J Comput Neurosci (2007) 22:297–299
DOI 10.1007/s10827-006-0013-7

BRIEF COMMUNICATION

From grids to places

M. Franzius · R. Vollgraf · L. Wiskott

Received: 28 August 2006 / Revised: 3 November 2006 / Accepted: 13 November 2006 / Published online: 29 December 2006

Abstract Hafting et al. (2005) described *grid cells* in the dorsocaudal region of the medial entorhinal cortex (dMEC). These cells show a strikingly regular grid-like firing-pattern as a function of the position of a rat in an enclosure. Since the dMEC projects to the hippocampal areas containing the well-known place cells, the question arises whether and how the localized responses of the latter can emerge based on the output of grid cells. Here, we show that, starting with simulated grid-cells, a simple linear transformation maximizing sparseness leads to a localized representation similar to place fields.

Keywords Place cell · Grid cell · Hippocampus · Entorhinal cortex

As reported by Hafting et al. (2005) grid cells in the dMEC show spatial firing patterns in the form of hexagonal grids with frequencies within one octave (39 to 73 cm mean distance), random phase shifts, and random orientations. The firing patterns of place cells in the hippocampus, on the other hand, are localized spots (Muller, 1996). Our hypothesis is that the latter can be generated from the former simply by sparsification, which is consistent with evidence that firing patterns in hippocampal regions CA1 and CA3 are sparser than in entorhinal cortex (O'Reilly and McClelland, 1994).

To show that place fields can be derived from grid-cells by sparsification we simulated a fully connected linear two-layer network. The input units were 100 simulated grid-cells of a virtual rat with activity patterns synthesized by Gaussians arranged on a hexagonal grid (Fig. 1(A)). Some positional jitter, random anisotropy, and amplitude variation of the Gaussians was introduced and white noise was added to qualitatively match the slightly irregular experimental data.

Let $g_i(\vec{r})$ denote the activity of grid-cell g_i as a function of location $\vec{r}$. Given a virtual path $\vec{r}(t)$ of a rat within the enclosure, the input into the hippocampus coming from the grid-cells is $x_i(t) := g_i(\vec{r}(t))$. To achieve sparseness we applied independent component analysis (ICA) (Hyvärinen, 1999b) on a set of 200.000 time points on the full set of 100 inputs by subtracting the mean and using the CuBICA algorithm, which attempts to diagonalize the tensors of third and fourth order cumulants (Blaschke and Wiskott, 2004), but we have obtained similar results with other sparsification algorithms, such as FastICA (Hyvärinen, 1999a) or simply maximizing peak activity under a unit variance, zero mean, and decorrelation constraint. The sign of each output unit, which is arbitrary for ICA, was chosen such that the value with the largest magnitude is positive, and then constants c_j were added to ensure nonnegative values. This yielded an affine transformation with matrix T producing 100 output signals $y_j(t) := \sum_i T_{ji} x_i(t) + c_j$ that are maximally independent and significantly sparser than the input signals (kurtosis increased on average from 2.8 for the input units to 27.3 for the output units). The output-unit activities as a function of location are $p_j(\vec{r}) := \sum_i T_{ji} g_i(\vec{r}) + c_j$ and show localized place fields (Fig. 1(G)). We measured the number of peaks in a unit's output by counting the number of distinct contiguous areas containing pixels with at least 50% of the

Action Editor: Alessandro Treves

M. Franzius (✉) · L. Wiskott
Institute for Theoretical Biology, Humboldt-University, Berlin, Germany
e-mail: m.franzius@biologie.hu-berlin.de

