IDEA GROUP PUBLISHING
701 E. Chocolate Avenue, Suite 200, Hershey PA 17033-1240, USA
Tel: 717/533-8845; Fax 717/533-8661; URL-http://www.idea-group.com

ITJ3542

This paper appears in the publication, **International Journal of Cognitive Informatics and Natural Intelligence**, Volume 1, Issue 1 edited by **Yingxu Wang © 2007, Idea Group Inc.**

Representing an Intrinsically Nonmetric Space of Compass Directions in an Artificial Neural Network

Michael R.W. Dawson, University of Alberta, Canada

Patricia M. Boechler, University of Alberta, Canada

ABSTRACT

The purpose of the article is to train an artificial neural network to make a judgment that is intrinsically antisymmetric, and to determine how such a judgment is mediated by the network's internal representations. One key component of navigation is judging the distance from one location to another. A second key component of navigation is judging heading – that is, judging the direction from one location to another. Importantly, heading judgments do not preserve the metric properties of space. In particular, they are antisymmetric: the judged heading from location x to location y should be opposite to the judged heading from location y to location x. What kind of representation can mediate such nonmetric navigational judgments? To explore this question, we trained an artificial neural network to judge the various bearings amongst 13 different cities in Alberta. We then interpreted the internal structure of this network in order to determine the nature of its internal representations. We found that he artificial neural network had developed a coarse directional code, and that one of the advantages of such coding was its ability to represent antisymmetric spatial regularities.

Keywords: *neural networks; spatial representation*

INTRODUCTION

Cognitive informatics is an interdisciplinary field of research in which formal aspects of information processing problems are related to biological mechanisms involved in such processing (Wang, 2003). One kind of information processing medium of interest to cognitive informatics is that of artificial neural networks (e.g., Anderson, 2003; Dawson & Zimmerman, 2003). The current article illustrates an example of artificial neural network research in the cognitive informatics tradition. The formal property of interest is the representation of relations that are intrinsically antisymmetric. At issue are the kinds of internal representations developed by neural networks to deal with this formal property.

Representations that preserve the metric properties of space have been fundamentally important to the study of how humans and animals navigate (Kitchin, 1994). Behavioral studies have demonstrated that animal representations of space do indeed appear to preserve a good deal of its metric nature (for introductions, see Cheng & Spetch, 1998; Gallistel, 1990; Gallistel & Cramer, 1996). Since the discovery of place cells in the hippocampus (O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1978), many researchers have been concerned with identifying the biological substrates that encode metric space (Redish, 1999).

If a space is metric, then relationships between points in the space are constrained by three different principles (Blumenthal, 1953). The first is the *minimality principle*, which dictates that the shortest distance in the space is between a point x and itself. The second is the *symmetry principle*, which dictates that the distance in the space between two points x and y is equal to the distance between points y and x. The third is the *triangle inequality*, which dictates that the shortest distance in the space between two points y and x is a straight line. The fourth is the *non-negativity principle*, which dictates that every distance in the space between two points y and x must be greater than or equal to 0.

The assumption that some mental representations are metric spaces has had important theoretical and methodological impacts on cognitive science. With respect to theory, metric spaces have been used to explain a variety of cognitive phenomena, including models of similarity judgments (Shepard, 1974), analogical reasoning (Sternberg, 1977), judgments of metaphor aptness (Tourangeau & Sternberg, 1981, 1982), and transformations of mental images (Shepard, 1984). With respect to method, the assumption that some cognitive judgments are constrained by metric properties has led researchers to analyze a variety of behaviours using multidimensional scaling (Romney, Shepard & Nerlove, 1972; Shepard, Romney & Nerlove, 1972).

Navigation and Nonmetric Relationships

Importantly, the principles that characterize a metric space are all defined as properties of distance. Knowing the distance from a current location to a goal location is obviously an important component of navigation. However, knowing this distance is not by itself sufficient for successful navigation. In addition, it is also crucial to know the direction in which to travel: the bearing of the goal location from the current location. Interestingly, the notion of direction does not fit well with the metric properties of space. Consider a place on the map of Alberta, the city of Calgary. The town of Banff is west of Calgary. But this directional relationship violates the symmetry property of a metric space, because the direction from Banff to Calgary is east, not west. Indeed, For example, if one were to create a table of directions between cities, representing these directions as cosines, this table would be an *antisymmetric* matrix. That is, the value recorded in row x and column y of the table would be equal to the negative of the value recorded in row y and column x of the same table for every pair of off-diagonal cells. Given their intrinsically antisymmetric nature, how might a navigating system represent the bearings between pairs of locations? Of course, one possibility would be to represent the positions of the locations in a metric space, and then compute the required bearings.

However, it is not clear that this approach is appropriate for biological systems. O'Keefe and Nadel (1978) argued that the hippocampus represents the metric properties of space in a cognitive map. Importantly, a number of researchers have argued against this position, and against the proposal that the hippocampus provides a metric, systematic, and cohesive cognitive map. First, place cells are not organized topographically; the arrangement of place cells in the hippocampus is not isomorphic to the arrangements of locations in an external space (Burgess, Recce & O'Keefe, 1995; Eichenbaum, Dudchenko, Wood, Shapiro & Tanila, 1999; McNaughton, Barnes, Gerrard, Gothard, Jung & Knierim, 1996).

