NMFS FELLOW - CONTINUING REPORT

NMFS Fellows - 2016-2017 Cole Monnahan NMFS Population Dynamics Sea Grant Graduate Fellowship: Optimizing Bayesian analysis in data-rich stock assessments and management resources in data-limited fisheries. Trevor Branch in support of Cole Monnahan E/I-21 Submitted On: 03/16/2017 10:01:24 AM

METRICS & MEASURES

Metric/Measure	Value	Note
Acres of coastal habitat	0	
Fishermen and seafood industry personnel	0	
Communities - economic and environmental development	0	
Stakeholders - sustainable approaches	0	
Informal education programs	0	
Stakeholders who receive information	0	
Volunteer hours	0	
P-12 students reached	0	
P-12 educators	0	

REQUESTED INFORMATION

Publications

Faster estimation of Bayesian models in ecology using Hamiltonian Monte Carlo Publication Type: Peer-reviewed: Journals (incl. articles), Books, Proceedings, and Other **Documents** Publication Year: 2016 **Publication Authors:** Publisher Info: Methods in Ecology and Evolution Notes: **Related URLs: Keywords: Publication URLs:** Abstract: Citation: Monnahan, C.C., J.T. Thorson, and T.A. Branch. (2017). Faster estimation of Bayesian models in ecology using Hamiltonian Monte Carlo. Methods in Ecology and Evolution 10.1111/2041-210x.12681.10.1111/2041-210x.12681 **Citation for Coverpage:** SG can post PDF online?: Yes Uploaded File: Monnahan_et_al_2017_Faster_estimation_of_Bayesian_models.pdf

Implications of process error in selectivity for approaches to weighting compositional data in fisheries stock assessments

Publication Type: Peer-reviewed: Journals (incl. articles), Books, Proceedings, and Other **Documents** Publication Year: 2016 **Publication Authors:** Publisher Info: Notes: **Related URLs:** Keywords: **Publication URLs:** Abstract: Citation: Stewart, I.J. and C.C. Monnahan. (2016). Implications of process error in selectivity for approaches to weighting compositional data in fisheries stock assessments. Fisheries Research 10.1016/j.fishres.2016.06.018.10.1016/j.fishres.2016.06.018 Citation for Coverpage: SG can post PDF online?: No Uploaded File: WSG-TA 16-30 Stewart.pdf

Students Supported

No Students Supported information reported

Narratives

Hamiltonian Monte Carlo for stock assessments Uploaded File: Monnahan_progress_report_2017.docx

Partners This Period

No Partners This Period information reported

STANDARD QUESTIONS

Community Hazard Resilience

No Community Hazard Resilience information reported

Economic Impacts

No Economic Impacts information reported

Impacts and Accomplishments

No Impacts and Accomplishments information reported

Leveraged Funds

No Leveraged Funds information reported

Meetings, Workshops, Presentations

(1)

Type of Event

Public or professional presentation

Description	Hamiltonian Monte Carlo in ADMB and TMB: current status and future directions. ADMB/TMB Developers' Workshop. Seattle, WA. 6/23/2016.Discuss my plans to implement HMC into the popular ADMB and TMB software platforms with the developers. Developed a plan for how to do this.
Event Date	06-23-2016
Number of Attendees	15

(2)

Type of Event	Sea Grant-sponsored/organized event
Description	 Faster estimation of Bayesian models in ecology using Hamiltonian Monte Carlo. NOAA Fisheries Sea Grant Fellowship Meeting. Southwest Fisheries Science Center. Santa Cruz, CA. 6/28/2016. Explain basics of HMC to help disseminate their use in the wider scientific community.
Event Date	06-28-2017
Number of Attendees	30

(3)

Type of Event	Public or professional presentation
Description	Introducing the no-U-turn MCMC sampler in ADMB and TMB: faster run times for large, complex fisheries models. UW School of Aquatic and Fisheries Science: Fisheries Think Tank. Seattle, WA. 2/7/2017 Demonstrate how to use these algorithms in both ADMB and TMB, using my newly-developed R package adnuts.
Event Date	02-07-2017
Number of Attendees	20

Tools, Technologies, Information Services / Sea Grant Products

No Tools, Technologies, Information Services / Sea Grant Products information reported

Project details

Cole Monnahan, monnahc@uw.edu

Project number: E/I-21, award number: NA13OAR4170109

Project title:

Optimizing Bayesian analysis in data-rich stock assessments and management resources in data-limited fisheries

Project update:

My work since June 2016 falls into three categories: (1) Incorporating the no-U-turn sampler into TMB and ADMB and writing an R package for easier dissemination of these methods, (2) Testing these methods for fisheries stock assessments, and (3) Publishing relevant peer-reviewed papers.

Hamiltonian algorithms

The "no-U-turn" sampler by Hoffman and Gelman (2014) is considered state of the art for Bayesian models. Previously it was only available in the software platform Stan (Stan Development Team 2015). In the past year I have written this algorithm into the source code for ADMB, and written an R package (adnuts) which acts as a front end for analysts to use these methods more effectively. Additionally, I have given several presentations demonstrating the theoretical underpinning of these methods, but also how to use them in practice.

In my remaining time in graduate school I will conduct more thorough tests to ensure the algorithms are coded correctly. I will then run several stock assessment models to demonstrate how this algorithm is beneficial for the stock assessment community. In addition, I will more formally quantify how much run time has been reduced over the default methods used currently. I will continue writing up the results for publication in a peer-reviewed scientific journal.

Implications of process error in selectivity for approaches to weighting compositional data in stock assessments

In a continuing collaboration with my committee member, Dr. Ian Stewart, we published a peerreviewed paper in the journal *Fisheries Research*. Our simulation study investigates how data weighting and process error interact in a stock assessment model. The results provide guidance for authors conducting assessments that need to "weight" their input data

Writing chapters

My first chapter was published this year in Methods in Ecology and Evolution. Currently my second chapter, looking at the effects of hooking spacing in catch rate standardization for longline fishers, is in review. My third chapter consists of software and an accompanying manual. I plan to submit this as an "application" to Methods in Ecology and Evolution.

The bulk of my remaining work is testing stock assessment methods. I will be writing that up for publication in *The Canadian Journal of Fisheries and Aquatic Science*.

Summary of progress

I have published my first chapter, my second is in review, my third is essentially done, leaving my last chapter to do. I have a defense date set for 6/9/2017, so I will be finishing these tasks over the new 3 months.

WSG-TA 16-30

ARTICLE IN PRESS

Fisheries Research xxx (2016) xxx-xxx

ELSEVIER

Contents lists available at ScienceDirect

Fisheries Research



journal homepage: www.elsevier.com/locate/fishres

Implications of process error in selectivity for approaches to weighting compositional data in fisheries stock assessments

Ian J. Stewart^{a,*}, Cole C. Monnahan^b

^a International Pacific Halibut Commission, 2320 West Commodore Way, Suite 300, Seattle, WA 98199-1287, United States ^b Quantitative Ecology and Resource Management, University of Washington, Seattle, WA 98195-5020, United States

ARTICLE INFO

Article history: Received 27 April 2016 Received in revised form 15 June 2016 Accepted 16 June 2016 Available online xxx

Keywords: Stock assessment Selectivity Data weighting Simulation testing Observation error

ABSTRACT

Lack-of-fit in a stock assessment model can be related to both data weighting and the treatment of process error. Although these contributing factors have been studied separately, interactions between them are potentially problematic. In this study we set up a simple simulation intended to provide general guidance to analysts on the performance of an age-structured model under differing assignments of compositional data weight and process variance. We compared cases where the true sample size is under-, 'right-' or over-weighted, and the degree of process variance (in this case temporal variability in selectivity) is under, correctly, or overestimated. Each case was evaluated with regard to estimation of spawning biomass, and MSY-related quantities. We also explored the effects of the estimation of natural mortality, steepness, as well as incorrectly specifying process error in selectivity when there is none. Results showed that right-weighted estimation models assuming the correct degree of process error performed best in estimating all quantities. Underweighting produced larger relative errors in spawning biomass, particularly when too much process error was allowed. Conversely, overweighting produced larger errors mainly when the degree of process error was underestimated. MSY-related quantities were sensitive to both the estimation of natural mortality, and particularly steepness. We suggest that data weighting and the treatment of process error should not be considered independently: estimation is most likely to be robust when process error is allowed (even if overestimated) and when compositional data are not excessively down-weighted.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Integrated statistical fisheries stock assessments are commonly used for the management of many important fish stock around the world (Fournier and Archibald, 1982; Hilborn and Walters, 1992; Maunder and Punt, 2013; Megrey, 1989; Quinn and Deriso, 1999). Modern integrated models can be complex, considering multiple sources of data, with alternative error assumptions and relative weights (Maunder and Piner, 2014; Maunder and Punt, 2013). Two recent workshops have highlighted the importance of how selectivity, relative data weights and process error are treated in integrated assessment models (Maunder et al., 2014 summary for data weighting workshop).

There are many sources of process error which may be important contributors to bias and imprecision in integrated stock

* Corresponding author.

E-mail addresses: lan@iphc.int (I.J. Stewart), monnahc@uw.edu (C.C. Monnahan).

http://dx.doi.org/10.1016/j.fishres.2016.06.018 0165-7836/© 2016 Elsevier B.V. All rights reserved. assessments, including recruitment variability, mortality rates, growth, selectivity and catchability (e.g., Hurtado-Ferro et al., 2014; Johnson et al., 2015; Linton and Bence, 2008; Ono et al., 2015). Recent research has highlighted the inherent complexity in the treatment of selectivity, historically considered to be 'nuisance' parameters, but increasingly acknowledged to be important for unbiased estimates of stock size and trend (Maunder et al., 2014). Misspecification of selectivity has long been known to produce bias in statistical catch-at-age models (Kimura, 1990), and several recent studies have shown that time-varying selectivity may be expected in many cases (Sampson, 2014; Sampson and Scott, 2012). Misspecifying selectivity has also been implicated in retrospective patterns (Hurtado-Ferro et al., 2014; Stewart and Martell 2014).

