What makes a great poster?

Daniel Farewell

RSS Conference Co-Chair 2018; IBS CNC Conference Chair 2017
(but have only ever made one poster myself)

12 August 2021

What does success look like?

What does success look like?

- best poster prize

What does success look like?

- best poster prize
- one great conversation
- several good conversations

What does success look like?

- best poster prize
- one great conversation
- several good conversations
- getting help with a problem
- getting people to read your paper

What does success look like?

- best poster prize
- one great conversation
- several good conversations
- getting help with a problem
- getting people to read your paper
- ... easy to make? easy to print? easy to transport?

success = best poster prize?

visual trump all else

- less is much more
- self-explanatory content
  (because judging is often done outside poster sessions)

success = best poster prize?

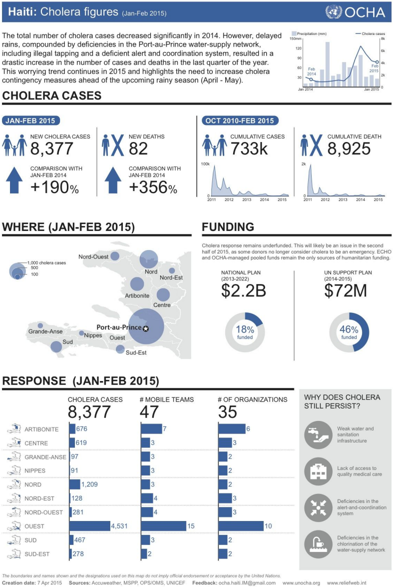

sucess = one great conversation?

(some) detail needed

- links to related work
  (especially if others in the field may be at the meeting!)
- but what's the big idea?

sucess = one great conversation?

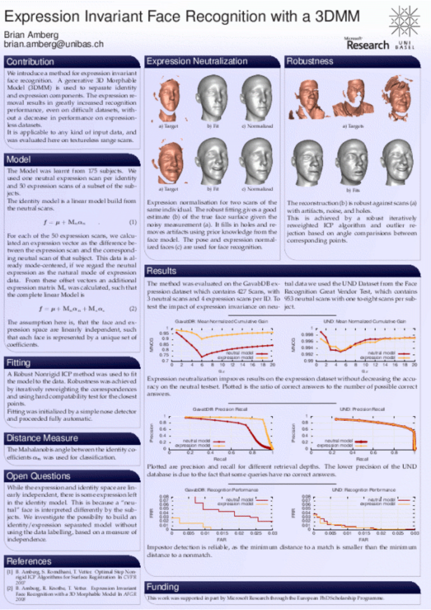

success = several good conversations?

draw people in with wide applicability

- e.g. like sports? read this poster!
- be controversial or surprising
  (e.g. do you believe in the placebo effect?)

success = several good conversations?

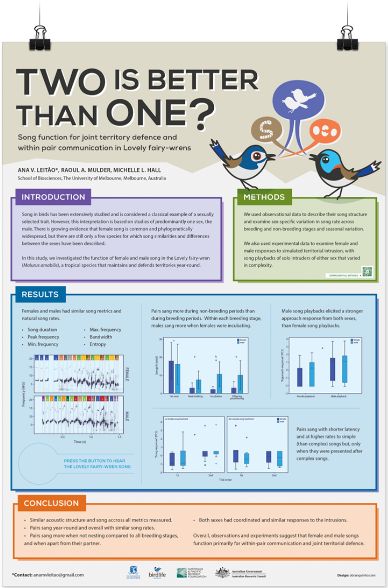

success = getting help with a problem?

ask questions

- e.g. can you complete this proof that Bayes is uniformly optimal?
- e.g. I can't explain the correlation between sandy beaches and sandwich consumption — can you?

success = getting help with a problem?

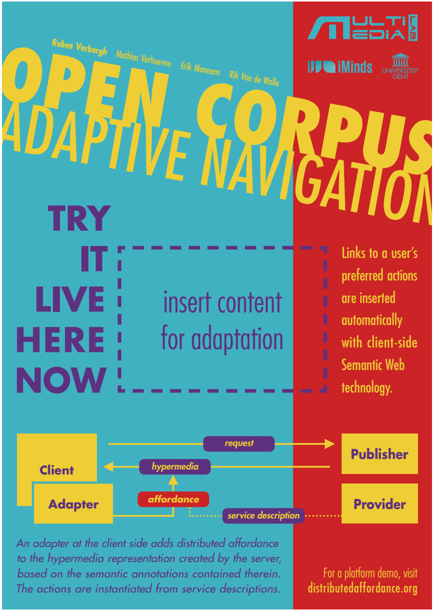

success = getting people to read your paper?

tease, but don't reveal too much!

