Comments

A hierarchical model for spatial capture–recapture data: comment

Tiago A. Marques,[1,2,4] Len Thomas,[1,] and
J. Andrew Royle[3]

The development of spatially explicit capture recapture (SECR) methods has received considerable attention in the last few years (e.g., Borchers and Efford 2008, Royle et al. 2009), and it is therefore likely that applied studies using these methods will become widespread.

Royle and Young (2008) present a method that uses SECR data to estimate the abundance and density of animals, within a Bayesian estimation framework. Animals within a known search area D are detected over T sampling occasions, with constant detection probability p given the animal is within the sampling area. Detected animals are assumed to be uniquely identified (e.g., via tagging or natural marks), and the location of each detection is recorded. Animal movement is modeled using a bivariate Gaussian with parameter σ. The model for animal locations assumes animals are randomly distributed over an area S which contains D, such that all animals potentially captured in D have their home range center within S. The information contained on the spatial location of detected animals allows the direct estimation of the area effectively covered by the search area. Therefore one can obtain not only abundance estimates (in an unknown area), as is the case in conventional capture–recapture studies, but actual density estimates. The conclusions include that the method "... provides an alternative to conventional methods (capture–recapture and distance sampling) that allows for temporary emigration and enables sampling on arbitrary plot sizes." and note "... suggested tolerable levels of small-sample bias (5–15%) in the estimator (posterior mean) even under

situations for which typical data sets contain few individuals captured more than one time."

Despite the approach looking promising in general, there are a number of issues that require clarification, and which affect the two conclusions reproduced above. Some of the simulation results originally presented are incorrect, and these are correctly reproduced here. All the notation used is taken directly from Royle and Young (2008).

The information on Table 1 from Royle and Young (2008) presented incorrect values for the scenarios with movement parameter $\sigma = 2$. Instead of reproducing here the correct information (but note the corrected table is given in the Appendix), we present closely related and more informative results in Table 1, namely the number of expected animals with a given number of captures for each of the scenarios used in the original simulation exercise in Royle and Young (2008). From this table we can see that the scenarios considered were extreme, with overall a very small number of recaptures, and hence little information available on the movement parameters and abundance. Therefore, it would be unrealistic to expect good performance from the methods in most simulated scenarios.

This led us to implement again the original simulations reported in Royle and Young (2008), as the results seem to some extent too good to be true. We used a discrete uniform prior for N with upper limit equal to three times the true number of animals present in the area S. Further, and following some investigation of mixing and convergence, we increased both the burn in and the number of iterations used for inference (6000 burn in, inferences based on the following 10 000 iterations). The results are presented in Table 2. These should correspond to those in Royle and Young (2008: Table 2), but are in fact quite different, with far more severe bias than that originally reported, especially for medium and large movement parameters ($\sigma = 2$ and 4). We note in passing that the low coverage observed for the 95% CI reported in Royle and Young (2008) was not observed for the $\sigma = 1$ scenarios. This suggests that, for the scenarios were the methods work well (low bias), they also might provide adequate estimates of variance.

It is possible that one of the reasons for the observed bias is that the posterior distributions are skewed, because with such small sample sizes the likelihoods are also skewed. In this case the mode might be a better summary of the posterior distribution than the mean. The choice of posterior point estimator (say mean, mode or median) is somewhat arbitrary, and the choice of the mean to summarize the results in Royle and Young (2008) might have been the worst option due to the skewness of the posterior distributions. The comparison

TABLE 1. The mean number of animals with a given number of captures for each simulation scenario (i.e., as a function of the movement parameter σ and density).

σ	Density	Number of captures				
		1	2	3	4	5
4	0.234	22	3	0	0	0
4	0.352	32	5	1	0	0
4	0.469	43	7	1	0	0
4	0.586	53	8	1	0	0
4	0.781	71	11	1	0	0
2	0.234	16	5	1	0	0
2	0.352	24	7	2	0	0
2	0.469	32	10	2	0	0
2	0.586	40	12	3	0	0
2	0.781	54	16	3	0	0
1	0.234	13	5	1	0	0
1	0.352	19	8	2	0	0
1	0.469	25	11	3	0	0
1	0.586	32	14	4	1	0
1	0.781	42	19	5	1	0

across these point estimates for the simulated scenarios is shown in Table 3. The relation between the three is for all cases what would be expected given skewed posteriors (mean > median > mode), with the mode being closest to truth.

