

Mini-review

# Transitive inference in non-human animals: An empirical and theoretical analysis

Marco Vasconcelos\*

*Department of Psychological Sciences, Purdue University, 703 Third Street, West Lafayette, IN 47907-2081, United States*

Received 12 October 2007; received in revised form 22 February 2008; accepted 27 February 2008

## Abstract

Transitive inference has long been considered one of the hallmarks of human deductive reasoning. Recent reports of transitive-like behaviors in non-human animals have prompted a flourishing empirical and theoretical search for the mechanism(s) that may mediate this ability in non-humans. In this paper, I begin by describing the transitive inference tasks customarily used with non-human animals and then review the empirical findings. Transitive inference has been demonstrated in a wide variety of species, and the signature effects that usually accompany transitive inference in humans (the serial position effect and the symbolic distance effect) have also been found in non-humans. I then critically analyze the most prominent models of this ability in non-human animals. Some models are cognitive, proposing for instance that animals use the rules of formal logic or form mental representations of the premises to solve the task, others are based on associative mechanisms such as value transfer and reinforcement and non-reinforcement. Overall, I argue that the reinforcement-based models are in a much better empirical and theoretical position. Hence, transitive inference in non-human animals should be considered a property of reinforcement history rather than of inferential processes. I finalize by shedding some light on some promising lines of research.

© 2008 Elsevier B.V. All rights reserved.

*Keywords:* Transitive inference; Symbolic distance effect; Serial position effect; Cognitive models; Value transfer theory; Reinforcement-based models

## 1. Introduction

For thousands of years, the transitive inference (TI) problem has been one of the hallmarks of human logical deductive reasoning. This task, known to the ancient Greeks, and actually one of the milestones of Aristotelian logic, was first introduced into modern psychology by Burt (1911, 1919a,b) in order to evaluate the reasoning abilities of young children. Similarly, Piaget (1928) became interested in syllogisms as a means to understand the development of basic cognitive abilities.

As an illustration, suppose you are given the following propositions: “David is taller than John” and “John is taller than Peter”. Afterwards, you are asked “Who is taller, David or Peter?” The answer to this question is quite obvious despite the fact that this information was not directly presented. This simple task is a verbal instantiation of the TI task and the verbal report “David” (in the case above) is evidence for the TI ability. Although most human adults readily solve such syllogisms, the ability

of children and non-human animals to solve this kind of problem was, and in some instances still is, questioned. For example, in Piaget’s view (1928, 1955, 1970), the TI competence emerges only with the concrete operational stage, which presumably precludes children younger than about 7 years old from a correct resolution of these tasks. Implicit in this position is the assumption that only after the acquisition of the rules of logic will a child correctly solve such syllogisms.

Formally, TI can be defined as a form of reasoning in which, given preliminary information (the premises), the subject deduces a logical conclusion. In general, a relation  $R$  between any two objects “ $a$ ” and “ $c$ ”,  $aRc$ , is said to be transitive if, given  $aRb$  and  $bRc$ , it follows that  $aRc$ . In a word, the ordinal relation between two elements must be inferred from premises that establish the relations of those two elements to a third (e.g., Halford, 1984). Given this definition, it should come as no surprise that TI has long been regarded as a hallmark of human cognitive abilities. Its apparent reliance on the rules of formal logic and its usual verbal implementation helped foster this somewhat anthropomorphic position.

It was in this long-established environment in which reasoning was seen as a general ability to correctly manipulate

\* Tel.: +1 765 494 6881; fax: +1 765 496 1264.

E-mail address: [marcov@psych.purdue.edu](mailto:marcov@psych.purdue.edu).

propositions according to the rules of logic that the question of whether non-human animals are able to correctly solve TI tasks was first raised. Needless to say, the question itself was at odds with the dominant theoretical standards. This scenario, however, was dramatically changed by a seminal paper by Bryant and Trabasso (1971). Using a semi-verbal instantiation of the TI task, they showed that failures of TI performance in young children reflected deficits of memory rather than deficits of logic or of any heuristic. Children were trained with pairs of rods of varying lengths and color, although only color was informative (the same length was visible for all rods). On each trial, children were presented with a pair of rods: *A* and *B*, *B* and *C*, *C* and *D*, or *D* and *E* and were asked which one was bigger or smaller. For example, red might be bigger than blue, blue might be bigger than green, and so forth. After learning the premises on a trial-and-error basis, children faced critical tests in which *B* and *D* were presented. The finding was that when children were tested only after they had clearly memorized the premises, they were indeed capable of TI, including the youngest, 4-year-olds (see also Riley and Trabasso, 1974). The ability to solve such syllogisms previous to the development of the cognitive abilities thought to sustain such performance strongly suggested that TI might be based on simpler mechanisms than previously acknowledged. Since then, the logical account fell out of favor and most research seems to support a mental model approach in humans (e.g., Byrne and Johnson-Laird, 1989; Carreiras and Santamaria, 1997; for details, see Section 4.1. below). Importantly, the door for the provocative hypothesis that non-human animals might be capable of TI was opened.

Since Bryant and Trabasso's (1971) paper, evidence for TI in a variety of species has steadily accumulated. The first to demonstrate TI in non-humans were McGonigle and Chalmers (1977) with squirrel monkeys, but after their demonstration many others followed: Boysen et al. (1993) and Gillan (1981) with chimpanzees; Buckmaster et al. (2004), McGonigle and Chalmers (1977, 1992), Rapp et al. (1996), and Treichler and Van Tilburg (1996) with monkeys; Davis (1992b), Dusek and Eichenbaum (1997), Roberts and Phelps (1994), and Van Elzaker et al. (2003) with rats; Lazareva and Wasserman (2006), Siemann et al. (1996b), Steirn et al. (1995), von Fersen et al. (1991), Weaver et al. (1997), and Wynne (1997) with pigeons; Bond et al. (2003), and Paz-y-Miño et al. (2004) with pinyon jays; Bond et al. (2003) with scrub-jays; Lazareva et al. (2004) with hooded crows; and Grosenick et al. (2007) with fish (*Astatotilapia burtoni*).

The ability of non-human animals to solve transitive syllogisms may have an important adaptive value (e.g., Wynne, 1995). Rank estimation in social animals provide a good example (e.g., Cheney and Seyfarth, 1986, 1990). In species in which dominance is not evident from physical traits, an animal incapable of TI would face a combinatorial explosion from having to engage in potentially dangerous interactions with every single member of the group. For example, if three items have to be ranked, the three possible pairs have to be experienced; if six items have to be ranked, the 15 possible pairs have to be experienced, and so forth.

Suppose a new member joins a social group. One way to estimate rank is for each member of the group to interact with

the newcomer. A more economical strategy would be to observe the newcomer's interactions with some members of the group and, based on the previous knowledge of the existing members' rank, estimate the relative rank of the newcomer (von Fersen et al., 1991). Field observations suggest this is indeed the case (e.g., Altmann, 1962) and laboratory studies seem to corroborate such an assertion. For example, Paz-y-Miño et al. (2004) showed such an effect with pinyon jays (for examples with fish and hens, see Grosenick et al., 2007; and Hogue et al., 1996, respectively). Furthermore, the less social but closely related western scrub jays exhibit less transitive behavior (Bond et al., 2003). Whether or not this ability depends on the complexity of the animals' social system is still uncertain, despite suppositions that such complexity might have driven the evolution of several cognitive abilities (e.g., Jolly, 1966; Kummer et al., 1997; Paz-y-Miño et al., 2004; Shettleworth, 2004).

In this paper, I review and analyze TI data in non-human animals, focusing on the proximal mechanisms proposed to explain this ability. I start by describing the tasks customarily used in TI studies with non-human animals, followed by the typical empirical findings. I analyze the results obtained with these tasks, focusing not only on the TI test itself but also on some interesting effects observed during training. Next, the models that have been proposed to explain TI in non-human animals are presented and analyzed. Of particular interest is the fit between the predictions of each model and the empirical findings, but careful attention will also be given to the models' theoretical robustness. I end by shedding some light on what seem to be promising lines of research.

## 2. The transitive inference tasks

Most research on TI in non-human animals has used the so-called *n*-term series task. This task consists of the presentation of successive pairs of stimuli. Originally developed for children by Bryant and Trabasso (1971), it was modified into a fully non-verbal version by McGonigle and Chalmers (1977) in a well-known study with squirrel monkeys. In its simplest form, the task involves five different stimuli. Fig. 1 presents a schematic of the task.

Each pair represents a simultaneous discrimination in which one of the stimuli is reinforced (+) and the other is not (–).




Fig. 1. Schematic of the *n*-term series task. Each pair represents a simultaneous discrimination; (+) reinforced choice; (–) non-reinforced choice. Left panel: Simultaneous discriminations presented during training. Right panel: Simultaneous discrimination presented during testing.

Depending on species, the stimuli are typically either visual or olfactory. They are also arbitrarily assigned, and customarily labeled from *A* through *E* for ease of presentation. In this example, five stimuli are involved, so this is a 5-term series task. Note, also, that the stimuli for each discrimination partially overlap, and that the stimuli *B*, *C*, and *D* are reinforced when presented in one discrimination and non-reinforced when presented in the other. Supposedly, this training creates the linear order  $A > B > C > D > E$ . In testing, *B* vs. *D* is presented.<sup>1</sup> This discrimination was never presented during training. Given that both *B* and *D* were reinforced in one discrimination and non-reinforced in another (viz., both were partially reinforced), choice of stimulus *B* suggests the animals have performed a TI-like operation.

An important procedural restriction is that at least five stimuli must be used. Only in this way will the test pair be composed of non-adjacent stimuli and not include one of the end-anchor stimuli (*A* and *E*). In the example above, any test composed of *A* vs. *C*, *A* vs. *D*, *A* vs. *E*, *C* vs. *E*, or *B* vs. *E* is uninformative because each of these discriminations contains a stimulus that was either always reinforced (*A*) or always non-reinforced (*E*). In any such test, *A* would be preferred and *E* would be avoided because of their direct reinforcement history. Naturally, increasing the series length allows us to test more than one pair of stimuli. If, for example, six stimuli were used (a 6-term series) three different test pairs would be possible: *B* vs. *D*, *B* vs. *E*, and *C* vs. *E*.

This relatively simple task has been implemented in a variety of different ways. Most research has used what I call *intermixed* training in which all discriminations are trained concurrently (e.g., Boysen et al., 1993; Buckmaster et al., 2004; Dusek and Eichenbaum, 1997; Gillan, 1981; Higa and Staddon, 1993; Lazareva et al., 2004; Lazareva and Wasserman, 2006; McGonigle and Chalmers, 1977, 1992; Rapp et al., 1996; Roberts and Phelps, 1994; Siemann et al., 1996b; Treichler and Van Tilburg, 1996; Van Elzaker et al., 2003; von Fersen et al., 1991; Wynne, 1997). Alternatively, some studies have used *sequential* training in which each discrimination is trained by itself to criterion, usually 90% correct or above (Benard and Giurfa, 2004; Davis, 1992b; Steirn et al., 1995; Weaver et al., 1997). These two procedures typically require quite different amounts of training: With intermixed training more extended training is needed, and is not uncommon for animals to be dropped from these studies due to persistent difficulties during acquisition (whether or not transitive performances following these two different training regimens emerge from the same underlying mechanism(s), I leave to a forthcoming section). Actually, most studies using intermixed training include a phase in which animals are initially exposed to sequential training in order to speed-up training. In what follows, studies are classified as involving intermixed training when all premises were concurrently trained to a high level of accuracy in the phase immediately before testing (i.e., independently of

whether or not premises were sequentially trained in preceding phases).

Intermixed and sequential training entail different testing procedures, too. Typically, with intermixed training, infrequent probe (test) trials are interspersed among regular training trials and testing is usually conducted over several sessions. Sequential training, on the other hand, “forces” test sessions to include only test trials because the animal never experienced more than one trial type per session during training. Furthermore, because the test trials are non-differentially reinforced following sequential training (in order not to induce a preference during test itself) only a small number (20–30) of trials are included in preference analyses.

Despite these testing drawbacks, sequential training permits another interesting independent variable to be manipulated: The direction of the series. Training can proceed in a *forward* (i.e.,  $A+B-$ ,  $B+C-$ ,  $C+D-$ ,  $D+E-$ ) or *backward* direction (i.e.,  $A-B+$ ,  $B-C+$ ,  $C-D+$ ,  $D-E+$ ), allegedly resulting in opposite linear orders.

Another training variation is whether or not feedback other than reinforcement vs. non-reinforcement is provided after each choice (see Lazareva et al., 2004; Lazareva and Wasserman, 2006). When this additional feedback is used, each stimulus (from *A* to *E*, when a 5-term series is used) is associated with some post-choice feedback stimulus of progressively smaller or bigger size (for example, five circles of different diameters might be used). After a choice has been made on a particular trial (e.g., *A* vs. *B*), the corresponding pair of post-choice feedback stimuli is presented for a fixed amount of time. For instance, a big circle might be presented at the location where *A* was presented and a somewhat smaller circle at the location where *B* was presented and similarly for the other premise pairs. This procedure was developed to test a particular theoretical approach (to which I will return later) but it also successfully avoids an important criticism of the *n*-term series task—the trained relations are *not* transitive.

To illustrate the argument, suppose you were given the following premises: “John likes Jessica” and “Jessica likes David”. Does it follow that John likes David? Obviously not. Consider now a simple 3-term series task as it would be used with animals: “*A* is reinforced, *B* is not”, “*B* is reinforced, *C* is not”. Does it follow that *A* is reinforced and *C* is not if they were to appear together? The core question here is the nature of the relationship being trained. In order to be transitive, a relation must have an underlying scale, and reinforcement does not possess a scale (for a detailed account, see Markovits and Dumas, 1992). Nonetheless, the foregoing argument is somewhat silenced if we carefully analyze the meaning of a transitive-like conclusion in such a task. Suppose, for the sake of argument, that the subject *does* conclude that “John likes David” or that “*A* is reinforced and *C* is not”. What does this imply? One possibility is that the subject did not detect the non-transitive nature of the relation and solved the task using the same process(es) that underlie TI when the relation is in fact transitive (Wynne et al., 1992). In other words, the inference may be incorrect but it is *transitive*. If this is the case, the criticism, although pertinent, is innocuous.

<sup>1</sup> In some rare instances, researchers have used triadic tests by presenting, for example, *B* vs. *C* vs. *D*.

Another frequent criticism is that because the trained relations are not ordered along a particular scale, the  $n$ -term series task is more susceptible to associative, non-logical solutions (e.g., Russell et al., 1996). The argument is that TI is an instantiation of logical processes and therefore should be studied only when such processes are possible. I argue for the opposite course: TI should be accepted on the grounds of the observed performance and only then should the underlying processes be hypothesized. Logic is only one of the possibilities and perhaps not the strongest one. Even if logical reasoning is one underlying process, the possibility remains that such logical underpinnings can emerge in the  $n$ -term series task (e.g., when the non-transitive nature of the relation is not detected).

Albeit rare, some studies have used tasks other than the  $n$ -term series to study TI. One example, by Paz-y-Miño et al. (2004), involved two groups of wild-caught pinyon jays. After dominance relationships had been established between the previously unknown jays, individuals from each group observed two jays, one from its own group and one from the other group, interact over a peanut. The critical test involved a social interaction between the observer and a member of the other group that was previously observed being dominant over a dominant member of the observer's group. The prediction was that if jays infer social status transitively, then the observer should behave submissively in the critical encounter. In fact, they did. Hogue et al. (1996) and Grosenick et al. (2007) used a similar task with hens and fish, respectively, and found similar results. Although a possible alternative to the  $n$ -term series task, this task is useful only with species in which rank cannot be inferred from physical traits.