Second, place cell receptive fields ("place fields") are at best *locally* metric (Touretzky, Wan & Redish, 1994), and as a result a good deal of spatial information (e.g., information about bearing) cannot be derived from place cell activity. Third, the responses of different place cells do not appear to be related to one another in a coherent and holistic spatial framework (Eichenbaum et al., 1999). Fourth, hippocampal cells appear to be sensitive to a wide variety of nonspatial variables (Best, White & Minai, 2001; Eichenbaum, 2002). Evidence of this sort has led some to abandon the view that the hippocampus provides a spatial cognitive map (Bennett, 1996).

Artificial neural networks represent an alternative approach to representing spatial relationships. For example, Dawson, Boechler, and Valsangkar-Smyth (2000) trained a network on a distance judgment task that violated the minimality principle. Interestingly, the internal components of this network(i.e., the "hidden units" in the network) were not very map-like. The responses of each hidden unit were at best locally metric. However, when several different patterns of such distorted representations were combined, an accurate spatial judgment (i.e., a distance rating) could be generated. Dawson, Boechler, and Valsangkar-Smyth called this representational scheme *coarse allocentric coding*. Clearly such coding is capable of representing spatial relations in a nonmetric space.

The purpose of the current article is to explore the capability of an artificial neural network to make directional judgments based on the antisymmetric data that was presented in Table 1. There are two questions of interest. First, is a network capable of making this set of judgments? Second, if a network can learn to respond in this way, then what is the nature of its internal representations?

METHOD

Problem Definition

A network was trained to make judgments about the bearing from one city to another. We chose thirteen different locations in the province of Alberta: Banff, Calgary, Camrose, Drumheller, Edmonton, Fort McMurray, Grande Prairie, Jasper, Lethbridge, Lloydminster, Medicine Hat, Red Deer, and Slave Lake. We took all possible pairs from this set to create a set of 169 different stimuli, each of which could be described as the question "In the context of a compass with 16 different directional petals, what bearing is City 1 from City 2?" where City 1 served as an origin to which a reference heading vector pointing north was assigned. Because all possible pairs of place names were used, thirteen of the stimuli involved rating the direction from one city to itself (e.g., the bearing from Banff to Banff). As well, for different place names a rating would be obtained for both orders of places (e.g., the bearing from Banff to Calgary would be judged, as would be the bearing from Calgary and Banff).

The desired ratings for each stimulus were created as follows. First, the cosines of each bearing were computed, as was described earlier when Table 1 was introduced. Second, these cosines were converted into discrete directions. This was done by converting each cosines into the nearest corresponding petal of the 16 petal compass rose that is illustrated in Figure 1. If a stimulus involved rating the direction from one place to itself, the rating was assigned a value of 0, which did not correspond to any of the 16 petals. For these stimuli, the network was trained not to make any directional judgment.

The resulting set of compass petal assignments is provided in Table 1. Because only positive numbers were used when petals were assigned, this matrix is not antisymmetric. That is, it is *not* the case that the value recorded in row x and column y of is equal to the negative of the cosine recorded in row y and column x for every pair of off-diagonal cells in Table 1. However, the table is strongly asymmetric. Following accepted procedures (Dawson & Harshman, 1986), the proportion of antisymmetric variance in the matrix was 0.593. In other words, symmetric variance makes up less than half of

Figure 1. The 16 petal compass rose that was used to define how the network judged bearing. When asked to judge the bearing from one city to another, the network was trained to turn on an output unit associated with the closest petal on this compass.