The reason for the lower bias presented in the original paper is that the prior for N (amount of data augmentation used) was constraining the posterior for N. In fact, for many of the results reported here, even three times N was not enough, in the sense that increasing data augmentation further would lead to further apparent bias. This suggests that the observed bias is a function of the amount of data augmentation used. Naturally one cannot report this as the bias of the

method, as effectively the bias is constrained by the prior used for N, which bounds the upward bias. In that sense, one could expect bias of corresponding maximum likelihood estimators as well.

Based on preliminary simulations, increasing (1) the number of sampling occasions, (2) the density, (3) the sampling area and (4) the detection probability (results not shown), it seems that a minimum number of recaptures is required to reach a setting in which the methods work well, i.e., when increasing data augmentation does not really increase bias. However, it also seems that a reasonably large number of recaptures is not a sufficient condition for the methods to work. If the scale of the problem is such that the sampling area used is very small compared to the movement of the animals, considerable bias is still present in settings with far more recaptures than those observed for the σ = 1 scenarios reported here (cf. Table 1). In such cases, the methods struggle even with a fair number of recaptures. An example was obtained considering σ = 4 and the highest density, and increasing the number of sampling occasions from the original five in Royle and Young (2008) to 10. This led to the expected number of capture histories of 1, 2, 3, 4, and 5+ recaptures being, respectively, 100, 30, 9, 2, and 0 (rounded to the nearest integer). Yet the mean bias observed was still around 34.5% (29.5% if the mode is used as the posterior point estimator). Therefore, while Royle and Young (2008) suggested that arbitrary plot sizes might be used, it seems that small sample areas compared with the scale of animal movement may lead to problems. The observed small sample bias per se is not unexpected, as that is often the case in statistical applications. The methods should work well given large enough sample

TABLE 2. Replica of Table 2 in Royle and Young (2008).

True density	σ = 1				σ = 2				σ = 4			
	Mean	Med.	SD	CI	Mean	Med.	SD	CI	Mean	Med.	SD	CI
0.234	0.283	0.261	0.086	0.95	0.305	0.302	0.098	0.89	0.353	0.364	0.104	0.90
0.352	0.393	0.389	0.087	0.98	0.437	0.423	0.122	0.92	0.562	0.580	0.154	0.83
0.469	0.517	0.502	0.099	0.97	0.562	0.548	0.137	0.89	0.742	0.740	0.203	0.76
0.586	0.623	0.602	0.118	0.93	0.687	0.674	0.178	0.92	0.947	0.920	0.246	0.72
0.781	0.824	0.825	0.121	0.97	0.915	0.890	0.171	0.90	1.142	1.100	0.300	0.74

Note: For each scenario we show the mean, median (Med.), and standard deviation (SD) across the 100 replicates of the Bayesian point estimate (posterior mean), and CI is the percentage of 95% posterior intervals that contained the true density.

TABLE 3. Scenarios as in Table 2, now with means of three different posterior point estimators: mean, median, and mode.

True density	σ = 1			σ = 2			σ = 4		
	Mean	Median	Mode	Mean	Median	Mode	Mean	Median	Mode
0.234	0.283	0.270	0.246	0.305	0.288	0.258	0.353	0.339	0.307
0.352	0.393	0.379	0.355	0.437	0.419	0.385	0.562	0.547	0.516
0.469	0.517	0.502	0.476	0.562	0.541	0.502	0.742	0.722	0.684
0.586	0.623	0.609	0.585	0.687	0.667	0.631	0.947	0.924	0.904
0.781	0.824	0.811	0.790	0.915	0.893	0.858	1.142	1.113	1.059

TABLE 4. The mean of three posterior Bayesian estimates (mean, median, and mode) for the estimated abundance inside area D ($N = 60$) and the corresponding bias for seven different scenarios with increasing information, compared with the original scenario where $\sigma = 1$, detection probability $p = 0.25$, and number of sampling occasions $T = 5$.