### 3. Evidence for transitive inference in non-human animals

Most studies on TI in non-human animals have yielded a relatively consistent pattern of findings within and across species. I start by characterizing the evidence for TI (performance at test) and then analyze training performance. I end this section by summarizing studies using what have been termed *circular* series. This will, in turn, allow a clearer understanding of the proposed explanatory models that heavily rely on training features.

#### 3.1. Test performance: the transitive choice

While human studies customarily use accuracy and latency data, animal studies have focused almost exclusively on accuracy data—D'Amato and Colombo (1990), McGonigle and Chalmers (1992), Rapp et al. (1996), and Wynne (1997) are the exceptions. Table 1 summarizes the results obtained with intermixed 5-term series tasks. As can be seen, a vast majority of the experiments report strong transitive preferences. The lowest average preference reported in a large set of experiments, including Boysen et al. (1993, Experiments 1, 3, and 4), Buckmaster et al. (2004, control group only), Dusek and Eichenbaum (1997, control group only), Lazareva et al. (2004, Experiment 1, ordered feedback group), Lazareva and Wasserman (2006), McGonigle and Chalmers (1977; 1992, Experiment 1), Rapp et al. (1996, aged subjects), Siemann et al. (1996b), Treichler and Van Tilburg

(1996, Experiment 1, Test 1), Van Elzaker et al. (2003, Experiment 1), and von Fersen et al. (1991, Experiment 1), is 80% despite procedural differences across studies, some of which are imposed by the species studied. For example, studies with rats usually employ olfactory stimuli, whereas studies with birds and primates typically use visual stimuli. If the ability to infer non-trained relations from potentially informative and partially overlapping premises is relatively general within the animal kingdom, then transitive behavior should be observed independently of sensory modality per se, provided the animal is perceptually sensitive to the stimuli being used. The data in Table 1 certainly suggest this is the case.

The reinforcement contingencies on test trials vary considerably, too. In some experiments, these trials were always reinforced or reinforced 50% of the time independently of the animal's choice (Boysen et al., 1993, Experiment 1; Gillan, 1981, Experiment 1A; Lazareva et al., 2004, Experiment 1; McGonigle and Chalmers, 1977, 1992; Rapp et al., 1996); in others, only the transitive choice was reinforced (differential reinforcement; Boysen et al., 1993, Experiments 2, 3, and 4; Buckmaster et al., 2004; D'Amato and Colombo, 1990; Dusek and Eichenbaum, 1997; Treichler and Van Tilburg, 1996, Experiment 1; Van Elzaker et al., 2003, Experiment 1); and, finally, in others, they were non-reinforced (viz., extinction; Higa and Staddon, 1993; Siemann et al., 1996b; von Fersen et al., 1991; Wynne, 1997). Considering the results, the test-trial reinforcement contingency does not seem to have a noticeable impact on performance. This strengthens the contention that transitive performance emerges from the training regimen and is not simply an artifact of testing itself.

Other procedural differences are noteworthy. For example, Lazareva and associates used the previously mentioned post-choice feedback procedure in which one group was exposed to informative, ordered feedback and another was exposed to constant feedback (Lazareva et al., 2004; Lazareva and Wasserman, 2006). Informative feedback should, if anything, foster the transitive nature of the relation being trained. Hence, several predictions are possible depending on what features of the situation the animal is sensitive to. If animals detect the non-transitive nature of the relation being trained, then the constant feedback group should not show TI. Conversely, to the extent that ordered feedback transforms this relation into a transitive one, the ordered feedback group should behave transitively. However, if animals do not detect the non-transitive nature of the training relations, then both groups should behave transitively on test trials (or, in a more pessimistic scenario, neither of them). If this argument holds, under no circumstance should the constant feedback group exhibit TI in the absence of TI in the ordered feedback group.

Presently, the evidence seems to support the argument. Lazareva et al. (2004) found evidence that hooded crows exhibit TI only when ordered feedback is provided (viz., by their interpretation, only when the trained relations are truly transitive), but, curiously enough, pigeons show TI independently of the transitive nature of the relations (viz., both when ordered or constant feedback are provided; Lazareva and Wasserman, 2006). The reason(s) for this difference is (are) still unknown. Nonethe-

Table 1  
Experiments using the 5-term series task with intermixed training

Source	Subjects	N <sup>#</sup>	Test contingency	%Transitive choice <sup>*</sup>
Boysen et al. (1993), Experiment 1	Chimpanzees	3 (3)	All choices reinforced	83, 100, 100
Boysen et al. (1993), Experiment 2	Chimpanzees	3 (3)	Differential reinforcement	86, 56, 64
Boysen et al. (1993), Experiment 3	Chimpanzees	1 (1)	Differential reinforcement	93
Boysen et al. (1993), Experiment 4	Chimpanzees	3 (3)	Differential reinforcement	100
Buckmaster et al. (2004)	Cynomolgus monkeys	6 (5)	Differential reinforcement	100 <sup>a</sup>
D'Amato and Colombo (1990), Experiment 2	Capuchin monkeys	3 (3)	Differential reinforcement	? <sup>b</sup>
Dusek and Eichenbaum (1997)	Rats	10 (8)	Differential reinforcement	81 <sup>a</sup>
Gillan (1981), Experiment 1A	Chimpanzees	3 (3)	Non-differential reinforcement, 50%	100, 58.3, 41.7 (83.3 <sup>c</sup> )
Higa and Staddon (1993), Experiment 1	Pigeons	4 (4)	Extinction	(34, 30, 59, 54) <sup>d</sup>
Higa and Staddon (1993), Experiment 2	Pigeons	4 (4)	Extinction	(47, 51, 57, 64) <sup>d</sup> (60, 67, 30, 83) <sup>e</sup>
Higa and Staddon (1993), Experiment 3	Pigeons	4 (4)	Extinction	(93, 80, 85, 56) <sup>d</sup> (77, 83, 22, 39) <sup>e</sup>
Lazareva et al. (2004), Experiment 1	Crows	8 (8)	Non-differential reinforcement, 50%	83.1 <sup>f</sup> , 43.1 <sup>g</sup>
Lazareva and Wasserman (2006)	Pigeons	12 (12)	All choices reinforced and extinction	82.1
McGonigle and Chalmers (1977), Experiment 1	Squirrel monkeys	8 (7)	All choices reinforced	90
McGonigle and Chalmers (1992), Experiment 1	Squirrel monkeys	6 (5)	All choices reinforced	98
Rapp et al. (1996), Phases 1–5	Rhesus monkeys	17 (16)	Non-differential reinforcement, 50%	79.1 (young subjects) 83.3 (aged subjects)
Siemann et al. (1996b)	Pigeons	10 (7)	Extinction	90.5, 94 <sup>c</sup>
Treichler and Van Tilburg (1996), Experiment 1, Test 1	Macaque monkeys	6 (6)	Differential reinforcement	94
Van Elzakker et al. (2003), Experiment 1	Rats	8 (6)	Differential reinforcement	87.5
von Fersen et al. (1991), Experiment 1	Pigeons	6 (4)	Extinction	87.5
Wynne (1997)	Pigeons	8 (3)	Extinction	95, 70, 55

<sup>#</sup> Number of subjects in the experiment (inside parenthesis is the actual number of subjects tested or included in statistical analyses).

<sup>\*</sup> When variability in results justifies, and the results are available, individual preferences are shown. When numerical values are not available, values are estimated from figures.

<sup>a</sup> Only the control group is considered. In the other group(s), the hippocampus was disconnected from either its cortical or subcortical pathways.

<sup>b</sup> Only latency data are reported.

<sup>c</sup> After retraining of the premises.

<sup>d</sup> Auto-run procedure with response-based V values.

<sup>e</sup> Auto-run procedure with time-based V values.

<sup>f</sup> Group with bias reversal training and ordered post-choice feedback.

<sup>g</sup> Group with bias reversal training and constant post-choice feedback.

less, as Lazareva and Wasserman argue, because of their relative linear social hierarchy (see Ritchey, 1951), pigeons may be less prone to distinguish between transitive and non-transitive relations than crows, whose social hierarchy is more complex. In the absence of further data, this appeal to between-species differences arising from different ecological demands is purely speculative. Nevertheless, the pattern of results observed in hooded crows and pigeons does not violate the above-mentioned predictions. Either both groups or only the ordered feedback group behave transitively.

Of the 21 experiments shown in Table 1, 14 report strong evidence for TI in non-human animals. Of the remaining seven, three also report some evidence for TI but the results clearly vary across subjects (in particular, Boysen et al., 1993, Experiment 2; Gillan, 1981, Experiment 1A, before retraining; Wynne, 1997). At least one of these cases seems to be related to the deterioration of training performance (see Gillan, 1981, Experiment 1A), given that after premise retraining the animal did prefer the transitive alternative.

With the exception of one experiment, Higa and Staddon (1993) repeatedly failed to observe TI. Their studies used, however, a very different training procedure in which pigeons were

reinforced after pecking the correct stimulus (*S*+) for a variable interval of time [variable interval 6.5 s (VI 6.5 s)] and the effect of pecking the incorrect stimulus (*S*−) was simply to reset the VI timer (i.e., to postpone reinforcement). In addition, these authors used an *auto-run* procedure to produce rapid acquisition and to control for premise presentation order during training. Instead of pseudo-randomly selecting the premise to present on a given trial, the *auto-run* procedure employed an adaptive rule that selected the premise with the poorest history (as given by a value continuously updated from the outset of training). In some instances, this value was response-dependent (i.e., the difference between the total number of pecks on the *S*+ and the total number of pecks on the *S*−); in other cases, the value was time-dependent (i.e., the total exposure time to each particular premise, independently of the distribution of pecks). In either case, overall, there was little evidence for TI.

At least two features of Higa and Staddon's (1993) procedure may have contributed to their null findings. First, birds never experienced non-reinforcement. The only consequence of pecking the *S*− was to reset the VI timer. If non-reinforcement is an important factor in reducing the relative value of each stimulus in order to create an order of differential values, then TI

Table 2  
Experiments using the 5-term series task with sequential training

Source	Subjects	N <sup>#</sup>	Test contingency	%Transitive choice <sup>*</sup>
Benard and Giurfa (2004), Experiment 1	Honeybees	12 (12)	Extinction	45
Benard and Giurfa (2004), Experiment 2	Honeybees	12 (12)	Extinction	50
Benard and Giurfa (2004), Experiment 3	Honeybees	8 (8)	Extinction	57
Davis (1992b), Experiment 1	Rats	4 (3)	Extinction	(100, 87.5, 83.3) <sup>a</sup>
Davis (1992b), Experiment 2	Rats	3 (3)	Extinction	(90, 85, 65) <sup>a</sup>
Steirn et al. (1995), Experiment 1	Pigeons	8 (8)	Non-differential reinforcement, 50%	85.4 <sup>a</sup> , 89.6 <sup>b</sup>
Steirn et al. (1995), Experiment 2	Pigeons	8 (8)	Non-differential reinforcement, 50%	87.5 <sup>a</sup> , 82.3 <sup>b</sup>

<sup>#</sup> Number of subjects in the experiment (inside parenthesis is the actual number of subjects tested or included in statistical analyses).

<sup>\*</sup> When variability in results justifies, and the results are available, individual preferences are shown. When numerical values are not available, values are estimated from figures.

<sup>a</sup> Backward training.

<sup>b</sup> Forward training.

might not emerge. Second, the most compelling evidence against TI seems to come from their first experiment. However, as the authors acknowledge, training trials and probe trials had significantly different durations in that experiment. Assuming birds were sensitive to this difference, they could refrain from pecking as soon as they detected that was a non-reinforced trial.

Despite these null findings, studies using an intermixed 5-term series overwhelmingly suggest that several animal species are capable of TI, a conclusion strengthened by studies using sequential training. Table 2 summarizes such studies. Again, the reinforcement contingency operating during testing varied (extinction or non-differential reinforcement) as did stimulus modality (olfactory stimuli with rats and visual stimuli with pigeons and honeybees), but these seemed to have little or no systematic impact on transitive performance.

In their first experiment, both Davis (1992b) and Steirn et al. (1995) employed a more elaborate training procedure: Instead of just using the four standard premises as in their second experiments, they added the discriminations AC, CE, and AE to encourage the development of TI. Specifically, animals were

successively trained in the following pairs: AB, BC, AC, CD, DE, CE, and AE. Davis used backward training, whereas Steirn and colleagues used both backward and forward training, but these manipulations were inconsequential given the evidence for TI in all groups. Their second experiments returned the typical transitive preference, too.

Benard and Giurfa (2004) used also forward and backward training in their series of experiments. Had their experiments been successful, they would constitute the first demonstration of TI in an invertebrate. Unfortunately, the honeybees were unable to learn the premises despite their proven capacity to learn visual patterns (e.g., Bitterman, 1996; Menzel and Giurfa, 2001). A factor that might have contributed to this failure was the reduced number of training trials (40 per premise). Steirn et al. (1995), for example, used a minimum of 288 trials per premise. Absent learning of the premise pairs, no conclusions about the ability of any animal to behave transitively are possible.

The *n*-term series task has not been restricted to five terms and, in fact, evidence for TI using longer series has been obtained (e.g., Bond et al., 2003, with corvids; Gillan, 1981,

Table 3  
Experiments using an *n*-term series task with more than five terms

Source	Subjects	N <sup>#</sup>	#Terms	%Transitive choice <sup>*</sup>
Bond et al. (2003), Experiment 2	Pinyon jays	5 (5)	7	86
	Scrub-jays	5 (4)		77
Gillan (1981), Experiments 1B, 2B, and 3	Chimpanzees	1 (1)	6	83.3 (Experiment 1B), 83.3 (Experiment 2B), 94.4 (Experiment 3)
Rapp et al. (1996), Phases 6–7	Rhesus monkeys	17 (16)	7	76.2 (young subjects)
				75.2 (aged subjects)
Roberts and Phelps (1994), Experiment 1	Rats	14 (14)	6	74.3 (linear order group) 47.9 (random order group)
Treichler and Van Tilburg (1996), Experiment 1, Test 2	Macaque monkeys	6 (6)	10	? <sup>a</sup>
Van Elzakker et al. (2003), Experiment 2	Rats	10 (9)	6	55 (BD test pair) 85 (BE test pair)
von Fersen et al. (1991), Experiment 2	Pigeons	4 (4)	7	82

<sup>#</sup> Number of subjects in the experiment (inside parenthesis is the actual number of subjects tested or included in statistical analyses).

<sup>\*</sup> When numerical values are not available, values are estimated from figures.

<sup>a</sup> In the original report, data are presented collapsed over all trial types (premises and transitive pairs). Average performance was 81%, but ANOVA showed that accuracy was significantly higher in the transitive pairs when compared with the premises.

with chimpanzees; Rapp et al., 1996, with rhesus monkeys; Roberts and Phelps, 1994, and Van Elzakker et al., 2003, with rats; Treichler and Van Tilburg, 1996, with macaque monkeys; von Fersen et al., 1991, with pigeons). Table 3 summarizes these studies. Again, there was ample procedural variation, such as the species studied, the number of terms (from 6 to 10), and the reinforcement contingency on probe trials (a particular feature omitted in Table 3 because it was inconsequential). Overall, the evidence for TI is quite convincing. Excluding one of the groups in Roberts and Phelps (1994, Experiment 1), whose expressed purpose was to eliminate TI, and one test pair in Van Elzakker et al. (2003, Experiment 2) the lowest average preference for the transitive alternative reported in these studies was 74.3%.

