

Marathon Race Planning: A Case-Based Reasoning Approach

Barry Smyth, Pádraig Cunningham

Insight Centre for Data Analytics

School of Computer Science, University College Dublin, Ireland

barry.smyth@ucd.ie, padraig.cunningham@ucd.ie

Abstract

We describe and evaluate a novel application of case-based reasoning to help marathon runners to achieve a personal best by: (a) predicting a challenging, but realistic race-time; and (b) recommending a race-plan to achieve this time.

1 Introduction

This work aims to help marathoners to achieve a new personal best (*PB*) by using case-based reasoning [Smyth, 2007; de Mántaras *et al.*, 2005; Bridge *et al.*, 2005] to address two problems faced by runners: (a) *predicting* an achievable finish-time; and (b) *recommending* a tailored pacing plan to help the runner achieve this time.

While health applications of AI are not new [Peek *et al.*, 2015; Buchanan and Shortliffe, 1984; Bichindaritz *et al.*, 2008; Wiesner and Pfeifer, 2014; 2010], always-on smartphones and wearables are creating an even greater opportunity for novel preventative, proactive, and personalised interventions [Ohlin and Olsson, 2015; Geleijnse *et al.*, 2011]. The sports world has embraced the data-centric world of sensors [Campbell *et al.*, 2008] and mobile apps, as athletes explore the power of data to optimise training and performance [Lewis, 2004; Kelly *et al.*, 2012; Leijdekkers and Gay, 2015; Möller *et al.*, 2011; Hermens *et al.*, 2014].

We focus on helping recreational runners to plan their race strategy. A key concept is *pace*, which is measured in minutes per mile/km – the inverse of *speed* – so that higher values correspond to slower speeds and vice versa [Trubee, 2011; Foster *et al.*, 1994; Deaner, 2006; Haney Jr, 2010]. There are three basic pacing strategies. 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, running it *faster* than the first. Many elites and disciplined runners aim for even or slightly negative-splits [Abbiss and Laursen, 2008]. Recreational runners typically run positive-splits, slowing during the second-half of the race, and sometimes even *hitting the dreaded wall* [Foster *et al.*, 1994].

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 and pacing: the finish-time they aim for; whether they opt for positive, negative or even splits; whether they avoid hitting the wall etc. This is where we believe there is a significant opportunity to support marathon runners, by advising them on a suitable target finish-time, and by providing them with a concrete race-plan, one that is personalized to their ability and tailored to a specific marathon course.

2 Problem Definition

Our objective is to support a runner who has completed at least one previous marathon and so has a race record to serve as a starting point.

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 a 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 stages or *segments* (0-5km, 5km-10km, ..., 35km-40km), plus a final 2.2km segment (40km-42.2km). A race-plan, or *pacing plan*, consists of a sequence of average paces (measured in minutes per km) for segments. For example, Figure 1(a) illustrates a race-record for a runner who completed a marathon in 4 hours and 13 minutes. The pacing profile shows relative paces — that is segment paces relative to the runner’s mean race pace — and indicates a positive-split. They started their race (the first 5km and 10km segments) 10% faster than their average pace, before slowing in the second-half, to finish 7% slower than their average pace in the final segment. We wish to generate *tailored* pacing plans, so that runners can benefit from a plan that is suitable for their goal time, their personal fitness level and ability, and that reflects the peculiarities of a given marathon course.

3 Using CBR to Achieve a Personal Best

One insight of this work is to use a CBR approach to predict suitable target finish-times and recommend pacing plans, based on the runner's own race experience and the experiences of similar runners. 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(a). Each case contains two race records; see Figure 1(b). One of these race records corresponds to a *non personal best (nPB)* race, the other corresponds 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

Each case in the case base corresponds to a single runner, r , with a *nPB* part and a *PB* part; see Equation 1. To be represented in the case base, r must 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 best finish-time is designated the personal best, and it is paired with the remaining $n - 1$ 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 paired with the two *nPB* records (m_1 and m_3) to produce two 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, because physiological differences between men and women have a material impact on marathon performance.

