Journal of Neurophysiology, 79: 1017-1044 (1998)

Interpreting neuronal population activity by reconstruction: A unified framework with application to hippocampal place cells

Kechen Zhang,[1] Iris Ginzburg,[1] Bruce L. McNaughton,[2] Terrence J. Sejnowski [1,3]

[1]Howard Hughes Medical Institute
Computational Neurobiology Laboratory
The Salk Institute for Biological Studies
La Jolla, California 92037

[2]Arizona Research Laboratories
Division of Neural Systems, Memory, and Aging
and Department of Psychology
University of Arizona
Tucson, Arizona 85724

[3]Department of Biology
University of California, San Diego
La Jolla, California 92093

Abstract

Physical variables such as the orientation of a line in the visual field or the location of the body in space are coded as activity levels in populations of neurons. Reconstruction or decoding is an inverse problem in which the physical variables are estimated from observed neural activity. Reconstruction is useful first in quantifying how much information about the physical variables is present in the population, and second, in providing insight into how the brain might use distributed representations in solving related computational problems such as visual object recognition and spatial navigation. Two classes of reconstruction methods, namely, probabilistic or Bayesian methods and basis function methods, are discussed. They include important existing methods as special cases, such as population vector coding, optimal linear estimation and template matching. As a representative example for the reconstruction problem, different methods were applied to multi-electrode spike train data from hippocampal place cells in freely moving rats. The reconstruction accuracy of the trajectories of the rats was compared for the different methods. Bayesian methods were especially accurate when a continuity constraint was enforced, and the best errors were within a factor of two of the the information-theoretic limit on how accurate any reconstruction can be, which were comparable with the intrinsic experimental errors in position tracking. In addition, the reconstruction analysis uncovered some interesting aspects of place cell activity, such as the tendency for erratic jumps of the reconstructed trajectory when the animal stopped running. In general, the theoretical values of the minimal achievable reconstruction errors quantify how accurately a physical variable is encoded in the neuronal population in the sense of mean square error, regardless of the method used for reading out the information. One related result is that the theoretical accuracy is independent of the width of the Gaussian tuning function only in two dimensions. Finally, all the reconstruction methods considered in this paper can be implemented by a unified neural network architecture, which the brain could feasibly use to solve related problems.

1 Introduction

Reconstruction is a useful strategy for analyzing data recorded from populations of neurons, in which external physical variables such as the orientation of a light bar on a screen, the direction of hand movement in space or the position of a freely moving animal in space are estimated from brain activity. Reconstruction is sometimes called decoding, where "coding" describes the tuning of neuronal responses to the variables of interest. A wide range of neural coding and decoding problems have been studied previously (Optican & Richmond, 1987; Bialek, Rieke, de Ruyter van Steveninck & Warland, 1991; Seung & Sompolinsky, 1993; Zohary, Shadlen & Newsome, 1994; Salinas & Abbott, 1994; Abbott, 1994; Snippe, 1996; Zemel, Dayan & Pouget, 1997).

Two main goals for reconstruction are approached in this paper. The first goal is technical and is exemplified by the population vector method applied to motor cortical activities during various reaching tasks (Georgopoulos, Schwartz & Kettner, 1986; Georgopoulos, Lurito, Petrides, Schwartz & Massey, 1989; Schwartz, 1994) and the template matching method applied to disparity selective cells in the visual cortex (Lehky & Sejnowski, 1990) and hippocampal place cells during rapid learning of place fields in a novel environment (Wilson & McNaughton, 1993). In these examples, reconstruction extracts information from noisy neuronal population activity and transforms it to a more explicit and convenient representation of movement and position. In this paper, various reconstruction methods that are theoretically optimal under different frameworks are considered; the ultimate theoretical limits on the best achievable accuracy for all possible methods are also derived.

Our second goal for reconstruction is biological. Since the brain extracts information distributed among the activity of populations of neurons in order to solve various computational problems, the question arises as to which algorithms might feasibly be used in the brain. Contrary to previous views, the various reconstruction methods including the Bayesian methods can be implemented by a simple feedforward neural network. This demonstrates that biological systems have the resources needed to implement the most efficient reconstruction algorithms, which can reach the Cramér-Rao lower bound under suitable conditions. In addition, our theoretical results on minimal reconstruction error become immediately relevant for understanding the accuracy of population coding in biological systems and its dependence on the tuning parameters of individual cells.

In this paper hippocampal place cells serve as a primary example of the reconstruction problem and are used for testing different methods. Some interesting biological properties of place cells are revealed by the reconstruction, which illustrates the power of this approach for studying populations of neurons.

2 Description of Place Cell Firing

2.1 A representative example of reconstruction

Position reconstruction based on hippocampal place cell activity provides an excellent example that captures all major aspects of reconstruction (Wilson & McNaughton, 1993). The task is to infer the position of the rat's head based on the patterns of spiking activity from simultaneously recorded place cells. As illustrated in Figure 1, a place cell is often quiet and fires maximally only when the animal's head is within certain restricted regions in the environment, called the place fields (O'Keefe & Dostrovsky, 1971; Olton, Branch & Best, 1978; McNaughton, Barnes & O'Keefe, 1983; Muller, Kubie & Ranck, 1987).

Successful simultaneous recordings from many cells in the hippocampus of a freely behaving rat make it possible to study population activity and reconstruction in single trials. In most other systems, such as recordings from motor cortex during reaching, neurons are usually recorded sequentially during the same task and the population is examined later under the assumption that the cells would have similar statistics if recorded simultaneously. In single-trial reconstruction, phenomena such as erratic jumps in reconstructed trajectory can easily become apparent (Figure 9).

It is important to realize that the general formulation of the reconstruction problem does not rely on any unique properties of place cells. In other words, the same approach used for place cells could be applied to other reconstruction problems without essential modification. The vector $\mathbf{x}$, which is interpreted in these recordings as the position of the animal in the maze, could in general be interpreted as *any external physical variables with which neural activity is correlated.*

2.2 Place cell data

The experimental data analyzed in this paper were obtained by methods that have been described by Wilson & McNaughton (1993), Skaggs, McNaughton, Wilson & Barnes (1996), and Barnes, Suster, Shen & McNaughton (1997). The accuracy of spike timing data of neurons from areas CA1 or CA3 of the right hippocampus was 0.1 msec. The head position was tracked by infrared diodes on the head stage that were sampled at 20 Hz and at a resolution of 256×256 pixels (about 111 cm $\times$ 111 cm). Several channels of hippocampal local field potentials were recorded at 200 Hz or 1000 Hz.

