estimating CdA with a power meter

R. Chung
rechung@gmail.com

with good data, field testing works

- Validated power model

  Martin, et al. (1998), “Validation of a mathematical model for road cycling power”, J App Biomech 14(3)

- Estimating drag area with good data collected in field

  Martin, et al. (2006a), “Modeling sprint cycling using field-derived parameters and forward integration”, MSSE 38(3):592-597

  Martin, et al. (2006b), “Aerodynamic drag area of cyclists determined with field-based measures”, Sportscience 10: 68-9

  Snyder, J.; and T. Schmidt (2004), “Determination of drag parameters utilizing a bicycle power meter”, HPeJ issue 1

what happens if the data aren't good?

the classic approach

- constant speed runs on flat windless roads
  - some alternatives: coast down tests, velodrome runs
  - often, results averaged over runs taken in opposite directions
  - occasionally, a few other adjustments and variations
- for constant speed on flat windless roads, power-drag equation simplifies to
  - $\text{watts} = k_0 v + k_1 v^3$, or $\text{watts}/v = k_0 + k_1 v^2$
- regress drag force (i.e., watts/v) on v^2
  - the regression intercept (k_0) is related to Crr
  - the regression slope (k_1) is related to CdA

the challenge

- I recorded power and speed during a ride consisting of a number of laps around a closed course
  - power was not constant
  - speed was not constant
  - the course was not flat
  - the wind was blowing weakly but (I believe) consistently and from the same direction during the entire ride
- how good of an estimate of CdA is it possible to get with these (lousy) data?
  - using usual approach, not good at all
  - using approach described here, not bad at all
  - with non-lousy data, you can get very good results

the data

data were collected at 1.26-second intervals with a Power Tap hub. The plot shows that neither speed nor power were constant

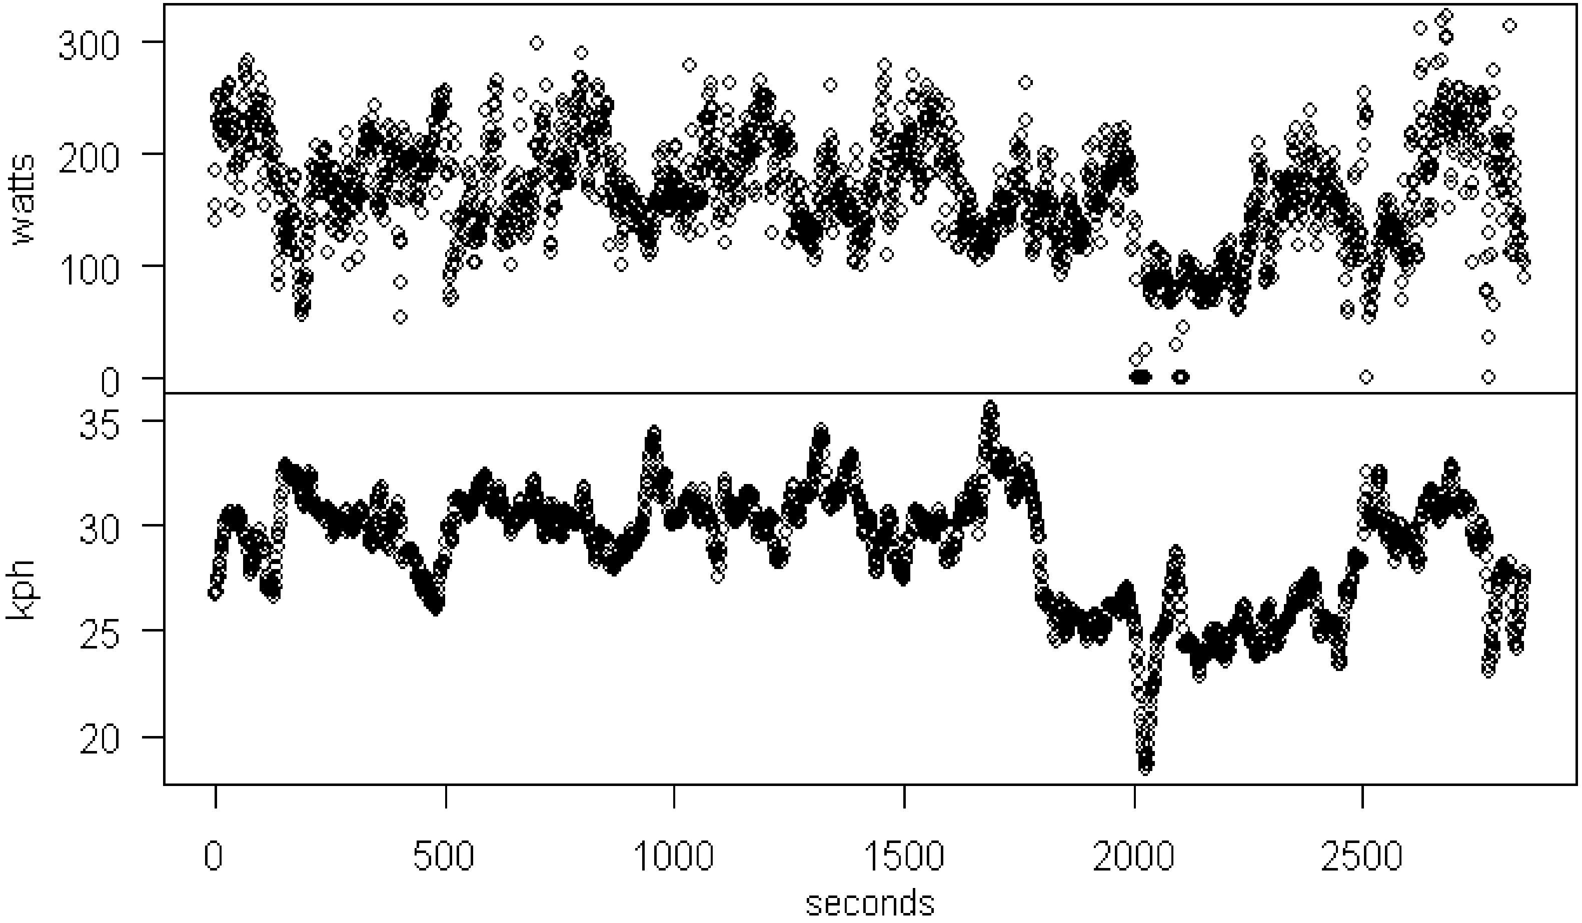

the problem

- flat, windless venues are hard to find (some have tried airplane hangars and building hallways)
- the regression approach is not robust to changes in speed or conditions

    using a 15 minute subset of the data produces a highly statistically significant regression slope that is *negative*, meaning negative CdA

    usual methods don't work well with these data at all

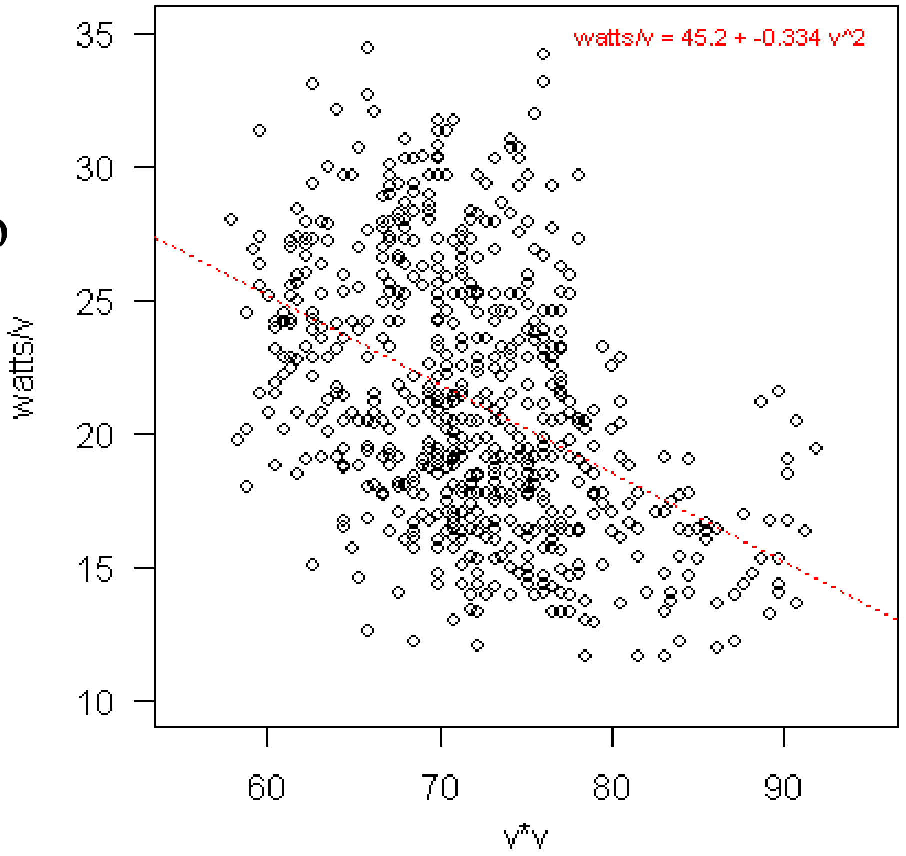

a different approach

- record speed and power from a series of laps on the same route
  - route need not be flat
  - speed and power need not be constant
  - hold position and don't use brakes
  - wind should be as close to zero as possible
- construct an elevation profile for the ride as a function of known power, speed, mass, and air density, and initial guesses at CdA and Crr. Plot the elevation profile against distance
- since each lap must start and end at the same place, find the value of CdA that produces zero net elevation gain over each lap (this means the estimated CdA applies over a lap). One (but not the only) way to do this is to try different values until the laps "line up"

Q: did we correctly identify laps?

this approach provides a self-check: it should identify the correct number of laps. On these data, we show seven-and-a-half laps, with about 10 meters of elevation change per lap. Was that right?

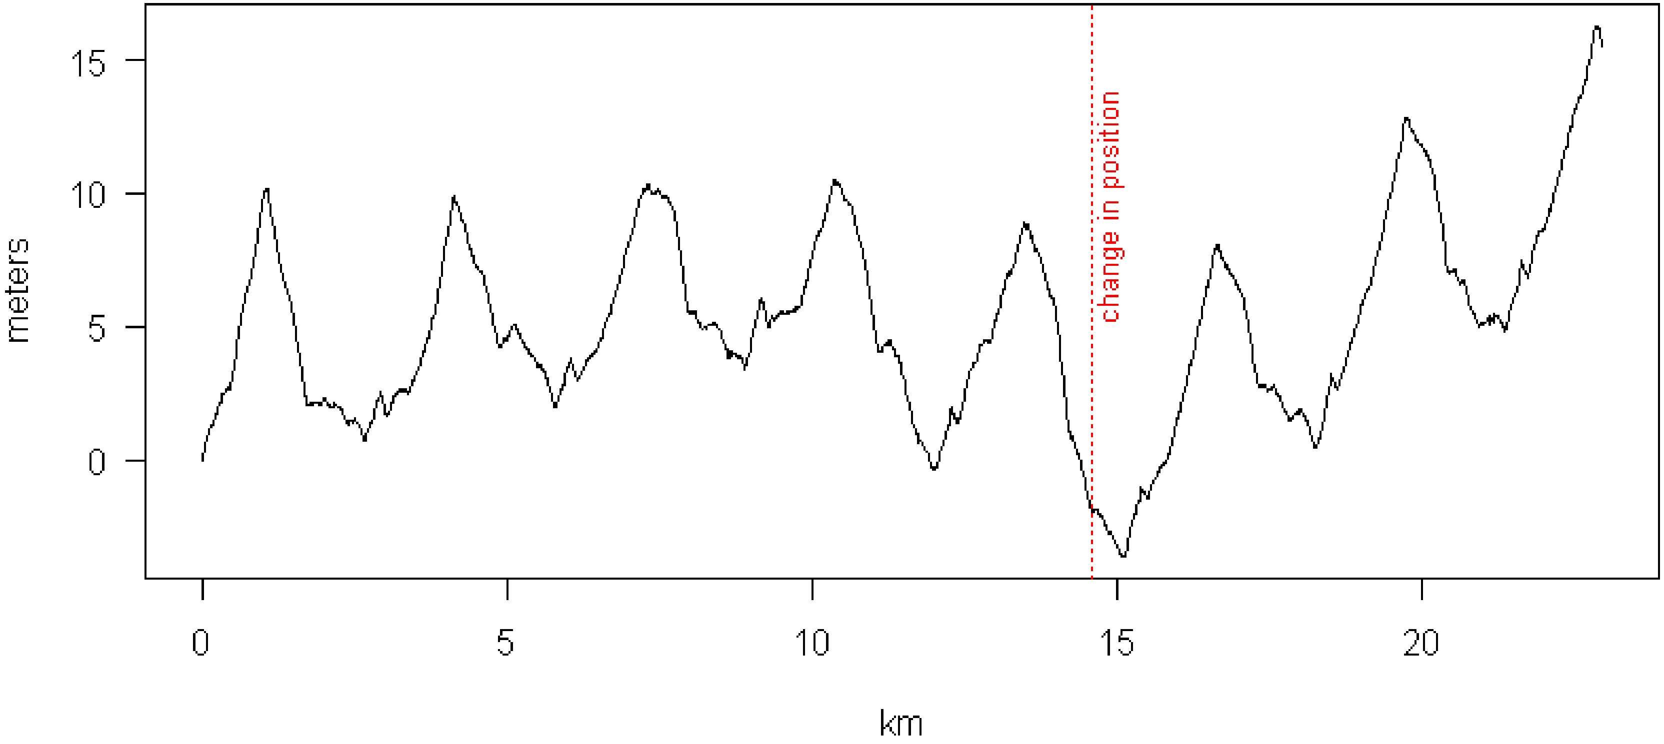

A: pretty much, yes

using *only* speed and power, we identify key features of the ride

- correct number of laps? ***yes***
- correct lap length? ***yes*** (3.12km)
- entered on one side of course and exited on other? ***yes***
- entered at “bottom” of course and exited at “top”? ***yes***
- 10 meter elevation change over each lap? ***close*** – I believe it's closer to 15 or 16 meters
- correctly identified high and low points within laps? ***yes***
- shows conditions were not constant (i.e., change of position during last two-and-a-half laps)? ***yes***
- brief use of brakes on third lap? ***yes***

lap lengths are well identified

- different guesses about CdA (or Crr or mass or air density) have only a small effect on the estimate of lap lengths

  changes in the parameters move the curve up-and-down but not left-and-right

- small changes in wind don't affect lap length much so lap lengths are relatively robust

  however, big changes in wind may

- relatively robust identification of the lap lengths means that it's feasible (though not always wise) to impose the “zero net elevation gain” constraint on each lap
- if you're interested, a Google map of the course is here: http://tinyurl.com/yq9r76

why plot?

- I could have done (and will later show) this algebraically
- however, for now it's easier and perhaps more instructive to plot graphs and show what's happening
    - algebraic solutions generally look for a parameter that minimizes some overall measure of fit
    - in this case, there's more pedagogic value in showing specific areas of fit and misfit rather than overall fit
    - the graphical approach makes it easy to find lap length and knowing lap length will be useful
    - *perhaps most importantly*, the plots give us a generalizable way to diagnose lousy estimates

so what was the CdA?

- hold your horses. First we have to check the assumptions and calculations. To do that, you need to know how to do them.
- we'll start from the beginning, with the power-drag equation, and split the analysis into two parts:
  - assuming no wind
  - assuming some wind, but wind which is consistent in speed and direction

a (simplified) power equation

w = watts needed to propel bike at speed v

= watts to account for rolling resistance +
watts to account for change in elevation +
watts to account for changes in speed +
watts to account for air resistance

$$= w_{rr} + w_{PE} + w_{KE} + w_{aero}$$

simplified power equation, continued

$$w = w_{rr} + w_{PE} + w_{KE} + w_{aero}$$

$$= Crr\ m\ v\ g + s\ m\ v\ g + a\ m\ v + CdA\ \rho\ v_{air}^2 v / 2$$

where

v = speed in m/s (i.e., “ground” speed)

m = total mass (kg) of rider + bike

$g = 9.81\ m/sec^2$

Crr = coefficient of rolling resistance

s = slope

a = acceleration

ρ = air density

v_{air} = “air” speed of bike

CdA = drag area

no wind approach

- assume $v_{air} = v$ and solve for slope as a function of other variables

  $s = w/(m\,g\,v) - Crr - a/g - (\rho\, CdA\, v^2)/(2\,m\,g)$

- use this formula to estimate point-by-point slopes from the data, supplemented by initial guesses at Crr and CdA. Ballpark guesses for starting values might be Crr = .005 and CdA = 0.3. v is in meters per second, so convert v = kph/3.6. A simple and not too terrible estimator for the acclerations, a, is the changes in v divided by 1.26 (these data were collected at 1.26-second intervals)

  standard approach assumes accelerations = 0. This approach calculates and uses them

- use estimated slopes to construct elevation change for each 1.26-second interval: elev.change = s*v*1.26

produce an “elevation” profile

- cumulate the “elevation” changes and plot against distance to produce a “virtual elevation” profile
- make guesses at CdA until the plotted laps line up (or solve algebraically for the CdA that achieves that)

  these steps may seem daunting but they only take a few commands in any spreadsheet, or programming environment like R or Matlab

estimated CdA should level the profile

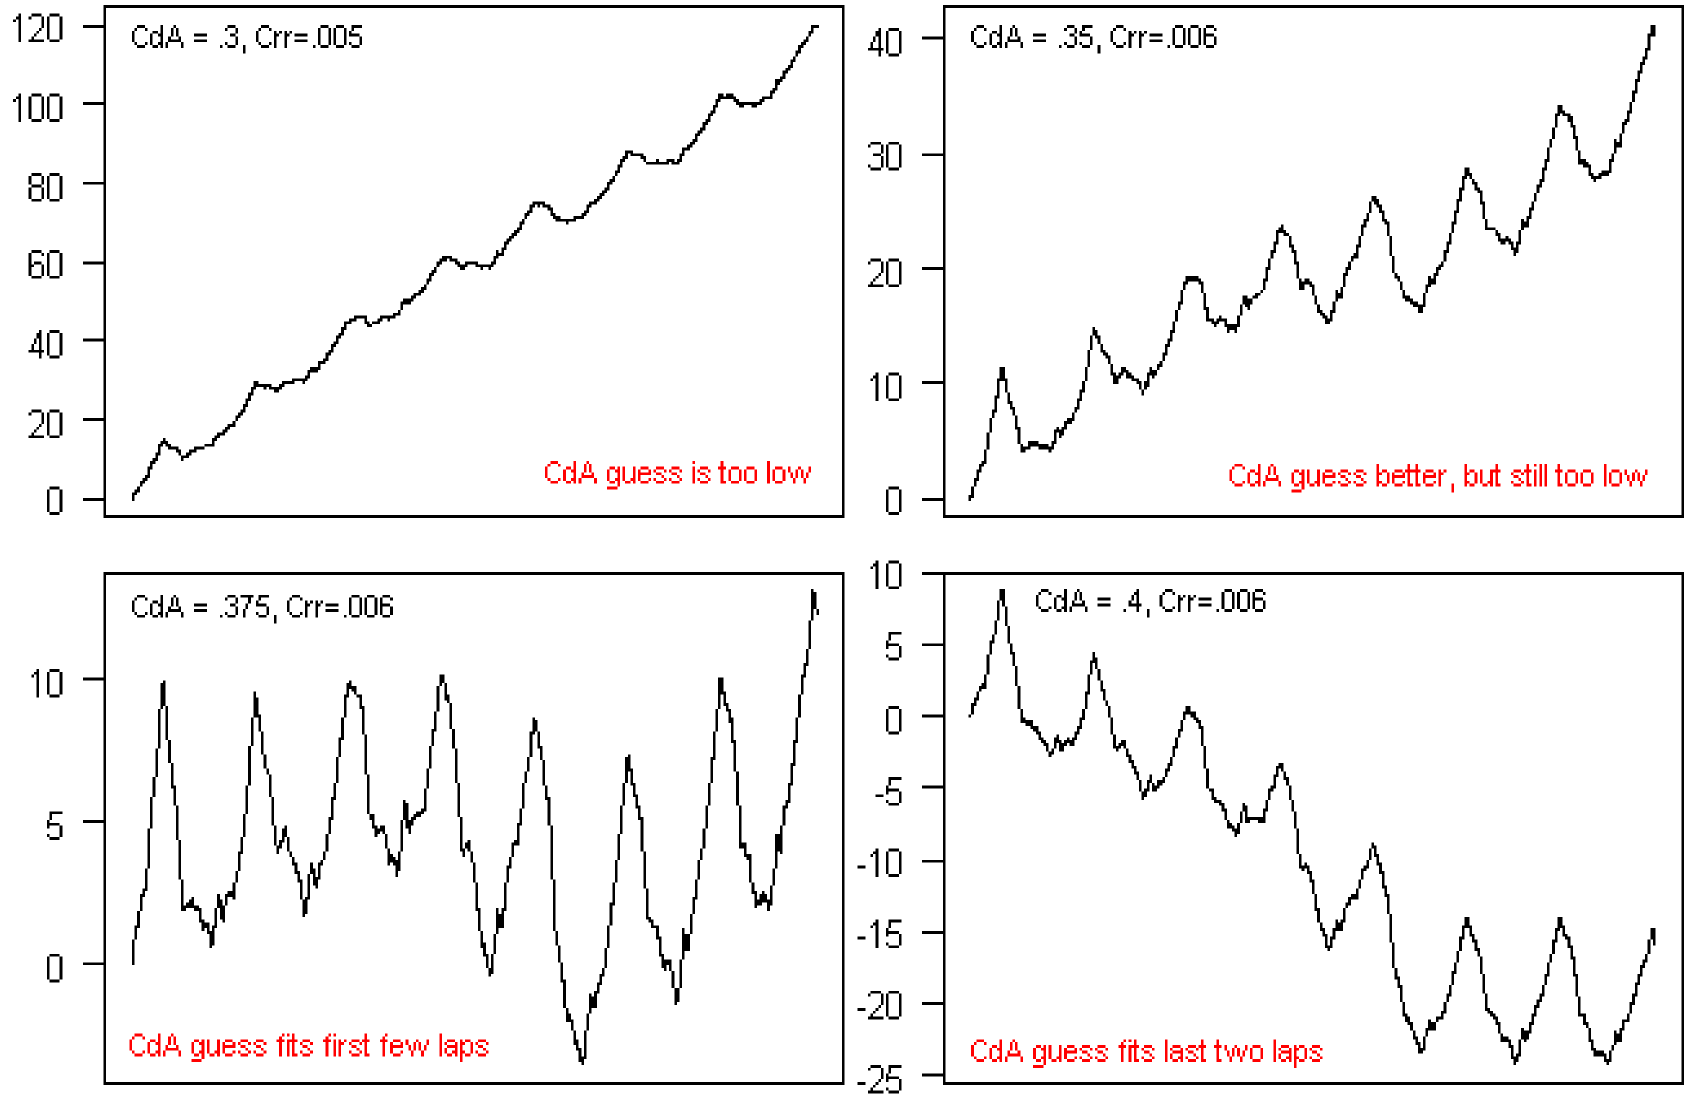

CdA = .3, Crr=.005

CdA guess is too low

CdA = .35, Crr=.006

CdA guess better, but still too low

CdA = .375, Crr=.006

CdA guess fits first few laps

CdA = .4, Crr=.006

CdA guess fits last two laps

why does this work?

classic approach regresses average drag force on (avg v)^2 and minimizes sum of squared errors. Instead, we minimize the sum of a more complex form of the error: it integrates the point-by-point elevation changes across distance, then imposes the constraint that the elevation gain across laps must net to zero

> laps are extra information that the classic approach ignores. In addition, the data were sequential; sequencing the data means accelerations can be calculated and included instead of assuming they are zero. There are other ways to produce a “solution system” but the elevation profile is a convenient way to maintain the sequence of the data and to allow for the additional constraint on fit.

when doesn't this work?

