Research Repository UCD

Title	Running with Cases: A CBR Approach to Running Your Best Marathon
Authors(s)	Smyth, Barry, Cunningham, Pádraig
Publication date	2017-06-21
Publication information	Smyth, Barry, and Pádraig Cunningham. "Running with Cases: A CBR Approach to Running Your Best Marathon." Springer, 2017.
Series	Lecture Notes in Computer Science
Publisher	Springer
Item record/more information	http://hdl.handle.net/10197/9164
Publisher's statement	The final publication is available at www.springerlink.com .
Publisher's version (DOI)	10.1007/978-3-319-61030-6_25

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Running with Cases: A CBR Approach to Running Your Best Marathon

Barry Smyth & Pádraig Cunningham

Insight Centre for Data Analytics
School of Computer Science
University College Dublin, Ireland

Abstract. Every year millions of people around the world train for, and compete in, the marathon. As race-day approaches, and training schedules begin to wind down, many participants will turn their attention to their race strategy, as they strive to achieve their best time. To help with this, in this paper we describe a novel application of case-based reasoning to address the dual task of: (1) predicting a challenging, but achievable, personal best race-time for a marathon runner; and (2) recommending a race-plan to achieve this time. We describe how suitable cases can be generated from pairs of race histories and how we can predict a personal best race-time and produce a tailored race-plan by reusing the race histories of similar runners. This work is evaluated using data from the last six years of the London Marathon.

Keywords: case-based reasoning, recommender systems, marathon running

1 Introduction

Running a marathon is hard. It depends on months of dedicated training and the discipline to follow a time-consuming and exhausting schedule, but a successful marathon also needs an appropriate race-plan for the event itself. The runner must select a suitable target finish-time and carefully plan the pacing of their race, segment by segment, kilometer by kilometer. The aim of this work is to help marathon runners to achieve a new personal-best during their next race. We do this by using case-based reasoning to solve these dual problems by: (1) *predicting* a best achievable finish-time; and (2) *recommending* a tailored pacing plan to help the runner achieve this time on race-day.

This work sits at the intersection between personal sensing, big data, and machine learning. An explosion of wearable sensors and mobile devices has created a tsunami of personal data [1], and the promise that it can be harnessed to help people to make better decisions and live healthier and more productive lives [2, 3]. Indeed, within the case-based reasoning community there has been a long history of applying case-based, data-driven methods to a wide range of healthcare problems [4]. Recently, the world of sports and fitness has similarly

embraced this big data vision as teams and athletes embrace the power of data to optimise the business of sports and the health and training of athletes; see for example [5, 6].

In this work we focus on recreational marathon running, helping runners to plan their race strategy for optimal performance. A key concept is that of *pace*, measured in minutes per mile/km, so that higher pace means slower speed. There is a growing body of research that explores the various factors that influence pacing during the marathon and other endurance events. For example, the work of [7] looks at the effect of age, gender, and ability on marathon pacing and concluded that female runners are typically more disciplined (running with less pace variation) than male runners; see also [8, 9]. Related to pacing is the runner’s *pacing strategy*: how the runner plans to adjust their pace during the race; see [10]. There are 3 basic pacing strategies, based on how the pace of a runner varies between the first and second-half of a race. For example, we say that a runner completes an *even-split* if their pace is even throughout the race. Running a *positive-split* means the second-half of the race is *slower* (higher pace) than the first-half of the race, whereas a *negative-split* means the runner speeds-up in the second-half of the race, running it *faster* than the first-half. Many elites and disciplined runners will aim for even or slightly negative-splits; see [11]. Recreational runners typically run positive-splits, slowing during the second-half of the race, sometimes significantly. And for some this means *hitting the wall*, referring to the sudden onset of debilitating fatigue and a near-complete loss of energy that can occur late in a race; see [12]. At best, this temporarily slows even the swiftest of runners; at worst it reduces them to a shuffling gait for the remainder of the race.

All this is to say that running the marathon is a challenge, and the difference between a good day and a terrible day, training aside, may well come down to how carefully a runner plans their race: what finish-time to aim for; positive vs. negative vs. even splits; avoiding hitting the wall etc. This is where we believe there is a significant opportunity to support marathon runners, by advising them on a suitable a target finish-time and providing them with a concrete race-plan, one that is personalized to their ability and tailored to the marathon course.

