Generating Weekly Training Plans in the Style of a Professional Swimming Coach Using Genetic Algorithms and Random Trees

Rikard Eriksson[1†], Johan Nicander[1†], Moa Johansson[1], and C. Mikael Mattsson[2,3,4]

[1] Department of Computer Science & Engineering, Chalmers University of Technology, Gothenburg, Sweden,
rikarderiksson@outlook.com, johan.nicander@hotmail.com,
moa.johansson@chalmers.se
[2] Department of Physiology & Pharmacology, Karolinska Institutet, Stockholm, Sweden
[3] Health and Sports Technology Initiative, Blekinge Institute of Technology, Karlskrona, Sweden
[4] Silicon Valley Exercise Analytics (svexa), Menlo Park, CA, USA,
mikael@svexa.com

Abstract. Optimal training planning is a combination of art and science, a time-consuming task that requires expert knowledge. As such, it is often exclusively available to top tier athletes. Many athletes outside the elite do not have access or cannot afford to hire a professional coach to help them create their training plans. In this study, we investigate if it is possible to use the historical training logs of elite swimmers to construct detailed weekly training plans similar to how a specific professional coach would have planned. We present a software system based on machine learning and genetic algorithms for generation of detailed weekly training plans based on desired volume, intensity, training frequency, and athlete characteristics. The system schedules training sessions from a library extracted from training plans written by a professional swimming coach. Results show that the proposed system is able to generate highly accurate training plans in terms of training load, types of sessions, and structure, compared to the human coach.

Keywords: Swimming, Training Planning, Training Plan Generation, Machine Learning, Exercise Intelligence

1 Introduction

The goal of training in any sport is to improve an athlete's physical, technical, and psychological attributes. At elite levels, it is also important to reach peak performance at the right time, usually a competition. In endurance sports, these goals are achieved through periodization, that is, cycling the training load, and

---

[†]Contributed equally to the article

tapering [2, 11]. Periodized training plans consider multiple sub-plans at different levels of abstraction. The sub-plans range from coarse long term plans for an entire season down to detailed plans specifying what should be done each session of a week.

To create an optimal training planning, sport-specific demands must be taken into consideration, such as the athlete's individual profile, long- and short-term planning, and daily readiness [13, 7]. This is a very time-consuming task and something that is often exclusively available to professional athletes at the highest level.

Athletes and coaches at the highest level often log details of training sessions, resulting in substantial amounts of data. Previous work has both quantified training load and modeled performance based on said load. Examples of these models are the Banister and Busso models [1, 10, 3]. The performance models has since been used to find the optimal training load over an entire season and the most effective patterns of tapering [4, 14, 15, 8]. Kumyaito et al. extends this work to create macro-plans including physiological constraints [9]. They use adaptive particle swarm optimization to find an optimal training plan for road cyclists which takes monotony and chronic training load ramp rate into account. This ensures that the training is both varied and not increased too fast, both of which have may lead to overtraining. The plans were then compared to plans from British cyclists and the results show an increased performance according to the Busso model.

There is little work that focuses on how to create a detailed plan for a week of training, and the existing work has focused on training plans for general fitness rather than sport specific demands. [12] uses numerical planning to create a training plan that covers general fitness for kick-boxers in terms of what muscle groups should be trained. However, training load is not considered. A study by Fister et al. utilizes different stochastic algorithms to create training plans from training logs, according to constraints specified by the authors [6]. Although these works aim to generate training plans that are similar to what a human coach would create, they not try to mimic the style of a specific coach. Further, to the extent of our knowledge, no attempt has been made to automate the generation of training plans for swimmers.

2   Aim

This study aims to create a system that can produce detailed weekly training plans for swimmers in the style, philosophy, and likeness of one specific professional swimming coach. For this task, the system utilizes historical plans from the coach, a specification of the training week to plan, and information about the athlete for which the plan is made. The specification for the week includes information about the number of training sessions to be performed, how much training should be performed in each intensity zone, the number of sessions for each particular area of focus, and the time left of the cycle. The athlete information includes the athlete's age, gender and competition type (specialty swimming

stroke and distance). To be able to compare different models and human coaches we can define a set of objectives that define when two training plans are considered similar. These objectives are used as evaluation metrics for the described system and result of experiments (for the mathematical definitions see [5]):

1. e_D, The training plans have a similar distribution of training load over the sessions.
2. e_Z, The training plans have a similar distribution of training load over the intensity zones.
3. e_T, The training plans contain similar orderings of session types.
4. e_S, The training plans have the same number of sessions of each type.
5. e_F, The training plans have the same number of sessions suited for each specialty.