- this method models point-by-point power as a function of point-by-point speed and point-by-point changes in speed but *everything else* gets tossed into the slope term. That's why what we get is a “virtual” elevation profile
- if 1) there are errors in measurement, or 2) the unmodeled parts of the power equation (like wind or brake usage) are large relative to the modeled parts, or 3) CdA changes because you didn't hold your position, then the virtual elevation profile will differ from the true elevation profile

we'll see more discussion of this when we talk about wind

absolute and relative CdA

- in this example I made guesses about both CdA and Crr. Look at the equation—an increase of .001 in Crr looks like an increase in the slope of .001 (=0.1%)
- so *with these data*, we appear to have pretty good relative accuracy but unless we know what Crr is, we won't have good absolute accuracy

  good relative accuracy means we can spot small *changes* in CdA even if (with these data) we can't nail down CdA itself. Sometimes you'll want do specific additional tests that will let you nail down both CdA and Crr
- I was bad and didn't measure air density (though I have a ballpark idea about what it was)

so what was CdA?

since there were so many things I didn't record, the best we can do *with these data* is to calculate CdA assuming different values of Crr and air density. We get:

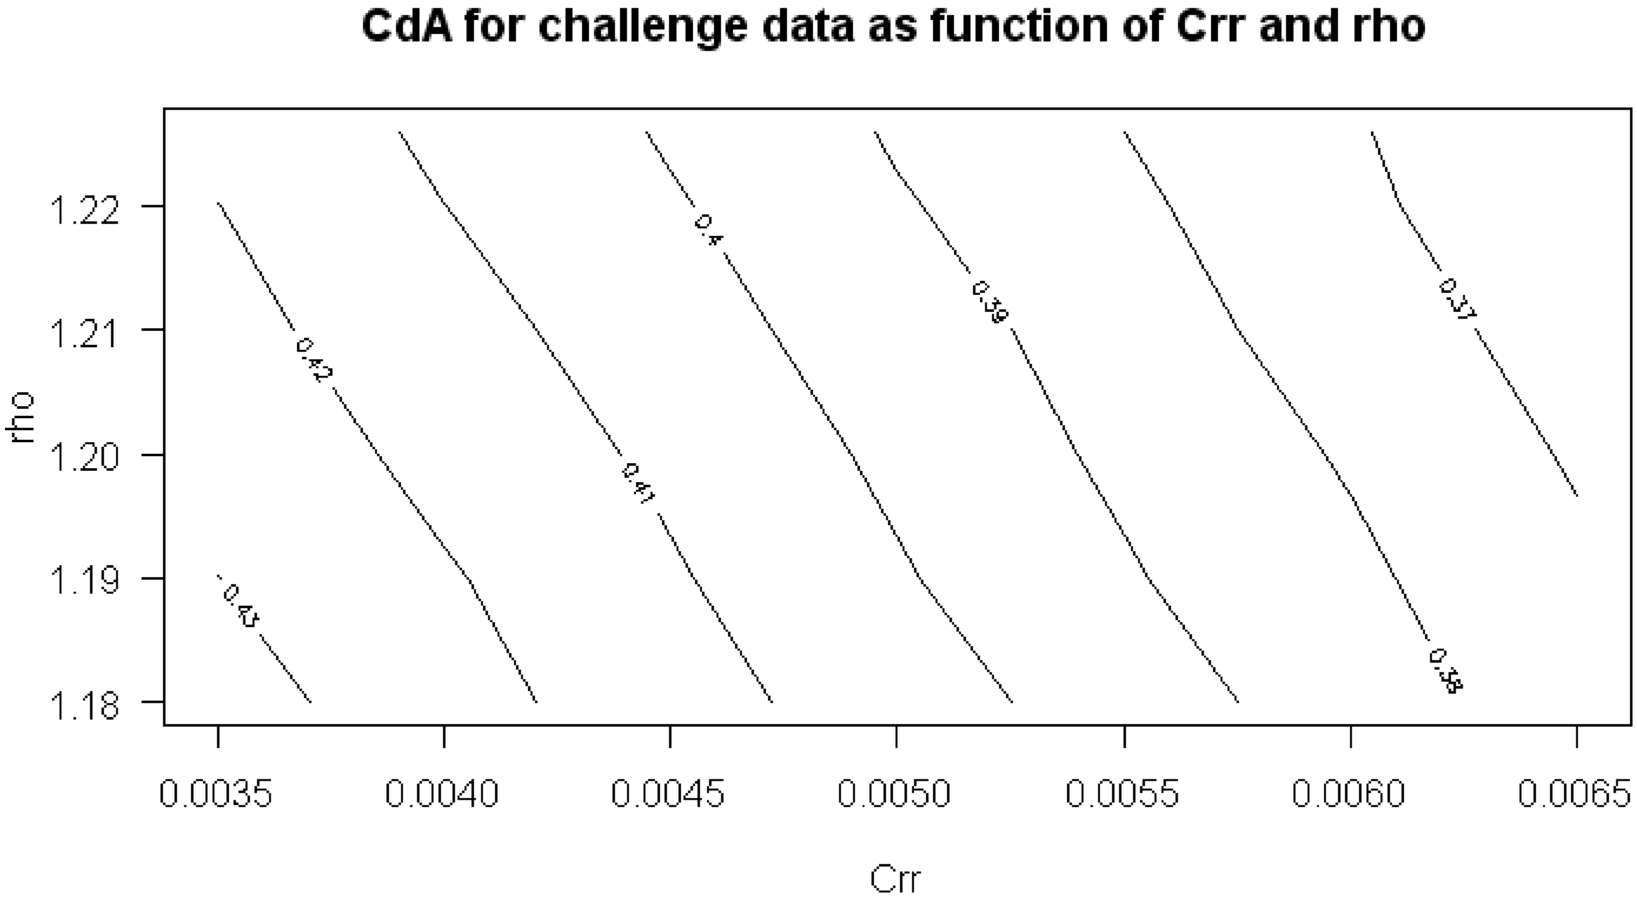
CdA for challenge data as function of Crr and rho

what about lap amplitude?

you may have noticed that the “amplitude” of the estimated elevation differed across laps. Could it be related to speed?

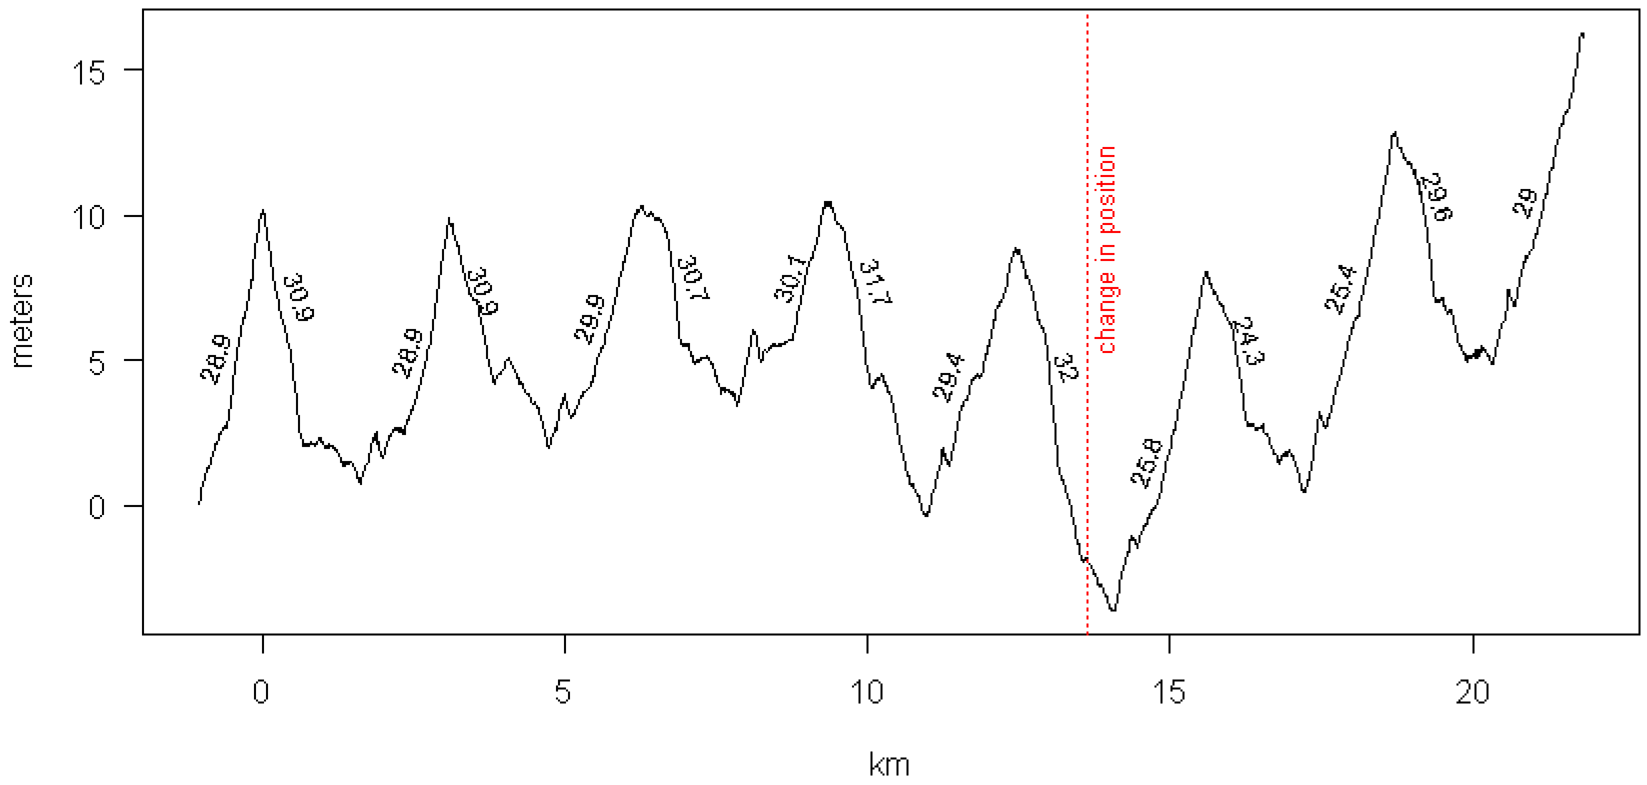

kph and estimated elevation profile

ground speed and air speed

- there is rough evidence that for these data the elevation profiles are speed-dependent
  - increased speed in the downhill direction increased elevation change
  - increased speed in uphill direction decreased elevation change
- could it be unmeasured wind?
  - up to this point, we've assumed no wind (i.e., ground speed = air speed)
    - recall that the challenge included the information that there was an unmeasured amount of wind, but that I thought it was light and from a consistent direction

a handy diagnostic

now you can see why I start with a graphical approach: it provides a handy diagnostic for whether the model assumptions are met

- unmeasured variables affect the profiles in recognizable ways
  - unmeasured wind typically makes the profiles speed dependent
  - unmeasured braking typically appears as a sudden jump in the estimated elevation

what about the wind?

is it possible to say anything about the wind from the data we have?

- we'll try adding a (small) non-zero tailwind for the downhill direction and an equivalent headwind for the uphill; then switch
    - note that this is only a rough correction: the actual course was not a straight out-and-back so adding a small amount of tailwind and headwind is a simplification – the actual course was closer to a right triangle

downhill tailwind, 1 m/s

here's a new estimated profile, assuming a consistent 1 m/s tailwind in the downhill segment and a 1 m/s headwind in the uphill segment. Notice that the laps don't have the same shape

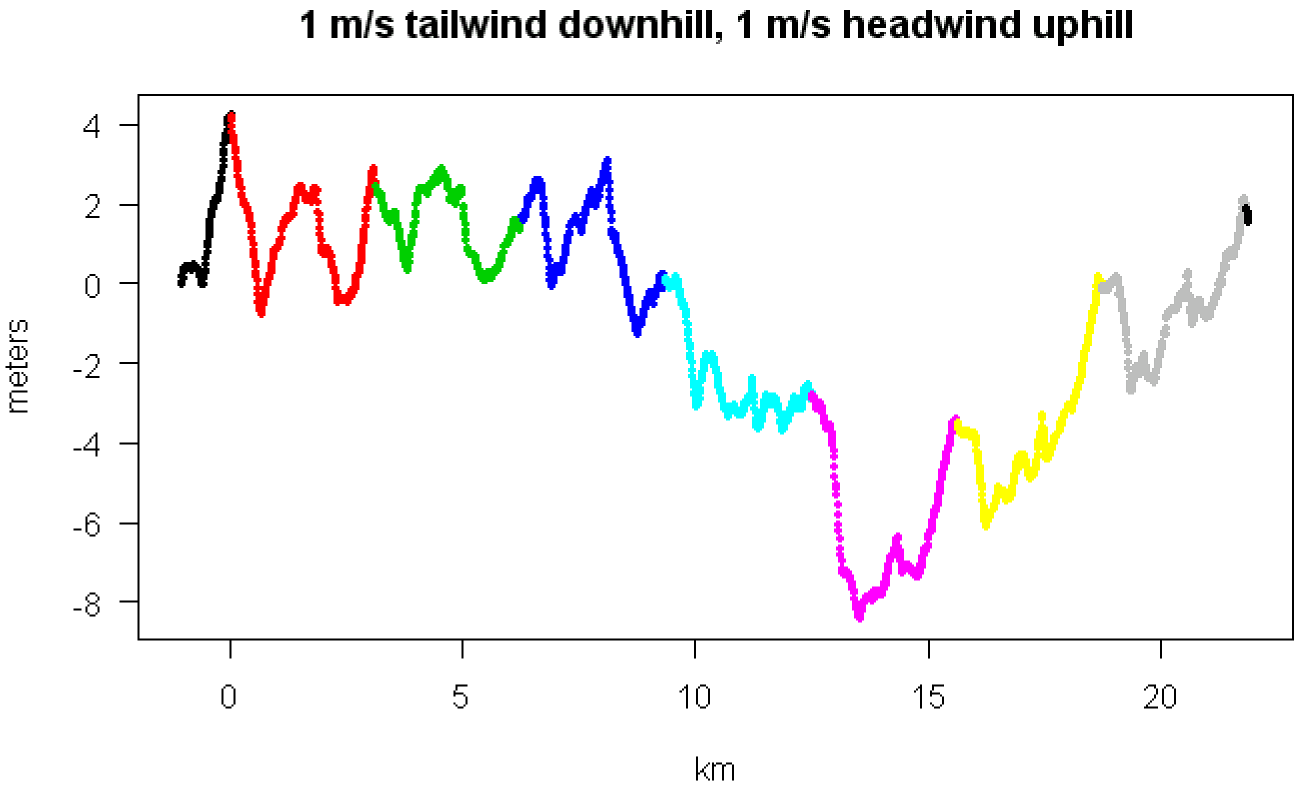

downhill headwind, 1 m/s

the lap amplitudes and profiles are much closer, and total elevation gain over each lap appears to be around 16 meters

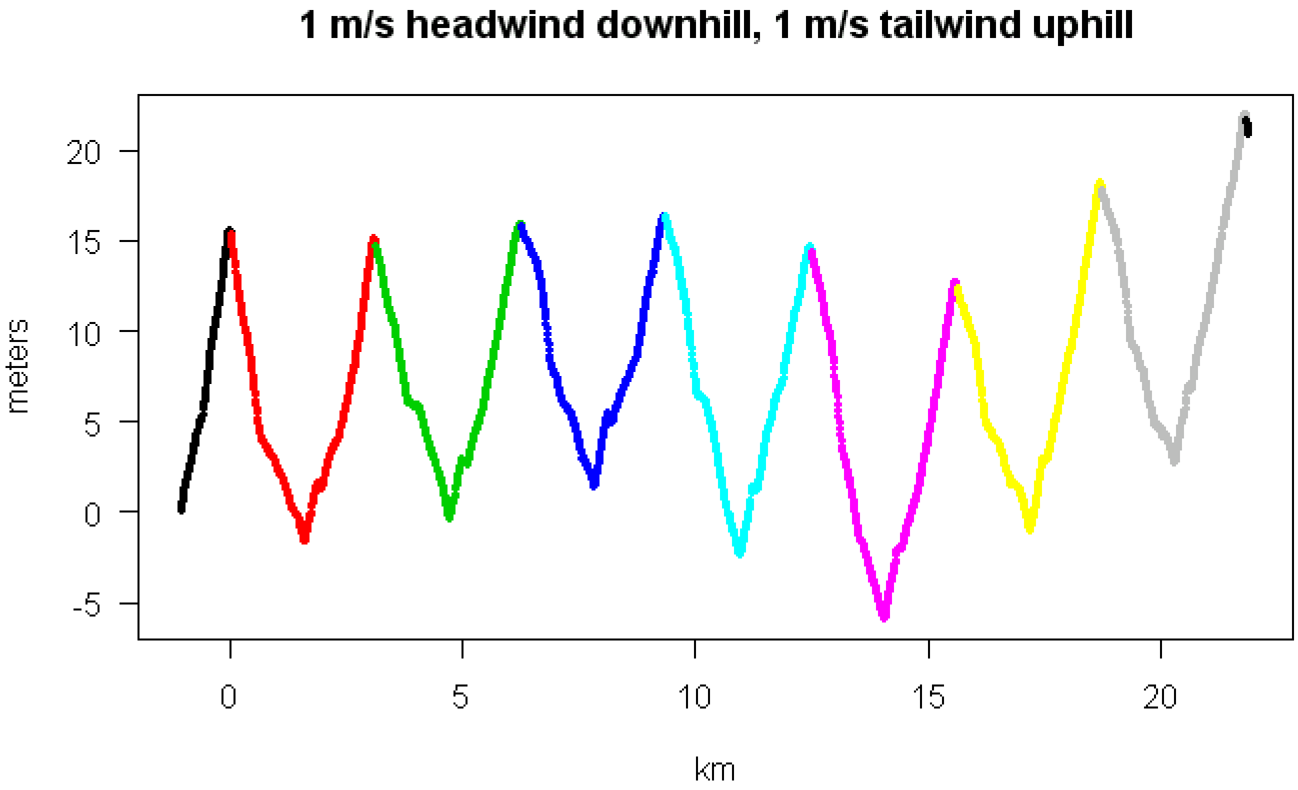

what about virtual wind?

- virtual elevation assumed zero wind. For virtual wind, assume zero elevation change and figure out what the wind must have been
- can you see change in position for last two-and-a-half laps in the virtual wind plot, or the slight use of brakes at “top” of lap 3?
- VE is much smoother than VW

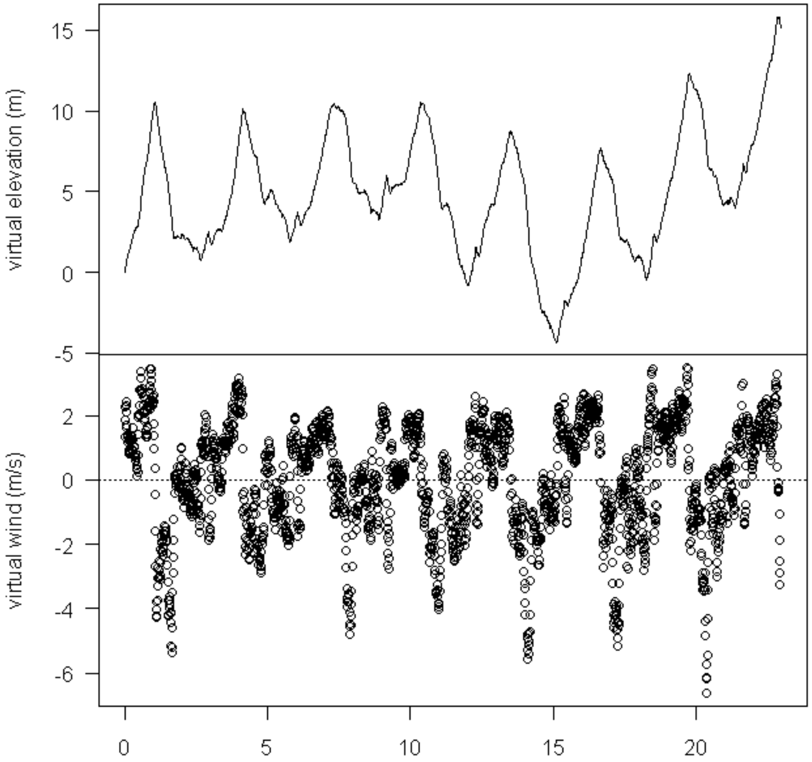

will it work with other examples?

- would I be showing these to you if I thought it didn't?
- when Crr is known, this method matches wind tunnel and classic field tests to within +/- 1%
- the following three examples illustrate the method with data not collected by me
  - three laps at Fiesta Island under windy conditions
  - a (flat, windless) race on the track
  - Dede Demet's Montreal World Cup win: a hilly road race

flat TT

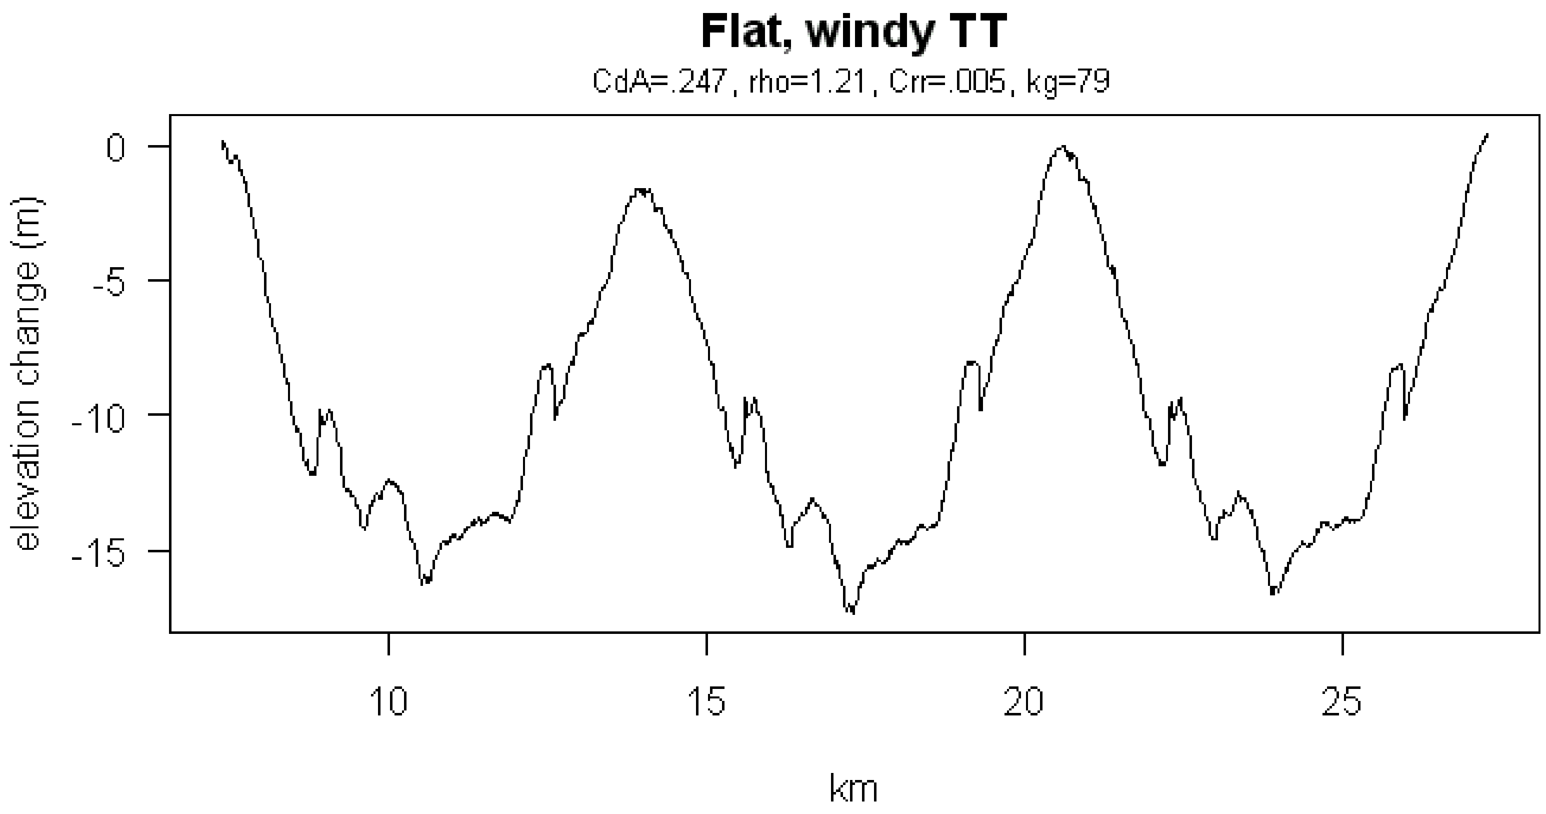
Flat, windy TT
CdA=.247, rho=1.21, Crr=.005, kg=79
elevation change (m)
0
-5
-10
-15
10
15
20
25
km

track race

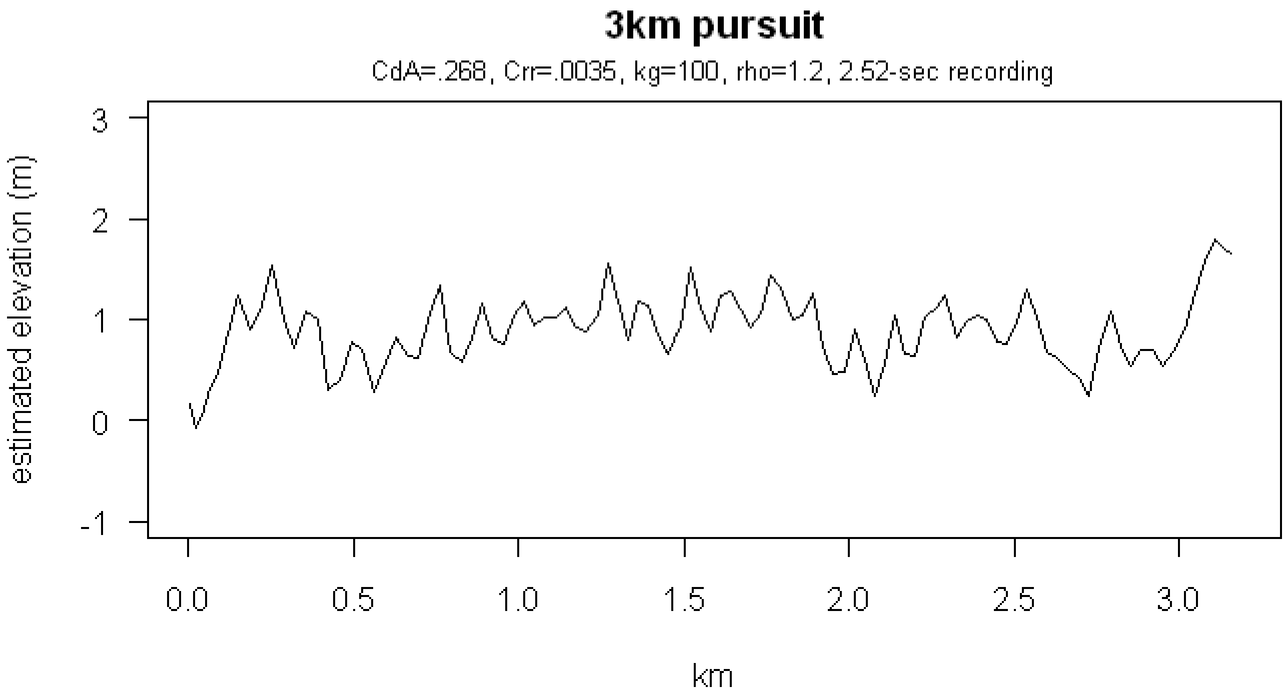

3km pursuit

CdA=.268, Crr=.0035, kg=100, rho=1.2, 2.52-sec recording

world cup road race

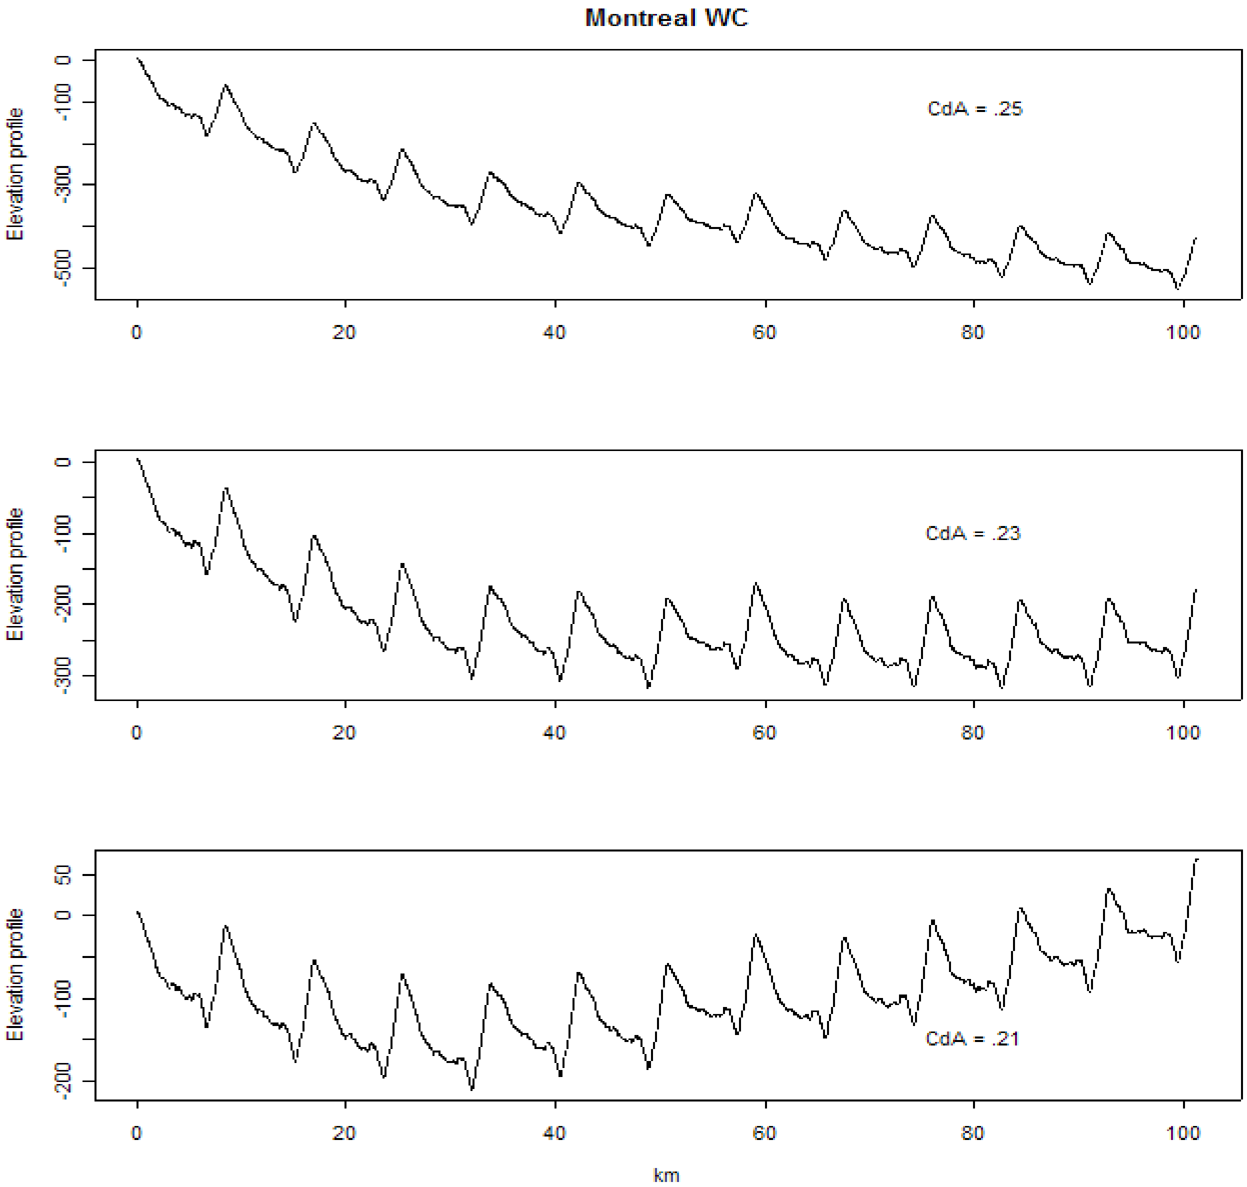

virtual altimeter

---

the next two examples show that the method appears to be both reasonably accurate and consistent

