Psychology 452
Week 13: Spatial Representations In PDP Networks

- **Hippocampus As A Cognitive Map**
- **Networks Learn Metric Spaces**
- **Networks Learn Nonmetric Spaces**

Course Trajectory

When	What
Weeks 1-3	Basics of three architectures (DAM, perceptron, MLP)
Weeks 4-6	Cognitive science of DAMs and perceptrons
Week 7	Music and networks
Weeks 8-10	Interpreting MLPs
Weeks 11-13	**Case studies (interpretations, applications, architectures)**

Psychology and Space

- **Spatial behavior, or spatial reasoning, have long been studied by psychologists**
- **Some of the earliest cognitive proposals are found in Tolman's studies of spatial behavior**
- **Tolman introduced the term 'cognitive map' long before the cognitive revolution occurred**


Edward Tolman

The Cognitive Map

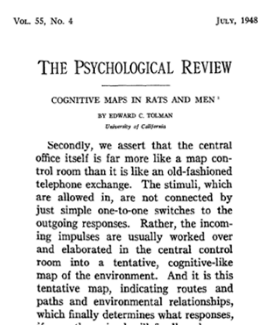
VOL. 55, NO. 4 JULY, 1948

THE PSYCHOLOGICAL REVIEW

COGNITIVE MAPS IN RATS AND MEN [1]

BY EDWARD C. TOLMAN
University of California

Secondly, we assert that the central office itself is far more like a map control room than it is like an old-fashioned telephone exchange. The stimuli, which are allowed in, are not connected by just simple one-to-one switches to the outgoing responses. Rather, the incoming impulses are usually worked over and elaborated in the central control room into a tentative, cognitive-like map of the environment. And it is this tentative map, indicating routes and paths and environmental relationships, which finally determines what responses, if any, the animal will finally release.

Example Evidence

- Rats, when finding a route to a goal blocked, will find a related route that will take them towards the goal whose learned path is not available
- "As a result of their original training, the rats had, it would seem, acquired not merely a strip-map to the effect that the original specifically trained-on path led to food but, rather, a wider comprehensive map to the effect that food was located in such and such a direction in the room"

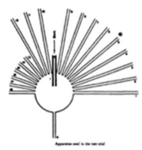

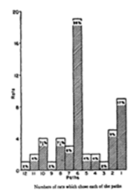

Critique

- Tolman's cognitivist musings were critiqued
- In a 1935 text on learning, Guthrie noted: "In his concern with what goes on in the rat's mind, Tolman has neglected to predict what the rat will do. So far as the theory is concerned the rat is left buried in thought; if he gets to the food-box at the end that is [the rat's] concern, not the concern of the theory" (p. 172)
- This is a variant of Ryle's regress being used to attack functionalism and cognitivism
- Needed: an architecture for the cognitive map!

Edwin Ray Guthrie

The Hippocampus

- The hippocampus is a component of the limbic system
- Earlier we saw that there are reasons to believe that it is a locus of Hebb-like learning
- There is also substantial evidence to suggest that it provides the cognitive map

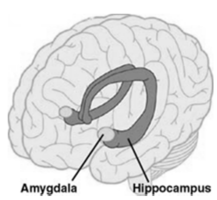

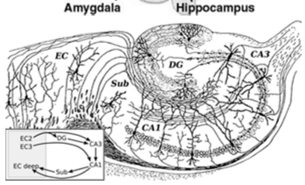

The Hippocampus And Place

- In the 1970s, single-cell recordings from the hippocampus revealed a biological basis for the 'cognitive map'
- Place cells fire when an animal's head is at a particular location in the world
- A variety of other spatial location cells have been found in the hippocampus since the discovery of place cells

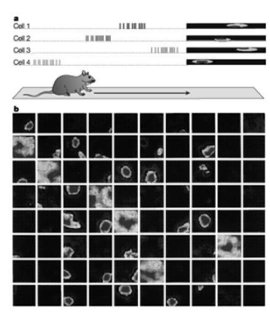

The Brain And Head Direction

- Further explorations of space in the brain discovered other interesting cells that encoded spatial information
- For example, head direction cells spike when an animal's head is pointed in a specific direction

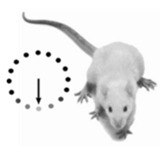

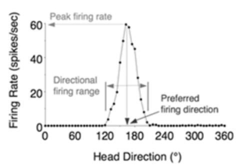

The Hippocampus As A Cognitive Map

- 1978 saw the publication of a seminal book detailing the results of such work
- 2014 saw John O'Keefe, May-Britt Moser and Edvard Moser win the Nobel Prize in Physiology or Medicine for this work