Roberts and Phelps (1994, Experiment 1) trained two groups of rats in two different variations of a 6-term series task. During premise training, both groups were trained on five odor discriminations. The linear order group had to discriminate between stimuli located in adjacent boxes with each particular box always presented in the same spatial location, so as to produce a linear arrangement from A through E. Rats in the random order group had to discriminate between stimuli located in adjacent boxes, too, but the boxes were placed at random locations. Test trials (B vs. D) were then given in a Y-maze. With this procedure, Roberts and Phelps found evidence for TI only in the linear order group.

As previously stated, an important advantage of longer series is that more transitive tests are possible—although only the average across transitive tests is presented in Table 3. Also, with longer series, the test pairs can vary in the number of intervening terms. For example, the test pairs BD and CE have only one intervening term (C and D, respectively), whereas the pair BE has two (C and D). In humans, response accuracy increases and reaction time decreases as the number of intervening terms between the two test stimuli increases (e.g., Bryant and Trabasso, 1971; McGonigle and Chalmers, 1984). This phenomenon is known as the *symbolic distance effect* (SDE). The number of studies analyzing this particular phenomenon in non-human animals is rather limited and, in some of them, the SDE is unfortunately confounded with *first-item effects* which refer to the finding that within equal symbolic distances, the pair closer to the always reinforced end of the series will be solved more accurately (von Fersen et al., 1991; Wynne, 1997).

Fig. 2 illustrates this confound as well as other potential problems with a 7-term series task. Suppose that after training the six premises of a 7-term series task, all non-adjacent pairs are presented during test sessions. Six of these pairs would be truly transitive (BD, BE, BF, CE, CF, and DF; shaded pairs in the figure) because they are non-adjacent and do not include any of the end-anchor stimuli (A or G). The remaining nine are non-adjacent, too, but not truly transitive because each includes at least one of the end-anchor stimuli.

Each column in Fig. 2 represents a symbolic distance and each row a different distance to the always reinforced end of the series (in this case, A). The *modus operandi* to analyze the SDE has been to average results within a particular symbolic distance (i.e., within columns) and then compare the averages across symbolic distances. This procedure usually yields the typ-

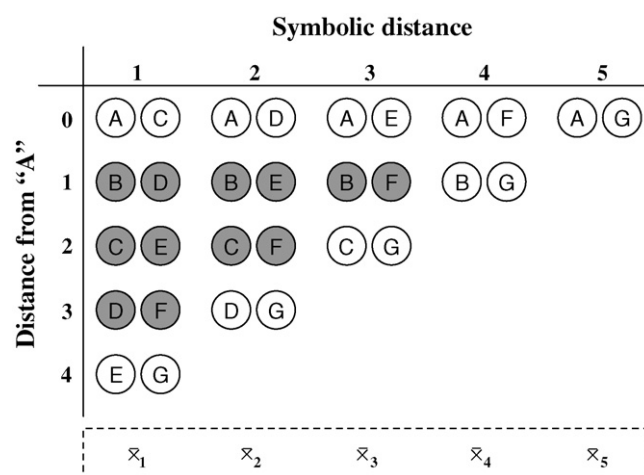


Fig. 2. Schematic of a typical test session after training a 7-term series task. Shaded pairs indicate truly transitive discriminations.

ical SDE ( $\bar{X}_1 < \bar{X}_2 < \bar{X}_3 < \bar{X}_4 < \bar{X}_5$  for accuracy data, and the reverse for reaction-times). Note, however, that the mean of each symbolic distance includes pairs from different rankings in the orthogonal dimension reflecting the first-item effects (i.e., distance from the end-anchor "A" stimulus). Unconfounded depictions of the SDE can be obtained by analyzing separately the pairs involving the second-ranked item (BD, BE, and BF) and/or the pairs involving the third-ranked item (CE and CF). Ideally, symbolic distance averages should only be calculated for symbolic rankings including two or more truly transitive pairs and should exclude any pairs including the end-anchor stimuli, although this is not the customary procedure.

An additional concern about the typical procedure used to analyze the SDE is that the mean for each symbolic distance is based on a different number of data points—the longer the symbolic distance, the smaller the number of available data points. This problem can be circumvented, too, by analyzing the SDE only within a given distance from the always reinforced end of the series (e.g., BD vs. BE vs. BF).

Mindful of these common confounds, the SDE in non-human animals is reliable, at least when based on accuracy data. For example, Treichler and Van Tilburg (1996) used a particularly potent procedure to detect the SDE by employing what functionally can be considered a 10-term series task. In particular, macaque monkeys were consecutively trained in two 5-term series tasks (from A to E and from F to J). After acquisition, these series were linked by training the pair EF. During testing which included both within- and between-series pairs, a clear trend toward more accurate performance with increasing number of intervening terms was observed.

Evidence for the SDE was also obtained by von Fersen et al. (1991, Experiment 2) with pigeons. They used a 7-term series task which allowed for six different non-adjacent test pairs. Their procedure was exactly like the one schematized in Fig. 2. Evidence for TI was observed in all pairs as was a tendency toward better performance in non-adjacent pairs as the number of intervening terms increases (although they, too, did not control for the first-item effects). Overall, the slope of the regression line

accounted for 75% of the variance, but this did not reach statistical significance mostly due to the limited number of classes (three). Nonetheless, their results are in the correct direction and make a compelling case for the SDE in non-human animals.

Statistically significant results were, however, reported by Bond et al. (2003) and Rapp et al. (1996), although only the former controlled for the first-item effects. In particular, Bond and colleagues obtained an unconfounded SDE in scrub jays for pairs involving the second-ranked item (BC, BD, BE, and BF). The absence of the effect in pinyon jays might have been due to ceiling effects because, as the authors acknowledge, two out of five pinyon jays made no errors in the non-adjacent pairs. With pairs involving the third-ranked item (CD, CE, and CF) a similar, although non-significant, trend was observed in both groups. Rapp and colleagues, on the other hand, obtained both an accuracy- and latency-based SDE in rhesus monkeys (although only young monkeys showed a SDE based on latency). Further accuracy-based evidence was obtained by Van Elzaker et al. (2003) in a 6-term series with rats. Using a between-subject design, they observed better transitive performance on the BE pair than on the BD pair (indeed, performance on the latter was at chance).

Studies using 5-term series have provided some evidence for the SDE, too. These findings should, however, be interpreted cautiously because of the confounds mentioned earlier. These studies include Higa and Staddon (1993, Experiments 2 and 3, only for some of the birds), Lazareva and Wasserman (2006), Siemann et al. (1996b) and von Fersen et al. (1991, Experiment 1). Particularly noteworthy are the studies reported by D'Amato and Colombo (1990, Experiment 2), McGonigle and Chalmers (1992), and Wynne (1997) because latency data were used to assess the SDE. Curiously, the results are conflicting: D'Amato and Colombo, and McGonigle and Chalmers found an inverse relationship between symbolic distance and latencies, whereas Wynne obtained a flat relationship despite observing a clear SDE based on accuracy.

Some null findings have also been reported. In Experiments 1B, 2B, and 3, Gillan (1981) used a 6-term series which allowed for three different non-adjacent test pairs (e.g., BD and CE with one intervening term, and BE with two intervening terms). Across experiments and test pairs, transitive choices averaged 87% overall but there was no obvious effect of test pair. However, because the same chimpanzee participated in all three experiments and exhibited optimal transitive performance in an immediately preceding experiment (Experiment 1A in Table 1), this may have created a ceiling effect which precluded observing a SDE.

The bulk of empirical support for TI and associated effects in non-human animals has hinged on the  $n$ -term series task. Despite variations in experimental procedures, studies have frequently obtained preference data that establishes TI as a real phenomenon in non-human animals. Additional support has come from rare studies using different tasks, such as Hogue et al. (1996), Paz-y-Miño et al. (2004), and Grosenick et al. (2007). Only latency data seem inconclusive, although their relevance should be somewhat deflated in face of their indirect nature.

All this evidence raises the question: What processes underlie TI in non-human animals? Several models have been proposed, but to fully appreciate their core assumptions, a clear depiction of premise performance during training and testing is needed. I turn to this next.

### 3.2. Premise performance

How rapidly do animals learn the premises of an  $n$ -term series task? The answer to this question varies widely across species but, generally, training proceeds rather slowly.<sup>2</sup> This fact, by itself, is not particularly surprising because a 5-term series task can be conceptualized as three concurrent and partially overlapping ambiguous-cue problems (e.g., Fletcher et al., 1968; Hall, 1980). An ambiguous-cue problem consists of two simultaneous discriminations in which the  $S+$  for one discrimination serves as the  $S-$  for the other. For example, consider only the premises  $A+B-$  and  $B+C-$ . Here, stimulus  $B$  is the ambiguous cue because it signals reinforcement or non-reinforcement depending on the other stimulus with which it is presented. By contrast, both  $A$  and  $C$  have unambiguous status as signals for reinforcement and non-reinforcement, respectively, in such a two-premise task. The typical finding is that asymptotic accuracy is lower and reached more slowly for the discrimination in which the ambiguous-cue is non-reinforced (e.g., Fletcher et al., 1968; Hall, 1980; Leary, 1958; Richards, 1973). With a 5-term series task, there are *three* ambiguous cues ( $B$ ,  $C$ , and  $D$ ), hence the slow acquisition rate comes as no surprise given their participation in the BC and CD discriminations (Urcuioli and Michalek, 2007).

Even though this slow acquisition rate is not always observed (e.g., it is particularly prevalent in avians, less so with primates), it has prompted several attempts to design training procedures to yield faster acquisition. One common strategy has been to include a phase of *sequential* training at some point during acquisition. In fact, only Higa and Staddon (1993), and Roberts and Phelps (1994) did not include a sequential phase in their training protocol. This was also the rationale for Higa and Staddon's (1993) auto-run procedure and for other procedural variations (e.g., Buckmaster et al., 2004; Siemann et al., 1996b; Wynne, 1997).

Although most animals eventually perform significantly above chance on all premises, how accurate are they on the various trained premises? The optimal scenario would be for accuracy to stabilize at high and comparable levels for all premises. This is seldom the case, though. Frequently, accuracy levels follow a U-shaped function, with the pairs at each end of the series better solved than the "interior" ones. By analogy with the typical retention functions obtained in immediate serial recall, this has been called the *serial position effect* (Bryant and Trabasso, 1971; Woocher et al., 1978). As will be shown in a forthcoming section, this feature is of critical importance for the reinforcement-based accounts of TI. Fig. 3 illustrates an ideal-

<sup>2</sup> This is true only for intermixed training. When a sequential procedure is used, training proceeds rapidly.



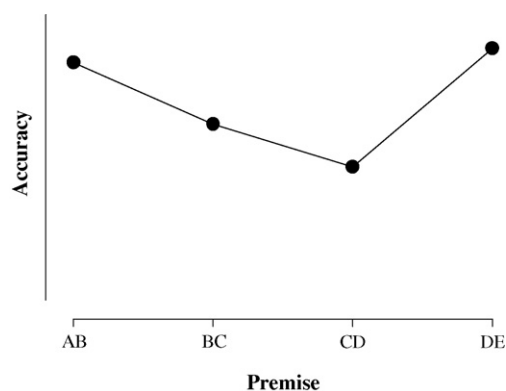


Fig. 3. Idealized serial position curve after 5-term series training.

ized serial position curve when animals are trained in a 5-term series task.

A variety of studies has reported the serial position effect in several species. In fact, the effect was present in the original demonstrations of TI in young children (Bryant and Trabasso, 1971; for an example in human adults, see Greene et al., 2001) as well as in the later demonstrations in non-human animals (e.g., Dusek and Eichenbaum, 1997; Gillan, 1981, Experiment 1A; Higa and Staddon, 1993; Lazareva et al., 2004; Lazareva and Wasserman, 2006, particularly in the constant feedback group; McGonigle and Chalmers, 1977, 1992, Experiment 1; Treichler and Van Tilburg, 1996; von Fersen et al., 1991, Experiments 1 and 2; Wynne, 1997). Overall, the effect is almost as prevalent in TI studies as transitive performance itself. Only Roberts and Phelps (1994), and Siemann et al. (1996b) did not obtain the U-shaped function, although both studies found that accuracy increased the closer a premise was to the never-reinforced end of the series (for a somewhat intermediate form, see Van Elzakker et al., 2003).

### 3.3. The circular series and the circular test

Most views about TI implicitly assume that this ability will emerge if and only if the series is linear (e.g.,  $A > B > C > D > E$ ). To test this assumption, several researchers have implemented circular series. To illustrate, suppose that after training the four premises shown in Fig. 1, the premise  $A-E+$  is introduced. With such training, the spatial arrangement of the stimuli is circular,<sup>3</sup> as shown in Fig. 4. Note, however, the logical impossibility created by such a series (e.g.,  $A$  cannot be both “larger” and “smaller” than  $E$ ).

This circular series task, known as the *transverse pattern-ing problem*, has actually been investigated outside of the TI realm using series with less than five terms. In general, these studies show that animals are capable of solving the task (e.g., Alvarado and Rudy, 1992, 1995; Couvillon and Bitterman, 1996; Thompson, 1953; Wynne, 1996). Studies directly related to TI have returned relatively similar findings, although with lower




Fig. 4. Circular arrangement of the stimuli when a 5-term series is closed into a loop.

overall accuracy caused by the increased length of the series used. Briefly, studies with pigeons (von Fersen et al., 1991, Experiment 3), rats (Davis, 1992b, Experiment 3; Roberts and Phelps, 1994, Experiment 2), and chimpanzees (Gillan, 1981, Experiment 2A) show that circular series training yields no TI but, despite some disturbances observed during premise training, subjects are usually able to learn the premises above chance levels. These particular results are of primary importance for evaluating the adequacy of the proposed models to explain TI in non-human animals. I analyze these models next.

## 4. Models of transitive inference

Of the several models proposed to explain TI, some are qualitative but most are quantitative in nature; some are clearly cognitive, using concepts such as mental representations and mental lines, but others are more behaviorally oriented focusing on such things as reinforcement histories and relative frequency of past events.

### 4.1. Cognitive models

Accounts of TI in humans traditionally ascribe this ability either to the application of the rules of formal logic or to the use of some kind of *mental model* (Byrne and Johnson-Laird, 1989; Johnson-Laird and Byrne, 1991). According to Piaget (1953), transitive syllogisms are solved through the coordination of the premises via application of logical rules commonly based on language. In this view, reasoning is literally synonymous with propositional calculus. Thus, it was thought that children become capable of TI only when they are old enough to successfully apprehend and use the rules of logic (viz., when they reach the concrete operational stage). Recently, however, theorists have begun to doubt the validity of the logical accounts in everyday problem solving and have proposed instead that subjects form some kind of mental model of the task at hand (for evidence in humans, see Byrne and Johnson-Laird, 1989; Carreiras and Santamaria, 1997; for a review, see Johnson-Laird and Byrne, 1991). In particular, instead of relying on logical principles,

<sup>3</sup> The spatial terminology is used only as a metaphor for ease of presentation. Sometimes, however, this metaphor is taken too literally, as the reader will notice in a forthcoming section.

spatial *paralogical* devices have been proposed (Acuna et al., 2002; Breslow, 1981; De Soto et al., 1965; Huttenlocher, 1968; McGonigle and Chalmers, 1986; Sternberg, 1980; Trabasso and Riley, 1975). These paralogical models assume that in an  $n$ -term series task, the subject integrates the independently presented premises into a (or series of) mental representation(s), supposedly spatial in nature, which can then be used to correctly solve any premise and test pair. Hence, a successful transitive decision for any test pair depends on determining which of the members of the pair is nearer one of the ends of the series, thus eliminating the need for formal deductions.

