Foraging as sampling without replacement: a Bayesian statistical model for estimating biases in target selection

Alasdair D. F. Clarke¹, Amelia R. Hunt², and Anna E. Hughes¹

¹*University of Essex, Department of Psychology, Colchester, UK, CO4 3SQ*

²*University of Aberdeen, School of Psychology, Aberdeen, UK, AB24 3FX*

Abstract

Foraging entails finding multiple targets sequentially. In humans and other animals, a key observation has been a tendency to forage in ‘runs’ of the same target type. This tendency is context-sensitive, and in humans, it is strongest when the targets are difficult to distinguish from the distractors. Many important questions have yet to be addressed about this and other tendencies in human foraging, and a key limitation is a lack of precise measures of foraging behaviour. The standard measures tend to be run statistics, such as the maximum run length and the number of runs. But these measures are not only interdependent, they are also constrained by the number and distribution of targets, making it difficult to make inferences about the effects of these aspects of the environment on foraging. Moreover, run statistics are underspecified about the underlying cognitive processes determining foraging behaviour. We present an alternative approach: modelling foraging as a procedure of generative sampling without replacement, implemented in a Bayesian multilevel model. This allows us to break behaviour down into a number of biases that influence target selection, such as the proximity of targets and a bias for selecting targets in runs, in a way that is not dependent on the number of targets present. Our method thereby facilitates direct comparison of specific foraging tendencies between search environments that differ in theoretically important dimensions. We demonstrate the use of our model with simulation examples and re-analysis of existing data. We believe our model will provide deeper insights into visual foraging and provide a foundation for further modelling work in this area.

1 Author Summary

Foraging has been well-studied in many species that rely on widely distributed food sources, such as bees and birds. Less well understood is how humans approach foraging tasks, and whether there are general policies we can identify that describe how we search for different categories of objects that can vary in quantity and distribution. We present a way to model foraging behaviour as a generative sampling without replacement procedure, implemented in a Bayesian multilevel model. This allows us to break down behaviour into a number of independent biases that influence target selection, including the proximity of targets, a bias for selecting targets in runs and a bias for a particular target type, in a way that is not dependent on the number of targets present. We believe this tool can open the door for foraging to become a standard task for refining our understanding of attention, working memory, prospective memory, learning, planning and decision-making.

12 2 Introduction

13 Foraging typically refers to the process of searching for and gathering dispersed food. As would be expected
14 given the survival value of this skill, a well-studied question in animal behaviour is the extent to which
15 foraging can be thought of as efficient (e.g. (Pyke, 1984; Brown, 1988)). That is, does the forager's path and
16 food source selection maximize energy intake and minimize energy expenditure and risk? For contemporary
17 humans, most foraging for food takes place in supermarkets, but foraging can be more broadly defined
18 as search for multiple instances of different categories of targets. Many tasks involve this more general
19 definition of foraging (e.g. a security officer monitoring a crowded event; a radiologist searching for all signs
20 of cancer in an X-ray (Wolfe, 2013)).

21 The human foraging literature has been directly inspired by animal behaviour research, including research
22 on how humans find patches (Hills et al., 2013) and how they terminate their search and move onto another
23 patch (Cain et al., 2012; Fougne et al., 2015; Hills et al., 2012; Wolfe, 2013; Wiegand and Wolfe, 2021). In the
24 latter context, an influential modelling framework has been marginal value theorem (MVT) (Charnov, 1976),
25 which predicts that an optimal forager should leave a patch when the "instantaneous rate of return" from a
26 given patch drops below the average rate of return across all the patches (including travel time, during which
27 no elements can be collected). Recent work has shown that human observers can forage optimally (Wolfe,
28 2013), although interestingly this may not hold in all cases e.g. older adults appear to show non-optimal
29 behaviour, staying too long in each patch to adhere to MVT (Wiegand and Wolfe, 2021), and other studies
30 have also found deviations from optimality (Kristjánsson et al., 2020a). Thus, when considering search
31 termination, there are good theoretical models of behaviour that can be used to understand the cognitive
32 processes underlying performance.

33 Another well-established feature of foraging is that animals tend to search for food in 'runs' of one
34 particular food type, particularly when prey are cryptic (Dawkins, 1971; Heinrich, 1976). (See Figure 1 for an
35 example.) Pollinators exhibit a tendency known as 'flower constancy', where they selectively visit flowers of
36 a particular category and ignore equivalently rewarding flowers from other categories. This tendency seems
37 to be context-sensitive in many species, suggesting it is a behavioural adaptation (Grüter and Ratnieks,
38 2011, e.g.). These observations led to research into the concept of the 'search image', where animals
39 direct selective attention to particular features of prey in order to facilitate fast detection (Pietrewicz and
40 Kamil, 1979; Dukas and Ellner, 1993). A similar concept of 'attentional templates' (the representation of the
41 search goal in working memory) has also been proposed for human visual search (Duncan and Humphreys,
42 1989). Target switching is important because it presents a promising model task for understanding the
43 coordination of cognitive and physical effort. Selecting targets based on proximity, that is, selecting all
44 the nearby targets before moving to a new area, minimizes the distance of travel, but the cost is the high
45 cognitive load associated with switching between multiple target templates in working memory. Selecting all
46 the targets of one category before moving on to the next category will minimize the cognitive load, at the
47 cost of increased distance of travel between targets. Which of these two strategies is more efficient will
48 depend on features of both the environment (the distance between targets, the number and complexity of
49 the different types of targets) and the individual (their physical as well as cognitive capabilities). Whether
50 and how humans and other species adjust their strategies in a way that is sensitive to these constraints is
51 a question with broad implications for understanding the efficiency of behaviour more generally.

