

Redundancy and Blocking in the Spatial Domain: A connectionist model

I.P.L. McLaren*

University of Cambridge

How can the observations of spatial blocking (Rodrigo, Chamizo, McLaren & Mackintosh, 1997) and cue redundancy (O'Keefe and Conway, 1978) be reconciled within the framework provided by an error-correcting, connectionist account of spatial navigation? I show that an implementation of McLaren's (1995) better beta model can serve this purpose, and examine some of the implications for spatial learning and memory.


In this paper I tackle an issue in spatial navigation. How can a system for navigation be devised so that it combines the ability to use cue combinations in a redundant fashion (O'Keefe and Conway, 1978) at the same time as allowing for the spatial equivalent of blocking (Rodrigo *et al.*, 1997). I start by briefly describing both sets of results before offering a connectionist implementation of the approach taken in McLaren (1995) as a partial solution to the problem.

O'Keefe and Conway's (1978) demonstration of redundancy in spatial navigation is illustrated in Figure 1. O'Keefe and Conway were able to show that there were some hippocampal cells (units), that they termed place cells, that fired when the animal was in a specific location on the maze used in training and testing. These regions were defined by the four landmarks placed around the maze, and O'Keefe and Conway made sure that this was the case by rotating the maze relative to extra-maze cues from time to time. The key result was that any two of these landmarks were sufficient for at least some of the place cells to fire when the animal entered the appropriate region.

This is a demonstration of redundancy because the coding of spatial location is not critically dependent on any given cue, either singly or in combination with other cues. In this case any permutation of two cues from four was sufficient to enable navigation, making the system serving as the basis for navigation robust to cue removal or alteration.

Rodrigo, Chamizo, McLaren, & Mackintosh (1997) were the first to demonstrate the spatial equivalent of blocking. The apparatus used is illustrated in Figure 2. This shows a swimming pool in a rectangular room, with four potential landmarks placed around the pool in the positions shown.

* Department of Experimental Psychology, Downing Street, Cambridge, CB2 3EB, UK.
E-mail address: iplm2@cus.cam.ac.uk



Remove all 4 cues: 6/8 units lose ability to discriminate place field

Remove any 2 cues: 3/8 units lose ability to discriminate place field

Conclusion: For some units, cues are used in combination to designate a given region such that any 2 cues will define it.




Figure 1. The top panel illustrates the training regime using four cues for this study. In the bottom 2 panels two different tests with two cues are shown. An example of the type of region on the maze that would lead to firing of a specific place cell is shown in black.

The basic design of the experiment was for the experimental animals (Blocking group) to be pre-trained for a number of trials with landmarks A, B and C in place, before being trained with all the landmarks in place (i.e. X added). A Control group was trained on landmarks A, B, C, X without pre-training (in later experiments pre-training on landmarks P, Q, R was used for the controls). The pattern of results obtained is shown below in Figure 3.




Figure 2. The swimming pool, platform (in black) and landmarks used in Rodrigo et al (1997).

Testing the animals with the cue combination ACX without the platform present produced better performance from the controls than from the experimental animals (lower score reflecting the larger amount of time spent in the target quadrant where the platform was on training trials). The experimental animals were, if anything, somewhat better than controls when tested on ABC, however, and as a result performance on ABC was significantly better than on ACX in the Blocking group. Thus it would appear that the pre-training to ABC has in some way blocked learning to the additional cue, X. The parallel with typical blocking results (e.g. Kamin, 1968) is obvious. In this type of study training A+ followed by AX+ reduces responding to X when compared to controls only trained on AX+ (or pre-trained on B+).

We are now in a position to appreciate the conundrum posed by the combination of spatial blocking in spatial learning with cue redundancy in navigation. Consider one class of explanation of the phenomenon of blocking, that which postulates cue competition in an error correcting system as in Rescorla and Wagner's (1972) influential model of Pavlovian conditioning. The reason why X may not have very much associative strength for the US on this explanation is that the competitor cue, A, has already developed near

asymptotic associative strength to the US and so reduced the error term governing learning to near zero. As a consequence learning to X is slow, and produces little by way of responding on test. If now we take a similar approach to the spatial blocking result of Rodrigo *et al* then we would need to say that learning to ABC had reduced the error score to near zero and that this blocked learning when X was added. The problem here is that when we take O'Keefe and Conway's demonstration of redundancy into account then a paradox arises. How can a system accommodate both blocking, allowing some cues to prevent learning to others, and redundancy, allowing any combination of (effective) cues to be equally effective? The type of error-correcting system that explains Kamin's results does not easily lend itself to this purpose. It operates by dividing up some associative strength amongst the cues used in the experiment. Such a system will tend to degrade in performance as (effective) cues are removed. We know that the cues in O'Keefe and Conway's experiment are all effective, but performance did not degrade (for some cells / units) as cues were deleted. Thus it would seem that there is a tension between the two sets of results that argues against any error-correcting account of spatial navigation.




