Running with Recommendation

Jakim Berndsen
Insight Centre for Data Analytics,
University College Dublin, Belfield,
Dublin 4
jakim.berndsen@insight-centre.org

Aonghus Lawlor
Insight Centre for Data Analytics,
University College Dublin, Belfield,
Dublin 4
aonghus.lawlor@insight-centre.org

Barry Smyth
Insight Centre for Data Analytics,
University College Dublin, Belfield,
Dublin 4
barry.smyth@insight-centre.org

ABSTRACT

We examine the feasibility of a collaborative recommender system in the exercise domain targeted specifically at runners. By using a large dataset of over 600000 runners' finish times we explore the contrasts between casual and elite runners and hypothesise how a recommender system may be used to mitigate some of these differences. We also briefly discuss some of the challenges faced by such a recommendation task and suggest how these challenges could be addressed.

CCS CONCEPTS

• **Information systems → Recommender systems**; • **Human-centered computing → Social recommendation**;

KEYWORDS

Sports Analytics, Explanations

ACM Reference format:
Jakim Berndsen, Aonghus Lawlor, and Barry Smyth. 2017. Running with Recommendation. In *Proceedings of Second International Workshop on Health Recommender Systems co-located with ACM RecSys 2017, Como, Italy, August 2017 (RecSys'17)*, 4 pages.

1 INTRODUCTION

Running is one of the most popular forms of exercise on the planet. In 2015 over 17 million race finishers were recorded in the United States alone [1]. The proliferation of online resources promoting running shows the growing popularity of the sport and is a persuasive demonstration of the participants' interest in improving their performance, and the desire to improve the performance over time, while avoiding injury, is a key motivation for us to address as a recommender system problem [10].

Despite the popularity of running and the availability of training resources, it remains a difficult pastime. Completing a race is the end result of weeks or months of meticulous training and planning. Selecting the best balance of hard work, recovery, and different training types remains a concern that requires an in-depth knowledge of running and human performance. While runners will exhibit natural performance increases simply by starting to run and by improving their fitness, it is extremely difficult to further optimise performance without access to a coach or extensive time investment in training methodologies. The availability of coaches is limited and usually reserved for competitive athletes, rather than those simply trying to improve themselves and maximise their personal performance. There also exist many conflicting training methodologies and thus despite extensive research an unsuitable training programme may be selected by a runner.

In an attempt to counteract these problems we propose a novel collaborative recommender system. Our system is designed to recommend training plans and race strategies to a runner and thus alleviate the need for a coach and reduce the requirement for extensive research to devise a personalised training plan best suited to the runner. By mining the data on training plans of runners with similar histories we can suggest a training plan to the user that they can follow with little cognitive effort and that will allow them to make significant gains over a period of time.

In this paper, we show that elite runners progress faster than casual runners and we surmise that this is due to the fact they have access to coaches, research, and have a tendency to approach training in a more strategic manner. Imparting the expertise that elite runners have garnered to a casual runner through the use of a recommender system will lead to casual runners exhibiting increased rates of improvement, similar to that of the elite runner. We will also highlight the suitability of a collaborative approach to recommendation, outline the time scales involved in such recommendation and briefly suggest some approaches towards recommendation of material in this domain. While the work presented in this paper focuses on the marathon distance, the results are generalisable across different race distances, and as such running as a whole.

2 RELATED WORK

The use of machine learning for athletic performance is still in its early phases. Much of the work previously undertaken in this field has focused on the problem of prediction - the question of how fast a runner would run a race given a previous race they have run at a different distance.

One of the earliest works in this regard was undertaken by Peter Riegel [9] in 1981. Riegel examined world record times for various activities, such as running and swimming, and found there to be a linear relationship between the log of the time taken and the log of the distance of the event. He thus fit an exponential equation of the form $t = ax^b$ to the world record times of different running events, where t is the time taken, x is the distance and a and b are constants. In this equation b can be considered the slowdown coefficient, or fatigue factor, and was found to take a value of 1.06 for running through the regression analysis. Riegel's formula proves to be very effective for predicting times for races of distances shorter than the marathon and is also most effective for the category of elite marathon runners. However, it struggles to predict the times for casual runners in the marathon, especially

[1] http://www.runningusa.org/statistics

for times slower than 230 minutes (as can be seen in Figure 2). 230 minutes is significant as it is in events that take longer than this that the linear relationship between log time and log distance no longer holds due to the body switching to different energy systems. The Riegel formula underestimates times drastically for the category of slower runners, which can have a disastrous effect on their race performance if these predictions are used to inform race tactics.