Models closely related to the application of the rules of formal logic as well as models based on spatial paralogic have been proposed to explain TI in non-human animals. For example, Harris and McGonigle (1994) proposed a model of TI based on the successive application of “if . . . then . . .” rules akin to the earlier logical models. This type of model has been successfully applied to cognitive processes like problem solving and reasoning (Newell and Simon, 1972), human learning (Anderson, 1993), and reinforcement-driven learning in rats (Holland et al., 1986). According to the Harris and McGonigle model, a stack of four rules like the following is sufficient to produce optimal performance in a 5-term series task:

1. If stimulus *A* is present, select it
2. If stimulus *E* is present, avoid it
3. If stimulus *B* is present, select it
4. If stimulus *C* is present, select it

Surveying these rules *in sequence* and terminating the decision process as soon as any of the prescribed actions is taken will ensure optimal performance (Harris and McGonigle, 1994). The particular stack used by each animal is supposedly learned on a trial-and-error basis during premise training, although trial-by-trial dynamics have not been explicitly proposed. With one additional assumption, the model can predict performance on *triadic* tests in which three test stimuli are presented instead of two. Consider, for example, two triadic tests: ABD and BDE. The first one can be easily solved by applying rule 1 in the stack above, but the second test requires two rules (2 and 3). To correctly mimic published data, the model further assumes that when a triad requiring the application of more than one rule is encountered (such as the second example), the animal applies the first rule and then chooses randomly between the remaining two stimuli on a certain (empirically estimated) percentage of trials. On the remaining trials, the animals resume the application of the rules. Using these assumptions, Harris and McGonigle showed that the model is able to accommodate both dyadic and triadic choice tests as well as all of the other main results reported in McGonigle and Chalmers (1977, 1992), including a latency-based SDE.

The model, however, suffers from serious shortcomings. For example, even assuming that animals learn these rules and order them appropriately, the model is unable to accommodate sequential training. To do that, one would have to assume that the animals first learn each rule and, only then, order them sequentially. But this presupposes that animals have formed a linear

representation of the stimuli, a concept suffering from its own problems, as I will show next. Also, in their modeling, Harris and McGonigle (1994) used only the 16 stacks that perform correctly on all premises, but this necessarily excludes the serial position effect and error in binary transitivity tests from the model predictions. In addition, none of the proposed stacks allows learning of circular series—contrary to evidence. Finally, and most importantly, what determines each rule and its order? Is it continuous reinforcement, partial reinforcement, and non-reinforcement? If so, then the model is unnecessarily complex, given that there are conditioning models that have been successfully applied to TI (see below).

Recently, Bryson and Leong (2006) proposed a two-tier learning model inspired by Harris and McGonigle’s (1994) model that overcomes some of the latter’s difficulties. One tier prioritizes stimuli and the other associates actions (select or avoid) with stimuli. In both tiers, priority is given by a weight that is learned on a trial-by-trial basis through reinforcement. Hence, on any given trial in which a choice between familiar stimuli is given, the animal attends to the more salient stimulus (i.e., the stimulus with highest priority) and then either selects or avoids it depending on the action weights for that particular stimulus on the second tier. The model has the advantage of specifying trial-by-trial dynamics, predicting the serial position effect, and fitting primate performance in triadic tests. Nevertheless, without further assumptions, it still cannot predict the SDE.

Spatial representations have been proposed as fundamental cognitive structures in non-human animals (e.g., Gallistel, 1990), so the possibility that animals integrate the premises from a TI task into an ordered series of mental representations has been entertained (e.g., D’Amato, 1991; D’Amato and Colombo, 1990; Davis, 1992a,b; Dusek and Eichenbaum, 1997; Gillan, 1981; McGonigle and Chalmers, 1986; Rapp et al., 1996; Roberts and Phelps, 1994; Terrace, 1986; Terrace and McGonigle, 1994; Treichler and Van Tilburg, 1996). These models share the fundamental assumption of premise integration into some symbolic spatial representation but differ, for example, with respect to when integration occurs (Breslow, 1981; Trabasso and Riley, 1975). For example, Bryant and Trabasso’s (1971) model proposes that during training the subject forms a separate representation of each premise and only during testing does the subject integrate the separate representations into a mental line. Conversely, most spatial representation models, such as the spatial integration model (e.g., De Soto et al., 1965), suggest that training directly promotes the integration of the premises into a linear series.

But the more fundamental questions are: (a) How is this mental line formed? and (b) How does the mental line get “translated” into the signature performance effects? One proposal is that the always reinforced and always non-reinforced stimuli are located first at the ends of the metaphorical mental line. With training, the other stimuli are then progressively ordered between the ends. Once formed, the line is used to solve any discrimination during training or testing via a spatial search along the imaginary line. Also, the end-inward construction of the mental representation is supposed to explain both the serial position

effect and the SDE. Because the end-anchor stimuli are unambiguous, a smaller number of errors will occur in those pairs that include one of them. More errors will occur in the pairs involving only non-end-anchor stimuli because their reinforcement status is ambiguous. Likewise, the SDE arises because the closer the stimulus representations are in a series, the more similar they are spatially, hence the more difficult the discrimination between them.

Several experiments have investigated the viability of such models in non-human animals. For example, Roberts and Phelps' (1994) experiments with linear and nonlinear (random or circular) arrangements were a test of this type of model and, at first glance, their results lend it some credibility. Namely, their findings of TI only when the stimuli were linearly arranged in a consistent fashion make a strong argument for spatial coding of the premises. Alternatively, their results may only mean that rats use spatial cues, when available, to supplement other cues (e.g., olfactory) in choosing the correct alternative in each premise. Such an interpretation does not imply any linear representation, although it must then explain the null findings from the non-linear groups, where spatial cues were made uninformative by systematically varying them across premise pairs. Successful acquisition of the premises by the non-linear groups can easily be explained by the presence of relevant olfactory cues, but the absence of TI is problematic. The reasons for this failure are not clear but other research shows that spatial cues are not required to observe TI in rats (e.g., Davis, 1992b; Dusek and Eichenbaum, 1997; Van Elzaker et al., 2003).

In another attempt to test the spatial representation hypothesis, Lazareva et al. (2004) and Lazareva and Wasserman (2006) used a post-choice feedback procedure to enhance a linear representation of the stimuli. In the authors' opinion, including the feedback stimuli provided an explicit scale to an otherwise non-orderable set of stimuli. However, the results were inconsistent between species and this difference was attributed to differences in the social hierarchies of the species. Also, if the post-feedback stimuli do indeed provide such a scale, it is not clear why the ordered feedback groups did not learn the premises faster than the constant feedback groups.

Dusek and Eichenbaum (1997) hypothesized that the hippocampal region is critical to TI and, in fact, to any task involving a hierarchical relation among stimuli (see also Eichenbaum, 2000, 2004; Rapp et al., 1996; Wallenstein et al., 1998). To test their hypothesis, they trained and tested three groups of rats in a 5-term series task. One group had a transection of the fornix, another had the perirhinal and entorhinal cortices removed, and the remaining one was a sham-operated control group. They found evidence for TI only in the control group, which suggests that the hippocampus plays a pivotal role in TI (for a similar manipulation in monkeys, see Buckmaster et al., 2004). However, despite the known hippocampal contribution to spatial learning and memory (e.g., Jarrard, 1993; Kesner, 1991; O'Keefe and Nadel, 1978), these data do not necessarily imply that animals form a linear representation of the premises. Alternative accounts (e.g., Frank et al., 2003; Van Elzaker et al., 2003) seriously question the role of the hippocampus in the flexible integration of independently presented premises. In

fact, the hippocampal explicit memory system seems to hinder transitive performance in some cases. For example, Frank et al. (2006) have shown that the benzodiazepene midazolam, a drug that has been argued to disrupt hippocampal function (e.g., Curran et al., 2006), actually enhances transitive performance in humans while resulting in serious memory deficits in a free recall task.

Overall, then, models appealing to some kind of spatial representation seem to have some qualitative agreement with the bulk of non-human data, although the results are open to alternative interpretations. Besides, the problems faced by such accounts detract from their attractiveness.

First, circular series are problematic because animals do learn, to some extent, the premises of a circular series (although failing the transitive tests). To accommodate these results, these models must assume that the shape of the mental representation is determined by the nature of the relations being learned, contrary to at least some proponents of these accounts (e.g., De Soto et al., 1965) who assumed that the linear arrangement is a cognitive *good figure*. Are we to assume that the end result of training a circular series is the formation of a mental circle?

Second, it is not entirely clear how these accounts accommodate sequential training because, with this training regimen, one of the end-anchor stimuli is the last to be introduced. This clearly violates the end-inward construction postulated by most models, unless the separate premise representations are integrated into a mental line only during testing as proposed by a minority of these models.

Third, independently of when the linear representation forms, it is not clear why transitive performance is not perfect. Most working memory models assume that, once retrieved, information is readily available; hence, every search along the mental line should yield perfect performance. In other words, these models are usually not probabilistic and do not include any assumption that adds noise to the decision-making process. In fact, some attempts have been made to add such noise, but the results have clearly been far from perfect. For example, McGonigle and Chalmers (1977) proposed a *binary sampling model* that yields less than perfect performance. Specifically, they proposed that on a given test trial, the animal identifies the stimulus missing between the test stimuli (on a BD test trial, C is missing; note the linear representation assumption), retrieves all possible combinations of the stimuli (in the case above, BC, CD, and BD), and finally randomly chooses one pair. If the chosen pair is a familiar premise from training (BC or CD), the correct stimulus is chosen. On the other hand, if the test pair (BD) is chosen, the animal selects one of the stimuli at random. Although correctly predicting a less than perfect transitive performance, this model also predicts the reverse of the SDE: The farther apart the test stimuli are, the more unfamiliar non-adjacent pairs are available for the animal to choose, hence leading to a greater proportion of random choices. Also, the model does not predict the serial position effect.

Finally, most of these models are static; they are silent about important procedural details. For example, why is the line constructed in an end-inward fashion? In what way does this guarantee that premises and test pairs including one of

the end-anchor stimuli are better solved? Why is it assumed that the “inner” stimuli are closer together in the representation to explain the SDE? Why does this proximity entail a greater number of errors? All these questions remain unanswered.

#### 4.2. Value transfer theory

Value transfer theory (VTT), initially proposed by von Fersen et al. (1991), assumes that associative value transfers from one stimulus to the other in a simultaneous discrimination. Subsequent research has shown that value transfers only from the reinforced stimulus ( $S+$ ) to the non-reinforced stimulus ( $S-$ ) and not in the opposite direction (Clement et al., 1998).

According to VTT, the total value of a stimulus  $i$  is the sum of two components: its direct value ( $R_i$ , conferred by direct reinforcement) and its indirect value ( $V_{i+1}$ , transferred from any  $S+$  with which stimulus  $i$  is paired). Eq. (1) shows the details.

$$V_i = R_i + a * V_{i+1} \quad (1)$$

where “ $a$ ” ( $0 < a < 0.5$ ) determines the amount of value transferred from the  $S+$  to the  $S-$  (for details, see von Fersen et al., 1991).

If we consider a 5-term series task, value will transfer from the  $S+$  to the  $S-$  of each pair, but the amount of value transferred will depend on the total value of the positive member. Remember, for example, that  $B-$  is presented with  $A$  (which is always reinforced) but  $D-$  is presented with  $C$  (reinforced when presented with  $D$ , but non-reinforced when presented with  $B$ ). Therefore, to the extent that the value of  $A$  is greater than the value of  $C$ , the value transferred from  $A$  to  $B$  will exceed the value transferred from  $C$  to  $D$ . By the end of training, then, the total value of  $B$  should exceed the total value of  $D$  because even though their direct values are equal (each is partially reinforced), their indirect values favor  $B$ . Thus,  $B$  should be preferred over  $D$ —the typical transitive choice.

Following von Fersen et al. (1991), I assume, for simplicity, that the direct value of a stimulus is either 2, 1, or 0, depending on whether the stimulus is always reinforced, reinforced with a probability of .5 or never reinforced, respectively. By the end of training, the different stimuli of the series will attain different and therefore orderable values (Table 4 shows the details). The final order of values is independent of the actual value of “ $a$ ”, and can be represented as:

$$V_A > V_B > V_C > V_D > V_E$$

This order has two important features: (a) on test trials ( $B$  vs.  $D$ ) it predicts a preference for  $B$ , or, in other words, it predicts a transitive-like behavior, and (b) if we assume that the choice in

Table 4  
Final value for the 5-term series’ stimuli according to value transfer theory (from von Fersen et al., 1991)

$V_A = 2$
$V_B = 1 + 2a$
$V_C = 1 + a + 2a^2$
$V_D = 1 + a + a^2 + 2a^3$
$V_E = a + a^2 + a^3 + 2a^4$

any of the discriminations is a function of the difference between the values of the presented stimuli, it also predicts the serial position effect and the SDE (for detailed computations, see von Fersen et al., 1991). Furthermore, these predictions are valid for any  $n$ -term series and for any transitive test trials.

Independent evidence for value transfer was provided in a rather simple and elegant experiment by Zentall and Sherburne (1994). Pigeons were trained on two simultaneous discriminations with no overlapping stimuli. In one of the discriminations, the  $S+$  was always reinforced ( $A+B-$ ); in the other the  $S+$  was reinforced only 50% of the time ( $C\pm D-$ ). During testing, probe trials with both  $S-$  stimuli ( $B$  vs.  $D$ ) were interspersed among regular training trials. Note that both  $B$  and  $D$  had the same direct history of non-reinforcement, hence the same direct value ( $R_B = R_D = 0$ ). If no value transfer occurs, there should not be a preference for either stimulus. However, if value transfer occurs,  $B$  should be preferred over  $D$  because  $B$  was presented with an always reinforced  $S+$  ( $A+$ ) whereas  $D$  was presented with a partially reinforced  $S+$  ( $C\pm$ ). As a result, the indirect value ( $V_{i+1}$ ) transferred from the  $S+$  to the  $S-$  will be larger in the  $A+B-$  discrimination than in the  $C\pm D-$  discrimination. This prediction was confirmed by Zentall and Sherburne in two different experiments (1994; see also Cohen et al., 2001; Siemann et al., 1996a; Zentall et al., 1996b). Since then, value transfer has also been implicated in other discrimination phenomena (e.g., Urcuioli, 2006; Urcuioli and Michalek, 2007).

But how does value transfer? One possible mechanism is second-order conditioning (Davis, 1992a). Assuming that when two stimuli are presented simultaneously, the animal observes them sequentially (Wright and Sands, 1981), then the  $S-$  is functionally followed by the  $S+$  on roughly one half of the trials. In Pavlovian terms,  $CS_2$  is followed by  $CS_1$  and then by the US. Such second-order conditioning should impart conditioned reinforcing properties (“value”) to  $CS_2$  (Zentall, 2004; Zentall et al., 1996a).