Figure 3. Results from Rodrigo *et al* (1997); dependent variable = mean rank so lower scores are better. T1 is the test with ACX, T2 is the test with ABC. G1 is the Blocking group, G2 the Control group.

The account favoured by O'Keefe and Nadel (1978), which postulates some map-like representation by way of explanation of the results of O'Keefe and Conway's experiments, fares no better when we take the spatial blocking result into consideration. This is because the result to be expected on a "cognitive map" hypothesis is that pre-training to ABC should, if anything, facilitate learning to X when it is added as a cue. The pre-training will have

engaged the map-making process, and the addition of a novel cue should then trigger appropriate exploratory behaviour and lead to its rapid integration into the pre-existing map. As the exact opposite is observed experimentally, i.e. learning is slower to X in the pre-trained group, then the result is difficult to accommodate in terms of any map.

Model

The aim in what follows is to show how an error-correcting system for spatial navigation can be devised such that it reconciles the two data sets considered here. The example of a guidance system given here follows on from McLaren's (1995) generalisation of the Beta model (Zipser, 1986) to the two landmark case. The necessary result given in that paper is straightforward, and is illustrated in Figure 4. If the line joining two landmarks makes an angle \emptyset with the animal's current X axis (assuming it to be employing Cartesian co-ordinates for its egocentric frame of reference), and the X co-ordinate of one of the landmarks (A) is X_a , then the X co-ordinate of some goal location (the platform, P, in this case) in the current frame of reference is given by $X = U\cos\emptyset + V\sin\emptyset + X_a$. Similarly the Y co-ordinate is given by $Y = U\sin\emptyset - V\cos\emptyset + Y_a$, where U and V are constants that depend only on the goal location. This result applies wherever the animal is in the environment, and whatever its current egocentric frame of reference. There is one value for each of U and V which corresponds to a given goal location in a given environment. Thus, U and V act as a representation of the goal location when taken in conjunction with perceptual information derived from the appropriate landmarks, and this representation is rotation and translation invariant. The problem, then, is to specify how the animal learns the correct values for U and V at the goal location, and how it learns to link these with the correct landmarks. If it can do this, then it can always make its way back to the goal if the landmarks remain unchanged, and can use U and V to stand for the goal.

If we set up a simple connectionist module employing units with linear activation functions, we can make the activity of the X_1 unit a weighted combination of the activities of the units coding for $\cos\emptyset$, $\sin\emptyset$, and X_a , as illustrated in the enlargement at the top left of Figure 5. In effect, the network instantiates the equation for X given above: it merely needs to learn the appropriate weights U and V to give the correct result (the weight from X_a is fixed at 1). If the animal is at the goal then X_1 can be set to zero and we can use an error correcting learning algorithm (e.g. delta rule) to adjust U and V. The problem we now face, however, is that for a given set of values of X_a and \emptyset there are many values of U and V that satisfy the requirement that X_1 be zero. The solution is to allow the animal to 'look around' at the goal, thus varying X_a and \emptyset , and to keep running the error correction procedure with every new 'look'. Only the correct values of U and V will give X_1 as zero for all X_a and \emptyset combinations, and the network will gradually settle on these. Once U and V have been learnt at the goal location, the animal can be placed anywhere in the environment, and, as long as it can estimate X_a and \emptyset , the

network will give the correct X co-ordinate of the platform as the activity of X_1 . Since a similar scheme gives the Y co-ordinate, the animal will be able to generate the vector from its current position to the goal location. If, however, the animal has had insufficient time at the goal location to develop the correct U and V then its specification of X and Y will be subject to error.




Figure 4. A schematic illustration of the quantities involved in McLaren's (1995) model of spatial navigation. A, B and C are landmarks. X and Y denote the current x and y axes in the animals egocentric reference frame.