Data weighting is also a particularly problematic aspect of stock assessment (Francis 2011), and the focus of the recent CAPAM workshop (Ref. summary for this issue). Data weighting is an issue because, although they arise from separate processes, the true underlying error distribution and sampling variances for most (perhaps all) fisheries data are unknown (Maunder, 2011). Further, the likelihoods that are used in most stock assessment models are often

I.J. Stewart, C.C. Monnahan / Fisheries Research xxx (2016) xxx-xxx

known to be convenient approximations (e.g., the multinomial), despite far greater known complexity in the underlying processes. The primary inputs for data weighting are the variance estimates assigned to indices of abundance and the sample sizes (or variances, depending on the likelihood function used) assigned to length- or age-composition data (hereafter 'composition data'). For instance, a research survey may collect lengths for thousands of fish, sampled from hundreds of hauls, but an effective sample size of 20 may ultimately be used in the assessment (Francis 2011; Stewart and Hamel, 2014). The results of stock assessments are sensitive to data weighting, and the choice of method used is most consequential when there is model misspecification (Punt, 2016). Several recent papers, and some in this issue, explore alternative methods for deriving weights after the model is assumed to be correctly specified (Citations pending other manuscripts in this issue).

However, existing research has largely focused on either the treatment of process error (e.g., estimation of time-varying fishery selectivity) or data weighting, but not both (although, see Thompson, this issue). Assuming one is known allows for several relatively simple methods (Maunder and Harley, 2011; Thompson and Lauth, 2012; Thorson and Taylor, 2014), but when both are uncertain the problem becomes much more difficult. Lack-of-fit cannot necessarily be objectively assigned to either observation variance or process variance in a stock assessment (and there are other sources of both process and observation error, Linton and Bence, 2008) based on model fit alone.

In this study, we structured a simulation approach to evaluate the performance of a simple age-structured model under combinations of different compositional data weighting and process error assumptions, and whether natural mortality, steepness of the stock-recruit curve, or both were estimated. This simple simulation is intended to provide general guidance to analysts, not in terms of which data weighting rules or approaches should be applied, but whether it is preferable to over- or under-estimate the weight placed on compositional data and whether the treatment of process error in fishery selectivity affects this choice.

2. Simulation approach

We used the ss3sim package (Anderson et al., 2014, 2015) implemented in R (R Core Team, 2016), to create an operating model with a very simple, but groundfish-like life history. This package uses the flexible software stock synthesis (Methot and Wetzel, 2013), coded in AD Model Builder (Fournier et al., 2012) for both the operating and estimation models in simulations (i.e., self-testing instead of cross-testing; Deroba et al., 2014). As in other published studies using ss3sim, this framework is well suited to relatively simple simulation experiments intended to target the effects of one to several factors with representative, but not necessarily realistic levels of model complexity (e.g., Hurtado-Ferro et al., 2014; Monnahan et al. 2015; Ono et al., 2015). All code used to produce this study is based on open source tools and available online (Appendix A in Supplementary material).

The operating model was used to generate replicate data sets including variability in annual recruitment, which were then fit with a range of estimation models; a summary of both models is provided below. Results were summarized via the distribution of relative error in the estimated spawning biomass across the time series, and a few key metrics related to Maximum Sustainable Yield (merely to illustrate the potentially differing effects on reference point estimation). We compared cases where the true sample size is under-, 'right-' or over-weighted, and the degree of process error (in this case temporal variability in fishery selectivity) is under, correctly, or overestimated. We also explored the effects of the estimation of natural mortality, steepness, as well as incorrectly specifying process error in fishery selectivity when there is none. A test was run with extremely large sample sizes to verify that model dynamics were performing as expected, and that the estimation models produced unbiased results when correctly specified.

2.1. Operating model specifications

The operating model was structured to be similar to many simple stock assessments with a moderately long time-series of observations, a single fishery responsible for the catch and a single fishery-independent survey (Table 1). Population dynamics are governed by fishing and natural mortality (M=0.2), with annually variable recruitment centered on a Beverton-Holt stockrecruitment function with an intermediate value for steepness (h = 0.65). Selectivity for both the fishery and survey was asymptotic, using a simple parametric logistic form (implemented using the ascending side of the double-normal parameterization in stock synthesis; assuming asymptotic selectivity likely reducing the potential for confounding with natural mortality in all estimation models). Growth was specified as a von Bertalanffy curve with moderately rapid growth and a clear asymptote at older ages. Fishing effort, growth parameters, and all other parameters were fixed among simulations, with the only variability arising from recruitment deviations unique to each replicate, and stochastic generation of data.

The fishery and fishery-independent survey each produced compositional data (lengths and ages) for a subset of the time-series years (Table 1, Appendix A in Supplementary material) based on a multinomial distribution matching the true population (i.e., unbiased) with samples sizes roughly consistent with those observed in 'data-rich' stock assessments. Two operating model scenarios were evaluated: 1) time-invariant fishery selectivity (no process error) and 2) a temporal trend in fishery selectivity (process error). The fishery selectivity trend corresponded to a linear decrease in the parameter defining the first size with 100% selectivity (as parameterized in stock synthesis), followed by a linear increase in that parameter over most of the informed time-series (Fig. 1). This pattern in fishery selectivity was selected over a simple 'white-noise' (independent random deviations from a mean) to mimic trends suspected in some stock assessments with directional changes in fishing behavior, biology or both (Stewart and Martell, 2014). This type of change is more likely to produce systematic bias in the estimated demographics of the removals, and therefore potentially more import for analysts to consider.

The operating model time series resulted in a stock fished down from unexploited equilibrium relatively rapidly then recovering to slightly higher levels at the very end of the time-series, characteristic of some histories observed in actual stock assessments (Fig. 2). For each combination of operating model scenario and estimation model case (described below), 300 replicate time series were randomly generated.

2.2. Estimation models

Estimation models were correctly specified (matching the operating model) for all model parameters except: fishery selectivity, virgin recruitment, survey catchability and the set of specific factors explored: the degree of process error, the degree of observation error, the value for natural mortality, and the value for steepness of the stock-recruitment function (Table 2; we did not evaluate the more complicated case of natural mortality and steepness estimated simultaneously). The sample sizes were not tuned in the estimation model, but were specified as either too small (0.1x), right (1x), or too large (10x), relative to the true simulated sample size (x). Fishery selectivity was either assumed to be constant, or allowed to vary as an additive random walk over time (Methot,

Please cite this article in press as: Stewart, I.J., Monnahan, C.C., Implications of process error in selectivity for approaches to weighting compositional data in fisheries stock assessments. Fish. Res. (2016), http://dx.doi.org/10.1016/j.fishres.2016.06.018

2

I.J. Stewart, C.C. Monnahan / Fisheries Research xxx (2016) xxx-xxx

Table 1

Operating model summary.

Specification	Value	Comment
Structure Time series length Catches Fleets Survey and fishery selectivity shape Stock-recruit function Recruitment deviations Biology	100 years Years 27–100 2 Asymptotic Beverton-Holt Annual Single-sex	Based on fixed vector of Fs (Fig. 2) Fishery and survey Two-parameter logistic Parameterized with R_0 and h Randomly generated from lognormal
Model parameters $Log(R_0)$ Steepness (h) Recruitment variability (σ_r) Natural Mortality (M) Survey catchability (Q) Length at age-1 Asymptotic length Brody growth coefficient (k) CV of length-at-age Survey selectivity slope Survey selectivity peak length Fishery selectivity peak length Selectivity process error(scenario 2) Implied selectivity sigma	18.7 0.65 0.4 0.2 1 20 cm 132 cm 0.2 0.1 5.2 41.8 5.1 Variable Trend up and down in peak parameter 0.5	Of course Constant across all ages Log (width); constant over time Constant over time Log (width); constant over time Constant or with process error (Fig. 1) When treated as deviations, this trend results in the implied sigma below.
Data Generation Survey index data Survey index sigma Survey length and age data Survey length and age sample size Fishery length and age data Fishery length and age sample size	Year 76–100 0.2 Year 76–100 500 (each) Triennially from year 36–72, annually thereafter 100	Biennial In log space; constant across years Biennial Generated from a multinomial Generated from a multinomial

Table 2

Estimation model factors; each combination across all levels of each was analyzed (except M and h were not simultaneously estimated).

Process error	Data weight	Natural mortality (M)	Steepness (h)
S0: None (Sigma = 0) S1: Sigma = 0.5	D ₁ : Under-weighted (x0.1) D ₂ : Right-weighted (x1)	M0: Fixed M1: Estimated	h0: Fixed h1: Estimated
S2: Sigma = 1.0	D ₃ : Over-weighted (x10)		

2015) in the estimation model (as if a trend might be expected by the analyst). The deviations for the random walk were constrained by a fixed sigma corresponding to either too little (sigma = 0), correctly specified (sigma = 0.5), or too large (sigma = 1.0). Preliminary analysis showed that allowing a random walk provided a parameterization that was capable of generally mimicking the true pattern (Fig. 3).

Each combination of factors (parameters estimated, compositional data weighting, and treatment of process error) represented a single case for the estimation model. Each case was fit to all 300 replicates for each of the two scenarios of the operating model. Fitting was performed via penalized maximum likelihood, with the sigmas for fishery selectivity and recruitment deviations specified as described above. In order to maintain the central tendency of the stock-recruitment function given lognormal variability in annual deviations, the bias correction for recruitment deviations (Methot and Taylor, 2011) was adjusted to match the conditions of each case and then held constant across all replicates. This was performed by estimating the correction over the first 10 time-series, then using the average of these for all replicates within each case (this procedure has become relatively standard for simulation experiments; see Anderson et al., 2014 for more information). Convergence of the estimation model for all replicates was evaluated based on inversion of the Hessian matrix.