In a direct test of VTT, Steirn et al. (1995, Experiment 3) trained pigeons in a task akin to the  $n$ -term series task. However, the simultaneous discriminations were arranged such as to produce differential transferred value without any transitive relation between the test stimuli. In particular, birds were trained in the following discriminations:  $A+B-$ ,  $E+C-$ ,  $C+D-$ , and  $A+E-$ . The value inequalities produced by such training can be expressed as  $A > B$  and  $A > E > C > D$ . With this procedure, the test stimuli ( $B$  and  $D$ ) are located in two different arrays, thus precluding any transitive relation. However, the associative processes envisioned by VTT can still operate and produce transitive-like performances. That was exactly what Steirn and colleagues found:  $B$  was preferred over  $D$ .

VTT was one of the first quantitative accounts of TI although without initial specification of the underlying mechanism. A plausible mechanism has since been identified and value transfer has proven conceptually useful for a variety of discrimination phenomena. Still, VTT has its own difficulties.

For example, although originally proposed in the context of intermixed training, when applied to sequential training its operating assumptions seem questionable. Consider, for example, forward sequential training (Fig. 5, left panel) again assuming

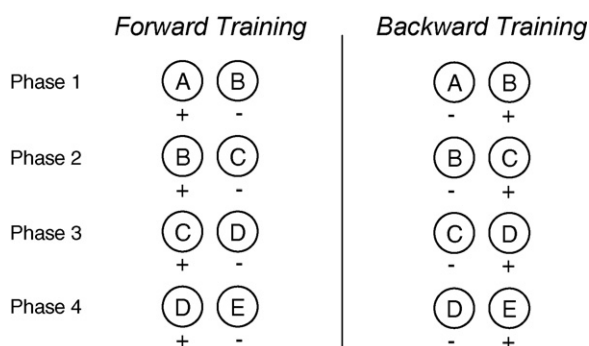


Fig. 5. Training order in a forward (left panel) and backward (right panel) sequential 5-term series task.

direct values of 2, 1, or 0, depending on whether the stimulus is always reinforced, reinforced with a probability of .5, or never reinforced, respectively. By the end of phase 1, stimulus *A* should have a total value of 2 and stimulus *B* should have a total value of  $0.2 = 0 + (0.1 \times 2)$  if “*a*” equals 0.1. But what happens during phase 2? Should the value of *B* be increased by 2 because *B* is now always reinforced? If so, the overall value of *B* would be 2.2 which is greater than the overall value of *A*. Extending this reasoning to phases 3 and 4 implies that, by the end of training, the order of stimulus values would be precisely the opposite of that needed to predict TI. The only way to predict TI with sequential training is to attend only to the phases in which the test stimuli, *B* and *D*, are non-reinforced (phases 1 and 3, respectively): When non-reinforced, *B* is presented with an always reinforced stimulus (*A*) and *D* with an overall partially reinforced stimulus (*C*). But this ignores the sequential characteristic of this training procedure and the question of how to justify *C* as a partially reinforced stimulus in phase 3. A rational solution is not obvious.

Similar problems arise when we analyze backward sequential training (Fig. 5, right panel). Here, the test stimuli (*B* and *D*) begin by being continuously reinforced and are then shifted to extinction. In this situation, VTT predicts that the overall value for all intermediate stimuli will be the same (2.2 assuming the same parameters as before). Naturally, no TI should emerge with such training but, in fact, it does. To be able to predict TI, one must again focus on the phases in which the test stimuli are non-reinforced. When non-reinforced, *B* is presented with an overall partially reinforced stimulus and *D* with an always reinforced stimulus.<sup>4</sup> But here an additional twist is needed: *B* is non-reinforced when presented with *C* (phase 2), but, at that moment, *C* has never been non-reinforced and therefore it functionally constitutes an always reinforced stimulus (*C* is non-reinforced only in phase 3 of training). So, in order to consider *C* a partially reinforced stimulus, the next training phase must be taken into account which then retroactively reduces the value already transferred to *B*.

Despite these difficulties, some of these *ad hoc* interpretations of VTT may, in fact, be plausible (e.g., Colwill and Rescorla, 1985a,b, 1990; see suggestions for future research below). A

critical test, though, involves partially reinforcing *A* (i.e.,  $A \pm B -$ ) in forward training or partially reinforcing *E* (i.e.,  $D - E \pm$ ) in backward training. In such a case, both test stimuli would be paired with an (overall) partially reinforced stimulus so the value transferred to each of them would be the same. Given identical direct and indirect values, no preference should emerge in the TI tests. Indeed, such a test has been performed using forward sequential training and, contrary to predictions, animals showed consistent transitive choices (Steirn and Weaver, 1995; Weaver et al., 1997, Experiment 1, Group E0). An alternative interpretation of these results is based on the fact that the non-transitive alternative is actually paired with an always non-reinforced stimulus which could, theoretically, detract from its value. If this were the case, value transfer could theoretically predict TI even when the transitive alternative is paired with a partially reinforced stimulus. Weaver et al. (1997) tested this hypothesis and found no evidence for such *negative* value transfer (see also Clement et al., 1998).

VTT also has difficulties explaining why animals can learn the premises of a circular series. Because all stimuli in such a series are partially reinforced overall, their direct value would be 1 and the indirectly transferred value would be equal for all stimuli (in particular, it would be  $a + a^2 + a^3 + \dots$ ). Hence, all stimuli should acquire the same overall value and, therefore, TI should not be observed on test trials. And in fact, it's not. But VTT also predicts that subjects should be unable to *learn* the premises, when in fact they do.

In short, value transfer is a real phenomenon (Zentall and Clement, 2001) apparently based on second-order conditioning. However, when applied to TI, VTT does not pass some critical tests and has some obvious problems with sequential training and circular series.

#### 4.3. Reinforcement-based models

Several stochastic linear operator models have been proposed to explain TI in non-human animals. In all these models, choice is assumed to be based on a common judgment scale, with each alternative representing a point in such a scale. The general assumption is that reinforcement and non-reinforcement have opposite effects on the probability of a subsequent response: The former increases it and the latter decreases it—an idea that goes back to Thorndike (1898).

This general assumption can be captured by Reward/Non-reward ratios. Suppose, for instance that, before testing, we calculate the number of Rewards (#*R*) and Non-rewards (#*N*) for each stimulus. Then, the Reward/Non-reward ratio ( $r = \#R/\#N$ ) for each particular stimulus can be used (as an estimate) to compare the values of the two test stimuli presented on any test trial. To obtain a transitive choice, the stimuli must have different histories of reinforcement and non-reinforcement in training (actually, the serial position effect suggests this is the case).

This particular approach has already been used by Delius and Siemann (1998) to successfully account for von Fersen et al.'s (1991) results without postulating value transfer. Although generally consistent, keep in mind that the models that follow are

<sup>4</sup> In this regimen, a preference for *D* over *B* would be evidence for TI.

dynamic and, therefore, their predictions do not always match the ratios' predictions (for an example, see Lazareva et al., 2004).

I will present the details and predictions of each particular model first and, since most models are empirically necessitated adaptations of their predecessors, will then discuss their strengths and vulnerabilities jointly. I omit the Bush–Mosteller model (Bush and Mosteller, 1955) as proposed by Couvillon and Bitterman (1992) largely because the Rescorla–Wagner model and Wynne's (1995, 1998) configural model were subsequently proposed to overcome its limitations.

#### 4.3.1. Rescorla–Wagner model

Wagner and Rescorla (1972) suggested that their well-known classical conditioning model could be easily applied to instrumental situations. To that purpose, the value of a stimulus should be incremented or decremented by some quantity every time it is presented; this quantity, in turn, should depend on the difference between a reinforcer-based asymptote and the combined value of the chosen stimulus and the remaining context. What follows is an extension of their famous equation, adapted by Wynne (1995) into a model of TI.

Suppose that on a particular trial, stimuli  $X$  and  $Y$  are presented, and  $X$  is reinforced and  $Y$  is not. Suppose also that the animal correctly chooses  $X$ . Then, according to this model, the value of  $X$  should be updated according to the following rule:

$$V(X)_{i+1} = V(X)_i + \beta * (1 - [V(X)_i + V(Z)_i]) \quad (2)$$

where  $V(X)_{i+1}$  is the updated value of  $X$  after the current trial,  $V(X)_i$  is the value of  $X$  before the current trial,  $\beta$  is a learning parameter,  $V(Z)_i$  includes both the context and any commonalities between the two presented stimuli, and the reinforcer-based asymptote equals 1.

Suppose now that the animal incorrectly chooses  $Y$ . Then, the value of  $Y$  would be updated according to the following rule:

$$V(Y)_{i+1} = V(Y)_i - \beta * [V(Y)_i + V(Z)_i] \quad (3)$$

Thus, in any given discrimination problem, the model assumes the continuous updating of a stimulus' value and that the animal chooses the stimulus with the highest relative value according to Luce's (1959) ratio rule:

$$r = \frac{V(X) + V(Z)}{V(X) + V(Y) + 2V(Z)} \quad (4)$$

To derive response probabilities, a scaling parameter ( $\alpha$ ) has to be added to the rule:

$$P(X|XY) = \frac{1}{1 + e^{-\alpha(2r-1)}} \quad (5)$$

where  $P(X|XY)$  is the probability of choosing  $X$  when presented along with  $Y$ .

As training proceeds, the value of each stimulus is continuously updated according to Eqs. (2) and (3), the values of the presented stimuli are compared using Luce's (1959) rule, and the stimulus with the highest relative value is chosen. The same applies to test trials.

Based purely on its mathematical assumptions, the serial position effect is a necessary consequence of this model and,

more specifically, of Luce's ratio. Because the series' last stimulus is never reinforced, its value will approach zero with training, which in turn predicts that the other stimulus in the premise will be chosen with probability 1. Also, because the first stimulus in the series is always reinforced, its value will be larger than the value of any other stimulus, leading to excellent performance on the first premise (Wynne, 1995). In a similar vein, the SDE effect is a direct consequence of the fact that, after training, different stimuli have different and orderable values, which in turn predicts that the greater the distance between stimuli (and hence the greater the numerical difference between their values), the greater the probability of choosing the correct stimuli (for more details, see Wynne, 1995). Finally, an important feature of such a model is that the final value of each stimulus critically depends on the precise sequence of trials and choices during training.

#### 4.3.2. Wynne's configural model

This model, proposed by Wynne (1995, 1998), is largely based on the Rescorla–Wagner model, but it makes the additional assumption that the stimuli are, to some extent, bound to the context in which they are presented (see Rescorla, 1972, 1973; Whitlow and Wagner, 1972). To clarify, the Rescorla–Wagner model assumes that each stimulus is the same independently of the context in which it is presented (e.g., stimulus  $B$  is perceived as being the same whether it is presented with  $A$  or with  $C$ ); on the contrary, the configural model assumes, for reasons that will become clear in the next section, that each nominal stimulus is not functionally the same when presented in different discriminations.

Of particular importance is the fact that a model based only on configural values will not be able to solve any transitive test, because the context at test is unknown (viz., the two test stimuli were never seen together). Thus, the model retains what Wynne (1995) called the elemental stimulus values from the Rescorla–Wagner model and adds configural values weighted by a new parameter,  $\gamma$ . When this parameter is low, the model behaves as the Rescorla–Wagner model but when the value is high, the model is able to solve any training pair and loses its ability to solve any test pairs (Wynne, 1998).

The elemental stimulus values [what we formerly called  $V(X)$  and  $V(Y)$ ] are just like in the Rescorla–Wagner model updated on a trial-by-trial basis using Eqs. (2) and (3). The configural values [ $V(X|XY)$  and  $V(Y|XY)$ ] are also updated on a trial-by-trial basis according to the following equations:

$$V(X|XY)_{i+1} = V(X|XY)_i + \beta * [1 - V(X|XY)_i] \quad (6)$$

$$V(Y|XY)_{i+1} = V(Y|XY)_i - \beta * V(Y|XY)_i \quad (7)$$

after reinforced and non-reinforced trials, respectively. Because the configural values are equal to zero on any test trial, the probability of choosing a particular stimulus will be given by the previously presented Eqs. (4) and (5). During premise training, however, choice is determined by the modified Luce's (1959) rule:

$$r = \frac{V(X) + V(Z) + \gamma V(X|XY)}{V(X) + V(Y) + 2V(Z) + \gamma[V(X|XY) + V(Y|XY)]} \quad (8)$$

To derive response probabilities, Eq. (5) is used again. For the same reasons pointed out for the Rescorla–Wagner model, the configural model predicts both the serial position effect and the SDE. The stimulus values, both elemental and configural, are also highly dependent on the exact sequence of trials and choices during training.

#### 4.3.3. Siemann–Delius (or eta-kappa) model

This model is simply a modification of Luce’s (1959) beta learning operator. Luce’s model will be introduced first followed by the eta-kappa ( $\varepsilon\kappa$ ) modification.

Luce’s model proposes that given a simultaneous discrimination  $X+Y-$ , each stimulus value updates, on a trial-by-trial basis, according to Eqs. (9) and (10) after reinforcement and non-reinforcement, respectively. In these equations,  $\beta_+$  is a learning parameter corresponding to reinforcement and  $\beta_-$  is the corresponding parameter for non-reinforcement. Hence, in any given  $X+Y-$  discrimination, the probability of choosing  $X$  is given by Eq. (11). With these three equations, the model is able to simulate individual performances in a TI task.

$$V(X)_{i+1} = V(X)_i + \beta_+ V(X)_i \quad (9)$$

$$V(Y)_{i+1} = V(Y)_i - \beta_- V(Y)_i \quad (10)$$

$$P(X|XY) = \frac{V(X)}{V(X) + V(Y)} \quad (11)$$

The model can also be applied to a population of subjects. Suppose that a functionally infinite number of subjects participates in any given trial. A proportion of these subjects,  $[P(X|XY)]$ , will choose  $X$  and the remaining ones,  $[P(Y|XY) = 1 - P(X|XY)]$ , will choose  $Y$ . The population average updates are then given by Eqs. (12) and (13).

$$V(X)_{i+1} = V(X)_i + \beta_+ V(X)_i P(X|XY) \quad (12)$$

$$V(Y)_{i+1} = V(Y)_i - \beta_- V(Y)_i P(Y|XY) \quad (13)$$

This particular instantiation of Luce’s model predicts transitive performance on any  $n$ -term series task as well as the serial position effect and the SDE (Siemann and Delius, 1998). Still, it has difficulties dealing with the situations in which, after successful premise acquisition, the animal fails to behave transitively—the transitive choice is a necessary mathematical consequence of premise acquisition. In fact, as soon as premise performance is above 50%, the model predicts transitive choices on any non-adjacent transitive pair. To overcome these difficulties, Delius and Siemann (1998; see also Siemann and Delius, 1998) proposed the  $\varepsilon\kappa$  modification of Luce’s model. As the reader will notice, the main assumption is common to Wynne’s configural model.

