

Memory in Clark's Nutcrackers: A Cognitive Model for Corvids

Elske van der Vaart (e.e.van.der.vaart@rug.nl)^{1,2}
Rineke Verbrugge (l.c.verbrugge@rug.nl)¹
Charlotte Hemelrijk (c.k.hemelrijk@rug.nl)²

¹Artificial Intelligence, University of Groningen, P.O. Box 407
9700 AK Groningen, the Netherlands

²Theoretical Biology, University of Groningen, P.O. Box 14
9750 AA Haren, the Netherlands

Abstract

Computational modeling has rarely been used to study questions in animal cognition, despite its apparent benefits. In this paper, we aim to demonstrate the value of this approach by focusing on work with Clark's nutcrackers. Like all corvids, these birds cache and recover food, by burying it under ground and returning to it later. With our computational model, we successfully replicate three laboratory experiments investigating this behavior. In the process, we provide the first integrated computational account of several behavioral effects of memory observed in corvid caching and recovery, in addition to a new explanation for a known empirical result.

Keywords: Computational model; animal cognition; corvid; Clark's nutcracker; *Nucifraga columbiana*; caching; memory.

Introduction

Computational models are a favored instrument in the cognitive science toolbox (Sun, 2008). Yet, there is an area of cognitive science where they are rarely used: That of animal cognition research (Penn, Holyoak, & Povinelli, 2008). Although many computational models are built to study other animals (Grimm & Railsback, 2005), the focus tends to be on ecological questions, such as 'what causes dominance hierarchies to form?' or 'how do individuals decide when to migrate?'. In contrast, cognitive questions, that concern animal memory, learning, or problem solving, are seldomly subjected to this approach. This despite the fact that computational models of animal cognition issues can be very useful (Penn et al., 2008).

What we are interested in, is the cognition underlying the caching and recovery behavior of corvids. This family of birds, which includes crows, jays, and nutcrackers, hides food under ground, saving it for later. Recovery can occur after hours, days, or months have passed, and depends on memory for individual cache sites. This behavior has been extensively researched in the laboratory, with a strong focus on its cognitive aspects (de Kort, Tebbich, Dally, Emery, & Clayton, 2006). All these cache and recovery experiments use the same basic paradigm: The birds are offered a bowl of food, a discrete set of sites to cache in, and the presence or absence of a conspecific, and very little else. Nevertheless, the questions asked and the data gathered are diverse, and may concern topics ranging from basic memory mechanisms to higher-level skills, such as future planning and social cognition (de Kort et al., 2006).

From a computational modeling perspective, this is excellent: It means that a single computational architecture of corvid cache and recovery cognition can be used to investigate a wide variety of cognitive phenomena.

In this paper, we present a step in that direction by focusing on three experiments with Clark's nutcrackers. These North American corvids are completely dependent on stored food in the winter months, and a single bird may bury up to 33,000 pine seeds a year, spread over thousands of different sites. Observational studies suggest that Clark's nutcrackers may recover their caches up to eleven months after making them, with a recovery accuracy of over 80%. One of the earliest laboratory experiments with these birds demonstrated the role of memory in this process: Like all corvids, a Clark's nutcracker cannot relocate caches by scent or by search, but only by remembering their location (see Kamil & Balda, 1985, for a review).

Since then, other laboratory experiments, in particular by Alan Kamil and Russell Balda, have investigated many more features of the Clark's nutcracker memory system, and it is three of these experiments that we replicate with our computational model (Balda, Kamil, & Grim, 1986; Kamil & Balda, 1990). All three have the same basic setup: The birds are tested in an experimental room, with 180 holes in the floor. These are spaced in a rectangular grid, 12 x 15 in size. Every hole can contain either a sand-filled cup, suitable for caching in, or a wooden plug, rendering it inaccessible. All subjects are always tested individually, and all sand-filled cups are smoothed over between sessions. Every experiment consists of a sequence of caching and recovery sessions. On caching sessions, the birds are offered a bowl of seeds to cache; on recovery sessions, the birds are hungry, and can only eat by recovering the seeds they have previously hidden in the experimental room.

