

# Corvid Caching: Insights From a Cognitive Model

Elske van der Vaart, Rineke Verbrugge, and Charlotte K. Hemelrijk  
University of Groningen

Caching and recovery of food by corvids is well-studied, but some ambiguous results remain. To help clarify these, we built a computational cognitive model. It is inspired by similar models built for humans, and it assumes that memory strength depends on frequency and recency of use. We compared our model's behavior to that of real birds in previously published experiments. Our model successfully replicated the outcomes of two experiments on recovery behavior and two experiments on cache site choice. Our "virtual birds" reproduced declines in recovery accuracy across sessions, revisits to previously emptied cache sites, a lack of correlation between caching and recovery order, and a preference for caching in safe locations. The model also produced two new explanations. First, that Clark's nutcrackers may become less accurate as recovery progresses not because of differential memory for different cache sites, as was once assumed, but because of chance effects. And second, that Western scrub jays may choose their cache sites not on the basis of negative recovery experiences only, as was previously thought, but on the basis of positive recovery experiences instead. Alternatively, *both* "punishment" and "reward" may be playing a role. We conclude with a set of new insights, a testable prediction, and directions for future work.

**Keywords:** computational model, Clark's nutcracker, Western scrub-jay, caching, memory

**Supplemental materials:** <http://dx.doi.org/10.1037/a0022988.supp>

Like many animals, corvids cache food items, saving them for later (vander Wall, 1990). Some species, like the Clark's nutcracker, bury thousands of pine seeds, and rely on them for months, surviving on little else throughout the winter (vander Wall & Balda, 1981); other corvids, like the common raven, hide just a few scraps of meat for a duration of hours, to get their fair share of a carcass (Heinrich & Pepper, 1998). Corvid caching and recovery has been extensively studied in the laboratory, and has been shown to depend on memory (for a review, see de Kort, Tebbich, Dally, Emery, & Clayton, 2006). However, despite a wealth of empirical data, some results are difficult to interpret.

In this paper, we attempt to clarify two such ambiguous results, one concerning the recovery of caches, the other concerning choice of cache sites. We do this using a new approach: That of computational cognitive modeling. Although many computational models have been built to study animals (Grimm & Railsback, 2005), so far, the focus has tended to be on ecological questions, such as

'what causes dominance hierarchies to form?' (Hemelrijk, 1999) or 'how do individuals decide when to migrate?' In contrast, cognitive questions, that concern animal memory, learning, or problem solving, are seldom subjected to this approach (Penn, Holyoak, & Povinelli, 2008).

However, computational models of *human* cognition have proven to be very useful, both for uncovering basic mechanisms and for predicting performance on practical tasks (Anderson, 2007; Sun, 2008). A single cognitive architecture, such as ACT-R, the Adaptive Control of Thought–Rational (Anderson et al., 2004), can be used to investigate topics ranging from list memory (Anderson, Bothell, Lebiere, & Matessa, 1998) to interval timing (Taatgen, van Rijn, & Anderson, 2007), and from learning mathematics (Anderson, 2005) to driving cars (Salvucci, 2006). Therefore, it seems likely that applying the same technique to corvids will also produce new insights.

The first issue that we explore with our computational model is the fact that, in the laboratory, Clark's nutcrackers become significantly less accurate as recovery progresses. As they find more and more of their own caches, they make more and more errors, over the course of just a few days (Balda, Kamil, & Grim, 1986; Kamil & Balda, 1990). This is surprising, as in the wild, these birds accurately recover thousands of seeds, months after caching them. Balda and Kamil et al. explain this apparent discrepancy by hypothesizing that Clark's nutcrackers remember some caches better than others, and recover their best-remembered caches first. However, attempts to find out what might make certain sites more memorable than others have so far not been successful (Kamil, Balda, & Good, 1999).

The second curious result that we focus on comes from a set of experiments by de Kort, Correia, Alexis, Dickinson, and Clayton

---

Elske van der Vaart, Artificial Intelligence and Theoretical Biology/Behavioural Ecology and Self-Organisation, University of Groningen, Groningen, the Netherlands; Rineke Verbrugge, Artificial Intelligence, University of Groningen; Charlotte K. Hemelrijk, Theoretical Biology/Behavioural Ecology and Self-Organisation, University of Groningen.

This article merges, updates, and expands work previously reported in unarchived conference proceedings (van der Vaart, Hemelrijk, & Verbrugge, 2009; van der Vaart, Verbrugge, & Hemelrijk, 2009). This research was supported by the Netherlands Organisation for Scientific Research, in the form of TopTalent Grant 021.001.089, awarded to Elske van der Vaart, and Vici Grant 277-80-001, awarded to Rineke Verbrugge.

Correspondence concerning this article should be addressed to Elske van der Vaart, Artificial Intelligence, University of Groningen, P.O. Box 407, 9700 AK, Groningen, the Netherlands. E-mail: e.e.van.der.vaart@rug.nl

(2007). In this paper, the question is whether Western scrub jays, another corvid species, can learn to adjust their choice of cache sites in response to their experiences at recovery, and if so, what the mechanism is: Are they reacting to *reward*, preferring to cache in locations where they have previously found worms, or are they influenced by *punishment*, avoiding locations where they have previously failed to find worms? Here, the problem is that two experiments seem to offer contradictory answers to this question: In one, the birds' behavior seems to indicate that reward is at work, while in the other, punishment seems to be the best explanation.

Our computational model tackles these two issues in the following way. First, we implement a basic set of assumptions about corvid cache and recovery cognition as a working computer program, creating a kind of "virtual bird." This set of basic assumptions is built on the basis of existing literature, both on corvids and on cognitive models. Then, we expose our virtual bird to simulated versions of the empirical setups, and see how its behavior compares to that of its biological counterparts.

We start from three basic assumptions: First, that when an individual caches or recovers, it encodes one, integrated memory of the event (Clayton, Yu, & Dickinson, 2001). In the model, this is accomplished by creating a *chunk*, which encodes the event's location and success. Second, that the fresher a memory is, and the more often it has been used, the easier it will be to recall (Anderson & Schooler, 1991). In the model, this is done by giving every chunk an *activation*, which decays with time, but increases with every update. And third, that a bird's memories of its earlier cache and recovery experiences strongly affect where it will cache and recover next.

In our simulations, we investigate whether this basic model is sufficient to replicate the different patterns evident in the recovery behavior experiments (Balda et al., 1986; Kamil & Balda, 1990). In addition to declining recovery accuracy, the patterns considered are the frequency of revisits to emptied sites and the lack of correlation between caching and recovery order. Furthermore, by systematically varying the kind of learning in the model, we study whether scrub jays are most likely motivated by reward or by punishment in the cache site choice experiments (de Kort, et al., 2007).

The paper is organized as follows. First, we describe the computational model and its underlying assumptions. Then, we provide an introduction to the recovery behavior experiments conducted with Clark's nutcrackers (Balda et al., 1986; Kamil & Balda, 1990), and show our model's fit of each. This is followed by a description of the cache site choice experiments done with scrub jays (de Kort, et al., 2007) and our model's replications of each. We conclude with a general discussion of our model, its findings, and future work.

## Method

The original corvid experiments under consideration all share much of the same structure: They consist of a sequence of *caching* and *recovery* sessions, where the birds are tested individually. In the caching sessions, the birds are generally sated, and offered food items to cache; in the recovery sessions, they are usually hungry, and the only food available is what they have cached in previous sessions. Another common feature is that the subjects are always offered a discrete set of cache sites to choose from. In the

recovery behavior experiments (Balda et al., 1986; Kamil & Balda, 1990), this set of cache sites consists of a series of cups in the floor of an experimental room. Each cup is filled with sand, and placed in a regular, rectangular pattern. In the experiments on cache site choice (de Kort, et al., 2007), the available sites are the individual sections of  $2 \times 8$  ice cube trays, filled with corn kibble. These ice cube trays are individually marked by a unique arrangement of colored blocks, and are kept the same across trials. In addition, every ice cube tray is always placed in the same location in a bird's home cage, where testing occurs.