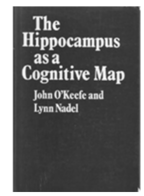

Studying Spatial Learning

- Why is the map in the hippocampus? This permits spatial learning to occur
- Morris (1984) invented the water maze to study spatial learning of rats
- Rats swim to escape the water, and discover a hidden platform
- On repeated trials, the time taken to find the platform decreases, indicating that the agents are learning its position in the maze

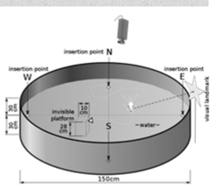

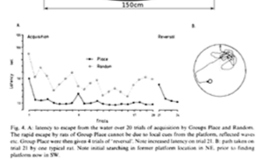

Blocking Spatial Learning

- Morris et al. (1986) studied the role of LTP in spatial learning
- Rats were placed in the Morris water maze
- In the experimental condition, animals were treated with D, L-AP5, known to block NMDA receptors in the hippocampus
- Rats treated in this way were unable to learn the position of the platform in the maze

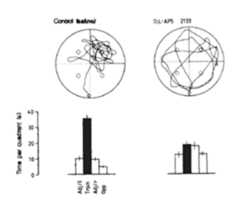

Paths taken in the maze by control animals on the left and by experimental animals on the right. Lack of preference for a particular quadrant indicates experimental animals did not learn the position of the platform

Critiquing The Map

- Some of the ideas about the hippocampus as a cognitive map have been challenged via results from neuroscience
- It is not topographically organized
- It is at best locally metric
- Let us consider whether PDP networks can contribute anything to this discussion

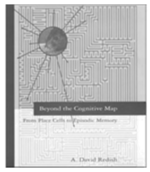

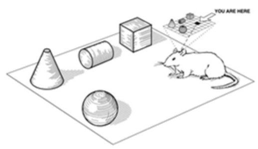

A Metric Space

- A metric space has three general properties true of the relations between points:
- Minimality
  - Distance from a point to itself is a minimum
- Symmetry
  - d(x,y) = d(y,x)
- Triangle Inequality
  - d(x,y) + d(y, z) ≥ d(x,z)
- If a network learned such a space, how would it represent it?

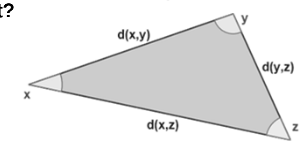

Alberta As A Metric Space

- Dawson, Boechler and Valsangar-Smyth (2000) trained a network to estimate the distances between cities in Alberta
- If it can learn such a task, then it has – by definition -- internalized a metric space

Distance Table

- The training set was created by using a table of distances between cities to determine output ratings

Table 1. (A) Distances between cities of Alberta, measured in kilometres. (B) The distances from Table 1a converted into ratings on a 0 to 10 scale

	BANFF	CALGARY	CAMROSE	DRUM-HELLER	EDMONTON	FORT MCMURRAY	GRANDE PRAIRIE	JASPER	LETH-BRIDGE	LLOYD-MINSTER	MEDICINE HAT	RED DEER	SLAVE LAKE
BANFF	0	128	381	263	401	840	682	287	342	626	419	253	652
CALGARY	128	0	274	138	294	733	720	412	216	519	293	145	545
CAMROSE	381	274	0	182	97	521	553	463	453	245	429	129	348
DRUMHELLER	263	138	182	0	279	703	735	547	282	416	247	165	530
EDMONTON	401	294	97	279	0	439	456	366	509	251	526	148	250
FORT MCMURRAY	840	733	521	703	439	0	752	796	948	587	931	587	436
GRANDE PRAIRIE	682	720	553	735	456	752	0	397	935	701	982	586	318
JASPER	287	412	463	547	366	796	397	0	626	613	703	413	464
LETHBRIDGE	342	216	453	282	509	948	935	626	0	605	168	360	760
LLOYDMINSTER	626	519	245	416	251	587	701	613	605	0	480	374	496
MEDICINE HAT	419	293	429	247	526	931	982	703	168	480	0	409	777
RED DEER	253	145	129	165	148	587	586	413	360	374	409	0	399
SLAVE LAKE	652	545	348	530	250	436	318	464	760	496	777	399	0

