert, P., Petrich, W., and Hamprecht, F. A. (2009). A comparison of random forest and its Gini importance with standard chemometric methods for the feature selection and classification of spectral data. *BMC Bioinformatics*, 10(1):213.

Ming, Y. and Jun, J. (2008). Heart rate prediction based on physical activity using feedforward neural network. *Proceedings - 2008 International Conference on Convergence and Hybrid Information Technology, ICHIT 2008*, pages 344–350.

Ni, J., Muhlstein, L., and McAuley, J. (2019). Modeling heart rate and activity data for personalized fitness recommendation. *The Web Conference 2019 - Proceedings of the World Wide Web Conference, WWW 2019*, 2:1343–1353.

Nieman, D. C. (1994). Exercise, infection, and immunity. *International journal of sports medicine*, 15(3):S131.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., and Dubourg, V. (2011). Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 12:2825–2830.

Smith, L. N. (2017). Cyclical learning rates for training neural networks. In *Proceedings - 2017 IEEE Winter Conference on Applications of Computer Vision, WACV 2017*, pages 464–472.

Surmenok, P. and Mackenzie, J. (2017). keras_lr_finder. *GitHub repository*.

Xiao, F., Chen, Y. M., Yuchi, M., Ding, M. Y., and Jo, J. (2010). Heart rate prediction model based on physical activities using evolutionary neural network. *Proceedings - 4th International Conference on Genetic and Evolutionary Computing, ICGEC 2010*, pages 198–201.