Statistical Thinking Comment Policy

Comments are welcomed. Be informed, informative, respectful, and criticize ideas, not people. Please read our Comment Policy before commenting.

Comments on this entire site are premoderated (only moderators can see this message). Change site settings.

10 Comments Statistical Thinking Disqus' Privacy Policy Frank Harrell

Favorite 1 Tweet Share Sort by Best

Join the discussion…

Ewout_Rotterdam • 4 years ago

This is a great post. I'm borrowing from it for a presentation in Stresa, Italy, for cardiologists, specifically re the concept of 'risk magnification'. Question: should we relabel this as 'benefit magnification'? That may be easier to link to the impact of treatment, which has benefit that depends on risk.
We may also want to relabel 'Absolute Risk Reduction' to 'Absolute Benefit' in the graph.

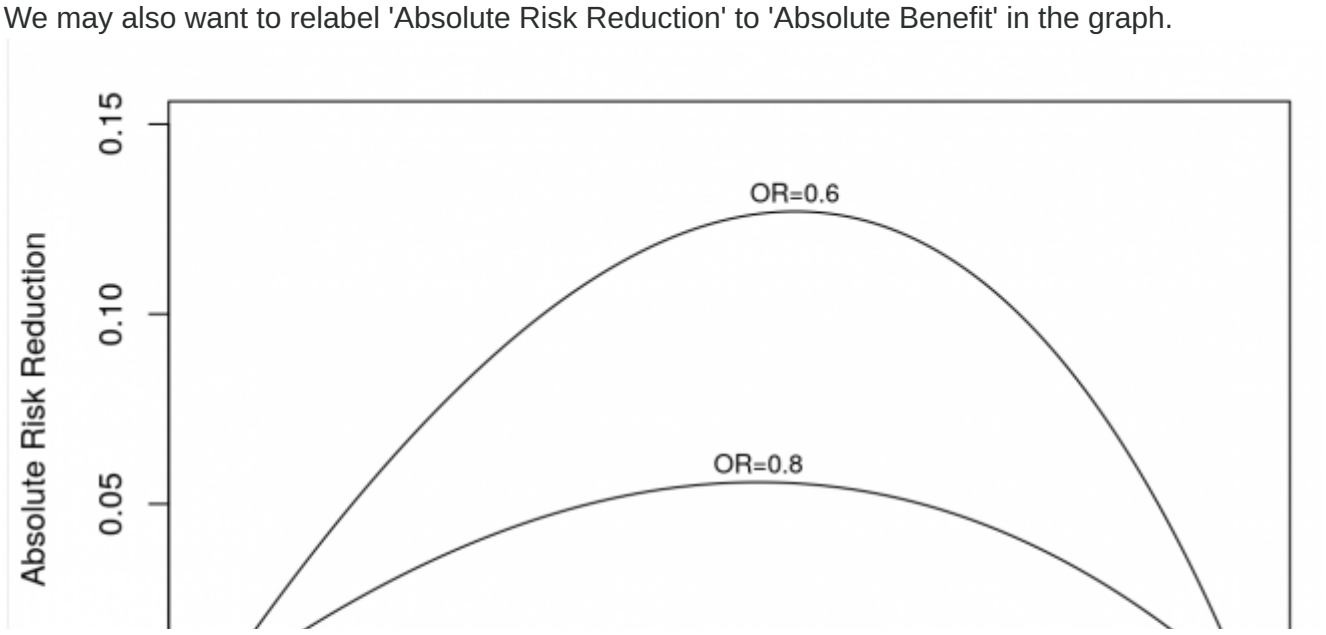

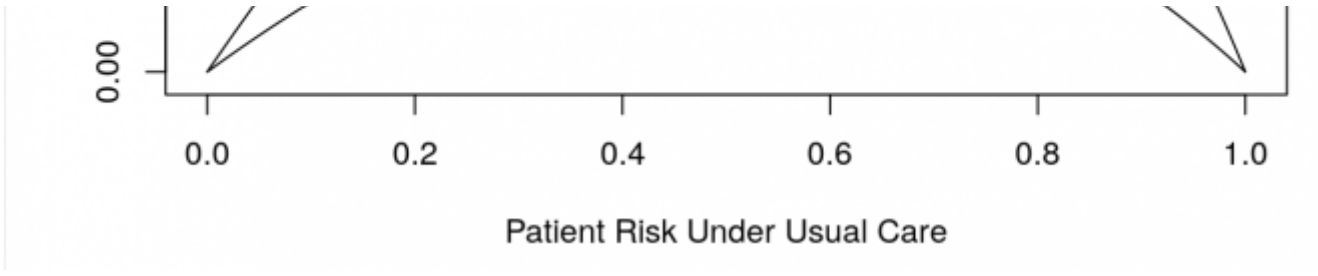

• Reply • Share ›

This comment was marked as spam.