From these three experiments, four patterns are apparent: A decrease in accuracy as recovery progresses, occasional return to already emptied sites, a lack of correlation between caching and recovery order, and a slight preference for re-caching in previously used cups. In this paper, we describe a computational model that successfully reproduces all four of these patterns. Its core component is *memory*, for cache and recovery events. To store these, we draw inspiration from the ACT-R (Anderson, 2007) cognitive architecture, and in particular, from its account of rational memory (Anderson & Schooler, 1991). What we use, is ACT-R's concept of

chunks: A chunk is a small piece of information, with an *activation* that depends on its own history of use, as well as that of related chunks. Essentially, what our model does, is to encode a bird's *options* for caching and recovery as chunks, and to compute their total activation based on the bird's *memory* of where it has cached and recovered before. With noise in the activation values of chunks, this mechanism is enough to replicate the outcomes of all three Clark's nutcracker experiments under consideration. In this way, we provide the first integrated computational account of different behavioral effects of memory in corvid caching and recovery, and a new explanation for the experimental finding that the recovery accuracy of Clark's nutcrackers declines across sessions. What further strengthens the validity of our model, is that we have extended it to replicate a second set of patterns, concerning cache site choice in the scrub jay, another corvid species (van der Vaart, Hemelrijk, & Verbrugge, to appear). Thus, the idea of constructing a single computational architecture of corvid cache and recovery cognition appears to be a fruitful one.

Model

Our implementation of the Clark's nutcracker experiments consists of two main components: A *simulator* and a *cognitive model*. The simulator runs the experiments, while the cognitive model is a computational theory of the cognitive processes under concern. Motivational processes that govern whether the birds want to cache or recover at all, are not considered; we simply assume that the birds want to cache in caching sessions and recover in recovery sessions.

The Basics of Chunks

Our model features two types of chunks: *Option* chunks and *memory* chunks. Option chunks represent the locations that are *available* for the bird to cache or recover in; memory chunks represent the *actual* cache or recovery events that the bird has experienced. Every chunk has two features: An *identifier* and an *activation*. A chunk's identifier specifies which cup it represents within the experimental room, as determined by its x and y location. A chunk's activation A_i consists of three parts: *Base-level activation* B_i , *spreading activation* S_i , and *noise*; see Equation 1.

A chunk's base-level activation B_i is computed according to Equation 2, following ACT-R's equation for base-level learning (Anderson, 2007). Here, t_j represents the elapsed time t since use j of chunk i , while d is a decay parameter. The weighing factor w_i is determined by chunk i 's type, and is considered in detail further on. The effect is that a chunk's base-level activation depends on its frequency and recency of use, and the kind of event it codes for. A chunk's spreading activation S_i depends on the activation of other chunks, and is discussed later. A chunk's noise value is re-computed every time it is evaluated, according to Equation 3, taken from ACT-R, where n is a parameter that we tune, and r is a random value between 0 and 1.

$$A_i = B_i + S_i + noise \quad (1)$$

$$B_i = w_i \cdot \sum_j t_j^{-d} \quad (2)$$

$$noise = n \cdot \ln\left(\frac{1-r}{r}\right) \quad (3)$$

$$S_{co_i} = -B_{cm_i} \quad (4)$$

$$S_{ro_i} = B_{cm_i} - B_{rm_i} \quad (5)$$

For the purpose of computing the activations of chunks, time is measured in steps. Every cache or recovery event counts as one step, and every non-experimental day counts as t steps, where t is a parameter that we tune. This simulates the flow of time *outside* of the experimental sessions.

The Structure of Caching and Recovery Sessions

At the beginning of every caching session, the simulator informs the cognitive model which cups are available to cache in. This is our equivalent of a Clark's nutcracker sitting on its perch, overseeing the room and registering its options. Then, every time the cognitive model starts to cache, it computes the activation of all its *cache option chunks*, according to Equation 1, and selects the most active one. This counts as a use of that chunk, and represents a bird's decision to cache in a particular cup. Once the cognitive model has selected its cache site, it caches there, and the corresponding *cache memory chunk* is given a use. Caching continues until the simulator asks the cognitive model to stop; this is determined by the number of caches made by the real birds in the original experiment.

A recovery session works in exactly the same way, except that it revolves around *recovery option chunks* and *recovery memory chunks*. The simulator ends a recovery session when the cognitive model has successfully retrieved as many caches as the real nutcrackers in the corresponding experiment are allowed to do.

The Memorability of Events

For the purpose of calculating a chunk's base-level activation B , according to Equation 2, cache memory chunks are given a weight w_{cm} of 5, while recovery memory chunks are given a weight w_{rm} of 2. This is inspired by the fact that Clark's nutcrackers probe a cup with their beaks about five times when making a cache, but only about twice when attempting to recover (Kamil, Balda, & Good, 1999). Option chunks, regardless of type, always carry a weight w_o of 1, representing the idea that *deciding* to cache or recover is less memorable than *actually* caching or recovering.

