



Using Case-Based Reasoning to Predict Marathon Performance and Recommend Tailored Training Plans

Ciara Feely¹(✉), Brian Caulfield², Aonghus Lawlor², and Barry Smyth²

¹ ML Labs, University College Dublin, Dublin, Ireland
Ciara.Feely@ucdconnect.ie

² Insight Centre for Data Analytics, University College Dublin, Dublin, Ireland
{brian.caulfield,aonghus.lawlor,barry.smyth}@ucd.ie

Abstract. Training for the marathon, especially a first marathon, is always a challenge. Many runners struggle to find the right balance between their workouts and their recovery, often leading to sub-optimal performance on race-day or even injury during training. We describe and evaluate a novel case-based reasoning system to help marathon runners as they train in two ways. First, it uses a case-base of training/workouts and race histories to predict future marathon times for a target runner, throughout their training program, helping runners to calibrate their progress and, ultimately, plan their race-day pacing. Second, the system recommends tailored training plans to runners, adapted for their current goal-time target, and based on the training plans of similar runners who have achieved this time. We evaluate the system using a dataset of more than 21,000 unique runners and 1.5 million training/workout sessions.

Keywords: CBR for health and exercise · Marathon running · Race-time prediction · Plan recommendation

1 Introduction

With the advent of wearable and mobile devices it has become increasingly routine for runners to track their training using apps such as Strava, RunKeeper, and MapMyRun. Researchers are harnessing this data to learn about how people exercise [1, 2], to provide personalised training advice [3–6] and motivational support [7–11], to predict their performance potential [12, 13], and even to provide them with real-time advice and guidance as they compete [14].

This work focuses on *recreational* (non-elite) marathon runners, although the ideas described should be equally applicable to other running distances (ultras, half-marathons, 10 km’s etc.) and endurance sports (cycling, triathlon, skiing, speed skating etc.). Its main technical contribution is to support marathon runners as they train, in two ways. Firstly, we predict a runner’s target race-time, based on their current training progress. This is important because it helps to

set appropriate race-day expectations for runners, helping them to better plan their race, but it also allows them to calibrate and fine-tune their training. Secondly, if runners wish to adjust their training – perhaps by targeting a faster or slower marathon time – then we describe a technique to generate a tailored training plan based on their current training habits and their new goals. In what follows, we describe and evaluate how both of these tasks can be fulfilled using case-based reasoning (CBR) by leveraging a case-base of more than 1.5 million training sessions logged by more than 21,000 marathoners. CBR is an appropriate method for these tasks as, for race-time prediction, the training completed by runners can be seen as the problem part, and race-time as the solution. Conversely, for training plan recommendation, the desired marathon finish-time is the problem, while the training plan is the solution.

2 Related Work

Fitness and exercise applications are popular targets for machine learning research, in part because of the volume of data that is now available, as people track their activities online, but also because of the wealth of interesting problems that exist when it comes to helping people to exercise safely and train effectively. The world of sports and fitness has been exploring the data captured by wearable sensors to solve a variety of tasks related to exercise, personalised training, motivation, and athlete performance [1–11]. Recently, case-based reasoning and other machine learning techniques have been utilised to support marathoners on race-day by providing them with real-time pacing advice [14].

A key task in this work is to predict future marathon times using training/workout data. This task is not new, but previous approaches have focused on either using a full complement of training/workout data or past race-times to generate predictions; see [12, 13, 15, 16]. Instead, we predict future race-times at various points during a training programme using incomplete training/workout data. Recently, the work of [17–20] used case-based reasoning ideas to accurately predict marathon performance but required runners to have completed at least one recent marathon. This means that these approaches are not suitable for first-time marathoners or novices. A key objective of the present work is to address this shortcoming, by using training/workout data, which even first-timers will generate at scale, instead of past marathon times.

Our second task involves recommending new training plans to runners. Such a virtual coaching assistant has long been discussed in the literature [6, 21, 22] but progress has been limited to some notable early efforts [23]. It is a challenging problem because generating a training plan depends on a complex mix of physiological and sport-specific factors as well as personal preferences. But this is precisely why a CBR approach is appealing: by reusing existing training plans (or parts of existing plans) from similar runners, we can provide a runner with tailored training recommendations without the need for an explicit domain model.