Despite its limitations Riegel's formula is the most commonly used method for predicting race times today. It is the equation used by many well known websites, most notably RunnersWorld (www.runnersworld.com).

A further regression analysis was performed by David Cameron [3] in 1997. This study fit a regression through the 7 fastest times in each event at the time. This regression offers a slightly better fit but suffers from the same problems as the Riegel formula. It also suffers from an under prediction problem at times greater than 230 minutes. Due to the relative simplicity of the Riegel compared to the Cameron Time Equivalence Model and their highly comparative results the Riegel formula is often preferred to the Cameron model by various running resources.

A study [14] in conjunction with slate.com aimed to address some of the inadequacies of these systems. A survey asked runners to report their own race times with the aim of building a better marathon predictor. The survey responses led to a dataset of 2164 usable responses. While again using linear regression analysis, the model was able to utilise further information about a runner's race history. The feature set was comprised of the two previous races run and the reported weekly training load of a runner. The results show that such a model outperforms the Riegel model, especially for more casual runners. However, the model still has limitations. The dataset size is relatively small and thus it is difficult to utilise more advanced machine learning techniques that require large amounts of data. The prediction improvement seen by using a second race suggests that additional information about a runner's history is beneficial in making predictions, yet is in itself limited as it does not take account that a runner may have run many hundreds of races previously. Despite these limitations, the decreases in prediction error have seen this model adopted by many runner's resource websites, including RunnersWorld. The model's results show the promise of what data analysis and machine learning techniques can achieve in this domain and that further strides are possible.

Moving away from prediction into the field of recommendation, Smyth [12][13] proposes a case-based recommendation system for recommending a personal best marathon time. Using information of a runner's previous marathon history and a case base of other runner's times Smyth recommends not only a time that runner should be capable of, but also a race plan to achieve the time taking into account the difficulty and terrain of the course on which the runner will compete. The recommendation made is important, as it should be difficult enough to feel tested but not be so fast that a runner ends up hitting the wall. By testing against the actual personal best run by a runner the study found that the system was able to predict the personal best time of a runner to within 5% accuracy and generate race plans that are more than 90% similar to the actual personal best split times run. These results show that there is scope for making recommendations in the domain of distance running.

3 RECOMMENDATION

3.1 Dataset

Table 1: Dataset Description

Number of Runners	618318	
Num Races	8212756	13.28 (per runner)
#Races	8212756	13.28 (per runner)
Marathon Runners	522943(M: 62%)	324165(F: 38%)
#Marathons	852157	2.97 (per runner)
Marathon Times	260±63	
Mean Time Between Marathons	352 (days)	
#Different Races per Runner	4.19	
#Mean Time Between Races	83 (days)	

As mentioned in the related work section the current methods used for marathon prediction rely on data ranging from a single race at each distance (i.e. the Riegel Model) to a few thousand self reported race times. In contrast, we have built a dataset scraped from various sources including race results tables and websites allowing athletes to self declare race times. This dataset contains over 600000 runners and their entire race histories. This is the first dataset of this scale that has been collected and allows for the first large scale data analysis and machine learning approach towards making predictions for a marathon time. While we do not have complete training histories for runners, which would give greater resolution in a machine learning problem, we attempt to approximate their training schedules by looking at the frequency of the races of various non-marathon distances the athlete has run and the improvements they make as they run them. Additional features, such as age and gender, allow us to further distinguish between runners in a way that has been neglected by many of the previous prediction models mentioned in Section 2. These features, particularly gender, have significant effects when it comes to distance running [1, 5].