As the experimental setup of the recovery behavior and the cache site choice studies was similar, we can simulate them with a single computational model. It consists of three parts: A *cognitive model*, a *setup model* and a *simulator model*. The cognitive model is a computational theory of the cognitive processes under concern; the virtual bird part of our work. The setup model keeps track of the state of the "physical world" in the original experiments: What the potential cache sites are, how many items are currently cached there, and so on. The simulator model, finally, is what actually runs the experiments: It ensures that the cognitive model and the setup model are initialized, that the virtual birds are offered the right number of caching and recovery sessions, and that data is collected for further analysis.

## The Cognitive Model

Our cognitive model focuses specifically on *how* corvids decide *where* to cache or recover. This means that motivational issues are not considered; in caching sessions, every virtual bird is simply instructed to cache as often as the real birds in the corresponding experiment did, and in recovery sessions, it continues to recover until the simulator model tells it to stop. This can occur when the virtual bird has recovered all its caches, or when it has recovered a specific number of caches, or when it has made the maximum number of recovery attempts,  $m_{ra}$ . For the recovery behavior experiments,  $m_{ra}$  is set to 150; for the cache choice experiments,  $m_{ra}$  is set to 40; in each case, this is about twice the sites available. As we do not know the maximum number of recovery attempts made in each set of experiments, we aimed only to choose reasonable values; they were not fit to the experimental data.

Every time a virtual bird is asked to cache or recover, it evaluates all its possible options—all the discrete cups or ice cube tray sections that are on offer. For each of these options, it estimates the "attractiveness" of caching or recovering there, depending on the type of session it is currently in. In caching sessions, the cognitive model estimates "cache attractiveness"; in recovery sessions, it estimates "recovery attractiveness." This attractiveness is a number, which the virtual bird calculates on the basis of its experiences. These experiences are explicitly encoded in memory, in *chunks*. Chunks, and the way they determine the behavior of our virtual birds, are explicitly described below, with a flowchart summarizing all the component processes in Figure 1.

**The basics of chunks.** A *chunk* encodes an event's *type* and *location*; "type" refers to whether the event was a caching or recovery event, and "location" refers to the associated cache site. In the case of a recovery event, a chunk also records *success*, which refers to whether or not a cache was actually found. Thus, a chunk is a memory of a particular kind of *experience*—caching in a site, successfully recovering there, or unsuccessfully recovering

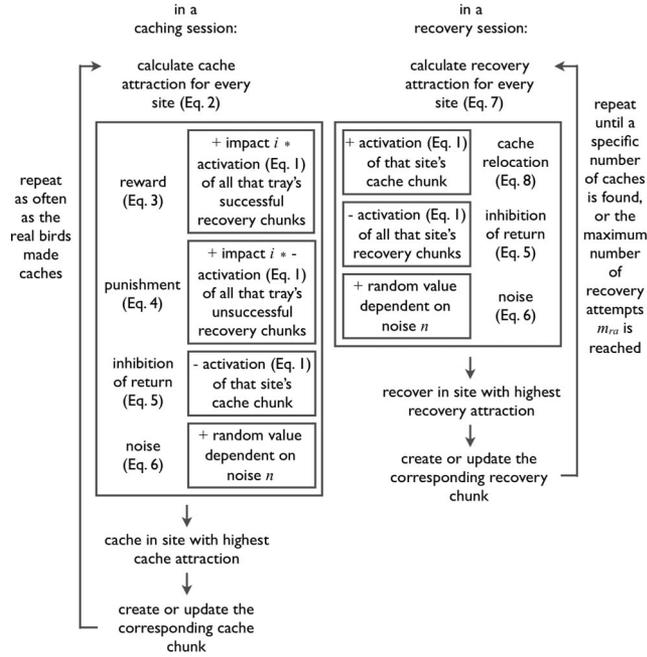


Figure 1. Flowchart summarizing the cognitive model's component processes.

there. Every time the cognitive model experiences one of these events, it creates the appropriate chunk, and encodes it in memory. If the appropriate chunk already exists, that chunk receives an *update* instead. These updates help determine the chunk's *activation*, which represents the memory's 'strength.' A chunk  $h$ 's activation  $A_h$  depends on its *recency* and *frequency* of use, and is calculated according to Equation 1, adapted from the cognitive architecture ACT-R (Anderson, 2007; Anderson & Schooler, 1991). In Equation 1,  $t_j$  represents the elapsed time  $t$  since use  $j$  of chunk  $h$ , while  $d$  is a decay parameter.

$$A_h = w_h \cdot \sum_j t_j^{-d} \quad (1)$$

The weighing factor  $w_h$  is determined by chunk  $h$ 's type. Cache chunks have weight  $w_c$ , successful recovery chunks have weight  $w_{r1}$ , and unsuccessful recovery chunks have weight  $w_{r0}$ . This weighing factor is included because when Clark's nutcrackers are caching, they probe the substrate of cups considerably more often than when they are recovering (Kamil et al., 1999). In addition, they probe cups slightly more when they are successfully retrieving nuts, than when they fail to find anything (Kamil et al., 1999). We have taken this as a proxy measure of memorability—the more a bird is engaged in a particular act, the more memorable it is likely to be—and have set the weights  $w_c$ ,  $w_{r1}$ , and  $w_{r0}$  to reflect the mean number of bill probes reported for caching, successful, and unsuccessful recovery events, respectively (Kamil et al., 1999); see Table 1.

For the purpose of computing the activations of chunks, time is measured in steps. Every cache or recovery event counts as one step, and time *outside* of the experimental sessions is not considered. Adding a fixed number of steps to represent every real time

day, so that memory strength decays between sessions, does not improve the model's fit. However, it does require an additional parameter that must be estimated, which we consider undesirable. Although real birds definitely do experience memory loss over time, our approach still seems reasonable, as the recovery accuracy of both Clark's nutcrackers and Western scrub jays appears to be insensitive to retention interval within the brief time frames used in these experiments (Bednekoff, Balda, Kamil, & Hile, 1997; Clayton, Yu, & Dickinson, 2003).

**Deciding where to cache.** When the cognitive model is deciding where to cache next, every site's cache attractiveness  $C_k$  is determined by the effects of prior experience plus *noise*. The site with the highest value of  $C_k$  is where the virtual bird will cache.  $C_k$  is computed according to Equation 2. Here,  $L_k$  stands for the influence of *learning*, which can be turned on or off, while  $I_k$  is an *inhibition of return* factor, which makes it less attractive for the virtual bird to return to the same site with the same aim.

$$C_k = L_k - I_k + noise \quad (2)$$

**Learning.** In the experiments on cache site choice (de Kort, et al., 2007), the question is whether scrub jays are responsive to "reward" or to "punishment." Accordingly, our cognitive model is capable of learning from either *positive* or *negative* experiences at recovery. This is captured by a learning effect  $L_k$  on the attractiveness of every cache site  $k$ .