52 When focusing on target switching, behaviour similar to that seen in nonhuman animal studies has now

Figure 1: *left*: Re-creating Figure 2 of Dawkins (1971): the first 100 grains taken by one chick in experiment 1. *right*: An example of how foraging tasks can be thought of as spatial, using example data from Kristjánsson et al. (2014): the 'path' taken by the forager is indicated by the white line, with the numbers indicating the order of targets taken (first three selections only).

been replicated multiple times in human adult participants (Kristjánsson et al., 2014; Johannesson et al., 2016; Wolfe et al., 2016; Kristjánsson and Kristjánsson, 2018; Clarke et al., 2018; Wolfe et al., 2019; Kristjánsson et al., 2020b; Ólafsdóttir et al., 2020): in a task where participants need to collect as many targets as possible (while ignoring distractors), they tend to switch target type regularly when the task is relatively easy (e.g. when the targets can be distinguished from the distractors by a single feature dimension, such as colour, in an analogous manner to feature search in classical visual search paradigms). However, when participants are completing a more difficult search task (e.g. a conjunction search, where the target can only be distinguished from the distractor using two or more feature dimensions) they tend to search in 'runs', exhaustively searching for one target type before switching to another.

Human foraging studies have also highlighted that there can be surprising heterogeneity between participants in their search strategies. Some participants have been termed 'superforagers', showing no difference in foraging strategy for feature and conjunction type searches (Kristjánsson et al., 2014; Johannesson et al., 2016; Thornton et al., 2019), potentially implying that they are able to hold multiple target templates simultaneously in working memory. However, recent work has also suggested that this group of people may be better described as employing a different strategy; working memory performance is not a good predictor of these individual differences (Jóhannesson et al., 2017), and in fact, it may not be a stable trait. Instead, people may be able to flexibly adjust their strategy in accordance with task demands (Thornton et al., 2019; Prpic et al., 2019; Thornton et al., 2020). It has even been suggested that 'superforagers' is a misnomer, as these participants may show more errors than 'normal' foragers, and instead these strategies may be sub-optimal responses to the task demands (Tagu and Kristjánsson, 2020). Interestingly, children under the age of 12 find conjunction foraging very difficult, indicating the important role of executive functions in this task, and highlighting that our understanding of the cognitive processes underlying run behaviour remains incomplete (Ólafsdóttir et al., 2020).

2.1 Problems with run statistics

An episode of foraging entails a sequence of target-detection responses which have both a spatial and a target identity dimension. With each target selection response in the sequence, the distribution of remaining

79 targets changes on both of these dimensions. In this respect, foraging data presents a unique analytical
80 challenge relative to single-target search.

81 To study target switching, many previous studies have used a set of similar, highly correlated dependent
82 measures: the target switch rate, the average number of runs per trial, or the average run length per trial. An
83 issue with all of these measures is that the ground truth of the search area will influence target switching;
84 the relative proximity and number of targets of the other category makes switches more likely. This is
85 problematic for comparing run statistics across different conditions or studies that vary the distributions
86 of target locations and categories. More fundamentally, run statistics do not adequately represent the
87 underlying cognitive processes that might be determining foraging behaviour. Ideally, we would to model
88 behaviour with respect to the relative contribution of parameters such as the tendency towards target
89 constancy or switching target type, in conjunction with possible biases for one target type over another.

90 Run statistics also leave aside the intrinsic spatial aspects of foraging tasks. The effect of the distribution
91 of foraging targets relative to the forager is central to understanding foraging efficiency: for example, patch
92 leaving is affected by distance between patches (Wolfe, 2013), and heavy bees have a stronger proximity bias
93 than lighter ones (Greenleaf et al., 2007). Despite this, relatively little work has looked at target switching in
94 direct relation to the effect of target proximity (although see Le et al. (2021) for one recent approach). This
95 is especially important given that the spatial distribution of targets varies not only between conditions and
96 studies, but also within a single episode of foraging, as each consecutive target is found and removed and
97 the distribution becomes sparser.

98 2.2 Our proposed approach

99 We propose an alternative strategy for analysing visual foraging data, which involves modelling the process
100 as a generative sampling without replacement procedure, implemented in a Bayesian multilevel model. A
101 strength of our approach is that it leads to measurements (model parameters) that are independent of the
102 number of targets presented on screen and the number of targets that the participants are tasked with
103 finding. In addition, we can estimate the latent biases that led to the run statistics.

104 We begin by demonstrating the modelling approach for a single trial of visual foraging data from Dawkins
105 (1971), a classic study from the animal foraging literature. From this single trial, we can estimate a parameter
106 that represents the probability that a forager will stick to the same target type or switch to a different type.
107 We will then extend the model to multiple trials and conditions, adding in a bias for target type, and show
108 that our model is able to detect small differences in these parameters that would not be picked up using
109 the dependent variables of maximum run length or number of runs. We then extend to multi-level models
110 to incorporate individual differences, and demonstrate the use of these models using previous data from
111 Tagu and Kristjánsson (2021).

112 Finally, we present a version of the model that incorporates a proximity bias, allowing us to estimate how
113 likely participants are to select a target close to the previously selected target. We use this to demonstrate
114 how the proximity bias affects the calculation of the other biases using example data from Clarke et al.
115 (2018).

3 Pulling balls out of a bag

We are modelling visual foraging as a process of weighted sampling without replacement, in what we call a 'bag foraging model'. Targets are selected one at a time and we define $t(i) = t_i$ as the i th target selected in a trial, $i = 1, 2, \dots$. Our model is binary: targets can be one of two classes, $t_i \in \{a, b\}$, with n_a and n_b counting the number of each class of targets that have not yet been selected. $n_T = n_a + n_b$ is the total number of eligible targets remaining in trial. We assume that any biases in target selection strategy are constant within, and over, trials.