3.2 Runner Improvement

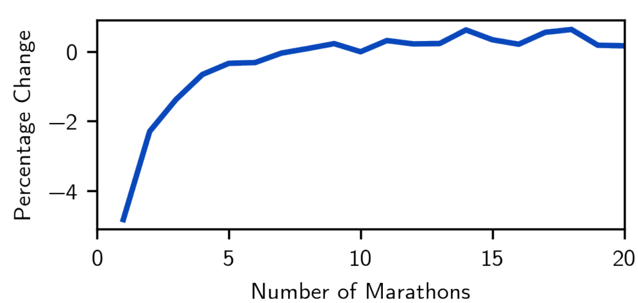

Figure 1: Average percentage change in finish time of runners from one race to the next as exhibited by runners in the dataset

Figure 1 depicts how a runner improves from one marathon race to the next. Typically a runner will see a large and steady improvement for the first 3-4 races they run, before witnessing a

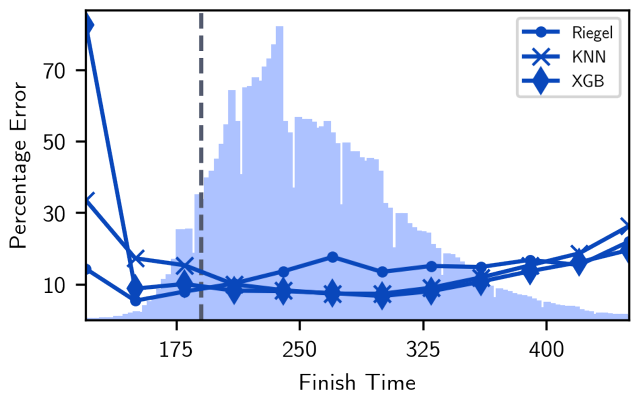

Figure 2: Average errors of Riegel, KNN and XGB models over time

plateau in performance. This natural performance improvement is a first indication of the potential benefits of a recommender system. Most runners do not have access to advanced training methods and lack the motivation or finances to employ a running coach, yet they still exhibit significant improvement from one race to the next. Finding suitable training plans requires a significant amount of research as there are many different methodologies runners have used in order to get results. A recommender system designed to assist runners could act in the role of a professional running coach and help personalise training plans based on data mining the successful training approaches of similar runners. The recommender system removes the time burden required to map out an adequate training plan and helps the runner to improve at the fastest rate.

3.3 Suitability of Collaborative Filtering

After creating a basic user profile for each runner comprising average race finish times at various non-marathon distances, we trained two different models to predict marathon times. These models were a simple K-Nearest Neighbours (KNN) model [6] and an Extreme Gradient Boosting (XGB) model [4]. The total percentage errors of these models are 10.45% (KNN) and 9.04% (XGB), which compare favourably to the error exhibited by the Riegel model of 12.8%.

Figure 2 shows how these prediction errors change for athletes with different race finish times. The Riegel model is indeed more accurate than KNN or XGB for elite runners (finish times < 190mins), but the vast majority of runners have slower finish times than this. For these slower runners, the KNN and XGB models outperform the Riegel model. This suggests that the use of user profiles and computing similarity between them is a good approach for describing how a runner is likely perform. This provides some justification that a collaborative filtering system would provide an adequate basis for building a recommender system for runners. The success of Smyth's work [12] in recommending personal best times outlined in Section 2 also appears to corroborate this finding.

It could also be pointed out here that the simpler Riegel model out-performs both the KNN and XGB models at various points in the distribution. This is certainly the case for the quickest runners, which is not surprising as the Riegel model is based on world record times and we have little data at these points from which to build a similarity model. Similarly, for the very slowest runners, we also have a data sparsity issue. However, for the purposes of building a running recommender this is not a problem. The elite runners, with whom our model struggles, tend to be professional athletes or passionate runners. These runners tend to be well coached and well informed on their training and thus a recommender system is expected to be of limited value to them. Runners at the very slowest end of the spectrum tend to be one-off runners that are running purely for fun and they are also unlikely to see any value in using a recommender system.