Scenario (changes)	σ	p	T	Mean		Median		Mode	
				Parameter	Bias (%)	Parameter	Bias (%)	Parameter	Bias (%)
Original	1.0	0.25	5	68.97	14.95	65.43	9.05	59.52	−0.80
σ	0.5	0.25	5	66.45	10.74	64.42	7.37	61.14	1.90
p	1.0	0.50	5	62.78	4.63	62.01	3.35	60.53	0.88
T	1.0	0.25	10	63.69	6.16	62.74	4.57	61.01	1.69
σp	0.5	0.50	5	60.91	1.51	60.51	0.85	59.78	−0.37
σT	0.5	0.25	10	61.11	1.86	60.70	1.17	59.96	−0.07
pT	1.0	0.50	10	60.71	1.18	60.19	0.32	59.22	−1.31
σpT	0.5	0.50	10	60.77	1.28	60.55	0.92	60.07	0.11

sizes, and provided the area sampled is not too small compared with the animal movement range. Using the original lizard application in Royle and Young (2008) as an example, results seem approximately unbiased (mean observed bias 1.55%; $T = 14$, $\sigma = 0.15$, $p = 0.122$, $D = 8.784$). The likely key factors making it work under this setting are the large number of sampling occasions coupled with a small movement parameter compared to the sampling area (0.15 vs. 3 units side square), which leads to a considerable number of recaptures. If one departs from the small movement, low density scenario and doubles p and/or T, and/or halves σ, we also obtain scenarios under which bias is small (Table 4; note again how the mode seems a better point estimator). Additional simulation studies are required to know under which exact conditions the methods work well.

We note in passing that the mixing for some of the parameters, namely the movement parameters, was extremely slow, and that autocorrelation in Markov chain Monte Carlo (MCMC) samples was high. We were focusing our inferences on abundance, and it seems like the number of iterations used was enough for that purpose. If the movement parameters were the focus of the investigation, longer runs should have been used. Therefore we reinforce the usual precautionary advice for practitioners to assess convergence under the data sets being analyzed. In a simulation study that assessment is slightly harder to do, because of the large number of runs involved, but nonetheless a sufficiently large number of iterations should be used.

The key point in this comment is that despite representing an interesting approach, and unlike what was presented in Royle and Young (2008), the methods

might not work well with such low sample sizes or with arbitrary sampling units, and therefore carefully thinking about sample sizes and the relationship between plot size and animal home range size is necessary when using the methods in practice to estimate animal density or abundance. At the very least, one should always pay close attention to whether the amount of data augmentation used was enough. This can be done by plotting the distributions of estimated N or ψ, and making sure these are not bounded above. It is a good hint that problems might be present if the results keep changing indefinitely as data augmentation is increased. All the R code (R Development Core Team 2009) and the WinBUGS model file used in implementing the simulations and building the tables presented are available in the Supplement. We recommend that, at the very least, potential practitioners looking to use similar methods build on this material to simulate scenarios that mimic their settings. This way, potential problems might be anticipated and, if possible, survey design might be modified to avoid them.

Literature cited

Borchers, D., and M. Efford. 2008. Spatially explicit maximum likelihood methods for capture–recapture studies. Biometrics 64:377–385.

R Development Core Team. 2009. R: a language and environment for statistical computing. R Project for Statistical Computing, Vienna, Austria. ⟨www.r-project.org⟩

Royle, J. A., K. U. Karanth, A. M. Gopalaswamy, and N. S. Kumar. 2009. Bayesian inference in camera trapping studies for a class of spatial capture–recapture models. Ecology 90: 3233–3244.

Royle, J. A., and K. V. Young. 2008. A hierarchical model for spatial capture–recapture data. Ecology 89:2281–2289.

APPENDIX

A table that corrects Table 1 in Royle and Young (2008) (*Ecological Archives* E092-045-A1).

SUPPLEMENT

R code for reproducing all the results presented in the paper (*Ecological Archives* E092-045-S1).