The model as shown in Figure 5 has X_1 represented as a unit whose activation carries the value for X, and then passes this to a final X Module. A viable scheme for representing this final value is to use a distributed pattern of activation across a set of units as shown in Figure 6.

The idea here is that it is the pattern of activation that represents the X value, not the activation of any one individual unit. This is done by 'tuning' each of the units so that it responds to a certain value, v , maximally and declines in response away from this value. To give a Gaussian distribution, a function of the form $a_v = Ae^{-k(x-v)^2}$ is used for this, where a_v is the activation of the unit at value v on the dimension, x is the X value, and A and k are

constants. The connection from X_1 to each of the units is taken to be of weight 1.

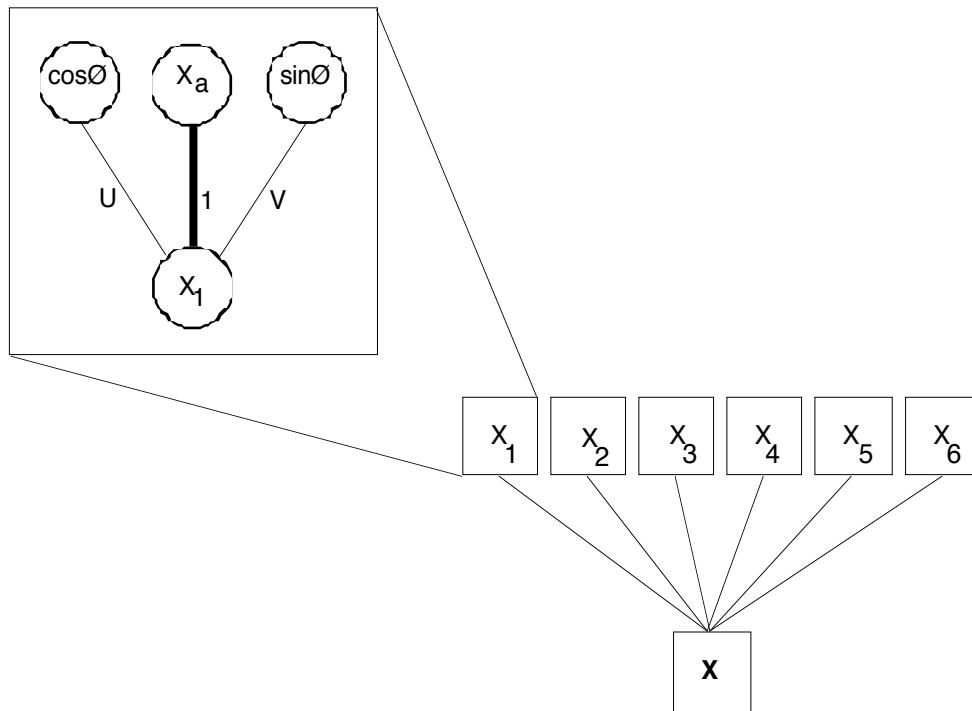


Figure 5. The basic architecture of the connectionist implementation of the Beta model. See text for details.

The point of this scheme is that it tackles the problem of generalising this approach to environments containing more than two cues. The answer adopted here is to assume that each pairwise combination of cues is dealt with separately by the type of network module just described, but that the results of these computations then combine to give the final \mathbf{X} (the lower right portion of Figure 5 shows how this might be done). But there is then a danger that fluctuations in the number of cues considered could lead to inappropriate variation in the output value. This would certainly be the case if we just summed the various outputs from the various X_i into one unit activation. But the problem does not arise if, instead, each pairwise combination of landmarks present and attended to will generate an X_i (which should be zero at the goal location) which is then passed on by fixed links to the final \mathbf{X} module (which controls behaviour) as constituted in Figure 6. This is because it is the pattern of activation over units that counts at this level of the system, and the overall effect is to give not only the right answer, but also one that is independent of the number of landmarks currently present / attended to in the environment, since variation in this number will only affect the total activation of the units,

not the distribution of activation across them. Thus, this system will be robust to cue removal, and will still be able to generate the correct goal location provided that at least two landmarks are available, though given any noise in the system its output will be subject to greater error when based on fewer landmarks.




Figure 6. A distributed pattern of activation representing a value on a dimension. This allows for a coding of X that is robust to variation in the number of landmarks used in the computation of X .