L. Wiskott
e-mail: l.wiskott@biologie.hu-berlin.de

R. Vollgraf
Department of Computer Science, Technical University of Berlin, Germany
e-mail: vro@cs.tuberlin. de

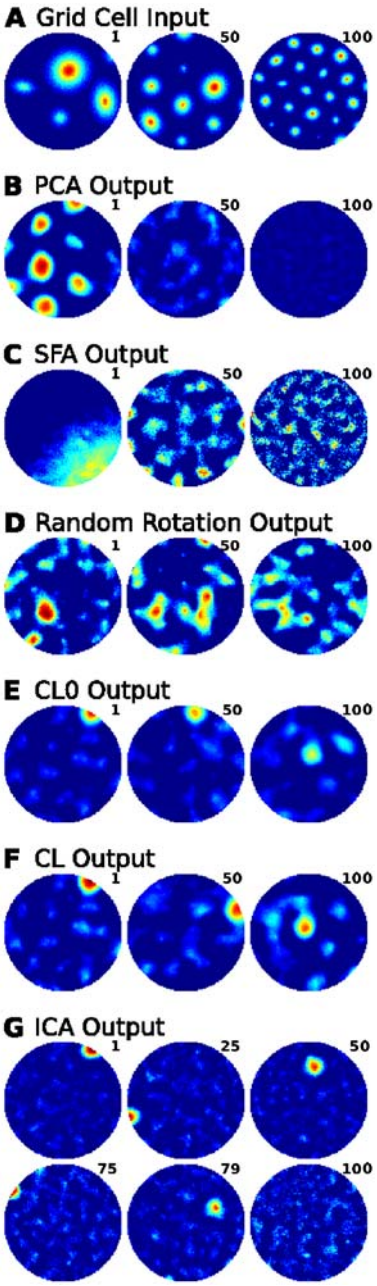

Fig. 1 (A) Spatial firing pattern (SFP) of the input units representing grid-cells. Three out of 100 units are shown. (B) SFP of 1st, 50th, and 100th output computed by principal component analysis, ordered by eigenvalues. (C) SFP of 1st, 50th and 100th output computed by Slow Feature Analysis, ordered by slowness. (D) SFP of three out of 100 typical output units computed by random mixtures of the inputs. (E) SFP of 1st, 50th and 100th output after initialization with sample vectors. Units are ordered by decreasing sparseness (kurtosis). (F) 1st, 50th and 100th output after competitive learning, ordered by kurtosis. (G) SFP of six out of 100 output units computed by independent component analysis as a means of sparsification, ordered by kurtosis. Place fields of sparser units tend to have higher peak activity and are more often located at the border of the enclosure, whereas less sparse units tend to have multiple place fields. Activities are color coded: red-high, green-medium, blue-zero activity. The full set of results can be viewed at http://itb.biologie.hu-berlin.de/~franzius/gridsToPlaces/

unit's maximum activity. A large proportion of output units (75%) show a single spot of activity (Fig. 1(G), units 1, 25, 50, 75), some units (6%) show few spots (Fig. 1(G), unit 79), both being consistent with the patterns of physiological place-cells. Only few output units (19%) show patterns of activity without clear structure (Fig. 1(G), unit 100). The size of the resulting place fields is similar for most units and comparable to the size of the smallest grid-cell fields, but it also depends on the number of grid cell inputs. More inputs lead to more localized output fields while too few inputs can increase the number of fields per output unit (note that throughout this paper the number of output unit is the same as the number of input units and the connectivity is complete).

There are different ways of achieving sparseness and localized place fields. We have used ICA here and have obtained similar results by maximizing peak activity. For a biologically more plausible implementation, we use competitive learning (CL). The weights of the units are initialized with the firing rate of the grid cells at a particular location, with a different location for each unit. This is to avoid "dead units", i.e. units that never win the competition and thus never learn, but since in any given environment a significant proportion of place cells is inactive, a random initialization leading to some "dead units" might be considered realistic as well. In our case, the resulting code already is fairly sparse and localized (mean kurtosis: 9.9, number of units with single peak: 49, see Fig. 1(E)). After competitive learning, kurtosis increases to 10.2 and the number of units with single peaks increases to 60 (Fig. 1(F)). Furthermore, the output units are less correlated after competitive learning than before (mean absolute correlation 0.014 vs. 0.189).