3.1 Single Trial

We start with only one trial of data and assume targets of class a and b are equally likely to be selected. We model the preference, $p_s \in [0, 1]$, for selecting a target of the same class as the previously selected target. Higher values of p_s will lead to longer runs with the same class of target being selected repeatedly, while values of $p_s < 0.5$ indicate a preference for alternating between target classes.

For the first target selection, t_1 , the probability of selecting a target of class a , is simply the proportion of a 's present in the stimulus:

$$Pr(t_1 = a) = \frac{n_a}{n_a + n_b} \quad (1)$$

For all subsequent target selections, $t_2, t_3, \dots$, the probabilities $Pr(t_i = a)$ and $Pr(t_i = b) = 1 - Pr(t_i = a)$ depend on the identity of the previously selected target, t_{i-1} :

$$Pr(t_i = a) = \begin{cases} \frac{p_s n_a}{p_s n_a + (1-p_s) n_b} & \text{if } t_{i-1} = a \\ \frac{(1-p_s) n_a}{(1-p_s) n_a + p_s n_b} & \text{if } t_{i-1} = b \end{cases} \quad (2)$$

An easy way to think about p_s is to consider the special case when there are equal numbers of both classes of targets left. In this case, p_s can be interpreted as:

$$p_s = Pr(t_i = t_{i-1} | n_a = n_b) \quad (3)$$

The model was coded in Stan and we use logit link functions to ensure that our probabilities will lie between 0 and 1. I.e., we fit $b_s = \text{logit}(p_s)$. We define our prior as $b_s \sim N(0, 1)$, which translates to a prior prediction that p_s will fall within 0.1 and 0.9 approximately 95% of the time. The code is presented in Appendix A and available on the Open Science Framework¹. We ran the model using R (v4.0.3) and rStan, v.2.21.2, (Stan Development Team, 2020) with four chains and 1000 iterations. Further details of model checking procedures can be found in the Supplementary Material (part 1, section 1).

3.1.1 Chicken foraging

We imported chicken foraging data from Dawkins (1971), which is often referenced as one of the classic papers in the foraging literature. The data is from one chick, choosing between two types of food (grains of rice), and is shown in Figure 1.

¹<https://osf.io/7yuaz/>

Figure 2: Prior and posterior probability distributions for p_s for the chicken foraging data from Dawkins (1971).

After running the model on the data, we inspected traceplots and model summary statistics ($\hat{R}$, n_{eff}) to check that for successful convergence. These can be found in the Supplementary material (part 1, section 1). The posterior is shown in Figure 2. By calculating the 53% and 97% HPDIs (highest posterior density intervals)² around the preference to select the same target class, we obtain intervals of $[0.83, 0.88]$ (more likely than not) and $[0.78, 0.92]$ (very likely).

3.1.2 Stable estimates over trial length and number of targets

How much data is required to get a stable estimate of these biases? Dawkins (1971) only looked at the first 100 of 2000 targets found (5%). We ran a simulation where the chickens found greater proportions of the grains (10% or 100%) and found that our model does improve with more data, but only marginally: there is no great benefit to a 20-fold increase in grains in this simple simulation. We also ran a simulation where the chickens always found 100 grains, but the total number of grains varied, and again found this had only small effects on the estimates of p_s . Our method appears to offer robust model fits even with relatively small numbers of targets (see part 1, section 1.4 of the Supplementary Materials for further details).

3.2 Multiple Trials & Conditions

One of the issues with previous measures used in the foraging literature is that they are not able to address why run lengths may differ. That is, what are the underlying cognitive processes that drive a switch to a different target? Our method allows us to consider multiple biases: for example, in addition to a forager's tendency to stick with one target type or to switch (p_s) we can add in a second parameter, p_a , which reflects the bias towards one target type or another (e.g. if one target has a higher value). Increasing either or both parameters will lead to a smaller number of longer runs, but for different reasons.

We therefore expanded our model to work over multiple trials and conditions, and allowed participants to have a preference for one target class over the other:

$$p_a = \text{Pr}(t_i = a | n_a = n_b) \quad (4)$$

²We use these intervals over the more common, but equally arbitrary, 95% for the following reasons: i) 53% represents a little over half of the distributions while 97% represents most of it. ii) there are many cognitive biases associated with the interpretation of probabilities and it is possible that 50% intervals could be viewed as being more likely to contain the “true” value than not. iii) 97% offers a more conservative choice than 95% while avoiding numerical difficulties in obtaining stable estimates out in the tails of the distribution.

We estimate p_s and p_a independently for each condition, under the assumption that they are constant over trials from the same condition. Our model then becomes:

$$Pr(t_1 = a) = \frac{p_a n_a}{p_a n_a + (1 - p_a) n_b} \quad (5)$$

$$Pr(t_i = a) = \begin{cases} \frac{p_a p_s n_a}{p_a p_s n_a + (1 - p_a)(1 - p_s) n_b} & \text{if } t_{i-1} = a \\ \frac{p_a (1 - p_s) n_a}{p_a (1 - p_s) n_a + (1 - p_a) p_s n_b} & \text{if } t_{i-1} = b \end{cases} \quad (6)$$

See Appendix A.2 for the Stan code; the model was run in a similar manner to the one trial model.

3.2.1 Misattribution example: a p_a bias with no p_s bias

A strength of our model is it allows us to more accurately determine the biases that underlie foraging patterns. We can compare a condition with no biases ($p_a = 0.5$ and $p_s = 0.5$, a 'neutral' condition) with one where the observer prefers one target over the other ($p_a = 0.8$ and $p_s = 0.5$, which we call a 'target bias' condition, as in this case one target is preferred over the other e.g. because it has been given a higher value in the experiment, or because of its visual properties). An ANOVA (the standard statistical method used in foraging papers) can detect a statistically significant difference between these two conditions in terms of both the number of runs ($F = 91.1, p < 0.001$) and the maximum run length ($F = 99.1, p < 0.001$): see Figure 3, panel A). However, it is not clear from this analysis what is driving these effects. The posterior probabilities from our model are presented in Figure 3 (panel C) and clearly distinguish the underlying processes i.e. that there is no difference between p_s in the two conditions but a clear difference between p_a : if we repeat the procedure 100 times, we find a difference in p_a 100% of the time, and a difference in p_s only 2% of the time.