- up and down Palomar Mountain
    - black shows altitude as reported by an altimeter, red shows estimated altitude, emphasizing that braking is ignored
- two years of the San Bruno Hill Climb
    - elevation profiles calculated from power and speed files for the same rider in two consecutive years, showing that the method captures consistent features of the profile

palomar mountain

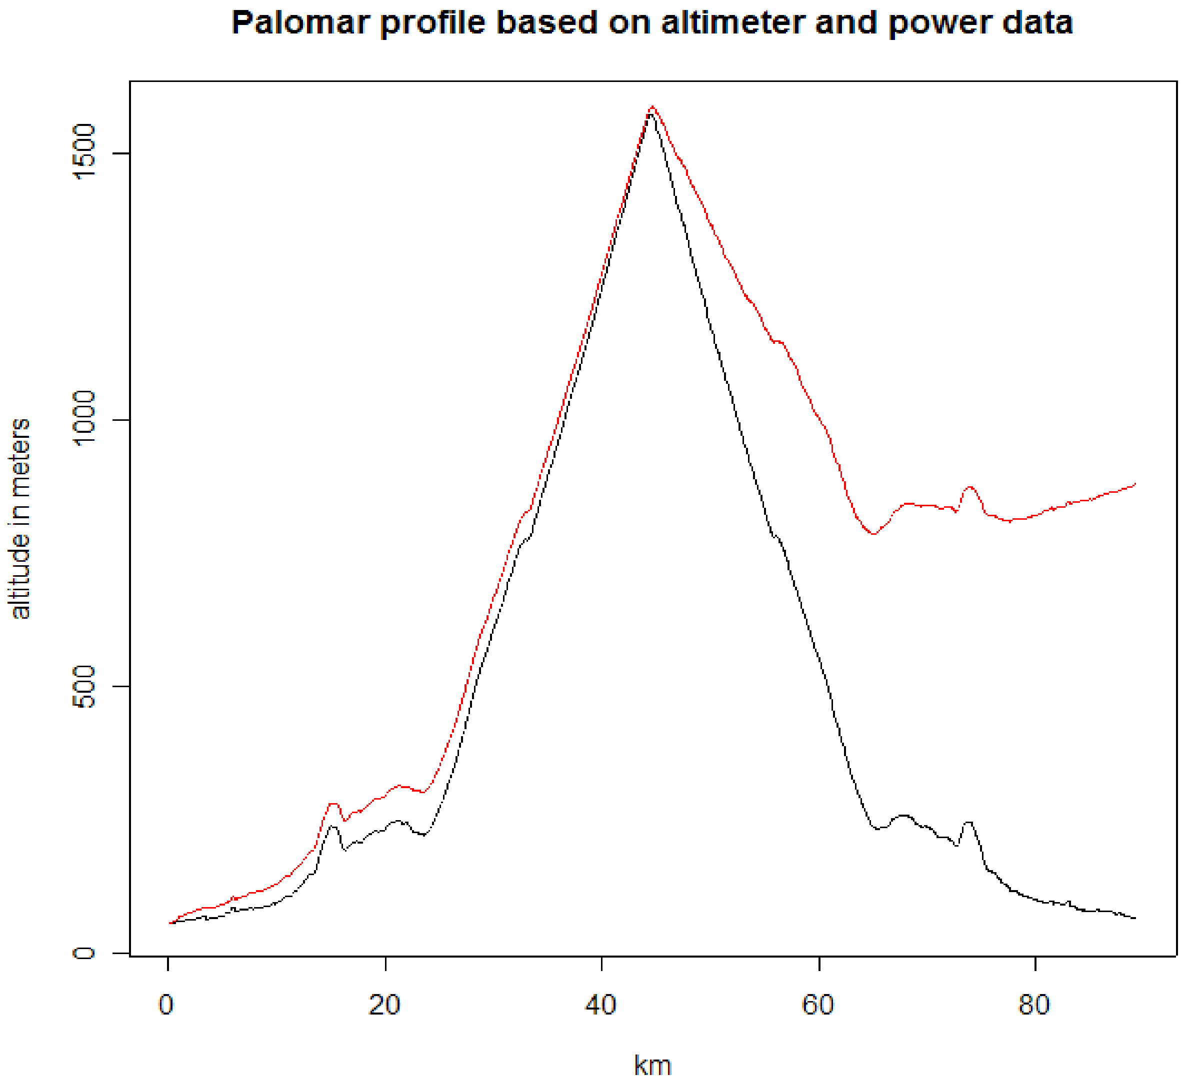

san bruno hill climb

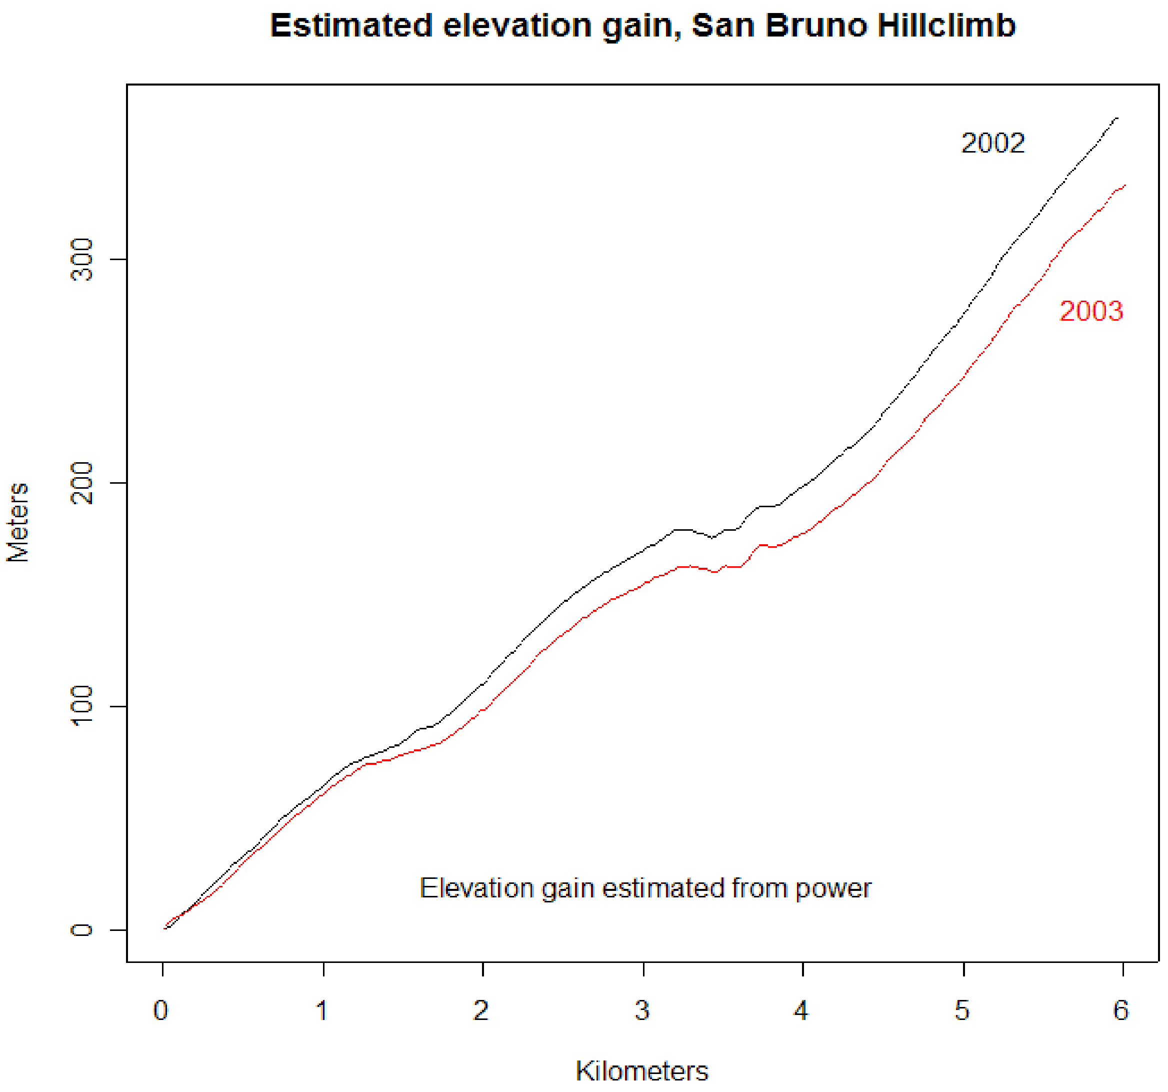

other examples

- growing number of examples of this method being applied by other people on their own data
    - lots of flexibility on venue: out-and-backs, especially for "U"-shaped courses can be good
- spreadsheets exist to simplify calculation

do we need to know Crr?

- often, we don't

  often (though not always) we're more interested in knowing how a change in position or equipment affects CdA. If we test using the same tires on the same roads then Crr is a constant and can be “removed” as a separately estimated variable.

- sometimes, we do

  in those cases, we need a way to estimate both CdA and Crr

prying apart CdA and Crr

- we've been starting with a guess at Crr. Is it possible to estimate Crr separately?
  - sometimes, yes
- first, understand why we've needed to do this
  - the usual field test approach requires flat roads (or constant grade). This approach allows you to use roads that aren't flat so we've lost a constraint. We need some extra information to compensate.
    - one possibility: make known changes and repeat
    - another: know the true elevation profile

extra information demands extra care

- requires very careful test runs
  - test on windless day
  - knowing true elevation profile helps a lot
- two examples
  - same hill, different speeds
  - flat course, monotonic speed change (aka the Shen method)

same hill, different speeds

- I rode up the same hill twice: once slow and once fast(er)
  - first ~ 170 watts, second ~ 245 watts
  - from topo maps climb known to be 114 meters
  - (this time) checked weight and air density
  - almost no wind
- thus far, we have only been solving for (Crr, CdA) pairs
  - if Crr overestimated then calculated CdA will be too low
  - if Crr underestimated then calculated CdA will be too high
  - more importantly, each (Crr, CdA) pair implies a different amount of total climbing
- knowing total elevation gain adds another constraint

Crr and CdA constrain elevation gain

- for a given Crr and speed, you can always find a CdA that matches a given total elevation change
- but for two different speeds there is only one (Crr, CdA) pairing that matches a given total elevation change at both speeds

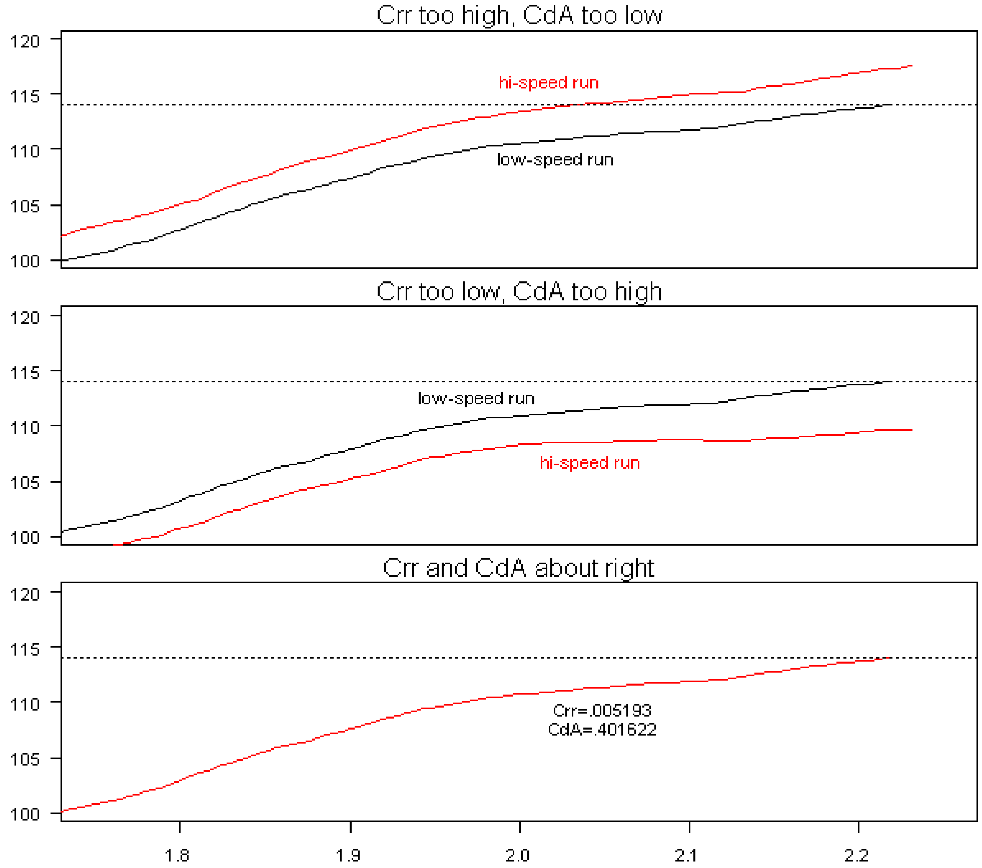

flat course, monotonic speed change

- increasing speed on a velodrome
- only one (Crr, CdA) pairing flattens *both* the overall and individual lap profiles

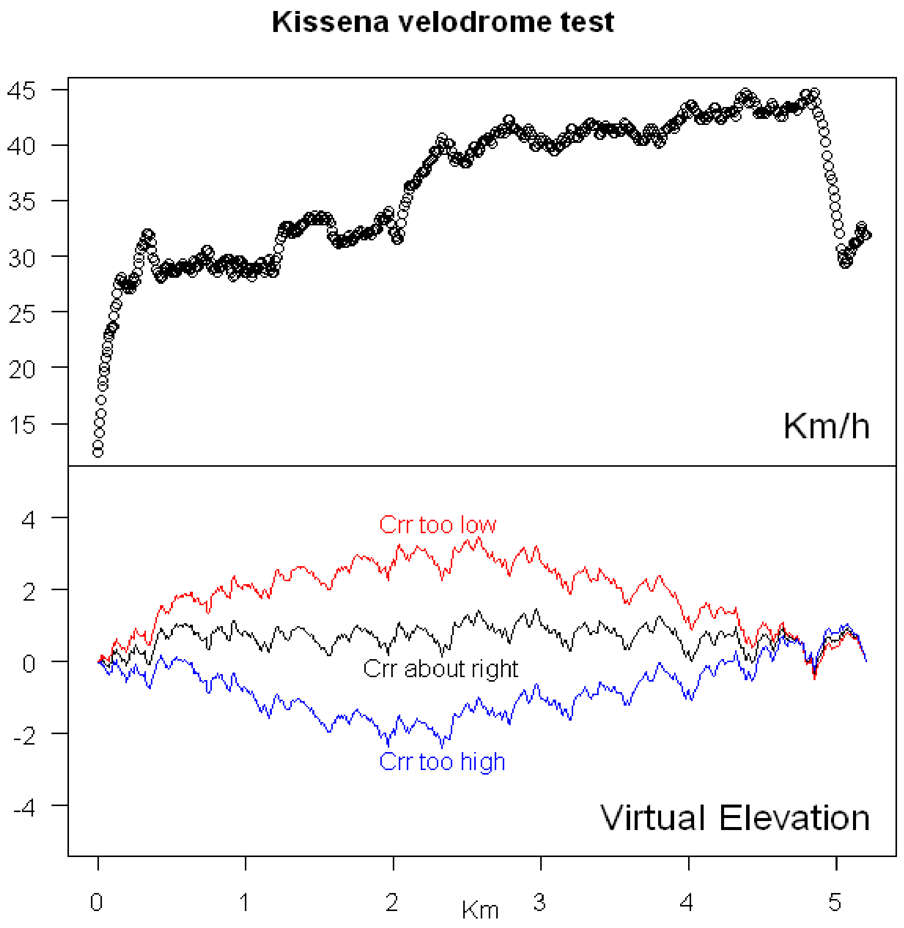

can we generalize?

- Crr and CdA constrain total elevation gain – but they also constrain elevation gain over any segment
- if we know true elevation profile over the entire course we can fit to arbitrary segments
  - this can come in handy for velodrome laps since we know the true profile

summary

- knowing only speed and power still provides an impressive amount of information when data are collected over laps
    - with these data, small changes in CdA are estimable
    - it's possible to examine how these estimates are affected by air density, wind speed, and wind direction
- knowing speed, power, and a little about the course provides even more information
    - you can tune the model not only to line up the profiles but also to match total elevation gain
- in some cases, knowing a bit more info can help you to get separate estimates of Crr and CdA

main conclusion

up to now, if you had an on-bike power meter, most field test protocols required that you find a flat and windless venue. Using this method greatly expands the list of appropriate field test sites.

- a loop like you might use for an industrial park crit could work
- I've used this approach on a long residential block with a small dip in the middle where it crossed a creek
  - ideal venue could be a bowl-shaped street that lets you speed up and still slow down at the ends to make the turnaround
- almost any loop where you don't have to use your brakes could work
- a single out-and-back up a slight hill could work

recommendations

- follow the recommendations in Martin et al. (2006a, 2006b) if you can
- if you can't follow Martin, do laps
  - shorter laps let you do more of them
- don't hold speed constant
  - the wider the spread across laps the easier it is to isolate separate effects
    - a small amount of elevation change can help increase speed variation as long as it's not so steep you need to brake
- measure air density, don't use your brakes, and if you need precision don't do this on a windy day

estimating CdA
with a power meter

R. Chung
rechung@gmail.com

Cyclists interested in going faster want to reduce the amount of drag that they must overcome. Cd is the coefficient of aerodynamic drag. A is front surface area. Their product, CdA, is also known as “drag area” and is typically measured in square meters. CdA represents the fraction of aerodynamic drag compared to a cube with front surface area A held perpendicular to the direction of motion. Sometimes you'll see CxA instead of CdA to specify drag in the direction of travel even if the wind is not from directly ahead. Cd usually ranges between 0 and 1, though direct measurement of Cd is difficult to do so we usually talk about the product, CdA (as it turns out, directly measuring A isn't a piece of cake, either, though from a history of science perspective some measurements of A have demonstrated quite a bit of ingenuity) and it is measured indirectly by drag force. An average-sized cyclist may have a CdA that ranges from around 0.2 m^2 (in a very good aerodynamic time trial position) up to perhaps 0.8 m^2 (on an upright “city” type bicycle).

The gold standard for measuring CdA is the wind tunnel but wind tunnels appropriate for testing aero drag on a bike are relatively expensive ($500 - $1000/hr). Since the demand for knowing CdA (or, at a minimum, determining differences in CdA between alternative positions or equipment) is high, riders have turned to field testing to try to estimate CdA. Field testing introduces its own set of problems: notably, that at best it can measure total drag (=aero drag + rolling resistance drag) so the aero component must be separated out. Traditionally, rolldown or coastdown tests were used though high variability in the estimates has been observed. With the introduction of on-bike power meters results have been both more accurate and more precise—but finding the right venue and collecting the right data to achieve good estimates is still challenging.

This presentation looks at real-world data collected by an on-bike power meter to assess how well these data can estimate CdA. In particular, it shows that the usual approach is not robust if data collection is less than perfectly controlled, presents an alternative indirect method that can produce good results under more general conditions, and uses these findings to make recommendations about how to minimize critical errors and to improve estimates.

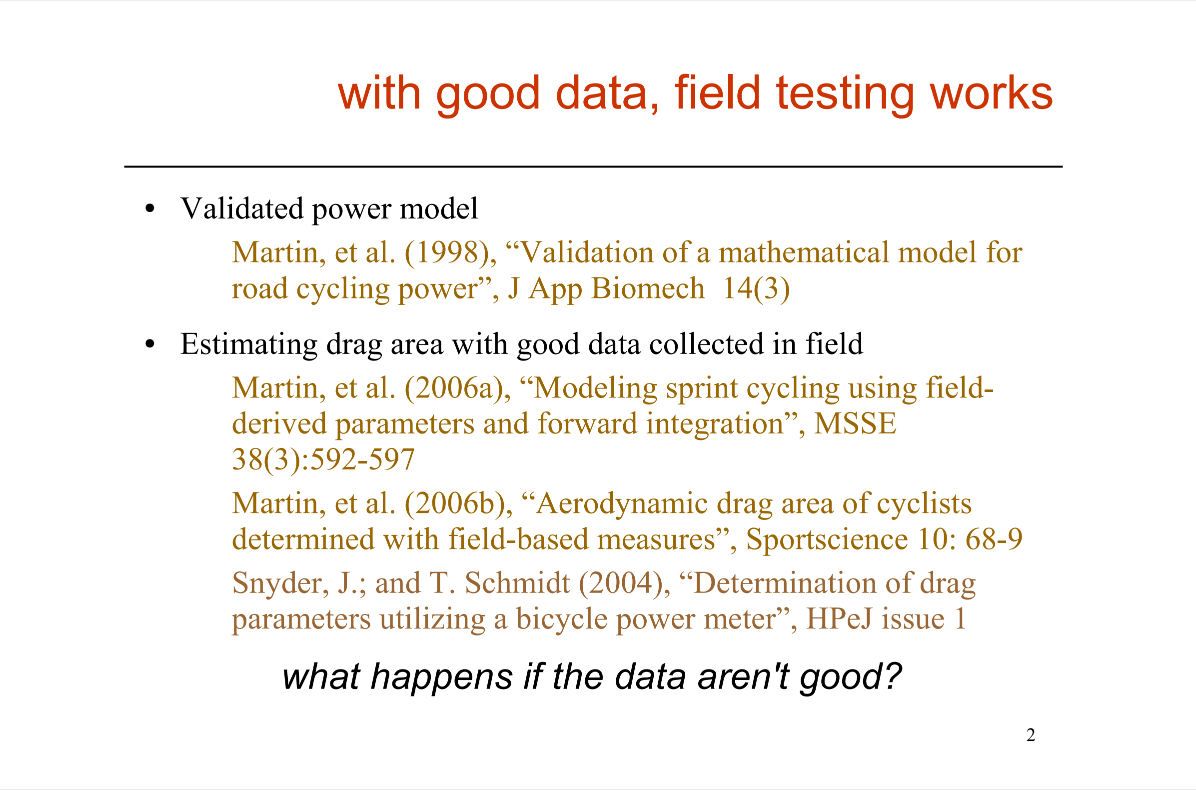

with good data, field testing works

- Validated power model
  - Martin, et al. (1998), “Validation of a mathematical model for road cycling power”, J App Biomech 14(3)
- Estimating drag area with good data collected in field
  - Martin, et al. (2006a), “Modeling sprint cycling using field-derived parameters and forward integration”, MSSE 38(3):592-597
  - Martin, et al. (2006b), “Aerodynamic drag area of cyclists determined with field-based measures”, Sportscience 10: 68-9
  - Snyder, J.; and T. Schmidt (2004), “Determination of drag parameters utilizing a bicycle power meter”, HPeJ issue 1

what happens if the data aren't good?

Sometimes we can directly observe and measure the thing we're trying to estimate. In cases where we can't, sometimes we can measure other related variables and calculate the thing we want by using known relationships among the variables. All field-based methods for estimating CdA are indirect. We'll discuss two: the “classic” regression method and one introduced here.

Field-testing, when carefully done, can produce good results – so why the heck would anyone care about using less-than-perfect data? The main reason is that sometimes, no matter how hard you try, the data you collect aren't perfect. We have three questions:

1. How lousy can field testing be and still produce reasonable results?

2. Are some methods better with less-than-perfect data than others?

3. What can we learn from this that can help improve field test estimates?

I come from a field where, for the most part, we can't run experiments and the data we work with tends to be expensive and difficult to replicate. Because of this we've developed a toolbox of techniques to salvage “dirty” data using robust indirect methods. Although the usual goal is to salvage bad data, perhaps the most valuable lesson is helping us to improve data collection when we get the chance to do it by learning which data elements are critical and how sensitive the final result will be to errors, and how to evaluate that sensitivity by developing a method to measure goodness-of-fit. That's important, so let's repeat it: finding a way to make an estimate is only half the job; the other half is finding a way to tell when the estimate is lousy, and by how much.

