Numenta

Massachusetts Institute of Technology

Representing N-dimensional cognitive variables with grid cells

Mirko Klukas [1]
Marcus Lewis [1]
Ila Fiete [2]

[1] Numenta (mklukas@numenta.com)
[2] MIT Department of Brain and Cognitive Sciences

Summary

Representation in a grid module seems inherently 2-dimensional

- Grid cells are defined functionally by their spatially periodic response patterns in 2D space.
- They also represent variables other than location in 2D physical space, e.g. location of visual gaze on 2D image (Killian et al., 2012), locus of covert attention (Wilming et al., 2018), "stretchy bird" (Constantinescu et al., 2016)
- Responses to 1D tracks correspond to slices of the 2D representation (Yoon et al., 2016).
- Anatomically, the cells in a grid module appear to be topographically arranged in cortical sheets in a 2D physical grid that mirrors their tuning to 2D external space (Tank, Dombeck)

Can we go beyond 2 Dimensions?

- Structural constraints seem limiting, grid modules seem to be confined to be inherently 2D
- How generally useful can the grid code be if the autonomous states of a grid module are inherently two-dimensional?
- Raises natural question of what dimensions, theoretically, it is possible for grid cells to represent.

Path-integrating N-dimensional variables

- We show here that a family of grid cell modules, each with 2D responses, can generate unambiguous representations of variables of much larger dimension
- Given an arbitrary N-dimensional velocity vector through a flat space, each module generates a unique linear projection to 2D and processes it independent of the other modules.
- Just as the conjunction of multiple band cells can encode a 2-dimensional location, we describe how M grid modules can represent variables of dimension up to 2M

Capacity.

- We define *capacity* in terms of the side length W of a cubic environment that can be *uniquely* encoded. We show numerically that the capacity grows exponentially in the number of modules M divided by N:

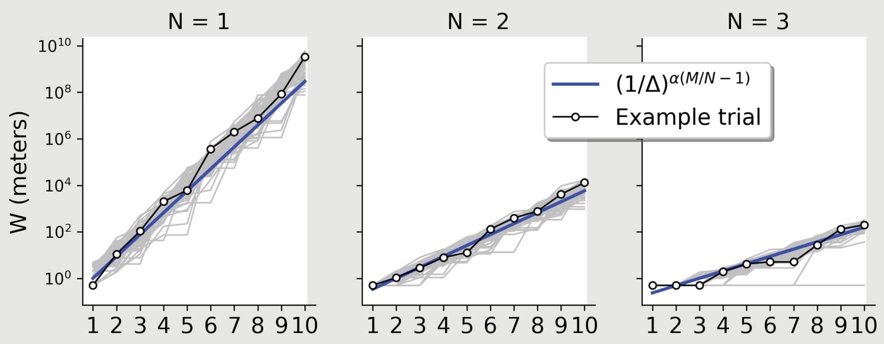

Properties of grid cells

The grid code in 2D

Mammalian grid cells are defined by the arrangement of their firing fields in 2D environments: they fire at multiple locations, with their firing fields lying at the vertices of an equilateral triangular lattice that tiles the explored space. Grid cells are organized into discrete *modules* of cells with a common spatial period and a common orientation. These co-modular cells differ in their 2D spatial phases.

Path-integration

According to the prevailing hypothesis and key models of grid cell dynamics, grid cell modules update their internal state (i.e., the currently active phase φ) by integrating estimates v of animal motion dx/dt in 2D space over time:

$$\frac{d\phi}{dt} = A \cdot v$$

One can think of A as a velocity projection, which relates changes in position in the environment to changes in phase; more concretely, it is the differential of the encoding map from location to phase.

Once the module is *anchored* to a position in the environment the module can generate a grid code for any other location x in the environment reached by a path γ:

$$\phi(x) = \phi_0 + \int A \cdot \tfrac{d\gamma}{dt}\ dt \mod 1.$$

Unique code in 2D

Taken individually, each module provides a highly ambiguous readout of location. It is only possible to determine where the animal is by considering multiple modules with different periods. The range over which grid cells of multiple modules provide a unique code for space grows exponentially with the number of modules (Fiete et al., 2008).