3 A CBR Approach to Marathon Training

Training for a marathon requires 12–16 weeks of dedicated effort, with most runners following carefully scripted training programmes based on their goals and ability. A typical week involves 3–6 training sessions, usually different types of runs: some short (5–10 km), some longer (15–30 km), some slow, some fast. Some runs introduce hills to build strength while others focus on stamina or recovery. As training progresses, new types of sessions encourage the physiological adaptations necessary for race-day. In other words, training for a marathon involves a complex mixture of workouts carefully balanced with rest and recovery.

By harnessing workout data, we provide runners with feedback as their training progresses. Predicting their likely marathon time will help runners to evaluate their progress, while the ability to make training recommendations will help them to adapt their otherwise *one-size-fits-all* training plan. In what follows, we will describe how we do this, but first we need to transform the time-series data from training sessions into a suitable representation for case-based reasoning.

3.1 From Training/Workout Sessions to Cases

The dataset used in this work includes approximately 1.5 million training activities by over 21 thousand marathon runners (73% male, 27% female) who completed either Dublin, London, or New York Marathons during the period 2014 – 2017; see Table 1. The anonymised dataset was produced by users of the popular mobile and web-based running app, Strava,¹ which has been made available as part of a data sharing agreement with the authors. The activities in the dataset all occur during a 16-week period directly before a marathon. This period was chosen as marathon plans are typically 12–16 weeks however, it is possible that some runners trained for less or more than 16 weeks. Each activity includes timing, distance, and elevation data sampled at 100 m intervals.

More formally, for a runner, r , we denote their training data as $T(r)$, a time-ordered sequence of training activities; see Eq. (1).

$$T(r) = \{A_1(r), A_2(r), \dots, A_n(r)\} \quad (1)$$

Each activity, $A_i(r) = (d, P)$, includes the number of days before the race (d) and a list of paces at 100 m intervals for the activity (P). A runner’s activities can be aggregated by week to extract key weekly features, including:

1. The *number of sessions* in the current week;
2. The *total weekly distance* in kms;
3. The *mean pace* for the week in mins/km;
4. The *longest run distance*;
5. The *fastest/slowest 10 km/5 km/1 km paces*.

¹ www.strava.com.

These features were chosen as they have been found to capture important aspects of marathon training in the past [16]. For example, the number and duration of *long-runs* is often cited as an important success criteria while, long-distance pacing typically correlates with marathon times.

In addition to these features that represent the current week of training, we also calculate the corresponding features for the training period up to and including the current week (e.g. *longest run distance* to date). Thus, for each runner r , we can generate a feature-based description for training week w , $F(r, w)$. Figure 1 demonstrates how the training of a runner in week 12 is transformed into a suitable feature representation.

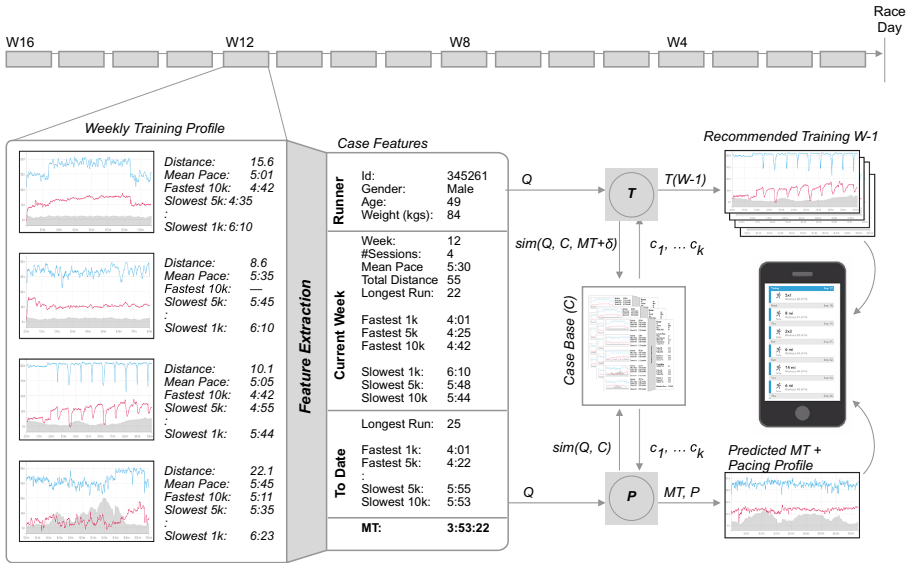


Fig. 1. An overview of a case-based reasoning system for supporting marathoners during their training by predicting (P) their estimated marathon time and by recommending (R) tailored training plan for an adjusted marathon time.