What I'm presenting here may seem like the long way 'round the problem but it's done this way for a reason: I'll show one way of many to make an estimate but, more importantly, how to construct diagnostics that you can use to tell when the estimate is lousy.

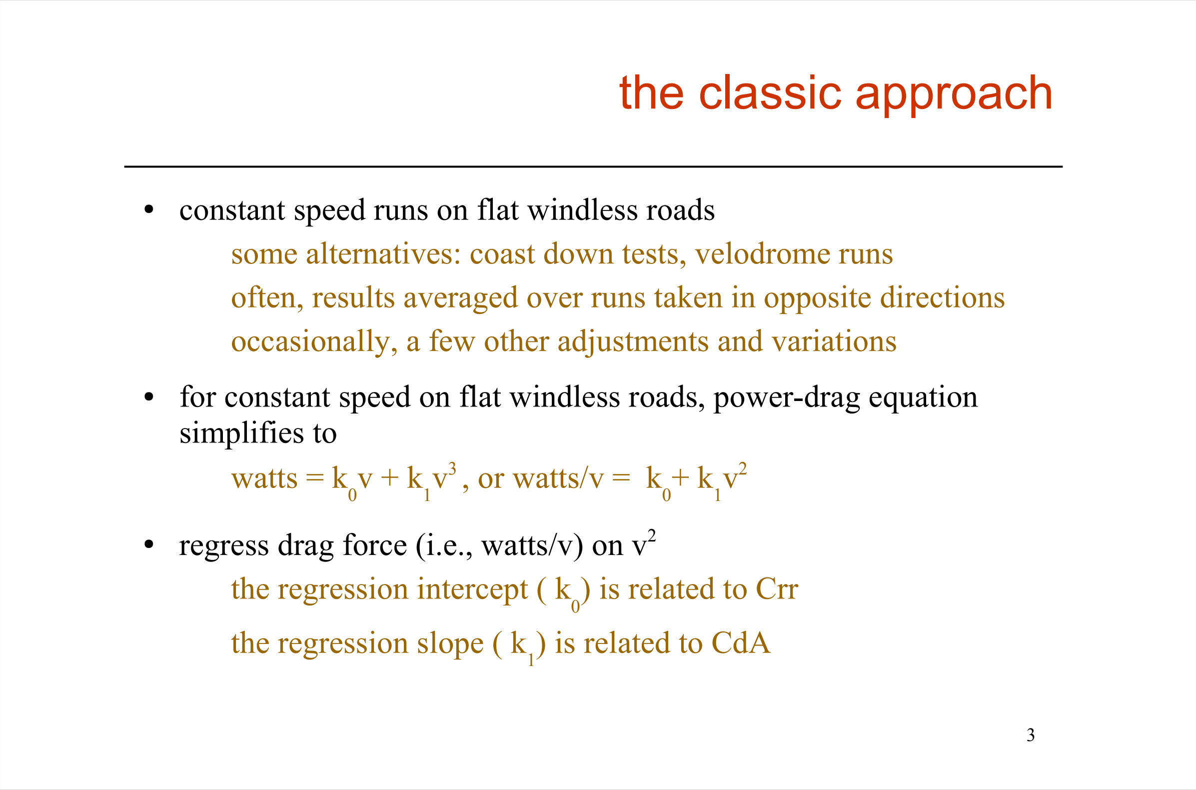

the classic approach

- constant speed runs on flat windless roads
  - some alternatives: coast down tests, velodrome runs
  - often, results averaged over runs taken in opposite directions
  - occasionally, a few other adjustments and variations
- for constant speed on flat windless roads, power-drag equation simplifies to
  - watts = $k_0 v + k_1 v^3$, or watts/v = $k_0 + k_1 v^2$
- regress drag force (i.e., watts/v) on v^2
  - the regression intercept (k_0) is related to Crr
  - the regression slope (k_1) is related to CdA

Other approaches include John Tetz' coastdown method, popular on the HPV circuit. See http://www.recumbents.com/mars/pages/proj/misc/coastdown.html

In the Human Power eJournal, John Snyder and Theo Schmidt present a method similar to the classic method here: http://www.hupi.org/HPeJ/0005/0005.htm

H.W. Schreuder presents a high precision coastdown method in http://www.xs4all.nl/~cp4trml/metingen/measurements.html using a datalogger similar in spirit to the one described here: http://www.hupi.org/HPeJ/0012/0012.html

Candau used coastdowns in a hallway timing trap using electric “eyes” similar to the ones used to ring a bell when a customer walks into a shop. See Candau, R. et al. (1999) “Simplified deceleration method for assessment of resistive forces in cycling.” MSSE 31(10): 1441.

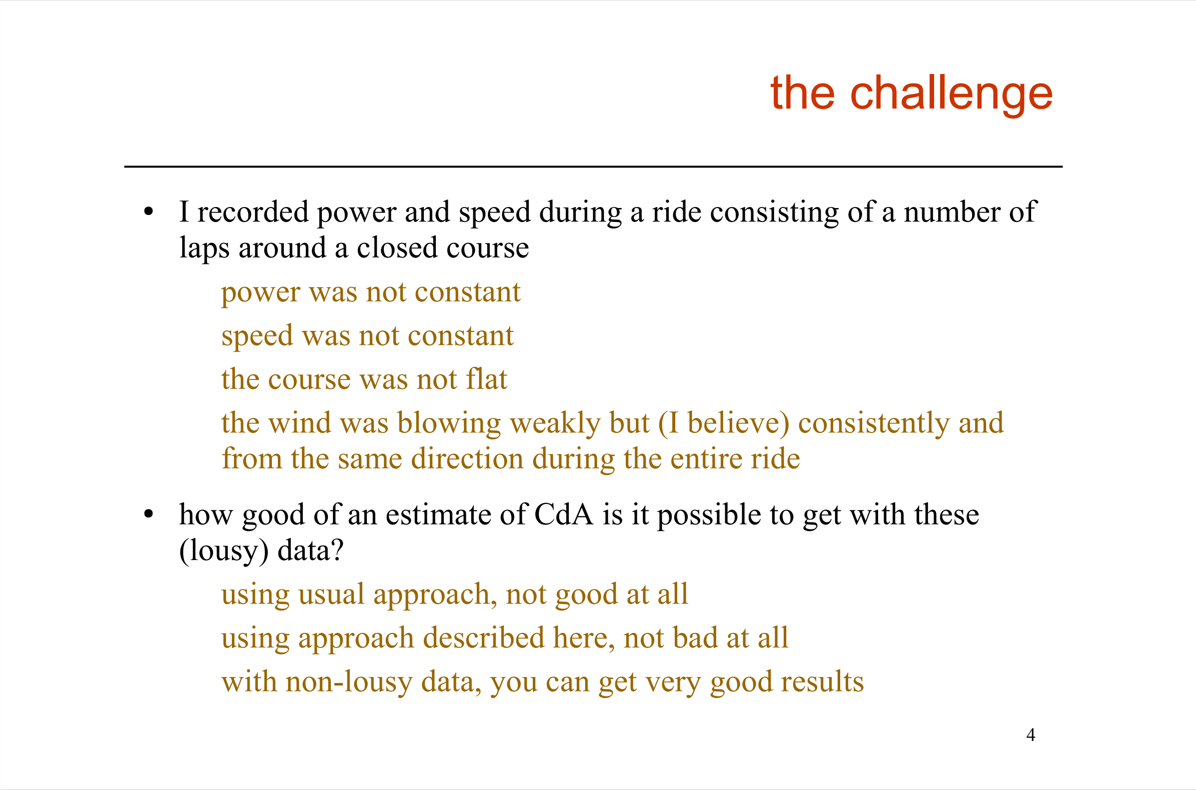

the challenge

- I recorded power and speed during a ride consisting of a number of laps around a closed course
    - power was not constant
    - speed was not constant
    - the course was not flat
    - the wind was blowing weakly but (I believe) consistently and from the same direction during the entire ride
- how good of an estimate of CdA is it possible to get with these (lousy) data?
    - using usual approach, not good at all
    - using approach described here, not bad at all
    - with non-lousy data, you can get very good results

At first glance, it seems these data are neither controlled enough nor detailed enough to produce a reasonable estimate of CdA. In fact, it's even worse: unlike most cyclists who would attempt CdA field testing, I knew neither my exact weight nor that of my bicycle, nor did I measure wind speed or the elements I needed in order to calculate air density. Given these circumstances, most people would dismiss the estimation as unmanageable.

However, not all methods used for estimating parameters of a model are equal. Here, I describe a method that, under certain broad conditions, can be used to make a reasonable estimate of CdA from data such as these – in certain narrow situations, one that is accurate, consistent, and with high discrimination and repeatability. Of perhaps even greater importace, I present diagnostic procedures to recognize when those conditions don't apply, so you won't be misled into thinking the estimate is accurate or consistent when it isn't. I will also discuss some of the weaknesses of this approach, and will show that the challenge data reveal characteristics that allow a reasonable estimate even though all we know are speed, power, and that the course consisted of a series of laps. I conclude with recommendations for field testing of aerodynamic drag.

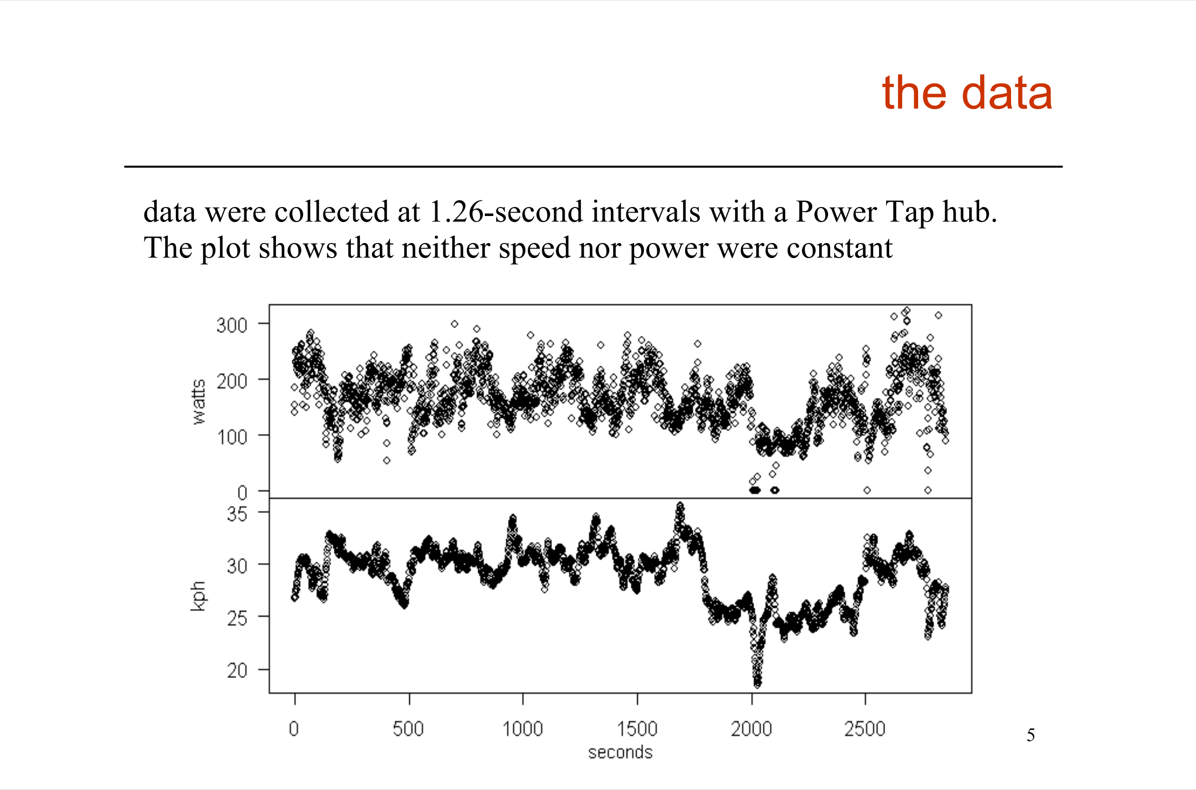

the data

data were collected at 1.26-second intervals with a Power Tap hub.
The plot shows that neither speed nor power were constant

You can download the data and experiment with them yourself:

http://anonymous.coward.free.fr/wattage/cda/field-cda-challenge.csv

The first few lines of the data file look like this:

```
"secs","watts","kph"
1.26,142,26.8
2.52,154,26.7
3.78,185,26.7
5.04,253,26.8
6.3,247,27.5
```

i.e, all we know are speed, power, the time interval, and that the data are in seqential order.

Torque was zeroed before the ride according to manufacturer's recommendation. The hub was checked with a static calibration test against known weights about 3 months before the data were collected and found to be (well) within specified limits for accuracy. That is, these data were collected with a standard hub in ordinary working condition using nothing special that isn't available to anyone who has a Power Tap hub.

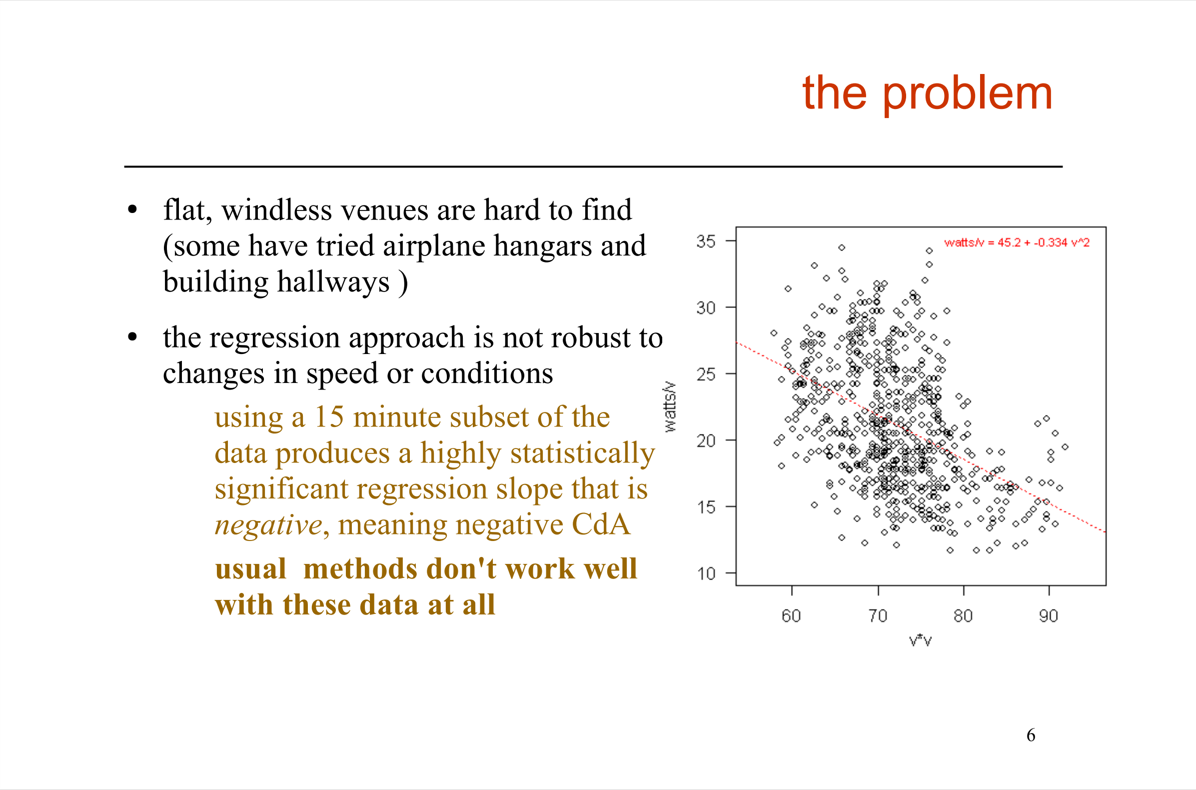

the problem

- flat, windless venues are hard to find (some have tried airplane hangars and building hallways)
- the regression approach is not robust to changes in speed or conditions

  using a 15 minute subset of the data produces a highly statistically significant regression slope that is *negative*, meaning negative CdA

  usual methods don't work well with these data at all

Remember, CdA is an *area.* You can't have negative area.

In a deeper sense, we started off knowing that these data didn't fit the standard assumptions of constant speed on a flat course, so we shouldn't be surprised that the fit isn't good. However, this shows that the fit isn't just not good, it's *catastrophically terrible* – and that in this context the usual regression approach isn't at all robust to failures of the assumptions. Part of the reason why regresssion has become such a popular technique is because when it fails it often fails gracefully. Alas, we've just demonstrated that "often" does not mean "always," and explains why testers try to control the conditions as carefully as possible – yet they still often don't get a good estimate of CdA.

Why is this non-robust approach the usual one? I suspect, but do not know, that the answer is historical: before the introduction of devices that could record speed and power, analyzing detailed second-by-second data simply wasn't an option. The usual approach works well when all one can analyze are averages collected over flat runs at constant speed; at that, it was an improvement over coast-down or roll-down tests. Candau et al. (1999 MSSE 31(10):1441-7) used coast-downs in building hallways to control slope and crosswinds. Jim Papadopoulos, who helped revise the latest edition of *Bicycling Science*, suggests testing on flat running tracks with an extra cycling computer set to show average speed so the rider can control speed even more strictly. In addition, performing regressions over a series of averages is relatively easy so at a time when data were sparse, computers were rare, and speed and power recorders were nonexistent, it was a clever approach that simplified data analysis.

You may be able to discern a philosophical conflict here: the experimental approach is to tightly control conditions which simplifies the analysis. I come from a field where we can't run controlled experiments so we develop (slightly?) complex analytical tools, then ask what happens if the conditions were poorly controlled.

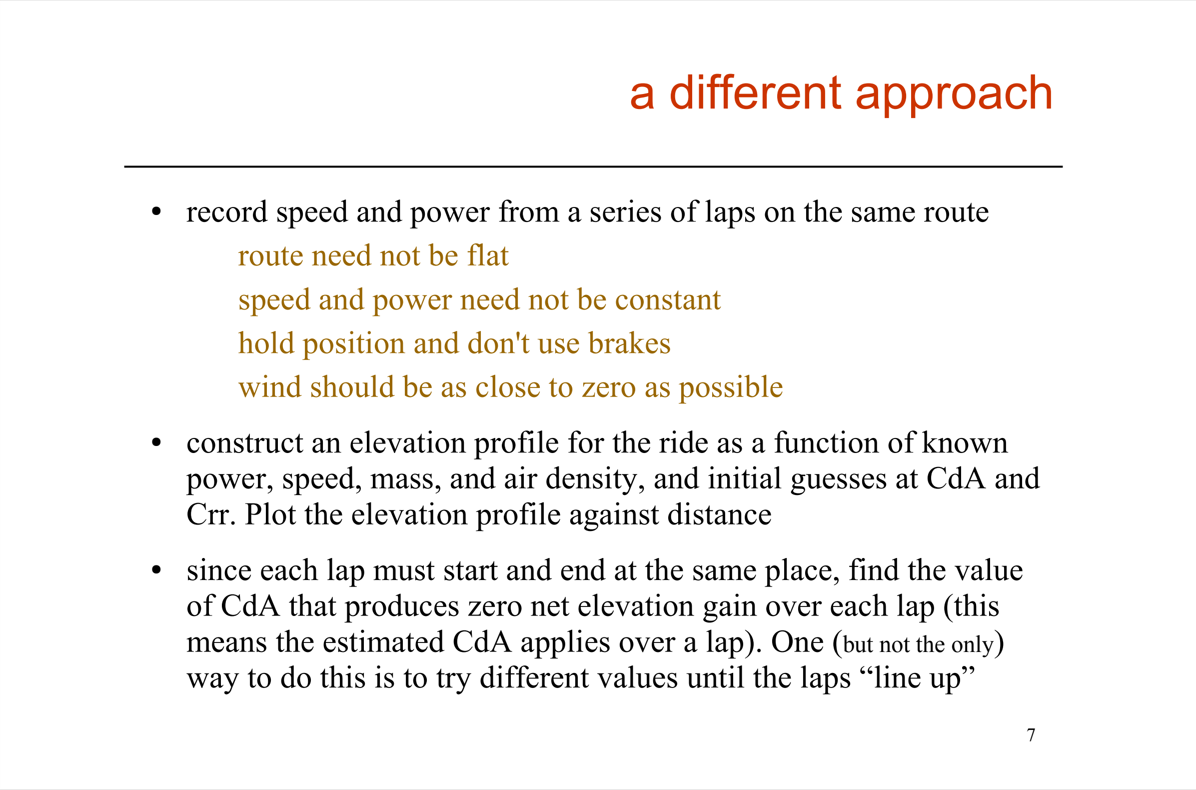

a different approach

- record speed and power from a series of laps on the same route
  - route need not be flat
  - speed and power need not be constant
  - hold position and don't use brakes
  - wind should be as close to zero as possible
- construct an elevation profile for the ride as a function of known power, speed, mass, and air density, and initial guesses at CdA and Crr. Plot the elevation profile against distance
- since each lap must start and end at the same place, find the value of CdA that produces zero net elevation gain over each lap (this means the estimated CdA applies over a lap). One (but not the only) way to do this is to try different values until the laps "line up"

7

The third point is key. Nothing comes for free and this approach adds an additional constraint that is absent in the conventional approach: elevation gain must net to zero for each lap. (Later, I'll show an approach that lets you use a course with a known true elevation gain.)

When I first tried this approach in 2003, it was prompted by an attempt to examine how well the HAC4 could estimate power from elevation change. A side effort was to "reverse-HAC" power data to see how well I could back out elevation change. I described it at

http://anonymous.coward.free.fr/wattage/altimeter/pseudohac4.html and

http://anonymous.coward.free.fr/wattage/altimeter/altimeter.html

In contrast to the usual approach, this approach demands much more data and a bit more calculation—but demands much less control over the conditions of the test. In essence, it depends on hundreds of times as many observations in order to correct for deviations from controlled conditions: it clearly would have been impractical to do before the advent of modern power meters and analytical tools.

Note that coast down tests are a subset of these conditions: the route isn't flat, speed isn't constant, you must hold position, don't use the brakes – but power is constant at zero.

There is another irony: although in my real job I usually teach statistical research methods, this particular approach doesn't use standard statistical techniques as does the conventional approach. Here, we exploit physics modeling to estimate one of the model parameters on a record-by-record basis, then examinethe behavior of that parameter. There are several model parameters we could have modeled but here we choose to model slope. The virtue of this particular parameter is that it lets us evaluate the course profile as an overall diagnostic of fit. We then tune the course profile so it matches from lap to lap.

There are ways to assess fit and precision more formally and we'll discuss those a bit later.

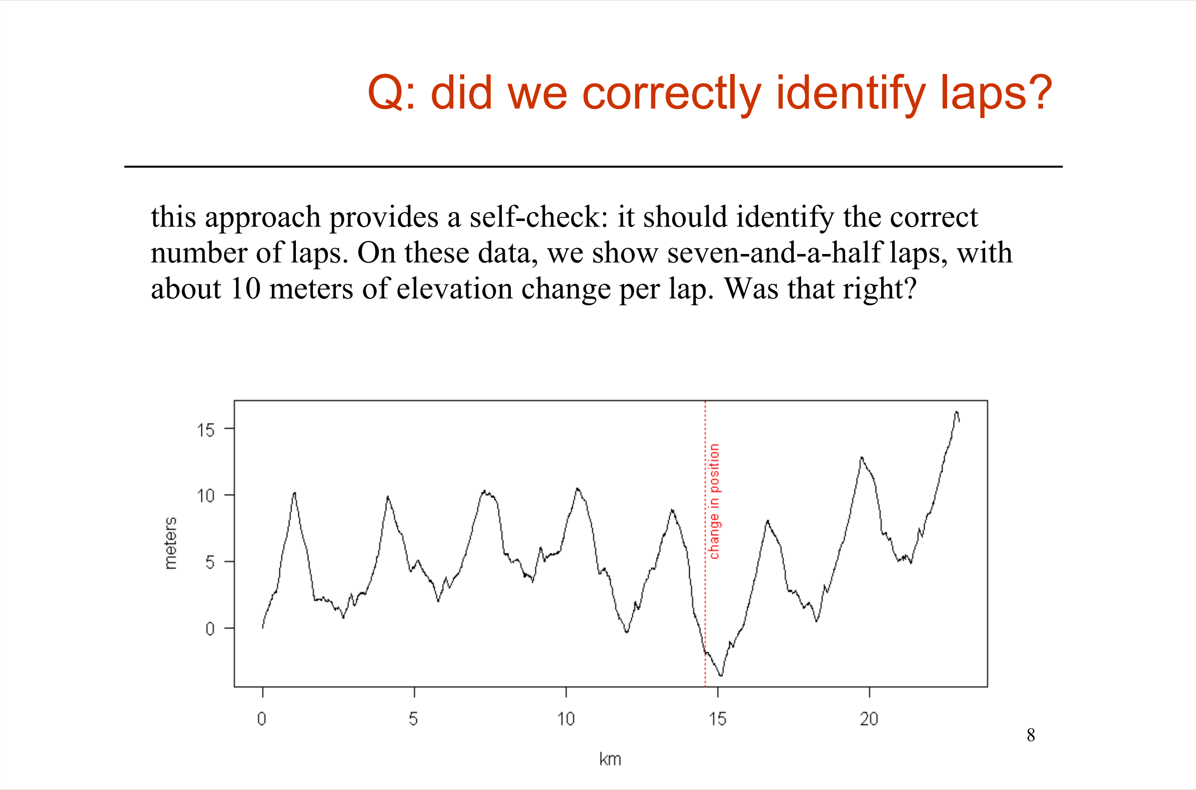

Q: did we correctly identify laps?

this approach provides a self-check: it should identify the correct number of laps. On these data, we show seven-and-a-half laps, with about 10 meters of elevation change per lap. Was that right?

8

There are two things I didn't mention in the original challenge, and they both emphasize why I begin by describing this method as a graphical tuning: first, it appears that there may have been an extra half lap; and second, it's easy to see that the first few laps lined up, while the last couple of laps seem like they were collected under different conditions. Both of those turn out to be true: I entered the course in a different spot than I exited, and I sat up for the last couple of laps. The dotted red line shows when I changed position from brake hoods to bar tops. On my bike, I have the brake hoods in “classic” position (i.e., lower than is now popular) so there is a height difference between being on the hoods and being on the tops. On the other hand, when I'm on the tops, I usually hold my hands nearer the stem, so I'm narrower. I did not expect that the difference between being lower but wider vs. higher but narrower would be so easily spotted.

The take-home message is that while it's possible to do this analytically (and I will, later in the presentation) it's important to start graphically because the graph is a key diagnostic tool.