Data were obtained from two young rats while running continuously in familiar mazes. Animal 1 was running on the elevated tracks of a figure-8 maze (Figure 1), with clockwise movement in the upper loop and counterclockwise movement in the lower loop. The outer dimensions of the maze were 93.5 cm $\times$ 80.7 cm, and width of the tracks was 6.4 cm (Figure 1). Animal 2 was running on the tracks of a rectangular maze, physically identical to the upper loop of the figure-8 maze. Both animals stopped briefly at the corners of the mazes for food reward during the recordings.

In reconstruction, the first 7 minutes of data from animal 1 were used to compute firing rate maps and visiting probabilities, and the subsequent 7 minutes of data were used for reconstruction. For animal 2, the sampling and testing intervals were 10 minutes each. For comparison, we also used the second halves of data for sampling and the first halves for testing. In another analysis, the data were divided into consecutive 3-minute time blocks, with only a single block chosen for sampling, leaving all remaining blocks for testing.

2.3 Basic description

The goal of this section is to quantify the basic tuning property of a place cell. Reconstruction only needs two distribution functions derived from the data. The first one is the probability $P(\mathbf{x})$ for the animal to visit each spatial position $\mathbf{x} = (x, y)$, a vector that denotes the position of the animal's head on the horizontal plane. In our dataset, $\mathbf{x}$ was discretized on 256×256 grid and sampled at 20 Hz. Let $N(\mathbf{x})$ be the number of times the animal was found at position $\mathbf{x}$ during a time period of, say several minutes; then the probability of finding the animal at $\mathbf{x}$ can be estimated as

$$P(\mathbf{x}) = \frac{N(\mathbf{x})}{\sum_{\mathbf{x}} N(\mathbf{x})}. \tag{1}$$

This probability is sometimes called *spatial occupancy.*

The second distribution needed by reconstruction is the average firing rate $f(\mathbf{x})$ of a place cell for each position $\mathbf{x}$. The firing rate distribution $f(\mathbf{x})$ is sometimes called the *tuning function* or the

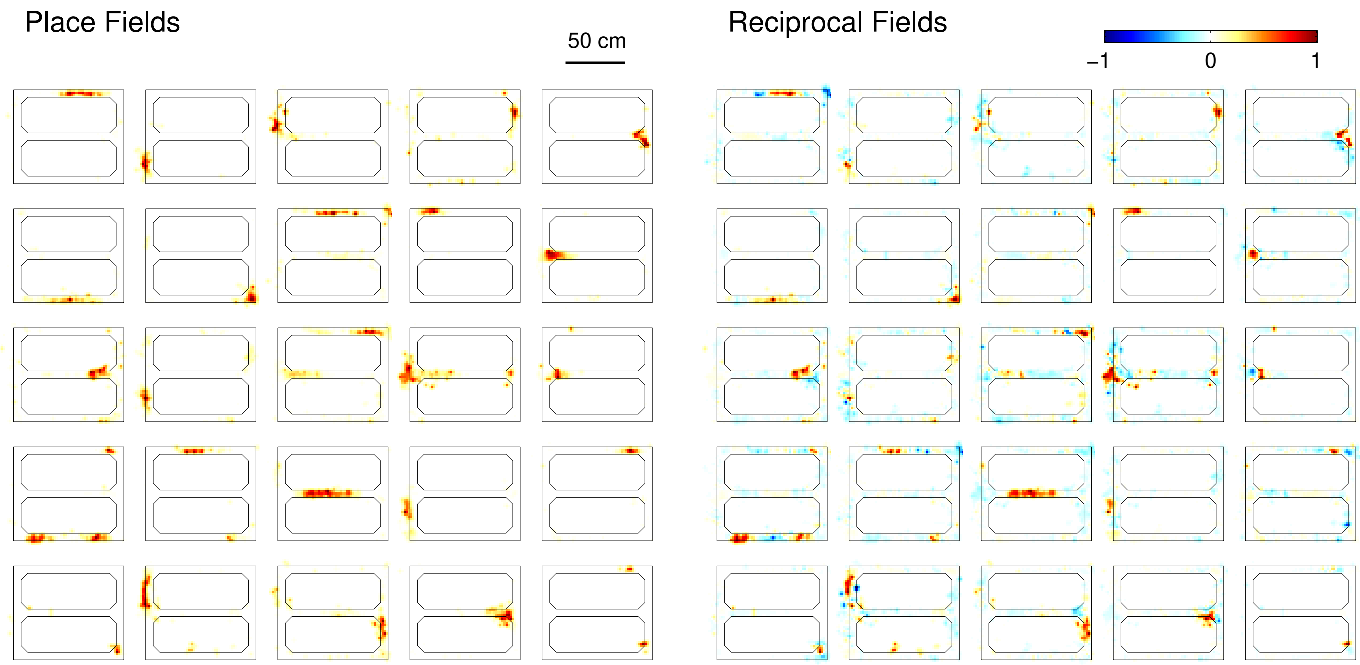

Figure 1: Firing rate maps of 25 hippocampal place cells simultaneously recorded in a rat (animal 1) running on the elevated track of a figure-8 maze, with clockwise movement in the upper loop and counterclockwise movement in the lower loop, and slowing down at the corners for food. The restricted regions with high firing rates are called the place fields. Another 32 simultaneously recorded cells in the same animal were active during sleep but virtually silent on this maze, and therefore were not shown. The maps were computed from 7 minutes of continuous data. The reciprocal fields were obtained from the firing rate maps using a pseudoinverse. In each plot, the scaling is linear with zero value corresponding to 0 in the color map and maximum positive value corresponding to 1. Here the reciprocal fields roughly resemble the place fields, but have negative values in surrounding regions.

firing rate map (Figure 1). Let $S(\mathbf{x})$ be the total number of spikes collected while the animal was at location $\mathbf{x}$, then the average firing rate at $\mathbf{x}$ is estimated as

$$f(\mathbf{x}) = \frac{S(\mathbf{x})}{N(\mathbf{x})\Delta t}, \tag{2}$$

where Δt is the time interval of position tracking so that $N(\mathbf{x})\Delta t$ is the total time spent at position $\mathbf{x}$. The derived firing rate map $f(\mathbf{x})$ was smoothed by convolving it with a Gaussian kernel (Parzen window) with a standard deviation of 1 cm. In theory the time average of the firing rate should equal to the spatial average $f_{\text{mean}} = \sum_{\mathbf{x}} f(\mathbf{x})P(\mathbf{x})$. This is no longer exactly true after the aforementioned smoothing, and can therefore be used as a cross check.

