Prune redundant models

Brian C. O'Meara

A recent article by colleagues (Tarasov and Uyeda (2026)) created a more complex model (HiClaSSE) and then showed that there can be various ways to construct different simplified models that make different conclusions about biology but have identical likelihood. Using their HiClaSSE framework, they demonstrate how different simplifications of a complex model can imply different biological interpretations while remaining indistinguishable from the perspective of the likelihood.

In this world where a lot of diversification models have come, rightly, under suspicion, reading through this paper is worthwhile for all users of `hisse`, especially if you want to build upon it for more complexity. However, it was reassuring to see that Tarasov and Uyeda (2026) note that:

> Ideally, a computational algorithm should be developed to automatically test irreducibility. In the meantime, since lumpability conditions (L1-L3) are complex, we provide practical tips: (1) Models without hidden states are always irreducible. (2) Constraining models helps—for example, avoiding dual transitions or giving each hidden trait its own speciation rate. (3) Hidden characters can also be treated as diversification-regime categories, as adopted in HiSSE framework, where each hidden category receives parameters for (i) speciation/extinction, (ii) within-category transitions, and (iii) between-category transitions. We did not detect any reducible models resulting from this construction.

Louca and Pennell (2020) has inspired a lot of discussion of congruence in diversification models. This is important when it affects the inferences we make. For example, given the results of Louca and Pennell (2020), we (Jeremy & Brian) put very little stock in *any* analysis seeking to look at change in speciation, extinction, and/or diversification rate over time using modern data ("pulled rates" are justifiable, but biological interpretation can be hard). However, as we showed (O'Meara and Beaulieu 2021; note this is a preprint; reviewers found it interesting but wanted a more focused message), and have an R package and website to demonstrate (http://flippedcoin.info), **even coin flipping has multiple congruent models**. Because a model of decreasing chance of heads fits as well as the standard binomial model does not mean coin flips can no longer be modeled, but rather that we have to understand the limits of what we are choosing to model. [Note that it is *not* that the Tarasov and Uyeda (2026) article argues for tossing models when there are other congruent models; we're just concerned about a misreading of this sort].

However, congruent models can already affect analyses in `hisse` (and in many other areas in phylogenetics). We first noticed this when developing `MiSSEGreedy()`, a way of examining various MiSSE models for estimating diversification rate parameters (similar use case to BAMM, CLaDS, and the DR statistic, and with good performance). Sometimes we can get two models with the same likelihood, or models with the same parameter values but different labels. For this reason, in 2021 we added a `PruneRedundantModels()` function to `hisse`. More about its use is discussed in Vasconcelos et al. (2022). But here we have an example:

```
phy <- ape::read.tree("https://raw.githubusercontent.com/thej022214/hisse/
                      refs/heads/master/vignettes/whales_Steemanetal2009.tre")

greedy_run <- hisse::MiSSEGreedy(phy, n.cores=max(1, -1+parallel::detectCores()))
```

We can look inside PruneRedundantModels() to see how it identifies models for deletion. Basically, it goes through four steps:

Compute the number of free parameters per model:

```
models <- greedy_run
mod.nparameters <- simplify2array(lapply(lapply(models, "[[", "starting.vals"),length))
```

Reorder models from simplest to most complex:

```
models <- models[order(mod.nparameters, decreasing=FALSE)]
```

Extract the negative log likelihoods from the sorted models

```
mod.loglik <- simplify2array(lapply(models, "[[", "loglik"))
```

And then go through all the models to see if the likelihood is the same (within the specified precision) of a previous model (which will have the same or fewer parameters). If so, it selects the model for deletion.

So we start with 60 examined models, and then using the following command:

```
greedy_pruned <- hisse::PruneRedundantModels(greedy_run)
```

reduce this to 30 models. This is especially useful for removing models in which additional parameters provide little or no improvement in likelihood. For example, imagine a model in which extinction rates for hidden states A, B, and C are estimated to be essentially zero, despite being free to vary. Simpler models in which some or all of these extinction rates are constrained to be equal or fixed at zero may achieve nearly identical likelihoods. Although these models differ in their numbers of free parameters, they may represent effectively the same fitted diversification process.

When such models are retained separately, they can collectively receive substantial weight during model averaging simply because the same fitted process is represented by multiple parameterizations. `PruneRedundantModels()` addresses this practical problem by retaining the simplest model among those with effectively similar likelihoods. However, identical likelihoods do not necessarily establish mathematical equivalence. Pruning should therefore be understood as a pragmatic way to reduce redundancy among fitted models, rather than a formal test of model congruence.