When *reward* is turned on, the cognitive model checks to see if any recovery chunks  $m$  exist that represent *successful* recovery attempts anywhere in the tray of site  $k$ . If such recovery chunks  $m$  exist, the learning effect  $L_k[\text{reward}]$  on site  $k$ 's attractiveness is equal to the total activation  $\sum A_m$  of those chunks  $m$ , modified by an impact factor  $i$ . This is captured by Equation 3. In other words, the stronger the model's memories of successfully recovering from a particular tray, the more attractive it finds it to cache there again.

$$L_k[\text{reward}] = i \cdot \sum A_m \quad (3)$$

Conversely, when *punishment* is turned on, the cognitive model looks for any recovery chunks  $n$  that represent *unsuccessful* recovery attempts anywhere in the tray of site  $k$ . If such recovery chunks  $n$  exist, the learning effect  $L_k[\text{punishment}]$  on site  $k$ 's attractiveness is equal to minus the total activation  $\sum A_n$  of those chunks  $n$ , modified by an impact factor  $i$ , as shown in Equation 4. Put differently, the stronger the model's memories of unsuccessfully recovering from a particular tray, the less attractive it finds it to cache there again.

$$L_k[\text{punishment}] = i \cdot \sum A_n \quad (4)$$

**Inhibition of return.** To calculate the inhibition of return  $I_k$  associated with a site  $k$ , the cognitive model checks whether any chunks  $l$  exist of the current session's type, that refer to the same

Table 1  
Fixed Settings

	$m_{ra}$	$w_c$	$w_{r1}$	$w_{r0}$
Experiment 1	150	5.02	1.82	1.44
Experiment 2	40	5.02	1.82	1.44

site. If any such chunks  $l$  exist, the effect of inhibition of return,  $I_k$ , associated with site  $k$  is equal to Equation 5, where  $\sum A_l$  is the sum of the activations of all such chunks  $l$ . Thus, the stronger the model's memories of already having cached or recovered in a particular location, the less likely it is to visit there again with the same purpose.

$$I_k = \sum A_l \quad (5)$$

**Noise.** Every site's attractiveness always has a *noise* component, representing sources of transient error. The noise term is computed according to Equation 6, taken from the cognitive architecture ACT-R (Anderson, 2007), where  $n$  is a parameter that we tune, and  $r$  is a random value between 0 and 1.

$$\text{noise} = n \cdot \ln \frac{1-r}{r} \quad (6)$$

**Deciding where to recover.** For the cognitive model, deciding where to recover is similar to deciding where to cache. Every site's recovery attractiveness  $R_k$  is equal to the effects of prior experience plus *noise*. Whichever site currently has the highest attractiveness  $R_k$  will be selected for recovery, as decided by Equation 7. Here,  $F_k$  is a cache relocation effect, which helps the virtual bird recover at sites where it has previously cached, while the *inhibition of return*  $I_k$  and *noise* are the same as described previously.

$$R_k = F_k + I_k + \text{noise} \quad (7)$$

**Cache relocation.** While recovering, the cognitive model calculates the cache relocation effect  $F_k$  associated with a site  $k$  by looking for an existing cache chunk  $o$  referring to the same site. If such a chunk is found, the cache relocation effect associated with site  $k$  is equal to that chunk's current activation  $A_o$ , according to Equation 8. Thus, the stronger the model's memory of having cached somewhere, the more attractive it finds it to recover there.

$$F_k = A_o \quad (8)$$

**Putting it all together.** Let us imagine that the virtual bird is instructed to choose a cache site, and that its memory contains three chunks: A cache chunk coding for site  $A_1$  in tray A, created on the previous time step, a cache chunk coding for site  $B_1$  in tray B, created during the previous trial, and a successful recovery chunk coding for site  $B_1$  in tray B, also created during the previous trial. Now, to decide where to cache next, the virtual bird calculates the *cache attractiveness* of all available sites, according to Equation 2.

For site  $A_1$  in tray A, the virtual bird has an associated cache chunk, so its attractiveness is dampened by *inhibition of return*, as given by Equation 5. The amount of inhibition of return depends on the activation of the cache chunk, as given by Equation 1. In addition, the site's attractiveness is modified by *noise*, as given by Equation 6. For the other sites in tray A, the virtual bird has no relevant memories, and so their attractiveness is determined by noise only.

For site  $B_1$  in tray B, the virtual bird has both a cache chunk and a successful recovery chunk from the previous trial. Therefore, inhibition of return is active, as is *reward learning*, according to Equation 3, if it is turned on. Inhibition of return causes the attractiveness of caching in site  $B_1$  to *decrease* with the activation of the cache chunk, while reward learning causes the attractiveness

of caching in site  $B_1$  to *increase* with the activation of the recovery chunk, modified by the impact factor  $i$ . Thus, in this case, the negative effect of inhibition of return is partially reduced by reward learning. However, because cache chunks are more heavily weighted than recovery chunks (see Table 1) inhibition of return is likely to be stronger than reward learning, depending on the chunks' relative ages, and the impact factor  $i$ . However, the cache attractiveness of all other sites in tray B is likely to be high, because inhibition of return does not generalize across trays, while reward learning *does*. This is because Equation 3 looks for all successful recovery chunks corresponding to a site's *tray*, not just the site itself. Finally, the cache attractiveness of all sites is modified by noise.

Whichever site is now estimated to be most attractive, is where the virtual bird will cache. Because of the effects of noise, this can still be any of the available sites. However, on *average*, the unused sites in tray B should now be most attractive to cache in, because of reward learning, while site  $A_1$  in tray A should be least attractive, because of inhibition of return, with the cache attractiveness of all other sites between these two extremes. The cognitive model will continue to make this calculation for each subsequent cache choice, based on its memories, until the simulator model tells it to stop. A recovery session works in exactly the same way, except that the virtual bird is comparing the *recovery attractiveness* of sites, according to Equation 7.

**Comparing to empirical data.** All model results are based on a sample size of 100 virtual birds for each experimental condition, using the parameter settings of Table 2. Because of the noise in the model, smaller sample sizes produce more variable outcomes. For the decay parameter  $d$  and the noise parameter  $n$ , we selected the settings that most closely matched the recovery accuracy of the real birds in the first experiment, Experiment 1a, and then used the same settings throughout the paper. For the impact parameter  $i$ , which is only used in Experiment 2, we chose the value that produced the best fit of Experiment 2a, and then kept it for Experiment 2b, without any further fitting. Within the range of values in Table 2, all possible combinations were tried. More extreme values were not considered, as these always resulted in recovery that was not accurate enough. However, the model's qualitative results do not depend strongly on either its exact parameter settings or the sample size used (see the online Supplementary Material).

## Experiment 1: Recovery Behavior

To investigate what insights our model can provide into the recovery behavior of Clark's nutcrackers, we compare its performance to that of the real birds in the first experiments of Balda et al. (1986) and Kamil and Balda (1990).

Table 2  
Adjustable Parameters

	Decay, $d$	Noise, $n$	Impact, $i$
Experiment 1	0.2	2	—
Experiment 2	0.2	2	0.9
Range evaluated	0–1, by 0.1	0–5, by 0.5	0–2, by 0.1

### Experiment 1a: Revisits to Emptied Sites (Experiment 1 in Balda et al., 1986)

The main objective of this experiment was to discover why Clark's nutcrackers tend to return to already emptied cups, termed "revisits" (Kamil & Balda, 1985). In the wild, when corvids recover their caches, this leaves visible signs of digging. In the laboratory, however, the sand is smoothed over between recovery sessions. This means that the birds have no visual reminders of where they have already recovered their caches.

Balda and coauthors hypothesized that perhaps this is what causes Clark's nutcrackers to make so many revisit errors in the laboratory. They tested this hypothesis by running an experiment with two conditions: The *marked* and *unmarked* condition. In the marked condition, the scattered sand of earlier digging was there at recovery; in the unmarked condition, it was not. Both conditions consisted of one caching session followed by three recovery sessions, which started ~10 days later, and occurred on alternate days. In each caching session, the birds were allowed to make about 20 caches; in each recovery session, they were allowed to recover approximately a third of what they had hidden. Four birds were used in the experiment, all of which participated in both conditions.