A

Table 1. Continued

	BANFF	CALGARY	CAMROSE	DRUM-HELLER	EDMONTON	FORT MCMURRAY	GRANDE PRAIRIE	JASPER	LETH-BRIDGE	LLOYD-MINSTER	MEDICINE HAT	RED DEER	SLAVE LAKE
BANFF	0	2	4	3	5	9	7	3	4	7	5	3	7
CALGARY	2	0	3	2	3	8	8	5	3	6	3	2	6
CAMROSE	4	3	0	2	1	6	6	5	5	3	5	2	4
DRUMHELLER	3	2	2	0	3	8	8	6	3	5	3	2	6
EDMONTON	5	3	1	3	0	5	5	4	6	3	6	2	3
FORT MCMURRAY	9	8	6	8	5	0	8	8	10	6	10	6	5
GRANDE PRAIRIE	7	8	6	8	5	8	0	4	10	8	10	6	4
JASPER	3	5	5	6	4	8	4	0	7	7	8	5	5
LETHBRIDGE	4	3	5	3	6	10	10	7	0	7	2	4	8
LLOYDMINSTER	7	6	3	5	3	6	8	7	7	0	5	4	5
MEDICINE HAT	5	3	5	3	6	10	10	8	2	5	0	5	8
RED DEER	3	2	2	2	2	6	6	5	4	4	5	0	4
SLAVE LAKE	7	6	4	6	3	5	4	5	8	5	8	4	0

B

The Alberta Network

- Network of value units
- Local coding of input cities and output ratings
- Smallest network that worked used 6 hidden units
- 169 training patterns

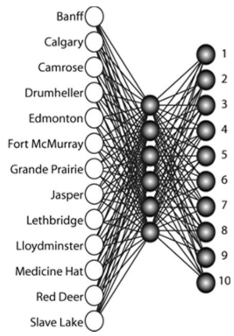

Projected 1D Maps

- Network interpretation was not straightforward
- It was eventually realized that each hidden unit was a one-dimensional map
- A map had a particular orientation
- Connection weights were correlated with distances projected onto this map!

Table 2. Results of relating Alberta map distances between cities and hidden units the values of the connection weights feeding into the hidden units in Simulation 1

Hidden Unit	Hidden Unit Latitude	Hidden Unit Longitude	Correlation Between Map Distances And Incoming Weights
H0	51.72	113.55	0.88
H1	50.84	113.63	0.59
H2	50.88	117.70	0.72
H3	53.39	115.05	-0.54
H4	53.91	113.42	0.79
H5	51.17	115.57	-0.48

Positioning Hidden Units

- Excel Solver was used to place each hidden unit on the map
- A hidden unit is the origin of a 1D map
- Solver also oriented the hidden unit's map
- Solver maximized correlation between weight and distance
- High correlations indicates that this is a plausible account of hidden unit function

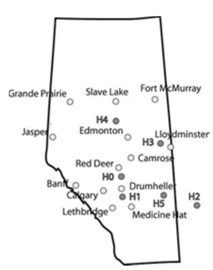

Hidden Unit Properties

- Hidden units were metric
- Individual hidden units, though, had a very inaccurate internal map
- Accuracy of space came from *coarse allocentric coding*!

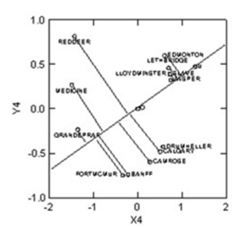

Place Cells And Coarse Allocentric Coding

- Place cells have been criticized as not being 'map-like'
- Our network is not map-like either, but has internalized a map of Alberta
- Hidden units are like place cells
- Perhaps the hippocampus is a PDP map, using coarse allocentric coding – which acts like a map, but doesn't look like one!

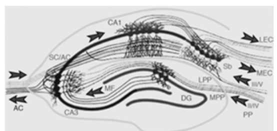

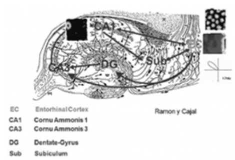

Advantages Of Non-Map-Like Maps

- Why might it be advantageous to have a cognitive map that is not map-like in structure?
- Some spatial reasoning may involve spaces that are not metric
- Map-like representations might fail on such tasks

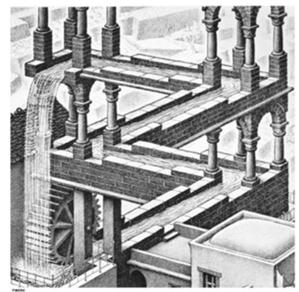

Antisymmetry And Direction

- One can violate metric space property by using antisymmetric relations instead of symmetric ones
- d(x,y) = -d(y,x)
- Compass directions are antisymmetric
- The direction from Edmonton to Calgary is the opposite of the direction from Calgary to Edmonton

$$A = \begin{bmatrix} 0 & -c & b \\ c & 0 & -a \\ -b & a & 0 \end{bmatrix}$$

A Direction Task

- Dawson and Boechler (2007) trained a network to judge the direction between pairs of Albertan cities
- This task is intrinsically nonmetric