The firing rate distribution $f(\mathbf{x})$ can be interpreted in terms of conditional probability as follows:

$$f(\mathbf{x}) = f_{\text{mean}} \frac{P(\mathbf{x}|\text{“spike”})}{P(\mathbf{x})}. \tag{3}$$

This is equivalent to Eq. (2), as is readily verified by using Eq. (1) and the definition of two additional quantities, namely, the conditional probability

$$P(\mathbf{x}|\text{“spike”}) = \frac{S(\mathbf{x})}{\sum_{\mathbf{x}} S(\mathbf{x})} \tag{4}$$

for *an arbitrarily chosen spike* to fall into the spatial position $\mathbf{x}$, and the overall average firing rate

$$f_{\text{mean}} = \frac{\sum_{\mathbf{x}} S(\mathbf{x})}{\sum_{\mathbf{x}} N(\mathbf{x})\Delta t}, \tag{5}$$

with $\sum_{\mathbf{x}} S(\mathbf{x})$ being the total number of spikes and $\sum_{\mathbf{x}} N(\mathbf{x})\Delta t$ being the total time of recording. The firing rate distribution $f(\mathbf{x})$ describes the intrinsic tendency for the cell to fire at each position $\mathbf{x}$, and is independent of how often the animal visits this position.

Our method of computing the firing rate map is similar to that in Muller, Kubie & Ranck (1987). An alternative method with adaptive binning has been used by Skaggs, McNaughton, Wilson & Barnes (1996). Additional consideration of the heading direction and the intrinsic directionality of place cells (McNaughton, Barnes & O'Keefe, 1983; Muller, Bostock, Taube & Kubie, 1994) may be formulated by running velocity modulation of firing rate as in the next section.

2.4 Firing rate modulation

The description in the preceding section can be broadened to incorporate the modulation of firing rate by additional variables other than the position $\mathbf{x}$, such as running speed, the heading direction, and the theta rhythm. For our dataset, including firing rate modulation only slightly improved reconstruction accuracy.

Running velocity

The firing rate of a place cell typically increases with running speed and also depends on the heading direction when movement is confined to a narrow track (McNaughton, Barnes & O'Keefe, 1983). Both the effects of speed and direction can be modeled as multiplicative modulation of firing rate by the running *velocity*. For our dataset, the animals run mostly in one direction, and we need only to consider the firing rate modulation by speed.

In general, suppose the average firing rate $f(\mathbf{x}, \mathbf{v})$ of a cell at position $\mathbf{x}$ also depends on the instantaneous running velocity $\mathbf{v}$. Consider multiplicatively separable firing rate modulation of the form:

$$f(\mathbf{x}, \mathbf{v}) = F(\mathbf{x})G(\mathbf{v}), \tag{6}$$

where functions $F(\mathbf{x})$ and $G(\mathbf{v})$ depend on $\mathbf{x}$ and $\mathbf{v}$ only, respectively. The simplest case is linear speed modulation:

$$f(\mathbf{x}, \mathbf{v}) = F(\mathbf{x})(av + b). \tag{7}$$

That is, $G(\mathbf{v}) = av + b$ is a linear function of the speed $v = |\mathbf{v}|$, with a and b being constant.

Equation (7) of linear speed modulation is a reasonable approximation for place cells, at least when averaged over a population [McNaughton, Barnes & O'Keefe (1983); see also Figure 8]. A similar modulation by speed occurs for cells in the motor cortex of monkeys. Schwartz (1993, 1994) has shown that adding up the population vector head-to-tail approximately reproduces the hand trajectory. This implies that the length of the population vector is roughly proportional to reaching speed. Thus, to a first approximation, the reaching speed modulates the average firing rate linearly also in the motor cortex.

Theta rhythm

The timing of spike firing relative to the phase of the theta waves in the hippocampal local field potential encodes additional information about the location of the rat (O'Keefe & Recce, 1993; Skaggs, McNaughton, Wilson & Barnes, 1996; Tsodyks, Skaggs, Sejnowski & McNaughton, 1996; Jensen & Lisman, 1996; Samsonovich & McNaughton, 1996; Brown, Frank & Wilson, 1996). Using phase information would be equivalent to splitting a single place field into several smaller regions while preserving the total number of spikes. How this might affect reconstruction accuracy is described quantitatively by Eq. (53). However, information from field potential recordings was not used here because with 20–30 simultaneously recorded cells, the optimal time window for reconstructing position was around 1 second (Figure 6), which spanned around 8 theta cycles and averaged out the phase information from a single cycle.

3 Linear Approach: Basis Functions

3.1 Direct basis and reciprocal basis

A basis function method uses a linear combination of fixed templates, or basis functions, with the coefficients in the linear combination proportional to the activity of the neurons. The template functions can be chosen arbitrarily. Define n_i as the number of spikes fired by cell i within the time window, and $\phi_i(\mathbf{x})$ as an arbitrary basis function or template function associated with this cell. The basic computation is the linear sum

$$\sum_i n_i \phi_i(\mathbf{x}), \tag{8}$$

which is a distribution function over two-dimensional physical space indexed by $\mathbf{x}$. The peak position of this distribution can be taken as the reconstructed position $\hat{\mathbf{x}}$ of the animal; that is,

$$\hat{\mathbf{x}}_{\text{basis}} = \arg\max_{\mathbf{x}} \sum_i n_i \phi_i(\mathbf{x}), \tag{9}$$

where arg max means the value of $\mathbf{x}$ that maximizes the function.

