Estimating variability and uncertainty in predatory relationships

A unified Bayesian framework for stable isotopes and fatty acid profiles

Philipp Neubauer and Olaf Jensen

13 August, 2014

DRAGONFLY
Data Science

To eat and be eaten

Interactions are complex

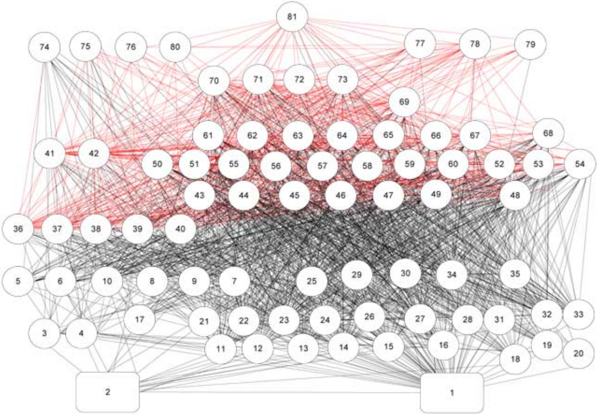

Figure 22. Food web of the northeast shelf large marine ecosystem (NES LME). Adapted from Link 2002.

Interactions are hard to quantify!

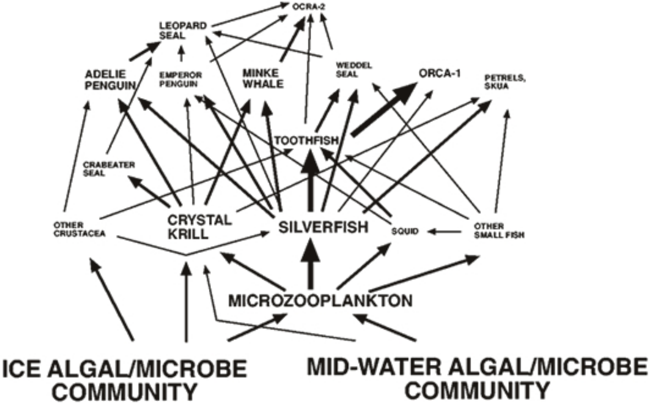

Figure: Landcare Research

Quantifying predation

Diet markers – Pros and Cons

- **Stomach contents & fecal matter**: Straightforward, BUT different digestion rates, only a snapshot
- **Stable Isotopes (SI)**: Integrate over weeks to months, trophic enrichment, sophisticated analysis methods BUT enrichment depends on prey concentration, few isotopes often means poor discrimination.
- **Fatty Acids (FA)**: Lots of potential markers, some characteristic markers. BUT metabolised differently depending on prey and predator species. Compositional data: sum to 1 constraint.

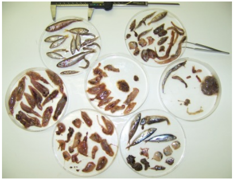

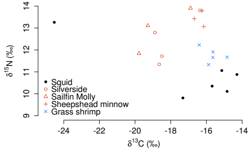

DRAGONFLY
Data Science

fastinR

Bayesian diet estimates from SI and FA

- Develop a common framework for stable isotopes and fatty acids: Mixing-model based framework to estimate diet proportions.
- **Challenge 1**: Deal with proportion data: things get tricky on the simplex!
- **Challenge 2**: Pick from a huge amount of FAs that could be included, but make the analysis hugely inefficient.
- **Challenge 3**: Deal with metabolic conversion of FAs.
- Make assumption that the two markers reflect the same diet (not necessarily true!) - combine likelihoods for each marker. Or analyse separately.
- Estimate diet proportions in a Bayesian framework.

Squid experiment (Gabi Stowasser et al. 2006)

Shrimp ~ 30 d

SC ~ 15 d

Fish ~ 30 d

SF ~ 15 d

fastinR

Challange 1

Deal with a compositional aspect of the data:

- Arises from measurement process (relative area under peaks from Gas Chromatography)
- Use appropriate likelihood (logistic normal)
- Makes the model somewhat inefficient with large number of FAs, but see next challenge…
- Difficult to parametrise the prior variance (tinkering is often needed to get the models to converge…)

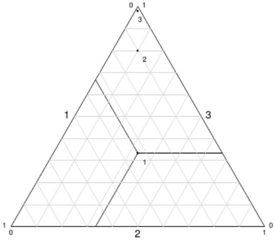

fastinR

Challange 2

Find a subset of FAs:
contribution to source
separation in multivariate space.