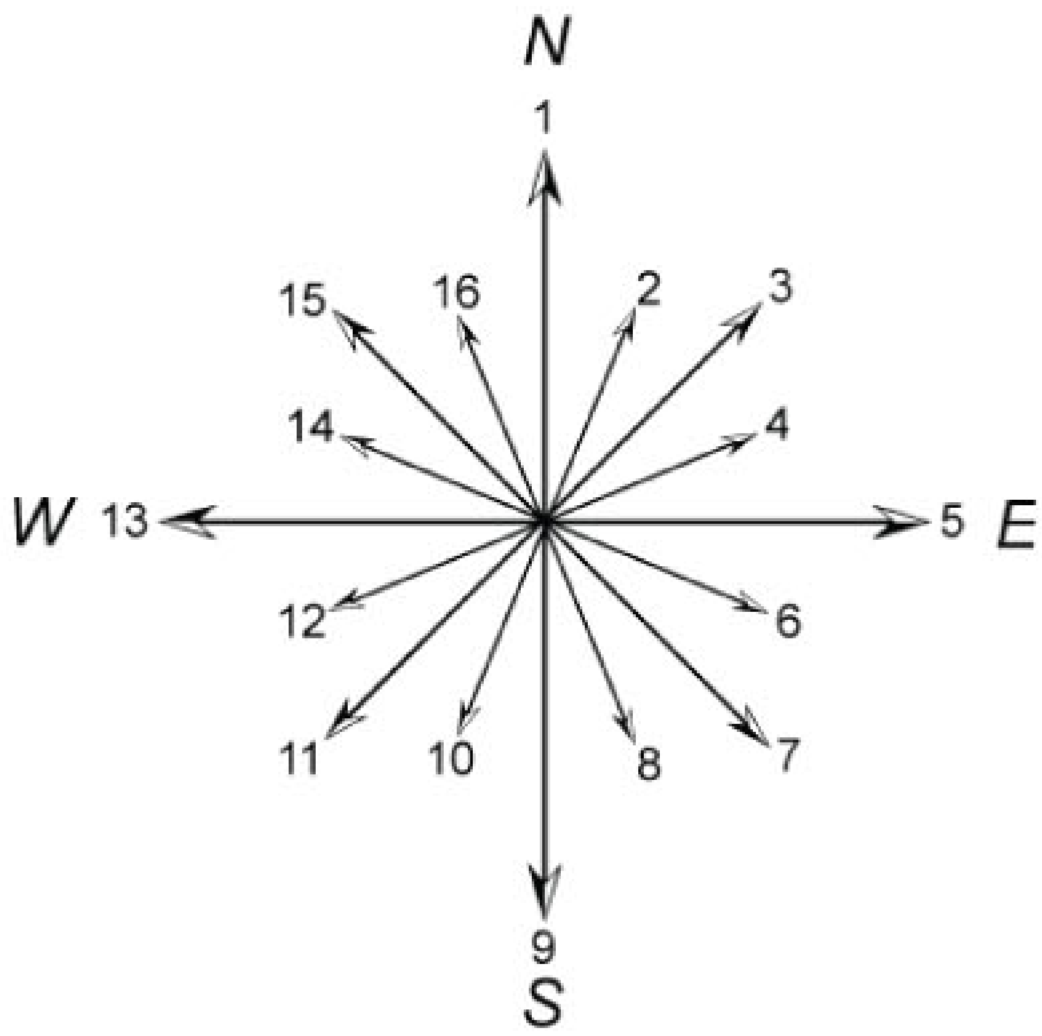

the overall variance in Table 1, indicating that this table is extremely asymmetric. This also indicates that this table strongly violates the symmetry assumption of metric space.

Network Architecture

Artificial neural network architectures can differ from one another along a variety of dimensions. One important distinction between different classes of architectures involves the activation function that converts a unit's total input into an internal level of activity (Duch & Jankowski, 1999). Typically, neural networks are composed of units whose activation function is a sigmoid-shaped curve that is defined by the logistic equation; such units have been called *integration devices*_(Ballard, 1986). However, not all networks are composed of integration devices. Some networks use processors that are tuned to activate to a small range of net inputs, and generate weak responses to net inputs that are either too small or too large to fall in this range. Ballard calls such units *value units.* The network that was trained in the current study had 16 output units, 7 hidden units, and 26 input units. All of the output units and all of the hidden units were value units that used the Gaussian activation function $G(net_j) = exp\ (-\pi(net_i - \mu_j)^2$, where net_i is the net input to the unit, and μ_j is the mean of the unit's activation function (Dawson & Schopflocher, 1992).

Input Unit Representation

Twenty-six input units were used to define a local code for pairs of cities to be compared. Each input unit represented a place name; the first input unit represented "Banff," the second input unit represented "Calgary," and so on alphabetically.

Every stimulus presented to the network was a request to judge the bearing "from City A to City B." Thus two different city names had to be represented in an input pattern: the name of the "from" city (i.e., the first name in a pair of names, and the origin from which the bearing was being judged), and the name of the "to" city (i.e., the second name in a pair of names). The first thirteen input units represented the name of the "from" city, while the remaining thirteen input units represented the name of the "to" city. As a result, each pattern that was presented to the network involved

Table 1. Assignment of compass petals to each city pair. See text for explanation.

	BANFF	CALGARY	CAMROSE	DRUMHELLER	EDMONTON	FORT MCMURRAY	GRANDE PRAIRIE	JASPER	LETHBRIDGE	LLOYDMINSTER	MEDICINE HAT	RED DEER	SLAVE LAKE
BANFF	0	5	3	5	3	3	15	15	6	4	6	4	1
CALGARY	13	0	2	4	1	2	15	14	7	4	6	1	1
CAMROSE	11	10	0	9	15	2	14	13	9	5	8	11	15
DRUMHELLER	13	12	1	0	16	1	14	14	9	3	7	14	16
EDMONTON	11	9	7	8	0	3	14	13	9	5	7	9	15
FORT MCMURRAY	11	10	10	9	11	0	13	12	9	8	9	10	12
GRANDE PRAIRIE	7	7	6	6	6	5	0	8	7	6	6	6	5
JASPER	7	6	5	6	5	4	16	0	6	5	6	5	3
LETHBRIDGE	14	15	1	1	1	1	15	14	0	3	5	16	16
LLOYDMINSTER	12	12	13	11	13	16	14	13	11	0	9	12	14
MEDICINE HAT	14	14	16	15	15	1	14	14	13	1	0	15	15
RED DEER	12	9	3	6	1	2	14	13	8	4	7	0	1
SLAVE LAKE	9	9	7	8	7	4	13	11	8	6	7	9	0

turning two input units on, one that belonged to the first set of thirteen input units, and one that belonged to the second set of thirteen input units. For example, to ask the network to judge the bearing from Banff to Calgary, the first and the fifteenth input units would be turned on, and all other input units would be turned off. Similarly, to ask the network to judge the bearing from Calgary to Banff, the second and the fourteenth input units would be turned on, and all other input units would be turned off.