2 Problem Definition

The problem addressed by this work is how to help a marathon runner achieve a new personal best in their next race. We will assume we are dealing with a runner who has completed at least one previous marathon and so has a race record to serve as a starting point. As a reminder, this problem involves two related tasks: the *prediction* of a suitable target finish-time and the *recommendation* of an appropriate pacing plan.

2.1 Best Achievable Finish-Time

We start by assuming our runner wants to beat their current best-time, but by how much? If they are too conservative they will chose a finish-time that does not

fully test them, and may leave them disappointed if they finish too comfortably on race-day. If they are too ambitious they may select finish-time that is beyond their ability and risk sabotaging their race; aiming for an overly ambitious target time is one sure way to end up hitting the wall later in the race. The point is that selecting a *best achievable* finish-time is non-trivial and getting it wrong can have a disastrous effect on race-day.

2.2 Race Plans & Pacing Profiles

Given a best achievable finish-time, the next task is to devise a suitable race-plan. We will assume the marathon is divided up into 8x5km segments (0-5km, 5km-10km, ..., 35km-40km), plus a final 2.2km segment (40km-42.2km); most big-city marathons measure split-times in 5km segments and so race records typically provide these *split-times*. For the purpose of this work a race-plan, or *pacing plan*, will consist of a sequence of average paces (measured in kms per minute) for each of these race segments. For the avoidance of doubt we will refer to a pacing *plan* as a *proposed* set of paces for the runner and a pacing *profile* as the *actual* set of paces achieved by the runner in their race. For example, Figure 1(a) illustrates a race-record for a runner who completed a marathon in 4 hours and 13 minutes; that's an average pace of 6 mins/km. The corresponding pacing profile shows they ran a positive-split, starting their race 10% faster than their average pace during the first 5k and 10k segments of the race, and slowing in the second-half, to finish 7% slower than their average pace in the final segment.

Most conventional pacing plans are fairly simple and assume the runner will run an even-split. For example, if our runner wishes to run a 4-hour (240 minute) marathon then this suggests an average pace of 5 minutes and 41 seconds per km for the duration of the race. Some plans may account for positive or negative splits, and thus advise the runner to run the first half of the race more quickly or slowly, respectively. However, conventional plans are not especially tailored for the runner or the course, and so they leave considerable room for improvement, especially when it comes to achieving a new personal-best. Thus, in this work we argue for a more *personalized* and *customized* approach to pacing, so that runners can benefit from a pacing plan that is suitable for their goal time, their personal fitness level and ability, and that is tailored for the peculiarities of a given marathon course.

3 Using CBR to Achieve a Personal-Best in the Marathon

The key insight of this work is that we can use a CBR approach to both predict suitable target finish-times and recommend tailored pacing plans, based on the runner's own race experience and the experiences of similar runners who have achieved personal-bests on the target course. To do this we will rely on a case base of *race pairs*, representing a pair of race records for a single runner. Each race record contains a pacing profile and a finish-time for a completed race; see

Figure 1(b). One of these race records corresponds to a *non personal-best* (*nPB*) race, the other to a *personal-best* (*PB*) race; the *nPB* plays the part of the *case description* while the *PB* is the *case solution*. Given a target/query runner (q), and their own recent race record (finish-time and pacing profile), we generate a finish-time prediction and pacing plan recommendation based on the *PB*'s of cases that have a similar *nPB* to q as summarised in Figure 1(c).

3.1 Case Generation

As mentioned above, each case in the case base corresponds to a single runner, r , with a *nPB* part, which represents a non personal-best race record, and a *PB* part as the personal-best race record; see Equation 1. To be represented in the case base r must obviously have at least two race records and in general may have $n > 2$ race records if they have run many races; for example, in Figure 1(b) we highlight 3 race records for r (for marathons m_1, m_2, m_3).

$$c_{ij}(r, m_i, m_j) = \langle nPB_i(r, m_i), PB(r, m_j) \rangle \quad (1)$$

The race record with the fastest finish-time is considered as the personal-best, and it is paired with the remaining $n - 1$ non personal-best records, producing $n - 1$ cases. As per Figure 1(b), r 's best race is m_2 , with a finish-time of 236 minutes. This is then paired with the 2 *nPB* records (m_1 and m_3) to produce 2 cases, $c(r, m_1, m_2)$ and $c(r, m_3, m_2)$ as shown.

