

From: John Mount jmount@win-vector.com
Subject: Re: Car rating data/model
Date: January 9, 2015 at 4:16 PM
To: John Mount jmount@win-vector.com
Bcc: John Mount jmount@win-vector.com



Hi John,

That looks good. I think I can reproduce your result in R.

To do this:

- 1) Load the data into R: (type) `cars <- read.table('http://www.win-vector.com/dfiles/car.data.csv',header=TRUE,sep=',')`
- 2) Take a look at the data: (type) `View(cars); summary(cars)`
- 3) Build the safety/rating pivot table: (type) `with(cars,table(safety,rating))`

I thought about this some more, and I also noticed no two-seater gets an acceptable rating. In fact I think we can recover a pretty accurate decision diagram by doing the following:

- 1) (paste):
`library(rpart); library(rpart.plot)`
`model <- rpart(rating ~ buying + maint + doors + persons + lug_boot + safety, data=cars,`
`control=rpart.control(maxdepth=6))`
`rpart.plot(model,extra=4)`
`levels(cars$rating)`

What this tells us is that there are four ratings in the magazine (in alphabetic order): "acc" "good" "unacc" "vgood" and a procedure for mapping measured car features to ratings is given by the following diagram (which we could print and use. For example: a car that seats 4 people, medium safety, high price and low maintenance would be should be rated unacceptable. Each node on the diagram is showing in the data we know about what proportion of the examples in that part of the diagram/tree have a given rating. For example the node ".83 .00 .07 .10" is 83% acceptable, 0% good, 7% unacceptable, and 10% very good. So our procedure for that leaf node would be to say "acceptable" and we would expect to be right about 83% of the time (at least on the original data).

I think we could even circulate the diagram (with some training) to our representatives and they could make the prediction on paper.