We generate a case ($C(r, w)$), representing r 's training during week w , by associating $F(r, w)$ with their marathon time, $MT(r)$, and also a pointer to their next week of training, $C(r, w - 1)$; see Eq. (2). These cases can be used in two ways: (a) to predict a runner's marathon time at week w , using the MT components of similar cases; and (b) to recommend next week's training, using the $C(r, w - 1)$ component of similar cases for a revised goal-time ($MT + \delta$).

$$C(r, w) = \{F(r, w), MT(r), C(r, w - 1)\} \quad (2)$$

When building a case-base of training activities we separate male and female runners because the physiological differences between men and women have a

significant bearing on training and performance. We also generate separate case-bases for each week of training, based on the feature-based description for a training week, $F(r, w)$, previously described. The marathon time $MT(r)$ for a case $C(r, w)$ encodes r 's marathon time in w weeks time and relates this to a specific week (week w) of training. It would not be appropriate to reuse such a case at a very different point in their training cycle, even for a similar runner.

3.2 Task 1: Predicting Goal Race-Times

The use-case for the first task is common: runner r in week w of training wishes to estimate their likely marathon time for race-day; the estimated time is not their *current* marathon time but rather their expected *future* marathon time, w weeks from now, based on their training to date. This is useful to know for a number of reasons. It helps r set appropriate race-day expectations and provides some level of confidence that their training is on-track, depending on whether the predicted time matches their goal. In addition, many marathon training programmes are parameterised with respect to a runner's goal marathon time – e.g., a long run session might include 5–10 km at *marathon pace* – so it is important to have an accurate estimate to work with.

To predict the marathon time of a runner r in week w , we use r 's current week of training as a query, and compute a standard Euclidean distance metric to identify the k most similar cases to r in the appropriate case-base (based on gender and training week). The predicted marathon time is the weighted average of the times for these similar runners; see P in Fig. 1. It is worth noting, but not discussed further here, that we can also recommend a suitable pacing plan to help the runner achieve this time on race-day, by reusing pacing profiles of the marathons completed by the k most similarly trained runners as in [17, 20].

3.3 Task 2: Recommending Tailored Training Programmes

To understand the use-case for the second task, imagine runner r has completed week 10 of their training plan and their predicted marathon time is 245 min. Given how well their training has gone so far, they decide that they want to break the iconic 4-h finish-time. Should they change their training plan to improve their chances of finishing faster? If so, how? What would a 4-h plan look like for them? Alternatively, if r 's training is proving to be too much of a challenge, they may wish to reduce their expectations and look for a training plan that suits a 4.5 h finish. What might this plan look like with 10 weeks of training still to go?

Instead of using r 's current training as a query to predict a marathon finish-time, we instead use their current training *and* their *revised* target time as a query to identify a new case, $C(r', w)$ from a runner r' who achieved the new target time (± 1 min), such that $C(r', w)$ is maximally similar to $C(r, w)$. Then, we can recommend $C(r', w-1)$ from the $C(r', w)$ case as r 's next week of training.

Note, for this task we focus on a single most similar case for r , rather than retrieving and reusing k similar cases. The main reason for this is that since runners can be following different types of training plans, it may not make sense

to try and combine these training plans from a recommendation perspective. That being said, it may make sense to offer r a choice of similar runners and therefore a choice of possible training for the following week.

3.4 From Single Weeks to Multiple Weeks

So far the focus has been on matching runner cases based on a single current week of training. Since many marathon programmes are designed around 4-week training blocks – during which training intensity ramps-up and then down to allow for recovery before the next block – it is also worth considering a longer, 4-week training period during prediction and recommendation. One way to do this is to extend our representations so that each case encodes the features of the previous 4 weeks of training.