3.2 Case Retrieval

Retrieval is a three-step process, as shown in Algorithm 1. Given a query race record (q) –that is a runner, a finish-time, and a *nPB* pacing profile – we first filter the available cases (CB) based on their finish-times, so that we only consider cases for retrieval if their finish-times are within t minutes of the query finish-time. This ensures that we are basing our reasoning on a set of cases that are somewhat comparable in terms of performance and ability.

Next, we filter on the basis of gender, only considering cases for retrieval if they have the same gender as the query runner. The reason for this is that physiological differences between men and women have a material impact on marathon performance. On average women tend to finish about 12% slower than men [7], and thus men and women with the same approximate finish-times will be racing at very different levels of maximum effort or ability; in relative terms, a 3-hour female finisher will be out-performing a 3-hour male finisher, for example. Hence we separate male and female cases for more effective matching, during retrieval.

Finally, we perform a standard, distance-weighted *kNN* retrieval over the remaining candidate cases C , comparing q 's pacing profile to their *nPB* profiles. These pacing profiles are real-valued vectors and, for now, we use a simple Euclidean-based similarity metric for similarity assessment. We select the top k most similar as the retrieved cases, R .

Fig. 1. Races, cases, predictions, and recommendations. (a) An example race record for a runner, showing a finish-time and a pacing profile containing pacing data for each of the 5km race segments; (b) Converting race records into cases; (c) An overview of the CBR process: given an nPB race record as a query, the system retrieves a set of k cases with *similar* nPB parts, and combines these to generate a personal-best finish-time prediction and a pacing plan to achieve this finish-time.

Algorithm 1: Outline CBR Algorithm.

Data: Given: q , query race record; CB , case base; k , number of cases to be retrieved; t , finish-time threshold.

Result: pb , predicted finish-time; pn , recommended pacing profile.

```

begin
   $C = \{c \in CB : Time(q) - t < Time(c) < Time(q) + t\}$ 
   $C = \{c \in C : c.gender == q.gender\}$ 
  if  $len(C) \geq k$  then
     $R = sort_k(sim(q, c) \forall c \in C)$ 
     $pb = predict(q, R)$ 
     $pn = recommend(q, R)$ 
    return  $pb, pn$ 
  else
    return  $None$ 
  end
end

```

3.3 Personal-Best Finish-Time Prediction

Given a set of similar cases, R , we need to estimate the best achievable finish-time for q . Each case in R represents another runner with a similar nPB to q who has achieved a faster personal-best on the same marathon course. The intuition is that since these PBs were achievable by these similar runners, then a similar PB should be achievable by the query runner.

For the purpose of this work we test three straightforward prediction approaches in which the predicted PB time of the query runner is a simple function of the personal-best times for the retrieved cases. With each approach the predicted PB finish-times are weighted based on the relative difference between the query runner's finish-time and the corresponding nPB finish-time of a retrieved case; see Equation 2. This adapts the PB times based on whether they were achieved by similar runners with slightly slower or faster nPB times.

$$w(q, c) = \frac{q(nPB).finish}{c(nPB).finish} \quad (2)$$

Best PB . In Equation 3 the predicted finish-time is the weighted PB finish-time of the single fastest retrieved case C_{best} ; this is the case with the minimum PB finish-time. For example, if the query runner has a finish-time of 245 minutes, and the nPB finish-time for the best retrieved case is a slightly faster 240 minutes, then the weighting factor will be 1.02 (245/240). If the PB finish-time of this retrieved case is 232 minutes then the predicted PB time for q will be just over 236 minutes.

$$PB_{best}(q, C) = w(q, C_{best}) \bullet Time(C_{best}(PB)) \quad (3)$$

Mean PB. Our second prediction approach calculates the *weighted mean* of the *PB* finish-times of retrieved cases as in Equation 4. This will obviously tend to provide a slower prediction finish-time than the *Best PB* approach above but it will be more representative of the personal-bests achieved by the similar runners.

$$PB_{mean}(q, C) = \frac{\sum_{\forall i \in 1..k} w(q, C_i) \bullet Time(C_i(PB))}{k} \quad (4)$$

Even PB. Finally, our third approach (*Even PB*) is based on the idea that more consistent or even pacing is better than more varied or erratic pacing. By measuring the coefficient of variation (*CoV*) of the relative paces in the *PB* pacing profile for each retrieved case, we can estimate their pacing variation, and select the one (C_{even}) with the lowest *CoV* value. The *PB* time of this C_{even} case is used as the predicted time for q .