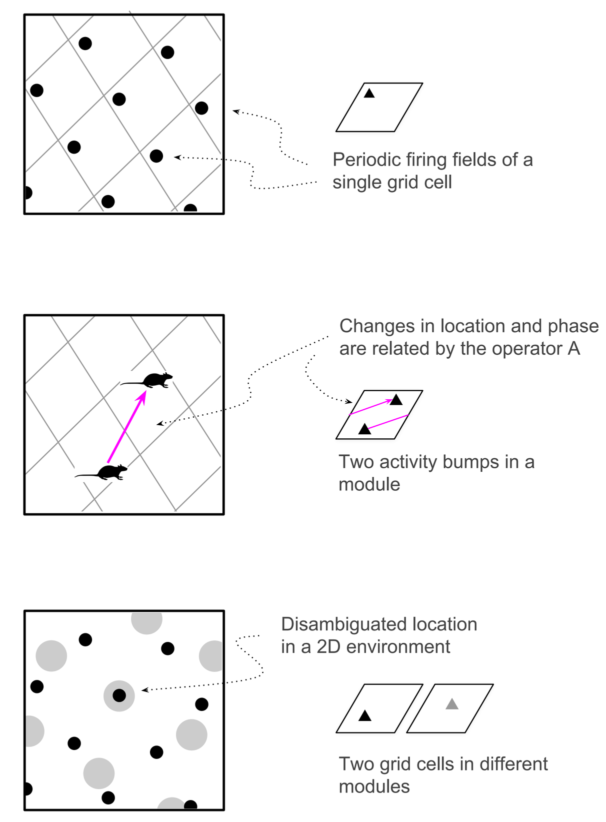

How grid cells can encode N-dimensional variables

How should a 2D module process 3D updates

Suppose an agent leaves the xy-plane and moves in the z-direction. How should a module update its activity? No matter what update rule we follow, there will be a 1-dimensional family of ambiguous positions that can't be distinguished by a single module.

Given an arbitrary *N*-dimensional velocity vector, each module generates a unique linear projection to 2D and processes it independent of the other modules.

That means one could simply replace the velocity projection operator A by one of shape Nx2, thus projecting the N-dimensional velocity down to two dimensions. In this case, given a displacement or velocity associated with the N-dimensional variable, grid cells would faithfully integrate those components that fall along the 2 projected dimensions.

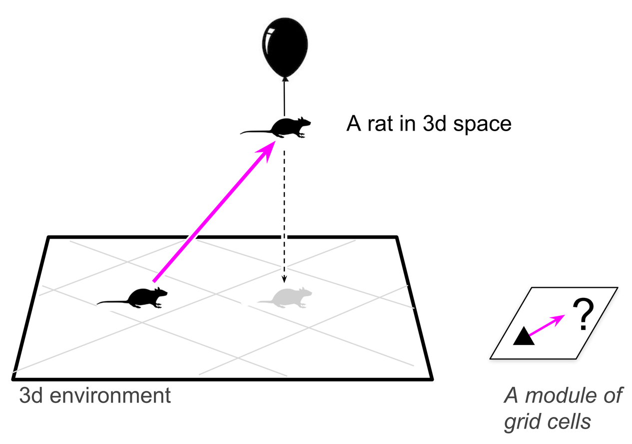

Relation to "Band Cells"

Cells whose spatial responses resemble parallel bands across a 2-dimensional environment, perpendicular to a preferred direction, and at a fixed distance from each other, are colloquially called *band cells*.

Note that, in the same way that we were not tied to inputs from 2D environments, we are not tied to the use of 2D grid modules either and it is not hard to see that all the above applies if we replace the 2D grid modules by 1D versions. The resulting spatial responses of these 1D modules will resemble parallel bands across a 2D environment as well.

Resolving ambiguities

The two types of spatial ambiguities of the 2D grid code, resulting from from the periodicity of the codes per module and from the projection of an N-dimensional variable down to two dimensions, can be simultaneously resolved by combining the codes of multiple modules with different periods and distinct choices for the set of velocity projection operators A.