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 .

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 , but who has gone on to

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

achieve a faster personal best on the same marathon course. We test three prediction approaches, with the predicted *PB* finish-times 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.

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

Best PB. In Equation 3 the predicted time is the weighted *PB* finish-time of the single fastest retrieved case, C_{best} .

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

Mean PB. Our second prediction approach calculates the *weighted mean* of the *PB* finish-times of the retrieved cases, as in Equation 4.

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

Even PB. *Even PB* is based on the idea that more even pacing is better than more varied pacing. By measuring the coefficient of variation (CoV) of the relative paces in the *PB* pacing profiles, we can 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

We 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 uses the relative pacing from C_{best} and maps its relative paces to the average pace for the predicted *PB* time for the query runner. For example, if the predicted *PB* time

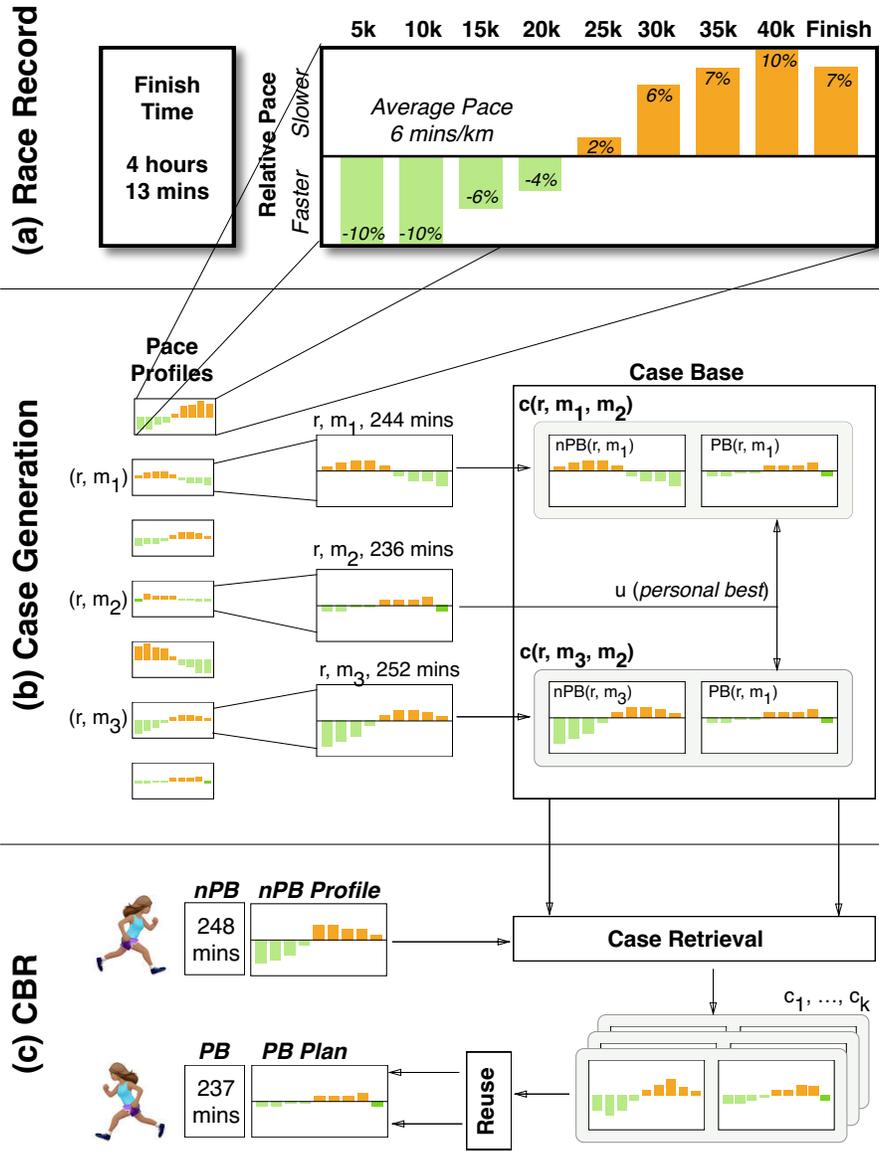


Figure 1: Races, cases, predictions, and recommendations.