Inhibition of Return

To prevent the model birds from returning to recently visited sites, every memory chunk spreads *negative activation*, or *inhibition*, to the corresponding option chunk. See Equation 4 for cache chunks, and Equation 5 for recovery chunks. To work out the case of Equation 4:

The higher the base-level activation B of the cache memory chunk cm_i , the lower the spreading activation S of the cache option chunk co_i , and the smaller the odds that the model bird will return to that cache site.

Knowing Where to Recover

What allows the cognitive model to relocate its caches, is the fact that every cache memory chunk spreads *positive activation* to the recovery option chunk that codes for the same location; this is included in Equation 5. This has the effect that the cognitive model is *more* likely to try and recover in cups where it has actually cached items.

Experiments

To validate our cognitive model, we test it against three experiments with Clark’s nutcrackers: Experiment 1 from Kamil *et al.* (1986), and Experiments 1 and 2 from Balda & Kamil (1990). Here, we describe both the nutcracker experiments and our model’s replications of them. Model results are the average of 1000 runs, using the parameters of Table 1; see the Model section for an explanation of each.

Table 1: Parameter values used in the experiments.

d	n	f	w_{cm}	w_{rm}	w_o
0.1	2	10	5	2	1

Experiment 1 (Experiment 1 in Kamil *et al.* (1986))

In this experiment, the authors measure two aspects of the recovery behavior of Clark’s nutcrackers: Their decreasing accuracy as recovery continues, and their tendency to revisit already emptied sites. To this end, four birds are allowed to cache in the experimental room, until they store seeds in about twenty cups. Approximately ten days later, three recovery sessions are held, on alternate days. In each of these sessions, every bird may recover about a third of its caches.

Empirical Results, Kamil *et al.* (1986) To calculate results, recovery accuracy is defined as the total number of caches recovered divided by the total number of cups visited. As Figure 1A shows, the birds’ average accuracy declines significantly across recovery sessions, starting at about 55% in session 1 and ending at about 15% in session 3.




Figure 1: Results of Experiment 1, real birds, Kamil *et al.* (1986) and computational model; 1A: Average recovery accuracy, 1B: Average revisits to previously emptied sites, with standard errors.

Discussion In this experiment, the focus is on the repeat visits to previously emptied sites. Kamil *et al.* (1986) present two possible explanations for this: Either the birds remember their cache sites but not their recovery attempts, *or* they remember both, but continue to make revisits for some reason. For our model birds, the answer lies somewhere in the middle. When deciding where to recover, they follow the recovery option chunk that is currently most active. When determining the activation of a recovery option chunk, a corresponding cache memory chunk *raises* its activation, while a corresponding recovery memory chunk *lowers* it; see Equation 5. In this calculation, the uses of cache memory chunks are weighted five times, while the uses of recovery memory chunks are only weighted twice (Equation 2, Table 1). As stated previously, this is based on empirical observations of caching and recovery events (Kamil *et al.*, 1999). As a consequence of this, recovery option chunks representing already visited cache sites tend to be less active than recovery option chunks representing not yet visited cache sites, but they also tend to be *more* active than recovery option chunks representing sites where the model bird never cached at all. This is what causes the model birds to make revisits.

Experiment 2 (Experiment 1 in Kamil & Balda (1990))

Here, Kamil & Balda (1990) investigate *why* Clark’s nutcrackers become less accurate across recovery sessions. Given that these birds successfully locate their caches up to eleven months after making them, it seems unlikely that the two-day delay between recovery sessions is causing their accuracy to decline. Instead, the authors argue, what may be happening is that the birds remember some cache sites better than others, for whatever reason. Then, if they retrieve these ‘best remembered’ cache sites first, this explains why recovery accuracy drops. To test this idea, ten Clark’s nutcrackers are exposed to two experimental conditions: The *quarters* and the *free* condition.

In the quarters condition, the birds are forced to recover their caches by room quarter, while in the free condition, they can recover at will. Each condition consists of one caching session, followed by four recovery sessions. In both conditions, during the caching session, only 32 cups are available for caching, eight in every quarter of the room. The birds may store seeds until they have created at least three caches in every quarter. A week later, recovery sessions begin, conducted on successive days. This is where the conditions vary: In the quarters condition, only one quarter of the room is available for recovery during each session, while in the free condition, all cups are always open. In the quarters condition, the birds may continue to recover until they have retrieved all caches created in the available quarter; in the free condition, they are chased out of the experimental room after they have recovered 25% of their caches. Now the reasoning is that if the birds remember some cache sites better than others, their recovery accuracy should stay the same across recovery sessions in the quarters condition, but decline in the free condition.