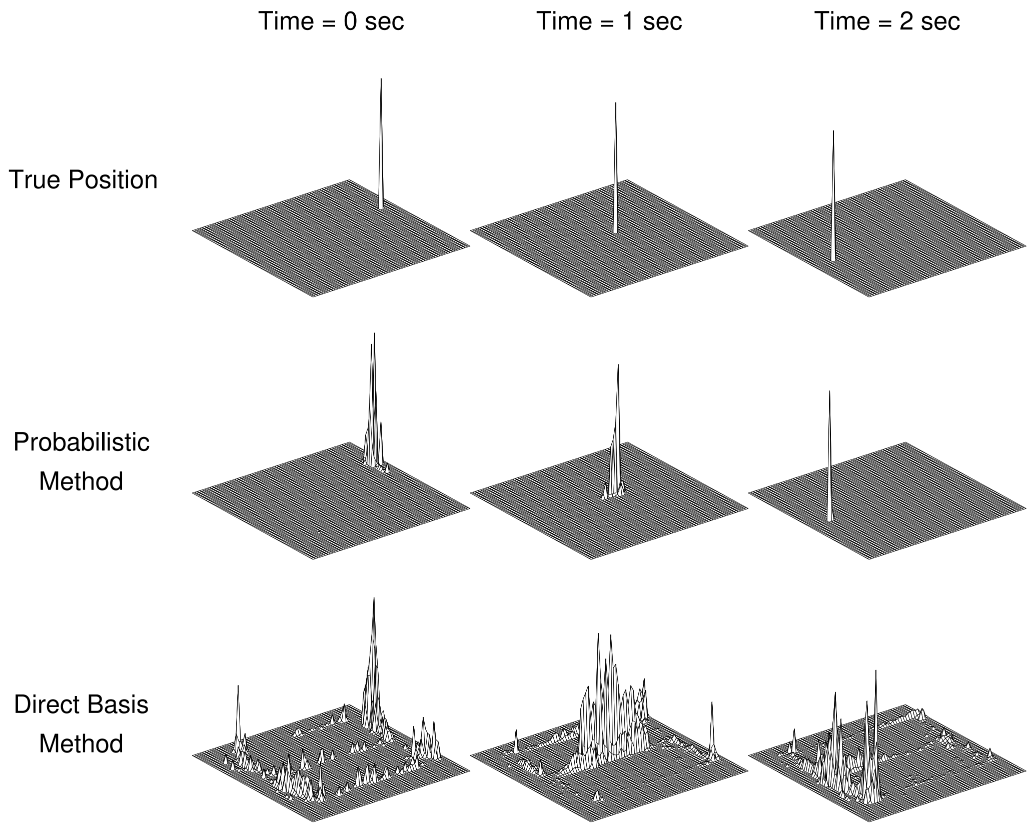

Figure 2: Snapshots of the reconstructed distribution density for the position of a rat (animal 1) as compared with its true position. The true position occupied a single pixel on the 64×64 grid, corresponding to $111\,\mathrm{cm} \times 111\,\mathrm{cm}$ in real space. The probabilistic method (Bayesian 1-step) often yielded a sharp distribution with a single peak, whereas the direct basis method typically led to a broader distribution with multiple peaks. The total numbers of spikes collected from all 25 cells for the three snapshots were 5, 22, and 31, respectively. The time window for each snapshot was 0.5 second.

The basis function framework encompasses existing methods such as the population vector method (Georgopoulos, Schwartz & Kettner, 1986), the template matching method (Lehky & Sejnowski, 1990; Wilson & McNaughton, 1993), and the optimal linear estimator (Salinas & Abbott, 1994) as special cases. A more detailed discussion of these methods follows in a later section.

The major question in this approach is how to choose template functions with minimal reconstruction error. It is useful to make a distinction between a direct basis and a reciprocal basis. A *direct basis* is identified with the firing rate distribution, which can be measured experimentally. A *reciprocal basis*, which arises naturally in minimizing the mean square error, is related to the Moore-Penrose pseudoinverse. Here reciprocal means that applying the pseudoinverse twice reverts back to the original basis.

To formalize the method, let $f_i(\mathbf{x})$ be the average firing rate of cell i when the animal is at position $\mathbf{x}$. Then, by definition, the expected number of spikes n_i this cell fires within a time window of unit length should be

$$n_i = \sum_{\mathbf{x}} \rho(\mathbf{x}) f_i(\mathbf{x}), \tag{10}$$

where the distribution function $\rho(\mathbf{x})$ describes the position of the animal and is 1 at the animal's current position and is zero elsewhere. Equation (10) remains valid even when $\rho(\mathbf{x})$ is an arbitrary probability distribution density.

The goal is to recover the position distribution $\rho(\mathbf{x})$ by combining some fixed template function $\phi_i(\mathbf{x})$ associated with each cell i. For an ideal reconstruction, we should have

$$\rho(\mathbf{x}) = \sum_{i=1}^{N} n_i \phi_i(\mathbf{x}) \tag{11}$$

for an arbitrary position distribution $\rho(\mathbf{x})$. In practice, we seek an approximate reconstruction by minimizing the error.

To find the optimal template functions $g_i(\mathbf{x})$, we use vector-matrix notation to rewrite Eq. (10) as

$$\mathbf{n} = \mathbf{F}^{\mathrm{T}} \mathbf{I}_k, \tag{12}$$

where $\mathbf{n} = (n_1, n_2, \ldots, n_N)^{\mathrm{T}}$ is a column vector, $\mathbf{F} = [\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_N]$ is a matrix with each column vector $\mathbf{f}_i$ obtained by concatenating all the pixels in the firing rate map $f_i(\mathbf{x})$. The exact method of the concatenation does not affect the final result, as long as it is used consistently. $\mathbf{I}_k$ is the k-th column of the identity matrix $\mathbf{I}$. The sole nonzero element of $\mathbf{I}_k$ represents the pixel of the animal's position after the concatenation. The requirement Eq. (11) for perfect reconstruction now becomes

$$\mathbf{I}_k = \mathbf{G}\mathbf{n} = \mathbf{G}\mathbf{F}^{\mathrm{T}} \mathbf{I}_k, \tag{13}$$

where in the last step Eq. (12) has been used, and matrix $\mathbf{G} = [\mathbf{g}_1, \mathbf{g}_2, \ldots, \mathbf{g}_N]$ is constructed in the same way as $\mathbf{F}$, and each column vector $\mathbf{g}_i$ corresponds to the basis function $g_i(\mathbf{x})$ to be determined. If Eq. (13) holds true for all k, that is, for all positions of the animal, then we must have $\mathbf{I} = \mathbf{G}\mathbf{F}^T\mathbf{I}$. Thus in general we should try to find $\mathbf{G}$ that minimizes

$$\left\| \mathbf{G}\mathbf{F}^{\mathrm{T}} - \mathbf{I} \right\|^2, \tag{14}$$

where $\|\mathbf{X}\|^2$ means the sum of the squares of all the elements of matrix $\mathbf{X}$. The solution to this least-mean-square problem is well-known:

$$\mathbf{G} = \mathbf{F}^{\dagger\mathrm{T}}, \tag{15}$$

where † denotes Moore-Penrose pseudoinverse and T denotes transpose. $\mathbf{F}$ and $\mathbf{G}$ are reciprocal in the sense that we also have

$$\mathbf{F} = \mathbf{G}^{\dagger \mathrm{T}}. \tag{16}$$

In other words, applying the pseudoinverse and transpose twice returns to the original basis. An example of reciprocal basis functions for place fields is shown in Figure 1.

The reciprocal nature of the pseudoinverse can be more clearly seen using singular value decomposition (Golub & Van Loan, 1996). Starting from the expansion $\mathbf{F} = \mathbf{UDV}^{\mathrm{T}}$, where $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices and $\mathbf{D}$ is a rectangular diagonal matrix, we immediately obtain $\mathbf{G} = \mathbf{UD}^{\dagger}\mathbf{V}^{\mathrm{T}}$, where $\mathbf{D}^{\dagger}$ is the diagonal matrix obtained from $\mathbf{D}^{\mathrm{T}}$ by replacing each nonzero element σ_i with its inverse $1/\sigma_i$, hence the reciprocal property. The familiar method of computing pseudoinverse from the normal equation

$$\mathbf{G}\mathbf{F}^{\mathrm{T}}\mathbf{F} = \mathbf{F} \tag{17}$$

by inverting the correlation matrix $\mathbf{F}^{\mathrm{T}}\mathbf{F}$ yields identical results, provided that the matrix is not singular or near singular.