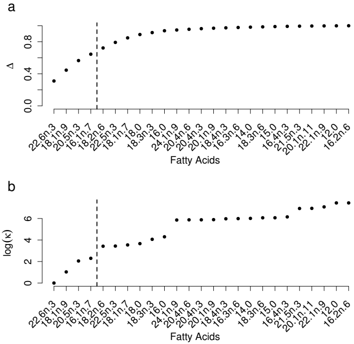

DRAGONFLY
Data Science

fastinR

Challange 2

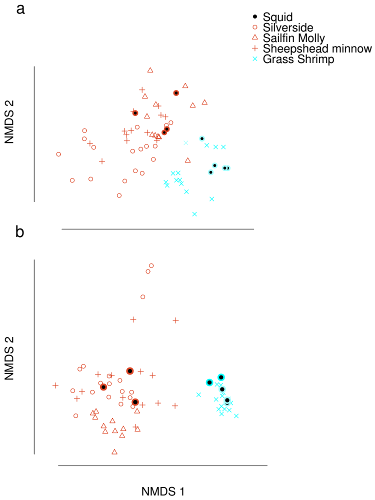

Find a subset of FAs:
contribution to source
separation in multivariate space.

DRAGONFLY
Data Science

fastinR

Challange 3

Metabolic conversion of FAs in the predator:

- Needs to account for compositional / proportion aspect of data: when the proportion of one FA goes up, another must go down!
- Just maths after that…
- Conversion coefficients (also fractionation for SI) estimated from 30 d trails with no switching.

fastinR

back to the action!