Balda et al. (1986) found no significant differences between the two conditions. Irrespective of whether there were signs of previous digging, two out of four Clark's nutcrackers always made more revisits than expected by chance, although only at the 10% significance level. In total, there were 16 recovery sessions where revisits could occur—four birds  $\times$  two conditions  $\times$  two sessions with emptied cups—so that 50% of sessions involved more revisits than expected by chance. Figure 2, left panel, shows the mean number of revisits made, collapsed across the "marked" and "unmarked" conditions.

In addition, the birds' mean accuracy declined significantly across recovery sessions, as defined by the total number of caches recovered divided by the total number of cups visited; see Figure 2, right panel. Nevertheless, of the 24 recovery sessions conducted with the birds—four birds  $\times$  two conditions  $\times$  three sessions—accuracy was significantly higher than chance in 19, which is 79% of the total.

**Model results.** We run 100 virtual birds only once, through one caching session followed by three recovery sessions. First, we

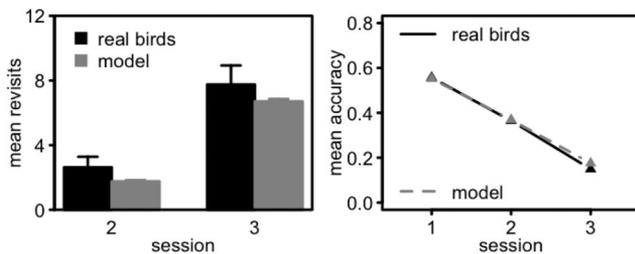


Figure 2. Results of Experiment 1a, real birds, Balda et al. (1986), and computational model. Left: Mean revisits to emptied sites, with SEs. Right: Mean recovery accuracy, as defined by the total number of caches recovered divided by the total number of cups visited. Adapted from "Revisits to emptied cache sites by Clark's nutcrackers (*Nucifraga columbiana*)" by R. P. Balda, A. C. Kamil, and K. Grim, 1986, *Animal Behaviour*, 34, 1289–1298. Copyright 1986 by Elsevier. Adapted with permission.

find that at the 10% significance level measured by Balda and colleagues, the virtual birds revisit significantly more emptied sites than expected by chance in 31% of recovery sessions, and that their mean revisit frequencies are similar to the empirical data; see Figure 2, left panel. Second, as shown in Figure 2, right panel, we also find a significant decrease in recovery accuracy,  $F(2, 198) = 230.33$ ,  $p < .001$ , and that recovery accuracy is greater than chance in 87% of sessions.

### Experiment 1b: Manipulating Recovery Order (Experiment 1 in Kamil and Balda, 1990)

Here, Kamil and Balda (1990) specifically investigated *why* Clark's nutcrackers become less accurate across recovery sessions. To test the idea that this might be because they remember some cache sites better than others, 10 Clark's nutcrackers were exposed to two experimental conditions: The *quarters* and the *free* condition.

In the quarters condition, the birds could only recover from a quarter of the room every recovery session; in the free condition, the whole room was available. Each condition consisted of one caching session, followed by four recovery sessions. In both conditions, during the caching session, only 32 cups were available for caching, eight in each quarter of the room. The birds were allowed to store seeds until they had created at least three caches in every quarter. A week later, recovery sessions began, conducted on successive days. This is where the conditions differed: In the quarters condition, only one quarter of the room was available for recovery during each session, while in the free condition, all cups were always open. In the quarters condition, the birds could continue to recover until they had retrieved all caches created in the available quarter; in the free condition, they were allowed to recover 25% of their caches each session. Seven birds successfully completed both conditions. Now, if Clark's nutcrackers remember some cache sites better than others, their accuracy should have stayed the same across recovery sessions in the quarters condition, but it should have declined in the free condition.

As can be seen in Figure 3, the results were as predicted. Specifically, there was no main effect of either condition or session, but there was a significant interaction between the two. Subsequent *t* tests revealed that this was because of a significant decline in accuracy in the free condition. However, this result was difficult to interpret, because chance levels of accuracy stayed the same in the quarters condition, but decreased in the free condition. In the quarters condition, a new set of cups and caches was available every recovery session; in the free condition, by contrast, the full set of cups was open from the start, but the number of caches left to be recovered kept decreasing. To compensate for this effect, Kamil and Balda calculated a 'modified accuracy score' that compensates for chance levels of accuracy, as derived from Olton and Samuelson (1976); see Equation 9. With these modified scores, Kamil and Balda again found a significant condition by session interaction, because of a significant decrease in accuracy in the free condition. In further analysis, no general relationship between caching and recovery order was found.

$$\text{modified accuracy} = \frac{(\% \text{ accuracy} - \% \text{ chance})}{(100 - \% \text{ chance})} \quad (9)$$

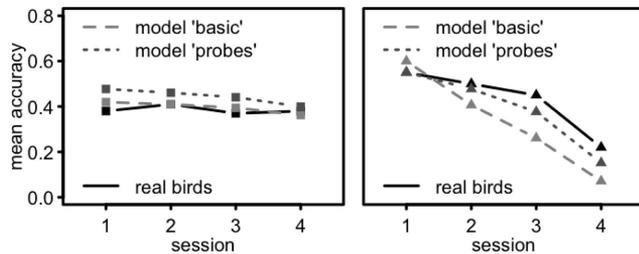


Figure 3. Mean recovery accuracy of Experiment 1b, real birds, Kamil and Balda (1990) and computational models; left: Quarters condition, right: Free condition. Here, the “probes” model makes 50% of its recoveries randomly, to reflect exploratory probes. Adapted from “Differential memory for different cache sites by Clark’s nutcrackers (*Nucifraga columbiana*)” by A. C. Kamil and R. P. Balda, 1990, *Journal of Experimental Psychology: Animal Behavior Processes*, 16, 162–168. Copyright 1990 by the American Psychological Association.

**Model results.** As can be seen in Figure 3, ‘basic model,’ the behavior of our virtual birds is similar to that of the Clark’s nutcrackers: Accuracy decreases in the free condition, but not in the quarters condition. However, because of our larger sample size, we find significant main effects of both condition,  $F(1, 99) = 28.054, p < .001$ , and session,  $F(3, 297) = 139.65, p < .001$ , as well as an interaction effect,  $F(3, 297) = 80.948, p < .001$ . When we calculate modified accuracy scores, we find no main effect of condition,  $F(1, 99) = 2.725, p = .102$ , but we do find an effect of session,  $F(3, 297) = 90.474, p < .001$ , as well as an interaction effect,  $F(3, 297) = 48.007, p < .001$ . Furthermore, like the real birds, our virtual birds show no systematic relationship between caching and recovery order in the free condition; of the 100 correlations, only six are significant,  $p < .05$ , and they occur in both directions.