3   Data

A library of 5313 historical training sessions and self-reported de-identified training logs from 435 training weeks for athletes trained by a professional coach was supplied by the company svexa. The data set contains information about the volume swum in different intensity zones, a detailed description of the session performed and additional meta-information on the type of the main set of the session as well as what competition type the session is suited for.

The data set was cleaned and pre-processed which resulted in a final library of 2702 training sessions after duplicate or malformed sessions were removed.

Aggregated weekly data was split into two parts: The first part, *weekly input*, consists of the data about the athlete (age, gender, specialty), goals (weeks left of macrocycle, focus area, weekly training load per intensity zone), and schedule constraints (number of sessions per week). The second part, the *detailed plan*, is the target for the model predictions and consists of a sequence of training sessions performed during the given week. Lastly, the data was split into an 80-20 train-test split based on randomly sampled calendar weeks.

4   Implementation

Our system, the Genetic and Random Trees training planner (GERT), consists of two modules: a machine learning system trained to infer a weekly load distribution and order of sessions, and a planning system that constructs a detailed plan of individual sessions selected from the session library implemented as a genetic algorithm [16]. An visual overview of GERT is shown in Figure 1. For the full implementation details and complete technical description see [5].

4.1   Learning the coaches' style

The first action of the GERT system is to extract from historical data how a coach would create a training plan given a weekly input. This is done in two
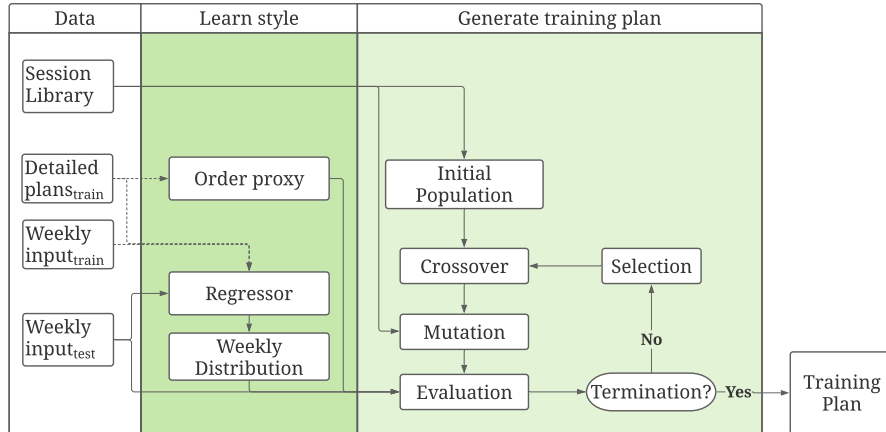
Fig. 1: Overview of GERT. In the left column (Data) all inputs to GERT are listed. In the middle column all learning stages are listed. In the right column stages used during the Generating training plans portion of execution are listed.

steps: first learning how to distribute the overall training load for the week over the sessions, and secondly how the desired load should be achieved for each session (e.g. duration and intensity). These two steps, in combination with the details of how each session should be performed, results in a plan that is typical for the coach in question. We refer to this as the coach's style. In this particular study we used data from one coach, and the process can be repeated to include several different coaches with different styles and training philosophies.

First, the system learns how the coach would plan the training load of each session during the week. This is done by fitting an internal regression model to the distribution of training load over the sessions in the historical training plans.

Secondly, GERT will learn what type of session each workout should be. In this work, the type of a sessions is distinguished by its dominant intensity zone. For example low intensity session have the majority of the workload in intensity zone I-II. This is done by extracting series of three consecutive session types from the historical plans, e.g. {I-II, VII, rest}. The collected series are then used as a proxy for good session orderings during the generation stage of the system.

For our system, the details of how each session is performed is solved by the use of a session library.

4.2   Generating training plans

The goal of this stage is to generate a training plan based on sessions from the session library such that the objectives specified in Aim are fulfilled. We implement this using a genetic algorithm, whereby a population of weekly training plans are constructed by randomly sampling sessions from the session library.

The population is iteratively improved over a number of generations. At each generation, the population is altered through crossover and mutation. The samples are then evaluated based on a loss function that is a weighted combination of the five metrics defined in section 2.

After the evaluation, a termination criteria is checked to see if the population is still improving. If so, the best samples are kept and the process starts over from the crossover stage. Otherwise the best scoring sample is return as the suggested training plan.

5   Results

The model was tested with the goal of producing a weekly training plan similar to that of the human coach. It was compared to a baseline: simply picking the closest matching existing weekly training plan from the training set using a nearest neighbor approach (BKNN). Here closeness of training plans is calculated using the metrics in section 2.