This representational scheme was chosen because it contains absolutely no information about the location of the different places on a map of Alberta. In other words, the input units themselves did not provide any information that the network could use to perform the ratings task. All possible pairs of names from the set of thirteen cities were used as training patterns. As a result, there were 169 different input patterns that were used to train the network.

Output Unit Representation

Sixteen output units were used to represent the network's rating of the bearing between the two place names presented as input. The output units were also value units. To represent a judgment of bearing, the network was trained to turn on one, and only one, of its output units. Each of these output units represented one of the compass rose petals from 1 to 16. For example, if the network turned output unit 13 on, this indicated that it was judging the bearing to be 270° (i.e., west). For the 13 cases in which the direction of a city to itself was to be judged, the network was trained to turn all of its output units off.

Number of Hidden Units

Seven hidden units were used in this network to solve the problem. Each of these units was a value unit. We selected this number of hidden units because pilot tests had shown that this was the smallest number of hidden units that could be used by the network to discover a mapping from input to output. When fewer than seven hidden units were used, the network was never able to completely learn the task. Previous research has suggested that forcing a network to learn a task with the minimum number of hidden units produces a network that is much easier to interpret, in comparison to a network that has more hidden units than are required to solve a problem (Berkeley, Dawson, Medler, Schopflocher & Hornsby, 1995).

Network Training

The network in the current article was trained using the variation of the generalized delta rule that has been developed for networks of value units (Dawson & Schopflocher, 1992). Prior to training, all of the connection weights were randomly assigned values ranging from –0.10 to +0.10. The biases of processing units (i.e., the means of the Gaussian activation functions, which are analogous to thresholds) were all initialized to a value of 0. The network was trained with a learning rate of 0.01 and zero momentum. During each epoch of training, each of the 169 stimuli was presented to the network, with the order of presentation of the set of 169 patterns randomized prior to beginning the epoch. The learning rule was used to update connection weights in the network after each stimulus presentation.

Training proceeded until the network generated a "hit" for every output unit on every pattern. A hit was operationalized as an activation of 0.90 or higher when the desired activation was 1.00 and as an activation of 0.10 or lower when the desired activation was 0.00. The network converged on a solution to the problem — generating a correct response for each of the 169 patterns — after 7645 epochs of training.

RESULTS

Asymmetry of Hidden Unit Responses

We noted earlier that the matrix of compass rose points (Table 1) is highly asymmetric. One question of interest is the degree to which the asymmetry of the desired responses to various city pairs is reflected in the responses and the connection weights of the seven hidden units in this network.

Our first approach in answering this question was to create a 13 X 13 matrix of hidden unit responses for each hidden unit. The rows and columns of each matrix corresponded to the thirteen cities, and each matrix entry was the response of a hidden unit to a particular city pair. After creating these matrices, we decomposed them into their symmetric and antisymmetric components in order to determine the degree to which hidden unit responses were asymmetric. The first column of numbers in Table 2 presents the proportion of antisymmetry for the activation matrix of each hidden unit. It can be seen from this table that all of the activation matrices had substantial antisymmetric components. In other words, the asymmetry of the direction judgment task is reflected in the asymmetry of each hidden unit's responses.

Given that hidden unit activations are strongly asymmetric, it would not be surprising to discover that the net inputs to the hidden units are strongly asymmetric as well. However, because hidden unit activations represent a Gaussian transformation of net inputs, it is not completely clear what the relationship between the structure of hidden unit activations and hidden unit net inputs should be. Therefore we computed a second set of 13 X 13 matrices for each hidden unit. These were identical to the first set, with the exception that each matrix entry was the net input to the hidden unit instead of the hidden unit's activation to a particular city pair. The proportions of antisymmetry of the net input matrices are also reported in Table 2. Again, these matrices have substantial antisymmetric components. Furthermore, there is a strong relationship between the degree of antisymmetry in the activation matrices and the degree of antisymmetry in the net input matrices. The correlation between the first two columns of numbers in Table 2 is 0.96.

The source of the strong asymmetry in hidden unit activations and net inputs must be

Table 2. Measures of asymmetries of the activations and net inputs of the hidden units to each pair of cities, and the correlation between the two sets of weights that feed into each hidden unit. See text for details.