Another option – and the one proposed here – is to use an ensemble approach to combine the predictions produced by similar cases for the 4 weeks including and preceding the current week. For example, for week $w = 10$, we generate 4 predictions using the case-bases for weeks 10, 11, 12, and 13, and the final prediction is produced from the median of these individual predictions.

One problem with the above approach is that runners who are on similar training plans can sometimes be out of sync with respect to their individual training weeks so that some weeks are “out of sequence”. To deal with this, we also implement a variation of this 4-week ensemble such that the case-bases used are produced by first *ordering* the 4 training weeks in ascending order of training-load (longest run distance for now). For example, for week $w = 10$, we use cases from weeks 10, 11, 12, 13, but we order them based on their longest run distance. So the $w - 3$ case-base contains the shortest training week for runners, the $w - 2$ case-base contains the next shortest training week etc. The advantage of this approach is that it facilitates a better alignment between the training weeks of runners over a 4-week period.

Obviously, the advantage of these ensemble approaches is that predictions are based on an extended view of training, rather than a single-week snapshot, which may lead to more accurate predictions. In what follows we will refer to the first ensemble approach as the *unordered* ensemble – to indicate that the weeks have not been ordered by training-load – and the second technique as the *ordered* ensemble.

Training plan recommendations can also take advantage of these extended approaches in a straightforward way, by using the ensemble methods to generate a (more accurate) race-time prediction and then using this single predicted time as the basis for the subsequent training plan recommendation as described previously.

4 Evaluation

We test the performance of our approach to race-time prediction and training plan recommendation using the Strava dataset referenced previously. In what

follows we describe this dataset in detail, and the evaluation methodology, before presenting key results for the prediction and recommendation tasks.

4.1 Setup

The details of the dataset used in this study are summarised in Table 1. It includes approximately 5,000 female runners who completed their marathon in 3–5 h and over 15,000 male runners who completed their marathons in up to 5 h; while the original dataset included some sub 3-h females and some slower (>5 h) males and females, these were relatively rare and excluded from this evaluation. Using this dataset we generate case-bases of weekly marathon training sessions for male and female runners, as previously described.

Each of the evaluations that follow adopt a similar, tenfold cross validation methodology, separating test and training data for the male and female case-bases for each week of training. During each iteration we extract 10% of the cases to use as test queries with case-bases constructed from the remaining cases.

For the prediction task we calculate the RMSE between the predicted marathon time and known marathon time for each test case. For the training-plan recommendation task we compare the recommended training plan to the corresponding plan for the test runner, to determine how its training load varies under different target time adjustments; we will discuss the details of this in due course.

In preparation for this evaluation we tested overall prediction accuracy for different values of k (the number of cases retrieved and reused) finding that accuracy improved (RMSE decreased) as k increased, before stabilising for $k \geq 15$. These results are not shown here for reasons of space but we use the $k = 15$ setting for the evaluations that follow.

4.2 Prediction Error by Training Week

One of the unique features of this work is the ability to generate marathon time predictions at any point in a runner’s training plan, not just at the completion of training. As such, it is important to understand how prediction accuracy changes as training progresses.

Figure 2 shows the results of this analysis for men and women and for each of the 3 CBR variants (single-week vs unordered 4-week vs ordered 4-week). As we might expect, prediction error falls steadily as training progresses, for men and women, and for each variant. A notable exception is one week before race-day for the single-week version, where RMSE increases slightly. This can be explained by the so-called *marathon taper* during which some runners significantly reduce their training load, so that they are rested for their race. Runners vary in when, how and even if they taper, so it is likely that the increase in error for the single-week representation exists because of a lack of taper consistency among the single-week cases, which is less problematic in the 4-week ensembles.

The 4-week variants produce more accurate predictions than the single-week approach, with the ordered variant consistently producing the most accurate

predictions overall, for each week and for men and women.² In each case, for men and women, the weekly differences in error between the ordered 4-week variant and both the single-week and un-ordered 4-week variants are all statistically significant (based on a one-sided t-test with $p < 0.01$). As a base-line, to further support the validity of the CBR approach, a linear regression model was fitted to each week of data. The results are omitted due to space constraints however,

Table 1. A summary of the dataset used in this study for runners of Dublin, London, and New York marathons in the period 2014–2017. The table includes gender and age information as well as mean (and standard deviation) data for age, race-time (minutes), number of weekly activities, and weekly distance).