It follows from Eq. (17) that the firing rate maps themselves become identical to the optimal template functions if they do not overlap with each other. That is, $\mathbf{F} = \mathbf{G}$ or $f_i(\mathbf{x}) = g_i(\mathbf{x})$ if

$$\mathbf{F}^{\mathrm{T}}\mathbf{F} = \mathbf{I}, \tag{18}$$

or equivalently, the tuning functions $f_1(\mathbf{x}), \ldots, f_N(\mathbf{x})$ are uncorrelated:

$$\sum_{\mathbf{x}} f_i(\mathbf{x}) f_j(\mathbf{x}) = C\delta_{ij}, \tag{19}$$

where C is a proportional constant and δ_{ij} is Kronecker delta which equals to 1 when $i = j$ and 0 otherwise. The conditions for orthogonality Eq. (18) or Eq. (19) are approximately satisfied for the recordings in Figure 1, which explains why the reciprocal fields look similar to the place fields (direct basis), and why the corresponding performances are also similar (Figures 3 and 5).

In the above analysis only the case where the animal is equally likely to visit any position $\mathbf{x}$ has been considered; that is, the probability distribution $P(\mathbf{x})$ is assumed to be uniform. In this case, the direct basis is the firing rate map $\phi_i(\mathbf{x}) = f_i(\mathbf{x})$, and the optimal template function that minimizes the mean square error is the reciprocal basis $\phi_i(\mathbf{x}) = g_i(\mathbf{x})$ given by Eq. (15). For nonuniform $P(\mathbf{x})$, instead of using the raw firing rate maps $f_i(\mathbf{x})$ as the basis functions, choose

$$\phi_i(\mathbf{x}) = f_i(\mathbf{x}) P(\mathbf{x}). \tag{20}$$

This weights the basis functions by the prior probability for the animal to visit each position. As expected, simulations confirm that this modification reduces reconstruction error. Similarly, we can use

$$\phi_i(\mathbf{x}) = g_i(\mathbf{x}) P(\mathbf{x}) \tag{21}$$

as the reciprocal basis functions. Other modifications of the reciprocal basis are considered in Appendix A and compared in Table 1. In our comparison of different methods in Figures 3 and 5, the simpler Eqs. (20) and (21) were used.

In summary, given the numbers of spikes $(n_1, n_2, \ldots, n_N)$ fired by all the cells during a given time interval, combine the chosen basis functions $\phi_i(\mathbf{x})$ linearly as $\sum_i n_i \phi_i(\mathbf{x})$, and then take the peak position $\hat{\mathbf{x}}$ as the reconstructed position of the rat as described by Eq. (9). The whole trajectory is constructed by sliding the time window forward (Figure 3).

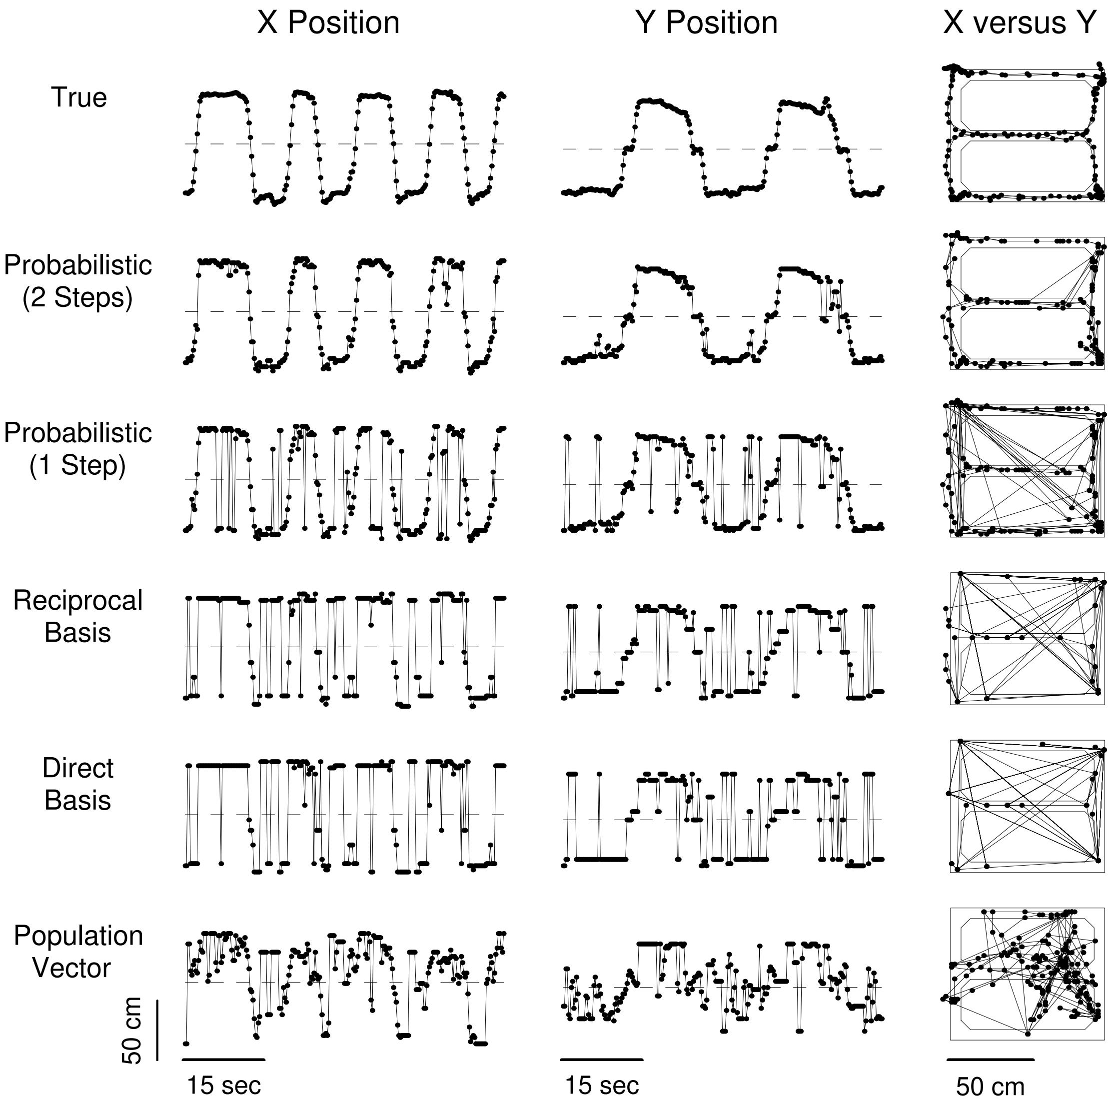

Figure 3: True X and Y positions of animal 1 running on the figure-8 maze as compared with the positions reconstructed by different methods using 25 place cells. The same 60-second segment is shown in all plots. The time window for reconstruction was 0.5 second, which was moved forward with a time step of 0.25 second. For a fair comparison of different methods, if none of the 25 cells fired within the time window, the reconstructed position at the preceding time step was used. The probabilistic or Bayesian methods were especially accurate, and the erratic jumps in the reconstructed trajectory were reduced by a continuity constraint using information from 2 consecutive time steps.