There are other linear transformations, however, that do not lead to localized place fields. As controls we have applied random mixtures, principal component analysis (PCA), and slow feature analysis (SFA; Wiskott and Sejnowski, 2002) to the grid cell input. The latter minimizes the mean squared time derivative of the outputs and has been chosen because Wyss et al. (2006) have presented a model based on the slowness principle that was able learn localized place cells. As one would expect, with random rotations of the input the results retain some grid structure but are less regular than the input (Fig. 1(D)) and no unit has one single or two peaks of activity. With PCA the first units (i.e. those with highest variance) are highly structured and have large amplitudes, much like the grid cells themselves, while the later low-variance units have low amplitudes and are noise-like (Fig. 1(B)). None of these units had a single or two peaks of activity. From the temporal slowness objective we would expect patterns with low spatial frequencies first, and high-frequency non-localized patterns later, when outputs are sorted by slowness (Fig. 1(C)). None of these outputs have only one or two peaks of activity. None of these three

alternative linear transformations (Fig. 1(B)–(D)) leads to localized place fields. Different starting conditions may lead to different results, but for all learning rules considered 5 out of 5 simulations showed the same qualitative behavior.

We conclude that sparse coding is a simple and efficient computational approach for the generation of place cells from grid cells. The mean kurtosis and percentage of localized place fields increase from 9.9 and 49% for the simple initialization with input vectors over 10.2 and 60% after competitive learning to 27.3 and 75% for the ICA algorithm, respectively. Other methods we have tested, such as random rotations, PCA, and SFA fail completely in generating localized place fields. The fact that SFA fails in our simulations is inconsistent with the results from Wyss et al. (2006). Possibly, their model contains some hidden mechanisms that favor sparseness in addition to slowness. The simple initialization with input vectors is extremely quick and already fairly efficient (cf. McNaughton et al., 2006). Such a simple mechanism might be a way for the almost instantaneous formation of place fields in a new environment. However, competitive learning still improves on that significantly while preserving many of the place fields chosen by the initialization process (Units 1 and 100 in Fig. 1(E) and (F) maintained their place field while Unit 50 did not). Thus, competitive learning (or any other sparsification method) could be used as a refinement. ICA once again improves on the results of competitive learning but is biologically less plausible. There is some indication that grid cells reshuffle their phases if the animal is placed in a new environment (McNaughton, 2006). We have found that this results in output units like those with random rotations even if the place fields were localized before the reshuffling. Thus, in our linear mode, for a successful remapping either the phases would have to change in some coherent way or the connectivity has to readapt. We believe the latter is more likely and we have seen above that it can be done rather quickly. However, even if sparseness is efficient in creating place fields from grid cells, the complexity of place-field formation is now only shifted to the computation of grid-cell behavior and still open for discussion.

Acknowledgments This research was funded by the Volkswagen Foundation (MF, LW) and the Wellcome Trust (RV: 10008261).

References

Blaschke T, Wiskott L (2004) CuBICA: Independent component analysis by simultaneous third- and fourth-order cumulant diagonalization. IEEE Trans. Signal Process. 52(5): 1250–1256.

Hafting T, Fyhn M, Molden S, Moser M, Moser EI (2005) Microstructure of a spatial map in the entorhinal cortex. Nature 436(7052):801– 806. 10.1038/nature03721.

Hyvärinen A (1999a) Fast and robust fixed-point algorithms for independent component analysis. IEEE Trans. Neural Netw. 10: 626–634.

Hyvärinen A (1999b) Survey on independent component analysis. Neural Comput. Surv. 2: 94–128.

McNaughton B, Battaglia F, Jensen O, Moser E, Moser M (2006) Path integration and the neural basis of the 'cognitive map'. Nat. Rev. Neurosci. 7: 663–678.

Muller R (1996) A quarter century of place cells. Neuron 17:813–822.

O'Reilly RC, McClelland JL (1994) Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. Hippocampus 4(6): 661–682.

Wiskott L, Sejnowski T (2002) Slow feature analysis: Unsupervised learning of invariances. Neural Comput. 14(4): 715–770.

Wyss R, König P, Verschure P (2006) A model of the ventral visual system based on temporal stability and local memory. PLOS Biol. 4(5): e120.