Learning, in this case error correction, is driven by the X module units whose pattern of activity is compared to the desired pattern corresponding to a value of 0 at the goal (shown in Figure 6) to give an error score. The maximum activation in the target pattern (A) is simply set to the maximum currently observed across these units, and the other target activations are scaled from this. The error score in this case is the sum of the squares of the differences between desired and actual activations. Error correction is implemented by taking the error of each X_i unit when at the goal to be this error score multiplied by -1 , ensuring that each unit's activation will (in the long run) be driven towards zero. This is because learning depends on error multiplied by activation, and as the error score as defined here will be always positive, then if X is positive the input to it will be decreased, but if it is negative it will be increased (and so X will become less negative).

Simulations

Using this system it is possible to train responding with an arbitrary number of landmarks, but the simulations reported here used one to four landmarks as in O'Keefe and Conway's (1978) experiment. When trained on four landmarks the system learned to generate a vector to the goal location very readily, and when then tested on any three or any two landmarks it succeeded in generating the same vector without significant degradation in performance. Thus it is able to show the kind of cue redundancy found by O'Keefe and Conway, though no claim is made here for any correspondence with the neural basis of those findings. The model considered here is purely computational in nature and does not seek to map onto specific brain mechanisms, only to explain behaviour.

We are now in a position to return to the issue of spatial blocking, and in doing so we find that the conundrum posed in the introduction to this paper undergoes the most remarkable transformation! The problem was that cue redundancy and spatial blocking might be incompatible features of a model of spatial navigation that used error-correction. The resolution offered here is that it is *because* of cue redundancy that our error correcting model will now deliver blocking. If a goal location had already been learned with three cues, and now a fourth is added, then the error signal would necessarily be relatively small, because of its dilution by the output from the many other x activations which will be at or near zero already, and learning would consequently be quite slow. In other words, the output pattern of activation will be near to the desired pattern corresponding to zero output, and so will give a low error score and produce slow learning. Consequently the U s and V s for the fourth landmark (in combination with the others) would take some time to develop. If, instead, all four cues were novel then the error term would tend to be larger, and learning would be rapid (at least initially). By the time the error term approaches the value for the blocking group the weights involving the fourth cue will be much nearer their correct values. Performance involving the fourth cue is thus predicted to be better in this group (given the right amount of training), and performance on the 'original' three cues may well be somewhat worse. In summary, blocking will depend on two main features of the model. First, that learning is error driven, with the error determined by all the landmarks present. Second, that performance is based on all the landmarks in a redundant fashion so that any subset of the landmarks will suffice.

Figure 7 shows some typical simulation results for the type of experiment considered in this paper. The Blocking group has been given extensive training with landmarks A,B,C prior to training with the complete set A,B,C,X. The Control group has been trained in exactly the same way on A,B,C,X after earlier training on P, Q, R (three other landmarks). The results are quite clear, performance (indexed by a score reflecting the error in X , with a higher score signifying worse performance) on A,C,X in the Blocking group is worse than that in the Control group, while performance on A,B,C shows some advantage for the animals pre-trained with this set of landmarks.

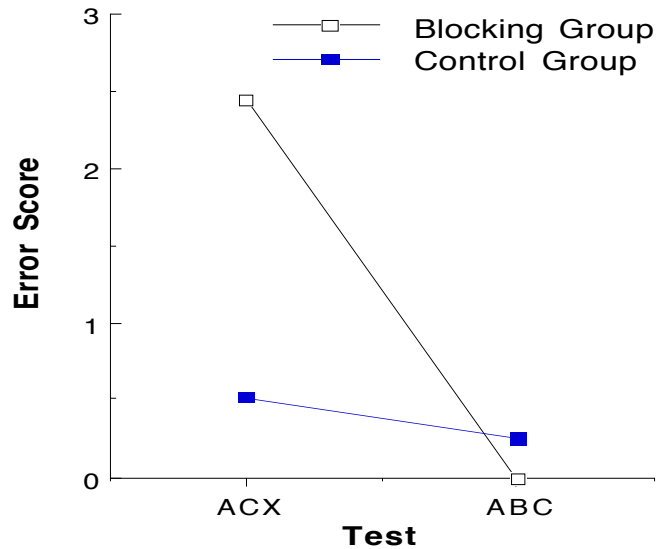


Figure 7. Simulation results for the spatial blocking preparation used by Rodrigo *et al* (1997).

Thus we have a model that can successfully show the required robustness in the face of cue deletion and demonstrate spatial blocking. I now consider some further implications of this approach to spatial navigation.