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 generates a new pacing plan based on the *mean* relative segment paces of the PB profiles from the k retrieved cases.

Even Profile generates a pacing plan from the PB profile of C_{even} , the case with the most even pacing.

4 Evaluation

We evaluate our approach using 12,968 cases from runners of the London Marathon (2011 - 2016), who have run at least

3 races. We use 10-fold cross-validation 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% for prediction and recommendation, repeating this 10 times and averaging the results. The nPB part of each test case is used as a query and the PB part is held back to evaluate the prediction and pacing plan recommendation.

We generate predictions/recommendations using the three CBR approaches (*Best*, *Mean*, *Even*) described earlier. 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

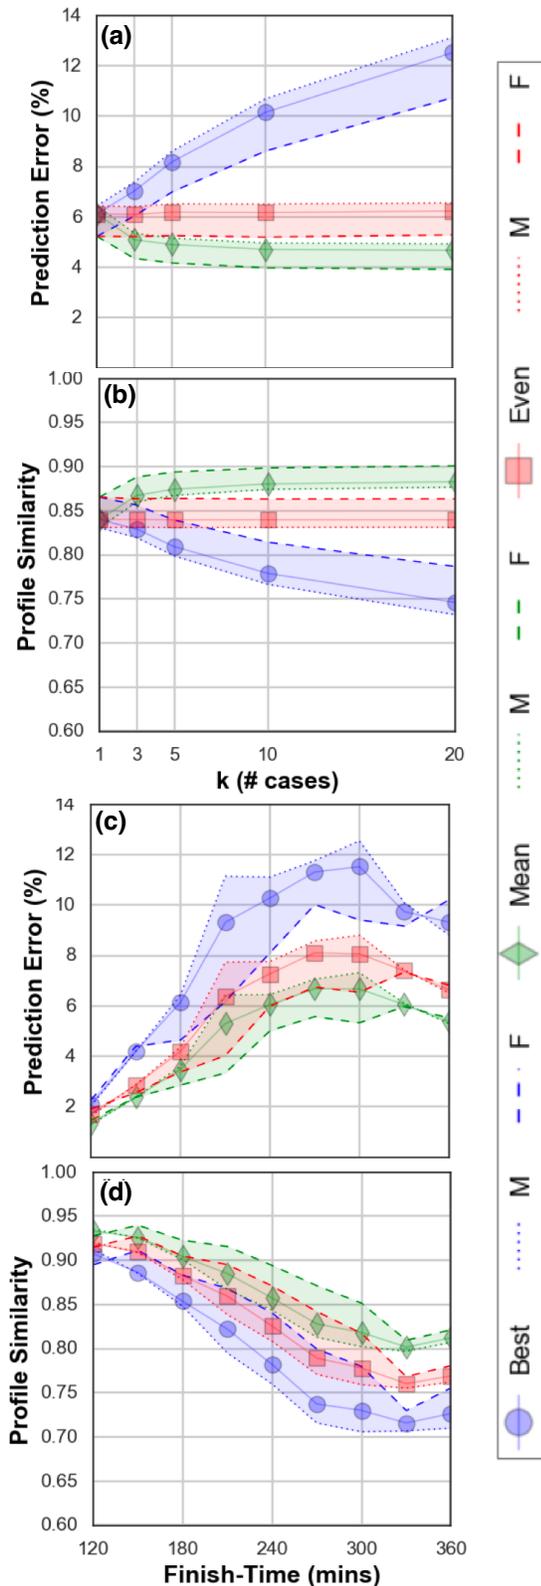


Figure 2: Prediction error and pacing similarity vs. k and nPB finish-time for *Best*, *Mean*, and *Even* strategies.

actual pacing profile that was also held back.