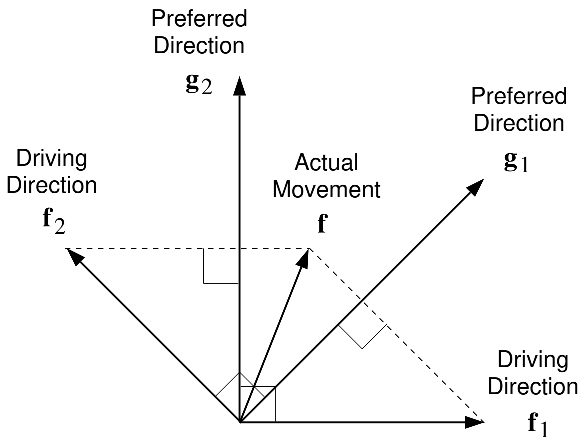

Figure 4: A hypothetical two-neuron motor system showing the distinction between the driving directions and the preferred directions, which form a pair of reciprocal vector bases. The apparent preferred directions differ from the true driving directions unless the system is orthogonal. Here reciprocity means that if vectors $\mathbf{g}_1$ and $\mathbf{g}_2$ are taken as the driving directions, then vectors $\mathbf{f}_1$ and $\mathbf{f}_2$ will become the apparent preferred directions.

3.2 A simple example of reciprocal basis

The reciprocal basis method is illustrated by a hypothetical example of coding in the motor system (Figure 4). Imagine two motoneurons each of which, when activated alone, will drive movement in the direction $\mathbf{f}_1$ or $\mathbf{f}_2$, respectively. We call $\mathbf{f}_i$ the driving direction of neuron i. Assume that the two neurons are activated randomly and independently and that the actual movement is the vector sum

$$\mathbf{f} = n_1\mathbf{f}_1 + n_2\mathbf{f}_2, \tag{22}$$

where n_1 and n_2 are the activity of the two neurons relative to their resting levels.

Suppose that both the activity of one neuron and the overall movement are being measured. How is the activity of this neuron correlated with the actual movement direction? Each neuron i will appear to have a cosine directional tuning with a fixed preferred direction $\mathbf{g}_i$, which, however, is different from the true driving direction $\mathbf{f}_i$. To see why this is the case, write the driving directions $\mathbf{f}_1$ and $\mathbf{f}_2$ as column vectors and define $\mathbf{F} = [\mathbf{f}_1, \mathbf{f}_2]$ as a 2×2 matrix. Equation (22) is equivalent to:

$$\mathbf{F}\begin{bmatrix} n_1 \\ n_2 \end{bmatrix} = \mathbf{f}, \tag{23}$$

from which activities n_1 and n_2 can be determined by inverting the matrix $\mathbf{F}$. Thus

$$\begin{bmatrix} n_1 \\ n_2 \end{bmatrix} = \mathbf{G}^{\mathrm{T}}\mathbf{f}, \tag{24}$$

where

$$\mathbf{G} = \mathbf{F}^{-\mathrm{T}} \tag{25}$$

is the transposed inverse. This is a special case of the reciprocal basis in Eq. (15) since the pseudoinverse is identical to the ordinary inverse for an invertible square matrix. Expressing $\mathbf{G} = [\mathbf{g}_1, \mathbf{g}_2]$ in terms of two column vectors $\mathbf{g}_1$ and $\mathbf{g}_2$, the solution Eq. (24) becomes

$$n_1 = \mathbf{g}_1 \cdot \mathbf{f} \qquad \text{and} \qquad n_2 = \mathbf{g}_2 \cdot \mathbf{f}. \tag{26}$$

Therefore, the activity n_i of each neuron is the inner product between the actual movement direction $\mathbf{f}$ and a fixed direction $\mathbf{g}_i$, which is the *preferred direction.* The driving directions and the preferred directions are reciprocal in the same sense as in the preceding section. They are identical only when the two driving directions are orthogonal. In general, Eq. (25) means $\mathbf{F}^{\mathrm{T}}\mathbf{G} = \mathbf{I}$ (unit matrix), or equivalently, $\mathbf{f}_1 \cdot \mathbf{g}_1 = \mathbf{f}_2 \cdot \mathbf{g}_2 = 1$ and $\mathbf{f}_1 \cdot \mathbf{g}_2 = \mathbf{f}_2 \cdot \mathbf{g}_1 = 0$. That is, in general, the preferred direction $\mathbf{g}_1$ of neuron 1 is always perpendicular to the driving direction $\mathbf{f}_2$ of neuron 2, and vice versa. The distinction between a pair of reciprocal vector bases is equivalent to that between contravariant and covariant tensors (Pellionisz, 1984; Pellionisz & Llinas, 1985).

The biological meaning of a reciprocal basis is not always as clear as in the simple example given above. The reciprocal field for each place cell depends not just on the place field of that cell, but also on all other cells in the population. If cells work together in groups or modules which are relatively independent of each other, then the reciprocal pairs should only include those cells in the same group. A different grouping can affect the shapes of the reciprocal fields. The reciprocal fields as shown in Figure 1 are meaningful only with respect to those 25 cells in question, and are therefore only a theoretical construct. Nonetheless, the concept of reciprocity they illustrate might have biological implications.