3. Results

Estimation model behavior across both operating model scenarios and all combinations of factors was robust, with all replicates converging for all scenarios.

In the simple operating model scenario without any process error in fishery selectivity, fitting estimation models assuming the correct value for steepness and natural mortality in all cases resulted in essentially unbiased and reasonably precise estimates of spawning biomass over the entire time-series (Median absolute relative error (MARE)=0.05; Table 3, Fig. 4). The cases allowing for process error in fishery selectivity were similarly precise (MARE = 0.05) when compared to those correctly specified, as the penalty (sigma) constrained the deviations toward zero in the absence of signal in the data. When the composition data were over-weighted (by a factor of 10) from the true sample sizes, the time-series also showed more variability (presumably the estimated fishery selectivity was following random variability due to sampling error) but remained unbiased.

In the second operating model scenario, with process error in fishery selectivity, more interesting patterns were observed. Specifically, when natural mortality and steepness were correctly specified as well as the degree of process error, the time-series of spawning biomass estimates remained unbiased and relatively precise (MARE = 0.04–0.06 depending on compositional data weighting; Table 3, Fig. 5). However, when the estimation models

I.J. Stewart, C.C. Monnahan / Fisheries Research xxx (2016) xxx-xxx

4

 Table 3

 Median absolute relative error over the entire estimated time series of spawning biomass for each combination of operating model scenario and estimation model case.

Process error	Natural mortality	Steepness	Under-weighted (D1)	Right-weighted (D2)	Over-weighted (D3)
Operating model with	out process error in selectivity				
SO	MO	h0	0.05	0.05	0.05
S0	MO	h1	0.11	0.09	0.13
S0	M1	h0	0.05	0.05	0.05
S1	M0	hO	0.05	0.05	0.05
S1	MO	h1	0.11	0.09	0.13
S1	M1	h0	0.05	0.05	0.05
S2	MO	h0	0.05	0.05	0.05
S2	MO	h1	0.11	0.09	0.13
S2	M1	h0	0.05	0.05	0.05
Operating model with	process error in selectivity				
SO	MO	h0	0.05	0.12	0.26
S0	MO	h1	0.40	0.45	0.51
S0	M1	h0	0.06	0.11	0.26
S1	M0	h0	0.05	0.04	0.06
S1	M0	h1	0.18	0.11	0.15
S1	M1	h0	0.05	0.04	0.06
S2	MO	h0	0.05	0.04	0.06
S2	MO	h1	0.12	0.11	0.15
S2	M1	h0	0.05	0.04	0.05





were misspecified to have no process error, the time series became systematically biased (MARE = 0.05–0.26), as the fishery removals from the dynamics were also misspecified. Right-weighting only partially ameliorated this bias, and underweighting the compo-



Fig. 2. Time series plot of median spawning biomass across replicates (a), from the operating models with and without process error in selectivity. The *y*-axis units are arbitrary and left off for clarity. Instantaneous fishing mortality (*F*), constant across all replicates (b).

sition data largely removed it (MARE=0.05), likely allowing the unbiased survey index to drive the estimated trend.

For cases where natural mortality was not assumed to be known without error, the degree of bias and imprecision was substantial

I.J. Stewart, C.C. Monnahan / Fisheries Research xxx (2016) xxx-xxx



Fig. 3. Median fishery selectivity parameter deviations estimated across all replicates, with (bottom row) and without (top row) process error in the operating model (natural mortality and steepness are fixed at the true values; M0 and h0). Columns represent the level of process error specified in the EM.



Fig. 4. Time series estimates of relative error in spawning biomass (shading indicates the 25, 50, 75, and 95th percentiles) for the operating model with no process error, and estimation models do not estimate natural mortality (M0) or steepness (h0).

(MARE = 0.06–0.26; Table 3, Fig. 6). The worst performing case was represented by misspecified fishery selectivity and over-weighted compositional data (Fig. 6c). For this case, many replicates resulted in estimates of steepness equal to 1.0. Reducing the weight on the composition data improved the precision, but not the bias, as the misspecification remained. For the correctly specified selectivity sigma cases, the worst precision occurred when the data were underweighted, and there was little cost in imprecision to over-

weighting the data. A similar pattern was also observed even when too much process error in fishery selectivity was allowed.

For estimation models where fishery selectivity was misspecified and steepness was not assumed to be known, performance of all data weighting was poor (MARE=0.40–0.51; Table 3, Figs. 7 and 8). Correctly specifying fishery selectivity, or allowing too much process error performed appreciably better across

6

ARTICLE IN PRESS

I.J. Stewart, C.C. Monnahan / Fisheries Research xxx (2016) xxx-xxx



Fig. 5. Time series estimates of relative error in spawning biomass (shading indicates the 25, 50, 75, and 95th percentiles) for the operating model with process error (scenario 2), and estimation models that do not estimate natural mortality (M0) or steepness (h0).



Fig. 6. Time series estimates of relative error in spawning biomass (shading indicates the 25, 50, 75, and 95th percentiles) for the operating model including process error, and estimation models estimating natural mortality (M1).

all data weightings, with right-weighted data performing best (MARE = 0.11).

Overall, the simulation results showed that right-weighted estimation models assuming the correct degree of process error performed best in estimating the time-series of spawning biomass (Table 4), as well as the spawning biomass corresponding to MSY (Table 5). Under-weighting produced larger relative errors in spawning biomass, particularly when too much process error was allowed. Conversely, overweighting produced larger errors mainly when the degree of process error was underestimated. MSY-related quantities were sensitive to both the estimation of natural mortality, and particularly steepness (Fig. 7).

I.J. Stewart, C.C. Monnahan / Fisheries Research xxx (2016) xxx-xxx



Fig. 7. Distribution of relative error in spawning biomass producing MSY across all cases of the estimation model, and both operating model scenarios.



Fig. 8. Distribution of relative error in spawning biomass over all years for all cases of the estimation model, and both operating model scenarios.

4. Discussion

The results of this study are consistent with other recent work suggesting there is little cost besides increased run times to estimating process error even if not present, but potentially substantial bias resulting from misspecification due to the overly simplistic assumption that fishery selectivity is constant (Martell and Stewart, 2014; Punt et al., 2014; Thorson and Taylor, 2014). However, more complex approaches to smoothing of fishery selectivity may perform somewhat differently (Maunder and Harley, 2011), and if used, may warrant additional investigation. The bias due to misspecification in fishery selectivity, except in the case where all other model parameters are known perfectly and there is an unbiased trend index, cannot be ameliorated by adjustments to composi-

I.J. Stewart, C.C. Monnahan / Fisheries Research xxx (2016) xxx-xxx

8

Table 4

Median absolute relative error in estimated SSB_{MSY} for each combination of operating model scenario and estimation model case.

Process error	Natural mortality	Steepness	Under-weighted (D1)	Right-weighted (D2)	Over-weighted (D3)
Operating model witho	ut process error in selectivity				
SO	MO	h0	0.05	0.05	0.05
S0	M0	h1	0.11	0.09	0.13
S0	M1	h0	0.05	0.05	0.05
S1	M0	h0	0.05	0.05	0.05
S1	M0	h1	0.11	0.09	0.13
S1	M1	h0	0.05	0.05	0.05
S2	M0	h0	0.05	0.05	0.05
S2	M0	h1	0.11	0.09	0.13
S2	M1	h0	0.05	0.05	0.05
Operating model with	process error in selectivity				
SO	MO	h0	0.05	0.12	0.26
S0	M0	h1	0.40	0.45	0.51
S0	M1	h0	0.06	0.11	0.26
S1	M0	h0	0.05	0.04	0.06
S1	M0	h1	0.18	0.11	0.15
S1	M1	h0	0.05	0.04	0.06
S2	MO	h0	0.05	0.04	0.06
S2	MO	h1	0.12	0.11	0.15
S2	M1	h0	0.05	0.04	0.05

tional data weighting. Put simply, analysts should be aware that they cannot weight their way out of a misspecified model!

Although purely objective methods for determining compositional data weighting and process error specification would be highly desirable for stock assessment analysts, we pragmatically assert that the specifics of any particular assessment may suggest tempering default approaches with a more subjective approach. There is a cost to down-weighting compositional data, except if the rest of the model is perfectly specified. Where uncertainty is likely to exist in other scaling parameters (such as natural mortality and steepness) excessive down-weighting should be avoided.

Our results suggest that data weighting and the treatment of process error should be considered together: estimation is most likely to be robust when process error is allowed (even if overestimated) and when data are not excessively down-weighted. We recognize that the population dynamics simulated in this study are simple, and not likely to be specifically representative of individual species or life-history groups for which assessments may be conducted. As such, a general simulation is no substitute for careful examination of model performance given a particular configuration of observed data and life-history characteristics. Further, the approach taken here considers only several simple estimation models; more complex models may exhibit differing behavior and should also be explored in future studies. However, we suspect the trends across treatment of data and process error may be similar. Our results should serve as a starting point for analysts conducting assessments: they provide general conceptual guidance for an approach when neither the true degree of process error, nor the correct data weighting is known.

Acknowledgements

We thank the participants in the CAPAM data weighting workshop for stimulating discussion on the topic. We also acknowledge the developers of ss3sim, who provided a very useful tool for rapid investigation of simple fisheries problems through simulation. CCM was funded through the National Sea Grant/NOAA Fisheries Population Dynamics Fellowship. This publication is partially funded by the Joint Institute for the Study of the Atmosphere and Ocean (JISAO) under NOAA Cooperative Agreement No. NA100AR4320148, Contribution No. 2683. This manuscript was improved based on helpful comments from G.G. Thompson and one anonymous reviewer.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.fishres.2016.06. 018.