Each stimulus is assumed to have not one but two values: the elemental and the configural value. The elemental values are updated according to Eqs. (14) and (15), after reinforcement and non-reinforcement, respectively.  $\varepsilon$  is a parameter determining

the weight of the elemental value changes.

$$V(X)_{i+1} = V(X)_i + \beta_+ V(X)_i P(X|XY)\varepsilon \quad (14)$$

$$V(Y)_{i+1} = V(Y)_i - \beta_- V(Y)_i P(Y|XY)\varepsilon \quad (15)$$

Likewise, the configural values are updated according to Eqs. (16) and (17), where  $\kappa = 1 - \varepsilon$ .

$$V(X|XY)_{i+1} = V(X|XY)_i + \beta_+ V(X|XY)_i P(X|XY)\kappa \quad (16)$$

$$V(Y|XY)_{i+1} = V(Y|XY)_i - \beta_- V(Y|XY)_i P(Y|XY)\kappa \quad (17)$$

Just like in the configural model, both elemental and configural values are updated on a trial-by-trial basis after each premise training trial. The choice probability on each of these trials is given by Eq. (18).

$$P(X|XY) = \frac{V(X) * V(X|XY)}{V(X) * V(X|XY) + V(Y) * V(Y|XY)} \quad (18)$$

During non-adjacent test trials, however, configural values are not considered and the choice probability equation reduces to Eq. (11).

Just like Wynne’s configural model, the  $\varepsilon\kappa$  model exhibits a tradeoff between premise and test pair accuracy (for an illustrative depiction see Siemann and Delius, 1998, Fig. 5). When  $\varepsilon = 0$ , training performance is perfect but no TI is predicted. As  $\varepsilon$  increases, TI predictions emerge without much deterioration of premise accuracy. Finally, when  $\varepsilon = 1$ , the model behaves exactly like Luce’s model. This algebraic learning model has also been translated into a neural network model (see Siemann and Delius, 1998; for another neural network model, see De Lillo et al., 2001).

#### 4.3.4. Evaluating the reinforcement-based models

Both the Rescorla–Wagner and Luce’s model can be considered precursors of the more elaborated and, to some extent, better fitting configural and  $\varepsilon\kappa$  models, respectively. Nonetheless, the Rescorla–Wagner model can still predict practically all of the signature TI effects.

Wynne (1995) exposed the Rescorla–Wagner model and his configural model to the exact sequence of 27,000 trials reported in von Fersen et al. (1991). With an intermixed 5-term series task (Experiment 1), both models produced a reasonably good fit to the observed premise performance, both before and after an intermediate phase of sequential training. When the series was extended to seven terms (Experiment 2), the Rescorla–Wagner model produced a good fit in the end pair FG but underestimated premise performance on the remaining ones (in particular, it predicted indifference on CD, DE, and EF). On the other hand, the configural model predicted better performance on the central pairs allowing for a much better fit than the Rescorla–Wagner model. Yet, both models correctly predicted the serial position effect. The most striking difference between the models emerged, however, when the 7-term series was closed into a circular series (Experiment 3). Despite the difficulty of this task, some of von Fersen et al.’s pigeons were able to perform above chance in all premises. This particular feature was only captured by the configural model; the Rescorla–Wagner model predicted

chance performance on all premises. The models are, nonetheless, very similar: Both produced reasonably good fits to the pigeon data, and also predicted the observed serial position effect and the SDE.

In a different experiment, Wynne (1997) used the parameters that best fit the results of von Fersen et al. to predict his pigeons' performance. Again, he found that both models accurately accounted both for training and testing performance, although the Rescorla–Wagner model produced a better fit overall.

Of particular interest is whether these models accurately predict transitive performance following sequential training. In fact, they do. Wynne (1995, 1998) showed that both models correctly account for the transitive behavior exhibited by Steirn et al.'s (1995) pigeons on the BD test pair, both with forward and backward training. Moreover, these fittings were obtained using the same parameter values ( $\alpha$ ,  $\beta$ , and  $\gamma$ ) used to fit von Fersen et al.'s (1991) results. Taken together, this seems to suggest that TI after intermixed and sequential training are indeed based on the same underlying mechanism(s).

Luce's model has also been used to simulate published human data. For example, Siemann and Delius (1993) showed that the model accurately fits both training and test performances of human subjects in a 6-term series task using polygons as stimuli. As previously stated, the disadvantage of this model is that it always predicts TI after the premises are learned—a prediction that is not always borne out in some animals. The  $\epsilon\kappa$  model (as well as Wynne's configural model), however, accommodates this empirical finding by assuming that premise training sometimes yields predominantly configural learning which should result in non-transitive performance despite accurate premise performance. Additionally, the model also predicts above-chance premise performance when a series is closed into a loop. In particular, Siemann and Delius (1998) showed that the  $\epsilon\kappa$  model accounts for the above-chance performance exhibited by von Fersen et al.'s (1991) pigeons in a circular series.

Overall, the models that incorporate both elemental and configural values (viz., the configural and the  $\epsilon\kappa$  models) accurately account both for the signature effects of TI in non-human animals and the typical findings with circular series. Still, in some instances, their predictions fail to fit the observed preferences on transitive tests. Lazareva et al. (2004) and Lazareva and Wasserman (2006) provided such an example.

Lazareva et al. (2004) trained crows in a 5-term series task using ordered or constant post-choice feedback (see Table 1). After acquisition, however, instead of immediately proceeding to testing, Lazareva and colleagues used the Reward/Non-reward ratios to estimate the associative strengths of both  $B$  and  $D$ . Birds for which the  $D$  ratio exceeded the  $B$  ratio proceeded to testing, but birds for which the  $D$  ratio was smaller than the  $B$  ratio were exposed to massed presentations of the DE premise (what the authors called a reversal phase). The rationale for this procedure was as follows: If TI in non-human animals is indeed a property of their reinforcement histories, then by making the  $D$  ratio greater than the  $B$  ratio the animals should prefer the *anti-transitive* alternative ( $D$ ) during testing (for a similar, but unsuccessful attempt, see Siemann et al., 1996b). In fact, the

constant feedback group was, on average, indifferent and the ordered feedback group strongly preferred  $B$ . Simulations using the configural and the  $\epsilon\kappa$  models were then run using the exact sequence of trials and choices experienced by each bird. Both models provided a very good fit for premise accuracy in both groups but failed to account for the observed preference for  $B$  over  $D$  in the ordered feedback group. Actually, both models predicted the mean preference in the transitive BD tests to be at or below chance for both groups.<sup>5</sup>

Lazareva and Wasserman (2006) replicated Lazareva et al.'s (2004) experiment with pigeons (see Table 1). The results differed somewhat from Lazareva et al.'s findings with crows (viz., both the constant and the ordered feedback group behaved transitively), but again both configural-cue models failed to fit the test results despite accurately fitting premise performance. Post hoc simulations also revealed a rather curious result: The models failed to account for the test results only when the entire sequence of events experienced by the birds was entered into the simulation. When the reversal phase was excluded, the models accurately predicted performance on the non-adjacent, transitive and non-transitive test pairs in both groups. Why would this happen? To understand this, one needs to consider the effects of the reversal phase on the associative strengths of each stimulus.

Both models assume that the elemental value of each stimulus is updated whenever that particular stimulus is presented. The configural value, on the other hand, is updated only when the particular stimulus is presented in a particular pair. During reversal training (massed presentations of pair DE), the elemental value of  $D$  increases, whereas the elemental value of  $B$  is left unchanged. By the end of reversal, the elemental value of  $D$  is usually near or actually exceeds the elemental value of  $B$ , and that is why the models predict indifference or a preference for  $D$  over  $B$ . To be able to account for the actual results ( $B$  preferred over  $D$ ), one would need to assume that the reversal phase leads to a predominantly configural update. The models would, then, account for Lazareva et al.'s (2004) and Lazareva and Wasserman's (2006) results because predictions for test performances are independent of configural values. Of course, why massive presentations of one training pair would lead to a predominantly configural update is unspecified, although this is not a novel proposal (e.g., Siemann and Delius, 1998). For instance, some findings suggests that animals use configural cues even when a task can be solved based solely on individual stimuli (e.g., Alvarado and Rudy, 1992; Wynne, 1996).

Be that as it may, recent neurobiological findings seem to support the existence of both elemental and configural values. Indeed, Frank et al. (2003) using a computational neural network model of the hippocampus and neocortex (O'Reilly and Rudy, 2001) found that the anchor-problems (AB and EF, in a 6-term series task) and the transitive tests can be solved based solely on elemental values, but that the inner problems require hippocampus-dependent "conjunctive representations" of the

<sup>5</sup> Simulations with modified versions of the configural and the  $\epsilon\kappa$  models incorporating a value transfer mechanism were also run. Good fits were found for training accuracy but not for testing.



pairs<sup>6</sup> (see also Greene et al., 2006). The model accurately fits the SDE observed in Van Elzakker et al. (2003) based on learning mechanisms similar to the ones described for the reinforcement-based models and a hippocampus-generated blocking effect (Kamin, 1968) that arises from the use of the Rescorla–Wagner rule in updating synaptic weights (for further details, see Frank et al., 2003; for another biologically motivated neural network, see Levy and Wu, 1997; Wu and Levy, 2001).

The previously mentioned Frank et al.'s (2006) study with a transient deactivation of the hippocampus by the drug midazolam underscores the importance of both configural and elemental cues: Consistent with the idea that performance on transitive tests depends only on elemental values, midazolam improved transitive performance despite its detrimental effects on explicit memory. Moreover, early in training Frank and colleagues found better performance on the anchor pairs of the 5-term series task than on the inner pairs suggesting that the transient deactivation of the hippocampus hinders the development of configural cues (see also Alvarado and Rudy, 1995). Elemental values and their differentiation, on the other hand, seem to depend on the basal ganglia–dopamine system (Frank et al., 2003; Frank et al., 2004).

In summary, the reinforcement-based models, particularly the ones including configural cues, are able to account for all the TI signature effects in non-human animals: preference for the transitive alternative, the SDE, and the serial position effect. They also accurately fit the above-chance training performance observed with circular series and are consistent with some findings in the neurobiology domain (but see Buckmaster et al., 2004; Dusek and Eichenbaum, 1997; Eichenbaum, 2000, 2004; Wallenstein et al., 1998, for an alternative interpretation of the hippocampus' function in TI tasks). For the configural model, these predictions hold independently of the intermixed or sequential nature of training and, in some instances, the parameters found with intermixed training accurately predict transitive performance with sequential training. The exceptions are the simulations run by Lazareva et al. (2004) and Lazareva and Wasserman (2006). As seen, the models do not cope correctly with reversal phases. Perhaps the abovementioned predominant configural updating is able to solve this problem but, at the moment, this proposal is speculative. Also, from the models' perspective, it is not entirely clear how or if these post-choice feedback stimuli should also have associative strengths updated on a trial-by-trial basis (the fact that the models can predict the birds' performance before the reversal phase in Lazareva and Wasserman suggests not, but the hypothetical existence of these values does not automatically preclude the models from continuing to do so).

#### 4.4. Summing up the models

As Killeen (1999) puts it, “models are go-betweens: They go between the data and our sense of understanding” (p. 276). They must predict the data and, at the same time, be conceptually

coherent and robust. In a sense, modeling is always a tradeoff between what must be predicted and what is theoretically plausible: Sometimes accounting for an additional detail in the data means another free parameter; sometimes strict theoretical positions prevent the acceptance of that additional parameter. The “game” must be played without losing contact either with the data or with our sense of understanding, to use Killeen's words. Do all models of TI in non-human animals have this quality? The answer is “No”.

Different models propose different underpinnings to the TI ability, from mental lines and logical rules to value transfer and reinforcement-driven mechanisms. Yet, none of them is in perfect agreement with the data. As seen, some fail to account for basic phenomena, some fail critical tests and still others fail less relevant tests, although all of them make contact, at least to some extent, with the most striking features of the data.

The cognitive models have serious difficulties dealing with sequential training and circular series. Also, the way they incorporate the serial position effect, the SDE and error in the typical TI tests seriously undermines their theoretical robustness. VTT suffers from similar problems: It does not cope well with sequential training and with circular series. Also, it fails critical tests suggested by its own assumptions (namely, it fails the test of partially reinforcing *A* in the AB premise). Finally, the reinforcement-based models, particularly the ones including configural cues, do not cope well with reversal phases.

Conceding that all models have drawbacks, the question is whether these drawbacks are of equal relevance. At the top of the pyramid are the three essential effects that must be predicted by any model: the serial position effect, the SDE and the transitive preference. Models failing to predict such effects can be safely discarded. Clearly, VTT and all reinforcement-based models are up to this task but some doubts remain about the cognitive models. This fact *per se* should be sufficient to reject such models. Left with VTT and the reinforcement-based models, an analysis of their predictions for sequential training and circular series undoubtedly favors the latter. To sum up, then, reinforcement-based models, in particular the configural and the  $\epsilon\kappa$  models, seem to be much more successful in making contact with the data.

Still, it is when we look at the models in terms of their plausibility that most fundamental problems arise, in particular for the cognitive models. As Wynne (1998) puts it, the fundamental problem cognitive models face is their own ontological status. It is simply unclear what such mental representations or logical rules are and how they are formed.

Models should be about processes, but no processes are specified in most cognitive models. Somehow, animals end up with some set of conveniently ordered logical rules or with some mental line that supposedly underlies TI. Some would argue that the process is obviously the experience of reinforcement and non-reinforcement. But if that is the case, why not stop there? Why do we need to add some set of rules or some mental line (or whatever mentalistic paraphernalia)? At the moment we focus our attention in the history of reinforcement and non-reinforcement we eliminate the need for such devices. It is as if we were adding

<sup>6</sup> In Frank et al.'s (2003) model, the “conjunctive representations” are supposedly explicit. I am not claiming this is the case in non-human animals.

another level of complexity when the previous level already does the job.

This taste for complexity seems to have arisen from the cognitive models' original purpose—to explain TI in humans. Some recent reports suggest, however, that such complexity might not even be needed for humans. In particular, successful transitive performance and awareness of the hierarchical relationship among the stimuli seem to be dissociable. For example, [Greene et al. \(2001\)](#) trained two groups of human subjects in a fully non-verbal 5-term series task. One group of subjects was informed that the stimuli could be ordered (the informed group) whereas the other group was told that they should try to learn the task by trial-and-error (the uninformed group). When given the BD test pair, both groups preferred *B*, although the informed participants were usually better. Most importantly, in the uninformed group, transitive performance did not correlate with awareness of the hierarchical relationship among the stimuli (as indicated by a post-experimental questionnaire). Furthermore, in a different experiment with only uninformed participants, [Greene and colleagues](#) interrupted the transitive tests whenever performance on the BD pairs reached asymptote and assessed awareness of the hierarchy at that point. The findings were twofold: First, the awareness levels were reduced compared to situations in which testing proceeded without interruption and, second, transitive performance was indistinguishable from previous experiments in which testing was not interrupted. [Frank et al.'s \(2006\)](#) study with midazolam-induced amnesia in which inactivation of the hippocampus caused severe explicit memory deficits but enhanced transitive performance confirmed these findings. In short, successful performance does not seem to depend on explicit knowledge of the underlying hierarchy (see also [Frank et al., 2005](#); [Greene et al., 2006](#); [Libben and Titone, 2008](#); [Siemann and Delius, 1993, 1996](#)). The simple observation that, under certain circumstances, humans do exhibit TI without conscious awareness of the hierarchy poses difficulties to any model based on logic or mental representations ([Martin and Alsop, 2004](#); [Moses et al., 2006](#); [Smith and Squire, 2005](#)).

Another way to look at the models is in terms of the mathematical precision of their predictions. Here, indeed, an order can be found. Cognitive models and VTT make only ordinal predictions. Reinforcement-based models, on the other hand, pinpoint exact preference values. It should not be taken lightly that the best-fitting predictions come from these latter models in which mismatches would certainly be most noticeable.