## Discussion

In these two experiments, our virtual birds reproduce two important aspects of Clark’s nutcracker recovery behavior: They frequently revisit emptied cache sites, and their recovery accuracy declines across sessions. For the former effect, three possible explanations have been offered: Either the birds remember their cache sites but not their recovery attempts, *or* they remember both, but continue to make revisits for unknown reasons, *or* they start to forget their cache sites as soon as they recover them, and make revisits because of ‘incomplete forgetting’ (Balda et al., 1986; Kamil, Balda, Olson, & Good, 1993). For our virtual birds, the answer lies halfway between the first two explanations: They remember both, but they remember cache sites *better*, because of the different weights used to calculate the activations of cache chunks and recovery chunks. As explained in the Method section, these different weights are inspired by the fact that Clark’s nutcrackers probe cups about five times when caching, but only twice when recovering (Kamil et al., 1999). When the model is deciding where to recover, it looks for cache chunks and recovery chunks corresponding to every possible location. The more active a cache chunk, the higher the attractiveness of recovering there; the more active a recovery chunk, the lower the attractiveness of recovering there. Initially, because the activations of chunks decay with time,

a newer, but lower-weighted recovery chunk will “cancel out” an older, but higher-weighted cache chunk. However, as the activation of the recovery chunk starts to decay also, the “cancelling out” starts to fail, creating revisits to emptied sites. We do not claim that this is the only possible explanation for the revisits made by Clark’s nutcrackers, but it does follow the behavioral evidence.

Our computational model also provides a different explanation for the second pattern; that is, the decline in recovery accuracy across sessions. Kamil and Balda (1990) concluded that this was because of Clark’s nutcrackers remembering some cache sites better than others, because the effect disappeared when recovery order was controlled. For our virtual birds, however, the memorability of all cache sites starts out the same. The activation of a cache chunk depends only on how often the virtual bird has cached in that particular site, and how recently it has done so. However, this cannot be causing the model’s decreasing recovery accuracy, because then we would find no difference between conditions in Experiment 1b, as well as a significant, negative correlation between caching and recovery order.

Instead, what explains our model’s fit of Experiment 1b is the differing chance levels between the quarters and the free condition. In the quarters condition, the number of caches that can be recovered remains the same across sessions, while in the free condition, it declines. Kamil and Balda discarded this as a possible factor because the birds’ accuracy actually declines *faster* than chance in the free condition. In the model, this decrease in chance levels is “magnified” by *noise*. When the virtual bird is deciding where to recover, it calculates the attractiveness of doing so in all possible cups. On average, cups where it has cached will be more attractive than cups where it has not cached, because of the positive effect of existing cache chunks. However, noise may cause an “incorrect” cup to temporarily be more attractive than all “correct” cups. The lower the ratio of correct to incorrect cups, the higher the odds of this occurring, and the higher the noise level, the larger the effect. In the quarters condition, the ratio of correct to incorrect cups remains the same, while in the free condition, it decreases across sessions; this, in combination with noise, is what explains our model’s decline in recovery accuracy.

One aspect of Clark’s nutcracker recovery behavior that our cognitive model currently fails to capture is the fact that many recovery attempts in cups that never contained caches are probably not true errors, but ‘exploratory probes.’ Evidence for this comes from the fact that when the costs of making a recovery attempt are increased, the number of errors made drops significantly (Bednekoff & Balda, 1997). In the model, on the other hand, all errors are ‘true errors,’ caused by noise. However, we *can* adapt the model so that it too can make “exploratory probes”—where it randomly chooses a recovery cup, instead of computing cups’ attractiveness on the basis of its memory—and this gives us the results of Figure 3, ‘probe model.’ In this run, 50% of recovery attempts were in fact exploratory probes, and *noise* was set to 1. This illustrates how the model can be used to investigate the consequences of different assumptions concerning the strategies used by real birds, as well as how incorporating such strategies might improve its fit of the empirical data.

## Experiment 2: Cache Site Choice

In de Kort et al. (2007), it was found that Western scrub jays are sensitive to the fate of their caches at recovery. When the birds

were offered one caching tray that was always pilfered, and another that was always returned full, they learned to preferentially cache in the latter tray. This experiment, however, left open the question of whether the birds were learning through *reward*, preferring to cache in the tray that was always reinforced, or if they were learning through *punishment*, avoiding the tray that was never reinforced. These two accounts were contrasted in de Kort et al.'s (2007) Experiment 4a and 4b, and it is these two experiments that we simulate.

However, as we built our model, we discovered that 'reinforcement,' or the lack thereof, can be interpreted in two ways. Either, it can relate specifically to a bird relocating its *own* caches in their expected locations, or it can relate to finding *any* worms, irrespective of whether the bird has cached them there. Therefore, we tried 'reward by recovering a cache,' 'reward by finding any worm,' 'punishment by finding a cache missing,' and 'punishment by any failed recovery attempt.' Of those combinations, only the first (reward by recovering a cache) successfully reproduced all the empirical data. Therefore, we present the results of this model in detail, and explain why the others failed to fit in the subsequent discussion. We thus run the model with the reward of Equation 3 turned on, the punishment of Equation 4 turned off, and the additional restriction that virtual birds only learn from recovery attempts directed at sites where they have previously cached. In the model, this is implemented by specifying that the learning effect  $L_k[\text{reward}]$  associated with a particular site is only affected by successful recovery chunks for which a corresponding cache chunk also exists.

### Experiment 2a: Return and Move (Experiment 4a in de Kort et al., 2007)

In this experiment, some caches were *moved* from one tray to another, in an effort to distinguish between the reward and punishment accounts. It consisted of two training trials, followed by one test trial. The training trials consisted of a caching session and a recovery session, one day apart. In the caching session, eight scrub jays were presented with two trays, A and B, but one of them, tray B, was unavailable for caching, as it was covered by a lid. However, during the recovery session, both trays were fully accessible. At this point, the birds in the *return condition* found their caches returned to them in tray A, where they had left them, while those in the *move condition* found their caches moved to tray B, with tray A now empty. Now, the question was where the birds in the two conditions would prefer to cache on the test trial, in which both trays were fully accessible. As shown in Figure 4, top left, the birds in the return condition cached predominantly in tray A, while those in the move condition cached predominantly in tray B, with a significant condition by tray interaction. However, subsequent analysis confirmed only that the "return birds" significantly preferred to cache in tray A, while failing to confirm that the "move birds" significantly preferred to cache in tray B. Thus, this experiment was considered inconclusive.

**Model results.** Given that our virtual birds learn only from recovering their own caches, we might expect them to only show tray preferences in the return condition. However, like the real birds, our virtual birds prefer tray A in the return condition, and tray B in the move condition; see Figure 4, middle left. Although this last preference seems numerically weaker in the cognitive

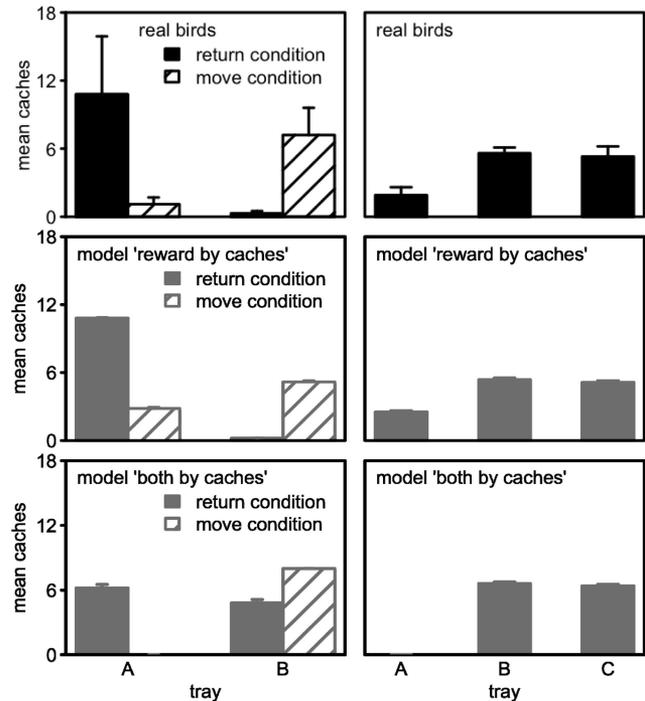


Figure 4. Average caches in trays in Experiment 2, with SEs; left column: Experiment 2a, right column: Experiment 2b; top row: Real birds, de Kort et al. (2007), middle row: Standard model, with reward by recovering a cache, bottom row: Alternative model, with both punishment and reward by recovering a cache, with  $\text{impact} = 1.9$ . Adapted from "The control of food-caching behavior by western scrub-jays (*Aphelocoma californica*)" by S. R. de Kort, et al., 2007, *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 361–370. Copyright 2007 by the American Psychological Association.