Hidden Unit	Proportion Asymmetry Of Activation Matrix	Proportion Asymmetry Of Net Input Matrix	Correlation Between "From" Weights and "To" Weights
H1	0.47	0.63	-0.27
H2	0.36	0.36	0.28
H3	0.51	0.49	0.03
H4	0.92	0.95	-0.91
H5	0.72	0.86	-0.76
H6	0.45	0.50	-0.01
H7	0.81	0.92	-0.86

the connection weights from the input units to the hidden units. Each hidden unit has two connection weights feeding into it from the same location: one weight for when the location is the "from" city, and the other weight for when it is the "to" city. For each hidden unit, we took each set of weights (i.e., the 13 "from" weights and the 13 "to" weights) and correlated them to find their relationship. The resulting correlations are reported in the last column of Table 2. It can be seen from this table that these correlations range from being weakly positive to strongly negative. This indicates that the source of asymmetry in hidden unit responses are differences between the "from" weights and the "to" weights that feed into a hidden unit from the same city. Furthermore, as the correlation between weights becomes more negative, there is a greater proportion of antisymmetry in both the activation and the net input matrices. Indeed, the correlation between the last column of Table 2 and the proportion of antisymmetry in the activation matrices (i.e., the first column of numbers) is –0.95. The correlation between the last column of Table 2 and the proportion of antisymmetry in the net input matrices (i.e., the middle column of numbers) is –0.99.

Directional Selectivity of Hidden Units

The previous section has shown that hidden unit responses reflect the intrinsic asymmetry of the direction judgment task. But how do hidden unit responses represent information relevant to city bearings? Figure 2 presents a polar bar chart for each hidden unit. These charts present the median hidden unit activity to city pairs as a function of the desired output (i.e., the appropriate compass rose petal) for each of the city pairs, excluding the thirteen stimuli to which the network was trained not to respond.

There are four important observations that can be made from the graphs in Figure 2. First, all of the hidden units can be described as being directionally selective. This is because all of the hidden units generate strong responses when input city pairs are associated with some directions, but generate weak responses when input city pairs are associated with other directions. Second, the directional selectivity of all of the hidden units is approximately confined to a hemisphere of the compass rose or less; for example, Hidden Unit 5's directional selectivity is confined to almost a quadrant of the compass. Third, taken as a whole, all of

Figure 2. Bar plots of the median hidden unit activity as a function of the petal on the compass rose, for each of the seven hidden units in the network. The 13 stimuli to which the network was trained not to respond have been excluded from each of these plots.

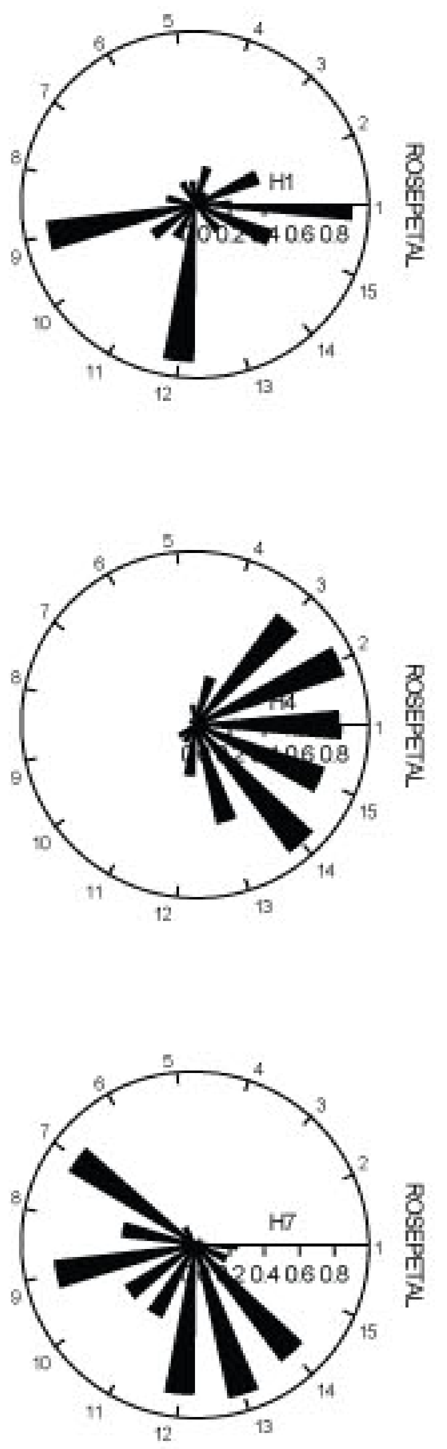

the compass directions generate a response in some of the hidden units. In other words, the combined directional selectivity of the entire set of hidden units covers the complete compass rose. Fourth, the directional selectivity of individual hidden units is sometimes systematic, but at other times is not. For example, Hidden Unit 4 exhibits systematic responses, because it generates strong responses to directions in the northern hemisphere of the compass (petals 3 through 14 in Figure 1). However, Hidden Unit 1 is much less systematic, generating strong responses to compass petals 1, 12, and 9, but not to other directions that lay between these petals (e.g., petals 10 and 11).