References

- Anderson, S.C., Monnahan, C.C., Johnson, K.F., Ono, K., Valero, J.L., 2014. ss3sim: an R package for fisheries Stock Assessment simulation with stock synthesis. PLoS One 9, e92725.
- Anderson, S.C., Monnahan, C.C., Johnson, K.F., Ono, K., Valero, J.L., Cunningham, C.J., Hicks, A.C., Hurtado-Ferro, F., Kuriyama, P.T., Licandeo, R., McGilliard, C.R., Rudd, M.B., C.C., Stawitz, Szuwalski, C.S., Taylor, I.G., Vert-Pre, K.A., Whitten, A.R., ss3sim: fisheries stock assessment simulation testing with stock synthesis. R package version 090.
- Deroba, J.J., Butterworth, D.S., Methot, R.D., De Oliveira, J.A.A., Fernandez, C., Nielsen, A., Cadrin, S.X., Dickey-Collas, M., Legault, C.M., Ianelli, J., Valero, J.L., Needle, C.L., O'Malley, J.M., Chang, Y.J., Thompson, G.G., Canales, C., Swain, D.P., Miller, D.C.M., Hintzen, N.T., Bertignac, M., Ibaibarriaga, L., Silva, A., Murta, A., Kell, L.T., de Moor, C.L., Parma, A.M., Dichmont, C.M., Restrepo, V.R., Ye, Y., Jardim, E., Spencer, P.D., Hanselman, D.H., Blaylock, J., Mood, M., Hulson, P.J.F., 2014. Simulation testing the robustness of stock assessment models to error: some results from the ICES strategic initiative on stock assessment methods. ICES J. Mar. Sci. 72, 19–30.
- Fournier, D.A., Archibald, C.P., 1982. A general theory for analyzing catch at age data. Can. J. Fish Aquat. Sci. 39, 1195–1207.
- Fournier, D.A., Skaug, H.J., Ancheta, J., Ianelli, J., Magnusson, A., Maunder, M.N., Nielsen, A., Sibert, J.A.D., 2012. Model Builder: using automatic differentiation for statistical inference of highly parameterized complex nonlinear models. Optim. Methods Softw. 27, 233–249.
- Francis, R.I.C.C., 2011. Data weighting in statistical fisheries stock assessment models. Can. J. Fish Aquat. Sci. 68, 1124–1138.
- Hilborn, R., Walters, C.J., 1992. Quantitative Fisheries Stock Assessment: Choice, Dynamics and Uncertainty. Chapman and Hall, London.
- Hurtado-Ferro, F., Szuwalski, C.S., Valero, J.L., Anderson, S.C., Cunningham, C.J., Johnson, K.F., Licandeo, R., McGilliard, C.R., Monnahan, C.C., Muradian, M.L., Ono, K., Vert-Pre, K.A., Whitten, A.R., Punt, A.E., 2014. Looking in the rear-view mirror: bias and retrospective patterns in integrated, age-structured stock assessment models. ICES J. Mar. Sci. 72, 99–110.
- Johnson, K.F., Monnahan, C.C., McGilliard, C.R., Vert-pre, K.A., Anderson, S.C., Cunningham, C.J., Hurtado-Ferro, F., Licandeo, R.R., Muradian, M.L., Ono, K., Szuwalski, C.S., Valero, J.L., Whitten, A.R., Punt, A.E., 2015. Time-varying natural mortality in fisheries stock assessment models: identifying a default approach. ICES J. Mar. Sci. 72, 137–150.
- Kimura, D.K., 1990. Approaches to age-structured seperable sequential population analysis. Can. J. Fish Aquat. Sci. 47, 2364–2374.
- Linton, B.C., Bence, J.R., 2008. Evaluating methods for estimating process and observation error variances in statistical catch-at-age analysis. Fish. Res. 94, 26–35.
- Martell, S., Stewart, I., 2014. Towards defining good practices for modeling time-varying selectivity. Fish. Res. 158, 84–95.

I.J. Stewart, C.C. Monnahan / Fisheries Research xxx (2016) xxx-xxx

- Maunder, M.N., Harley, S.J., 2011. Using cross validation model selection to determine the shape of nonparametric selectivity curves in fisheries stock assessment models. Fish. Res. 110, 283–288.
- Maunder, M.N., Piner, K.R., 2014. Contemporary fisheries stock assessment: many issues still remain. ICES J. Mar. Sci. 72, 7–18.
- Maunder, M.N., Punt, A.E., 2013. A review of integrated analysis in fisheries stock assessment. Fish. Res. 142, 61–74.
- Maunder, M.N., Crone, P.R., Valero, J.L., Semmens, B.X., 2014. Selectivity: theory, estimation, and application in fishery stock assessment models. Fish. Res. 158, 1–4.
- Maunder, M.N., 2011. Review and evaluation of likelihood functions for composition data in stock-assessment models: estimating the effective sample size. Fish. Res. 109, 311–319.
- Megrey, B.M., 1989. Review and comparison of age-structured stock assessment models from theoretical and applied points of view. Am. Fish. Soc. Symp. 6, 8–48.
- Methot, R.D., Taylor, I.G., 2011. Adjusting for bias due to variability in estimated recruitments in fishery assessment models. Can. J. Fish. Aquat. Sci. 68, 1744–1760.
- Methot Jr, R.D., Wetzel, C.R., 2013. Stock synthesis: a biological and statistical framework for fish stock assessment and fishery management. Fish. Res. 142, 86–99.
- Methot, R.D., 2015. User Manual for Stock Synthesis. Model Version 3.24s. NOAA Fisheries, Seattle, WA, pp. 152.
- Monnahan, C.C., Ono, K., Anderson, S.C., Rudd, M.B., Hicks, A.C., Hurtado-Ferro, F., Johnson, K.F., Kuriyama, P.T., Licandeo, R.R., Stawitz, C.C., Taylor, I.G., Valero, J.L., 2015. The effect of length bin width on growth estimation in integrated age-structured stock assessments. Fish. Res.

- Ono, K., Licandeo, R., Muradian, M.L., Cunningham, C.J., Anderson, S.C., Hurtado-Ferro, F., Johnson, K.F., McGilliard, C.R., Monnahan, C.C., Szuwalski, C.S., Valero, J.L., Vert-Pre, K.A., Whitten, A.R., Punt, A.E., 2015. The importance of length and age composition data in statistical age-structured models for marine species. ICES J. Mar. Sci. 72, 31–43.
- Punt, A.E., Hurtado-Ferro, F., Whitten, A.R., 2014. Model selection for selectivity in fisheries stock assessments. Fish. Res. 158, 124–134.
- Punt, A.E., 2016. Some insights into data weighting in integrated stock assessments. Fish. Res.
- Quinn, T.J.I., Deriso, R.B., 1999. Quantitative Fish Dynamics. Oxford University Press, New York, pp. 542.
- R Core Team, 2016. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria http://www.R-project. org/.
- Sampson, D.B., Scott, R.D., 2012. An exploration of the shapes and stability of population-selection curves. Fish Fish. 13, 89–104.
- Sampson, D.B., 2014. Fishery selection and its relevance to stock assessment and fishery management. Fish. Res. 158, 5–14.
- Stewart, I.J., Hamel, O.S., 2014. trapping of sample sizes for length- or age-composition data used in stock assessments. Can. J. Fish. Aquat. Sci. 71, 581–588.
- Stewart, I.J., Martell, S.J.D., 2014. A historical review of selectivity approaches and retrospective patterns in the Pacific halibut stock assessment. Fish. Res. 158, 40–49.
- Thompson, G.G., Lauth, R.R., 2012. Assessment of the Pacific cod stock in the Eastern Bering Sea and Aleutian Islands area. NPFMC Bering Sea and Aleutian Islands SAFE, 245–544 (Chapter 2).
- Thorson, J.T., Taylor, I.G., 2014. A comparison of parametric, semi-parametric, and non-parametric approaches to selectivity in age-structured assessment models. Fish. Res. 158, 74–83.

Methods in Ecology and Evolution 2017, 8, 339-348

RRITICI

doi: 10.1111/2041-210X.12681

COMMENTARY Faster estimation of Bayesian models in ecology using Hamiltonian Monte Carlo

Cole C. Monnahan¹*, James T. Thorson² and Trevor A. Branch³

¹Quantitative Ecology and Resource Management, University of Washington, Box 352182, Seattle, WA 98195, USA; ²Fisheries Resource Assessment and Monitoring Division, Northwest Fisheries Science Center, National Marine Fisheries Service, National Oceanic and Atmospheric Administration, 2725 Montlake Blvd. East, Seattle, WA 98112, USA; and ³School of Aquatic and Fishery Sciences, University of Washington, Box 355020, Seattle, WA 98195, USA

Summary

1. Bayesian inference is a powerful tool to better understand ecological processes across varied subfields in ecology, and is often implemented in generic and flexible software packages such as the widely used BUGS family (BUGS, WinBUGS, OpenBUGS and JAGS). However, some models have prohibitively long run times when implemented in BUGS. A relatively new software platform called Stan uses Hamiltonian Monte Carlo (HMC), a family of Markov chain Monte Carlo (MCMC) algorithms which promise improved efficiency and faster inference relative to those used by BUGS. Stan is gaining traction in many fields as an alternative to BUGS, but adoption has been slow in ecology, likely due in part to the complex nature of HMC.

2. Here, we provide an intuitive illustration of the principles of HMC on a set of simple models. We then compared the relative efficiency of BUGS and Stan using population ecology models that vary in size and complexity. For hierarchical models, we also investigated the effect of an alternative parameterization of random effects, known as non-centering.

3.. For small, simple models there is little practical difference between the two platforms, but Stan outperforms BUGS as model size and complexity grows. Stan also performs well for hierarchical models, but is more sensitive to model parameterization than BUGS. Stan may also be more robust to biased inference caused by pathologies, because it produces diagnostic warnings where BUGS provides none. Disadvantages of Stan include an inability to use discrete parameters, more complex diagnostics and a greater requirement for hands-on tuning.

