The Fitbit generation

From couch potato to marathon runner

Kaylea Haynes, a PhD student in statistics at Lancaster University, applies changepoint analysis to Fitbit data to optimise her exercise regime

I used to be a bit of a couch potato. From time to time I would flirt with the idea of going for a run, but I never went further than once round the block before heading home. In December 2014 I decided it was time for a change, and to motivate myself I set a goal of running a 10 km race by April 2015. Skip to the present day, and I am proud to say I have run five 10 km races, three half marathons and a marathon.

However, if I am honest, I have slacked off my training since running the Edinburgh marathon in May. I still try to run several times a week, swim twice a week, and cycle to university now and again, but I also spend a good portion of my day sitting at my desk. I know this because I have started to collect and analyse my heart rate data – and, using this information, I plan to reinvigorate my training regime.

Heart rate is a good indicator of how hard you are exercising, and to monitor this I use the Fitbit Charge HR, an activity-tracking wristband. There are different heart-rate zones that you can train in, and these are usually defined as "peak", "cardio" and "fat burn". Different zones enhance different aspects of your fitness. For example, a sprinter would aim to train more of the time in the "peak" zone, which improves performance and speed, whereas an endurance athlete would find it more beneficial to be working in the "cardio" zone, which improves aerobic fitness.

The heart rate training zones are typically given as a percentage of the maximum heart rate, which is approximated by:

Max Heart Rate = 220 – age[1]

Various sources give different percentage boundaries for each of the zones, but Table 1 lists the breakdowns I use as a guideline.

Table 1. A breakdown of the different heart-rate training zones

Peak 85%	High-intensity exercise (anaerobic) zone, which improves performance and speed.
Cardio 70–84%	Medium to high intensity zone, which improves aerobic fitness.
Fat burn 50–69%	Low to medium intensity zone, where a higher percentage of calories are burned from fat.
Out of zone < 50%	Heart rate is not high enough to be considered exercise.

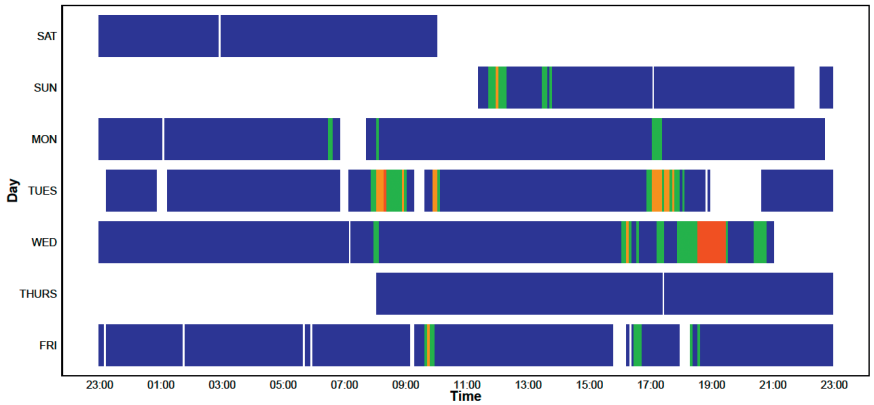

Figure 1. Heart-rate zones recorded over one week (white, no heart-rate recorded; blue, out of zone; green, fat burn zone; orange, cardio zone; red, peak zone)

In the zones

Figure 1 shows my heart-rate zones for a week, with red areas denoting "peak", orange "cardio", green "fat burn" and blue "out of zone" (i.e. I was not exercising).

The white areas are the times where no heart rate was recorded. This could be for a variety of reasons: perhaps I was swimming or having a shower, my activity tracker was out of charge, or the heart-rate sensors on the device were not positioned properly on my wrist.

For the week shown in Figure 1 you can see that I was very active on Tuesday; I did a track session in the morning and then went to a circuit training fitness class in the evening. On Wednesday I spent an hour or so in the red zone, which corresponds to me competing in a 10 km race. For the rest of the week, however, I was mostly out of zone.

Of course, if I want to improve my exercise regime I need to know more than just my heart-rate zones at various points in the day. I need to understand the factors – such as speed or distance covered during exercise – that lead to changes in heart rate.

To do this I use an area of statistics called changepoint analysis. A changepoint is essentially a point in a time series where we experience a change in the statistical properties: a change in the mean, for example. To detect such changes in my heart rate, I use the non-parametric method proposed by Zou *et al.* since I do not know the underlying distribution of the data (many changepoint detection methods are parametric and thus assume that you know the distribution).[2]

Figure 2 shows my daily heart rate for the Tuesday of the week illustrated in Figure 1. (I used the fitbitScraper R package to get my data from the Fitbit website into R.[3]) The dashed lines are the detected changepoints. The times when I was most active are highlighted by the grey blocks. Those in pink are the times when my activity tracker failed to record a heart rate.