$$PB_{even}(q, C) = w(q, C_{even}) \bullet Time(C_{even}(PB)) \quad (5)$$

3.4 Pacing Recommendation

Given a predicted personal-best finish-time it is straightforward to turn this into an average pace (mins/km) for q to run her race. However, as discussed previously, such an approach is unlikely to be successful, due to a variety of factors including ability, discipline, and course conditions. Therefore, in this section we describe how to use CBR to fulfill the requirement for a more personalized and tailored pacing plan with which to achieve this new personal-best finish-time. In fact, a key advantage of CBR in this regard, compared to a more global, eager learning approaches, is precisely that it retains access to the local cases and their pacing profiles. It means that we can use the *PB* profiles of retrieved cases as the basis for a pacing plan for q and, in what follows, we describe 3 different approaches as *companions* to our 3 prediction approaches.

Best Profile. The first recommendation approach (*Best Profile*) reuses the *PB* pacing profile of the case with the best *PB* time, C_{best} . In other words, we take the relative pacing from C_{best} and map its relative paces to the average pace for predicted *PB* time for the query runner. For example, if the predicted *PB* time is for 232 minutes, indicating an average pace of 5 minutes 30 second per km, and the *PB* profile in C_{best} calls for a first 5km that is 5% faster than average pace, then the generated pacing profile for q will advise running the first 5km at just over 5 minutes and 13 seconds per km.

Mean Profile. The second approach (*Mean Profile*) generates a new pacing plan based on the *mean* relative segment paces of the *PB* profiles from the k retrieved cases. So, if, on average, the *PB* profiles of the retrieved cases have their runners starting 6% faster than their average pace then our *Mean Profile* plan for a 232-minute finish will suggest that q begins her race at 5 minutes and 10 seconds per km.

Even Profile. Finally, we can generate our pacing plan based on the *PB* profile of C_{even} , the case with the most even pacing profile. If its profile calls for the first 5km to be run at a pace that is just 2% faster than average race pace then the pacing profile for q will suggest a pace of 5 minutes and 23 seconds per km.

3.5 A Worked Example

Up to now we have described an approach to generating cases from pairs of race records. We have shown how these cases can be reused (according to three different strategies, *Best*, *Mean*, *Even*) as the basis for finish-time predictions and pacing recommendations, to help q achieve a new personal best in their marathon. By design we have kept our case representation, retrieval, and reuse strategies straightforward, to make it easy to embed this approach in a practical application for runners.

A summary example of such an application is depicted in Figure 1(c) in which a (query) runner submits a (*nPB*) race record for a 248-minute finish-time with a pace profile that indicates a strong positive-split; they started out fast and completed the first-half of the race quite a bit faster than the second-half of the race. In terms of providing this information to the CBR system, it may be feasible for the runner to provide their race number and for the appropriate record to be pulled from the public race records in real-time. Large marathons such as London or Boston could readily accommodate this by integrating the service as part of their site offering.

In any event, having submitted a race record, the CBR system retrieves a set of k cases, $c_1 \dots c_k$ and, using one of the strategies above, generates a predicted *PB* time and suitable race plan for the runner. In this examples the predicted *PB* time indicates to the runner that she should be able to achieve a personal-best finish-time of 237 minutes, a fairly significant 11-minute improvement over her previous time. The recommended pacing plan suggests running the race with a slight positive-split, starting out a few percent faster than their target average pace of 5 minutes and 37 seconds per km, while running a slightly slower second-half of the race; it is a much more evenly paced race plan than their *nPB* race record. Importantly, this pacing plan provides our runner with a set of concrete paces for each of the race segments and because it is based on how other runners (with similar *nPB* records to the query runner) have run the same course; in short, it provides the runner with a more personalized and tailored race plan that might otherwise be the case.

4 Evaluation

To evaluate our approach we will focus separately on the tasks of finish-time prediction and profile recommendation using data from the last six years of the London Marathon (2011 - 2016).