Note the scale on the x-axis is in km and the scale on the y-axis is in meters – the vertical scale is exaggerated by two orders of magnitude. This is a flat course.

A final observation: although the course was closed to cars, there were other riders on the course. I was able to do the entire ride shown without using my brakes except for one brief moment. Notice the flattened top of the third peak? That's where a rider in front of me went wide and I feathered my front brake to avoid him. The other peaks are much more pointed. Does it make sense to you that a brief use of the brakes should appear as a flattening?

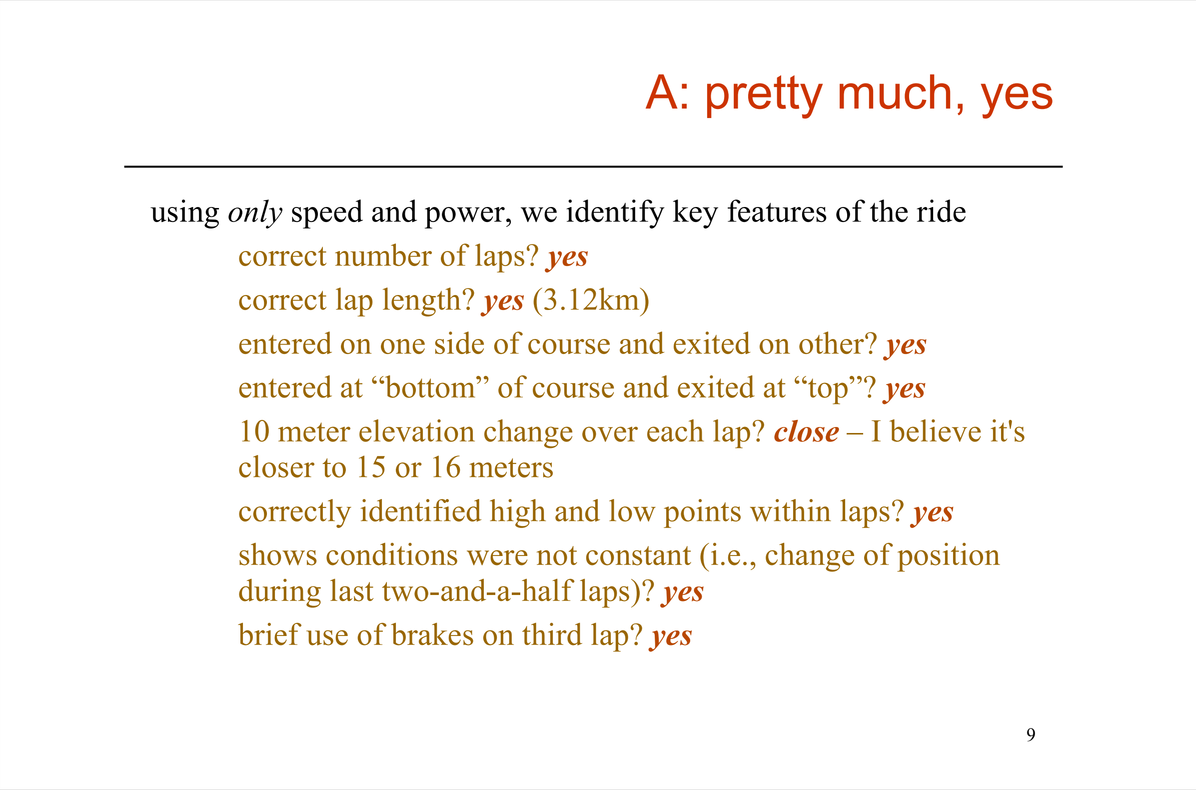

A: pretty much, yes

using *only* speed and power, we identify key features of the ride

- correct number of laps? ***yes***
- correct lap length? ***yes*** (3.12km)
- entered on one side of course and exited on other? ***yes***
- entered at “bottom” of course and exited at “top”? ***yes***
- 10 meter elevation change over each lap? ***close*** – I believe it's closer to 15 or 16 meters
- correctly identified high and low points within laps? ***yes***
- shows conditions were not constant (i.e., change of position during last two-and-a-half laps)? ***yes***
- brief use of brakes on third lap? ***yes***

This method identifies features of the laps quite clearly from just power and speed.

The method missed on elevation, but notice that it didn't show random fluctuations over the course, or 200 meters of elevation change: it correctly identifies the course as nearly flat, with the same features in the same places on each lap. Remember that we've assumed the wind was zero although we know it wasn't. This method tosses all unmodeled variations into the estimated slope so the wind shows up as “virtual” elevation and gives us a rough idea of how strong the wind was: it translated into about 5 fewer meters of elevation change for each 3.12 km lap.

This is a clue as to why it doesn't matter that much what my total mass is. This is a relatively flat course, so we'd expect that total drag force would be dominated by aero drag, not drag due to lugging my fat rear up a hill. As we will later see when we explore the sensitivity of the method, a 1% increase in total mass resulted in a 0.3% decrease in estimated CdA, i.e., I could be off by 7% on my weight and the error in the estimated CdA would only be about 2%. On this course the estimate of CdA is pretty inelastic with respect to mass. Keep this in mind when we discuss hills.

We now know the lap length to be 3.12 km. We'll use this information later. In the meantime, note that the wind doesn't appear to change the lap length very much in any of the data files we'll be looking at. If the wind is gusty and strong (relative to the actual gradient) and if it changes direction then the lap lengths may be harder to identify. This is a foreshadowing of a useful diagnostic that we'll discuss later.

lap lengths are well identified

- different guesses about CdA (or Crr or mass or air density) have only a small effect on the estimate of lap lengths
  - changes in the parameters move the curve up-and-down but not left-and-right
- small changes in wind don't affect lap length much so lap lengths are relatively robust
  - however, big changes in wind may
- relatively robust identification of the lap lengths means that it's feasible (though not always wise) to impose the “zero net elevation gain” constraint on each lap
- if you're interested, a Google map of the course is here: http://tinyurl.com/yq9r76

The route follows a clockwise course: northward up the Route Dauphine, eastward onto Route Royale de Beaute, southwest down the Route de Bourdon, and finishes with a hairpin back onto Dauphine. For these data, I started at the northeast corner and exited at the hairpin. According to Google Earth, the high point is at the hairpin, at 62m ASL; the low point is around the “e” in “Royale” at 46m. This surprised me—my “feeling” from riding the course was that the ride was a bit flatter, and that the low point was on Dauphine. As an aside, if you flip to the satellite view and zoom all the way in, you can see individual riders.

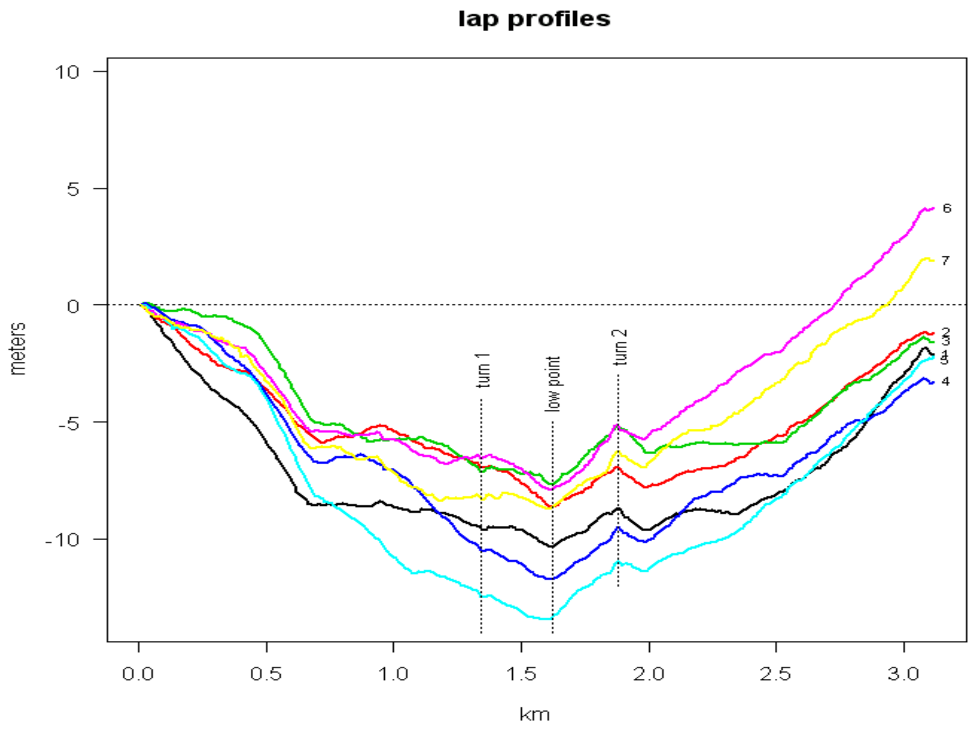

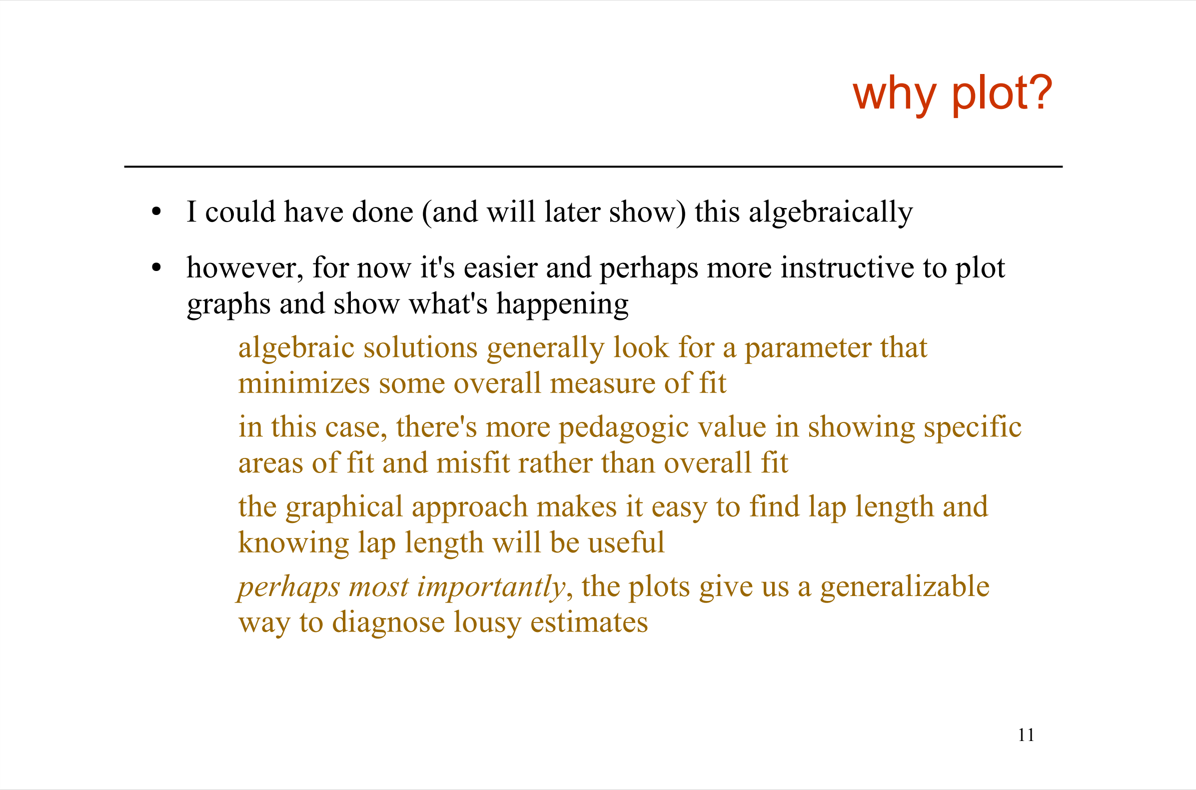

why plot?

- I could have done (and will later show) this algebraically
- however, for now it's easier and perhaps more instructive to plot graphs and show what's happening
    - algebraic solutions generally look for a parameter that minimizes some overall measure of fit
    - in this case, there's more pedagogic value in showing specific areas of fit and misfit rather than overall fit
    - the graphical approach makes it easy to find lap length and knowing lap length will be useful
    - *perhaps most importantly*, the plots give us a generalizable way to diagnose lousy estimates

One closed-form algebraic approach is to transform the power equation into a work equation, then solve for the CdA. To do this, integrate the power equation over time. I'll come back to this later.

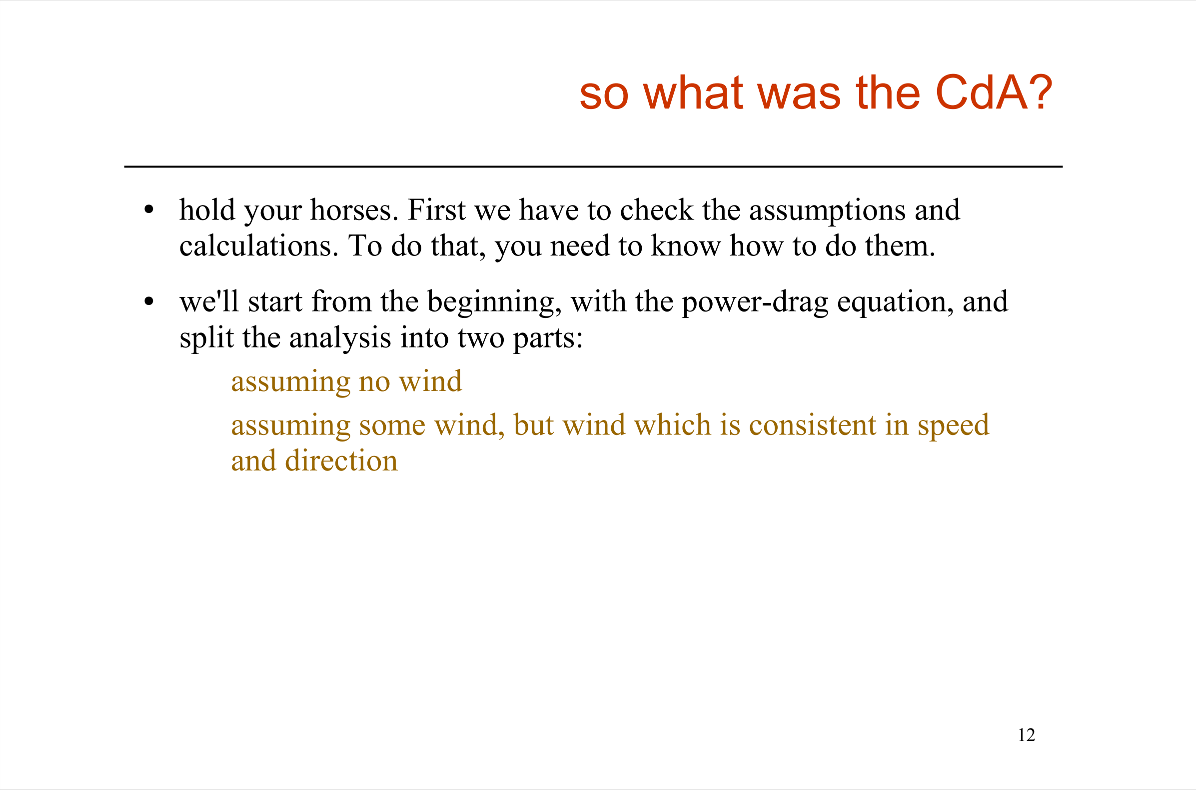

so what was the CdA?

- hold your horses. First we have to check the assumptions and calculations. To do that, you need to know how to do them.
- we'll start from the beginning, with the power-drag equation, and split the analysis into two parts:
  - assuming no wind
  - assuming some wind, but wind which is consistent in speed and direction

For these data, we don't know wind, Crr, air density, or total mass – though if you were going to collect your own data you should probably weigh yourself and check the weather service for barometric pressure, temperature, and your elevation above sea level (barometric pressures are often, but not always, normalized to sea level so you should double check for your location). I wasn't kidding when I said that these data were less than ideal.

Before we get to the estimate of CdA, we need to double-check the assumptions and how sensitive the method is to them.

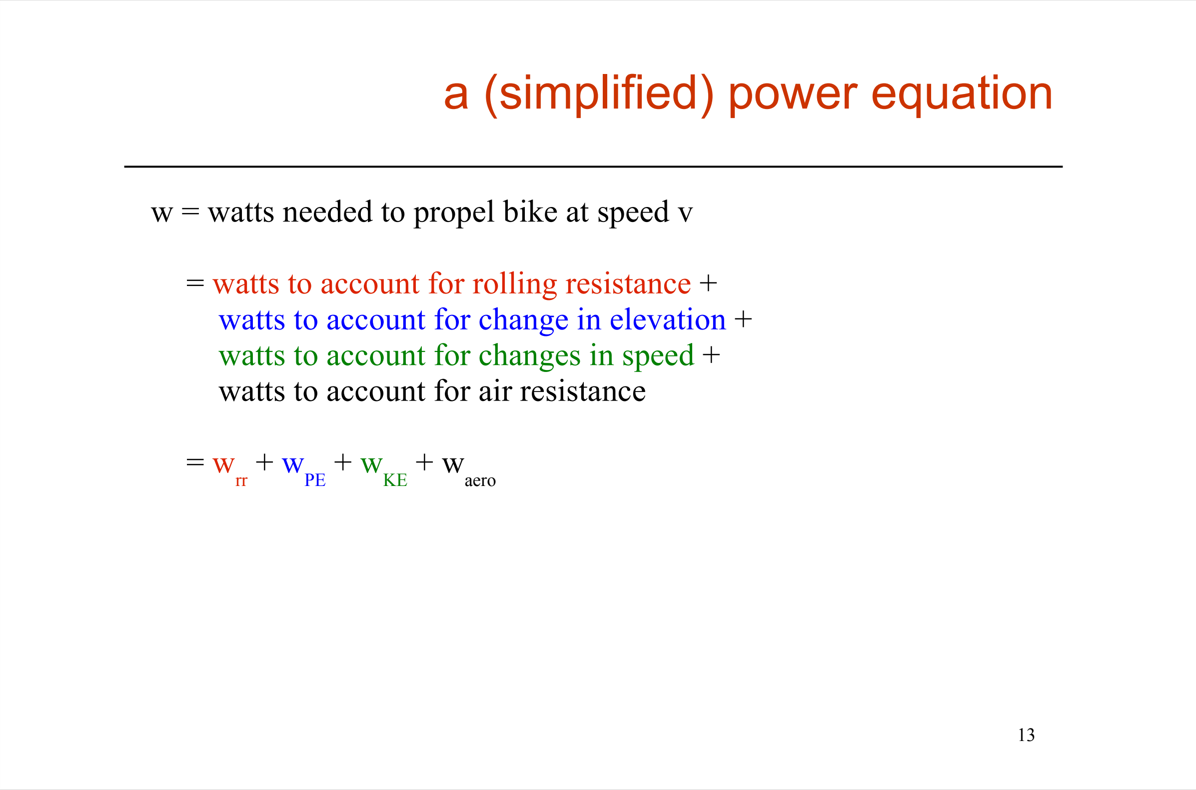

a (simplified) power equation

w = watts needed to propel bike at speed v

= watts to account for rolling resistance +
watts to account for change in elevation +
watts to account for changes in speed +
watts to account for air resistance

$= w_{rr} + w_{PE} + w_{KE} + w_{aero}$

Many models include a term for overall drivetrain efficiency, η, but all of the data files I'm looking at come from Power Taps which, in theory, should be downstream of drivetrain losses, i.e., η = 1. If you have an SRM, which measures power at the crank (i.e., upstream of drivetrain losses), you will want to decide how to model drivetrain losses. Martin et al. presumed a fixed percentage loss of 2.3% of power (i.e., η = 0.977). Other choices might include a fixed wattage loss, or loss with two components: a fixed amount and a fixed percentage.

PE is “potential energy” and represents change in elevation.

KE is “kinetic energy” and represents change in speed.

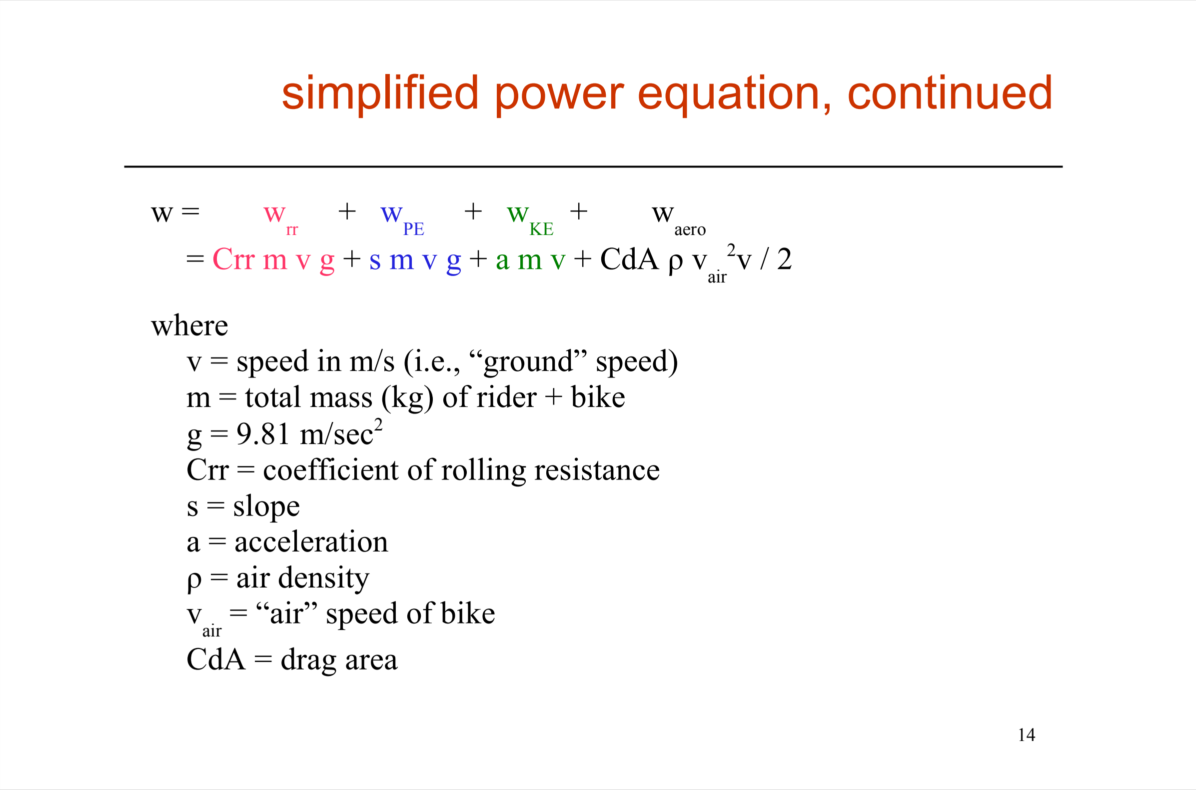

simplified power equation, continued

$$w = w_{rr} + w_{PE} + w_{KE} + w_{aero}$$

$$= Crr\, m\, v\, g + s\, m\, v\, g + a\, m\, v + CdA\, \rho\, v_{air}^2 v / 2$$

where

v = speed in m/s (i.e., “ground” speed)

m = total mass (kg) of rider + bike

$g = 9.81\ m/sec^2$

Crr = coefficient of rolling resistance

s = slope

a = acceleration

ρ = air density

v_{air} = “air” speed of bike

CdA = drag area

When there's no wind, ground speed=air speed. Then, if PE and KE =0, you can get regression approach.

One of the reasons this is a “simplified” model is because it ignores yaw angle, i.e., the angle at which the wind strikes the rider. CdA varies with yaw angle but if wind is low relative to rider speed, the yaw angle approaches zero. If there is no wind at all, the yaw angle is exactly zero.

Typically, you'd know your total mass and you'd record temperature, altitude, and humidity in order to calculate air density, ρ, so one way to solve the equation is to make an initial guess at either Crr or CdA and calculate the other.

Although it appears that we are assuming fixed Crr, a more precise statement is that we only need to assume that whatever changes in Crr occur, they do not change from lap to lap, e.g., the lap can have rough patches and smooth patches but they don't migrate randomly around the lap. We're assuming that each time you ride over the same spot the Crr in that spot will be the same as it was on previous laps.

Earlier I mentioned an algebraic approach. Integrate the power equation to convert it to units of work in joules. Then total work for *any* segment of a course is:

$J = J_{rr} + J_{PE} + J_{KE} + J_{aero}$

That's true for any segment. In particular it's true for specific segments of a course that correspond to laps, so the net elevation change is zero. Reformulating the power equation into a work equation lets us factor out Crr and CdA. Then, given a particular value of Crr, we can solve for CdA.

If we knew from external information what the actual change in elevation over a segment was, we could use that information to make estimates over that segment rather than whole laps. For example, suppose part of your lap included a hill for which you knew the altitudes at the bottom and the top. You could could use this method on that segment alone.

no wind approach

- assume $v_{air} = v$ and solve for slope as a function of other variables

  $s = w/(m\ g\ v) - Crr - a/g - (\rho\ CdA\ v^2)/(2\ m\ g)$

- use this formula to estimate point-by-point slopes from the data, supplemented by initial guesses at Crr and CdA. Ballpark guesses for starting values might be Crr = .005 and CdA = 0.3. v is in meters per second, so convert v = kph/3.6. A simple and not too terrible estimator for the acclerations, a, is the changes in v divided by 1.26 (these data were collected at 1.26-second intervals)

  standard approach assumes accelerations = 0. This approach calculates and uses them

- use estimated slopes to construct elevation change for each 1.26-second interval: elev.change = s*v*1.26

We'll assume wind speed is zero right now, but we'll see what effect it has later.

There are other ways to calculate the accelerations, a, but this method is relatively simple and reasonably close. If you're really careful with your data collection, youre probably the kind of person who'll want to use a slightly better estimate of acceleration, and to make a correction for the "exposure." One such estimate of the acceleration appropriate to the interval t is a = (v(t+1)–v(t-1))/ (2*Δt), and then you'd make "endpoint" adjustments for the first and last intervals. However, I'm indebted to Adam Haile for suggesting an even better alternative: a1 = ((v1-v0)/Δt)*((v1+v0)/2)/v1 (i.e., the "simple" acceleration term weighted by the ratio of average speed over the previous intervals to the speed of **this** interval). Adam's acceleration term simplifies to a1 = (v1^2 - v0^2)/(2*Δt*v1).