			Runners	Age	Race-Time	Activites/Wk	Distance/Wk
City	Year	Sex					
Dublin	2014	F	52	37±7	251.69±23.66	2.97±1.19	31.35±10.6
		M	305	38±7	225.86±32.11	3.53±1.75	39.61±17.67
	2015	F	81	38±7	252.6±25.02	3.43±1.34	35.61±12.76
		M	496	38±8	222.32±32.5	3.62±1.83	39.64±18.33
	2016	F	180	39±8	248.7±26.7	3.49±1.53	35.91±14.83
		M	918	39±7	222.56±30.08	3.59±1.73	40.0±17.42
London	2015	F	535	38±8	239.45±30.19	3.52±1.57	38.79±14.55
		M	2091	39±8	210.2±34.76	4.0±2.21	45.21±22.03
	2016	F	881	38±8	239.2±31.4	3.6±1.56	38.76±15.19
		M	3088	40±8	211.31±35.89	4.13±2.37	45.99±23.07
	2017	F	1427	38±8	243.76±30.19	3.57±1.78	37.68±14.41
		M	4056	40±9	213.91±36.62	4.24±2.39	46.33±23.13
NYC	2015	F	324	37±8	244.21±27.5	3.61±1.42	38.84±13.99
		M	1374	40±8	225.21±35.37	3.63±1.66	42.07±19.95
	2016	F	693	37±8	245.31±28.68	3.64±1.59	38.11±14.96
		M	2180	40±8	225.34±34.14	3.64±1.72	41.55±18.86
	2017	F	1193	36±8	245.67±29.35	3.62±1.52	39.21±16.61
		M	3274	40±8	225.01±33.94	3.68±1.66	42.27±19.18

² Prediction estimates are more accurate for women than for men, echoing similar findings by [17] when using previous marathon times to predict future PBs.




Fig. 2. The prediction error (RMSE in minutes) by training week for (a) men and (b) women using the weekly and 4-week variants.

the linear regression model was statistical significantly ($p < 0.01$) less accurate than the single-week variant (and therefore both 4-week variants).

Indeed, predictions made 10 weeks before race-day, by the ordered variant, are as accurate as the predictions made by the single-week variant 5–6 weeks later. This is an important difference because, as mentioned earlier, having an accurate estimate of marathon time helps to inform subsequent training; workouts are often expressed relative to marathon pace. Thus, the availability of more accurate marathon predictions, earlier in training, has the potential to significantly optimise training.

4.3 Prediction Stability

While accuracy is important, it is not the only consideration when it comes to selecting a variant to use in practice. For example, if predictions tend to vary from week to week, then runners may be less likely to trust in them and therefore less likely to heed the advice and recommendations being made. To evaluate this, in Fig. 3 we calculate the absolute difference in the predicted marathon times between consecutive weeks for each runner and present the average difference for male and females and for each week of training and CBR variant.

Figure 3 shows that, in addition to enjoying better prediction accuracy, the 4-week variants also produce significantly more stable predictions, week on week. For example, 8 weeks from race-day, the single week variant generates an average prediction that differs from the previous week by approximately 9–10 min. By comparison, the 4-week variants produce predictions that differ from the previous week by only about 4 min; a useful side-effect of the ensemble prediction approach. In this case, the unordered variant produces more stable predictions for men and women than the ordered variant. The differences between the 4-week variants and the single-week variant are statistically significant based on a one-sided t-test with $p < 0.01$.




Fig. 3. The absolute difference in consecutive weekly predictions by training week for (a) men and (b) women using the weekly and 4-week variants.

4.4 Prediction Error by Ability

Fig. 4 plots the prediction error by runner ability – using their actual marathon times as a proxy for ability – for men and women at 10, 6, and 2 weeks before race-day. For reasons of space, we only show the results for the 4-week ordered variant, which proved to be the most accurate overall.




Fig. 4. The prediction error (RMSE in minutes) by marathon time (mins) for (a) men and (b) women using the weekly and 4-week variants.

Error rates increase significantly for slower runners (males >225 min and females >240 min) with the most accurate predictions associated with finish times of 210 min for male runners and 240 min for female runners. This is at least partly due to the distribution of marathon times in the training data: most of the training data is for runners in the 3–4 h finish-time range with relatively fewer faster and slower runners, leading to a paucity of training cases at the extremes, and less reliable predictions as a result.