3.2.2 Differentiation example: two conditions which give the same run statistics

To demonstrate the power of our approach, we can look at a simulation example, comparing a condition with a small stick bias ($p_a = 0.5, p_s = 0.55$) to a condition with a small target type preference ($p_a = 0.65, p_s = 0.5$). We simulated 50 trials from an experiment with these two conditions, with each trial consisting of 40 targets. As can be seen in Figure 3 (panel B), these two conditions give rise to very similar distributions of maximum run lengths and number of runs. Indeed, when applying the standard ANOVA to these data, we fail to detect a statistically significant difference: for example, in one simulated dataset, $F = 0.02, p = 0.9$ for the maximum run lengths and $F = 0.41, p = 0.52$ for the number of runs.

In contrast, running our model on these data results in the posterior probability distribution seen in Figure 3 (panel D). From these, we can calculate the difference between conditions, which can be summarised with 97% HPDIs for p_s , $[-0.09, -0.02]$, and p_a , $[0.10, 0.17]$. These intervals do not cross 0, consistent with a difference between conditions for each measure separately. An alternative way to summarise these posterior distributions is to calculate $Pr(\text{difference} > 0 | \text{data})$. We can then carry out a power analysis by repeating this procedure 100 times. When we do this, we detect a significant difference in p_a and p_s in 100% and 83% of simulations, respectively. In contrast, applying the same power analysis using the standard run statistics, we find a statistically significant difference in the number of runs in just 8% of the simulations. Similarly, we find a difference 21% of the time for the maximum run length statistic. We detect a difference

Figure 3: A: boxplots showing the maximum run length and number of runs in each of our two simulated misattribution conditions. B: boxplots showing the maximum run length and number of runs in each of our two simulated differentiation conditions. C: density plots showing the p_a and p_s values calculated by our model for each of the two simulated misattribution conditions. D: density plots showing the p_a and p_s values calculated by our model for each of the two simulated differentiation conditions.

in either statistic, and hence detect a difference between conditions, only 29% of the time. Our approach is therefore not only more sensitive than run statistics to the difference between conditions, but is able to correctly attribute the difference to a combination of a bias to switch and a preference for one target over the other.

3.3 Multi-level: Tagu and Kristjánsson (2021)

The final version of the bag foraging model allows multiple participants in a multi-level framework. Both p_s and p_a are allowed to vary from one participant to the next.

Using simulated data, we can show that our model can recover both overall p_s and p_a biases, as well as the parameters from individual participants (see Supplementary Materials, part 1, section 3 for details).

Our model also works well with more complex examples. In a recent study, Tagu and Kristjánsson (2021) included a target value manipulation, where participants searched for three target types (among three different types of distractors). Participants had to earn a certain number of points in order to complete the trial, and in the 'value' condition, one of the target types was deemed high value, whereas the other two were low value. We would therefore expect this condition to have a higher p_a bias than the 'no value' control, where all targets had the same value. This is indeed what we see in Figure 4. This example also shows the flexibility of the model; in this experiment, the proportion of the different target types (high and

Figure 4: p_a and p_s biases across participants in multi-level re-analysis of Tagu and Kristjánsson (2021), showing both the value and no-value conditions.

low value) was not equal to 50:50, and trials were not necessarily equal lengths. Note, for simplicity, we considered only the blocks where participants used a computer mouse in our analysis, and only the trials where the targets had a random distribution (see Supplementary Materials).

3.4 Discussion

The bag foraging model we have presented above can account for run behaviour in foraging tasks. Alongside data from Dawkins (1971) and Tagu and Kristjánsson (2021), we have also fit the multi-level model to data from (Kristjánsson et al., 2014) (presented in Supplementary Material, part 1, section 3) and can replicate their key findings. For example, on average, participants have a lower p_s in the feature condition than in the conjunction condition, demonstrating that they are more likely to have longer ‘runs’ of one target type in the conjunction condition. We can also clearly identify the so-called ‘superforager’ participants, who show a much smaller gap between conditions. However, our method also allows us to see new patterns in the data. For example, the p_a values in the conjunction condition were often below 0.5. This indicates that participants had a preference for the targets labelled 0, suggesting that by chance, in this dataset, these targets were more salient: the Supplementary Material shows how the individual p_{Az} correlated well with the participants’ tendency to select a ‘0’ for the first target on each trial. Our method therefore allows deeper analysis of the underlying biases than has been possible with the metrics used in previous studies, while also correlating well with those measures (see Supplementary Material).

One limitation of our model is that it can give less accurate parameter estimates if there are very extreme biases e.g. if the two target types are found almost exclusively in disjoint, non-overlapping runs. If $p_s = 1$ (and $p_a = 0.5$), our simulated results suggest that the recovery of the p_s parameter will be accurate, but the posterior for p_a is shifted away from 0.5. If we take the converse example where $p_a = 1$ (and $p_s = 0.5$), we obtain a model fit that generates posterior predictions that are identical to the types of runs that would be generated by these parameters i.e. the model correctly predicts that all the targets of one type will be selected first. However, we see that both p_a and p_s estimates are actually close to one (see Supplementary Material, part 1, section 2). This is because an extremely strong target preference means that the participant will find all these preferred targets before the non-preferred, generating a trial where they appear never to switch. Thus, the estimated parameters are entirely consistent with the observed data.