## 5. Directions for future research

Currently, both empirically and theoretically oriented questions remain unanswered. For instance, it has been implicitly assumed by most researchers that the transitive test after intermixed and sequential training taps the same TI-like ability. This assumption is based on the fact that such a test usually yields similar preferences. In addition, the ability of reinforcement-based, configural-cue models to predict animals' performance both with intermixed and sequential training strengthened the assumption.




Fig. 6. A procedure to study the viability of value transfer with forward sequential training.

However, of the three signature effects, only the preference for the transitive alternative has been demonstrated with sequential training. Although the serial position effect is not expected to emerge with the sequential regimen due to the ordered nature of training (but see [Benard and Giurfa, 2004](#)), the SDE should clearly be found. This empirical shortcoming could be easily addressed by sequentially training animals in a 7-term series task and then testing them in a manner that avoids the typical confounds found in this type of analysis (see p. 7).

Also, the reinforcement-based models suggest that TI may be observed at relatively low levels of premise performance provided the test stimuli have acquired differential value. Despite this prediction, the relation between baseline (premise) performance and TI is not empirically clear given that TI has been evaluated only at steady-state, when animals exhibit asymptotic levels of premise performance. By inserting infrequent probe trials during premise training, a clearer depiction of TI acquisition would certainly emerge.

Some other questions related to the models of TI in non-human animals are unanswered, too. For example, research suggests that value transfer is a real mechanism that operates in simultaneous discriminations (e.g., [Siemann et al., 1996a](#); [Urcuioli, 2006](#); [Urcuioli and Michalek, 2007](#); [Zentall, 2004](#); [Zentall and Sherburne, 1994](#); [Zentall et al., 1996a](#)). But although value transfer may not be necessary to produce transitive choices ([Steirn et al., 1995](#)), differential value transfer may very well contribute to the effect when present (see [Lazareva et al., 2004](#)). However, its contribution to sequential training is questionable (although open to empirical inquiry).

For example, suppose animals are sequentially trained on three simultaneous discriminations (corresponding to forward sequential training in the typical *n*-term series task). [Fig. 6](#) shows the details. Consider, first, Group 1. In the first two phases, the animal is sequentially trained on two simultaneous discriminations (PA and AN). This is just a sequential implementation of the ambiguous-cue problem procedure. *P* is the positive stimulus because it is always reinforced, *A* is the ambiguous stimulus because it signals reinforcement or non-reinforcement depending on the other stimulus with which it is presented, and *N* is the negative stimulus because it is never reinforced. In phase 3, a new discrimination is trained in which *X* is reinforced and *Y* is not. During testing, non-differentially reinforced *N* vs. *Y* test

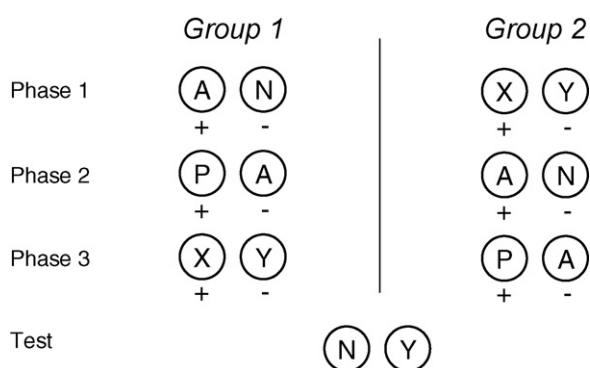


Fig. 7. A procedure to study the viability of value transfer with backward sequential training.

trials are presented. Group 2 just counterbalances the training order.

In the absence of value transfer neither the *N* nor the *Y* test stimulus should be preferred, although, in Group 1, animals might avoid *Y* and, in Group 2, animals might avoid *N* because those stimuli were the most recently non-reinforced. On average, however, the avoidance-based preferences would cancel each other out. But, if both groups exhibit a preference for *Y*, this would support a value transfer interpretation. VTT predicts this outcome if we assume that a stimulus non-reinforced in one phase of training and reinforced in another phase is, functionally, a partially reinforced stimulus. According to this interpretation, stimulus *A* in Fig. 6 is a partially reinforced stimulus and so transfers a smaller amount of value to *N* than *X* (an always reinforced stimulus) transfers to *Y*.

A similar experiment could be designed with backward sequential training (see Fig. 7). The only difference here is that the training order of the *PA* and *AN* discriminations is reversed. Again, in the absence of value transfer, preferences based on the most recently avoided stimulus might be observed. However, a preference for *Y* in both groups would support value transfer. Again, one would have to consider *A* a partially reinforced stimulus with the additional assumption that the amount of value that *A* transfers to *N* is some proportion of its final overall value and not its value when presented with *N*.

Finally, the ability of the reinforcement-based models to cope with reversal phases needs to be clarified, too. Whenever the non-transitive alternative has a richer reinforcement history than the transitive one, these models predict (in most cases) an *anti-transitive* preference. Up to now, this prediction has only been tested with procedures involving post-choice feedback stimuli, and the results failed to support it. However, this interpretation is clouded by the presence of the post-choice stimuli given that these models are silent about such stimuli. Thus, it is not clear whether these stimuli have some *value* of their own updated on a trial-by-trial basis or whether they influence the value of the premise stimuli. A less ambiguous test would involve training animals in a regular *n*-term series task (without post-choice feedback stimuli) with a reversal phase prior to testing. The results of this hypothetical experiment are crucial to evaluate the feasibility of the mechanisms proposed by the reinforcement-based models.

## 6. Concluding remarks

TI has been demonstrated in a variety of different species. Using series of different lengths, different reinforcement contingencies, different sensory modalities, different training regimens, etc., the published evidence establishes TI in non-human animals as a reliable phenomenon. Furthermore, training and test data obtained with non-human animals closely parallels the typical findings obtained with humans: preference for the transitive alternative, the serial position effect, and the SDE (although further research is needed on the latter to circumvent typical confounds).

To explain TI in non-human animals, psychologists have appealed to a variety of concepts and mechanisms. Among them are mental lines or spatial representations, propositional calculus, value transfer, and the history of reinforcement and non-reinforcement. These models are empirically distinguishable and, in general, conceptually incompatible. Empirically, they vary in the range of phenomena for which they account; conceptually, they vary in the degree of consistency and the number of *ad hoc* elements and mechanisms they must incorporate to account for the TI signature effects. Altogether, the reinforcement-based models have stronger empirical support and are conceptually much more parsimonious and robust than the proposed alternatives.

Up to now, the study of TI in non-human animals has relied almost exclusively on the *n*-term series task which allows for an extremely informative comparative analysis of performance. Nonetheless, it is not clear if the mechanisms proposed to explain TI in the *n*-term series task are applicable to transitive-like behavior of animals in complex natural settings (Allen, 2006), for instance in rank estimation (Cheney and Seyfarth, 1986, 1990). One can easily foresee several difficulties. For example, in an *n*-term series task, the animal encounters  $n - 1$  pairs (usually four of five) during training. An animal living in a group of 50 animals, encounters 1225,  $\binom{50}{2}$ , possible pairs if we include its own interactions with the other members and do not consider the possibility of alliances. Non-human animals already have some difficulties learning the usual four or five premises of an *n*-term series task. Moreover, social hierarchies are not static and the ability to quickly reorder them would be adaptive (e.g., Allen, 2006; Seyfarth and Cheney, 2002).

At first glance then, naturalistic instantiations of TI seem less suited to a reinforcement-based account, although the other proposed theoretical alternatives face similar difficulties. Whether or not such difficulties will be overtaken remains to be seen. Nonetheless, the assumption that TI plays a pivotal role in such naturalistic interactions is still questionable. The possibility remains that animals may use some explicit or implicit fitness index which eliminates the need for any inferential mechanism.

For the moment, the available evidence strongly suggests that TI in non-human animals is a property of the animal's reinforcement history. The burden of proof lies, then, with those who assume that animals perform operations akin to propositional

calculus or that they order premises along some mental line in some undefined mental space.

In summary, the TI ability, the serial position effect, and the SDE should not be taken as *prima facie* evidence of complex seriation processes such as the spatial representation of premises along an imaginary mental line. Instead, they should remind us that apparently complex behavioral patterns can sometimes arise from associative mechanisms. In keeping with Lloyd Morgan's canon, simpler mechanisms should always be a standard against which competing alternatives are judged.

## Acknowledgement

I thank John Capaldi, Jacky Emmerton, and George Hollich for comments on a earlier version of the paper. I am particularly grateful to Peter Urciuoli for his encouragement, guidance, and many critical readings.

## References

- Acuna, B.D., Sanes, J.N., Donoghue, J.P., 2002. Cognitive mechanisms of transitive inference. *Experimental Brain Research* 146, 1–10.
- Allen, C., 2006. Transitive inference in animals: Reasoning or conditioned associations? In: Hurley, S., Nudds, M. (Eds.), *Rational Animals?* Oxford University Press, New York, pp. 175–185.
- Altman, S.A., 1962. A field study of the sociobiology of rhesus monkeys, *Macaca mulatta*. *Annals of the New York Academy of Sciences* 102, 338–435.
- Alvarado, M.C., Rudy, J.W., 1992. Some properties of configural learning: an investigation of the transverse-patterning problem. *Journal of Experimental Psychology: Animal Behavior Processes* 18, 145–153.
- Alvarado, M.C., Rudy, J.W., 1995. Rats with damage to the hippocampal formation are impaired on the transverse-patterning problem but not on elemental discriminations. *Behavioral Neuroscience* 109, 204–211.
- Anderson, J.R., 1993. *Rules of the Mind*. Erlbaum, Hillsdale, NJ.
- Benard, J., Giurfa, M., 2004. A test of transitive inferences in free-flying honeybees: unsuccessful performance due to memory constraints. *Learning and Memory* 11, 328–336.
- Bitterman, M.E., 1996. Comparative analysis of learning in honeybees. *Animal Learning & Behavior* 24, 123–141.
- Bond, A.B., Kamil, A.C., Balda, R.P., 2003. Social complexity and transitive inference in corvids. *Animal Behaviour* 65, 479–487.
- Boysen, S.T., Berntson, G.G., Shreyer, T.A., Quigley, K.S., 1993. Processing of ordinality and transitivity by chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology* 107, 208–215.
- Breslow, L., 1981. Reevaluation of the literature on the development of transitive inferences. *Psychological Bulletin* 89, 325–351.
- Bryant, P.E., Trabasso, T., 1971. Transitive inferences and memory in young children. *Nature* 232, 456–458.
- Bryson, J.J., Leong, J.C.S., 2006. Primate errors in transitive 'inference': a two-tier learning model. *Animal Cognition* 10, 1–15.
- Buckmaster, C.A., Eichenbaum, H., Amaral, D.G., Suzuki, W.A., Rapp, P.R., 2004. Entorhinal cortex lesions disrupt the relational organization of memory in monkeys. *Journal of Neuroscience* 24, 9811–9825.
- Burt, C., 1911. Experimental tests of higher mental processes and their relation to general intelligence. *Journal of Experimental Pedagogy* 1, 93–112.
- Burt, C., 1919a. The development of reasoning in school children: I. *Journal of Experimental Pedagogy* 5, 68–77.
- Burt, C., 1919b. The development of reasoning in school children: II. *Journal of Experimental Pedagogy* 5, 121–127.
- Bush, R.R., Mosteller, F., 1955. *Stochastic Models for Learning*. Wiley, New York.
- Byrne, R.M.J., Johnson-Laird, P.N., 1989. Spatial reasoning. *Journal of Memory and Language* 28, 564–575.
- Carreiras, M., Santamaria, C., 1997. Reasoning about relations: spatial and nonspatial problems. *Tinking and Reasoning* 3, 191–208.
- Cheney, D.L., Seyfarth, R.M., 1986. The recognition of social alliances by vervet monkeys. *Animal Behaviour* 24, 1722–1731.
- Cheney, D.L., Seyfarth, R.M., 1990. *How Monkeys See the World*. University of Chicago Press, Chicago, IL.
- Clement, T.S., Weaver, J.E., Sherburne, L.M., Zentall, T.R., 1998. Simultaneous discrimination learning in pigeons: value of S– affects the relative value of its associated S+. *Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology* 51B, 363–378.
- Cohen, J.S., Drummond, C., Terrelonge, N., 2001. Value transfer in simultaneous object discrimination by rats. *Animal Learning & Behavior* 29, 326–335.
- Colwill, R.M., Rescorla, R.A., 1985a. Instrumental responding remains sensitive to reinforcer devaluation after extensive training. *Journal of Experimental Psychology: Animal Behavior Processes* 11, 520–526.
- Colwill, R.M., Rescorla, R.A., 1985b. Post-conditioning devaluation of a reinforcer affects instrumental responding. *Journal of Experimental Psychology: Animal Behavior Processes* 11, 120–132.
- Colwill, R.M., Rescorla, R.A., 1990. Effect of reinforcer devaluation on discriminative control of instrumental behavior. *Journal of Experimental Psychology: Animal Behavior Processes* 16, 40–47.
- Couvillon, P.A., Bitterman, M.E., 1992. A conventional conditioning analysis of "transitive inference" in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes* 18, 308–310.
- Couvillon, P.A., Bitterman, M.E., 1996. Transverse patterning in pigeons. *Animal Learning & Behavior* 24, 410–422.
- Curran, T., DeBuse, C., Wroch, B., Hirshman, E., 2006. Combined pharmacological and electrophysiological dissociation of familiarity and recollection. *Journal of Neuroscience* 26, 1979–1985.
- D'Amato, M.R., 1991. Comparative cognition: processing of serial order and serial pattern. In: Dachowski, L., Flaherty, C.F. (Eds.), *Current Topics in Animal Learning: Brain, Emotion and Cognition*. Erlbaum, Hillsdale, NJ, pp. 165–185.
- D'Amato, M.R., Colombo, M., 1990. The symbolic distance effect in monkeys (*Cebus apella*). *Animal Learning & Behavior* 18, 133–140.
- Davis, H., 1992a. Logical transitivity in animals. In: Honig, W.K., Fetterman, J.G. (Eds.), *Cognitive Aspects of Stimulus Control*. Erlbaum, Hillsdale, NJ, pp. 405–429.
- Davis, H., 1992b. Transitive inference in rats (*Rattus norvegicus*). *Journal of Comparative Psychology* 106, 342–349.
- De Lillo, C., Floreano, D., Antinucci, F., 2001. Transitive choices by a simple, fully connected, backpropagation neural network: implications for the comparative study of transitive inference. *Animal Cognition* 4, 61–68.
- De Soto, C.B., London, M., Handel, S., 1965. Social reasoning and spatial paralogic. *Journal of Personality and Social Psychology* 2, 513–521.
- Delius, J.D., Siemann, M., 1998. Transitive responding in animals and humans: exaptation rather than adaptation? *Behavioural Processes* 42, 107–137.
- Dusek, J.A., Eichenbaum, H., 1997. The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences* 94, 7109–7114.
- Eichenbaum, H., 2000. A cortical-hippocampal system for declarative memory. *Nature Reviews Neuroscience* 1, 41–50.
- Eichenbaum, H., 2004. Hippocampus: cognitive processes and neural representations that underlie declarative memory. *Neuron* 44, 109–120.
- Fletcher, H.J., Grogg, T.M., Garske, J.P., 1968. Ambiguous-cue problem performance of children, retardates, and monkeys. *Journal of Comparative and Physiological Psychology* 66, 477–482.
- Frank, M.J., Rudy, J.W., O'Reilly, R.C., 2003. Transitivity, flexibility, conjunctive representations, and the hippocampus. II. A computational analysis. *Hippocampus* 13, 341–354.
- Frank, M.J., Seeberger, L.C., O'Reilly, R.C., 2004. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science* 306, 1940–1943.
- Frank, M.J., Rudy, J.W., Levy, W.B., O'Reilly, R.C., 2005. When logic fails: implicit transitive inference in humans. *Memory & Cognition* 33, 742–750.

