

Question: Is there a big gap to convert from RStan to Nimble ?

Answer: Coming from WinBUGS or jags is immediate and straightforward. But coming from Stan does (unfortunately) mean re-writing your model code in the bugs/jags/nimble language. But that said, the languages are generally all the "same idea", but do use different syntaxes between Stan and the others

Question: So it is 'hidden', because the majority of the estimated nodes/states are not observed? Can these unobserved states be extracted from the model explicitly, given that the model will estimate them anyway?

Answer #1: They can be included in the "monitors" parameter so that you can see their estimation

Answer #2: Great question. Yes, the "H" for "hidden" means the (unobserved) "z" states, which represent the *true* underlying state, are not directly observed - we don't know what they are for sure - all we observe are the (noisy, inaccurate) observations, "y". For "extracting the states, yes, during the course of MCMC sampling, the latent states are sampled, thus we can obtain estimates of the posterior distributions of these states. As Guido just said, you can "monitor" these states (although they are not monitored by default), and obtain posterior samples for them, and hence approximate posterior distributions.

Question: Could you also include habitat covariates to model the probability of transition between sites?

Answer #1: I think that can be achieved with the multinomial logistic approach he mentioned. I hope he'll touch the subject.

Answer #2: Yes, you can have covariates on transition probabilities. Individual or time specific covariates. Habitat specific covariates are trickier as we usually have very few sites in multisite capture-recapture models.

Question: Ok thanks, I guess it is in the same way that you would include covariates in survival or detection?

Answer: Yes, same approach except for more than 2 sites - see e.g. <https://onlinelibrary.wiley.com/doi/full/10.1002/ece3.4827>. We cover the use of covariates in a longer version of this workshop, see https://oliviergimenez.github.io/bayesian-cr-workshop/slides/5_survival.html#1. See also https://oliviergimenez.github.io/bayesian-cr-workshop/worksheets/4_demo.html and <https://oliviergimenez.github.io/bayesian-cr-workshop/worksheets/exercises.html>

Question: Is it an assumption from all Hidden Markov models that all individuals are captured at the first occasion?

Answer: Yes, for all practical models that we use. Unless you have a complete census of the population (are aware of all individuals that could be captured) in advance, in which case you'd know which individuals are *not* captured on the first occasions. But without a complete census, we don't even know what individuals are out there in the population, so can't practically model non-detections (prior to the first detection) of individuals, since they wouldn't be marked (aka we aren't even aware of their existence).

Remark: Function `fortune()` from package `fortunes` can be good to pass the time while we wait

Question: Just for clarification: Did I understand correctly, that a different sampler can be used for different parameters within the same model run?? That's amazing!

Answer: Yes, for sure. Totally customizable. You can even apply multiple samplers to the same parameter, or specify particular samplers to operate more than once, on each MCMC iteration. See `"samplerExecutionOrder"`