Coarse Coding of Hidden Units

One key contribution of connectionism to general theories about representation is the notion of *coarse coding* (Hinton, McClelland, & Rumelhart, 1986). In general, coarse coding means that an individual processor is sensitive to a broad range of features, or at least to a broad range of values of an individual feature (e.g., Churchland & Sejnowski, 1992). However, if different processors have overlapping sensitivities, then their outputs can be pooled, which can result in a highly useful and accurate representation of a specific feature. There are two sources of evidence that indicate that the hidden units of the current network are involved in the coarse coding of city bearings. The first comes from the observations of Figure 2 that have already been made. The facts that each hidden unit is sensitive to a variety of different directions, and that more than one hidden unit exhibits sensitivity to the same direction, both are strong indicators that this network uses coarse coding.

Table 3. Activation produced in three different hidden units by the 15 city pairs that all cause the network to choose compass petal 6 as the response

"From" City	"To" City	H3	H5	H6
BANFF	LETHBRIDGE	0.96	0.41	0.01
BANFF	MEDICINE HAT	0.96	0.46	0.01
CALGARY	MEDICINE HAT	0.06	0.87	0.61
GRANDE PRAIRIE	CAMROSE	0.17	0.77	0.85
GRANDE PRAIRIE	DRUMHELLER	0.64	0.75	0.60
GRANDE PRAIRIE	EDMONTON	0.89	0.99	0.97
GRANDE PRAIRIE	LLOYDMINSTER	0.01	0.49	0.39
GRANDE PRAIRIE	MEDICINE HAT	0.66	0.72	0.23
GRANDE PRAIRIE	RED DEER	0.00	0.75	0.81
JASPER	CALGARY	0.72	0.86	0.49
JASPER	DRUMHELLER	0.37	0.57	0.41
JASPER	LETHBRIDGE	0.68	0.66	0.37
JASPER	MEDICINE HAT	0.35	0.61	0.38
RED DEER	DRUMHELLER	0.78	0.56	0.00
SLAVE LAKE	LLOYDMINSTER	0.61	0.76	0.53

The second source of evidence comes from taking a more detailed examination of directional sensitivity than can be made by simply examining the median levels of activity that are graphed in Figure 2. Let us take as an example compass petal 6 from Figure 1. There are 15 different city pairs that lead the network to choose this particular direction; each of these pairs is listed in Table 3. An examination of Figure 2 indicates that Hidden Units 3, 5, and 6 all exhibit either a moderate or a strong sensitivity to this direction. Table 3 also presents the activity of each of these hidden units produced by each of the 15 city pairs in the table.

Table 3 indicates that the hidden units are involved in coarse coding by demonstrating that their directional sensitivity is noisy. For example, consider the activity of Hidden Unit 3. It generates very strong activity to some of the city pairs (e.g., Banff and Lethbridge, Banff and Medicine Hat), but also generates near zero activity to other city pairs (e.g., Calgary and Medicine Hat, Grand Prairie and Lloydminster), even though all of these city pairs are associated with the same bearing. Similar examples exist for the other two hidden units. In other words, while a case can be made that these hidden units have a preference for compass petal 6 (as well as to a handful of other directions), it is also clear that this preference is mixed. None of these hidden units generate strong activity to every pattern associated with this compass direction.

A second property that one expects to see in coarse coding is a complementary pattern of activation. That is, if one hidden unit fails to respond to one stimulus that it might be expected to be sensitive to, then this failure should be compensated by the behavior of other hidden units involved in the coarse coding. Table 3 also provides evidence of this type. Consider for example the city pair of Red Deer and Drumheller. Hidden Unit 6 fails to generate any activity to this pattern at all. However, Hidden Unit 5 generates a moderate response to it, while Hidden Unit 3 generates an even stronger response. Indeed, there is no example in Table 3 of a stimulus pair that generates near zero activity in all three of these hidden units.

DISCUSSION

Relation to Head Direction Cells

Given the directional selectivity of the hidden units in the network, one obvious question concerns their relationship to head direction cells that have been extensively studied by neuroscientists (for overviews see Redish, 1999; Sharp, Blair & Cho, 2001; Taube, 1998). In general, head direction cells have a preferred direction of firing, and generate strong signals when an animal's head is pointing in this preferred direction. Usually, head direction neurons for a variety of different directions are available, and cooperative and competitive interactions between these different cells result in one direction being selected (and represented with neural activity in the appropriate cells). Most models of this type of processing use an attractor network (a self-organizing artificial neural network) with different head direction cells linked with excitatory and inhibitory connections to create a directional cell assembly (e.g., Redish, 1999).

The source of this directional sensitivity is complex. While the responses of head direction cells can come under the control of visual landmarks, these neurons are also able to maintain their directional heading when such landmarks are absent, most likely under the control of self-motion cues such as vestibular input (e.g., Knierim, Kudrimoti, & McNaughton, 1998). Such idiothetic cues as this latter type are an important source of input to head direction cells, and as a result these cells are well suited for performing navigation via path integration (e.g., Gallistel, 1990). Furthermore, there is an important temporal component to head direction cells. For example, head direction cells located in the postsubiculum represent head direction 10 ms in the past, while head direction cells located in the anterior dorsal thalamus anticipate head direction by 25 ms into the future (Taube & Muller, 1998).