DISCUSSION

The most obvious implication that can be derived from the model with very little by way of analysis, is that learning to navigate in an environment will be a function of the experience with that environment. In particular, the sort of restricted viewing conditions employed by Mazmanian and Roberts (1983), whereby the animal cannot see all the landmarks when at the goal, and has only limited experience of the possible variation in appearance (angle, distance) of those it can see must be expected to lead to poor learning and poor performance. This was the result found by Mazmanian and Roberts. It can be explained in model terms by appealing to the need to sample the landmark combinations learned about from a variety of different orientations when at the goal. If the viewing angle is restricted the possible variation will be limited and learning will take longer to reach asymptote; indeed, if too little variation is possible, it may never reach it.

Another prediction of the model follows from the fact that only one of the landmarks need be proximal, in the sense that the distance to it can be assayed. If we let landmark B in Figure 4 recede into the far distance then the equations given still apply, but now B will serve simply as a directional cue.

Collett, Cartwright, and Smith (1986) provide evidence relevant to this issue derived from experiments with gerbils. The animals were trained in an open arena containing distal cues and more proximal landmarks placed in the arena. Animals searched in the appropriate spot when trained to find seeds at a constant distance and direction from a single cylindrical landmark. The landmark and seed were translated about the arena during training, and the animal's start point was also varied. This suggests that distance information is provided by the landmark, but that the distal cues provide the necessary orientation, as this will not be greatly affected by the movement of objects and animals within the arena. Note, however, that it is not simply the case of some distant cue(s) being used as a "beacon" to orient the animal. This would lead to errors when the landmark was moved. The distal cues must be used to compute some reference direction that is relatively independent of location. This could then be used in conjunction with the more proximal landmark to give a vector to the goal, as the model considered here would predict. The idea is shown in Figure 8.

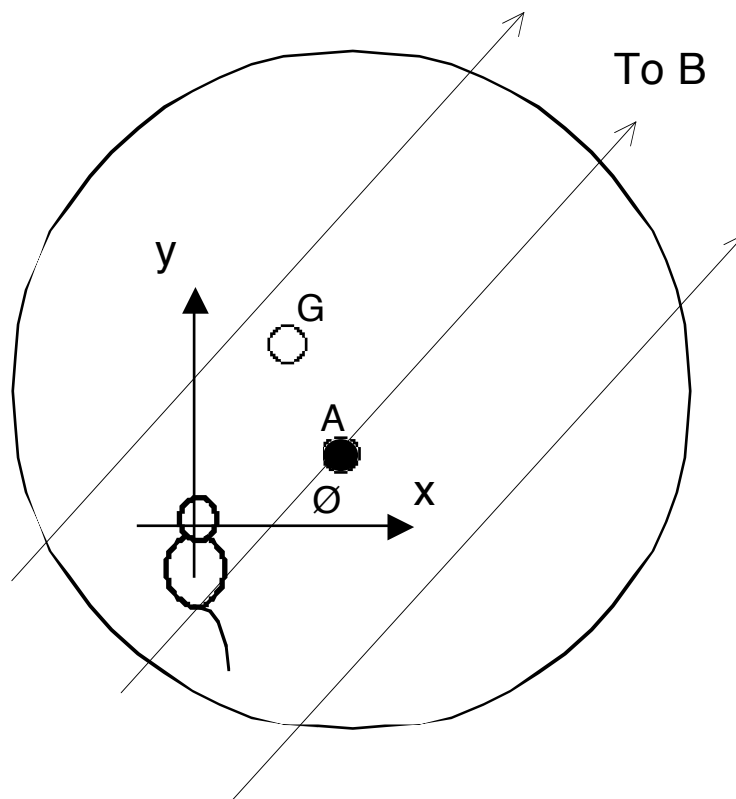


Figure 8. The cylindrical arena used by Collett *et al* (1986) with both proximal (A) and distal (B) cues. The model generalises well to this case as long as at least one cue is near enough for its distance to be estimated. The more distal cues give directional information (Ø).

