

Spatial Navigation and the Brain's GPS

Study Guide

by Richard J. Young, Ph.D.


richard@deepneuro.ai

Chapter 3 of "The NeuroAI Handbook"

Place cells, grid cells, and why the brain's map already looks like a graph and an attention head.

One-sentence thesis. The brain navigates by pairing a universal hexagonal ruler (grid cells) with sparse, meaningful pins (place cells), and those two ideas prefigure the graph and the attention mechanism at the heart of modern AI.

PART 1: The Navigation Problem and the Hippocampus



The hippocampal pipeline: input enters through entorhinal cortex, DG separates similar patterns, CA3 completes them from partial cues, CA1 binds and outputs.

The Miracle of Navigation

How do you remember the route to your coffee shop, or how does a squirrel relocate a buried nut after a winter? This ability feels effortless, yet it is one of the brain's most sophisticated feats, run by an internal GPS in the **hippocampal formation**. The chapter's payoff is that the brain's solution to mapping space turns out to rhyme with two workhorses of AI: the graph and the attention mechanism.

The Hippocampal Pipeline

The hippocampus binds the pieces of an experience, *what* happened, *where*, and *when*, into one memory,

and its wiring is a staged pipeline. Sensory input arrives through the **entorhinal cortex (EC)**, the gateway where grid cells live. The **dentate gyrus (DG)** performs *pattern separation*: it forces similar-looking places into distinct representations so memories do not blur together. **CA3** does the opposite, *pattern completion*, recalling a whole memory from a partial cue through its recurrent connections. **CA1** binds and outputs the result. Separation then completion, a build-up of structure from simple inputs, is a theme you will meet again in deep networks.

Place Cells: Pins on the Map

Discovered by John O'Keefe in 1971, a **place cell** is the easiest piece to grasp: a hippocampal neuron that fires only when the animal is in one specific spot, its *place field*. A few properties make it powerful. It **remaps**: enter a new room and the cells instantly build a fresh, independent map, so your "lobby" cell in Paris is not your "lobby" cell in Tokyo. It is **experience-driven**, sharpening for places you care about. And it is **sparse**, only a small fraction of cells are active anywhere, an efficient, robust code that AI prizes.

In symbols, a place field is well modeled as a Gaussian bump peaking at its center $\mathbf{c}$ and fading with distance:

$$r(\mathbf{x}) = A \exp\left(-\frac{\|\mathbf{x} - \mathbf{c}\|^2}{2\sigma^2}\right)$$


ANALOGY

Place cells are smart pins on a mental pinboard. Enter a cafe and a "cafe" pin lights up; walk to the park and a different pin fires. Each stays silent until you return to its spot, and the whole collection is your map, a map of *meaningful* places, not raw coordinates.

Grid Cells: The Universal Ruler

If place cells are the pins, what is the pinboard? Edward and May-Britt Moser found the answer in 2005 (a Nobel-winning discovery): **grid cells** in the entorhi-

nal cortex fire not in one spot but at many, arranged in a perfect **hexagonal lattice** tiling the whole environment.



One grid cell's firing peaks. A hexagonal lattice is the most efficient way to tile a 2D plane.

Three things make the grid special. It is **geometrically optimal**, hexagons tile a plane with the fewest cells. It is **multi-scale**: grid cells cluster into modules of different spacings, fine grids for precision and coarse grids for range. And it is **universal**, the same lattice reappears in every environment, a metric the brain can lay over any new space. The hexagon emerges from the interference of three plane waves set 60 degrees apart:

$$g(\mathbf{x}) = \sum_{k=1}^3 \cos(\mathbf{u}_k \cdot \mathbf{x})$$


The Supporting Cast

Place and grid cells do not work alone. **Head-direction cells** are an internal compass, each tuned to fire when the animal faces one direction regardless of location. **Border cells** fire near walls and edges, anchoring the map to boundaries. Together the four cell types form a complete, flexible cognitive map.

Property	Place cells	Grid cells
Location	Hippocampus	Entorhinal cortex
Firing pattern	One unique place field	Repeating hexagonal lattice
New environment	Remap: a fresh map each time	Stable, a universal metric
Discovered	O'Keefe, 1971	Moser & Moser, 2005 (Nobel)
Role	"You are here" pins	Coordinate ruler
AI echo	Sparse codes, graph nodes	Metric / positional code

The Cognitive Map Is a Graph

Step back and a cognitive map is literally a **graph**: places are nodes, and the paths between them are edges whose strength reflects how easily you can travel between them. That is exactly the object a **graph neural network** is built to reason over. Ask “from the cafe, what is the shortest route to the park?” and a GNN propagates information across the graph to find it, mirroring how the brain plans trajectories through remembered space.



Places as nodes, routes as edges: the cognitive map is a graph a GNN can search.


Path Integration and Its Drift

The brain also tracks position from self-motion alone, **path integration** (dead reckoning), summing velocity over time, with grid cells thought to be the engine:

$$\mathbf{x}(t) = \mathbf{x}_0 + \int_0^t \mathbf{v}(\tau) d\tau$$

The catch is that errors *accumulate*: integrate a slightly

wrong velocity long enough and the estimate drifts away from the truth. This is why an internal GPS alone is never enough.



Dead reckoning drifts: small velocity errors compound into a growing gap.

Landmarks as Attention

To cancel the drift, the brain resets against **landmarks**, a distinctive building, a river, a peak. But which cues matter? When you are lost you do not weigh every tree equally; you fix on the few informative ones. That selective weighting is exactly what an **attention** mechanism does in a Transformer, scoring which pieces of input deserve focus. An AI navigation agent can use the same trick to decide which sensory cues to trust when re-localizing.

ANALOGY

A lost hiker does not study every tree. She locks onto the ridge line and the river bend, two landmarks, and re-fixes her position from them. Attention in a Transformer is that same choice, made continuously: spend focus only where it pays.

Quick Self-Test

1. What does the dentate gyrus (pattern separation) do that CA3 (pattern completion) does not, and why do you need both?
2. When you enter a new room, place cells remap but grid cells stay stable. What computational role does each behavior serve?
3. Why is a hexagonal grid the brain’s choice for a coordinate system rather than a square one?
4. Explain the analogy “a cognitive map is a graph.” What are the nodes and edges, and what AI model consumes that structure?
5. Path integration drifts over time. What corrects it, and which Transformer component does that correction resemble?

If You Remember One Thing

The brain solves navigation by combining two representations, a stable universal ruler (grid cells) and a sparse map of meaningful places (place cells), and stabilizes the result with landmarks. Those are not just biological curiosities: the map-as-graph is a graph neural network, and landmark selection is attention, so the oldest navigation problem in nature keeps handing modern AI its blueprints.

Richard J. Young, Ph.D. · richard@deepneuro.ai · Senior AI Research Scientist, UnitedHealth Group · UNLV Lee Business School · ORCID 0000-0002-1109-7552