3.4 Benefits of Recommender System

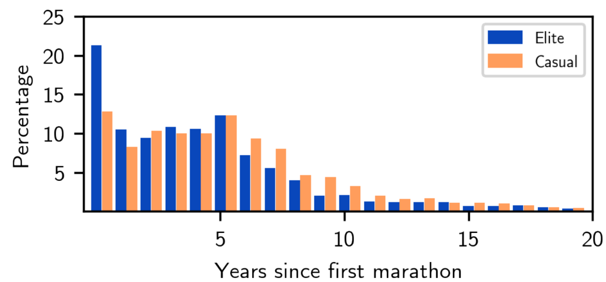

Figure 3: Proportion of runners running Personal Best performances for elite and casual runners based on time since the first marathon run

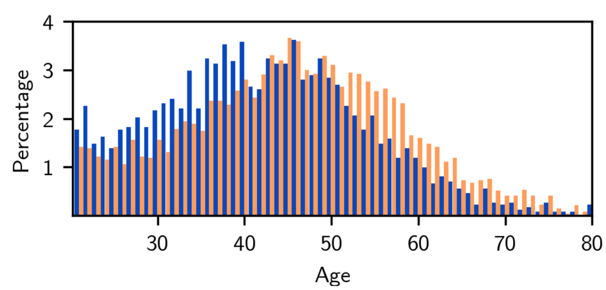

Figure 4: Proportion of runners running Personal Best performances for elite and casual runners based on age at time of personal best

We make the reasonable claim that elite runners either have access to coaching or are well-informed on training methods. As a result, they produce faster times but there are other phenomena that are a side effect of this optimised training. We define an elite runner as one who has run the qualifying standard for the Boston Marathon of 190 minutes, and then we compare the performance progression over time of elite runners and more casual runners. The Boston Marathon qualifying time was chosen as this is considered a goal time for many keen marathon runners and is a time that requires substantial training and effort to achieve. We also tell from Figure 2 that it corresponds roughly to the finish time at which the user based models begin outperforming the Riegel model which makes it a natural cut off.

In Figure 3 we show the point at which elite and casual runners first run their overall Personal Best (PB) time since running their first marathon. The elite runners clearly peak much earlier than casual runners, with nearly 20% of elites achieving a PB in the first year after they start marathon running. In contrast, a higher fraction of casual runners achieve PB's than elite runners, after a period of 5 years from their first marathon. In Figure 4 we show the age at which the personal best is achieved. Again, there is a strong difference between elite and casual runners, with elites much more likely to achieve the PB before the age of 40. Elite runners not only achieve their PB's at a younger age but also at an earlier stage of their running career. This finding affirms the notion that the extra knowledge elite runners have over casual runners is a significant advantage when it comes to making performance gains. It is clear that a recommender system would be useful in this field - an automated personalised training plan generated by a collaborative recommender system would mitigate the need for a professional coach or extensive knowledge of running training for a casual runner. Such recommendations should lead to faster performance gains from casual runners and would see casual runners maximise their potential earlier.

3.5 Methods of Recommendation

It is important to note at this point that not even all elite runners maximise their potential quickly. Many elite runners will not run their personal best until up to five years after their first marathon. This demonstrates the potential time scale involved in such a recommender system with improvements not being apparent for months or even years after the first interaction with the system.

Such a recommendation system poses a unique challenge. How does a recommender system motivate a runner to keep using a system for a period of years, especially when the benefits of use may not be instant? Training for a marathon is difficult and a recommender may recommend a training session that, while clearly beneficial, may not lead to enjoyment or satisfaction for the user. The recommendation system must therefore keep a user engaged for long periods and convince them to make potentially unwanted decisions in order for them to see benefit.

An important factor in achieving this goal is to provide the user with meaningful explanations. The ability of a system to make its reasoning transparent contributes significantly to the users acceptance of the recommendation [2] and improves their confidence in the recommendation [11]. Various training methodologies are already well documented and explained. The concept of nudging [8], to slowly adjust the user's behaviour, has been shown to be very useful in recommender systems. The use of personalised explanations can motivate a user to interact with the system and spur them on to do the sessions the system recommends. As demonstrated in Section 3.3 the system is capable of making predictions of the runner's finish time. As the recommender system nudges the user to a particular training strategy, accurate predictions of finish times and outcomes can be presented to the user to improve their motivation and engagement with the system. For runners to gain the maximum benefit from such a system it must be persuasive, easy to follow, and provide motivation so the recommender is engaging for long enough to have an effect on a runner's training and change their behaviour [7].