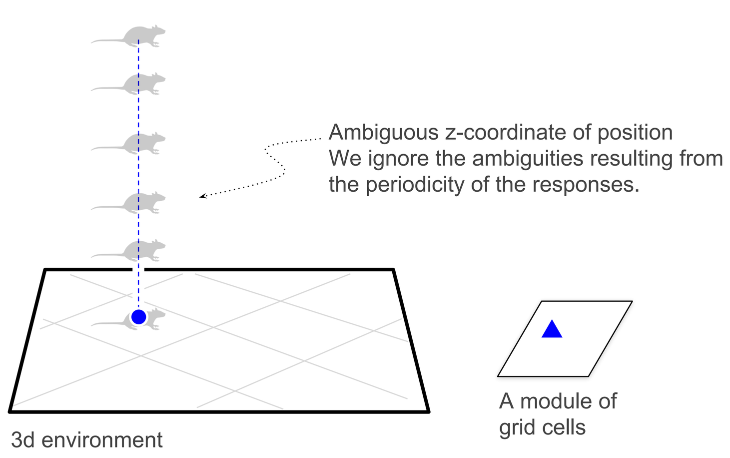

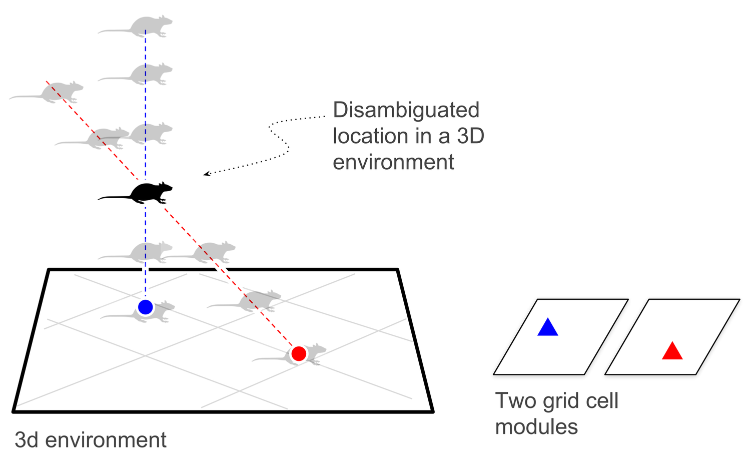

Capacity and tuning curves

The capacity grows exponential in module number

We define "capacity" to be the maximal side length w of a hypercube such that no two points are assigned the same grid code.

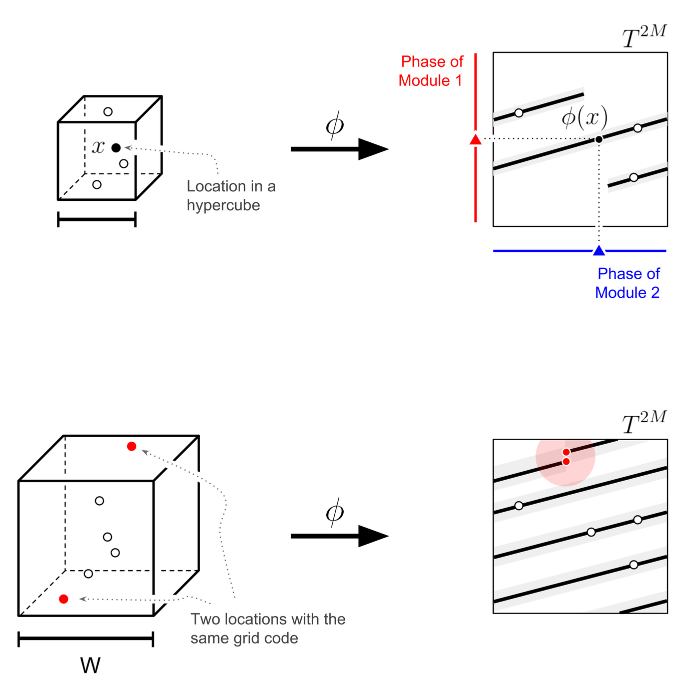

Suppose x were of dimension N>2, then one could simply divide the collection of grid modules into N distinct groups, each of which independently encodes one of the variable's N coordinates. A thorough analysis of the 1D case presented in (Fiete et al.,2008) shows that the 1D capacity grows exponentially with the number of modules. In consequence the capacity for this simple coding scheme above grows exponentially in the Nth fraction of the number of modules:

$$W \propto \left(\frac{1}{\Delta}\right)^{\alpha(M/N-1)}$$

Here Δ denotes the "phase resolution" and α is a parameter that depends on the scales of the modules.

We show numerically that the capacity for our proposed coding scheme behaves similar.

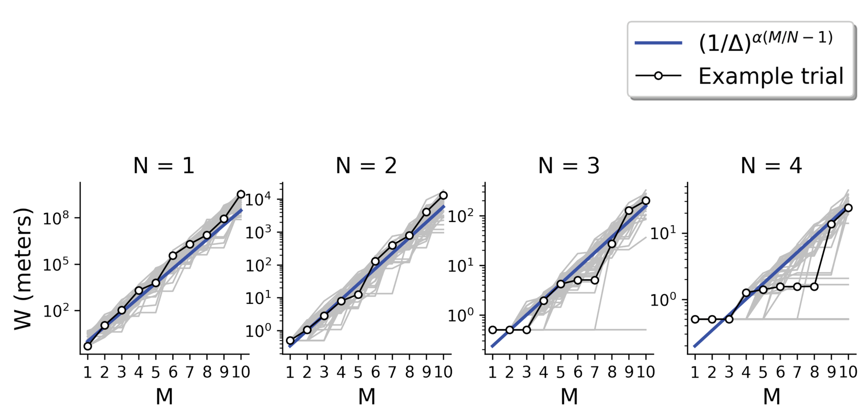

Predicted firing fields

Inspired by an experiment studying firing fields of "grid cells on steeply sloping terrain" carried out in the Jeffery lab (2015) we simulated firing fields on randomly tilted planes in 3 dimensional space induced by our coding scheme.