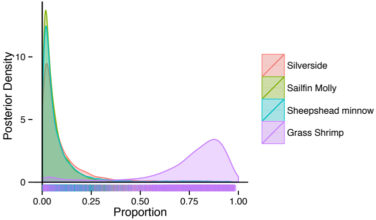

Stable isotopes, SC treatment

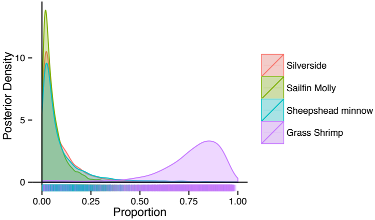

Fatty Acids, SF treatment

DRAGONFLY
Data Science

fastinR

combining markers

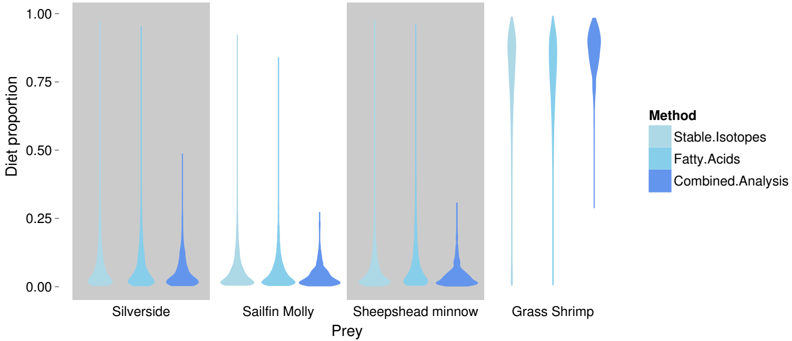

Combined markers, SC treatment

fastinR

combining markers - switched treatment

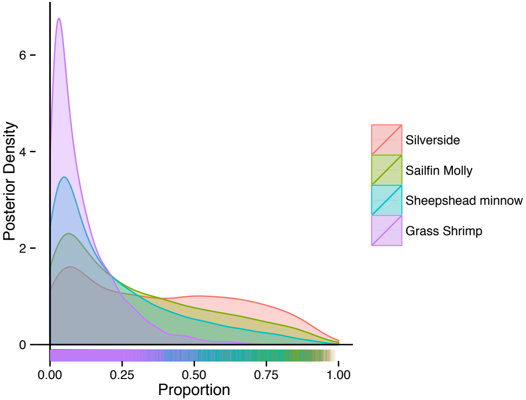

Combined markers, SF treatment

fastinR

combining markers - switched treatment

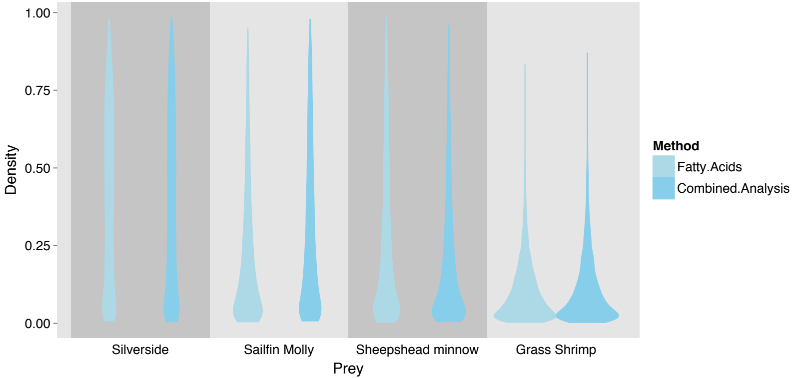

Combined markers, SF treatment

DRAGONFLY Data Science

fastinR

combining markers - switched treatment

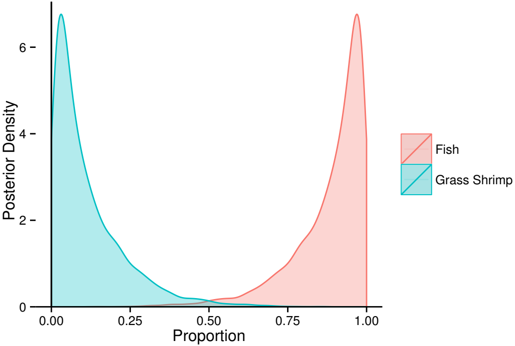

Combined markers, SF treatment

fastinR

Individual diet estimates

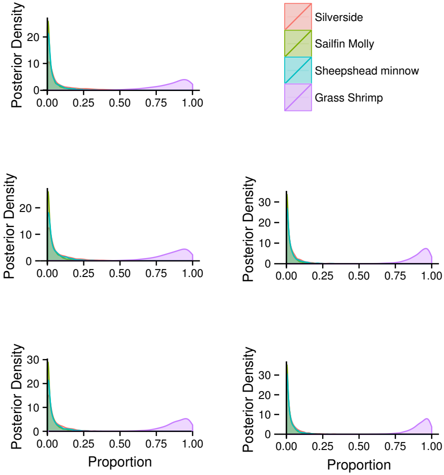

Combined markers, SC treatment, Individual predators

fastinR

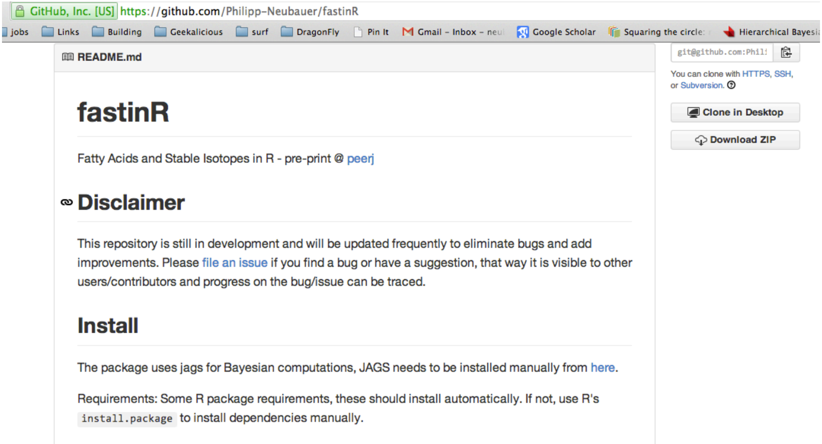

fastinR

Fatty Acids and Stable Isotopes in R - pre-print @ peerj

Disclaimer

This repository is still in development and will be updated frequently to eliminate bugs and add improvements. Please file an issue if you find a bug or have a suggestion, that way it is visible to other users/contributors and progress on the bug/issue can be traced.

Install

The package uses jags for Bayesian computations, JAGS needs to be installed manually from here.

Requirements: Some R package requirements, these should install automatically. If not, use R's `install.package` to install dependencies manually.

fastinR

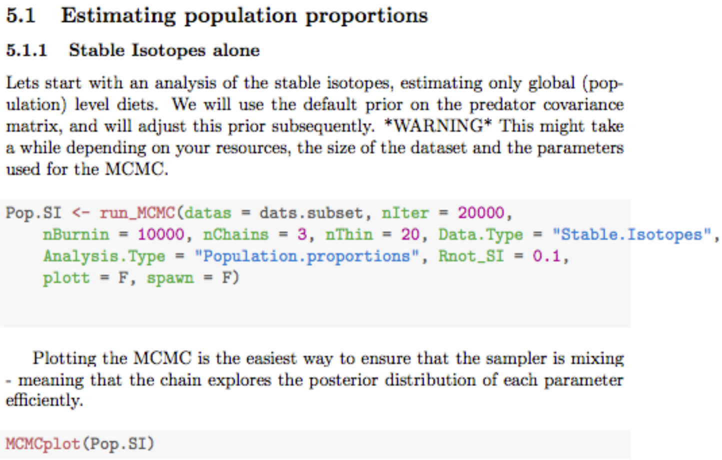

5.1 Estimating population proportions

5.1.1 Stable Isotopes alone

Lets start with an analysis of the stable isotopes, estimating only global (population) level diets. We will use the default prior on the predator covariance matrix, and will adjust this prior subsequently. *WARNING* This might take a while depending on your resources, the size of the dataset and the parameters used for the MCMC.

```
Pop.SI <- run_MCMC(datas = dats.subset, nIter = 20000,
    nBurnin = 10000, nChains = 3, nThin = 20, Data.Type = "Stable.Isotopes",
    Analysis.Type = "Population.proportions", Rnot_SI = 0.1,
    plott = F, spawn = F)
```

Plotting the MCMC is the easiest way to ensure that the sampler is mixing - meaning that the chain explores the posterior distribution of each parameter efficiently.

```
MCMCplot(Pop.SI)
```

fastinR

Use and abuse at will!

Open source so please contribute!