Figure 2: Results of Experiment 2, average recovery accuracy, real birds, Kamil *et al.* (1990) and computational model; 2A: Control condition, 2B: Quarters condition.

Empirical Results, Balda & Kamil (1990) As can be seen in Figure 2, the results are as expected: Recovery accuracy stays the same across recovery sessions in the quarters condition, but declines significantly in the free condition, and more quickly than would be expected by chance. In this experiment revisits to sites already emptied in previous recovery sessions are not counted as errors, because they can only occur in the free condition; instead, these revisits are ignored when calculating accuracy. In further analysis, Balda & Kamil (1990) look for a general relationship between caching and recovery order by calculating Spearman's rank order correlations for the 10 birds. Three of these are significant, but two are positive and one is negative, suggesting that no general relationship exists.

Model Results As can be seen in Figure 2, the behavior of our model birds is similar to that of the Clark's nutcrackers: Accuracy does not decrease in the quarters condition, but does decrease in the free condition. Like the real birds, our model birds also show no systematic relationship between caching and recovery order in the free condition; of the 1000 correlations, only 3 are significant.

Discussion From these results, Kamil & Balda (1990) conclude that, in fact, Clark's nutcrackers remember some cache sites better than others. For our model birds, however, the explanation is different. In principle, they remember all cache sites equally, with the exception that the activation of cache memory chunks created earlier will have decayed more than the activation of cache memory chunks created later. However, if this were the explanation for the model birds' decline in recovery accuracy, we would expect to see no difference between conditions in this experiment. So what explains the model birds' constant performance in the quarters condition, but not the free condition?

The answer lies in the fact that, in the quarters condition, the number of caches that can be recovered remains the same across sessions, while in the free condition, it declines. When the cognitive model is deciding where to recover, it calculates the activations of all its recovery option chunks. On average, recovery option chunks representing cache sites are more active than recovery option chunks not representing cache sites, due to the spreading activation

coming from cache memory chunks. This is what allows the cognitive model to make accurate recovery attempts, *most* of the time. However, noise may cause an 'incorrect' recovery option chunk to temporarily be more active than all 'correct' recovery option chunks. The lower the ratio of 'correct' to 'incorrect' recovery option chunks, the higher the odds of this occurring. In the quarters condition, the ratio of 'correct' to 'incorrect' chunks remains the same across sessions, because a fresh quarter of cups is available every time. In the free condition, by contrast, the ratio of 'correct' to 'incorrect' chunks decreases across sessions, because the birds continue to recover from the same set of cups. This is what explains our model's performance.

Experiment 3 (Experiment 2 in Balda & Kamil (1990))

In this experiment, Balda & Kamil (1990) further explore the idea of differential memory for different cache sites. They hypothesize that perhaps certain cache sites have physical attributes that make them particularly memorable, such as their placement near certain kinds of landmarks. If this is true, the authors argue, it predicts that if the birds are forced to repeatedly cache in the same sites, they should always cache and recover from them in the same order. After all, if specific sites have physical attributes that make them particularly attractive, they should always be preferred.

To test this theory, seven nutcrackers are exposed to an experiment with three stages. Each stage consists of a caching and a recovery session, with a week between the two, and a week between stages. In stage 1, the birds may freely make 15 to 18 caches. In stage 2, for every subject, only the cups used as cache sites in stage 1 of the experiment are available for caching. In stage 3, this set of cups is again available, together with a second set of cups, that is randomly selected and of equal size. In both stages 2 and 3, the birds are allowed to cache in about nine cups. In all three recovery sessions, the birds can freely recover. Now, the main question is whether or not the birds will demonstrate site preferences by always caching and recovering in the same order, thus indicating site preferences.

Empirical Results, Balda & Kamil (1990) To analyze whether or not the birds prefer specific cache sites, Balda & Kamil (1990) calculate four Spearman's rank order correlations: Between caching and recovery in stage 1, between caching and recovery in stage 2, between caching in stage 1 and in stage 2, and between recovery in stage 1 and in stage 2. If a physical attribute is making some sites more memorable or more preferable, the birds should consistently choose to cache and recover in those sites first, producing significant correlations. Instead, the authors find no general relationships between caching and recovery orders; for all measures, they find a few significant correlations for some birds, but they go in both directions.