4. Given these results, Stan is a valuable tool for many ecologists utilizing Bayesian inference, particularly for problems where BUGS is prohibitively slow. As such, Stan can extend the boundaries of feasible models for applied problems, leading to better understanding of ecological processes. Fields that would likely benefit include estimation of individual and population growth rates, meta-analyses and cross-system comparisons and spatiotemporal models.

Key-words: Bayesian inference, hierarchical modelling, Markov chain Monte Carlo, no-U-turn sampler, Stan

Introduction

Bayesian inference is used widely throughout ecology, including population dynamics, genetics, community ecology and environmental impact assessment, among other subfields (Ellison 2004). In the Bayesian paradigm, the likelihood of the observed data is combined with prior distributions on parameters, resulting in a posterior probability distribution of parameters, from which inference is made (Gelman *et al.* 2014). Expectations of posterior quantities, such as means or quantiles, are commonly approximated using numerical techniques, with Markov chain Monte Carlo (MCMC) being the most common (Brooks *et al.* 2011).

The popularity of Bayesian inference grew particularly fast with the development of generic and flexible software platforms, with the BUGS family (here defined as BUGS, WINBUGS, OPENBUGS and JAGS; see Appendix A, Supporting Information) being by far the most common (Fig. 1). For a given model, BUGS automatically selects an MCMC algorithm and arguments controlling its behaviour (i.e. tuning parameters), where necessary. The analyst can thus focus on the model and scientific questions, rather than the mechanics of the underlying MCMC algorithms. As such, these platforms have been the workhorse for Bayesian analyses in ecology and other fields for the last 20 years.

However, for certain models, the time required for inference (run-time) using BUGS is prohibitively long. Long run-times

^{*}Correspondence author. E-mail: monnahc@uw.edu

[[]Correction note: The abstract was originally omitted from this article when it was first published 14 November 2016, the abstract was added on 16 January 2017.]



Fig. 1. Citation patterns of Stan and the BUGS family of Bayesian software platforms, for all journals in all fields. Data are from ISI Web of Science Core Collection. The *y*-axis units are the same, despite variable ranges.

often occur in BUGS because the underlying MCMC algorithms are inefficient, which is further compounded when the model needs to run many times during development, model selection (e.g. cross-validation; Hooten & Hobbs 2015), or simulation testing. These issues remain despite the increasing power of computers because data sets are increasing in size and models are becoming more complex (Bolker *et al.* 2013). At the same time, hierarchical modelling is becoming increasingly popular, as this type of model is widely recognized as a natural tool for formulating and thinking about problems in many ecological subfields (Royle & Dorazio 2008; Cressie *et al.* 2009; Thorson & Minto 2014). Thus, there is a need for alternatives to BUGS that are faster across a range of model size, complexity and hierarchical structure.

A family of MCMC algorithms called Hamiltonian Monte Carlo (HMC; Neal 2011) promises improved efficiency over the algorithms used by BUGS, but until recently have been slow to be adopted for two reasons. First, HMC requires precise gradients (i.e. derivatives of the log-posterior with respect to parameters), but analytical formulas are rare and numerical techniques are imprecise, particularly in higher dimensions. Secondly, the original HMC algorithm requires expert, handson tuning to be efficient (Neal 2011). Both of these hurdles have recently been overcome, the first with automatic differentiation (e.g. Griewank 1989) and the second with an HMC algorithm known as the no-U-turn sampler (NUTS; Hoffman & Gelman 2014). These advances have been packaged into the open-source, generic and flexible modelling software Stan (Gelman, Lee & Guo 2015; Stan Development Team 2016, Carpenter et al. in press), which effectively aims to replace the

BUGs family and is quickly gaining traction across diverse fields (Fig. 1).

Despite the potential of HMC, and the availability of Stan, adoption has been slow in ecology, likely because ecologists are either unaware of its existence, or are unsure when it should be preferred over BUGS. Here, we illustrate the principles that underlie HMC and then compare the efficiency between Stan and a BUGS variant, JAGS (Plummer 2003), across a range of models in population ecology. Specifically, we test how HMC performance scales with model size and complexity, and its suitability for hierarchical models. Our goal is to explore the relative benefits of Stan and JAGS and to provide guidance for ecologists looking to use the power of HMC for faster and more robust Bayesian inference.

Principles of Hamiltonian Monte Carlo

The existing literature on HMC tends to focus on mathematical proofs of statistical validity and is accessible primarily to statisticians. We therefore first illustrate the principles of HMC using simple models, and contrast it with other MCMC algorithms.

Markov chain Monte Carlo algorithms sequentially generate posterior samples (i.e. vectors containing a value for each parameter), resulting in a finite number of autocorrelated samples which are used for inference (Gelman *et al.* 2014). Many algorithms *transition* between samples by proposing a new sample, based on the current sample and tuning parameters, and then accept it with known probability. If rejected, the current iteration is the same as the previous one.

For example, the widely used random walk Metropolis algorithm (Metropolis *et al.* 1953) typically proposes a multivariate normal sample, centered at the current sample and uses the proposed to current posterior density ratio to determine the acceptance probability. In this case, all parameters are proposed and updated simultaneously, and the covariance of the proposal distribution is tuned to achieve an optimal acceptance rate (Roberts & Rosenthal 2001). Other algorithms update a single parameter at a time, looping through each within a transition. This is the behaviour typically used by BUGS, which uses Gibbs sampling if possible, and alternatives if not.

If an algorithm cannot propose samples in regions of the posterior distant to the current state, then it exhibits random walk behaviour: multiple transitions are necessary to move between regions, leading to higher autocorrelation and slow mixing. HMC avoids this inefficient random walk behaviour because it can propose values (almost) anywhere in the posterior from anywhere else. It does this using a physical system known as Hamiltonian dynamics.

HAMILTONIAN DYNAMICS

A Hamiltonian system can be conceptualized as a ball moving about a frictionless surface over time (e.g. imagine a marble inside a large bowl). The ball is affected by gravity and its own momentum: gravity pulls it down while momentum keeps it



Fig. 2. Basics of Hamiltonian dynamics. (a) An example where a ball is dropped from the black point, it rolls down the surface over time (t), and momentum carries it up the other side where it reverses direction (red line), returning to where it started. The lines are offset to distinguish black and red paths. The position and momentum variables (b) and energies (c) over time corresponding to the path in (a). (d) Multiple paths for a 2d parabola. Grey dashed lines show posterior contours; initial positions and paths are red arrows and black lines. (e) Partial path (black line) on a posterior of a logistic population model with intrinsic growth rate (r) and carrying capacity (K). Red arrow shows initial position. (f) The energies for the trajectory in (e).

going in the same direction. A set of differential equations govern the movement of the ball over time (its *path*).

There are some important concepts associated with the ball. The *position* of the ball is its coordinate vector (i.e. where it is on the surface) and associated with each position variable is a *momentum* variable. The *potential energy* is the height of the surface at a given position. The *kinetic energy* is related to the momentum, assumed for now to be the sum of the squared momenta. Because the surface is frictionless, the total energy (potential plus kinetic), known as the *Hamiltonian* (*H*), remains constant over time. Later, we will see that, in the context of MCMC, the position vector corresponds to the model parameters and the potential energy to the negative log of the posterior density.

For now, consider the parabola $y = x^2$ (Fig. 2a), which has a single position variable (x) and thus a single momentum variable. We place the ball at position x = -1 and height (potential energy) y = 1, and let it go such that it has no initial momentum or kinetic energy. Gravity pulls it down, building speed over time as potential energy is converted to kinetic energy (Fig. 2b,c). Momentum carries it past position x = 0, where all potential energy has been converted into kinetic energy. As there is no friction, it stops exactly at x = 1 and y = 1, where the potential and kinetic energies return to their initial states (Fig. 2c). At this point, it will reverse course (Fig. 2a–c red lines) and oscillate forever with the energies varying but their sum (*H*) remaining constant.

Now consider a 2D parabola, $y = x_1^2 + x_2^2$ (i.e. a bowl shape). The position and momentum vectors are of length two, but the kinetic and potential energies are scalars. We place the ball as before, but this time we flick it, imparting momentum with a direction and magnitude (Fig. 2d). If flicked sideways, it will move in a circle of constant height. If flicked straight down, it will cross the bottom and go up the other side. An elliptical path occurs when flicking the ball at a downward angle. A more complex surface typical of a real model, such as a logistic growth model (see 'Case studies' below), leads to more complex paths (Fig. 2e,f), but



which obey the same principles and intuition as these simple examples.

The principles of Hamiltonian dynamics relate directly to MCMC by providing a way to generate efficient transitions. The ball could move (almost) anywhere given the right length of time and initial momentum, thus providing transitions with directed movement and avoiding inefficient random walk behaviour. MCMC algorithms that utilize Hamiltonian dynamics are generally referred to as HMC, and we briefly review two: static HMC and NUTS.

STATIC HMC

Static HMC was the first MCMC algorithm to utilize Hamiltonian dynamics (Duane *et al.* 1987). Although replaced by more advanced algorithms, static HMC is simpler to explain and contains most of the properties relevant for understanding NUTS. A static HMC transition occurs by simulating the ball from the current position with random momenta for a finite

Fig. 3. Examples demonstrating the basics of HMC. (a) The effect of different step sizes (ɛ) and number of steps (L) on trajectories. The blue and red trajectories approximate the same path (solid grey line), with the same initial position (red point) and trajectory length (EL), but opposite momentum. (b) Trajectories on a logistic posterior surface with identical initial position (black point) and momentum vectors. The black trajectory is slow to traverse the surface, while the red trajectory shows accumulating approximation errors, causing it to diverge. The blue trajectory utilizes a mass matrix, making the surface easier to traverse. (c) Multiple iterations of static HMC; black points are and accepted and intermediate steps (grey arrows) are discarded. (d) The acceptance ratios (α) of the trajectories in (b), with corresponding acceptance probability of $min(1, \alpha)$. Multiple draws from the same initial position using a random walk Metropolis (e) or NUTS (f) algorithm, with and without an appropriate mass matrix (colours).

length of time and proposing the state (position) at the end of this simulated, finite path.