Interval training: an example

Having run a marathon in May, I have decided to now focus on improving my speed. One way to do this is through interval training, which involves alternating between short, high-intensity runs and recovery periods. I recorded my heart rate during one such session – a "pyramid session" – in which I ran 100 m, 200 m, 400 m and 800 m, and then 400 m, 200 m and 100 m.

For this exercise, I swapped my Fitbit for a Garmin Forerunner GPS watch with a heart-rate monitor chest strap, so I could also record distance and speed as well as heart rate. The plot in Figure 3 shows how my heart rate

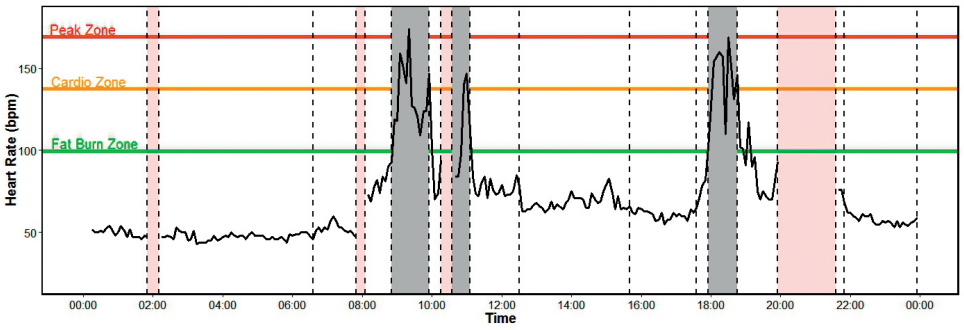

Figure 2. Detecting changes in daily heart rate (the grey blocks are the times where the author was most active and the pink blocks are times where no heart rate was recorded)

changed with distance. I have coloured the plot to show which of the heart-rate zones I was in. The dashed lines are the changes detected using changepoint analysis. These changes roughly correspond to the times where I go from a recovery interval to a high-intensity interval or vice versa, albeit with a bit of a lag.

At the start of each high-intensity interval my heart rate increased fairly steadily, before evening out over the period. At points where I came to a stop, my heart rate dropped suddenly – except for the time after the first 200 m run, where I continued to walk during the recovery period. In each of the running intervals I was working within the red zone. This indicates that I was pushing myself. Had I stayed in the orange zone during the intervals, this would have been a sign that I had the capacity to work harder.

Quantifying myself

Collecting and analysing my own data in this way has given me useful insights into how active I am currently and how I can improve my fitness. Of course, the data alone does not tell the whole story. My Fitbit or Garmin might be good at tracking my heart rate, but they cannot say anything about the activities I engage in, or the surreptitious slice of cake I have with my morning coffee.

The data points the way, but taking the time to think about the data in the context of everyday life – aided by an activity diary – has helped me formulate an appropriate exercise plan.

By training my body to use the anaerobic energy system (as indicated by the red heart-rate zone), I hope to be able to run faster and to keep my speed up for longer. And by monitoring my activity levels on a weekly basis, I intend to make sure I do not get sucked back into the life of a couch potato.

References

1. BrianMac Sports Coach (2015) Heart rate training zones. bit.ly/1vgNtER
2. Zou, C., Yin, G., Feng, L. and Wang, Z. (2014) Nonparametric maximum likelihood approach to multiple change-point problems. *Annals of Statistics*, **42**(3), 970–1002.
3. Nissen, C. (2015) fitbitScraper: Scrapes Data from "www.fitbit.com". R package version 0.1.4. bit.ly/1HJMHr7

Kaylea Haynes is a final year PhD student at the STOR-i centre for doctoral training at Lancaster University. You can follow her training blog at kayleasvitalstats.wordpress.com

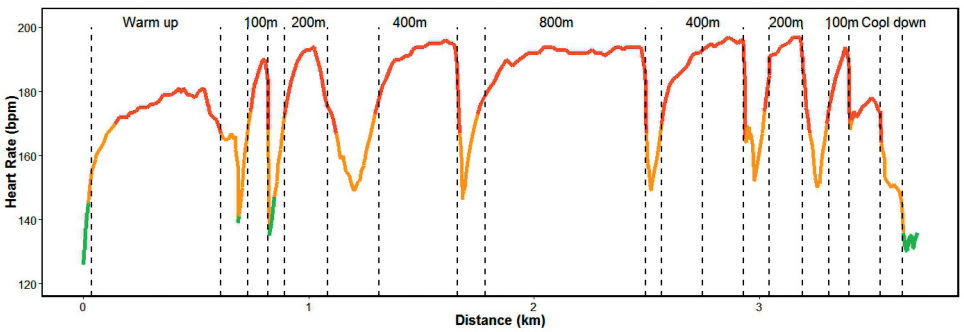

Figure 3. Recording of heart rate during an interval training session. Colours correspond to the heart-rate zones in Table 1