4.1 Prediction Error and Profile Similarity vs. k

We begin with prediction accuracy and recommendation quality versus k , the number of cases retrieved; see Figures 2(a) & (b). Our strategies behave differently for increasing k . *Mean* produces the lowest errors and benefits from increasing values of k , up to $k \geq 10$; *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 with k . The problem for *Best* is that more retrieved cases lead to faster, best finish-times as predictions. So *Best* tends to predict increasingly ambitious PB times as k increases, times that have not been achieved by our test runners.

For pacing similarity — our proxy for recommendation quality — we see a similar, albeit inverted pattern in Figure 2(b). Note too how the prediction accuracy (and profile similarity) for all strategies is better for women than for men, regardless of k . In short, women run more evenly paced (predictable) races [Trubee, 2011; Deaner, 2006]. Our results suggest that this also extends to the matter of predicting a personal best time and recommending a bespoke pacing plan.

4.2 On the Influence of Ability

It is also useful to consider the relationship between accuracy/quality and a runner’s ability level, in terms of finish-times. These prediction error and profile similarity results, for different finish-times, are presented in Figures 2(c) & (d); 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, and at different rates for 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.

The similarity of recommended profiles to the actual PB profiles falls as finish-times increase; see Figure 2(d). The *Mean* strategy continues to perform better than *Even* and *Best* and women enjoy more similar pacing profiles than men.

5 Conclusions

We have summarised a CBR solution to help marathon runners achieve a personal best. The results indicate that the *Mean* CBR approach is capable of accurately predicting personal best finish-times and of recommending high-quality pacing plans. A more comprehensive treatment and additional research can be found in [Smyth and Cunningham, 2017a] and [Smyth and Cunningham, 2017b].

Acknowledgments

This work is supported by Science Foundation Ireland under grant number SFI/12/RC/2289 and is an abridged version of a paper titled "Running with Cases: A CBR Approach to Running your Best Marathon", from the 25th International Conference on Case-Based Reasoning.