4.1 Dataset

Briefly, our London dataset includes 215,575 race records for 185,143 unique runners. Each race record contains, among other features, the 5km split-times and finish-times that we need as the basis for our case base. 37% of runners are female and just over 20% (or 37,704 unique runners) appear more than once in the dataset. This means we can construct at least 37,704 cases. However, since many of these repeat runners have only run 2 races this does limit the variation available for the identification of personal-best times and profiles. Consequently, in this study we limit our data to runners who have run 3 or more races; this way we can choose a personal best race from at least 3 races and every runner will be represented by at least 2 cases. There are 5,390 such runners and on average they have run 3.4 races each, leading to a case base of 12,968 individual cases with 5,390 personal-bests.

4.2 Methodology

For the purpose of this evaluation we use a standard 10-fold cross-validation methodology to test prediction accuracy and recommendation quality. Briefly, we randomly hold-out 10% of cases to act as a test-set and use the remaining 90% of cases as the basis for prediction and recommendation, repeating this 10 times and averaging the results. Thus, the *nPB* part of the each test case is used as a query and the *PB* part is held back to evaluate the finish-time prediction and pacing plan recommendation.

Evaluation Metrics. To evaluate prediction accuracy we calculate the average percentage error between the predicted personal-best time and the actual personal-best time held back from the test case. To evaluate the quality of the recommended pacing plan we estimate the similarity between this recommended plan and actual pacing profile that was also held back; our measure of profile similarity is based on the relative difference between the segment paces of the recommended plan and the segment paces of the actual pace profile.

Test Algorithms. For each prediction/recommendation session we generate predictions/recommendations using the 3 CBR approaches (*best*, *mean*, *even*) as described earlier. While in theory we could possibly *mix-and-match* between prediction and recommendation strategies, for the purpose of this study, we use a *like-with-like* approach and pair each prediction strategy with the corresponding recommendation strategy. Thus when we refer to the *Best* strategy, for example, we mean the *Best* prediction strategy combined with the *Best* recommendation strategy.

4.3 Prediction Error and Profile Similarity vs. k

To begin with we will look at prediction accuracy and recommendation quality versus k , the number of cases retrieved; see Figures 2 (a) & (b). Both graphs

contain *ribbon-plots* for each of the 3 test algorithms. Each ribbon-plot shows the prediction error (or profile similarity) for all runners, as the marked central lines, and separately for male and female runners, as the boundary lines. In this way we can get a sense of how gender influences prediction error and pacing profile similarity as k varies.

Fig. 2. Prediction error (a) and pacing profile similarity (b) versus k for *Best*, *Mean*, and *Even* strategies, and for all runners and men and women.

With respect to prediction error, we can see how our strategies behave quite differently for increasing k . *Mean* produces the lowest errors and benefits from increasing values of k , before stabilising for $k \geq 10$; *Mean* achieves an error of $\approx 4.5\%$, for all runners, and as low as 4% for women. In contrast, *Even* produces predictions with an average error of $\approx 6\%$ regardless of k , while the accuracy of *Best* deteriorates steadily with increasing k . The problem for *Best* is that more retrieved cases generally lead to faster best finish-times to use as a prediction. So *Best* tends to predict more and more ambitious personal-best times as k increases, times that have not been achieved by our test runners.

When it comes to pacing profile similarity – our proxy for recommendation quality – we see a similar, albeit inverted pattern in Figure 2(b). The *Mean* strategy out-performs the others, recommending pacing plans that are increasingly similar to actual *PB* pacing profiles (before stabilising at $k \approx 5$). The profile similarity of *Even* recommendations are largely unaffected by k , while the *Best* recommendations become less and less similar to the actual *PB* pacing profiles as k increases; once again more cases mean faster best C_{Best} 's with pacing profiles that are increasingly different from the personal-best race records of the test users.

It is also interesting to note how the prediction accuracy (and profile similarity) for all strategies is better for women than for men, regardless of k . This is something that is not so surprising in marathon running research, where recent studies have shown how female runners tend to run in a more disciplined manner than their male counterparts [7, 8]. In short, women run more evenly paced races, they are less likely to hit the wall, and so their races unfold in a more predictable fashion, compared to men’s races. Our results suggest that this also extends to the matter of predicting a personal-best time and recommending a bespoke pacing plan.

4.4 On the Influence of Ability

These results speak to the utility of the *Mean* strategy when it comes to making accurate personal-best predictions and recommending high quality pacing plans. But the difference between men and women suggests that all runners are not equal. Building on this idea, it is also useful to consider the relationship between accuracy/quality and a runner’s ability, in terms of finish-times. For example, is it easier to predict personal-best times for faster or slower runners?