- simplify (#betterposter)
- point to the paper for the details people need

success = getting people to read your paper?

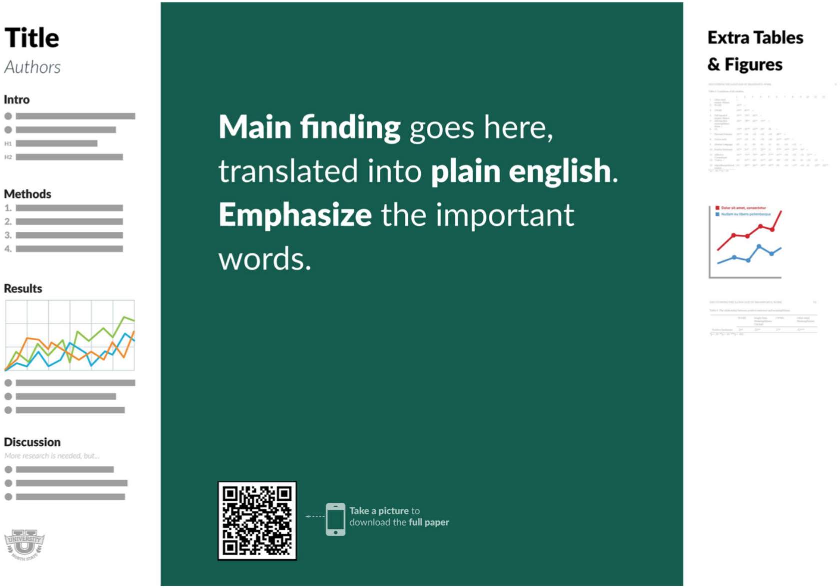

FIN

Ignorability

and unbalanced longitudinal data

DANIEL FAREWELL CHAO HUANG VANESSA DIDELEZ

No such thing* as missing data, authors claim

Marked point processes offer clearer perspective


*except when this happens

Most treatments of dropout within longitudinal studies rely on a missing data framework: usually, that of Rubin (1976). This implicitly presumes, however, that the missing data actually exist, and important notions such as 'ignorability' and 'missing at random' are based on their assumed properties.

But how are we to determine whether such assumptions are reasonable? Difficulties arise because (for instance) it is not clear what exactly is meant by 'the customer satisfaction score that would have been recorded, had this individual not died, but instead decided to attend their scheduled dental checkup'.

So can we define a condition equivalent to ignorability but without appealing to missing data? We think so, with careful consideration of the causal relationship between the points and marks of a marked point process. In a marked point process, we have the points (that is, the observation times)

$$t = (t_1, t_2, t_3, \dots)$$

and the marks (the observations)

$$y = (y_1, y_2, y_3, \dots)$$

and that's basically it. Once observations stop, it helps to think of the 'observation times' as infinite, with associated irrelevant marks. But there's no 'underlying longitudinal process'. No missing data.

Instead, formulate a causal model relating the points to the marks, and use it to write down the joint likelihood:

$$p(t, y)$$

Ignorability then boils down to a simple question: are there parts of this joint likelihood that can be ignored?

Here's a simple example

Suppose that the point process can depend on the values of previous marks. For instance, findings on medical checkups might be used to schedule subsequent checkup dates. Let's also assume that checkup findings are correlated, modelled using a random effect. Here's (part of) a causal DAG:

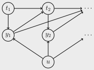

Use this to write out the joint likelihood of the points and the marks, and you'll get something like the expression on the right. Assuming no parameters in common, that first factor can be completely ignored, leaving just the familiar mixed model likelihood.

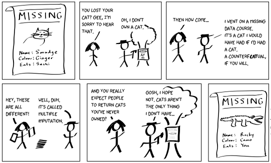

With thanks (and apologies) to xkcd

$$p(t, y) = \prod_j p(t_j \mid t_1, \dots, t_{j-1}, y_1, \dots, y_{j-1})$$

Look! It's the usual mixed effects model!

$$\times \sum_u p(u) \prod_j p(y_j \mid t_j, u)$$

farewelld@cardiff.ac.uk huangc12@cardiff.ac.uk didelez@leibniz-bips.de