242 However, assuming a slightly less extreme value of $p_a = 0.97$, or adjusting the priors, allows better recovery
 243 of the correct parameters (see Supplementary Material part 1, section 2.3 for full model information). In the
 244 case of the Kristjánsson et al. (2014) data, the p_a bias appears to be real, as manual inspection of the first
 245 five target selections per trial in the conjunction condition shows a similar bias towards the target labelled
 246 '0'. Thus, some caution should therefore be exercised in interpreting parameters where there are strong
 247 biases, but it is simple to check if there are model fitting problems.

248 4 Spatial Biases

249 Our initial model incorporates several biases which appear to help explain the findings in previously published
 250 research (Tagu and Kristjánsson, 2021). However, this model has no representation of the spatial location
 251 of targets. It seems plausible that people are more likely to select targets nearby to the one they have just
 252 selected. Proximity could be particularly important in feature search, where the targets are relatively easily
 253 distinguished so there are fewer additional constraints on selection; if so, perhaps proximity is relatively
 254 less important for conjunction search. We can also investigate the effect of direction/momentum: do
 255 observers incorporate some form of momentum to their target selections, or do they prefer to double back
 256 on themselves? Extending our bag foraging model to incorporate spatial biases allows us to see to what
 257 extent these hypotheses are true, and how the inclusion of a spatial bias will affect our estimation of the
 258 other biases.

259 For our second model, we take our stimuli to be a collection of targets, $t(i) = (t_i, x_i, y_i)$, where (x_i, y_i)
 260 gives the target's location and $t_i \in \{a, b\}$ gives the target's class. For mathematical convenience, we will
 261 take a and b to factor levels with $a = 1$ and $b = -1$. If we take $d(i, j)$ to be the Euclidean distance between
 262 targets (x_i, y_i) and (x_j, y_j) , we can define a proximity measure between the two items as:

$$\rho_d(i, j) = e^{-\sigma_d d(i, j)}$$

263 where σ_d is a proximity tuning parameter that controls how quickly ρ decreases from 1 to 0.

264 Similarly, we can define the angular difference $\theta(i, j)$ to be:

$$\theta(i, j) = \frac{f(\text{atan2}(j, j) - \text{atan2}(i - 1, i))}{\pi}$$

265 where $\text{atan2}(j, i)$ is the direction of travel from i to j and $f(\phi_1, \phi_2)$ calculates the angular difference:

$$f(\phi_1, \phi_2) = \min((\phi_1 - \phi_2) \% 2\pi, (\phi_2 - \phi_1) \% 2\pi)$$

266 This should give a value between 0 and 1, with 0 indicating a target that "in front" of the current target,
 267 0.5 represents a change of direction of 90° and 1 represented doubling back in the direction we just came
 268 from. Having calculated θ , we can now define:

$$\rho_\theta(i, j) = e^{-\sigma_\theta \theta(i, j)}$$

269 In the bag foraging model, all targets of the same class were interchangeable and so a simple Bernoulli
 270 process could be used to model each target selection. However, this no longer holds, as each target has a

271 unique spatial location that needs to be taken into account. Therefore we will use a categorical model with
272 the weights for each remaining target being defined as:

$$w(i) = g(b_a t_i + b_s m(t_i, t_{i-1})) \times \rho_d(i, i-1) \rho_\theta(i, i-1) \quad (7)$$

273 where g is the inverse logit function and m is an indicator function that equals 1 if $t_i = t_{i,1}$ and -1
274 otherwise. Targets that have already been selected earlier in a trial have their weights set to 0.

275 4.1 Simulation Examples

276 We present a simple example using simulated data in the Supplementary Material (part 2), and demonstrate
277 that we can recover the parameters used to generate the data. A more interesting example can be seen
278 in the case where the targets have spatial structure in the environment, being found in patches³. In our
279 simulation, we set b_a and b_s to be zero, and used a proximity tuning value of 15. If analysed without including
280 a proximity bias parameter, the model mistakenly generates a b_s bias: the clumping of targets does indeed
281 lead to them being collected in runs, simply because our simulated participant has a preference for nearby
282 targets (see Figure 5). However, the model including the proximity bias parameter is able to recover the
283 correct parameters. Our model is therefore able to distinguish between a preference for a *similar* target and
284 a preference for a *nearby* target, a distinction that is likely to be particularly important in more naturalistic
285 set ups where target clumping might be expected to be prevalent.

286 4.2 Multi-level: Clarke et al. (2018)

287 The data from Clarke et al. (2018) closely replicates Kristjánsson et al. (2014), yet has around four times as
288 many participants, giving us greater power to investigate individual differences in feature and conjunction
289 foraging. The fixed-effects of the model are shown in Figure 6. There is a small bias towards repeating target
290 selections during feature search, and this increases to approach ceiling in the conjunction condition. There
291 is also a very strong proximity bias in both conditions, with very little weight given to targets more than a
292 quarter stimulus width from the currently selected target. The parameter for the direction bias comes out
293 as negative, leading the model to favour doubling back on itself to select targets that may have been skipped
294 over (see Supplementary Material part 2, section 1.3 for simulated examples of how different direction biases
295 affect foraging behaviour). Interestingly, both of these biases are slightly weaker for conjunction foraging,
296 suggesting participants are relying less on spatial attention and emphasising feature-based attention.

297 As this model was fit using a multi-level framework, we can also look differences between participants.
298 Figure 7 shows how the parameters vary over participants and conditions. We can see that with the
299 exception of parameters linked to the spatial biases, there are little-to-no correlations between different
300 parameters across conditions, or within parameters across conditions. However, the two spatial parameters
301 appear to be correlated: participants with a stronger preference to select nearby targets also have a more
302 negative direction bias. Participants with a weaker proximity bias are more likely to have a positive directional
303 bias. While there is no clear partition of participants into subgroups, the two opposing ends of this range
304 could be interpreted as a strategy to search within local patches compared to a more global scanning
305 strategy.

³See Supplementary Material for example stimulus.