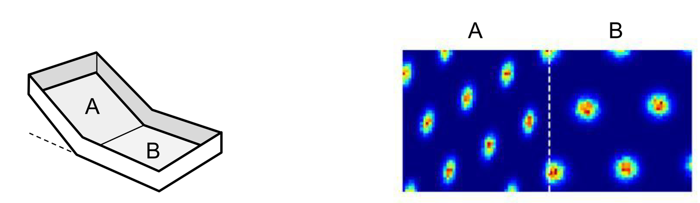

Jeffery Lab (2015)
Grid cells on steeply sloping terrain: evidence for planar rather than volumetric encoding

Firing fields predicted by our coding scheme

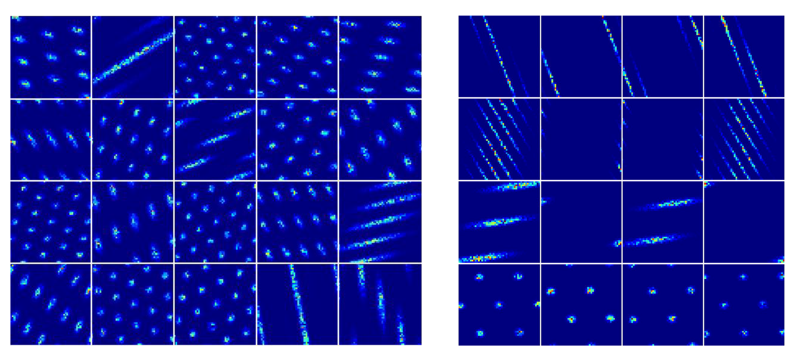

Simulated cell recordings from different grid cells in different modules in response to a N-dimensional variable $v=(v_1, \ldots, v_N)$ mapped against a 2D subspace spanned by two of its components v_i and v_j :

Same as on the left, but each row shows firing fields for 4 cells from the same module.

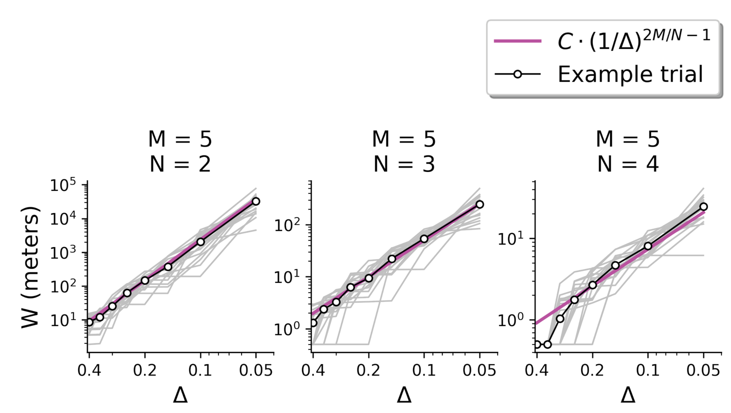

References

Constantinescu, O'Reilly, Behrens, "Organizing Conceptual Knowledge in Humans with a Grid-like Code", Science (2016).

Fiete, Burak, Brookings, "What Grid Cells Convey about Rat Location", Journal of Neuroscience (2008).

Hayman, Casali, Wilson, Jeffery, "Grid cells on steeply sloping terrain: evidence for planar rather than volumetric encoding", Frontiers in Psychology (2015).

Heys, Rangarajan, Dombeck, "The functional micro-organization of grid cells revealed by cellular-resolution imaging", Neuron (2014).

Killian, Jutras, Buffalo, "A map of visual space in the primate entorhinal cortex", Nature (2012).

Wilming, P. König, S. König, Buffalo, "Entorhinal cortex receptive fields are modulated by spatial attention, even without movement", eLife (2018).

Yoon, Lewallen, Kinkhabwala, Tank, Fiete, "Grid Cell Responses in 1D Environments Assessed as Slices through a 2D Lattice", Neuron (2016).