You'll generally know your mass and be conscientious enough to take the measurements you need to calculate ρ but for me on that day m=84kg and ρ=1.2 are close enough. You'll see later that for these data the estimate is relatively insensitive to m. Later, we'll re-examine the "relative" insensitivity to changes in mass.

produce an “elevation” profile

- cumulate the “elevation” changes and plot against distance to produce a “virtual elevation” profile
- make guesses at CdA until the plotted laps line up (or solve algebraically for the CdA that achieves that)

  these steps may seem daunting but they only take a few commands in any spreadsheet, or programming environment like R or Matlab

I put “elevation” in quotes because what we're really calculating is a “virtual” elevation. Basically, anything that is left unmodeled (like wind or using your brakes) is tossed into the slope, and we're using slope to calculate the elevation profile.

To get distance (in km) from speed (in m/s), cumulate v*1.26/1000

You may recognize that what we're doing is getting an elevation profile for the ride by integration of estimated elevation change over distance.

Note that this cumulative approach uses the raw and unsmoothed point-by-point slopes. Because of the limited resolution of the data (that is, the data are collected at 1.26-second intervals, and speed is recorded in tenths of km/h), you may think it's good to slightly smooth the calculated elevations. I have not bothered to do this because the calculated profiles are already smooth enough, and cumulating acts as a built-in smoother (you may recall from integral calculus that the integral of most functions is usually smoother than the function itself; and that the derivative of a function is usually rougher than the function itself).

Although I have implemented this method in a set of R functions, not everyone uses R. Others have created spreadsheets to do (some) of the analysis and they are far more accessible to most readers.. Here's a link to a spreadsheet created by Alex Simmons:

http://wattage.googlegroups.com/web/AeroTestVirtualElevation(Chung)Method.zip

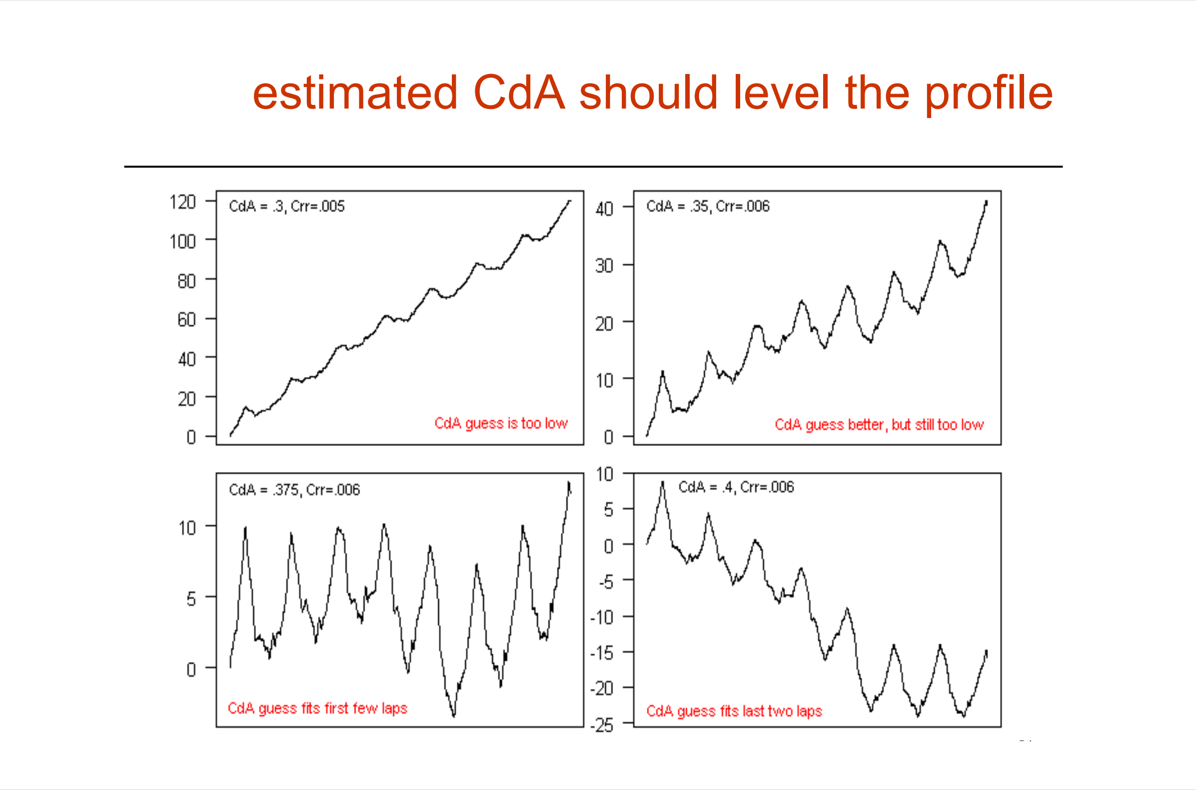

estimated CdA should level the profile

The upper left panel shows that the profiles creep upward, a clear symptom that the putative estimate of CdA is too low. Also note the y-axis vertical scale: it appears that over the course of what will soon be recognizable as seven-and-a-half laps, the total cumulative elevation gain was about 120 meters, or about 16 meters of net gain per lap. Nonetheless, you can already see the rough outlines of laps.

The upper right panel uses a higher estimate of CdA, and also of Crr (of which I will later say more). Now the number of laps is clear. The vertical scale has shrunk, indicating perhaps 5 meters of net gain per lap. Still, we need to increase the estimated CdA.

The lower left panel shows what appears to be a stationary pattern for the first three laps, i.e., zero net elevation gain over those laps. There appears to be a slight fall in lap amplitude, then a secular upward drift. This is a symptom that the CdA changed from the first few laps to the last couple of laps. As explained earlier, I sat more upright for the last couple of laps. This panel is the same image as the larger version you saw earlier.

The lower right panel raises the CdA once more, to a level where the last couple of laps appear stationary. This identifies a combination of CdA and Crr that could apply over the last part of the ride.

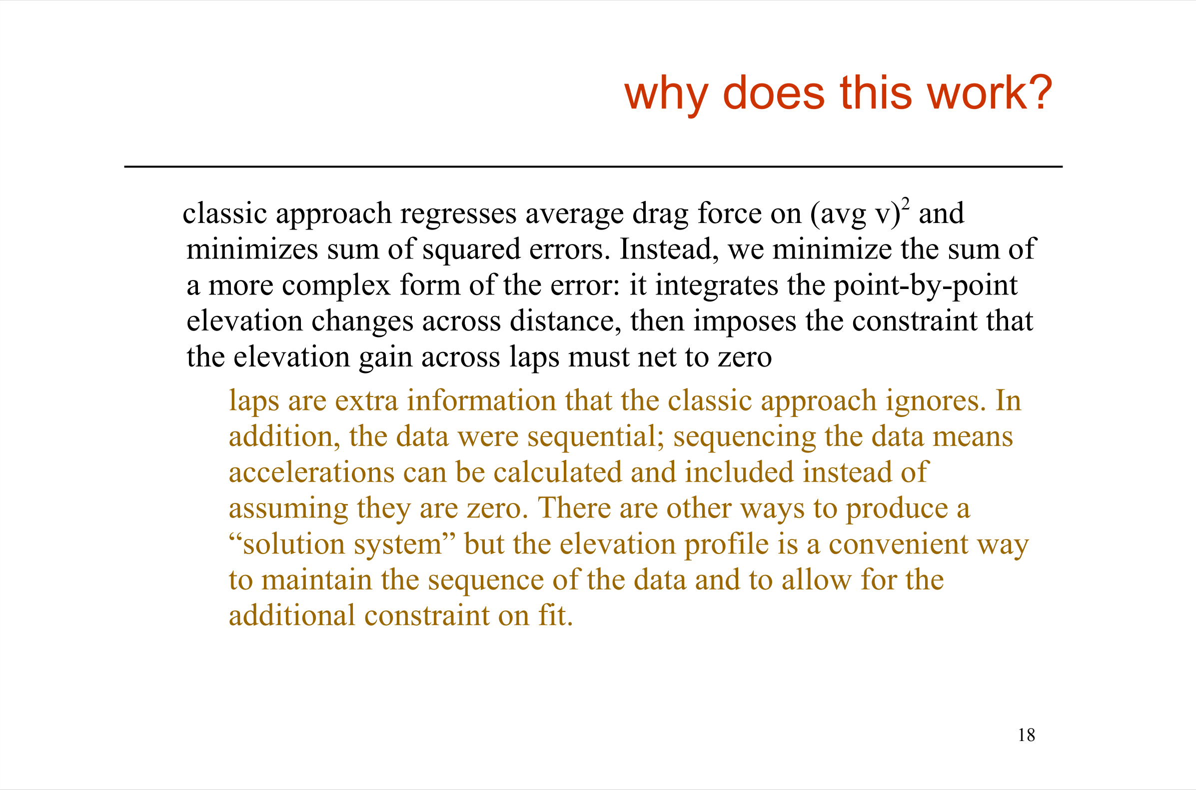

why does this work?

classic approach regresses average drag force on (avg v)^2 and minimizes sum of squared errors. Instead, we minimize the sum of a more complex form of the error: it integrates the point-by-point elevation changes across distance, then imposes the constraint that the elevation gain across laps must net to zero

- laps are extra information that the classic approach ignores. In addition, the data were sequential; sequencing the data means accelerations can be calculated and included instead of assuming they are zero. There are other ways to produce a “solution system” but the elevation profile is a convenient way to maintain the sequence of the data and to allow for the additional constraint on fit.

Minimizing squared error has certain very desirable properties from the point of view of statistical inference; however, if the method is not robust to error, statistical inference is unimportant.

In spirit, this is similar to the method of maximum likelihood; in this case, CdA is chosen to maximize the “likelihood” (loosely defined) of observing elevation profiles with zero net elevation gain from lap to lap.

As mentioned earlier, another way of thinking of this is as a generalization of a coastdown test. In a typical coastdown, you coast from a known speed down to another known speed on a surface of known slope. In that case, you're applying a known power: zero. In this case, you're doing a “coastdown” with known non-zero power, and using the recorded speed to tell you how quickly you're decelerating. See H.W. Schreuder (op. cit.) for a discussion of high precsion coastdowns.

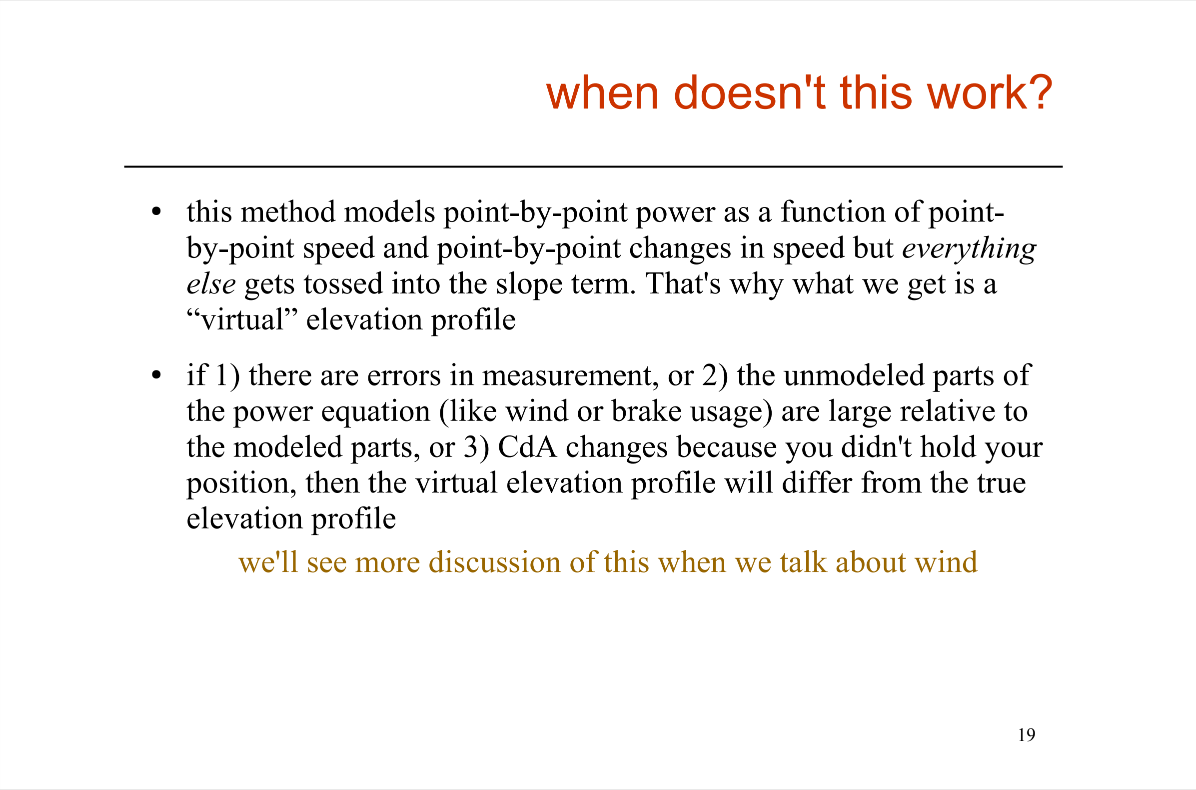

when doesn't this work?

- this method models point-by-point power as a function of point-by-point speed and point-by-point changes in speed but *everything else* gets tossed into the slope term. That's why what we get is a “virtual” elevation profile
- if 1) there are errors in measurement, or 2) the unmodeled parts of the power equation (like wind or brake usage) are large relative to the modeled parts, or 3) CdA changes because you didn't hold your position, then the virtual elevation profile will differ from the true elevation profile

we'll see more discussion of this when we talk about wind

Note that the flatter the course, the greater potential impact of unmeasured wind since then wind will be larger relative to the true elevation.

You want the errorsto be small relative to the modeled parts. In the usual approach, you tightly control speed, acceleration, and the slope and you choose windless days. In this approach, you don't have to control the speed and acceleration since they're measured well. However, you want a good spread of speeds and a reasonable amount of change in elevation to help “isolate” wind.

That is, if you know the true elevation profile this gives you a good way to assess how much the estimate was affected by unmeasured wind. This turns out to be useful: the usual approach is to wait for a wind-free day, to test on a flat (or constant slope) road, to hold speed constant (or, at least, to minimize changes in speed) but there is no simple way to tell if the measurements were tainted by wind, or changes in speed, or a small degree of slope.

If you ride laps, you can “overlay” them to see how similar the VE profiles are for each of the laps. If they're very different, you know it was too windy, or you didn't hold your position, or something else happened to the measurements.

Here's another way to think of it: we're trying to raise the “signal-to-noise” ratio. The classical approach to field testing tries to increase this ratio by decreasing the noise. Decreasing noise is always a good thing but another approach is to decrease noise and to increase the signal. This approach models accelerations and “sequences” the data in order to increase the signal, then re-casts the model in a way that lets us measure deviations from fit.

absolute and relative CdA

- in this example I made guesses about both CdA and Crr. Look at the equation—an increase of .001 in Crr looks like an increase in the slope of .001 (=0.1%)
- so *with these data*, we appear to have pretty good relative accuracy but unless we know what Crr is, we won't have good absolute accuracy
  - good relative accuracy means we can spot small *changes* in CdA even if (with these data) we can't nail down CdA itself. Sometimes you'll want do specific additional tests that will let you nail down both CdA and Crr
- I was bad and didn't measure air density (though I have a ballpark idea about what it was)

How small of a change can we spot? With these data, you can spot a change in estimated CdA of 1%. Using this method, one rider reported repeated estimates of CdA within .001 m^2; another added a 5cm x 5cm cardboard square (=.0025 m^2) to his bike and reported an estimated change in drag area of .003 m^2.

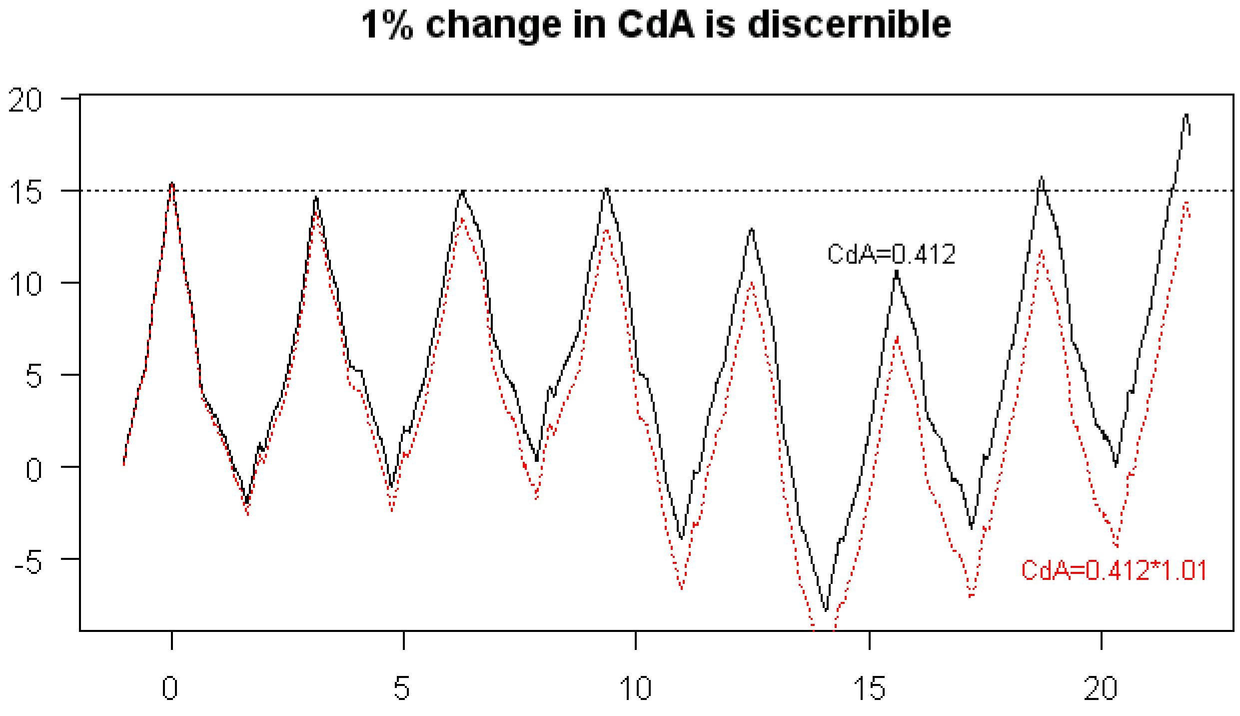
1% change in CdA is discernible

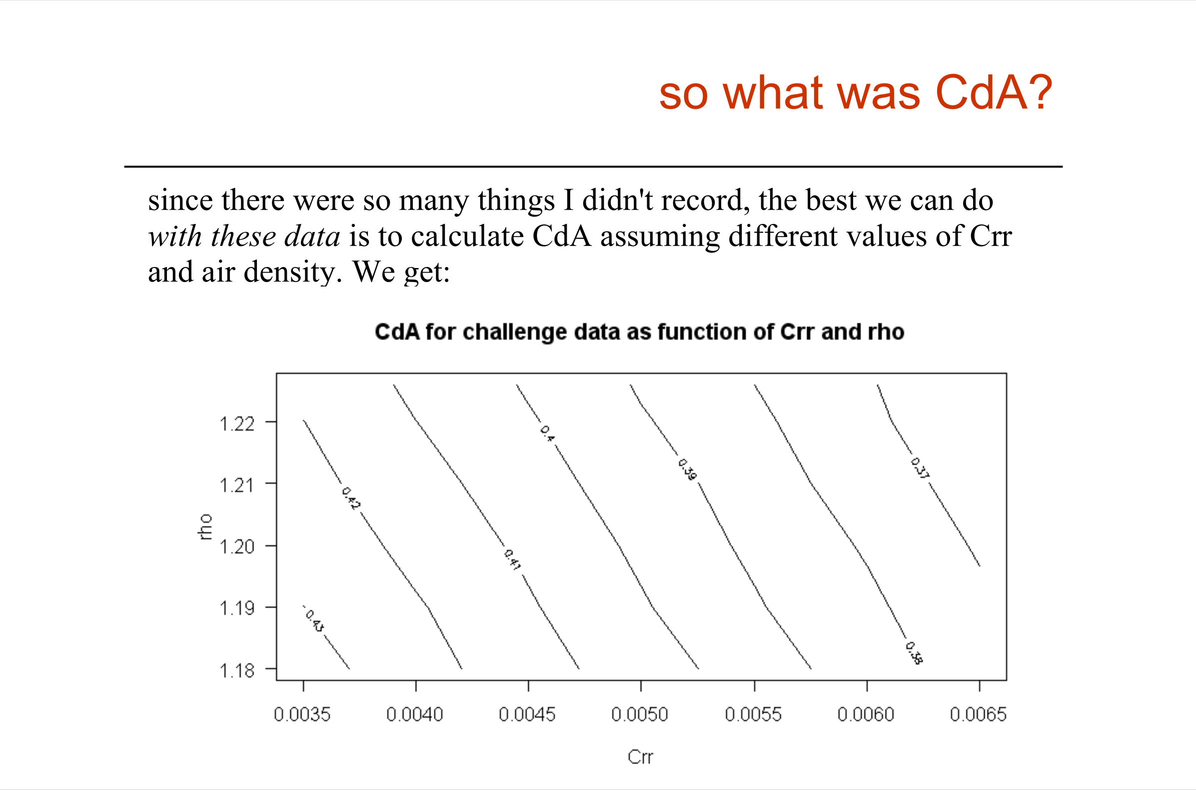

so what was CdA?

since there were so many things I didn't record, the best we can do *with these data* is to calculate CdA assuming different values of Crr and air density. We get:

Let's review: using only power and speed, we can show that the calculated profiles are relatively inelastic to mass. Most people doing field testing would at least make an attempt to measure air density but I didn't so the best we can do here is to produce an estimate for CdA that depends on Crr and air density.

For given Crr, increasing air density implies decreasing CdA, and a 1% change in air density implies around a 1.5% change in CdA.

For given air density, increasing Crr implies decreasing CdA, and a 1% change in Crr implies around a 0.3% change in CdA.

This may make it sound like Crr is less important than air density, but air density is easy to measure and it changes relatively slowly while Crr is hard to measure well, road surfaces can change quickly, and changing road surfaces can change Crr by much more than 1%. The bottom line is that although CdA is relatively less sensitive to changes in Crr than to change in air density, the magnitude of changes in Crr can be large so the overall effect is also large. Conversely, if you're off on air density by a little bit, it won't affect CdA that much. Bottom line, you should probably do your best to record air temperature and barometric pressure, but don't sweat too much about air density changing over the course of your runs.

Here's an important observation: for these data, the total elevation change doesn't appear to be that sensitive to changes in Crr. That's so for these data but it will turn out that this is not always the case; in fact, we'll exploit this difference later.

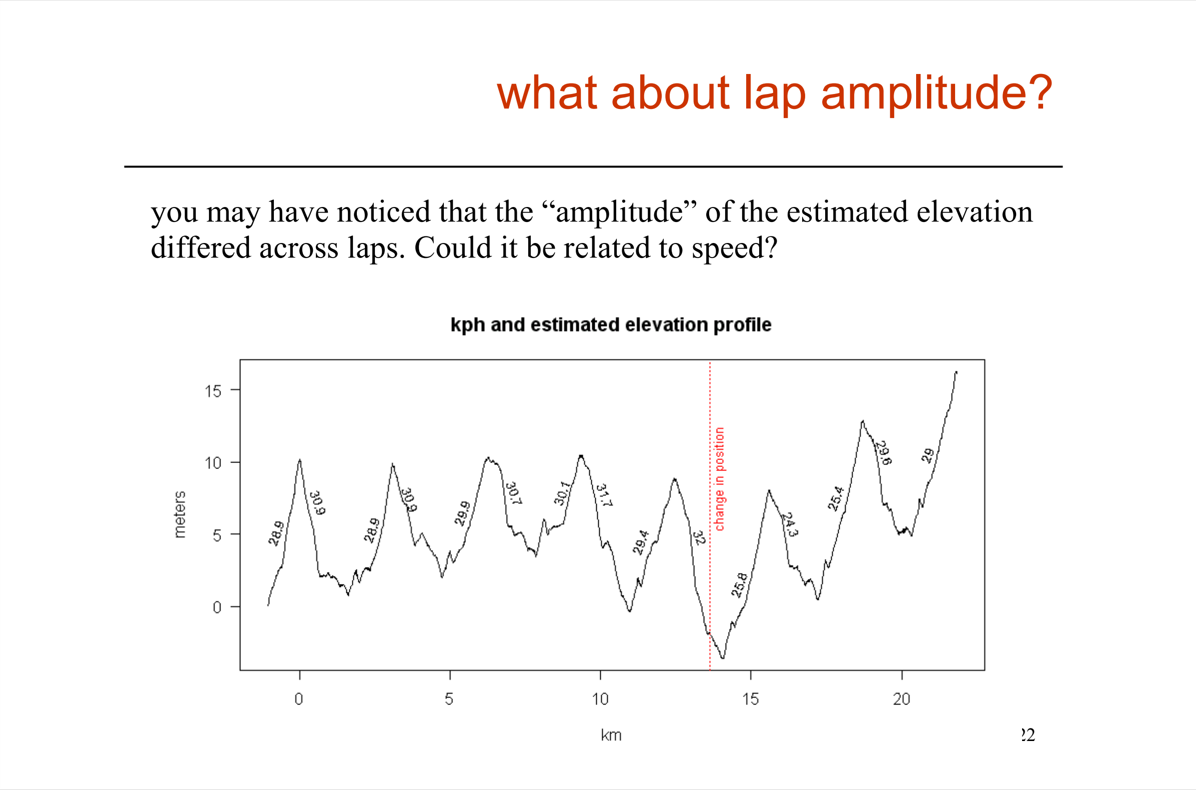

what about lap amplitude?

you may have noticed that the “amplitude” of the estimated elevation differed across laps. Could it be related to speed?

kph and estimated elevation profile

Being able to recover the elevation profile from the data lets us look at average speed over equivalent segments of the ride. Here I've split the ride into “downhill” and “uphill” portions and looked at the average speed over each segment. Note that I don't use the estimated CdA and Crr in this part of the calculation—we don't need them in order to identify lap lengths or the relative high and low points. Since we can identify matching parts of the laps by distance, we could have split the laps into arbitrarily many segments; I chose two for simplicity's sake.

ground speed and air speed

- there is rough evidence that for these data the elevation profiles are speed-dependent
  - increased speed in the downhill direction increased elevation change
  - increased speed in uphill direction decreased elevation change
- could it be unmeasured wind?
  - up to this point, we've assumed no wind (i.e., ground speed = air speed)
    - recall that the challenge included the information that there was an unmeasured amount of wind, but that I thought it was light and from a consistent direction

The profile we've estimated up to this point showed about a 10 m change in elevation over a distance of about 1.6 km, or an average slope of around 0.6%, so the “downhill” is not very downhill at all.

a handy diagnostic

now you can see why I start with a graphical approach: it provides a handy diagnostic for whether the model assumptions are met

- unmeasured variables affect the profiles in recognizable ways
  - unmeasured wind typically makes the profiles speed dependent
  - unmeasured braking typically appears as a sudden jump in the estimated elevation

Just because I often start with graphical analyses doesn't mean that's all I do. However, graphical approaches are often a good way to summarize models and are especially useful when they can tell you about model failures and poor fit.