Collett *et al* were also able to demonstrate that gerbils could encode goals in terms of proximal landmark arrays if they were the only reliable cue, i.e. if the array and goal were translated and rotated during training. They will use this configurational information even when the distal cues can be exploited, as they showed by performing an experiment that brings the two sets of cues into conflict. Gerbils trained to search for food in the centre of an equilateral triangle that was translated about the arena during training were then tested on the triangle rotated through sixty degrees. The result of this was that the animal searched in the centre of the triangle (correct according to array based information) and at three other locations just outside the triangles boundaries midway along its sides. The latter locations are those that correspond to partial matches between the vectors from the landmarks to the goal coded with respect to the reference direction provided by the distal cues. It may be that the model offered here is the only one that can accommodate both types of search result observed in this experiment.

A final consideration concerns studies of overshadowing in the spatial domain (e.g. Sanchez, Rodrigo, Chamizo & Mackintosh, 1999) and their relationship to the model outlined here. Overshadowing is often considered to be closely related to blocking as an empirical phenomenon, and both are often explained in terms of an error correcting learning algorithm. This is not the approach taken here, however, since the model developed in this paper will not predict overshadowing. Instead I propose to interpret the findings of studies of overshadowing in the spatial domain as relating more to the issue of landmark identification, an issue glossed over up to this point in our discussion of spatial learning. It is clearly a prerequisite of the model described here that landmarks can be accurately identified and linked to the appropriate coefficients if the animal is to navigate successfully. It would be disastrous to apply a set of coefficients to the wrong landmarks, as this would generate spurious goal information. I would expect the identification mechanism to differentiate well between similar landmarks, as confusion could lead to major errors in estimating goal location. With these considerations in mind, I note that studies that have demonstrated overshadowing in the spatial domain tend to use a manipulation of the landmarks that would fall foul of such a sensitive identification mechanism. In Sanchez *et al* (1999) one landmark had both visual and auditory components during training. On test one component was removed and this resulted in worse performance. An interpretation in terms of overshadowing is possible, but one in terms of generalisation decrement is also viable, and this would fit in with the explanation offered here. If the landmark had changed sufficiently that it was not considered to be the same landmark any more then it would no longer be of any use for navigation. Further studies which show that training involving a landmark with just a visual component say, before the addition of an auditory component on test does not result in a similar decline in performance are not necessarily a problem for this explanation. It may be that if a full feature match is obtained then the landmark identification mechanism is satisfied, and the presence of additional features is deemed irrelevant, or they are construed as another landmark. If that were so, then there is no reason to expect any generalisation decrement in such a case.

In conclusion, the model offered here is by no means complete. For example: it currently has little to say about landmark identification, or how that is used to engage appropriate learning to a given landmark. But it does show how a relatively sophisticated guidance system could, in principle, explain a number of basic phenomena with respect to spatial navigation. Given this, I hope that the reader will agree with me that it is an approach to modelling spatial navigation that warrants further development.

REFERENCES

- Collett, T.S., Cartwright, B.A. and Smith, B.A. 1986. Landmark learning and visuo-spatial memories in gerbils. *Journal of Comparative Physiology A*, 158, 835-851.
- Kamin, L.J. 1968. 'Attention-like' processes in classical conditioning. In M.R. Jones (Ed.) *Miami symposium on the prediction of behaviour: Aversive stimulation*. Univ. of Miami Press, 9-33
- Mazmanian, D.S. and Roberts, W.A. 1983. Spatial memory in rats under restricted viewing conditions. *Learning and Motivation*, 14, 123-139.
- McLaren, I.P.L. 1995. A better beta model of navigation.. *British Journal of Mathematical and Statistical Psychology*, 48, 51-55.
- O'Keefe, J. and Conway, D.H. 1978. Hippocampus place units in the freely moving rat: Why they fire where they fire. *Experimental Brain Research*, 31, 573-590.
- O'Keefe, J. and Nadel, L. 1978. *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford.
- Rescorla, R.A. and Wagner, A.R. 1972. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black & W.F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp.64-99). New York: Appleton-Century Crofts.
- Rodrigo, T., Chamizo, V.D., McLaren, I.P.L., & Mackintosh, N.J. 1997. Blocking in the spatial domain. *Journal of Experimental Psychology: Animal Behavior Processes*, 23, 110-118.
- Sanchez, J., Rodrigo, T., Chamizo, V.D., & Mackintosh, N.J. 1999. Overshadowing in the Spatial Domain. *Animal Learning and Behavior* 27, 391- 398.
- Zipser, D. 1986. Biologically plausible models of place recognition and goal location. In McClelland, J.L. and Rumelhart, D.E. (Eds.) *Parallel Distributed Processing*. Vol. II. Cambridge. Mass. Bradford Books.