3.3 Special cases of basis functions

Template matching

The template matching method for reconstruction has been used for modeling stereo hyperacuity (Lehky & Sejnowski, 1990) and place cells (Wilson & McNaughton, 1993; Fenton & Muller, 1997). This method is equivalent to the direct basis method. At each fixed position $\mathbf{x}$, the average firing rates of all the cells can be described by a vector $\mathbf{f} = (f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_N(\mathbf{x}))$. A vector template is stored for each position $\mathbf{x}$ and is later used to match the actual number of spikes fired by these cells within a time window: $\mathbf{n} = (n_1, n_2, \ldots, n_N)$. During reconstruction, the matching method searches for the position in which the dot product $\mathbf{n} \cdot \mathbf{f}$ is maximized. This is equivalent to maximizing $\mathbf{n} \cdot \mathbf{f} = \sum_i n_i f_i(\mathbf{x})$ using $f_i(\mathbf{x})$ as basis functions. In other words, the reconstructed position is

$$\hat{\mathbf{x}}_{\text{template}} = \arg\max_{\mathbf{x}} \sum_i n_i f_i(\mathbf{x}). \tag{27}$$

Uniform scaling by the numbers of spikes does not affect the peak position or the reconstruction result. Position-dependent scaling such as $s(\mathbf{x}) = 1/\sum_i f_i(\mathbf{x})^2$ can be accommodated by redefining each template function $f_i(\mathbf{x})$ as $f_i(\mathbf{x})s(\mathbf{x})$.

Vector methods

Vector reconstruction is widely applied to data from neural populations. Vector methods such as the population vector and the optimal linear estimator compute

$$\hat{\mathbf{x}}_{\text{vector}} = \sum_i n_i \mathbf{x}_i, \tag{28}$$

where $\mathbf{x}_i$ is a fixed vector for each neuron i, and n_i is its activity level, such as the total number of spikes during a time interval or the firing rate.

This method is different from the basis function method since at each time step the basis function method yields *a scalar distribution function* of the animal's probable position over the entire maze, as shown in Figure 2. By contrast, in the same situation a vector method would yield *a single vector* which specifies a single point on the maze. On the other hand, after being discretized and concatenated, the relation between the direct and reciprocal basis functions is similar to that between the population vector and the optimal linear estimator (Georgopoulos, Schwartz & Kettner, 1986; Sanger, 1994; Salinas & Abbott, 1994).

The basis function framework includes the population vector method as a special case in the sense that the *performance* of population vector decoding is identical to combining cosine functions as the template functions. The *direction* of the population vector $\hat{\mathbf{x}}_{\text{vector}}$ in Eq. (28) is identical to the peak position of the following scalar function:

$$\psi(\mathbf{x}) \equiv \sum_i n_i \phi_i(\mathbf{x}), \tag{29}$$

where vector $\mathbf{x}$ is a free variable in the same space as $\hat{\mathbf{x}}_{\text{vector}}$ but has unit length, and the template function

$$\phi_i(\mathbf{x}) \equiv \mathbf{x}_i \cdot \mathbf{x} \tag{30}$$

is a cosine function of the angle between the two vectors. In other words, the direction of the reconstructed vector defined by Eq. (28) can also be computed by

$$\hat{\mathbf{x}}_{\text{vector}} / |\hat{\mathbf{x}}_{\text{vector}}| = \arg\max_{\mathbf{x}} \psi(\mathbf{x}) \tag{31}$$

where $|\hat{\mathbf{x}}_{\text{vector}}|$ is the length of the vector $\hat{\mathbf{x}}_{\text{vector}}$. To see this, take the dot product of the both sides of Eq. (28) with the unit vector $\mathbf{x}$, and then use the definitions in Eqs. (29) and (30) to obtain $\psi(\mathbf{x}) = \mathbf{x} \cdot \hat{\mathbf{x}}_{\text{vector}}$, which is maximized only when the unit vector $\mathbf{x}$ is aligned with $\hat{\mathbf{x}}_{\text{vector}}$. Therefore, basis functions with cosine tuning can precisely implement a vector method [see also Salinas & Abbott (1994)].

Population vector with scaling

For place cells, the population vector needs to be scaled by the total activity so that the reconstructed position

$$\hat{\mathbf{x}}_{\text{vector}} = \frac{\sum_i n_i \mathbf{x}_i}{\sum_i n_i}, \tag{32}$$

where $\mathbf{x}_i$ is the center of the place field of cell i. This method has been considered by several authors (Muller, Kubie & Ranck, 1987; Burgess, Recce & O'Keefe, 1994; Blum & Abbott, 1996; Abbott & Blum, 1996). It will be shown later that the scaling in Eq. (32) can be justified by the Bayesian method.

The two forms of population vectors in Eqs. (28) and (32) have different physical meanings. A vector representing direction will be called a *directional* vector, and a vector representing spatial position will be called a *positional* vector. If we scale a directional vector by an arbitrary scalar, the resultant vector still points in the same direction, but scaling a positional vector will lead to a different position. Vector reconstruction methods are most suitable for reconstructing a directional vector. When reconstructing position in two-dimensional space, a vector method can produce implausible results, even after scaling by total activity (Figure 3). We return to this point later.

To further illustrate the difference between the basis function method and a population vector, consider the special case in which every place field has the same Gaussian tuning function $f_i(\mathbf{x}) = G(\mathbf{x} - \mathbf{x}_i)$, where $\mathbf{x}_i$ is the center of the place field of cell i. To find the peak position of $\sum_i n_i f_i(x)$, set its derivative with respect to $\mathbf{x}$ to zero. This leads to

$$\hat{\mathbf{x}}_{\text{basis}} = \frac{\sum_i n_i \mathbf{x}_i G(\hat{\mathbf{x}}_{\text{basis}} - \mathbf{x}_i)}{\sum_i n_i G(\hat{\mathbf{x}}_{\text{basis}} - \mathbf{x}_i)}, \tag{33}$$

which can be compared with Eq. (32). Because each $G(\hat{\mathbf{x}}_{\text{basis}} - \mathbf{x}_i)$ is small unless the estimated position $\hat{\mathbf{x}}_{\text{basis}}$ is sufficiently close to the place field center $\mathbf{x}_i$, the overall effect may be interpreted as a population vector with an additional weighting scheme in favor of those place fields centered close to the estimated position. Consequently, an occasionally active cell with a place field far away from the estimated position does not affect the final result of the basis function method as much as in the case of a population vector.

4 Probabilistic Approach: Bayesian Methods

4.1 Background

The Bayesian approach is a natural one for reconstruction. It is optimal within a probabilistic framework and, as shown here, yields the best reconstruction results for our dataset. Once the position dependence of the firing rates of a group of place cells is known, the Bayes formula directly addresses the inverse problem: given the firing rates of these cells, how can we infer the most probable position? Bayesian reconstruction has been applied by Földiák (1993) to visual orientation tuning, by Sanger (1996) to motor directional tuning, and by Ginzburg (Gerrard, Ginzburg, Barnes, Sejnowski & McNaughton, 1995) and Brown (Brown, Frank & Wilson, 1996) to place cells.