Figure 5: Posterior distributions for both a bag foraging model and a spatial model trained on patchy stimuli, where target types are clumped. The shaded ribbons for the proximity and direction weighting indicate 53%, 89% and 97% HPDIs. The spatial model gives more weight to nearby targets (i.e. where the distance is close to zero), and has no particular preference for any direction, as indicated by the fact that a value of 1 falls within the HPDI for all directions.

Finally, we can use the fitted model to simulate new data and compare run statistics and inter-target distances with the original empirical data. We do this for $n = 100$ samples from the fitted posterior probability distributions so that we can quantify the uncertainty in the model's predictions. The results (Figure 8) show a good correspondence between the posterior predictions and empirical data. Both the summary run statistics (the two statistics most commonly analysed in the visual foraging literature) are a close-to-perfect match between the model and the empirical data. The inter-target distances are more interesting. The model offers a reasonably good fit, especially in terms of the mean distances and in the later half of a trial. However, we can see two main differences: i) the variance in the human data is larger, suggesting that perhaps the negative exponential proximity function down-weights more distant targets too much; and ii) there is an interesting dynamic at the start of the trial in which the first few targets selected tend to have shorter inter-target distances than the model would expect. This raises the possibility that participants either chose an initial target that is in a relatively dense patch of potential targets, or that they are able to accurately plan out an efficient order for the first five or so target selections that minimises the distance that they would be required to travel. This observation highlights another use of the model, which is to detect deviations in human behaviour from the model's predictions. These deviations can guide the direction of further investigation.

Figure 6: Posterior distributions for our model when trained on data from Clarke et al. (2018). The shaded ribbons for the proximity and direction weighting indicate 53%, 89% and 97% HPDIs. Note that there is a higher weighting for directions closer to π i.e. completely reversing the direction, indicating negative momentum.

4.3 Discussion

The spatial model presented extends the simple bag foraging model to incorporate new parameters that further explain participant behaviour on foraging tasks by accounting for spatial attention. This section demonstrates the flexibility and modularity of the bag foraging model, which can be used as a starting point for exploring many questions about foraging behaviour. As well as presenting the model fit for the Clarke et al. (2018) dataset above, the Supplementary Material (part 2) presents fits for a number of other datasets, including Kristjánsson et al. (2014), Tagu and Kristjánsson (2021) and Thornton et al. (2020). Another key strength of our model demonstrated in the Supplementary Materials is that it is able to generate good parameter estimates with relatively little data (e.g. one trial per participant per condition), in stark contrast to using traditional run statistics, which reduces each trial to one summary number. This means that our model could allow analysis of foraging data from populations where it is difficult to carry out large numbers of trials, such as from clinical populations or when studying animal behaviour.

5 General Discussion

One of the biggest strengths of our approach is its relative simplicity. By taking the task structure of sampling without replacement into account, we need only four parameters to characterise the relatively complex sequences of runs and inter-target distances produced during foraging tasks. While approximating these parameters with summary statistics (e.g. the mean of the run lengths) in many cases will lead to broadly similar conclusions, we have demonstrated that in some instances, run statistics give misleading

Figure 7: Individual differences in Clarke et al. (2018). The first two rows show the correlations in the random effect structure between different parameters. We can see that p_s and p_s appear to be independent from one another, and independent from the spatial biases b_p and p_m . However, b_p and p_m are correlated with one another in both the feature and conjunction conditions. The bottom row shows the correlation between conditions for each parameter in our model.

results that our approach avoids. In addition, summary statistics require much more data per participant and condition to be able to make good estimates. By treating each selection instance as a separate observation, our approach reaches stable estimates of biases far more efficiently.

Our model is also extremely flexible, and could be adapted in many ways. The underlying assumptions of the model are all tuneable and testable: for example, it would be straightforward to compare and contrast different proximity functions, as our choice of using an isotropic negative exponential is relatively arbitrary, and other functions may offer improved performance for predicting behaviour. The momentum function also assumes a monotonic trend from going forwards in a straight line to going backwards, but again, there are other possibilities that could easily be tested. For example, it could also be explored whether functions that allow for a forward-backward preference over perpendicular movement are better. Interesting further questions could explore momentum and proximity biases across different modalities. Most of the data included here was based on finger or mouse-click foraging, but larger movements (such as moving around a large area) or smaller ones (such as eye movements) are likely to be characterised by different proximity and momentum functions.

More broadly, the model is extremely straightforward to extend. As it stands, our model works well with different types of foraging tasks, as we have demonstrated using both exhaustive and non-exhaustive foraging example datasets. However, while outside the scope of the current paper, it would be easy to

Figure 8: Posterior predictions for Clarke et al. (2018). (*top:*) Scatter plots between empirical and predicted summary run statistics. Error bars indicated 89% HPDI for the posterior predictions. (*bottom:*) Comparison of inter-target distances and how they vary over time. Shaded region gives 89% intervals.

also implement stopping rules for models of non-exhaustive foraging, and it would be possible to make quantitative predictions about distributions by assuming different stopping rules (for example, by assuming that after each target selection, there is some given probability to terminate the trial). This would allow modelling of patch-leaving rules, assuming that one trial acts as one patch. Within a trial, it would also be possible to model patches made up of multiple "clumped" targets on a single large display, by modifying the spatial bias component to include a categorical factor to indicate the different clumps of targets (that are therefore assumed to be perceptually 'further away' from each other). The p_a bias could also be broken down into sub-biases based on how preference changes based on stimulus properties such as colour, luminance and value. The model could also be extended to explore how biases change temporally (e.g. over multiple trials), or to include new biases, such as a bias for the initial target selection. Finally, it would also be possible to extend beyond two target categories: for example, for three categories, instead of fitting one value for $Pr(t_i = a)$ and then letting $Pr(t_i = b) = 1 - Pr(t_i = a)$, we would need to add an extra parameter, which could be parameterised $Pr(t_i = c) = 1 - Pr(t_i = a) - Pr(t_i = b)$.