Often on an out-and-back time trial on a straight course the only time you use the brakes is at the turnaround. The turnaround will be obvious in the virtual elevation profile because there is a sudden jump in the profile at that spot. Sometimes it's possible to “cut-and-splice”: cut out a little bit of the file surrounding the turnaround and splice the two half-profiles together.

And now you also can see why I started with this example: the wind was strong enough to affect the estimated CdA but this method makes it easy to see that this happened.

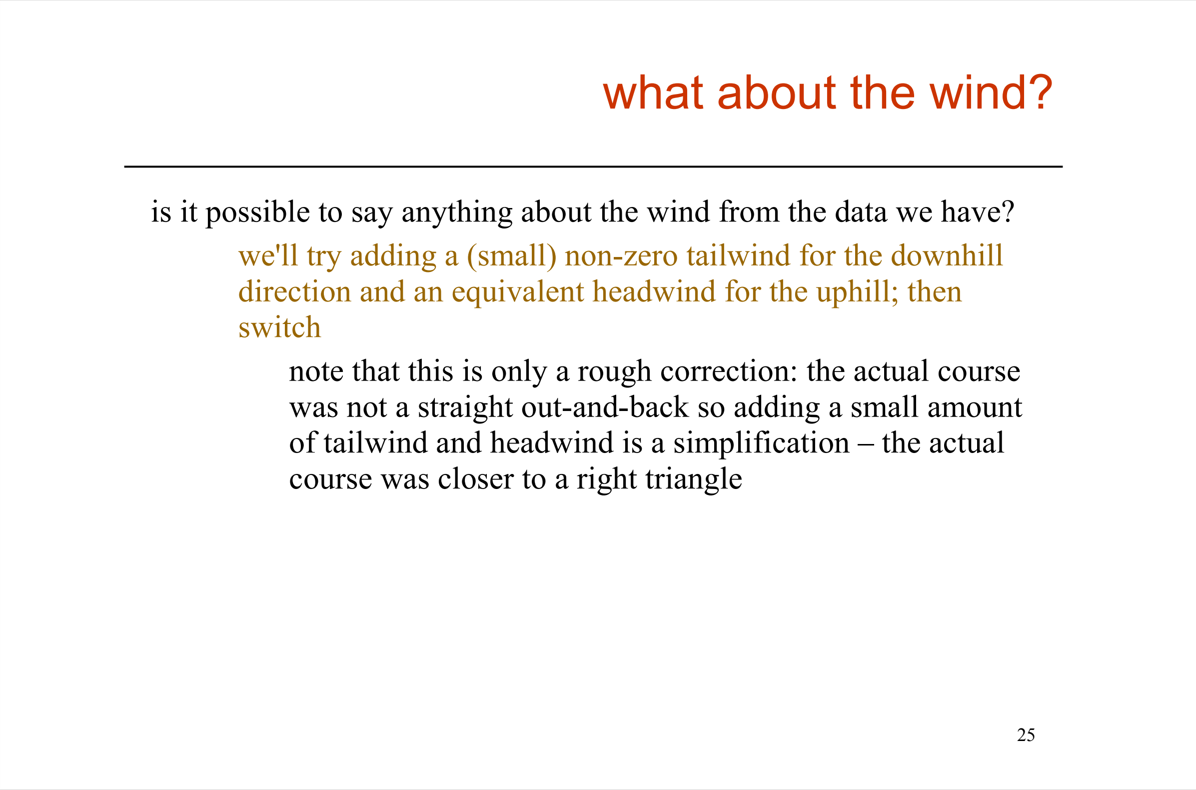

what about the wind?

is it possible to say anything about the wind from the data we have?

- we'll try adding a (small) non-zero tailwind for the downhill direction and an equivalent headwind for the uphill; then switch
  - note that this is only a rough correction: the actual course was not a straight out-and-back so adding a small amount of tailwind and headwind is a simplification – the actual course was closer to a right triangle

As previously mentioned, we could have split the laps into arbitrarily many segments, each with their own wind speed and direction. There is no particularly need to balance the wind in the downhill and uphill sections but we'll begin with this (obviously) simplified model. I gave a URL to the Google map for the course earlier.

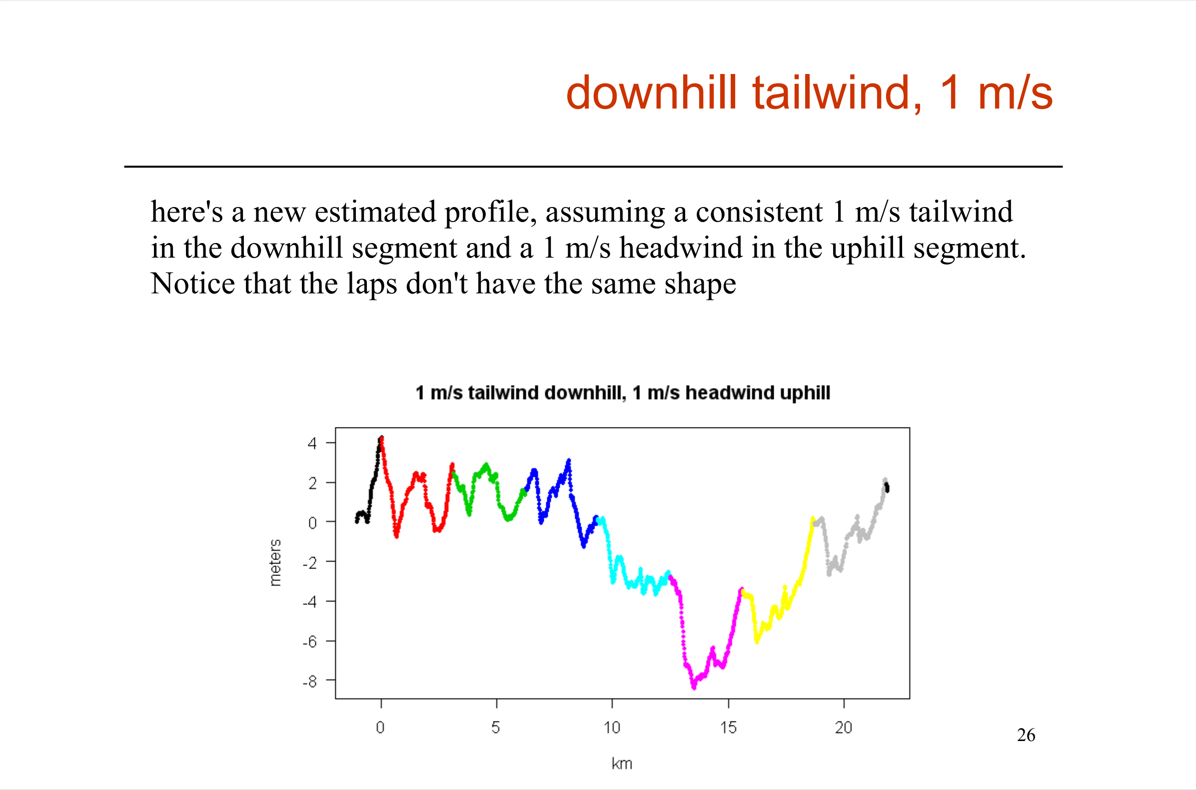

downhill tailwind, 1 m/s

here's a new estimated profile, assuming a consistent 1 m/s tailwind in the downhill segment and a 1 m/s headwind in the uphill segment. Notice that the laps don't have the same shape

I've used color so you can see the beginning and end of each lap. Notice that the total elevation gain across the first few laps is around 4 meters.

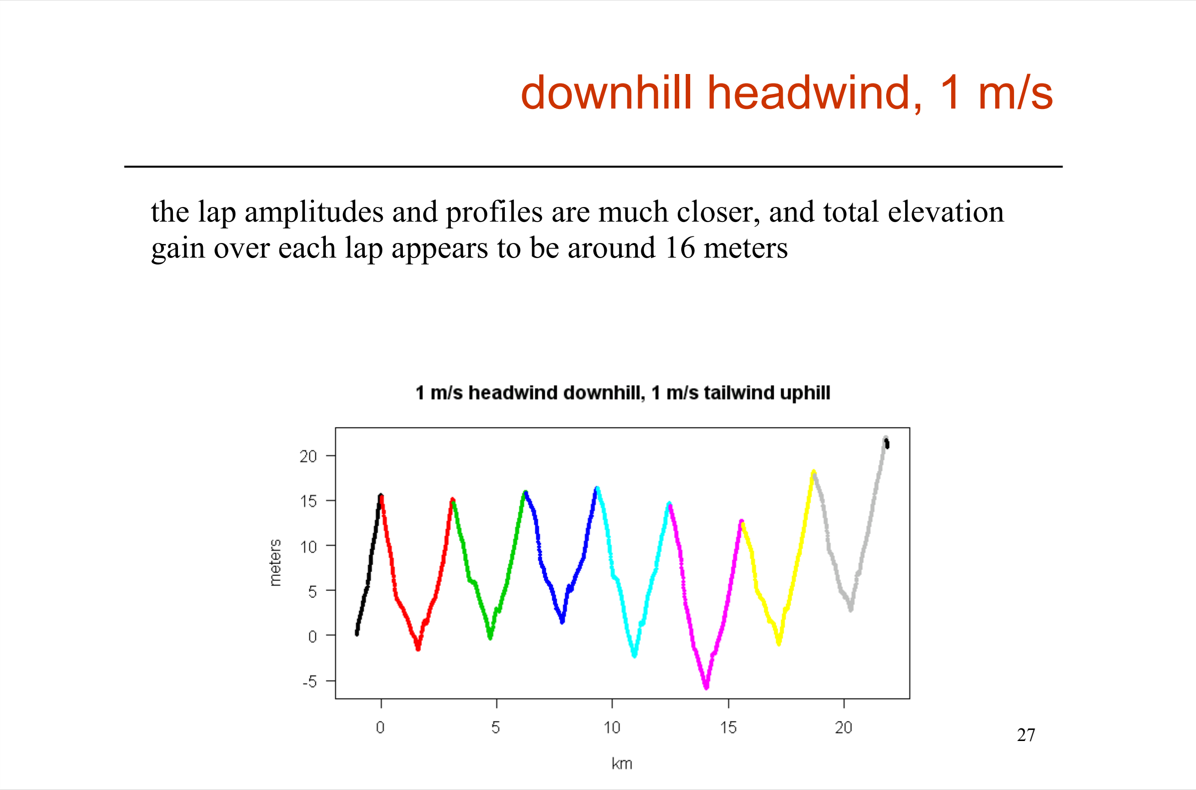

downhill headwind, 1 m/s

the lap amplitudes and profiles are much closer, and total elevation gain over each lap appears to be around 16 meters

Holy cow. The total elevation gained across each lap is almost exactly what was shown on the topo: 15 to 16 meters from high point to low.

Doubling the windspeed to 2 m/s increases the elevation gain but doesn't change the overall shape (much).

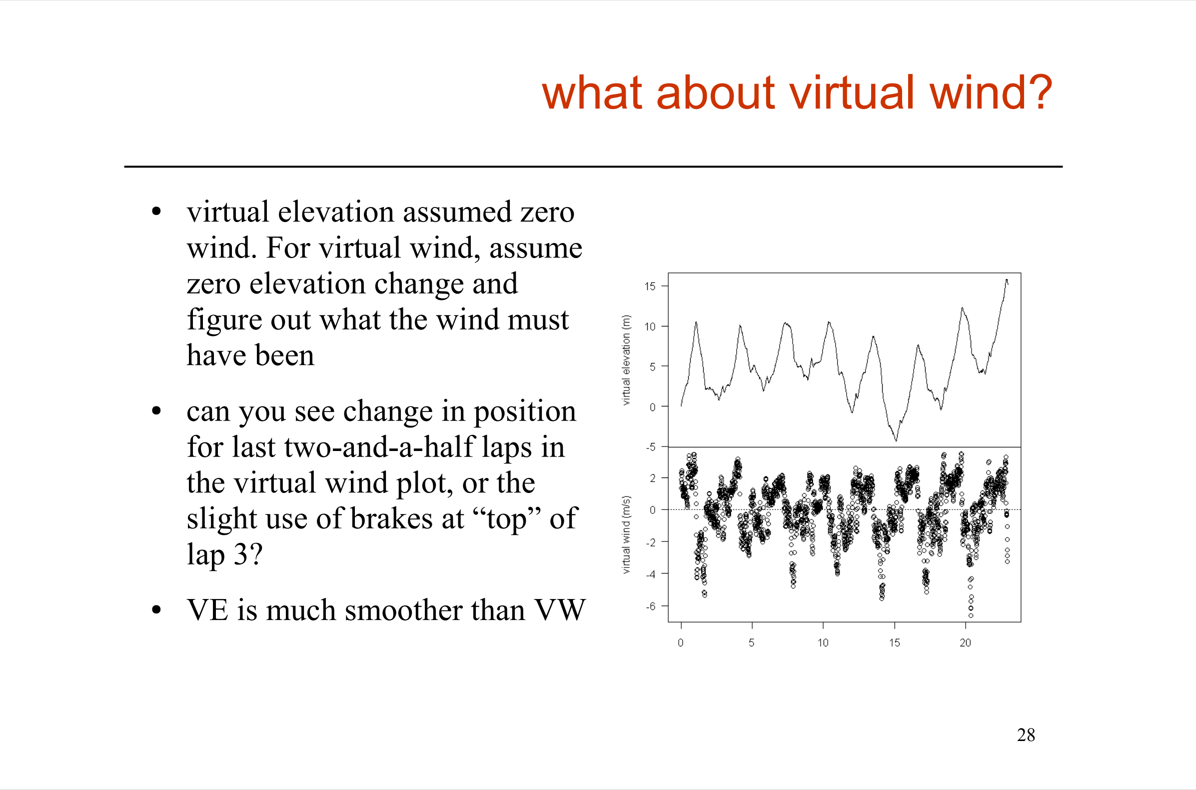

what about virtual wind?

- virtual elevation assumed zero wind. For virtual wind, assume zero elevation change and figure out what the wind must have been
- can you see change in position for last two-and-a-half laps in the virtual wind plot, or the slight use of brakes at "top" of lap 3?
- VE is much smoother than VW

For the virtual elevation calculation we assumed zero wind and figured out what the slope must have been. For the bottom panel, we do exactly the opposite: we assume the course was absolutely flat and figure out what the wind should have been. Of course, we know that the wind did blow and the course wasn't flat so our "true" profile and wind should be somewhere between the two.

If you plot virtual slope rather than virtual elevation (we calculated virtual slope in order to get virtual elevation), you'll see it looks very similar to virtual wind. Integrating slope over distance to get the elevation profile is the key – and integrating wind over distance doesn't have the same easy interpretation because while we know that once we get back to the start of a lap the elevation nets to zero, no such constraint applies to the wind.

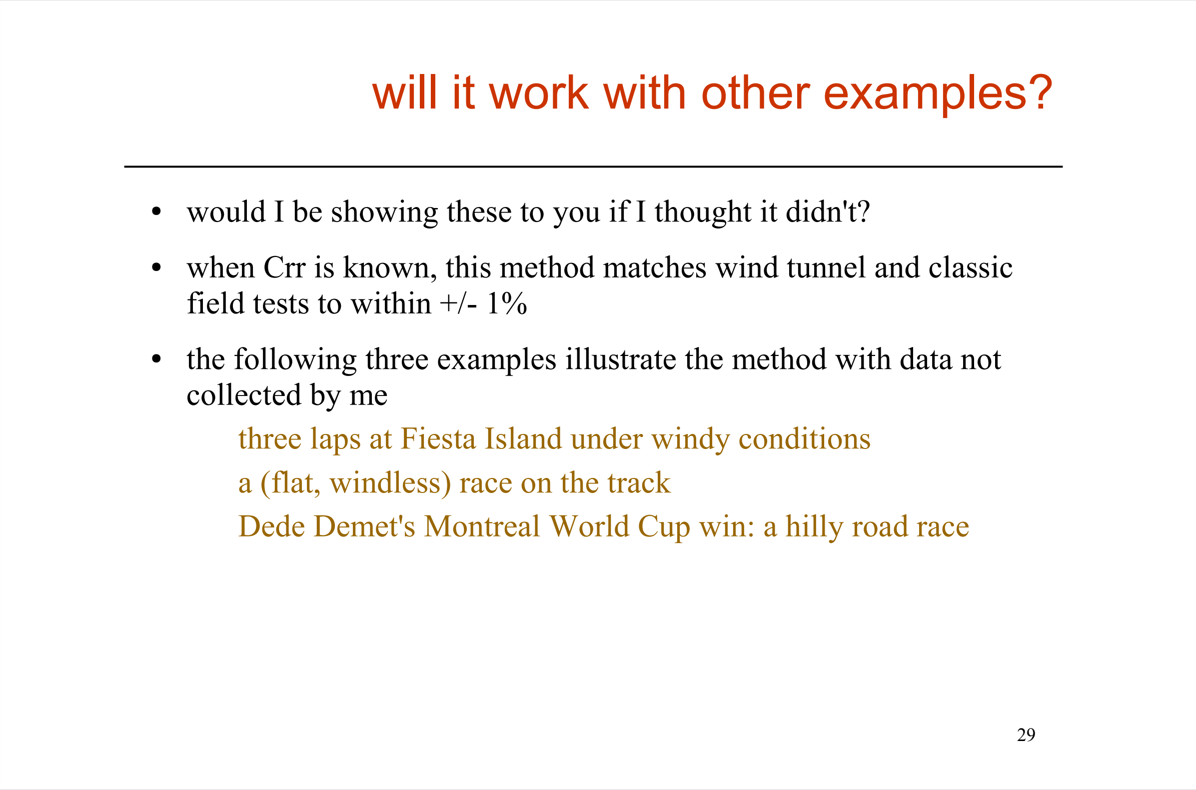

will it work with other examples?

- would I be showing these to you if I thought it didn't?
- when Crr is known, this method matches wind tunnel and classic field tests to within +/- 1%
- the following three examples illustrate the method with data not collected by me
  - three laps at Fiesta Island under windy conditions
  - a (flat, windless) race on the track
  - Dede Demet's Montreal World Cup win: a hilly road race

Before we move on to the other examples, here's a little bit of information about the conditions under which the challenge data were collected: I don't know the exact total mass but 84 or 85 kg is probably pretty close. The temperature was in the upper 50's or low 60's (F, or around 15 to 17 C), it rained that evening, and you know from the description of the course that this was in Paris, which is about 60m above sea level. A reasonable estimate of rho is in the range of 1.2 or 1.21. Most of the course is pretty smooth asphalt, though there are some slightly rougher spots. I was wearing arm warmers, leg warmers, and a wind vest, and I wasn't trying at all to be aero.

These data were not included with the challenge so I haven't discussed them. However, if you do use these data, along with a small amount of wind and an estimate of Crr of .0035, I get an estimated CdA for the first few laps of 0.415, and about .436 (5% higher) for the last two laps. It appears that with these data the method can discern a change in CdA smaller than that, down to 1%.

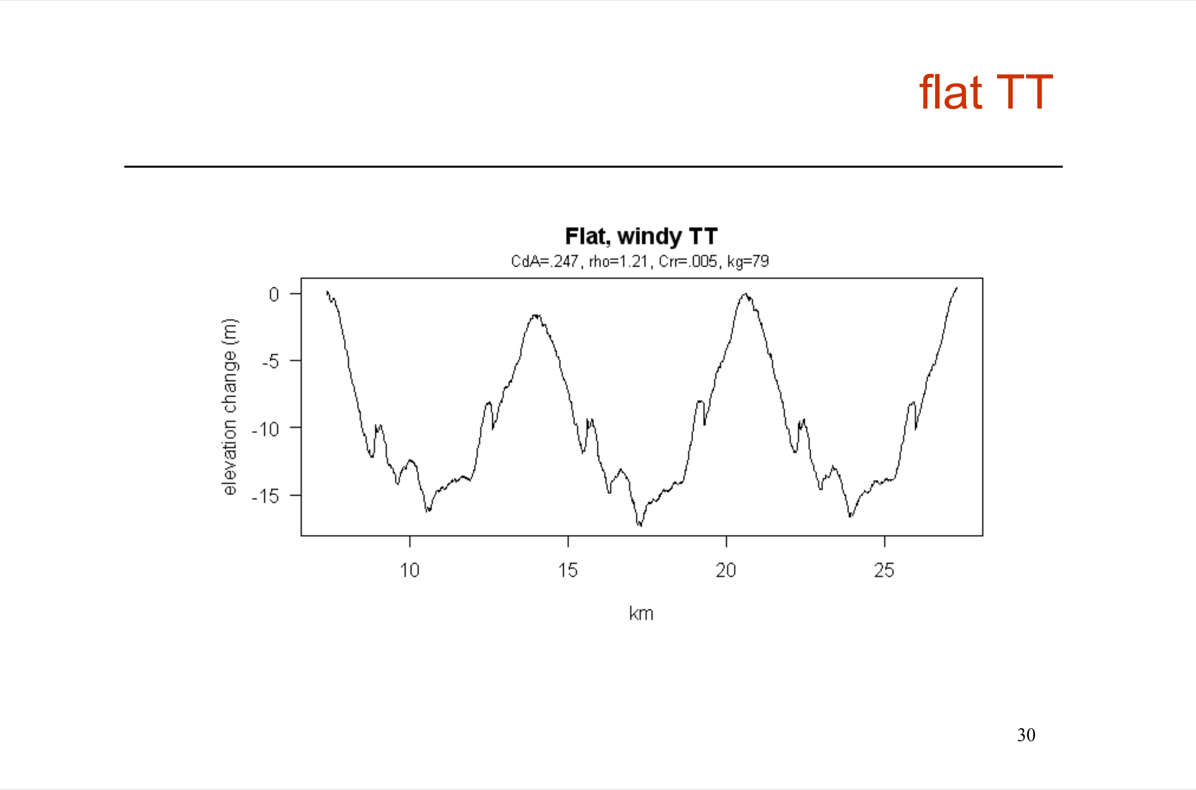

flat TT

The data were collected by Kraig Willett. Kraig posted these data on biketechreview.com with no other information than that the ride data were collected over three laps at Fiesta Island. The course is flat; reportedly, the actual elevation change is no more than a meter. Therefore, rather than a virtual elevation profile, you can interpret this as a virtual wind profile. The profiles across laps are quite consistent, suggesting that the wind was quite steady. The CdA estimate appears quite low. In addition, it appears that he started off in the downwind direction. Subsequent to my analysis, Kraig said that the wind was blowing 7 or 8 mph, he was in aerobars and using 140mm BMX cranks, and made a guess at his weight that day.

track race

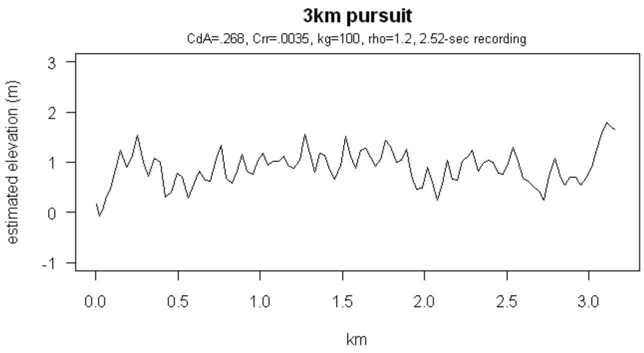

The track race was a 3km individual pursuit. The data were collected by Bob Schwartz and posted on rec.bicycles.racing. The method shows 24 bumps or dips in elevation over the race, with most of the bumps appearing to fall within a one meter band (except at the very beginning and very end of the race). The dips are equally-spaced and represent the turns – this suggests that the race was held on a 250m track. After looking at the graph, I asked Bob about the track: it was the 250m outdoor Blaine/NSC track in Minnesota. Bob says that although he doesn't remember the exact total mass that day, he believes my estimate was low – Bob is a big guy. Fortunately, adding 10 kg to total mass has almost no effect on the estimated CdA, as we might expect for a flat race. Perhaps surprisingly, an increase in Crr of 33% decreases the estimate of CdA by only 2.5%. A 1% increase in rho decreases the estimate for CdA by about 1%, as we would expect. The data were collected at 2.52-second intervals, as for the data in the next example. Because of the large initial acceleration from a standing start, I used the slightly more complicated acceleration term with an adjustment for the first record.

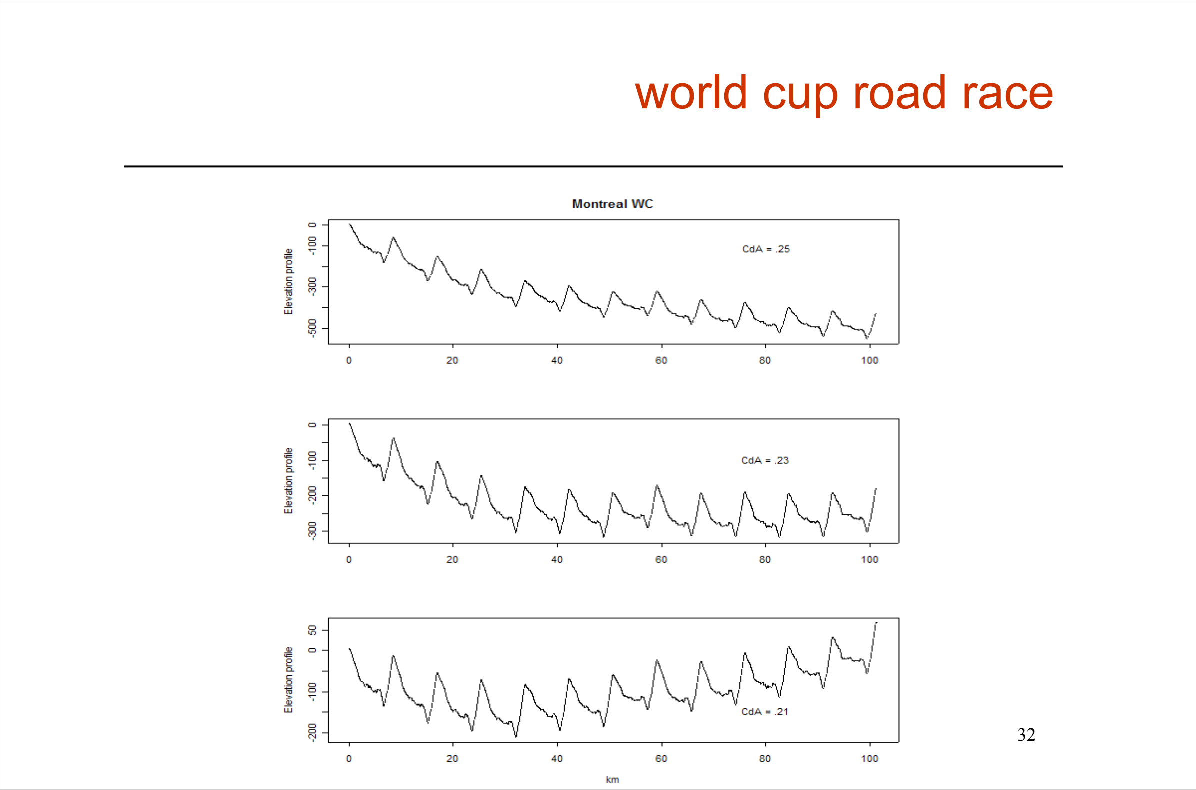

world cup road race

Dede Demet's data file used to be posted on the Power Tap web site. I don't know her weight, the wind conditions, or how she used her brakes but we can still pick up the 12 laps of the race. This data set was the first I analyzed using this approach, in April 2003. Note the low estimated CdA, and the two different “clearing” values for the early part and the late part of the race. This was a hilly race and her mass matters much more to the CdA estimate – and I don't know what it is. However, it appears to have been low enough to suggest that she was sheltered in the peloton.