There is a similar increase in error for faster (<210 min) females as there are relatively few of these in the dataset; the effect is less pronounced for faster males although still present. Generally speaking, we can also see how earlier predictions (week 10) tend to be less accurate regardless of gender or finish-time.

Another explanation for the significant increase in prediction error for the slower runners is that their training plans will tend to be less specific than those for faster runners and, as a result, may provide fewer or less reliable signals that can be used for prediction. For example, beginner training plans will tend to focus on helping a runner to *finish* the marathon distance, rather than achieve a particular time, and as such there will be less of a focus on pace, leading to less reliable ‘fastest pace’ features.

4.5 Evaluating Training Plan Recommendations

Evaluating training plan recommendations is less straight forward as there is no direct ground-truth to compare the recommendations to; after all, the aim is to suggest a training plan that is *different* (harder or easier) from the current plan for a given runner. Ideally, these recommended plans should be evaluated as part of a live-user trial – perhaps by obtaining user feedback on their desirability or suitability or by evaluating whether they lead to better outcomes, if and when users adopt them.

Such a study is beyond the scope of the present work. Instead, we propose a *plausibility* test by measuring how the *training load* of the recommended plans compares to the runner’s default plan: we compare their recommended next-week of training to their current next-week training plan. If a runner requests a plan for a marathon time that is faster ($\delta < 0$) than their current predicted marathon time, then the recommended plan should have a higher training load than their current plan, and vice versa if they request a plan for a slower ($\delta > 0$) marathon time. We use two measures of training load: (1) the average pace for the week; and (2) total weekly distance. Higher training loads should be associated with faster weeks or longer weeks or both. We calculate the percentage difference, with respect to the runner’s current plan, for distance and pace.

The results are shown in Figs. 5 and 6 for weeks 4, 6 and 8 of training using the ordered variant. We compare the recommendations produced when runners request plans that are associated with marathon times that are 5, 10, 15, and 20 min faster or slower than their current *predicted* marathon time. The results are generally consistent with expectations: when runners request training plans that are faster than their current predicted finish-time ($\delta < 0$) then mean weekly pace tends to speed-up (a negative % difference as in Fig. 5) while total weekly distance tends to increase (a positive % difference as in Fig. 6). The reverse is true when they request a plan for a slower marathon time.

The changes in pace exhibit a very strong correlation with δ ($R^2 > 0.92$ for men and women). The changes in weekly distance are also strongly correlated with δ for men ($R^2 > 0.90$), but less so for women ($R^2 > 0.66$ on average). The relative changes in distance tend to be greater (for a given δ) than the corresponding changes in pace. For example, for males to improve their predicted

time by 15 min, means they will have to increase their weekly distance by up to 5% and speed-up by 2–3%.

While not definitive, these results are encouraging. Recommending new training plans is a very challenging recommendation task; conventional recommendation techniques have largely focused on recommending simple, atomic items (books, music, movies) rather than complex items, such as training plans, which are made up of a complex mix of components and factors. The fact that we can generate training plan recommendations that are consistent with a runner’s modified goals is an encouraging start. And since these plans are based on the real training plans of similar runners, this increases the chances that they will be well received by runners.




Fig. 5. The difference in mean weekly pace (mins/km) for training plans based on adjusted goal-times for (a) men and (b) women during training weeks 4, 6, and 8. Note: $\delta < 0$ implies the goal-time is δ minutes *faster* than the runner’s current predicted time.




Fig. 6. The difference in mean weekly distance (km) for training plans based on adjusted goal-times for (a) men and (b) women during training weeks 4, 6, and 8. Note: $\delta < 0$ implies the goal-time is δ minutes *faster* than the runner’s current predicted time.

5 Conclusions

In this paper, we described an initial study about how raw training/workout data that is routinely collected by fitness apps can be used to support runners as they train for the marathon. We focused on two important tasks in particular: (a) race-time predictions, as training progresses; and (b) recommending tailored training plans to runners if their goals change during training. A number of CBR variations were described – reusing the training and racing experiences of similar runners – and evaluated. The results are promising. It was possible to predict marathon finish-times with a reasonable degree of accuracy and to recommend training plans that are consistent with a runner’s changing goals. Unlike the work of [17–20], which required runners to have run multiple marathons, this approach is suitable for novice and veteran runners alike, because it is based on current training data, with no requirement for previous marathon experience.