4 CONCLUSION

In this paper we have examined the opportunity for a recommender system for runners. Elite runners improve at a faster rate through knowledge gained from coaching and research. This knowledge can be imparted to casual runners through the use of a recommender system, leading to greater levels of improvement for such runners. We have gathered an adequately large database of runners and race times and through this have shown that marathon times can be predicted by simple neighbourhood models. The better prediction accuracy of these models highlights the feasibility of a collaborative filtering approach to such recommendation. Lastly, we highlighted some of the issues with such a recommender system, namely the time scale involved. We suggested methods as to how such recommendation could be presented to runners in order to keep them sufficiently motivated and allow adequate time for this recommendation to take effect.

In future we will implement such a recommender system. We expect that the results found will be generalisable to other endurance sports and as such we look to expand our research into activities such as swimming and cycling. The proliferation of wearable technology, such as heart rate monitors and GPS units, provides large quantities of data from training events and races. Such data will provide greater resolution for a machine learning approach and we will use such an approach to build a recommender system that can inform a user before, during, and after a training session or race.

ACKNOWLEDGMENTS

This work is supported by Science Foundation Ireland through the Insight Centre for Data Analytics under grant number SFI/12/RC/2289

REFERENCES

[1] R Beneke, R Leithauser, and M Doppelmayr. 2005. Women will do it in the long run. *British journal of sports medicine* 39, 7 (2005), 410.
[2] Mustafa Bilgic and Raymond J Mooney. 2005. Explaining recommendations: Satisfaction vs. promotion. In *Beyond Personalization Workshop, IUI*, Vol. 5. 153.
[3] "David F. Cameron". 1998. (1998). "http://www.cs.uml.edu/~phoffman/cammod.html"
[4] Tianqi Chen and Tong He. 2015. Xgboost: extreme gradient boosting. *R package version 0.4-2* (2015).
[5] Robert O Deaner. 2006. More males run fast: A stable sex difference in competitiveness in US distance runners. *Evolution and Human Behavior* 27, 1 (2006), 63–84.
[6] Daniel T Larose. 2005. k-Nearest Neighbor Algorithm. *Discovering Knowledge in Data: An Introduction to Data Mining* (2005), 90–106.
[7] Ting-Peng Liang, Hung-Jen Lai, and Yi-Cheng Ku. 2006. Personalized content recommendation and user satisfaction: Theoretical synthesis and empirical findings. *Journal of Management Information Systems* 23, 3 (2006), 45–70.
[8] Eoghan O'Shea, Sarah Jane Delany, Rob Lane, and Brian Mac Namee. 2014. NudgeAlong: A Case Based Approach to Changing User Behaviour. In *International Conference on Case-Based Reasoning*. Springer, 345–359.
[9] Peter S Riegel. 1981. Athletic Records and Human Endurance: A time-vs.-distance equation describing world-record performances may be used to compare the relative endurance capabilities of various groups of people. *American Scientist* 69, 3 (1981), 285–290.
[10] Hanna Schaefer, Santiago Hors-Fraile, Raghav Pavan Karumur, André Calero Valdez, Alan Said, Helma Torkamaan, Tom Ulmer, and Christoph Trattner. 2017. Towards Health (Aware) Recommender Systems. *Proc. of DH* 17 (2017).
[11] Rashmi Sinha and Kirsten Swearingen. 2002. The role of transparency in recommender systems. In *CHI'02 extended abstracts on Human factors in computing systems*. ACM, 830–831.
[12] Barry Smyth. 2017. A Novel Recommender System for helping Marathons to Achieve a new Personal-Best. *Proceedings of ACM RecSys 2017, Lake Como, Italy, August 2017* (2017).
[13] Barry Smyth and Padraig Cunningham. 2017. 'Running with Cases: A CBR Approach to Running Your Best Marathon'. *Proceedings of ICCBR 2017, Trondheim, Norway, June 2017* (2017).
[14] Andrew J Vickers and Emily A Vertosick. 2016. An empirical study of race times in recreational endurance runners. *BMC Sports Science, Medicine and Rehabilitation* 8, 1 (2016), 26.