However, three issues complicate this process. The first is how to simulate movement on arbitrary log-posteriors (i.e. generate paths). Simple models like a parabola have analytical solutions to the underlying differential equations; thus, exact, continuous paths are possible. However, for most models, the continuous paths must be approximated using a numerical method known as the leapfrog integrator (we refer to approximated paths as *trajectories*). A trajectory depends on the *step size* (ε) and the *number of steps* (*L*; Fig. 3a,b). The position vector at step *L* is the proposed sample for that transition, while the intermediate steps are discarded (Fig. 3c). Approximation errors cause the ball to deviate from the continuous path, and thus, *H* is not constant over time (Fig. 3d).

The next challenge is determining the optimal *trajectory length* (i.e. εL). If the trajectory length is too short, distant proposals are impossible, leading to an inefficient random walk. If

it is too long, the trajectory will retrace its steps (e.g. Fig. 3a), which is wasteful computationally. Thus, efficiency depends on the trajectory length, but the optimal length is difficult to determine and a crucial tuning step required for static HMC (Betancourt 2016b).

The last issue is determining the step size, given a trajectory length. The same length can be attained by taking fewer steps of larger size, or more steps of smaller size (Fig. 3a,b). As each step is computationally costly, the fewer the steps the faster the transition. However, there is a downside to large step sizes: they lead to more variation in H, and in some cases, the approximation error accumulates such that the total energy (H) goes to infinity, known as a divergent transition (red trajectory, Fig. 3b). A Metropolis acceptance step accounts for variation in H by accepting the proposed state with probability $\min(1, \alpha)$, where α is the exponential of the energy lost. Thus, proposals are always accepted if the total energy has decreased, whereas increased energy is accepted with a probability <1 (Fig. 3d). Increasing the step size reduces run-time, but increases approximation error, leading to more rejected states and divergent transitions, degrading the efficiency of the algorithm. Optimizing the step size is thus another crucial step in static HMC (Betancourt, Byrne & Girolami 2014a).

Given a step size and number of steps, the last step is to specify a kinetic energy function. In HMC, it is typically the log density of a multivariate normal random vector where the covariance matrix is known as the *mass matrix*. Previously, we assumed the kinetic energy was the sum of the squared momenta, corresponding to an identity mass matrix. The effect of the mass matrix is to globally transform the posterior to have a simpler geometry for sampling. The variances stretch the posterior so all parameters have the same scale, while the covariances rotate it so they are approximately independent. When successful, the transformed parameters have a scale of 1 and no correlations, resembling iid standard normal random variables (blue trajectory, Fig. 3b.)

The mass matrix is analogous to the covariance of the proposal function sometimes used in Metropolis-Hastings samplers, which can have substantial impacts on sampling (Fig. 3e). Depending on the model, HMC algorithms can be efficient with an identity mass matrix (Fig. 3f), but it will require more leapfrog steps per transition and more time (Fig. 3b). Thus, to get efficient sampling with HMC, the mass matrix should approximate the covariance of the posterior, but this information is often not known *a priori*.

Specifying an optimal trajectory length, step size and mass matrix is critical for static HMC to work efficiently, leading it to require expert hands-on tuning and *a priori* knowledge (Neal 2011). Fortunately, NUTS automates this process and provides efficient sampling with minimal or no tuning.

THE NO-U-TURN SAMPLER

No-U-turn sampler extends static HMC by automating tuning: neither the step size nor number of steps need be specified by the user. NUTS determines the number of steps via a sophisticated tree building algorithm, which we briefly describe here. A single NUTS trajectory is built by iteratively accumulating steps. In the first iteration, a single leapfrog step is taken from the current state so the trajectory has a total of two steps. Then, two more steps are added (total of four), then four more (total of eight), and so forth, with each iteration doubling the length of the trajectory. This doubling procedure repeats until the trajectory turns back on itself and a 'U-turn' occurs, or the trajectory diverges (i.e. *H* goes to infinity). The number of doublings is known as the *tree depth*. The key aspect of this tree building algorithm is that it automatically creates trajectories that are neither too short nor too long. In practice, this means trajectory lengths vary among transitions: it may take eight steps or 128, depending on the position and momentum vectors.

The no-U-turn sampler determines the step size by adapting it during the warm-up (burn-in) phase to a target acceptance rate (adapt_delta in Stan). The tuned step size is then used for all sampling iterations. In contrast to static HMC, NUTS does not use a Metropolis acceptance step, so an analogous statistic is used for adaptation. Betancourt, Byrne & Girolami (2014a) found this target acceptance rate should generally be between 0.6 and 0.9, with larger values being more robust in practice. Thus, NUTS effectively reduces static HMC to a single, user-specified tuning parameter: the target acceptance rate.

HMC IN PRACTICE

One disadvantage of HMC is that, unlike BUGS, only continuous parameters are possible because discrete parameters do not have gradients. A manual implementation could overcome this by alternating Gibbs updates and HMC (Neal 2011), and future versions of Stan may implement such a scheme. Alternatively, in some cases, they can be marginalized out manually by the user (Chapter 10 and 12, Stan Development Team 2016).

Another disadvantage is that HMC is developed using sophisticated mathematics and statistics (e.g. Betancourt *et al.* 2014b), making it difficult to develop a deep understanding or intuition about their behaviour. We provide implementations of the static HMC and NUTS algorithms, written in R (R Core Team 2016), in Appendix B. We encourage the interested reader to experiment with the samplers to further their understanding of HMC, while using the faster and more robust Stan implementation for inference of real problems.

No-U-turn sampler (and static HMC) is similar to other MCMC algorithms: valid inference is conditioned on a converged chain, but this is impossible to prove (Gelman *et al.* 2014). The analyst is responsible for assessing convergence before making inference, and for NUTS, this includes assessing adaptation. Information about step size, tree depths and mass matrix quantities are reported in the output of a Stan run, and they should be checked routinely. For example, the adapted step size should be consistent across multiple chains, postwarm-up divergences should be minimized (by increasing target acceptance rate) and the maximum tree depth increased if necessary. The user manual (Stan Development Team 2016) has more information, advice on fitting strategies and details of the adaptation procedure for the mass matrix and step size.

Key concepts that arise when using NUTS in Stan are summarized briefly below:

• Smaller step sizes have higher acceptance rates, but require more steps and thus time. Larger step sizes reject more states and can have more divergences. The optimal step size depends on the model and is tuned to achieve a *target acceptance rate* set by the user (adapt_delta), defaulting to 0.8, but higher values needed for more difficult posteriors.

• The number of steps is determined dynamically for each transition using a tree building algorithm, where the trajectory repeatedly doubles in length until a U-turn occurs. The number of doublings is known as the *tree depth*.

• If the mass matrix approximates the covariance of the posterior, the algorithm 'sees' a simpler surface and is more efficient. By default only the diagonal terms are estimated, accounting for differences in scales, but not correlations, between parameters. Mass matrices with nonzero covariance terms, referred to as *dense*, are available in Stan but are not commonly used.

• The optimal step size depends on the mass matrix, and the mass matrix cannot be well estimated without sampling from the entire posterior, which requires a reasonable step size. Thus, sufficiently long warm-ups are needed for effective adaptation and efficient sampling.

Case studies

We tested the efficiency of Stan and JAGS for simulated and empirical models from population ecology. To quantify efficiency, we used the minimum number of effective samples per unit time, $E = \hat{N}_{\text{ESS}}/t$, a standard approach to compare among algorithms and software platforms. Further details of how this was calculated can be found in Appendix C. This definition of efficiency (*E*) can be roughly thought of as the number of independent samples generated per unit time. We used matching parameterizations for Stan and JAGS, but explored two parameterizations for each hierarchical model and platform. MCMC efficiency for hierarchical models depends on the random effect parameterization, with the *centered* and *non-centered* complementary forms being useful for a broad class of models (Papaspiliopoulos, Roberts & Skold 2007; Betancourt & Girolami 2015). Briefly, the centered form models the random effects (τ) directly: $\tau \sim N(\mu, \sigma^2)$, while the non-centered form does it indirectly by letting $\tau = \mu + \sigma Z$, where $Z \sim N(0, 1)$ are the model parameters and implying $\tau \sim N(\mu, \sigma^2)$. See Appendix D for further information and references. We test both forms because the most efficient can depend on the amount of information about σ .

Initial values, random seeds and length of adaptation can have large impacts, particularly for HMC, so we ran 20 chains of length 40 000 without thinning, initialized from a random sample from a previously run long chain. We used the first half of each chain as a warm-up, discarding those samples but including warm-up time (but not compilation time) in the total run-time. We also did not include time to tune the target acceptance rate for Stan, as the analyst will often determine acceptable tuning parameters during model development. We used default settings for JAGS and Stan, except increasing the target acceptance rate from its default of 0-8 where needed (see Appendix E). We checked convergence, as is typically carried out for MCMC output, such as the potential scale reduction, \hat{R} , being close to 1 (Gelman *et al.* 2014), in addition to the specific diagnostics for NUTS.