- Frank, M.J., O'Reilly, R.C., Curran, T., 2006. When memory fails, intuition reigns: midazolam enhances implicit inference in humans. *Psychological Science* 17, 700–707.
- Gallistel, C.R., 1990. *The Organization of Learning*. Bradford Books/MIT Press, Cambridge, MA.
- Gillan, D.J., 1981. Reasoning in the chimpanzee: II. Transitive inference. *Journal of Experimental Psychology: Animal Behavior Processes* 7, 150–164.
- Greene, A.J., Spellman, B., Dusek, J.A., Eichenbaum, H.B., Levy, W.B., 2001. Relational learning with and without awareness: transitive inference using nonverbal stimuli in humans. *Memory & Cognition* 29, 893–902.
- Greene, A.J., Gross, W.L., Elsingher, C.L., Rao, S.M., 2006. An fMRI analysis of the human hippocampus: inference, context, and task awareness. *Journal of Cognitive Neuroscience* 18, 1156–1173.
- Grosenick, L., Clement, T.S., Fernald, R.D., 2007. Fish can infer social rank by observation alone. *Nature* 445, 429–432.
- Halford, G.S., 1984. Can young children integrate premises in transitivity and serial order task? *Cognitive Psychology* 16, 65–93.
- Hall, G., 1980. An investigation of ambiguous-cue learning in pigeons. *Animal Learning & Behavior* 8, 282–286.
- Harris, M., McGonigle, B., 1994. A model of transitive choice. *Quarterly Journal of Experimental Psychology* 47B, 319–348.
- Higa, J., Staddon, J., 1993. “Transitive inference” in multiple conditional discriminations. *Journal of the Experimental Analysis of Behavior* 59, 265–291.
- Hogue, M.E., Beaugrand, J.P., Lague, P.C., 1996. Coherent use of information by hens observing their former dominant defeating or being defeated by a stranger. *Behavioural Processes* 38, 241–252.
- Holland, J.H., Holyoak, K.J., Nisbett, R.E., Thagard, P., 1986. *Induction: Processes in Inference, Learning and Discovery*. MIT Press, Cambridge, MA.
- Huttenlocher, J., 1968. Constructing spatial images: a strategy in reasoning. *Psychological Review* 75, 550–560.
- Jarrard, L.E., 1993. On the role of the hippocampus in learning and memory in rats. *Behavioral and Neural Sciences* 60, 9–26.
- Johnson-Laird, P.N., Byrne, R.M.J., 1991. *Deduction*. Erlbaum, Hillsdale, NJ.
- Jolly, A., 1966. Lemur social behavior and primate intelligence. *Science* 153, 501–506.
- Kamin, L.J., 1968. “Attention-like” processes in classical conditioning. In: Jones, M.R. (Ed.), *Miami Symposium on the Prediction of Behavior: Aversive Stimulation*. University of Miami Press, Coral Gables, FL, pp. 9–32.
- Kesner, R.P., 1991. The role of the hippocampus within an attribute model of memory. *Hippocampus* 1, 279–282.
- Killeen, P.R., 1999. Modeling modeling. *Journal of the Experimental Analysis of Behavior* 71, 275–280.
- Kummer, H., Daston, L., Gigerenzer, G., Silk, J., 1997. The social intelligence hypothesis. In: Weingart, P., Richerson, P., Mitchell, S.D., Maasen, S. (Eds.), *Human by Nature: Between Biology and the Social Sciences*. Hillsdale, NJ, Erlbaum, pp. 157–179.
- Lazareva, O.F., Wasserman, E.A., 2006. Effect of stimulus orderability and reinforcement history on transitive responding in pigeons. *Behavioural Processes* 72, 161–172.
- Lazareva, O.F., Smirnova, A.A., Bagozkaja, M.S., Zorina, Z.A., Rayevsky, V.V., Wasserman, E.A., 2004. Transitive responding in hooded crows requires linearly ordered stimuli. *Journal of the Experimental Analysis of Behavior* 82, 1–19.
- Leary, R.W., 1958. The learning of ambiguous-cue problems by monkeys. *American Journal of Psychology* 71, 718–724.
- Levy, W.B., Wu, X., 1997. Simple, biologically motivated neural network solves the transitive inference problem. In: *Proceedings of the IEEE International Conference on Neural Networks*, 1, vol. I, pp. 368–371.
- Libben, M., Titone, D., 2008. The role of awareness and working memory in human transitive inference. *Behavioural Processes* 77, 43–54.
- Luce, R.D., 1959. *Individual Choice Behavior: a Theoretical Analysis*. Wiley, New York.
- Markovits, H., Dumas, C., 1992. Can pigeons really make transitive inferences? *Journal of Experimental Psychology: Animal Behavior Processes* 18, 311–312.
- Martin, N., Alsop, B., 2004. Transitive inference and awareness in humans. *Behavioural Processes* 67, 157–165.
- McGonigle, B.O., Chalmers, M., 1977. Are monkeys logical? *Nature* 267, 694–696.
- McGonigle, B.O., Chalmers, M., 1984. The selective impact of question form and input code on the symbolic distance effect in children. *Journal of Experimental Child Psychology* 37, 525–554.
- McGonigle, B.O., Chalmers, M., 1986. Representation and strategies during inference. In: Myers, T., Brown, K. (Eds.), *Reasoning and Discourse Processes*. Academic Press, New York, pp. 141–164.
- McGonigle, B.O., Chalmers, M., 1992. Monkeys are rational! *The Quarterly Journal of Experimental Psychology* 45B, 189–228.
- Menzel, R., Giurfa, M., 2001. Cognitive architecture of a mini-brain: the honeybee. *Trends in Cognitive Science* 5, 62–71.
- Moses, S.N., Villate, C., Ryan, J.D., 2006. An investigation of learning strategy supporting transitive inference performance in humans compared to other species. *Neuropsychologia* 44, 1370–1387.
- Newell, A., Simon, H.A., 1972. *Human Problem Solving*. Prentice-Hall, Englewood Cliffs, NJ.
- O'Keefe, J., Nadel, L., 1978. *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford, UK.
- O'Reilly, R.C., Rudy, J.W., 2001. Conjunctive representations in learning and memory: principles of cortical and hippocampal activation function. *Psychological Review* 108, 311–345.
- Paz-y-Miño, C., Bond, A.B., Kamil, A.C., Balda, R.P., 2004. Pinyon jays use transitive inference to predict social dominance. *Nature* 430, 778–781.
- Piaget, J., 1928. *Judgment and reasoning in the child*. Harcourt, Brace and Co., New York.
- Piaget, J., 1953. *Logic and Psychology*. Manchester University Press, Manchester.
- Piaget, J., 1955. *The Child's Construction of Reality*. Routedledge & Kegan Paul, London.
- Piaget, J., 1970. *Genetic Epistemology* (E. Duckworth, Trans.). Columbia University Press, New York.
- Rapp, P.R., Kansky, M.T., Eichenbaum, H., 1996. Learning and memory for hierarchical relationships in the monkey: effects of aging. *Behavioral Neuroscience* 110, 887–897.
- Rescorla, R.A., 1972. ‘Configural’ conditioning in discrete-trial bar pressing. *Journal of Comparative and Physiological Psychology* 79, 307–317.
- Rescorla, R.A., 1973. Evidence for unique ‘stimulus’ account of configural conditioning. *Journal of Comparative and Physiological Psychology* 85, 331–338.
- Richards, R.W., 1973. Performance of the pigeon on the ambiguous-cue problem. *Bulletin of the Psychonomic Society* 1, 445–446.
- Riley, C., Trabasso, T., 1974. Comparatives, logical structures and encoding in a transitive inference task. *Journal of Experimental Child Psychology* 17, 187–203.
- Ritchey, F., 1951. Dominance-subordination and territorial relationships in the common pigeon. *Physiological Zoology* 24, 167–176.
- Roberts, W.A., Phelps, M.T., 1994. Transitive inference in rats: a test of the spatial coding hypothesis. *Psychological Science* 5, 368–374.
- Russell, J., McCormack, T., Robinson, J., Lillis, G., 1996. Logical (versus associative) performance on transitive reasoning tasks by children: implications for the status of animals’ performance. *Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology* 49, 231–244.
- Seyfarth, R.M., Cheney, D.L., 2002. The structure of social knowledge in monkeys. In: Bekoff, M., Allen, C., Burghardt, G.M. (Eds.), *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*. MIT Press, Cambridge, MA, pp. 379–384.
- Shettleworth, S.J., 2004. Rank inferred by reason. *Nature* 430, 732–733.
- Siemann, M., Delius, J.D., 1993. Implicit deductive responding in humans. *Naturwissenschaften* 80, 364–366.
- Siemann, M., Delius, J.D., 1996. Influences of task concreteness upon transitive inference in humans. *Psychological Research* 59, 81–93.
- Siemann, M., Delius, J.D., 1998. Algebraic learning and neural network models for transitive and non-transitive responding. *European Journal of Cognitive Psychology* 10, 307–334.

- Siemann, M., Delius, J.D., Dombrowski, D., Daniel, S., 1996a. Value transfer in discriminative conditioning with pigeons. *Psychological Record* 46, 707–728.
- Siemann, M., Delius, J.D., Wright, A.A., 1996b. Transitive responding in pigeons: influences of stimulus frequency and reinforcement history. *Behavioural Processes* 37, 185–195.
- Smith, C., Squire, L.R., 2005. Declarative memory, awareness, and transitive inference. *Journal of Neuroscience* 25, 10138–10146.
- Steirn, J.N., Weaver, J.E., (May) 1995. Positive and negative value transfer in a transitive inference task with pigeons. Paper presented at the meeting of the Midwestern Psychological Association, Chicago.
- Steirn, J.N., Weaver, J.E., Zentall, T.R., 1995. Transitive inference in pigeons: simplified procedures and a test of value transfer theory. *Animal Learning & Behavior* 23, 76–82.
- Sternberg, R.J., 1980. Representation and process in linear syllogistic reasoning. *Journal of Experimental Psychology: General* 109, 119–159.
- Terrace, H.S., 1986. A nonverbal organism's knowledge of ordinal position in serial learning task. *Journal of Experimental Psychology: Animal Behavior Processes* 12, 203–214.
- Terrace, H.S., McGonigle, B., 1994. Memory and representation of serial order by children, monkeys, and pigeons. *Current Directions in Psychological Science* 3, 180–185.
- Thompson, R., 1953. Approach-avoidance in an ambivalent object discrimination problem. *Journal of Experimental Psychology* 45, 341–344.
- Thorndike, E.L., 1898. Animal intelligence: an experimental study of the association processes in animals. *Psychological Review Monographs* 2, No. 8.
- Trabasso, T., Riley, C.A., 1975. The construction and use of representations involving linear order. In: Solso, R.L. (Ed.), *Information Processing and Cognition*. The Loyola Symposium. Erlbaum, Hillsdale, NJ, pp. 381–410.
- Treichler, F., Van Tilburg, D., 1996. Concurrent conditional discrimination tests of transitive inference by macaque monkeys: list linking. *Journal of Experimental Psychology: Animal Behavior Processes* 22, 105–117.
- Urcuioli, P.J., 2006. When discrimination fails (or at least falters). *Journal of Experimental Psychology: Animal Behavior Processes* 32, 359–370.
- Urcuioli, P.J., Michalek, S., 2007. Value transfer contributes to ambiguous-cue discrimination learning. *Psychological Bulletin & Review* 14, 658–662.
- Van Elzakker, M., O'Reilly, R.C., Rudy, J.W., 2003. Transitivity, flexibility, conjunctive representations and the hippocampus. I. An empirical analysis. *Hippocampus* 13, 334–340.
- von Fersen, L., Wynne, C., Delius, J.D., Staddon, J., 1991. Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes* 17, 334–341.
- Wagner, A.R., Rescorla, R.A., 1972. Inhibition in Pavlovian conditioning: application of a theory. In: Boakes, R.A., Haliday, M.S. (Eds.), *Inhibition and Learning*. Academic Press, New York, pp. 301–336.
- Wallenstein, G.V., Eichenbaum, H., Hasselmo, M.E., 1998. The hippocampus as an associator of discontiguous events. *Trends in Neurosciences* 21, 317–323.
- Weaver, J.E., Steirn, J.N., Zentall, T.R., 1997. Transitive inference in pigeons: control for differential value transfer. *Psychonomic Bulletin & Review* 4, 113–117.
- Whitlow, J.W., Wagner, A.R., 1972. Negative patterning in classical conditioning: summation of response tendencies to isolable and configural components. *Psychonomic Science* 27, 299–301.
- Woocher, F.D., Glass, A.L., Holyoak, K.J., 1978. Position discriminability in linear orderings. *Memory & Cognition* 6, 165–173.
- Wright, A.A., Sands, S.F., 1981. A model of detection and decision processes during matching to sample by pigeons: performance with 88 different wavelengths in delayed and simultaneous matching tasks. *Journal of Experimental Psychology: Animal Behavior Processes* 7, 191–216.
- Wu, X., Levy, W.B., 2001. Simulating symbolic distance effects in the transitive inference problem. *Neurocomputing* 30–40, 1603–1610.
- Wynne, C.D.L., 1995. Reinforcement accounts for transitive inference performance. *Animal Learning & Behavior* 23, 207–217.
- Wynne, C.D.L., 1996. Transverse patterning in pigeons. *Behavioural Processes* 38, 119–130.
- Wynne, C.D.L., 1997. Pigeon transitive inference: tests of simple accounts of a complex performance. *Behavioural Processes* 39, 95–112.
- Wynne, C.D.L., 1998. A minimal model of transitive inference. In: Wynne, C.D.L., Staddon, J.E.R. (Eds.), *Models of Action: Mechanisms for Adaptive Behavior*. Erlbaum, Mahwah, NJ, pp. 269–307.
- Wynne, C.D.L., von Fersen, L., Staddon, J.E.R., 1992. Pigeons' inferences are transitive and the outcome of elementary conditioning principles: a response. *Journal of Experimental Psychology: Animal Behavior Processes* 18, 313–315.
- Zentall, T.R., 2004. Pavlovian processes in simultaneous discriminations. *International Journal of Comparative Psychology* 17, 185–202.
- Zentall, T.R., Clement, T.S., 2001. Simultaneous discrimination learning: stimulus interactions. *Animal Learning & Behavior* 29, 311–325.
- Zentall, T.R., Sherburne, L.M., 1994. Transfer of value from S+ to S– in a simultaneous discrimination. *Journal of Experimental Psychology: Animal Behavior Processes* 20, 176–183.
- Zentall, T.R., Sherburne, L.M., Roper, K.L., Kraemer, P.J., 1996a. Value transfer in a simultaneous discrimination appears to result from within-event Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes* 22, 68–75.
- Zentall, T.R., Weaver, J.E., Sherburne, L.M., 1996b. Value transfer in concurrent-schedule discriminations by pigeons. *Animal Learning & Behavior* 24, 401–409.