The results from the experiments are summarised in Table 1. The experiments show that GERT outperforms the baseline model on all scores.

Table 1: The mean value for the terms of the loss function and the loss function itself from the nearest neighbour baseline model (BKNN) and GERT. For all scores, lower is better and the lowest score possible is 0. The uncertainties in the results, reported in parenthesis, are twice the estimator for the standard error for each score.

model	e_D	e_Z	e_T	e_S	e_F	$\mathcal{L}$
BKNN	4.7(3)	2.8(3)	6.2(5)	0.9(2)	1.5(2)	21(2)
GERT	3.8(3)	0.32(6)	5.1(5)	0(0)	0(0)	6.6(5)

In Figure 2, the total training load decided by the baseline model and GERT respectively are plotted against the total training load decided by the coach for all training plans in the test set. GERT clearly outperforms the Baseline, doing a near-perfect job mimicking the coach's total training load.

Figure 3 shows a (slightly simplified) example of the output from GERT. The training plan generated by GERT follows a similar structure to that of a human coach in terms of training load distribution over the week, session types and intensity zone distribution. The two training plans only differ slightly in order of the session types, as can be seen in the morning sessions of the later sessions of the week.

6   Discussion

The inferred training load distribution is the source of much of the error, as this information is not given in the input. However, the load of the sessions in the
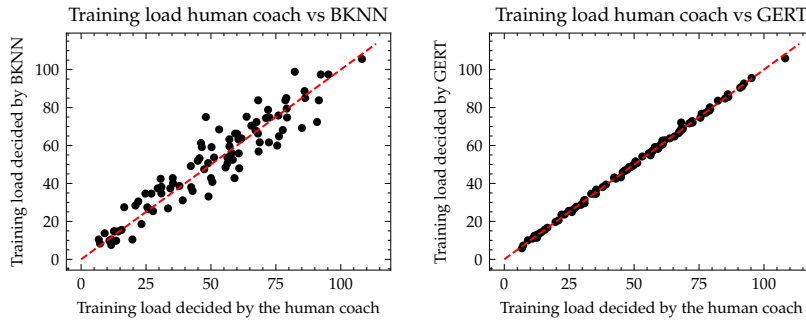
(a) Training load of a weekly program set by the human coach versus that training program generated by the Baseline (BKNN) for all data points in the held-out test set. $R^2 = 0.88$

(b) Training load of a weekly program set by the human coach versus that training program generated by GERT for all data points in the held-out test set. $R^2 = 1.00$

Fig. 2: Comparison of how well the two models Baseline (BKNN) and GERT can mimic the total training load set by the human coach.

generated training plan will almost always differ from the load inferred in the first step of the algorithm, this can be seen in Figure 3a. There are three main reasons for this: First, this could be a result of early stopping of the genetic algorithm. Secondly, it could be due to interplay between the different metrics in the loss function. For example, if the training plan from the human coach contains orders of session types that do not exist in the training set, GERT would prefer a different order. The third reason is that there might not exist a perfect solution.

During this work, we made the assumption that the training in the data set reported by the athletes is the same as the training the coach planned for. The effects of this assumption are hard to investigate without a complete data set of the coach's original planning, although it is reasonable to believe that most of the training happened according to plan.

We also make the simplifying assumption that the weekly plans are independent of each other given the macro- and meso-plans. This allows the creation of the weekly plans on a week-by-week basis, without the additional complexity of creating dependent plans.

7   Conclusion

The GERT system produced training plans close to perfectly mimicking the professional coach. We believe that our system can be easily adapted to produce training plans for other coaches, athletes, and endurance sports, by simply feeding it new data. Since our current data set only contain training sessions from a

Training load by session

Training load by intensity zone
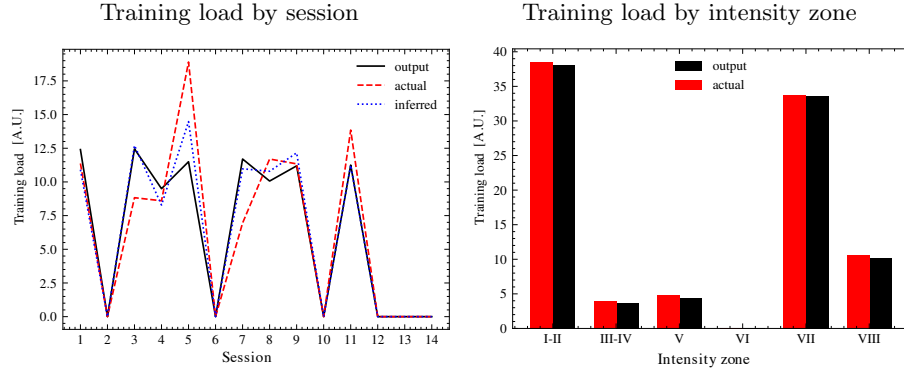
(a) The distribution of training load over the sessions of the week for the plan that was generated by GERT (output), the distribution that was predicted by GERT's regressor (inferred), and the distribution of the week written by the human coach (actual).