There are many opportunities to extend this research and improve the results obtained. We are currently developing a Strava companion app for providing predictions and recommendations to users based on their logged training sessions, making it possible to evaluate how users respond to this advice, and whether their performances improve as a result. Further representation improvements are also feasible, for example, by including heartrate data as a signal for effort and intensity during training, or by using time-series analysis techniques [24–26] to detect different types of training sessions. Another option is to employ feature analysis and selection techniques to determine which features are best predictors for race-time, as well as investigating different multi-week representations. Additionally, it is planned to transform the predictions and recommendations into a format that the runners could more easily interpret by providing upper and lower bounds, alongside average values for the race-time and weekly training completed by the similar runners retrieved from the case-base.

Finally, although the focus of this work has been exclusively on marathon runners, it is straightforward to adapt these techniques for other running distances, from shorter 5k, 10k and half-marathon races to longer ultra marathons, and it should also be possible to apply the work to other endurance sports such as cycling, triathlons, adventure racing, skiing and even skating.

Acknowledgments. Supported by Science Foundation Ireland through the Insight Centre for Data Analytics (12/RC/2289.P2) and the SFI Centre for Research Training in Machine Learning (18/CRT/6183).

References

1. Gasparetti, F., Aiello, L.M., Quercia, D.: Evaluating the efficacy of traditional fitness tracker recommendations. In: Proceedings of the 24th International Conference on Intelligent User Interfaces: Companion, IUI 2019, pp. 15–16, ACM, New York (2019)

2. Schneider, O.S, MacLean, K.E., Altun, K., Karuei, I., Wu, M.M.: Real-time gait classification for persuasive smartphone apps: Structuring the literature and pushing the limits. In: Proceedings of the 2013 International Conference on Intelligent User Interfaces, IUI 2013, pp. 161–172. ACM, New York (2013)
3. Cau, F.M., Mancosu, M.S, Mulas, F., Pilloni, P., Spano, L.D.: An intelligent interface for supporting coaches in providing running feedback. In: Proceedings of the 13th Biannual Conference of the Italian SIGCHI Chapter: Designing the next interaction, CHIItaly 2019, Aдова, Italy, September 23–25, pp. 6:1–6:5 (2019)
4. Mulas, F., Pilloni, P., Manca, M., Boratto, L., Carta, S.: Using new communication technologies and social media interaction to improve the motivation of users to exercise. In: Second International Conference on Future Generation Communication Technologies (FGCT 2013), London, United Kingdom, November 12–14, pp. 87–92 (2013)
5. Boratto, L., Carta, S., Iguider, W., Mulas, F., Pilloni, P.: Predicting workout quality to help coaches support sportspeople. In: Proceedings of the 3rd International Workshop on Health Recommender Systems, HealthRecSys 2018, co-located with the 12th ACM Conference on Recommender Systems (ACM RecSys 2018), Vancouver, BC, Canada, October 6, pp. 8–12 (2018)
6. Monteiro-Guerra, F.M. , Rivera-Romero, O., Luque, L.F., Caulfield, B.: Personalization in real-time physical activity coaching using mobile applications: A scoping review. *IEEE J. Biomed. Health Inform.* (2019)
7. Boratto, L., Carta, S., Fenu, G., Manca, M., Mulas, F., Pilloni, P.: The role of social interaction on users motivation to exercise: A persuasive web framework to enhance the self-management of a healthy lifestyle. *Pervasive Mob. Comput.* **36**, 98–114 (2017)
8. Pilloni, P., Piras, L., Carta, S., Fenu, G., Mulas, F., Boratto, L.: Recommender system lets coaches identify and help athletes who begin losing motivation. *IEEE Comput.* **51**(3), 36–42 (2018)
9. Buttussi, F., Chittaro, L., Nadalutti, D.: Bringing mobile guides and fitness activities together: A solution based on an embodied virtual trainer. In: Proceedings of the 8th Conference on Human-Computer Interaction with Mobile Devices and Services, MobileHCI 2006, pp. 29–36, ACM, New York (2006)
10. Hosseinpour, M., Terlutter, R.: Your personal motivator is with you: A systematic review of mobile phone applications aiming at increasing physical activity. *Sports Med.* **49**, 1425–1447 (2019)
11. Mulas, F., Carta, S., Pilloni, P., Manca, M.: Everywhere run: A virtual personal trainer for supporting people in their running activity. In: Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology, ACE 2011, Lisbon, Portugal, November 8–11, p. 70 (2011)
12. Bartolucci, F., Murphy, T.B.: A finite mixture latent trajectory model for modeling ultrarunners' behavior in a 24-hour race. *J. Quant. Anal. Sports* **11**(4), 193–203 (2015)
13. Keogh, A., Smyth, B., Caulfield, B., Lawlor, A., Berndsen, J., Doherty, C.: Prediction equations for marathon performance: A systematic review. *Int. J. Sports Physiol. Perform.* **14**(9), 1159–1169 (2019)
14. Berndsen, J., Smyth, B., Lawlor, A.: Pace my race: Recommendations for marathon running. In: Proceedings of the 13th ACM Conference on Recommender Systems, pp. 246–250. ACM (2019)