Fig. 3. Prediction error (a) and pacing profile similarity (b) vs. nPB finish-time for *Best*, *Mean*, and *Even* strategies, for all runners and men and women.

These prediction error and profile similarity results, for different finish-times, are presented in Figures 3(a) & (b); for simplicity the results have been averaged over all values of k . Clearly, finish-times have a significant impact on both prediction error and profile similarity. For example, the fastest (elite) runners benefit from very accurate personal-best predictions by all three strategies, but as finish-times increase so too do prediction errors, at different rates for the different strategies. Once again the *Mean* strategy benefits from the most accurate

predictions across all finish-times, and the difference between men and women generally persists. It is interesting too to note that error rates begin to fall again after the 300-minute mark, suggesting that the *PBs* of slower finishers are more predictable.

The similarity of recommended pacing profiles to the actual personal-best profiles falls as finish-times increase; see Figure 3. The *Mean* strategy continues to perform better than *Even* and *Best* and women enjoy more similar pacing profiles than men.

4.5 Personal-Best Improvements

One more factor that may influence prediction accuracy and recommendation quality is the ambitiousness of a personal-best. Are very ambitious personal-bests more or less difficult to predict and plan for? To test this we define the *PB Difference* as the relative difference between a runner's non personal-best time and their personal-best time; for example, a *PB* difference of 0.1 indicates that the *PB* time of a case is 10% faster than its *nPB* time.

Fig. 4. Prediction error (a) and pacing profile similarity (b) vs. *PB Difference* for *Best*, *Mean*, and *Even* strategies, and for all runners and men and women.

Figures 4(a) & (b) shows how prediction error and profile similarity are influenced by *PB Difference*. As we might expect, runners whose *PB* time are less ambitious (lower *PB Difference* values) enjoy more accurate predictions for *Mean* and *Even* than runners with larger differences between their *PB* and *nPB*; see Figure 4(a). Likewise, profile similarity is also highest for runners with lower *PB Difference* values; see Figure 4(b).

Once again the *Best* strategy performs poorly in general, for prediction error and profile similarity, with one notable exception. If *PB* times are more than 12%

faster than their *nPB* times, then the inherently ambitious *Best* strategy outperforms *Mean* and *Even*. That said, it is worth noting that a 12% improvement in marathon finish-time represents a very significant improvement, and is not so common in practice; in our dataset less than 20% of runners achieve such a personal-best improvement.

5 Discussion

These results point to *Mean* as the best strategy to use in practice, because it offers more accurate *PB* finish-time predictions and more similar *PB* pacing profiles when compared with the actual *PBs* of the test runners. *Even* also performs well but *Best* appears to be far too ambitious.

5.1 Personal-Best Improvements

With all this discussion about error rates and similarities it is easy to lose sight of the practical reason for the proposed system: to help marathon runners plan for, and achieve, new personal-best finish-times. And so, as we draw to a conclusion, it is worth forming a more concrete sense of what these new *PB* times are likely to be for runners' *nPB* times. These results are presented in Figure 5 for our most accurate strategy, *Mean*, as a graph of the predicted *PB* improvement in minutes versus the *nPB* time.

As we might expect, the predicted *PB* improvement increases with finish-time; it is 'easier' for a 5-hour marathoner to increase their time by 20 minutes than it is for a 3-hour marathoner. What is particularly encouraging here is the scale of the improvements on offer. For example, a 4-hour marathoner can expect to improve their *PB* time by just over 20 minutes on average (± 75 seconds based on the error associated with predictions for such runners). This should be very encouraging for most recreational marathoners, and the combination of these accurate predictions and sensible, tailored pacing profiles may help many enthusiastic marathoners to improve significantly in their next race.

5.2 Pushing Beyond the *PB*

Our various strategies have been evaluated using the actual *PB* races of test runners and it should be pointed out that it remains unclear whether these runners *could* have produced even better *PBs* had they received the right pacing advice when they ran these races; no doubt some did but most will have adopted fairly simply pacing strategies. Perhaps with the right pacing advice the more ambitious *Even* or *Best* predictions might be achievable on the day.

Certainly conventional marathon wisdom suggests that a more even pacing profile is better and thus in practice we may find an additional benefit to presenting runners with the *Even* pacing plan even though its corresponding finish-time prediction is marginally less accurate (faster) than the *mean*. By following this *Even* profile the runner may achieve a more ambitious finish-time. Testing this

Fig. 5. Prediction *PB* finish-times for runners based on their non *PB* finish-times.

hypothesis requires further work but it speaks to the potential benefits of presenting *Mean* and *Even* predictions and recommendations to runners, rather than focusing exclusively on just one strategy, such as *Mean*.