In many respects, the kind of information that head direction cells appear to encode differs from the kind of information to which the hidden units in the artificial neural network are sensitive. For instance, while individual head direction cells are sensitive to a unary directional feature (e.g., the compass direction of the animal's head), each hidden unit is sensitive to a more complex directional feature (e.g., several different headings emanating from an origin city). Furthermore, individual head direction cells represent an egocentric direction: that is, the direction from the animal's head to some location in space. In contrast, each hidden unit uses an allocentric frame of reference rather than an egocentric frame of reference. That is, each hidden unit is sensitive to directions from one location on a map to others, whether none of the map locations correspond to the hidden unit's location.

However, some general similarities between the two types of processors exist as well. For example, head direction is usually viewed as being represented by the activities of an ensemble of neurons, each sensitive to a different direction (e.g., Redish, 1999). This is functionally very similar to the notion of a coarse code that was revealed in the network. Head direction cells also often exhibit Gaussian-shaped tuning functions (Goodridge & Touretzky, 2000), although the shape of this function can be distorted as an animal moves.

Relation to Other Spatial Network Codes

We have studied other networks on tasks that can be related to the one described above, and it is interesting to consider the relationships between the current network and these other "navigational nets".

Dawson, Boechler, and Valsangkar-Smyth (2000) trained one network to make distance judgments between all possible pairs of the same 13 Albertan cities that were used in the current simulation. They discovered a representation in which hidden units occupied positions on the map of Alberta, and hidden unit connection weights represented transformed distances between city and hidden unit locations. One key similarity between this network and the current one is that both used coarse coding. A second similarity was that the coarse coding discovered by Dawson, Boechler and Valsangkar-Smyth was flexible enough to mediate distance judgments that violated the metric properties of space. A second network that they trained violated the minimality principle. This network required an additional hidden unit to deal with this violation, but all of the hidden units used the same representational code as did the hidden units in the network in which the minimality principle was not violated. However, a key difference between these networks and the current one was that a compressed input representation (13 input units) was used, and so it could not be trained on distance judgments that violated other metric principles, such as symmetry.

In a more recent study, a second network was trained to make distance judgments when all of the metric principles were preserved, but the input representation was identical to the one used in the current simulation (Dawson, Boechler & Orsten, 2005). This network also used coarse coding in its internal representations. Interestingly, even though this network was trained to make distance judgments, it used a coarse coding of direction. However, the representation of direction in its hidden units was quite different from that described above. Dawson and Boechler found that each hidden unit could be described as occupying a position on the map of Alberta, and had a particular point of view of the map from this position. The hidden unit then acted somewhat as a sextant, where connection weights were related to the angle between the heading of the hidden unit (i.e., its point of view) and the bearings of the cities.

AKNOWLEDGMENTS

This research was supported by NSERC and SSHRC research grants to MRWD, and by a SSHRC research grant to PMB.

REFERENCES

Anderson, J. A. (2003). Arithmetic on a parallel computer: Perception versus logic. *Brain & Mind, 4*, 169-188.

Ballard, D. (1986). Cortical structures and parallel processing: Structure and function. *The Behavioral And Brain Sciences, 9*, 67-120.

Bennett, A. T. D. (1996). Do animals have cognitive maps? *Journal of Experimental Biology, 199*(1), 219-224.

Berkeley, I. S. N., Dawson, M. R. W., Medler, D. A., Schopflocher, D. P., & Hornsby, L. (1995). Density plots of hidden value unit activations reveal interpretable bands. *Connection Science, 7*, 167-186.

Best, P. J., White, A. M., & Minai, A. (2001). Spatial processing in the brain: The activity of hippocampal place cells. *Annual Review of Neuroscience, 24*, 459-486.

Blumenthal, L. M. (1953). *Theory and applications of distance geometry*. London: Oxford.

Burgess, N., Recce, M., & O'Keefe, J. (1995). Spatial models of the hippocampus. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 468-472). Cambridge, MA: MIT Press.

Cheng, K., & Spetch, M. L. (1998). Mechanisms of landmark use in mammals and birds. In S. Healy (Ed.), *Spatial representation in animals* (pp. 1-17). Oxford: Oxford University Press.

Churchland, P. S., & Sejnowski, T. J. (1992). *The computational brain*. Cambridge, MA: MIT Press.

Dawson, M.R.W., Boechler, P.M., & Valsangkar-Smyth, M. (2000). Representing space in a PDP network: Coarse allocentric coding can mediate metric and nonmetric spatial judgements. *Spatial Cognition and Computation, 2*, 181-218.

Dawson, M. R. W., & Harshman, R. A. (1986). The multidimensional analysis of asymmetries in alphabetic confusion matrices: Evidence for global-to-local and local-to-global processing. *Perception & Psychophysics, 40*, 370-383.