In other aspects of foraging, such as the marginal value theorem for predicting when foragers should leave a patch (Charnov, 1976), formal mathematical modelling has been used to effectively represent the underlying structure of foraging behaviour. Our work provides a similar framework for target selection

during foraging, extending our theoretical models of foraging behaviour and giving a way for researchers to begin to understand the cognitive processes that underlie performance in these tasks. We envision future extensions bringing these and other models together into a complete understanding of all aspects of foraging. While a full account of human foraging is of course likely to be extremely complex, involving many factors (Kristjánsson et al., 2020a), we are optimistic that our approach will provide a framework to begin this challenging task. By making our code freely available, we hope to enable other researchers to update the model as they require, perhaps to take into account other cognitive factors that they think may be important in foraging tasks, or to enable them to use the model to study their own variants of the task.

There are a number of potentially important factors that we do not consider in the current model. For example, it does not currently allow for modelling inter-target selection times, which are thought to be important in determining foraging behaviour (Thornton et al., 2020). Similarly, it does not explicitly model the effect of background clutter or distractors, although implicit effects (such as decreasing the salience of one of the target classes) will be captured. As most foraging studies to date have not focused on distractor effects, information on the distribution of distractors is not routinely included in freely-available datasets, making it more difficult to model this aspect of the task. However, if such data were available, the model could be extended to treat distractors in the same way as targets, but with a sampling-with-replacement and inhibition of return process for when a distractor is selected.

5.1 Conclusion

Our Bayesian statistical model of visual foraging provides a simple and flexible way to estimate cognitively meaningful parameters underlying foraging tasks. We demonstrate that the model is able to reproduce the patterns found in existing data, and also provides new avenues for research questions, such as the role of salience in predicting target preferences and how proximity and direction biases may represent stable individual differences across different task variants. By considering the task structure in developing our model, we can go beyond simple linear models to generate more powerful predictions and deeper insights into foraging behaviour.

6 Author Contributions

ADFC, ARH and AEH jointly wrote the manuscript.

7 Acknowledgements

This research was supported by the Economic and Social Research Council (ESRC, ES/S016120/1). The authors would like to thank all researchers who share their data on OSF. We would also like to thank Jérôme Tagu and two anonymous reviewers for their helpful comments that greatly improved the manuscript.

References

Brown, J. S. (1988). Patch use as an indicator of habitat preference, predation risk, and competition. *Behavioral ecology and sociobiology*, 22(1):37–47.

- 407 Cain, M. S., Vul, E., Clark, K., and Mitroff, S. R. (2012). A bayesian optimal foraging model of human visual search.
408 *Psychological science*, 23(9):1047–1054.
- 409 Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical population biology*, 9(2):129–136.
- 410 Clarke, A. D., Irons, J. L., James, W., Leber, A. B., and Hunt, A. R. (2018). Stable individual differences in strategies
411 within, but not between, visual search tasks. *Quarterly Journal of Experimental Psychology*, page 1747021820929190.
- 412 Dawkins, M. (1971). Shifts of ‘attention’ in chicks during feeding. *Animal Behaviour*, 19(3):575–582.
- 413 Dukas, R. and Ellner, S. (1993). Information processing and prey detection. *Ecology*, 74(5):1337–1346.
- 414 Duncan, J. and Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological review*, 96(3):433.
- 415 Fougny, D., Cormiea, S. M., Zhang, J., Alvarez, G. A., and Wolfe, J. M. (2015). Winter is coming: How humans forage in
416 a temporally structured environment. *Journal of vision*, 15(11):1–1.
- 417 Greenleaf, S. S., Williams, N. M., Winfree, R., and Kremen, C. (2007). Bee foraging ranges and their relationship to body
418 size. *Oecologia*, 153(3):589–596.
- 419 Grüter, C. and Ratnieks, F. L. (2011). Honeybee foragers increase the use of waggle dance information when private
420 information becomes unrewarding. *Animal behaviour*, 81(5):949–954.
- 421 Heinrich, B. (1976). The foraging specializations of individual bumblebees. *Ecological monographs*, 46(2):105–128.
- 422 Hills, T. T., Jones, M. N., and Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological review*, 119(2):431.
- 423 Hills, T. T., Kalff, C., and Wiener, J. M. (2013). Adaptive lévy processes and area-restricted search in human foraging.
424 *PLoS One*, 8(4):e60488.
- 425 Jóhannesson, Ó. I., Kristjánsson, Á., and Thornton, I. M. (2017). Are foraging patterns in humans related to working
426 memory and inhibitory control? *Japanese Psychological Research*, 59(2):152–166.
- 427 Jóhannesson, O. I., Thornton, I. M., Smith, I. J., Chetverikov, A., and Kristjánsson, A. (2016). Visual foraging with fingers
428 and eye gaze. *i-Perception*, 7(2):2041669516637279.
- 429 Kristjánsson, Á., Björnsson, A. S., and Kristjánsson, T. (2020a). Foraging with anne treisman: Features versus conjunc-
430 tions, patch leaving and memory for foraged locations. *Attention, Perception, & Psychophysics*, pages 1–14.
- 431 Kristjánsson, Á., Jóhannesson, Ó. I., and Thornton, I. M. (2014). Common attentional constraints in visual foraging. *PloS*
432 *one*, 9(6):e100752.
- 433 Kristjánsson, T., Draschkow, D., Pálsson, Á., Haraldsson, D., Jónsson, P. Ö., and Kristjánsson, Á. (2020b). Moving foraging
434 into three dimensions: Feature-versus conjunction-based foraging in virtual reality. *Quarterly Journal of Experimental*
435 *Psychology*, page 1747021820937020.
- 436 Kristjánsson, T. and Kristjánsson, Á. (2018). Foraging through multiple target categories reveals the flexibility of visual
437 working memory. *Acta Psychologica*, 183:108–115.
- 438 Le, S. T. T., Kristjansson, A., and MacInnes, W. J. (2021). Bayesian approximations to the theory of visual attention (tva)
439 in a foraging task.
- 440 Ólafsdóttir, I. M., Gestsdóttir, S., and Kristjánsson, Á. (2020). Age differences in foraging and executive functions: A
441 cross-sectional study. *Journal of Experimental Child Psychology*, 198:104910.