Demet's Power Tap was set on 2.52-second recording that day. In that mode the PT records only every other 1.26-second record. Accelerations were estimated across the 2.52-second records as if they were complete. As a check, I have deleted every other 1.26-second record from other data sets – the results appear robust for the data sets I've tried. This means that the short cut I used to estimate acceleration isn't too terribly critical. In fact, note that the calculations I've used for this method are relatively crude. The method appears to be moderately robust.

virtual altimeter

the next two examples show that the method appears to be both reasonably accurate and consistent

- up and down Palomar Mountain
  - black shows altitude as reported by an altimeter, red shows estimated altitude, emphasizing that braking is ignored
- two years of the San Bruno Hill Climb
  - elevation profiles calculated from power and speed files for the same rider in two consecutive years, showing that the method captures consistent features of the profile

Both of these were previously presented at
http://anonymous.coward.free.fr/wattage/altimeter/altimeter.html

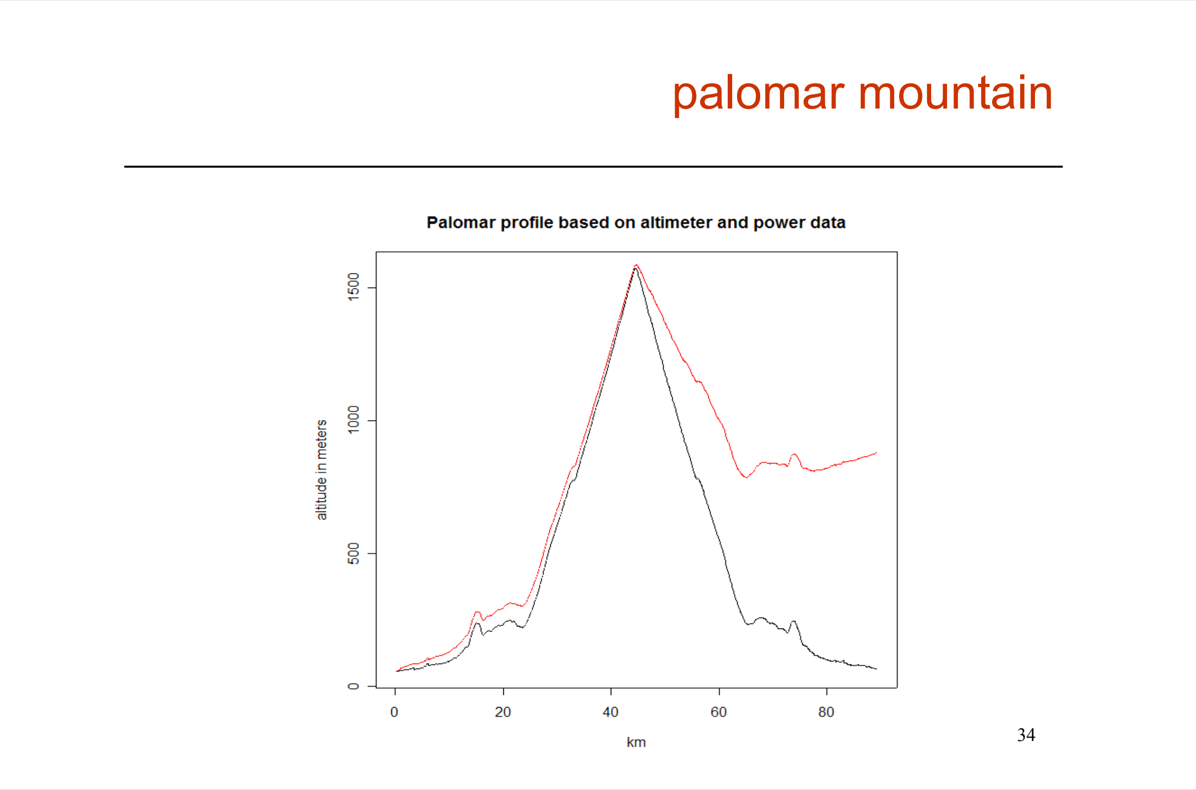

palomar mountain

Palomar profile based on altimeter and power data

The data were collected by Kraig Willett for his “three PM shoot-out.” The altimeter was the one on the Polar S710. The red line shows that the method works best when the brakes are not used, i..e., when power is well-modeled by the assumed speed-drag equation.

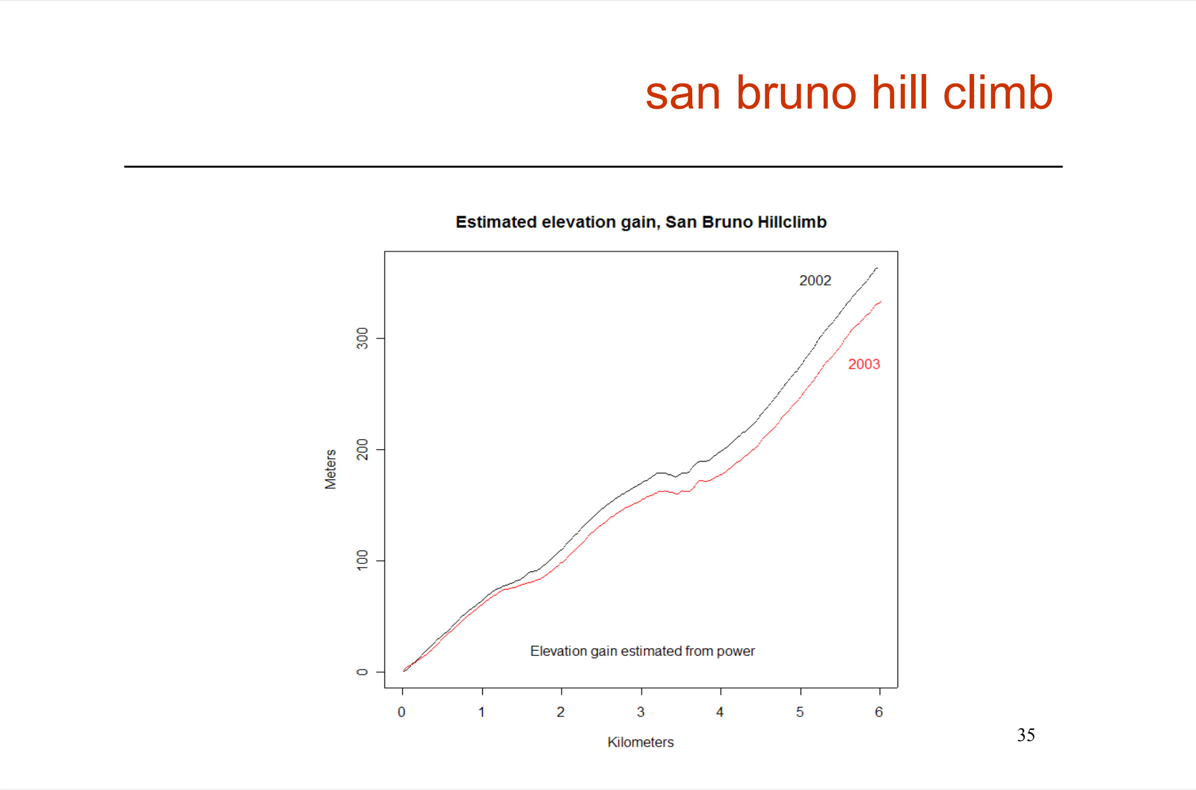

san bruno hill climb

Estimated elevation gain, San Bruno Hillclimb

The data were collected by Gary Gellin during the 2002 and 2003 New Year's Day races up San Bruno Mountain, south of San Francisco, and posted on the Wattage list. Gary reported that his total mass was the same both years.

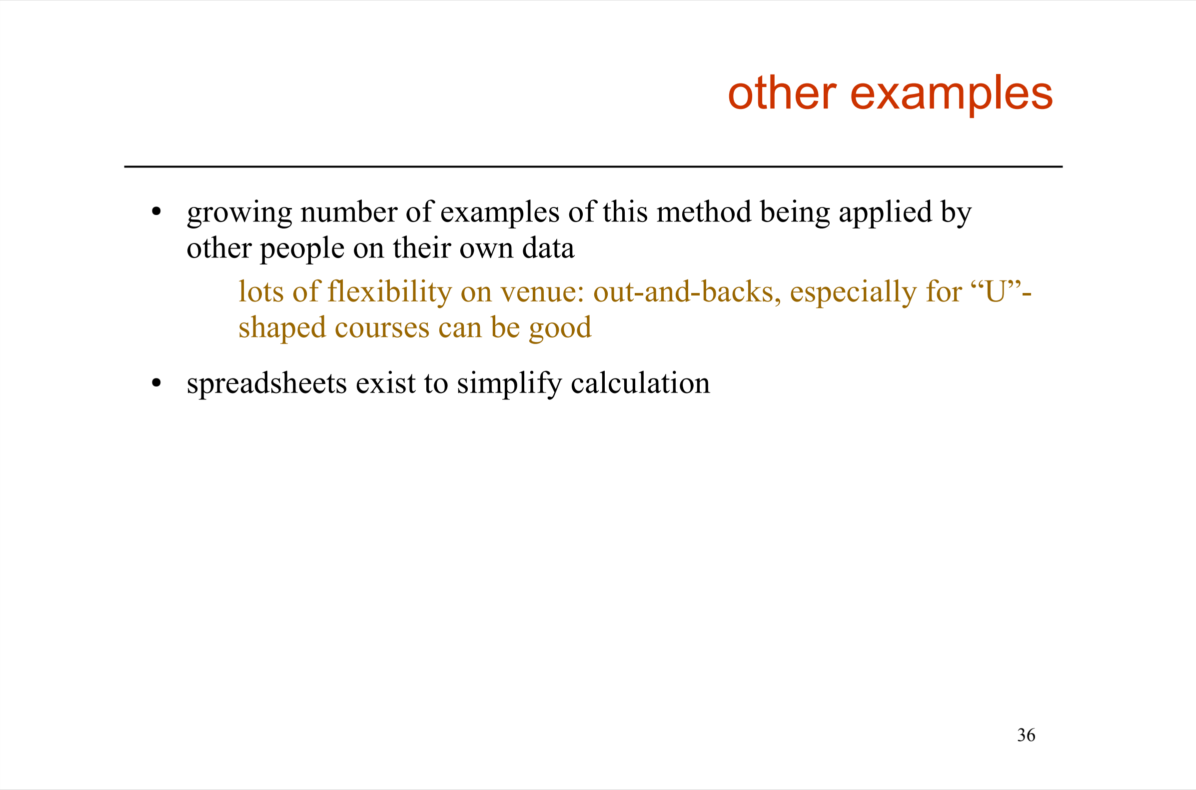

other examples

- growing number of examples of this method being applied by other people on their own data
  - lots of flexibility on venue: out-and-backs, especially for “U”-shaped courses can be good
- spreadsheets exist to simplify calculation

Here are a couple of links that show other examples:

http://alex-cycle.blogspot.com/2008/03/funky-things-with-power-meter-77.html

http://colinsbikingbits.blogspot.com/2009/09/time-trial-tri-bar-height-chung-test-3.html

http://forum.slowtwitch.com/cgi-bin/gforum.cgi?do=post_view_flat;post=1802183;page=1;mh=-1;

http://jasperga.blogspot.com/2009/11/chung-method-is-no-joke.html

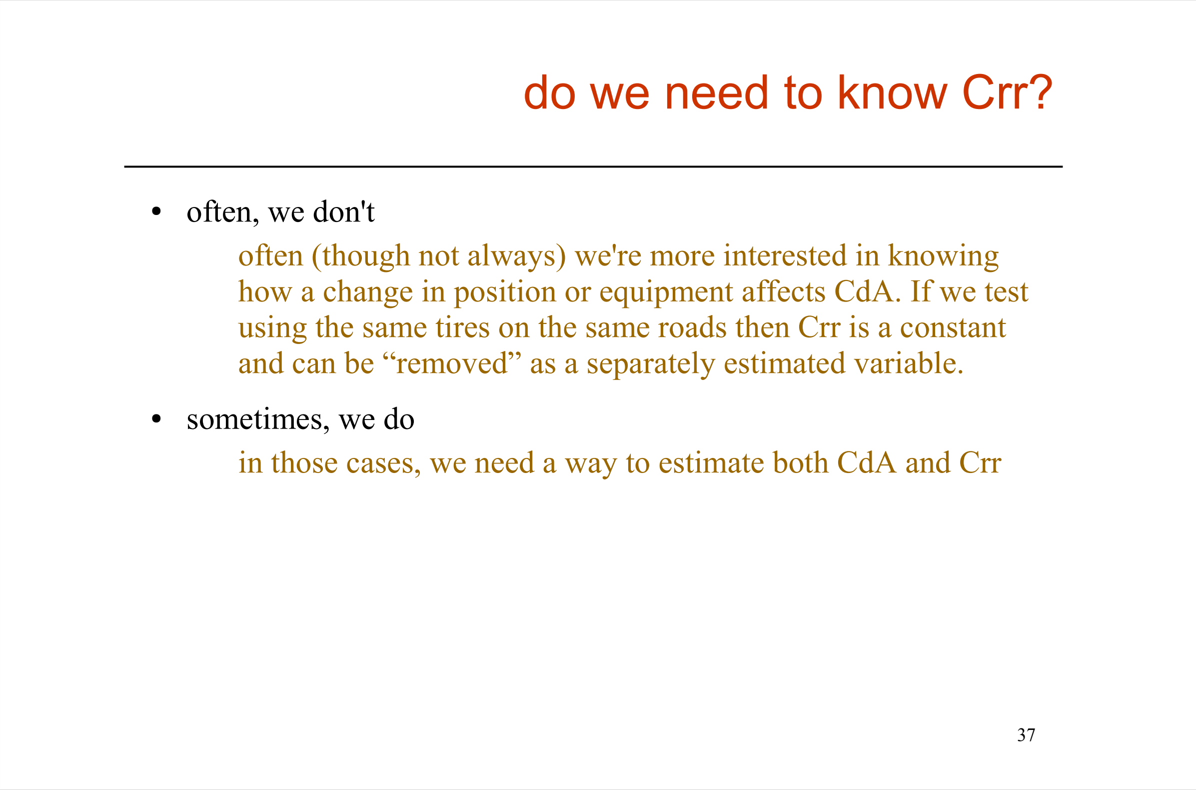

do we need to know Crr?

- often, we don't

    often (though not always) we're more interested in knowing how a change in position or equipment affects CdA. If we test using the same tires on the same roads then Crr is a constant and can be “removed” as a separately estimated variable.

- sometimes, we do

    in those cases, we need a way to estimate both CdA and Crr

Because this method can be less time-consuming than the classic approach, the need to get a separate estimate of Crr is reduced. That is, the classic field test approach lets you make separate estimates of Crr and CdA but it requires multiple passes down the same course at a wide range of speeds. This approach will let you estimate *changes* in CdA for a given Crr (or changes in Crr for a given CdA) much more quickly, especially if the laps are short.

Shorter laps have other advantages, too: they reduce the amount of time you're in the field so you less exposed to changes in conditions like changes in wind or passing cars.

prying apart CdA and Crr

- we've been starting with a guess at Crr. Is it possible to estimate Crr separately?
  - sometimes, yes
- first, understand why we've needed to do this
  - the usual field test approach requires flat roads (or constant grade). This approach allows you to use roads that aren't flat so we've lost a constraint. We need some extra information to compensate.
    - one possibility: make known changes and repeat
    - another: know the true elevation profile

There are several methods that allow you to separately estimate Crr and CdA, depending on the data collected. Although I discuss one way here, if I were really interested in getting precise estimates of Crr I'd follow the protocol laid out by Tom Anhalt and used by Al Morrison for tire testing using rollers and a power meter. You can find the results of Al's tests at biketechreview.com

extra information demands extra care

- requires very careful test runs
  - test on windless day
  - knowing true elevation profile helps a lot
- two examples
  - same hill, different speeds
  - flat course, monotonic speed change (aka the Shen method)

same hill, different speeds

- I rode up the same hill twice: once slow and once fast(er)
    - first ~ 170 watts, second ~ 245 watts
    - from topo maps climb known to be 114 meters
    - (this time) checked weight and air density
    - almost no wind
- thus far, we have only been solving for (Crr, CdA) pairs
    - if Crr overestimated then calculated CdA will be too low
    - if Crr underestimated then calculated CdA will be too high
    - more importantly, each (Crr, CdA) pair implies a different amount of total climbing
- knowing total elevation gain adds another constraint

The key is that a particular combination of Crr and CdA imply a certain range of elevation gain, i.e., in this case, we're not just “leveling” the lap profiles, we're also trying to match the total amplitude of the profiles.

You can find the data file here:

http://anonymous.coward.free.fr/wattage/cda/skylinefieldtest.csv

As an aside, this ride wasn't intended as a data collection ride; it was only after I'd started the ride that I realized the conditions were essentially windless and I could use the ride for data collection. Intervals 1 and 3 are the climbing portions; interval 2 is the descent but the route is twisty enough that I had to use the brakes. The temperature was in the low 40's F, I was wearing tights and a winter jersey, and I estimated the weather conditions when I got back home after the ride.

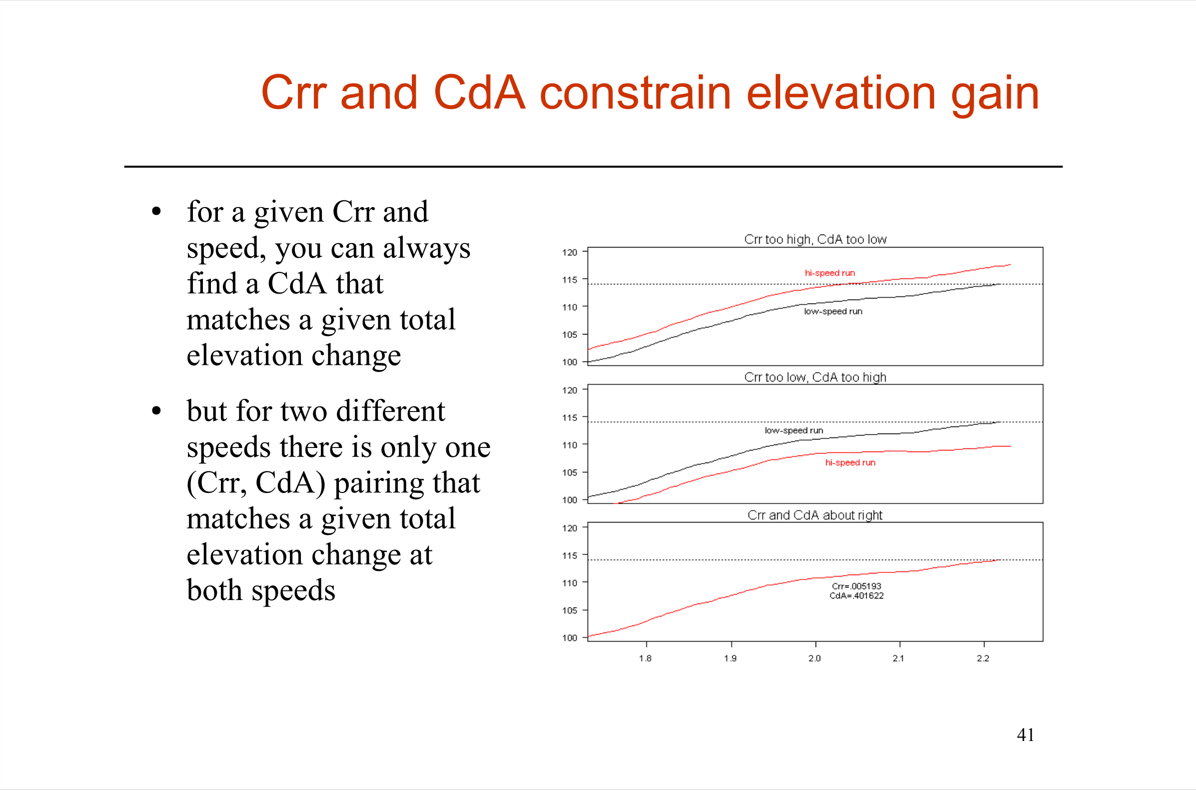

Crr and CdA constrain elevation gain

- for a given Crr and speed, you can always find a CdA that matches a given total elevation change
- but for two different speeds there is only one (Crr, CdA) pairing that matches a given total elevation change at both speeds

41

We can always find a “leveling” CdA for a given Crr but if the Crr is too low, the CdA will be too high and the VE profile will differ from the true lap profile; if the Crr is too high, the CdA will be too low and the VE profile will differ from the true lap profile *but in the opposite direction.* That's what we see here: in the top panel the red line (the higher power run) is above the black lin; in the middle panel the red line is below the black line. Only one (Crr, CdA) combination produced the same total 114 meter elevation gain for both runs; that's what's shown in the bottom panel.

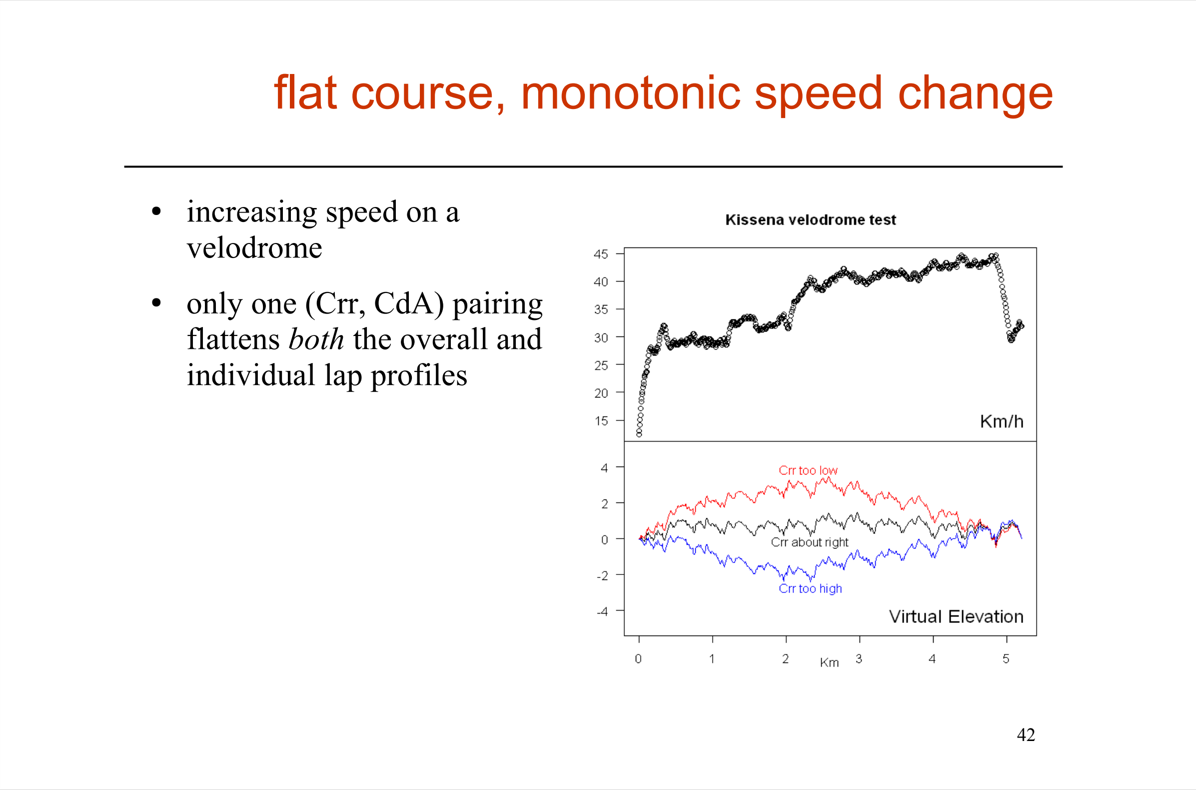

flat course, monotonic speed change

- increasing speed on a velodrome
- only one (Crr, CdA) pairing flattens *both* the overall and individual lap profiles

Data collected by Andy Shen, and discussed here:

http://velocitynation.com/content/coachingfitness/2009/shen-method

can we generalize?

- Crr and CdA constrain total elevation gain – but they also constrain elevation gain over any segment
- if we know true elevation profile over the entire course we can fit to arbitrary segments
  - this can come in handy for velodrome laps since we know the true profile

summary

- knowing only speed and power still provides an impressive amount of information when data are collected over laps
  - with these data, small changes in CdA are estimable
  - it's possible to examine how these estimates are affected by air density, wind speed, and wind direction
- knowing speed, power, and a little about the course provides even more information
  - you can tune the model not only to line up the profiles but also to match total elevation gain
- in some cases, knowing a bit more info can help you to get separate estimates of Crr and CdA

main conclusion

up to now, if you had an on-bike power meter, most field test protocols required that you find a flat and windless venue. Using this method greatly expands the list of appropriate field test sites.

- a loop like you might use for an industrial park crit could work
- I've used this approach on a long residential block with a small dip in the middle where it crossed a creek
  - ideal venue could be a bowl-shaped street that lets you speed up and still slow down at the ends to make the turnaround
- almost any loop where you don't have to use your brakes could work
- a single out-and-back up a slight hill could work

What's next? I've done some stuff with accuracy, precision, and a little bit of validation. In addition, I need to explain the analytical solution approach. That will have to wait.

recommendations

- follow the recommendations in Martin et al. (2006a, 2006b) if you can
- if you can't follow Martin, do laps
  - shorter laps let you do more of them
- don't hold speed constant
  - the wider the spread across laps the easier it is to isolate separate effects
    - a small amount of elevation change can help increase speed variation as long as it's not so steep you need to brake
- measure air density, don't use your brakes, and if you need precision don't do this on a windy day