Our tests included two simulated models and four models with real data (Table 1). The simulated models were a multivariate normal with random covariances (MVND) or repeated correlations (MVNC), both of which were easy to vary in the number of fixed effects and covariance structure. Our simulated nonlinear mixed effects somatic Growth model varied in

Table 1. Summary of case studies used to compare efficiency between Stan and JAGS. Further details are available in Appendix E. Latent parameters are those modelled as random effects

Model name	Description	Data	Parameters (Latent)	Hierarchical structure	Reference
MVND	Multivariate normal with covariances generated from inverse Wishart	Simulated	Varies: 2–200	None	Simulated
MVNC	Multivariate normal with all off- diagonals set to ρ	Simulated	Varies: 5–50	None	Simulated
Growth	Nonlinear somatic growth with repeated measures	Lengths at age	Varies: 16-406 (10-400)	Normal on growth rate and maximum length, in log space	Simulated; see Schnute (1981)
Redkite	Age-dependent survival probabilities	Mark–recapture of birds	5	None	Section 8.4 of Kéry & Schaub (2012)
Swallows	State-space survival and detection with environmental covariates	Mark–recapture of birds	177 (172)	Year and family effects for survival, family effects for detection	Section 14.5 of Korner-Nievergelt <i>et al.</i> (2015)
Logistic	State-space fisheries logistic population dynamics	Annual catch per unit effort; catches	28 (22)	Annual biomass dynamics deviations	Millar & Meyer (2000)
Wildflower	Binomial generalized linear model of flowering success	Stages, flower, and seed pod production	1101 (1072)	Year effects on intercept; crossed effects on intercept and slope for covariate	Bolker <i>et al.</i> (2013)



Fig. 4. Comparison of efficiency (*E*) for Stan and JAGS across simulated models. The means (points) and ranges (segments) are across 20 replicates. (a) A multivariate normal with increasing dimensionality (MVND), either independent or with random correlations from an inverse Wishart distribution. Ranges are too narrow to be visible. (b) A multivariate normal with repeated correlations on the off-diagonals for varying dimensions (MVNC). (c) A nonlinear mixed effects model with two latent parameters per individual (Growth); ranges were left out for visual clarity.

the number of individuals. The first two real data models were fit to mark–recapture data of birds and differed in their size and complexity: the Redkite model only estimates survival while the Swallows model estimates survival and detection probabilities using environmental covariates in a complex hierarchical state-space formulation. We also fit a state-space Logistic population dynamics model to fisheries data to estimate temporal trends in abundance. Lastly, our Wildflower model was a generalized linear mixed effects model with crossed random effects estimating flowering success. The case studies ranged from 5 to 1101 parameters and were a mixture of hierarchical and non-hierarchical models. Further details can be found in Appendix E, and model files for both Stan and JAGS in Appendix B. We did our analyses using R and the packages RSTAN and RJAGS.

Results

For the multivariate normal models (MVND and MVNC), the run-time of JAGS increased at a faster rate than Stan with increasing number of parameters, although the minimum effective sample size for a given run was similar between the two software platforms. Stan was more efficient by several orders of magnitude because its run-time for each sample was faster, and increasingly better with more parameters (Fig. 4a,b). For the growth model, Stan consistently outperformed JAGS at higher dimensions for both parameterizations. However, Stan had more variable efficiencies than JAGS with fewer individuals.

Stan was more efficient for the real-world models as well (Table 2), up to 63 times for the Logistic model in the noncentered form. JAGS was faster for the centered Swallows and Wildflower models, but for both the non-centered Stan model was the fastest option overall. Thus, Stan was faster for all models (using the optimal parameterization), although the variability in Stan's efficiency tended to be higher than for JAGS (results not shown), likely reflecting HMC's sensitivity to tuning compared to other algorithms.

We also found clear differences between software platforms in the effect of the parameterization for hierarchical models. For Stan, the non-centered form was consistently faster than the centered form for models with real data: 4-3 times faster for the Logistic, 2-8 times for the Swallows and 129 times for the Wildflower model. In contrast, JAGS was slower for all three: 0-90, 1-00 and 0-67, respectively. For the simulated Growth

Table 2. Case study results comparing efficiency of Stan and JAGS. Max correlation is the largest absolute pairwise correlation, calculated from converged samples. Efficiency (E) is the number of effective samples per time

Model	Random effects parameterization	Max correlation	Median $E_{\rm stan}$	Median E_{jags}	Median $E_{\text{stan}}/E_{\text{jags}}$ (Range)
Redkite	NA	0.83	1102.85	302.99	3.54 (1.14-10.03)
Logistic	Centered	0.96	12.35	0.98	12.2 (7.88-34.54)
Logistic	Non-centered	0.96	53.60	0.88	63.33 (18.25-132.02)
Swallows	Centered	0.90	0.12	0.10	0.94 (0-2.96)
Swallows	Non-centered	0.81	0.34	0.10	2.4 (0.1-10.04)
Wildflower	Centered	0.96	0.01	0.06	0.14(0.02-1.03)
Wildflower	Non-centered	0.96	1.29	0.04	34.2 (13.11-60.7)



Fig. 5. Effects of non-centering on divergences and bias for the random effects on growth rate in the Growth model with 10 individuals. τ is the deviation from the mean for an arbitrary individual and the parameters in the centered model, σ its standard deviation and $Z \sim N(0, 1)$ the parameters in the noncentered model. Samples from: (a) the centered model (target acceptance rate $\delta = 0.95$); (b) the non-centered model ($\delta = 0.80$); and (c) the transformed non-centered parameters, $\tau = \sigma Z$. Divergences in (a), shown in red, arise because the adapted step size is too large for the high gradients at low σ , creating an inaccessible region and leading to biased σ (i.e. no samples below $\log \sigma = -6$). The non-centered parameterization eliminates the curvature and hence the divergences and bias (c). (d) Median rate of divergent transitions using $\delta = 0.80$ for both parameterizations. As information increases about σ (i.e. more individuals) the marginal distribution of σ narrows, simplifying the geometry and lowering the rate of divergences.

Table 3. Summary of key differences between JAGS and Stan

	JAGS	Stan
Inference	Bayesian only (MCMC)	Bayesian (MCMC with NUTS and variational inference) and penalized maximum likelihood
Tuning	Automatic with no options	Automatic with options for target acceptance rate (adapt_delta), mass matrix (diagonal or dense)
Discrete parameters	Use directly	Incompatible-must be marginalized out analytically
General pros	Easy to use, no tuning, discrete parameters	Scales well with dimensionality, posterior complexity; suitable for hierarchical models, especially the non-centered form
General cons	Few alternatives to reduce run-time when prohibitively slow	No discrete parameters, more difficult modelling language and additional MCMC diagnostics to check
Potential pathologies	No feedback	Divergences and excessive tree depths warn of steep or flat curvature, respectively

model, the non-centered form was faster for Stan, but slower for JAGS across all dimensionalities (Fig. 4c).

Discussion

Hamiltonian Monte Carlo is a family of MCMC algorithms which utilizes the posterior geometry and properties of Hamiltonian dynamics to make directed MCMC transitions, minimizing the inefficient random walk behaviour that degrades the performance for many algorithms used by JAGS. HMC is available to ecologists in the form of Stan, a generic and flexible software package with a similar workflow to JAGS. Here, we demonstrated that Stan outperformed JAGS for all simulated and real-world models from population ecology across a range of dimensions and complexity. Stan was more sensitive to the parameterization of the random effects, suggesting analysts use non-centered parameterizations to improve performance (Appendix D).

Our findings corroborate studies from other fields (e.g. Grant *et al.* 2016), but come with caveats when trying to extrapolate. For example, our simulated models might not reflect nuances in real data, or might not be representative of typical models in other subfields of ecology. Fair comparisons between software are also difficult, because many factors influence performance, including, but not limited to, priors, tuning parameters, length of chains and parameterization chosen. For

instance, a model that is faster in Stan with a specific prior or parameterization may be faster in JAGS with alternatives. Nevertheless, the results from our case studies suggest that Stan will often be more efficient and thus provide faster inference.

Although our focus was on quantifying sampling efficiencies, the software platforms also behave differently for pathological models. Pathologies are properties of the posterior which obstruct an algorithm's ability to explore the entire posterior, resulting in biased inference of quantities of interest (Betancourt 2016a). For instance, posteriors with regions of very low or high curvature (gradients) can be pathological for HMC (section 6.6, Livingstone et al. 2016). Pathologies affect both Stan and JAGS, but Stan naturally diagnoses them: regions of high curvature are identified by divergences, and flat regions by excessive tree depths (Betancourt 2016a). JAGS provides no such feedback, and pathologies may not be apparent using traditional MCMC diagnostics. Pathologies using Stan occur in practice: centered hierarchical models can exhibit biased hypervariances due to high curvature (Fig. 5a). A Stan user can try to eliminate potential bias by reducing the step size, reparameterizing (e.g. non-centering, Fig. 5b-d), changing priors or restructuring their model. Thus, Stan is not only more efficient than JAGS, but it may also provide more robust inference because a user is more likely to detect and eliminate potential biases.

Despite its promise, HMC has some clear disadvantages, with the most critical that discrete parameters are disallowed, such as a discrete latent states or population numbers (e.g. Dail & Madsen 2011). HMC can still be used if the parameters can be marginalized out analytically, as in the binary states of the Swallows model, and this technique is often possible and can also make substantial improvements for JAGS as well (results not shown). HMC is also sensitive to tuning, despite the automation provided by NUTS. For instance, if warm-up periods are too short to effectively explore the entire posterior, then the step size and mass matrix will be suboptimal and efficiency may suffer. Users must also be more involved in assessing tuning for Stan, and be familiar with the principles of HMC to understand diagnostic output from Stan.

There are other HMC algorithms in addition to NUTS, and other gradient-based algorithms for Bayesian inference, which were not tested here. For instance, Riemann Manifold HMC varies the mass matrix along the trajectory (Girolami & Calderhead 2011; Betancourt 2013) and variational inference is a faster alternative to MCMC which approximates the posterior (Kucukelbir et al. 2016). There are also alternative software platforms not tested here, such as NIMBLE (de Valpine et al. 2016) and ensemble sampling (Goodman & Weare 2010), and future work comparing these to JAGS and Stan would be worthwhile. Stan is also not the only platform coupling automatic differentiation and HMC that is used by ecologists. Both AD Model Builder (Fournier et al. 2012) and Template Model Builder (Kristensen et al. 2015) have HMC capabilities, but neither are as well developed or mature as Stan (author CCM is a developer of them). Our results suggest improving HMC capabilities in these software programs would be worthwhile for their user bases.