(b) The distribution of training load over the intensity zones of the week for the plan that was generated by GERT (output) and the human coach (actual).

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
VII	I-II	V	VII	I-II	VII	
VII	I-II	VII	I-II	VII	V	
	I-II		I-II			
	I-II		I-II			

☐ GERT
☐ Human Coach

(c) The orderings of the sessions, with session type illustrated by dominant intensity zones, for GERT and the human coach. Differences between the two plans are highlighted with yellow.

I-II					
2x					
5x100 skpks					
8x50 kick @55					
15x100 fs i1 @1:25					
8x25 fly 15/15 @30					
6x200 fins bk @2:35					
3x					
250 pull fs tech					
3x50 work good turns					
3x100 swim down					

I-II	III-IV	V	VI	VII	VIII
6.2	0	0	0	0	0

(d) Example of a session type with dominant intensity zone I-II

Fig. 3: Example of a training plan generated by GERT.

single coach, athlete group and sport, we have only been able to investigate this specific setting. However, we have no reason to believe it would not be generalizable to new data. We suggest further investigation of other applications of the method presented in this work to see what generalizations are possible. Further, by combining GERT with previous works that optimize the training load for a macro-plan, such as the model by Kumyaito [9], it would be possible to create a pipeline that generates complete, detailed individualized training plans.

With GERT, we have taken the first step towards automating individualized training plans in the style of expert coaches and making top-tier, individualized training plans available to a broader audience.

References

1. Banister, E.W.: Modeling elite athletic performance. Physiological testing of elite athletes **347**, 403–422 (1991)
2. Bompa, T., Jones, D.: Theory and Methodology of Training: The Key to Athletic Performance. Kendall/Hunt Publishing Company (1983)
3. Busso, T., Denis, C., Bonnefoy, R., Geyssant, A., Lacour, J.R.: Modeling of adaptations to physical training by using a recursive least squares algorithm. Journal of applied physiology (1997)
4. Busso, T., Thomas, L.: Using mathematical modeling in training planning. International journal of sports physiology and performance **1**(4), 400–405 (2006)
5. Eriksson, R., Nicander, J.: Automated generation of training programs for swimmers. Master's thesis, Chalmers University of Technology (2021)
6. Fister, I., Fister Jr, I., Fister, D.: Generating training plans based on existing sports activities. In: Computational Intelligence in Sports, pp. 139–180. Springer (2019)
7. Issurin, V.B.: New horizons for the methodology and physiology of training periodization. Sports medicine **40**(3), 189–206 (2010)
8. Kumyaito, N., Tamee, K.: Intelligence planning for aerobic training using a genetic algorithm. In: International Symposium on Natural Language Processing, pp. 196–207. Springer (2016)
9. Kumyaito, N., Yupapin, P., Tamee, K.: Planning a sports training program using adaptive particle swarm optimization with emphasis on physiological constraints. BMC research notes **11**(1), 1–6 (2018)
10. Mujika, I., Busso, T., Lacoste, L., Barale, F., Geyssant, A., Chatard, J.C.: Modeled responses to training and taper in competitive swimmers. Medicine and science in sports and exercise **28**(2), 251–258 (1996)
11. Mujika, I., Chatard, J.C., Busso, T., Geyssant, A., Barale, F., Lacoste, L.: Effects of training on performance in competitive swimming. Canadian Journal of Applied Physiology **20**(4), 395–406 (1995)
12. Skerik, T., Chrpa, L., Faber, W., Vallati, M.: Automated training plan generation for athletes. In: 2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 3865–3870. IEEE (2018)
13. Smith, D.J.: A framework for understanding the training process leading to elite performance. Sports medicine **33**(15), 1103–1126 (2003)
14. Thomas, L., Mujika, I., Busso, T.: A model study of optimal training reduction during pre-event taper in elite swimmers. Journal of sports sciences **26**(6), 643–652 (2008)
15. Thomas, L., Mujika, I., Busso, T.: Computer simulations assessing the potential performance benefit of a final increase in training during pre-event taper. The Journal of Strength & Conditioning Research **23**(6), 1729–1736 (2009)
16. Whitley, D.: A genetic algorithm tutorial. Statistics and computing **4**(2), 65–85 (1994)