Pietrewicz, A. T. and Kamil, A. C. (1979). Search image formation in the blue jay (*cyanocitta cristata*). *Science*, 204(4399):1332–1333.

Prpic, V., Kniestedt, I., Camilleri, E., Maureira, M. G., Kristjánsson, Á., and Thornton, I. M. (2019). A serious game to explore human foraging in a 3d environment. *PLoS one*, 14(7):e0219827.

Pyke, G. H. (1984). Optimal foraging theory: a critical review. *Annual review of ecology and systematics*, 15(1):523–575.

Stan Development Team (2020). RStan: the R interface to Stan. R package version 2.21.2.

Tagu, J. and Kristjánsson, Á. (2020). Dynamics of attentional and oculomotor orienting in visual foraging tasks. *Quarterly Journal of Experimental Psychology*, page 1747021820919351.

Tagu, J. and Kristjánsson, Á. (2021). The selection balance: contrasting value, proximity and priming in a multitarget foraging task.

Thornton, I. M., de’Sperati, C., and Kristjánsson, Á. (2019). The influence of selection modality, display dynamics and error feedback on patterns of human foraging. *Visual Cognition*, 27(5-8):626–648.

Thornton, I. M., Nguyen, T. T., and Kristjánsson, Á. (2020). Foraging tempo: Human run patterns in multiple-target search are constrained by the rate of successive responses. *Quarterly Journal of Experimental Psychology*, page 1747021820961640.

Wiegand, I. and Wolfe, J. M. (2021). Target value and prevalence influence visual foraging in younger and older age. *Vision Research*, 186:87–102.

Wolfe, J. M. (2013). When is it time to move to the next raspberry bush? foraging rules in human visual search. *Journal of vision*, 13(3):10–10.

Wolfe, J. M., Aizenman, A. M., Boettcher, S. E., and Cain, M. S. (2016). Hybrid foraging search: Searching for multiple instances of multiple types of target. *Vision research*, 119:50–59.

Wolfe, J. M., Cain, M. S., and Aizenman, A. M. (2019). Guidance and selection history in hybrid foraging visual search. *Attention, Perception, & Psychophysics*, 81(3):637–653.

8 Legends

Figure 1: Re-creating Figure 2 of Dawkins (1971): the first 100 grains taken by one chick in experiment 1. *bottom:* An example of how foraging tasks can be thought of as spatial, using example data from Kristjánsson et al. (2014): the ‘path’ taken by the forager is indicated by the white line, with the numbers indicating the order of targets taken (first three trials only).

Figure 2: Prior and posterior probability distributions for p_s for the chicken foraging data from Dawkins (1971).

Figure 3: A: boxplots showing the maximum run length and number of runs in each of our two simulated misattribution conditions. B: boxplots showing the maximum run length and number of runs in each of our two simulated differentiation conditions. C: density plots showing the p_a and p_s values calculated by our model for each of the two simulated misattribution conditions. D: density plots showing the p_a and p_s values calculated by our model for each of the two simulated differentiation conditions.

Figure 4: p_a and p_s biases across participants in multi-level re-analysis of Tagu and Kristjánsson (2021), showing both the value and no-value conditions.

Figure 5: Posterior distributions for both a bag foraging model and a spatial model trained on patchy stimuli, where target types are clumped. The shaded ribbons for the proximity and direction weighting indicate 53%, 89% and 97% HPDIs. The spatial model gives more weight to nearby targets (i.e. where the distance is close to zero), and has no particular preference for any direction, as indicated by the fact that a value of 1 falls within the HPDI for all directions.

Figure 6: Posterior distributions for our model when trained on data from Clarke et al. (2018). The shaded ribbons for the proximity and direction weighting indicate 53%, 89% and 97% HPDIs. Note that there is a higher weighting for directions closer to π i.e. completely reversing the direction, indicating negative momentum.

Figure 7: Individual differences in Clarke et al. (2018). The first two rows show the correlations in the random effect structure between different parameters. We can see that p_s and p_d appear to be independent from one another, and independent from the spatial biases b_p and p_m . However, b_p and p_m are correlated with one another in both the feature and conjunction conditions. The bottom row shows the correlation between conditions for each parameter in our model.

Figure 8: Posterior predictions for Clarke et al. (2018). (*top:*) Scatter plots between empirical and predicted summary run statistics. Error bars indicated 89% HPDI for the posterior predictions. (*bottom:*) Comparison of inter-target distances and how they vary over time. Shaded region gives 89% intervals.

Supplementary materials part 1: supplementary materials for the bag foraging model.

Supplementary materials part 2: supplementary materials for the spatial foraging model.

Bag foraging model - one trial: the Stan code for the basic version of our sampling-without-replacement model for visual foraging.

Bag foraging model - multiple trials and conditions: the Stan code for the version of our sampling-without-replacement model for visual foraging, incorporating multiple trials and conditions.

Bag foraging model - multi level: the Stan code for the version of our sampling-without-replacement model for visual foraging, incorporating a multi-level random effects structure.

Spatial foraging model: the Stan code for the version of our sampling-without-replacement model for visual foraging, incorporating spatial biases.

Spatial foraging model - multi level: the Stan code for the version of our sampling-without-replacement model for visual foraging, incorporating spatial biases and a multi-level random effects structure.