The preferred software depends on the situation (Table 3), and JAGS will clearly remain a valuable tool when run-time is not prohibitive, but also likely in additional cases such as prototyping models or introducing Bayesian techniques. Stan is clearly the best option for highly parameterized models or smaller models with more difficult geometries (e.g. high or anisotropic correlations). One promising application for HMC is fisheries stock assessment models, which are often extremely large, nonlinear hierarchical models that rarely use Bayesian inference because of prohibitively slow run-times (e.g. Stewart *et al.* 2013). Many other fields likely have similar examples where Bayesian inference is currently infeasible, and we anticipate that HMC will make some of these problems tractable for the first time.

Increasingly large and complex data sets, and powerful software tools, allow analysts to investigate ecological processes which were previously infeasible. Here we demonstrated that Stan, which implements HMC in a flexible modelling platform, is a promising tool when status quo methods such as JAGS are prohibitively slow. We believe Stan should be in the methodological toolbox for every quantitative ecologist because it will extend the boundaries of feasible models for applied problems and lead to better understanding of ecological processes.

Acknowledgements

We thank Bob Carpenter and Michael Betancourt for insights on a variety of conceptual issues and constructive feedback on an earlier draft. Margaret Siple, Eric Buhle, Kevin See, Jim Hastie and two anonymous reviewers provided valuable feedback on an earlier version of this manuscript. This publication is partially funded by the Joint Institute for the Study of the Atmosphere and Ocean (JISAO) under NOAA Cooperative Agreement NA10OAR4320148 (2010–2015) and NA15OAR4320063 (2015–2020), Contribution No. 2016-01-23. This work was partially funded in part by a grant from Washington Sea Grant, University of Washington, pursuant to National Oceanic and Atmospheric Administration Award No. NA14OAR4170078. TAB was also funded by the Richard C. and Lois M. Worthington Endowed Professorship in Fisheries Management. The views expressed herein are those of the authors and do not necessarily reflect the views of NOAA or any of its subagencies.

Data accessibility

We used a combination of real data taken from previously published studies and simulated data using R scripts. All data, R scripts and results files are publicly available at https://github.com/colemonnahan/gradmcmc/tree/v1.0 and are also archived at Monnahan (2016).

References

- Betancourt, M. (2013) Generalizing the no-U-turn sampler to Riemannian manifolds. arXiv preprint arXiv:1304.1920.
- Betancourt, M. (2016a) Diagnosing suboptimal cotangent disintegrations in Hamiltonian Monte Carlo. arXiv preprint arXiv:1604.00695.
- Betancourt, M. (2016b) Identifying the optimal integration time in Hamiltonian Monte Carlo. arXiv preprint arXiv:1601.00225.
- Betancourt, M., Byrne, S. & Girolami, M. (2014a) Optimizing the integrator step size for Hamiltonian Monte Carlo. arXiv preprint arXiv:1411.6669.
- Betancourt, M. & Girolami, M. (2015) Hamiltonian Monte Carlo for hierarchical models. *Current Trends in Bayesian Methodology with Applications*, (eds S. K. Upadhyay, U. Singh, D.K. Dey & A. Loganathan), pp. 79–101. CRC Press, Boca Raton, FL, USA.
- Betancourt, M., Byrne, S., Livingstone, S. & Girolami, M. (2014b) The geometric foundations of Hamiltonian Monte Carlo. arXiv preprint arXiv:1410.5110.

- Bolker, B.M., Gardner, B., Maunder, M. et al. (2013) Strategies for fitting nonlinear ecological models in R, AD Model Builder, and BUGS. *Methods in Ecol*ogy and Evolution, 4, 501–512.
- Brooks, S., Gelman, A., Jones, G. & Meng, X.-L. (2011) Handbook of Markov Chain Monte Carlo. CRC Press, Boca Raton, FL, USA.
- Carpenter, B., Gelman, A., Hoffman, M. et al. (in press) Stan: a probabilistic programming language. Journal of Statistical Software.
- Cressie, N., Calder, C.A., Clark, J.S., Ver Hoef, J.M. & Wikle, C.K. (2009) Accounting for uncertainty in ecological analysis: the strengths and limitations of hierarchical statistical modeling. *Ecological Applications*, **19**, 553– 570.
- Dail, D. & Madsen, L. (2011) Models for estimating abundance from repeated counts of an open metapopulation. *Biometrics*, 67, 577–587.
- Duane, S., Kennedy, A.D., Pendleton, B.J. & Roweth, D. (1987) Hybrid Monte Carlo. *Physics Letters B*, 195, 216–222.

Ellison, A.M. (2004) Bayesian inference in ecology. Ecology Letters, 7, 509-520.

- Fournier, D.A., Skaug, H.J., Ancheta, J., Ianelli, J., Magnusson, A., Maunder, M.N., Nielsen, A. & Sibert, J. (2012) AD Model Builder: using automatic differentiation for statistical inference of highly parameterized complex nonlinear models. *Optimization Methods and Software*, 27, 233–249.
- Gelman, A., Lee, D. & Guo, J.Q. (2015) Stan: a probabilistic programming language for Bayesian inference and optimization. *Journal of Educational and Behavioral Statistics*, 40, 530–543.
- Gelman, A., Carlin, J.B., Stern, H.S. & Rubin, D.B. (2014) Bayesian Data Analysis. Taylor & Francis, Boca Raton, FL, USA.
- Girolami, M. & Calderhead, B. (2011) Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 73, 123–214.
- Goodman, J. & Weare, J. (2010) Ensemble samplers with affine invariance. Communications in Applied Mathematics and Computational Science, 5, 65–80.
- Grant, R.L., Furr, D.C., Carpenter, B. & Gelman, A. (2016) Fitting Bayesian item response models in Stata and Stan. arXiv preprint arXiv:1601.03443.
- Griewank, A. (1989) On automatic differentiation. Mathematical Programming: Recent Developments and Applications, 6, 83–107.
- Hoffman, M.D. & Gelman, A. (2014) The no-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15, 1593–1623.
- Hooten, M.B. & Hobbs, N.T. (2015) A guide to Bayesian model selection for ecologists. *Ecological Monographs*, 85, 3–28.
- Kéry, M. & Schaub, M. (2012) Bayesian Population Analysis Using WinBUGS: A Hierarchical Perspective. Academic Press, Amsterdam, Netherlands.
- Korner-Nievergelt, F., Roth, T., von Felten, S., Guélat, J., Almasi, B. & Korner-Nievergelt, P. (2015) Bayesian Data Analysis in Ecology Using Linear Models with R, BUGS, and Stan: Including Comparisons to Frequentist Statistics. Academic Press, Amsterdam, Netherlands.
- Kristensen, K., Nielsen, A., Berg, C.W., Skaug, H. & Bell, B. (2015) TMB: automatic differentiation and Laplace approximation. arXiv preprint arXiv:1509.00660.
- Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A. & Blei, D.M. (2016) Automatic differentiation variational inference. arXiv preprint arXiv:1603.00788.
- Livingstone, S., Betancourt, M., Byrne, S. & Girolami, M. (2016) On the geometric ergodicity of Hamiltonian Monte Carlo. arXiv preprint arXiv:1601.08057.
- Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H. & Teller, E. (1953) Equation of state calculations by fast computing machines. *The Journal* of Chemical Physics, 21, 1087–1092.

- Millar, R.B. & Meyer, R. (2000) Non-linear state space modelling of fisheries biomass dynamics by using Metropolis-Hastings within-Gibbs sampling. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 49, 327–342.
- Monnahan, C.C. (2016) Data and code for "Faster estimation of Bayesian models in ecology using Hamiltonian Monte Carlo". doi: 10.5281/ zenodo.159596
- Neal, R.M. (2011) MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, Vol. 2 (eds S. Brooks, A. Gelman, G. L. Jones & X.-L. Meng), pp. 113–162. CRC Press, Boca Raton, FL, USA.
- Papaspiliopoulos, O., Roberts, G.O. & Skold, M. (2007) A general framework for the parametrization of hierarchical models. *Statistical Science*, 22, 59–73.
- Plummer, M. (2003) JAGS: a program for analysis of Bayesian graphical models using Gibbs sampling. Page 125 in Proceedings of the 3rd international workshop on distributed statistical computing. Technische Universit at Wien, Wien, Austria.
- R Core Team (2016) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Roberts, G.O. & Rosenthal, J.S. (2001) Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16, 351–367.
- Royle, J.A. & Dorazio, R.M. (2008) Hierarchical Modeling and Inference in Ecology: The Analysis of Data from Populations, Metapopulations and Communities, 1st edn. Academic Press, London, UK.
- Schnute, J. (1981) A versatile growth model with statistically stable parameters. Canadian Journal of Fisheries and Aquatic Sciences, 38, 1128–1140.
- Stan Development Team (2016) Stan Modeling Language Users Guide and Reference Manual, version 2.12.0. Available at https://github.com/stan-dev/stan/ releases/download/v2.12.0/stan-reference-2.12.0.pdf.
- Stewart, I.J., Hicks, A.C., Taylor, I.G., Thorson, J.T., Wetzel, C. & Kupschus, S. (2013) A comparison of stock assessment uncertainty estimates using maximum likelihood and Bayesian methods implemented with the same model framework. *Fisheries Research*, **142**, 37–46.
- Thorson, J.T. & Minto, C. (2014) Mixed effects: a unifying framework for statistical modelling in fisheries biology. *ICES Journal of Marine Science*, **72**, 1245– 1256.
- de Valpine, P., Turek, D., Paciorek, C.J., Anderson-Bergman, C., Lang, D.T. & Bodik, R. (2016) Programming with models: writing statistical algorithms for general model structures with NIMBLE. *Journal of Computational and Graphi*cal Statistics, 1–28. doi:10.1080/10618600.2016.1172487.

Received 31 August 2016; accepted 6 October 2016 Handling Editor: Robert B. O'Hara

Supporting Information

Additional Supporting Information may be found online in the supporting information tab for this article:

Appendix S1 (Appendices A to E). Further details of the citation analysis, how MCMC efficiency was calculated, and the effect of noncentering.