5.3 PB Quality

Another factor worth considering is the method by which *PBs* have been identified. Currently they are simply the fastest race record for a runner. They may or may not reflect a genuine *PB*; i.e. a race that the runner ran to the best of her true ability and training. As part of future work we plan to explore this by using different techniques for selecting more reliable *PBs*. Limiting our case base to runners who have run many races is one way to do this. Another way might be to limit *PBs* to those race records that have more even pacing profiles. Either way these approaches will provide a way to at least partially evaluate the likely impact of *PB* quality.

6 Conclusions

We have proposed a case-based reasoning solution to help marathon runners achieve a personal-best at their next race. It represents a novel application of CBR to an important open problem for marathon runners, namely how to select a best achievable target-time and how to pace the race to achieve it.

We have described and evaluated three CBR variations using a large-scale, real-world dataset from the London Marathon. The results indicate that the *Mean* CBR approach is capable of accurately predicting personal-best finish-times and of recommending high-quality pacing plans. This may prove to be extremely valuable in helping runners to calibrate their expectations and plan effectively for their next race-day. By using training data from other marathons it will be possible to produce similarly personalised predictions and tailored recommendations for different marathon courses.

As part of future work we plan to explore other factors that likely impact prediction and recommendation performance, *PB* quality for example, as discussed, but also the length of time between a *nPB* and a *PB* in a case for instance; the race-pair represent races 5 years apart, for example, will this effect prediction or recommendation? Obviously we also plan to test this approach on other big city marathons such as Boston, Chicago, and Berlin. Some of these races have strict finish-time based qualification requirements, which aill no doubt effect the quality of the cases in their corresponding case bases. Will this also effect prediction or recommendation performance?

7 Acknowledgments

Supported by Science Foundation Ireland through the Insight Centre for Data Analytics under grant number SFI/12/RC/2289 and by Accenture Labs, Dublin.

References

1. A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson, H. Lu, X. Zheng, M. Musolesi, K. Fodor, and G.-S. Ahn, “The rise of people-centric sensing,” *IEEE Internet Computing*, vol. 12, no. 4, 2008.
2. V. Mayer-Schönberger and K. Cukier, *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt, 2013.
3. R. H. Ellaway, M. V. Pusic, R. M. Galbraith, and T. Cameron, “Developing the role of big data and analytics in health professional education,” *Medical teacher*, vol. 36, no. 3, pp. 216–222, 2014.
4. I. Bichindaritz, S. Montani, and L. Portinale, “Special issue on case-based reasoning in the health sciences,” *Applied Intelligence*, vol. 28, no. 3, pp. 207–209, 2008.
5. M. Lewis, *Moneyball: The art of winning an unfair game*. WW Norton & Company, 2004.
6. D. Kelly, G. F. Coughlan, B. S. Green, and B. Caulfield, “Automatic detection of collisions in elite level rugby union using a wearable sensing device,” *Sports Engineering*, vol. 15, no. 2, pp. 81–92, 2012.
7. N. W. Trubee, *The Effects of Age, Sex, Heat Stress, and Finish Time on Pacing in the Marathon*. PhD thesis, University of Dayton, 2011.
8. R. O. Deaner, “More males run fast: A stable sex difference in competitiveness in us distance runners,” *Evolution and Human Behavior*, vol. 27, no. 1, pp. 63–84, 2006.
9. T. A. Haney Jr, *Variability of pacing in marathon distance running*. PhD thesis, University of Nevada, Las Vegas, 2010.
10. C. Foster, M. Schrager, A. C. Snyder, and N. N. Thompson, “Pacing strategy and athletic performance,” *Sports Medicine*, vol. 17, no. 2, pp. 77–85, 1994.
11. C. R. Abbiss and P. B. Laursen, “Describing and understanding pacing strategies during athletic competition,” *Sports Medicine*, vol. 38, no. 3, pp. 239–252, 2008.
12. C. D. Stevinson and S. J. Biddle, “Cognitive orientations in marathon running and ”hitting the wall” .,” *British Journal of Sports Medicine*, vol. 32, no. 3, pp. 229–234, 1998.