model, it is still statistically significant, unlike in the empirical data, because of our much larger sample size. We find a significant condition by tray interaction,  $F(1, 198) = 2265.4, p < 0.001$ , as well as a significant preference for tray A in the return condition  $F(1, 99) = 11934, p < .001$  and for tray B in the move condition,  $F(1, 99) = 84.642, p < .001$ . In terms of the underlying mechanism, it turns out that in the move condition, our virtual birds are not really preferring tray B, but instead avoiding tray A, because of inhibition of return; see Equation 5. In the return condition, this effect is overpowered by reward learning. In the absence of any positive recovery experiences, however, inhibition of return drives our virtual birds to avoid sites where they have previously cached, and in the move condition, that includes most of tray A.

### Experiment 2b: Move, New Tray (Experiment 4b in de Kort et al., 2007)

The setup of this experiment again involved movement of caches, but now involving a third, neutral tray. It comprised two training trials, each consisting of a caching and a recovery session, and a single test trial, consisting of a caching session only. Every session, four birds were offered three trays, A, B, and C. During the training trials, caching in trays B and C was not possible, as access to these trays was blocked. However, at recovery, the scrub

jays found all their caches moved from tray A to tray B, leaving trays A and C empty. On the test trial, all three trays were available for caching, with the results of Figure 4, top right. There was a significant effect of tray on the number of worms cached, and this was because of the fact that both trays B and C contained more worms than tray A, while not differing from one another.

**Model results.** As can be seen in Figure 4, middle right, our virtual birds distribute their caches much as the real birds did: They cache primarily in trays B and C, and very little in tray A. This result is because of inhibition of return: Having already cached in tray A twice, the virtual birds find it unattractive to cache there again. We find a significant effect of tray on the number of worms cached,  $F(2, 198) = 80.657, p < .001$ , and Tukey post hoc comparisons confirm that trays A and B contain significantly different numbers of caches, as do trays A and C,  $p < .001$ , while trays B and C do not.

## Discussion

Of our original four hypotheses concerning learning—‘reward by recovering a cache’, ‘reward by finding any worm’, ‘punishment by finding a cache missing’, and ‘punishment by any failed recovery attempt’—our cognitive model provides support for the first, by showing that it is consistent with the empirical data in these two experiments. The idea that successful recoveries reinforce particular caching strategies is not new; Clayton et al. (2005) already propose it. However, de Kort et al. (2007) discard this explanation because it does not account for the results of the ‘move, new tray’ experiment. Yet, in our model, it *can*, because of the effects of a completely different mechanism: Inhibition of return. Originally included in the model simply to spread out caching within sessions—real birds do not cache in the same sites over and over—it unexpectedly turns out to account for all the results that the ‘reward by recovering a cache’ hypothesis cannot explain. In the move condition of the ‘return and move’ experiment, the virtual birds avoid tray A because they have previously cached there; in the ‘move, new tray’ experiment, the same thing happens. Thus, ‘reward by recovering a cache’, coupled with inhibition of return, can reproduce all of the empirical data.

However, within our present framework, ‘reward by recovering a cache’ also seems to be the only hypothesis that can be made to fit. The most interesting alternative to consider is that of ‘punishment by finding a cache missing’, as it is the one de Kort et al. (2007) seem to favor. They explain the birds’ behavior in these experiments as inhibited caching in pilfered locations, using a “memory at retrieval” process. The idea is that when the scrub jays attempt to recover a cache and find it missing, they recall the corresponding caching episode, and associate their current state of frustration with it, which discourages them from caching in that location again. As a result, the scrub jays learn to avoid trays that are pilfered, and place their caches in “safe” trays instead. Thus, according to de Kort and colleagues, the mechanism at work is *punishment*. However, this cannot account for the birds’ preference for tray A in the return condition of the ‘return and move’ experiment. Under this assumption, the virtual birds learn only from finding their own caches missing. Therefore, they cannot learn to distinguish between trays A and B, as they have had no negative experiences with either.

Similarly, we can also rule out the explanations ‘reward by finding any worm’ and ‘punishment by any failed recovery attempt’. In the ‘move, new tray’ experiment, the birds only find worms in tray B, so ‘reward by finding any worm’ predicts that they should also prefer to cache there, which they do not do. The reason for excluding ‘punishment by any failed recovery attempt’ is more subtle, and we only discovered it by actually running the model. It has to do with ‘move, new tray experiment’: At first glance, it seems that if the virtual birds learn only through punishment, they should always cache equally in trays B and C, because they have only lost caches from tray A. However, when we run an altered version of the model, where we let the virtual birds learn from *every* failed recovery attempt, whether or not they have actually cached there, they *always* show a significant preference for tray B. This happens because once the virtual birds fail to find their caches in tray A, they start probing the other trays. Then, they invariably experience more unsuccessful recoveries in tray C than in tray B, because their own caches have been moved to tray B, while tray C is empty. Only if recovery attempts are linked to their own caching events can the birds learn to avoid caching in tray A but treat trays B and C equally.

However, this does not mean that scrub jays cannot be learning from punishment. In particular, our cognitive model is also capable of reproducing the qualitative patterns evident in these experiments by assuming that it is active in *addition* to reward, though only in response to recovery attempts directed at their own caches. If we continue to fix all parameters at the settings derived for Clark’s nutcrackers, this model, ‘both reward and punishment by recovering caches’, cannot do better than Figure 4, bottom row. This result is similar to the empirical data, except that it shows a very weak preference for tray A in the return condition. This is because the virtual birds are now affected by conflicting information; every time they accidentally revisit a cache site that they have already emptied, tray A becomes less attractive to cache in. Nevertheless, many slight variations of the ‘both reward and punishment by recovering caches’ model do produce a robust preference for tray A. For instance, if we change our setting of *noise* so that the birds make less revisits to emptied sites, or if we assume that inhibition of return does not operate across sessions, then we can produce a very acceptable fit (see the online Supplementary Material).

Therefore, at present, our cognitive model cannot meaningfully distinguish between ‘reward by recovering a cache’ and ‘both reward and punishment by recovering a cache’. The first of these models depends on the assumption that inhibition of return works across trials, which has some empirical support: When Kamil and Balda (1985) forced Clark’s nutcrackers to reuse the exact same set of cache cups, the birds became agitated, and made an unprecedented number of caches in ‘out-of-cup’ locations, such as under landmarks. Other studies relevant to this question have been conducted with parids, the other major family of birds that stores food. When Male and Smulders (2007) tested coal tits, the birds avoided caching close to existing caches, but started reusing sites once they had successfully recovered there. This is consistent with our account of “reward learning” usually “overpowering” inhibition of return. On the actual question of reward versus punishment, the evidence from parids is equivocal; in one experiment black-capped chickadees shifted their caching from unsafe to safe sites, but the setup does not make it possible to determine the mechanism (Hampton & Sherry, 1994), while

in another, the birds failed to show any evidence of learning at all (Baker & Anderson, 1995), although admittedly they were only given one trial to do so.