Another measure of interest is recovery accuracy *within* sessions. If some sites are more memorable than others, recovery accuracy within sessions should decrease, as the better remembered sites are recovered first.




Figure 3: Results of Experiment 3, real birds, Kamil *et al.* (1990) and computational model; 3A: Average recovery accuracy, 3B: Average reuse of cache sites in stage 3, with standard errors, birds on the right, model on the left.

For each stage of the experiment, Kamil & Balda (1990) calculate the mean accuracy of the first three caches recovered, the middle three caches recovered, and the final three caches recovered. They find no significant differences *between* the three stages, but they do find a significant decrease in accuracy *within* the three stages; therefore, Figure 3A plots pooled accuracy results.

A final measure of cache site preference is the proportion of cache sites re-used in stage 3 of the experiment. If the birds choose to cache in particularly preferable sites in stage 1, we might expect them to still strongly prefer those sites in stage 3. Yet, this is not the case: As can be seen in Figure 3B, when offered their chosen set of sites from stage 1, and an equally sized set of new cups to cache in, they choose to cache in old cups only about sixty percent of the time.

Model Results Like the real Clark’s nutcrackers, our model birds produce no significant correlations between cache and recovery orders on any of the measures tested by Kamil & Balda (1990). Furthermore, we also find a decrease in accuracy within recovery sessions, in all three stages of the experiment. Pooled accuracy results are plotted in Figure 3A. Finally, as can be seen in Figure 3B, in stage 3, our model birds re-use cache sites at approximately the same levels as the real nutcrackers: They choose to place about 60% of their caches in old sites, and 40% in new sites.

Discussion From this experiment, Kamil & Balda (1990) conclude that Clark’s nutcrackers clearly do not have strong site preferences as dependent on *physical* attributes, or they would have consistently preferred to cache and recover from the same sites first. Our model birds show qualitatively the same patterns. The slight preference for old cache sites in stage 3 of the experiment can be explained by the fact that the cache option chunks corresponding to the cups that had already been chosen in stage 1 of the experiment already had two uses by this point, while the ‘new’ cups had none. As cache memory chunks only spread negative activation to cache option chunks *within* sessions, this means that the average activation of already-used cache option chunks is slightly higher than that of not-yet-used cache option chunks, explaining the model birds’ behavior.

General Discussion

Our computational model raises three main questions: First, what does it tell us about Clark’s nutcrackers? Second, how robust are its results? And third, how plausible is its design?

Implications of the Model for Clark’s Nutcrackers

One of the attractive aspects of our model is that it uses one main mechanism, but fits four different patterns. We assume that both a birds’ options and its choices are stored as chunks in memory, and that spreading activation between different chunk types takes care of the rest. This produces all four patterns apparent in the empirical data: A decline in accuracy both within and between recovery sessions, occasional return to already emptied sites, a lack of correlation between caching and recovery order, and a slight preference for re-caching in previously used cups.

In addition, we provide a new explanation for an observed result: The decline in accuracy as recovery proceeds. In a number of different papers, Kamil & Balda (1986; 1990) conclude that this is the result of differential memory for different cache sites, but in our cognitive model, the same effect arises as the result of chance. This seems to be a useful alternative theory, because the attempt to discover what might make certain sites more memorable than others has so far not been successful: As demonstrated by Kamil & Balda (1990) in the original version of our Experiment 3, Clark’s nutcrackers do not consistently prefer some sites, suggesting that physical characteristics of particular locations cannot be responsible for different memorability. The birds’ familiarity with particular cache sites is also an unlikely explanation, as Kamil, Balda & Good (1999) fail to find any predictors of recovery accuracy in the amount of time the birds spend making each cache.

However, several aspects of the model still need further work. One feature that seems particularly over-simplified is that all the cognitive model’s errors are “true errors” – failures to retrieve a correct cache site location. However, for the real Clark’s nutcrackers, it is probable that many errors are in fact acts of exploration. For instance, when the costs of making a recovery attempt are increased, the number of errors made drops significantly (Bednekoff & Balda, 1997). This is clearly an aspect of Clark’s nutcracker behavior that we should explicitly consider in future.