15. Claudino, J.G., Capanema, D.d.O., de Souza, T.V., Serrão, J.C., Machado Pereira, A.C., Nassis, G.P.: Current approaches to the use of artificial intelligence for injury risk assessment and performance prediction in team sports: A systematic review. *Sports Med. Open* **5**, 28 (2019)
16. Doherty, C., Keogh, A., Davenport, J., Lawlor, A., Smyth, B., Caulfield, B.: An evaluation of the training determinants of marathon performance: A meta-analysis with meta-regression. *J. Sci. Med. Sport* (2019)
17. Smyth, B., Cunningham, P.: Running with cases: A CBR approach to running your best marathon. In: Aha, D.W., Lieber, J. (eds.) ICCBR 2017. LNCS (LNAI), vol. 10339, pp. 360–374. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-61030-6_25
18. Smyth, B., Cunningham, P.: A novel recommender system for helping marathoners to achieve a new personal-best. In: Proceedings of the Eleventh ACM Conference on Recommender Systems, RecSys 2017, Como, Italy, August 27–31, pp. 116–120 (2017)
19. Smyth, B., Cunningham, P.: Marathon race planning: A case-based reasoning approach. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13–19, Stockholm, Sweden, pp. 5364–5368 (2018)
20. Smyth, B., Cunningham, P.: An analysis of case representations for marathon race prediction and planning. In: Cox, M.T., Funk, P., Begum, S. (eds.) ICCBR 2018. LNCS (LNAI), vol. 11156, pp. 369–384. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01081-2_25
21. Zahran, L., El-Beltagy, M., Saleh, M.: A conceptual framework for the generation of adaptive training plans in sports coaching. In: Hassanien, A.E., Shaalan, K., Tolba, M.F. (eds.) AISI 2019. AISC, vol. 1058, pp. 673–684. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-31129-2_62
22. Schneider, H.: Adapting at run-time: Exploring the design space of personalized fitness coaches. In: Proceedings of the 22Nd International Conference on Intelligent User Interfaces Companion, IUI 2017 Companion, pp. 173–176, ACM, New York (2017)
23. Fister Jr., I., Fister, I.: Generating the training plans based on existing sports activities using swarm intelligence. In: Patnaik, S., Yang, X.-S., Nakamatsu, K. (eds.) Nature-Inspired Computing and Optimization. MOST, vol. 10, pp. 79–94. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-50920-4_4
24. Senin, P., Lin, J., Wang, X., Oates, T., Gandhi, S., Boedihardjo, A.P., Chen, C., Frankenstein, S.: GrammarViz 3.0: Interactive discovery of variable-length time series patterns. *ACM Trans. Knowl. Discov. Data* **12**(1), 10:1–10:28 (2018)
25. Berlin, E., Laerhoven, K.V.: Detecting leisure activities with dense motif discovery. In: The 2012 ACM Conference on Ubiquitous Computing, Ubicomp 2012, Pittsburgh, PA, USA, September 5–8, pp. 250–259 (2012)
26. Cheng, H.-T.: Learning and recognizing the hierarchical and sequential structure of human activities. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA (2013)