### General Discussion

This paper demonstrates that a computational cognitive model of corvid caching and recovery can successfully replicate the outcomes of experiments concerning recovery behavior and experiments concerning cache site choice. Our virtual birds successfully reproduced declines in recovery accuracy, revisits to previously emptied cache sites, a lack of correlation between caching and recovery order, and a preference for caching in safe locations. By replicating four different experiments with one model—and keeping parameters constant—we show the integrative power of this approach. Of course, merely fitting empirical data is not very interesting. A computational model is only truly useful if it leads to new insights. We would claim that our model does that in four ways: It leads to new interpretations of existing results, it helps clarify the consequences of specific assumptions, it highlights the need for specific empirical data, and finally, it leads to testable hypotheses.

First, with regard to new interpretations. In Experiment 1, on recovery behavior, we show that the declining accuracy of Clark's nutcrackers might be because of chance, rather than because of differential memory for different cache sites, as was previously assumed (Balda et al., 1986; Kamil et al., 1993). This seems to be a useful alternative theory, because the attempt to discover what might make certain sites more memorable than others has not been successful. Both the physical characteristics of cache sites (Kamil & Balda, 1990) and the time spent making them (Kamil et al., 1999) have so far been ruled out as possible explanations.

Furthermore, in Experiment 2, on cache site choice, we show that it is likely that Western scrub jays react to positive recovery experiences when deciding where to cache. Originally, the data from this experiment was interpreted as evidence of scrub jays reacting to punishment, but not to reward (de Kort, et al., 2007). With our computational model, we show that the latter explanation, reward, *can* in fact account for all the empirical results, provided that the effects of inhibition of return are also taken into account. This appears to be a novel insight, while still having empirical support (Kamil & Balda, 1985).

Second, with regard to clarifying the consequences of different assumptions. In the 'move, new tray' experiment, it seems that scrub jays responding only to punishment should always cache equally in trays B and C, because they have only lost caches from tray A. However, running our simulations revealed a counterintuitive result: If the virtual birds learn from every unsuccessful recovery attempt, they invariably learn to prefer tray B over tray C, because their own caches have been moved from tray A to tray B, so that they experience more unsuccessful recoveries in tray C than in tray B.

Third, with regard to highlighting the need for specific experimental data. In Experiment 2, on cache site choice in Western scrub jays, we show that *either* reward, or reward and punishment working together, can explain the patterns found in the empirical data. However, the two accounts are distinguished by how they are affected by other aspects of the birds' behavior. If scrub jays make a lot of revisits to emptied sites, then learning by *both* punishment

and reward becomes very difficult. Too often, the birds incorrectly conclude that a site has been pilfered, when in fact they have already recovered the worms themselves. Thus, although "revisit tendency" and "learning" may seem to be two unrelated issues, our computational model shows that they are in fact intimately related. Our approach therefore identifies the empirical data necessary to distinguish different theories.

Finally, with regard to testable hypotheses. If our model's assumptions are correct, eliminating declining chance effects should eliminate declining recovery performance in Clark's nutcrackers. This was already done in Kamil and Balda (1990), our Experiment 1b, with the quarters and free condition, but here, eliminating declining chance effects *also* eliminated the birds' control over the order in which they recovered their caches. Thus, this experiment did not yet differentiate between our hypothesis, and the idea that Clark's nutcrackers become less accurate because they recover their best-remembered caches first. However, if the ratio of 'cups containing caches' to 'cups not containing caches' was kept constant across sessions, *without* manipulating recovery order, the two accounts offer different predictions. This could be done, for instance, by randomly closing more and more empty cups as the number of full cups dwindles. Then, our model predicts that recovery accuracy should remain constant, while the "differential memory" hypothesis predicts that it should decrease.

As far as we are aware, our model represents the first computational approach to the cognitive aspects of corvid caching. However, both Miller and Shettleworth (2007) and Dawson and coauthors (2010) present related models of spatial learning, where subjects must learn to relocate the rewarded corner of an enclosure. These two models take a very different approach from ours; Miller and Shettleworth (2007) simulate a single choice rule with changing associative strengths, while Dawson et al. (2010) apply a small neural network to simulate learning. In contrast, our model explicitly includes memory for past events, and can easily be used to test the consequences of different behavioral rules. These are complementary approaches; Miller and Shettleworth (2007) and Dawson et al. (2010) are drawing on formal models of classical conditioning, while we are inspired by the cognitive architectures built for humans.

Given the new insights already generated, further development of our model seems worthwhile. At present, it predicts only *where* corvids should cache or recover, without considering why they should cache at all. For Western scrub jays, it seems likely that two motivational systems operate in parallel: One that inflexibly drives them to cache, irrespective of the outcome, and one that is sensitive to hunger, learning, and specific satiety (Clayton & Dickinson, 1999). We can explicitly simulate such motivational systems, by extending our virtual birds with a "caching drive" and a "hunger drive," which together determine their propensity to cache. In this way, we can extend the model's applicability, and also use it to investigate other curious results, such as the fact that Western scrub jays continue to cache even if all their worms are pilfered (de Kort, et al., 2007).

However, there is also considerable potential to investigate completely new cognitive questions. Corvid caching has also been used to study landmark use (Gould-Beierle & Kamil, 1996), future planning (Raby, Alexis, Dickinson, & Clayton, 2007), and social-cognitive skills (Bugnyar & Heinrich, 2005), among others. Many of these experiments feature experimental setups similar to the

ones considered in this paper, with discrete sets of cache sites and relatively simple manipulations of the environment. And because our virtual birds already explicitly encode their caching and recovery experiences in chunks, we can easily extend our model to experiments where corvids must remember particular configurations of landmarks, or the identity of specific observers. Therefore, we aim to apply our cognitive model to such experiments in future.

Another interesting direction might be to use our model to look at the caching and recovery behavior of parids. Although the caching and recovery behavior of tits and chickadees is in many ways similar to that of corvids—they both scatter hoard, and they both use spatial memory to retrieve caches—there are also many intriguing differences. Parids, for instance, seem to have much shorter memories (Balda & Kamil, 1992; Brodin, 2005; Brodin & Kunz, 1997; Hitchcock & Sherry, 1990), and they do not appear to use observational memory to locate the caches of others (Baker et al., 1988; Hitchcock & Sherry, 1995), while corvids do (Bednekoff & Balda, 1996; Watanabe & Clayton, 2007). Attempting to replicate such results with our virtual birds might shed new light on the cognitive differences between corvids and parids, if any.