Robustness of the Model’s Results

When a computational model features free parameters, it is important to understand how strongly it *predicts* certain outcomes, and if there are any plausible alternatives that it *cannot* reproduce (Roberts & Pashler, 2000). For our cognitive model, this is certainly the case: As soon as *noise* is set higher than 0, in this type of experiment, it cannot fit anything but a decline in accuracy as recovery progresses (Experiment 2); constant performance is impossible, even if decay is set to 0. That constant performance is a plausible alternative, is demonstrated by Kamil & Balda’s (1985) original theory that this was true of Clark’s nutcrackers.

Plausibility of the Model's Implementation

Many of our model's core aspects are derived from the declarative memory module of the ACT-R (Anderson, 2007) cognitive architecture, lending it some initial validity. Of course, ACT-R was explicitly designed to model humans, so one might wonder whether our use of it for birds is appropriate. However, we do not think that is a problem in this case; as ACT-R's originator John Anderson himself notes (2007; page 18) many of the adaptive analyses on which ACT-R is based are not species-specific. There are many indications that, at a functional level, corvid memory may not be so different from ours; several experiments with scrub jays show that they are capable of flexibly integrating their memories, and have episodic-like "what, where, when" recall of past events (see de Kort *et al.*, 2006, for a review.)

Our adaptations of the architecture itself might be more problematic. While ACT-R allows for spreading activation, it is a fixed amount, and it spreads only from 'goal chunks' to 'target chunks', depending on the strength of the association between them. In our cognitive model, it is a chunk's *own* activation that spreads, and this activation can even be negative, inhibiting a chunk's retrieval. The main function of this mechanism is to prevent the model birds from repeatedly caching or recovering in the same location. One might wonder if such a mechanism is necessary at all; if, instead, the real birds might be using a behavioral strategy to avoid revisits, such as 'recover a cache, look away, attempt to recover a cache in the field of view now visible'. This, however, does not appear to be a likely explanation; after successfully retrieving seeds, Clark's nutcrackers fly back to a central perch to eat them (Kamil & Balda, 1985). This means that, when they are deciding where to recover next, a very large portion of the experimental room is visible to them. This makes it very difficult to think of a behavioral strategy that avoids revisits to the extent that the real Clark's nutcrackers do.

Of course, this does not imply that our technique of spreading negative activation is necessarily the best way of implementing an inhibition of return mechanism. It is possible that the same effect could be achieved by instead *increasing* the activations of all other chunks, but we believe our solution is computationally easier, and intuitively plausible. Interestingly, other recent ACT-R adaptations also make similar changes to the architecture: Van Maanen & van Rijn (2007) let activation spread between chunks of different types, and Juvina & Taatgen (2009) attach negative activations to chunks. Although the context and justification is different, this negative activation mechanism serves the same function as ours – suppression of repetition – and operates in a similar fashion, with inhibitory activation that decays over time.

Conclusions

In this paper, we have shown that our computational model of corvid cache and recovery cognition can successfully reproduce the outcomes of three experiments with Clark's nutcrackers, fitting four different patterns.

In addition, our computational model has provided a new explanation for the fact that Clark's nutcrackers become less accurate as recovery progresses.

Acknowledgments

We thank Alan Kamil and Russell Balda for further comments on their results, and NWO for funding Elske van der Vaart's doctorate, TopTalent grant 021.001.089.

References

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford: Oxford University Press.
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2, 396-408.
- 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. (1997). Clark's nutcracker spatial memory: Many errors might not be due to forgetting. *Animal Behaviour*, 54, 691-698.
- 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*. Oxford: Oxford University Press.
- Grimm, V., & Railsback, S. F. (2005). *Individual-based modeling and ecology*. Princeton, NJ: Princeton University Press.
- Juvina, I., & Taatgen, N. A. (2009). A repetition-suppression account of between-trial effects in a modified stroop paradigm. *Acta Psychologica*, 131, 72-84.
- 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(1), 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(2), 162-168.
- Kamil, A. C., Balda, R. P., & Good, S. (1999). Patterns of movement and orientation during caching and recovery by Clark's nutcrackers. *Animal Behaviour*, 57, 1327-1335.
- 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.
- Sun, R. (Ed.). (2008). *The Cambridge handbook of computational psychology*. New York: Cambridge University Press.
- van der Vaart, E., Hemelrijk, C. K., & Verbrugge, R. (to appear). Learning Where (Not) To Cache: A Cognitive Model for Corvids. In N.A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Meeting of the Cognitive Science Society*. Chicago: Lawrence Erlbaum Associates.
- van Maanen, L., & van Rijn, H. (2007). An accumulator model of semantic interference. *Cognitive Systems Research*, 8(3), 174-181.