Frank Harrell Mod ➚ Nadeem Riaz • 4 years ago

Nadeem I think your comments are perfect and add greatly to rounding out the discussion. Thanks for taking the time to write this. As you said, risk magnification is going to be most applicable when applying existing therapies to patient decision making.

1 • Reply • Share ›

Duncan • 4 years ago

You mention patient utility analysis and tradeoffs between multiple endpoints. I'd like to learn more, but I'm struggling to find a good place to start. Could you recommend something?

• Reply • Share ›

Frank Harrell Mod ➚ Duncan • 4 years ago

I don't have a best reference that comes to mind at the moment. It's a great question. Start here: http://www.citeulike.org/us...

• Reply • Share ›

John • 4 years ago

Thanks for the interesting discussion. When looking at the code for both the data hungry paper (Modern modelling techniques are data hungry: a simulation study for predicting dichotomous endpoints; see additional file 2) and the code posted at http://biostat.mc.vanderbil..., I noticed that both codes used unpruned rpart trees. This is not my interpretation of the intended use of the rpart function (see Section 4 in https://cran.r-project.org/... and I also think it is not implementing the original CART algorithm. Is there a reason for why unpruned trees are used?

• Reply • Share ›

Frank Harrell Mod ➚ John • 4 years ago

Thanks for raising this issue. It's my understanding that rpart automatically does 'optimal' pruning. If you see something in the documentation that doesn't support that statement please let me know.

• Reply • Share ›

John ➚ Frank Harrell • 4 years ago

Thanks for the quick reply. Everything in what follows is referring to the rpart vignette "An Introduction to Recursive Partitioning Using the RPART Routines" found at https://cran.r-project.org/... . Most of section 4 is describes how you do the pruning after rpart has been used to fit the initial (usually overfit) tree in section 3. For example, Section 4 starts by stating "We have built a complete tree, possibly quite large and/or complex, and must now decide how much of that model to retain". Furthermore, in Section 4.3 they do an example where they state "Looking at the table, we see that the best tree has 10 terminal nodes (9 splits), based on cross-validation. This sub tree is extracted with a call to prune and saved in fit9." Here the final tree built using rpart had 28 terminal nodes but the optimal tree, using the 1-SE error cross-validation criteria, has size 10.

As another example of a final tree outputted by rpart which is not optimally pruned, see Section 7.3 in the vignette. The cptable shows that the tree with one split has the smallest cross-validation error (xerror) but the final tree outputted by rpart is of size 10.

• Reply • Share ›

Frank Harrell Mod ➚ John • 4 years ago

This is extremely helpful and I need to address this better. But pruning will help with calibration and hurt predictive discrimination more than 1/2 of the time. Pruning makes tree structure more stable by making the tree try to do less. So ultimately, a single tree method is conservative when compared to the R^2 you'll get from regression.

• Reply • Share ›

Roy Tamura • 4 years ago

Thanks for a comprehensive set of notes of the conference. Interesting that PCORI was one of the

Thanks for a comprehensive set of notes of the conference. Interesting that PCORI was one of the sponsors. One of their 'must haves' in grant applications is a section on how you will detect patient heterogeneity in your clinical trial. I was helping a rare disease researcher design a small clinical trial for PCORI funding and wrote 'not applicable due to small sample size'. One reviewer was not happy with that response. For those writing up protocols, how would you recommend we discuss this topic (examination of subgroups)?

Frank Harrell Mod ↗ Roy Tamura • 4 years ago

What a terrific question. I do think PCORI has gone overboard on this. The statistical reality of being able to demonstrate evidence for a treatment x covariate interaction is that if the sample size is just big enough to detect an overall average effect, you may need 16 times that sample size to have adequate power to detect the interaction. So in many cases it's a lost cause - another reason for emphasizing risk magnification instead. The only rational approach I can think of regarding differential treatment effect (the ecological version of HTE) is to use Bayesian models and having priors for interaction effects. This allows one to borrow information, make best use of available data, and allow an interaction to be "half in and half out" of the model. For more about this see sim97bay.pdf in http://data.vanderbilt.edu/...

Related

- Implementation of the PATH Statement
- Implications of Interactions in Treatment Comparisons
- Assessing Heterogeneity of Treatment Effect, Estimating Patient-Specific Efficacy, and Studying Variation in Odds ratios, Risk Ratios, and Risk Differences
- EHRs and RCTs: Outcome Prediction vs. Optimal Treatment Selection
- RCT Analyses With Covariate Adjustment