References

- [Abbiss and Laursen, 2008] Chris R Abbiss and Paul B Laursen. Describing and understanding pacing strategies during athletic competition. *Sports Medicine*, 38(3):239–252, 2008.
- [Bichindaritz *et al.*, 2008] Isabelle Bichindaritz, Stefania Montani, and Luigi Portinale. Special issue on case-based reasoning in the health sciences. *Applied Intelligence*, 28(3):207–209, 2008.
- [Bridge *et al.*, 2005] Derek G. Bridge, Mehmet H. Göker, Lorraine McGinty, and Barry Smyth. Case-based recommender systems. *Knowledge Eng. Review*, 20(3):315–320, 2005.
- [Buchanan and Shortliffe, 1984] Bruce G. Buchanan and Edward H. Shortliffe. *Rule Based Expert Systems: The Mycin Experiments of the Stanford Heuristic Programming Project (The Addison-Wesley Series in Artificial Intelligence)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1984.
- [Campbell *et al.*, 2008] Andrew T Campbell, Shane B Eisenman, Nicholas D Lane, Emiliano Miluzzo, Ronald A Peterson, Hong Lu, Xiao Zheng, Mirco Musolesi, Kristóf Fodor, and Gahng-Seop Ahn. The rise of people-centric sensing. *IEEE Internet Computing*, 12(4), 2008.
- [de Mántaras *et al.*, 2005] Ramon López de Mántaras, David McSherry, Derek G. Bridge, David B. Leake, Barry Smyth, Susan Craw, Boi Faltings, Mary Lou Maher, Michael T. Cox, Kenneth D. Forbus, Mark T. Keane, Agnar Aamodt, and Ian D. Watson. Retrieval, reuse, revision and retention in case-based reasoning. *Knowledge Eng. Review*, 20(3):215–240, 2005.
- [Deaner, 2006] Robert O Deaner. More males run fast: A stable sex difference in competitiveness in us distance runners. *Evolution and Human Behavior*, 27(1):63–84, 2006.
- [Foster *et al.*, 1994] Carl Foster, Matthew Schrager, Ann C Snyder, and Nancy N Thompson. Pacing strategy and athletic performance. *Sports Medicine*, 17(2):77–85, 1994.
- [Geleijnse *et al.*, 2011] Gijs Geleijnse, Peggy Nachtigall, Pim van Kaam, and Lucienne Wijgergangs. A personalized recipe advice system to promote healthful choices. In *Proceedings of the 16th International Conference on Intelligent User Interfaces, IUI '11*, pages 437–438, New York, NY, USA, 2011. ACM.
- [Haney Jr, 2010] Thomas A Haney Jr. *Variability of pacing in marathon distance running*. PhD thesis, University of Nevada, Las Vegas, 2010.
- [Hermens *et al.*, 2014] Hermanus J. Hermens, Harm op den Akker, Monique Tabak, J.L.P Wijsman, and Miriam Marie Rosé Vollenbroek-Hutten. Personalized coaching systems to support healthy behavior in people with chronic conditions. 24(6):815–826, 2014.
- [Kelly *et al.*, 2012] Daniel Kelly, Garrett F Coughlan, Brian S Green, and Brian Caulfield. Automatic detection of collisions in elite level rugby union using a wearable sensing device. *Sports Engineering*, 15(2):81–92, 2012.
- [Leijdekkers and Gay, 2015] Peter Leijdekkers and Valerie Gay. *Improving User Engagement by Aggregating and Analysing Health and Fitness Data on a Mobile App*, pages 325–330. Springer International Publishing, Cham, 2015.
- [Lewis, 2004] Michael Lewis. *Moneyball: The art of winning an unfair game*. WW Norton & Company, 2004.
- [Möller *et al.*, 2011] Andreas Möller, Johannes Scherr, Luis Roalter, Stefan Diewald, Nils Hammerla, Thomas Plötz, Patrick Olivier, and Matthias Kranz. Gymskill: Mobile exercise skill assessment to support personal health and fitness. In *9th Intl. Conf. on Pervasive Computing (Pervasive 2011), Video, San Francisco, CA, USA, 2011*.
- [Ohlin and Olsson, 2015] Fredrik Ohlin and Carl Magnus Olsson. Intelligent computing in personal informatics: Key design considerations. In *Proceedings of the 20th International Conference on Intelligent User Interfaces, IUI '15*, pages 263–274, New York, NY, USA, 2015. ACM.
- [Peek *et al.*, 2015] Niels Peek, Carlo Combi, Roque Marin, and Riccardo Bellazzi. Thirty years of artificial intelligence in medicine (aime) conferences: A review of research themes. *Artificial Intelligence in Medicine*, 65(1):61 – 73, 2015.
- [Smyth and Cunningham, 2017a] Barry Smyth and Pádraig Cunningham. 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, 2017*, pages 116–120, 2017.
- [Smyth and Cunningham, 2017b] Barry Smyth and Pádraig Cunningham. Running with cases: A CBR approach to running your best marathon. In *Case-Based Reasoning Research and Development - 25th International Conference, ICCBR 2017, Trondheim, Norway, June 26-28, 2017, Proceedings*, pages 360–374, 2017.
- [Smyth, 2007] Barry Smyth. Case-based recommendation. In *The Adaptive Web, Methods and Strategies of Web Personalization*, pages 342–376, 2007.
- [Trubee, 2011] Nicholas William Trubee. *The Effects of Age, Sex, Heat Stress, and Finish Time on Pacing in the Marathon*. PhD thesis, University of Dayton, 2011.
- [Wiesner and Pfeifer, 2010] Martin Wiesner and Daniel Pfeifer. Adapting recommender systems to the requirements of personal health record systems. In *Proceedings of the 1st ACM International Health Informatics Symposium, IHI '10*, pages 410–414, New York, NY, USA, 2010. ACM.
- [Wiesner and Pfeifer, 2014] Martin Wiesner and Daniel Pfeifer. Health recommender systems: Concepts, requirements, technical basics and challenges. *International Journal of Environmental Research and Public Health*, 11(3):2580–2607, 2014.