## References

- Anderson, J. R. (2005). Human symbol manipulation within an integrated cognitive architecture. *Cognitive Science*, 29, 313–341.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford: Oxford University Press.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. L. (2004). An integrated theory of the mind. *Psychological Review*, 111, 1036–1060.
- Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language*, 38, 341–380.
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2, 396–408.
- Baker, M. C., & Anderson, P. (1995). Once-pilfered cache sites not avoided by black-capped chickadees. *Animal Behaviour*, 49, 1599–1602.
- Baker, M. C., Stone, E., Miller Baker, A. E., Shelden, R. J., Skillicorn, P., & Mantuch, M. D. (1988). Evidence against observational learning in storage and recovery of seeds by black-capped chickadees. *Auk*, 105, 492–497.
- Balda, R. P., & Kamil, A. C. (1992). Long-term spatial memory in Clark's nutcrackers, *Nucifraga columbiana*. *Animal Behaviour*, 44, 761–769.
- Balda, R. P., Kamil, A. C., & Grim, K. (1986). Revisits to emptied cache sites by Clark's nutcrackers (*Nucifraga columbiana*). *Animal Behaviour*, 34, 1289–1298.
- Bednekoff, P. A., & Balda, R. P. (1996). Observational spatial memory in Clark's nutcrackers and Mexican jays. *Animal Behaviour*, 52, 833–839.
- Bednekoff, P. A., & Balda, R. P. (1997). Clark's nutcracker spatial memory: Many errors might not be due to forgetting. *Animal Behaviour*, 54, 691–698.
- Bednekoff, P. A., Balda, R. P., Kamil, A. C., & Hile, A. G. (1997). Long-term spatial memory in four seed-caching corvid species. *Animal Behaviour*, 53, 335–341.
- Brodin, A. (2005). Mechanisms of cache retrieval in long-term hoarding birds. *Journal of Ethology*, 23, 77–83.
- Brodin, A., & Kunz, H. (1997). An experimental study of cache recovery by hoarding willow tits after different retention intervals. *Behaviour*, 134, 881–890.
- Bugnyar, T., & Heinrich, B. (2005). Ravens, *Corvus corax*, differentiate between knowledgeable and ignorant competitors. *Proceedings of the Royal Society B*, 272, 1641–1646.
- Clayton, N. S., Dally, J. M., Gilbert, J., & Dickinson, A. (2005). Food caching by western scrub-jays (*Aphelocoma californica*) is sensitive to the conditions at recovery. *Journal of Experimental Psychology: Animal Behavior Processes*, 31, 115–124.
- Clayton, N. S., & Dickinson, A. (1999). Motivational control of caching behaviour in the scrub jay, *Aphelocoma coerulescens*. *Animal Behaviour*, 57, 435–444.
- Clayton, N. S., Yu, K. S., & Dickinson, A. (2001). Scrub jays (*Aphelocoma coerulescens*) form integrated memories of the multiple features of caching episodes. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 17–29.
- Clayton, N. S., Yu, K. S., & Dickinson, A. (2003). Interacting cache memories: Evidence for flexible memory use by western scrub jays. *Journal of Experimental Psychology: Animal Behavior Processes*, 29, 14–22.
- Dawson, M. R. W., Kelly, D. M., Spetch, M. L., & Dupuis, B. (2010). Using perceptrons to explore the reorientation task. *Cognition*, 114, 207–226.
- de Kort, S. R., Correia, S. P. C., Alexis, D. M., Dickinson, A., & Clayton, N. S. (2007). The control of food-caching behavior by western scrub-jays (*Aphelocoma californica*). *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 361–370.
- de Kort, S. R., Tebbich, S., Dally, J. M., Emery, N. J., & Clayton, N. S. (2006). The comparative cognition of caching. In E. A. Wasserman & T. R. Zentall (Eds.), *Comparative cognition: Experimental explanations of animal intelligence* (pp. 602–618). Oxford: Oxford University Press.
- Gould-Beierle, K., & Kamil, A. C. (1996). The use of local and global cues by Clark's nutcrackers, *Nucifraga columbiana*. *Animal Behaviour*, 52, 519–528.
- Grimm, V., & Railsback, S. F. (2005). *Individual-based modeling and ecology*. Princeton, NJ: Princeton University Press.
- Hampton, R. R., & Sherry, D. F. (1994). The effects of cache loss on choice of cache sites in black-capped chickadees. *Behavioral Ecology*, 5, 44–50.
- Heinrich, B., & Pepper, J. W. (1998). Influence of competitors on caching behaviour in the common raven, *Corvus corax*. *Animal Behaviour*, 56, 1083–1090.
- Hemelrijk, C. K. (1999). An individual-orientated model of the emergence of despotic and egalitarian societies. *Proceedings of the Royal Society B*, 266, 361–369.
- Hitchcock, C. L., & Sherry, D. F. (1990). Long-term memory for cache sites in the black-capped chickadee. *Animal Behaviour*, 40, 701–712.
- Hitchcock, C. L., & Sherry, D. F. (1995). Cache pilfering and its prevention in pairs of black-capped chickadees. *Journal of Avian Biology*, 26, 187–192.
- Kamil, A. C., & Balda, R. P. (1985). Cache recovery and spatial memory in Clark's nutcrackers (*Nucifraga columbiana*). *Journal of Experimental Psychology: Animal Behavior Processes*, 11, 95–111.
- Kamil, A. C., & Balda, R. P. (1990). Differential memory for different cache sites by Clark's nutcrackers (*Nucifraga columbiana*). *Journal of Experimental Psychology: Animal Behavior Processes*, 16, 162–168.
- Kamil, A. C., Balda, R. P., & Good, S. (1999). Patterns of movement and orientation during caching and recovery by Clark's nutcrackers, *Nucifraga columbiana*. *Animal Behaviour*, 57, 1327–1335.
- Kamil, A. C., Balda, R. P., Olson, D. J., & Good, S. (1993). Returns to emptied cache sites by Clark's nutcrackers, *Nucifraga columbiana*: A puzzle revisited. *Animal Behaviour*, 45, 241–252.
- Male, L. H., & Smulders, T. V. (2007). Memory for food caches: Not just for retrieval. *Behavioral Ecology*, 18, 456–459.
- Miller, N. Y., & Shettleworth, S. J. (2007). Learning about environmental geometry: An associative model. *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 191–212.
- Olton, D. S., & Samuelson, R. J. (1976). Remembrance of places passed:

- Spatial memory in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, 2, 97–116.
- Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, 31, 109–178.
- Raby, C. R., Alexis, D. M., Dickinson, A., & Clayton, N. S. (2007). Planning for the future by western scrub-jays. *Nature*, 445, 919–921.
- Salvucci, D. D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors*, 48, 362–380.
- Sun, R. (Ed.). (2008). *The Cambridge handbook of computational psychology*. Cambridge, NY: Cambridge University Press.
- Taatgen, N. A., van Rijn, R., & Anderson, J. R. (2007). An integrated theory of prospective time interval estimation. *Psychological Review*, 144, 577–598.
- van der Vaart, E., Hemelrijk, C. K., & Verbrugge, R. (2009). Learning where (not) to cache: A cognitive model for corvids. In N. A. Taatgen & R. van Rijn (Eds.), *Proceedings of the 31st Annual Meeting of the Cognitive Science Society* (pp. 2420–2425). Austin, TX: Cognitive Science Society.
- van der Vaart, E., Verbrugge, R., & Hemelrijk, C. K. (2009). Memory effects in the Clark's nutcracker: A cognitive model for corvids. In A. Howes, D. Peebles, & R. Cooper (Eds.), *Proceedings of the Ninth International Conference on Cognitive Modelling* (pp. 412–417). Manchester, UK.
- vander Wall, S. B. (1990). *Food hoarding in animals*. Chicago: The University of Chicago Press.
- vander Wall, S. B., & Balda, R. P. (1981). Ecology and evolution of food-storage behavior in conifer-seed caching corvids. *Zeitschrift fur Tierpsychologie-Journal of Comparative Ethology*, 56(3), 217–242.
- Watanabe, S., & Clayton, N. S. (2007). Observational visuospatial encoding of the cache locations of others by western scrub-jays (*Aphelocoma californica*). *Journal of Ethology*, 25, 271–279.

Received March 16, 2010

Revision received December 13, 2010

Accepted December 20, 2010 ■

### Online First Publication

APA-published journal articles are now available Online First in the PsycARTICLES database. Electronic versions of journal articles will be accessible prior to the print publication, expediting access to the latest peer-reviewed research.

All PsycARTICLES institutional customers, individual APA PsycNET® database package subscribers, and individual journal subscribers may now search these records as an added benefit. Online First Publication (OFP) records can be released within as little as 30 days of acceptance and transfer into production, and are marked to indicate the posting status, allowing researchers to quickly and easily discover the latest literature. OFP articles will be the version of record; the articles have gone through the full production cycle except for assignment to an issue and pagination. After a journal issue's print publication, OFP records will be replaced with the final published article to reflect the final status and bibliographic information.