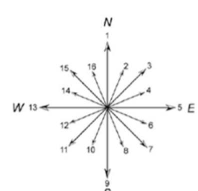

	BANFF	CALGARY	CAMROSE	DRUMHELLER	EDMONTON	FORT MCMURRAY	GRANDE PRAIRIE	JASPER	LETHBRIDGE	LLOYDMINSTER	MEDICINE HAT	RED DEER	SLAVE LAKE
BANFF	0	5	3	5	3	3	15	15	6	4	6	4	3
CALGARY	13	0	2	4	1	2	15	11	7	4	6	1	1
CAMROSE	11	10	0	9	15	2	14	13	9	5	8	11	15
DRUMHELLER	13	12	1	0	16	1	14	11	9	3	7	14	16
EDMONTON	11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

Coarse Coding

- A seven hidden unit network of value units learned the task
- Its hidden units coarse coded compass directions
- Not a topographic map of 'head direction'!

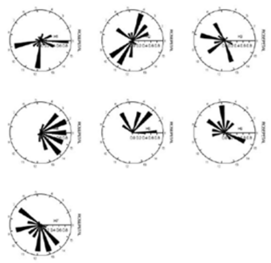

Asymmetry Of Wiretaps

- Dawson and Boechler wiretapped each hidden unit for each pair of input stimuli
- The resulting matrices of activity were highly asymmetric

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

Another Distance Network

- Dawson, Boechler and Orsten (2005) trained a different value unit network on the city distance problem

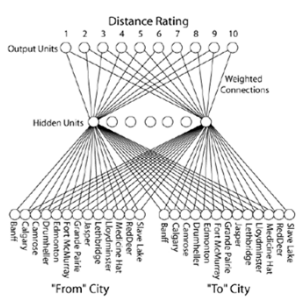

Weights Do Not Code Distance!

- The weights feeding into each hidden unit were highly systematic
- But weights could <u>not</u> be related to distances between cities!
- What systematic feature were hidden units detecting?

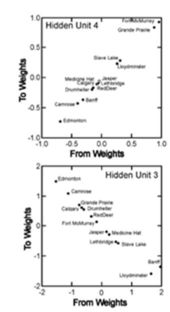

Hidden Units Encode Direction

- **Hidden unit weights encoded <u>direction</u>!**
- **When placed by Solver on the map, weights were correlated with direction from the origin and plane of the hidden unit's 1D map**

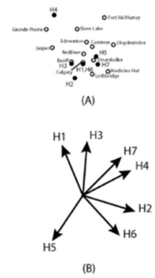

Hidden Sextants

- **Each hidden unit provided distance information by acting like a sextant**

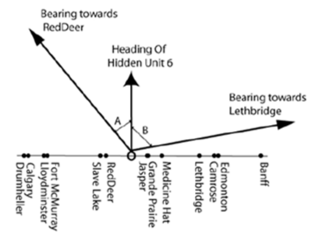

Figure 4. A graphical interpretation of the operation of hidden units in the network, using Hidden Unit 6 as an example. The horizontal line represents the positions of the thirteen cities on the unit's weight space. The circle represents the origin of this space. The unit's heading is given from this origin. Angle A is the bearing towards RedDeer relative to this heading, and Angle B is the bearing towards Lethbridge. The position of a city on the weight space is highly correlated with the cosine of the city's bearing.

Coarse Coding Again

- **As a sextant, each hidden unit is a highly inaccurate measure of distance**
- **Solution: coarse coding**
- **Seven sextant readings from different perspectives combine to measure distance accurately**

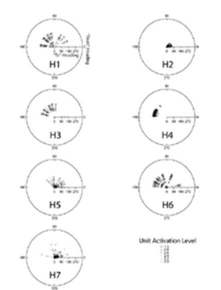

Figure 5. Bubble plots of the activity of each hidden unit as a function of the bearings of two cities (a *from* city and a *to* city) for each pair of cities in the training set. The size of the bubble reflects the degree of hidden unit activity. The coordinates of the bubble indicate the bearings of the two cities. The circular axis provides the bearing of the *from* city, and the horizontal axis provides the bearing of the *to* city.

Functional Maps

- **"The cognitive map is not a picture or image which 'looks like' what it represents; rather, it is an information structure from which map-like images can be reconstructed and from which behavior dependent upon place information can be generated" (O'Keefe & Nadel, 1978)**
- **Our network simulations provide several illustrations of this point: three very different representations that are not map-like, but which function like maps on spatial reasoning tasks of known structure**

John O'Keefe

Lynn Nadel