Dawson , M.R.W., Boechler, P.M., & Orsten, J. (2005). An artificial neural network that uses coarse allocentric coding of direction to represent distances between locations in a metric space. *Spatial Cognition and Computation, 5*, 29-67.

Dawson, M. R. W., & Schopflocher, D. P. (1992). Modifying the generalized delta rule to train networks of nonmonotonic processors for pattern classification. *Connection Science, 4*, 19-31.

Dawson, M. R. W., & Zimmerman, C. (2003). Interpreting the internal structure of a connectionist model of the balance scale task. *Brain & Mind, 4*, 129-149.

Duch, W., & Jankowski, N. (1999). Survey of neural transfer functions. *Neural Computing Surveys, 2*, 163-212.

Eichenbaum, H. (2002). *The cognitive neuroscience of memory: An introduction*. Oxford/New York: Oxford University Press.

Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., & Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron, 23*(2), 209-226.

Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.

Gallistel, C. R., & Cramer, A. E. (1996). Computations on metric maps in mammals: Getting oriented and choosing a multi-destination route. *Journal of Experimental Biology, 199*(1), 211-217.

Goodridge, J. P., & Touretzky, D. S. (2000). Modeling attractor deformation in the rodent head-direction system. *Journal of Neurophysiology, 83*(6), 3402-3410.

Hinton, G. E., McClelland, J., & Rumelhart, D. (1986). Distributed representations. In D. Rumelhart & J. McClelland (Eds.), *Parallel distributed processing* (Vol. 1, pp. 77-109). Cambridge, MA: MIT Press.

Kitchin, R. M. (1994). Cognitive maps: What are they and why study them? *Journal Of Environmental Psychology, 14*, 1-19.

Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1998). Interactions between idiothetic cues and external landmarks in the control of place cells and head direction cells. *Journal of Neurophysiology, 80*(1), 425-446.

McNaughton, B., Barnes, C. A., Gerrard, J. L., Gothard, K., Jung, M. W., Knierim, J. J., et al. (1996). Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *The Journal of Experimental Biology, 199*, 173-185.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Research, 34*, 171-175.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.

Redish, A. D. (1999). *Beyond the cognitive map*. Cambridge, MA: MIT Press.

Romney, A. K., Shepard, R. N., & Nerlove, S. B. (1972). *Multidimensional scaling: Theory and applications in the behavioral sciences. Volume II: Applications*. New York: Seminar Press.

Sharp, P. E., Blair, H. T., & Cho, J. W. (2001). The anatomical and computational basis of the rat head-direction cell signal. *Trends in Neurosciences, 24*(5), 289-294.

Shepard, R. N. (1974). Representation of structure in similarity data: Problems and prospects. *Psychometrika, 39*, 373-421.

Shepard, R. N. (1984). Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review, 91*, 417-447.

Shepard, R. N., Romney, A. K., & Nerlove, S. B. (1972). *Multidimensional scaling: Theory and applications in the behavioral sciences. Volume I: Theory*. New York: Seminar Press.

Sternberg, R. J. (1977). Component processes in analogical reasoning. *Psychological Review, 84*, 353-378.

Taube, J. S. (1998). Head direction cells and the neurophysiological basis for a sense of direction. *Progress in Neurobiology, 55*(3), 225-256.

Taube, J. S., & Muller, R. U. (1998). Comparisons of head direction cell activity in the postsubiculum and anterior thalamus of freely moving rats. *Hippocampus, 8*(2), 87-108.

Tourangeau, R., & Sternberg, R. J. (1981). Aptness in metaphor. *Cognitive psychol-*

ogy, 13, 27-55.

Tourangeau, R., & Sternberg, R. J. (1982). Understanding and appreciating metaphors. *Cognition, 11*, 203-244.

Touretzky, D. S., Wan, H. S., & Redish, A. D. (1994). Neural representation of space in rats and robots. In J. M. Zurada, R. J. Marks, & C. J. Robinson (Eds.), *Computational intelligence: Imitating life,* 57-68. New York: IEEE Press.

Wang, Y. (2003). Cognitive informatics: A new transdisciplinary research field. *Brain & Mind, 4*, 115-127.

Michael R.W. Dawson received his PhD in psychology from the University of Western Ontario in 1986, and is currently a full professor in the psychology department at the University of Alberta. His research interests include pure and applied research on artificial neural networks and the relationship of this research to empirical and theoretical issues in cognitive science. (Homepage: http://www.bcp.psych.ualberta.ca/~mike/)

Patricia Boechler received her PhD in psychology from the University of Alberta in 2002, and is currently a member of that university's educational psychology department. Her general research interests include cognition, memory, learning and developmental psychology. For the last few years her research has been centered around the study of cognition and learning in educational hypermedia. She has been investigating the effects of different types of interfaces on learning, taking into account individual differences in spatial and literacy skills. She is also interested in developing novel research and statistical methods for uncovering regularities in user navigation behaviours. She is applying her background in the study of spatial cognition using neural networks to the use of neural networks for understanding students' path patterns in educational hypermedia.