Two optional supporting models can be used with Bayesian reconstruction. The first is the tuning model, which is especially useful when data sampling is sparse. Brown, Frank & Wilson (1996) used Gaussian tuning model in their reconstruction, and we used a Gaussian tuning model for estimating the information-theoretical limit on reconstruction accuracy. In our reconstruction, the tuning functions were determined empirically without assuming an explicit analytical model. A second option for Bayesian reconstruction is the spike generation model. Because of the low firing rates of place cells, we used the Poisson firing model in our reconstruction. Bayesian reconstruction using Poisson spike statistics was first used by Sanger (1996). Alternatively, without using an explicit spike generation model, the stimulus-response relation can be obtained empirically as a look-up table (Földiák, 1993).

4.2 Basic method

Assume that using the methods in the preceding section we have measured the firing rate maps $f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_N(\mathbf{x})$ of a population of N place cells as well as the animal's position distribution $P(\mathbf{x})$ over a period of time. At an arbitrary later time t, given the numbers of spikes fired by these place cells within the time interval from $t - \tau/2$ to $t + \tau/2$, where τ is the length of time window, the goal is to compute the probability distribution of the animal's position at time t. Notice that what is to be computed here is a distribution of positions instead of a single position; we can always take the most probable position, which corresponds to the peak of the probability distribution, as the most likely reconstructed position of the animal (Figure 2).

Let the vector $\mathbf{x} = (x, y)$ be the position of the animal, and the vector $\mathbf{n} = (n_1, n_2, \ldots, n_N)$ be the numbers of spikes fired by our recorded cells within the time window, where n_i is the number of spikes of cell i. The reconstruction is based on the standard formula of conditional probability:

$$P(\mathbf{x}|\mathbf{n})P(\mathbf{n}) = P(\mathbf{n}|\mathbf{x})P(\mathbf{x}). \tag{34}$$

The goal is to compute $P(\mathbf{x}|\mathbf{n})$, the probability for the animal to be at position $\mathbf{x}$ given the numbers of spikes $\mathbf{n}$. $P(\mathbf{x})$ is the probability for the animal to be at position $\mathbf{x}$, which can be measured experimentally. The probability $P(\mathbf{n})$ for the numbers of spikes $\mathbf{n}$ to occur can be determined by normalizing the conditional probability $P(\mathbf{x}|\mathbf{n})$ over $\mathbf{x}$ and therefore does not have to be estimated directly.

The key step is to evaluate the conditional probability $P(\mathbf{n}|\mathbf{x})$, which is the probability for the numbers of spikes $\mathbf{n}$ to occur given that we know the animal is at location $\mathbf{x}$. It is intuitively clear that this probability is determined by the firing rate maps of the place cells. More precisely, if we assume that the spikes have Poisson distributions and that different place cells are statistically independent of one another, then we can obtain the explicit expression

$$P(\mathbf{n}|\mathbf{x}) = \prod_{i=1}^{N} P(n_i|\mathbf{x}) = \prod_{i=1}^{N} \frac{(\tau f_i(\mathbf{x}))^{n_i}}{n_i!} \exp(-\tau f_i(\mathbf{x})), \tag{35}$$

where $f_i(\mathbf{x})$ is the average firing rate of cell i at position $\mathbf{x}$, and τ is the length of the time window.

The final formula obtained from Eq. (34) reads

$$P(\mathbf{x}|\mathbf{n}) = C(\tau, \mathbf{n}) P(\mathbf{x}) \left(\prod_{i=1}^{N} f_i(\mathbf{x})^{n_i} \right) \exp\left(-\tau \sum_{i=1}^{N} f_i(\mathbf{x}) \right), \tag{36}$$

where $C(\tau, \mathbf{n})$ is a normalization factor, which can be determined by the normalization condition $\sum_{\mathbf{x}} P(\mathbf{x}|\mathbf{n}) = 1$ (Sanger, 1996). Optimality of Bayesian reconstruction is considered in Appendix C.

The Bayesian reconstruction method uses Eq. (36) to compute the probability $P(\mathbf{x}|\mathbf{n})$ for the animal to be at each position $\mathbf{x}$, given the numbers of spikes $\mathbf{n}$ of all the cells within the time window. In this probability distribution, the peak position or the most probable position is taken as the reconstructed position of the animal. In other words,

$$\hat{\mathbf{x}}_{\text{Bayes}} = \arg\max_{\mathbf{x}} P(\mathbf{x}|\mathbf{n}). \tag{37}$$

By sliding the time window forward, the entire trajectory of the animal's movement can be reconstructed from the time varying activity in the neural population.

4.3 Testing the assumptions

The Bayesian reconstruction as applied in Eq. (36) relies on two assumptions, namely, Poisson spike distribution and independent firing of different cells. These assumptions are reasonable and convenient approximations, although neither is precisely true.

Poisson spike statistics

The Poisson firing model is commonly used for neuronal spike statistics (Tuckwell, 1988). It is equivalent to assuming that the firing is as random as possible. If points are distributed randomly over a long time interval, with equal chance for each point to fall at any location, independently of each other, then the number of points contained in a small time window has a Poisson distribution. A direct test for a Poisson distribution is not easy for place cells because their mean firing rates are low and depend on the location, which an animal may visit only transiently. Nonetheless, the interspike intervals can be directly compared with the exponential distribution expected for a Poisson process. Consider cell i with firing rate map $f_i(\mathbf{x})$. According to the inhomogeneous Poisson model with mean firing rate depending on the position $\mathbf{x}$, the overall probability density for interspike interval s should be

$$p_i(s) = \sum_{\mathbf{x}} P(\mathbf{x}) f_i(\mathbf{x}) \exp(-s f_i(\mathbf{x})), \tag{38}$$

where $P(\mathbf{x})$ is the probability for the animal to visit position $\mathbf{x}$. The above formula implies that $dp_i(s)/ds < 0$ always holds true; that is, the interspike interval histogram must be monotonically decreasing.

Applying Eq. (38) to real place cell data, we found that it can approximately account for the distribution of intervals longer than 10 msec. But there are several problems. First, the actual occurrence of briefer intervals was much more frequent than expected. This means that the Poisson model fails to capture the tendency for hippocampal pyramidal cells to fire in bursts, called complex spikes (Ranck, 1973). Second, instead of decreasing monotonically, the actual interspike interval histograms tended to have a small bump within the theta frequency range (4–12 Hz), reflecting the periodic drive from the septal pacemakers (Stewart & Fox, 1990). Third, the Poisson spiking model does not have a refractory period to rule out extremely short intervals. A related finding by Fenton & Muller (1997) is that place cell firing for similar movement trajectories in an open environment was more variable that expected from a Poisson distribution. Despite these limitations, the inhomogeneous Poisson model proved to be adequate for the purpose of reconstruction.

Independence of different cells

The activities of two cells are statistically independent only if the joint probability of firing is a product of the individual probabilities. For two place cells 1 and 2, their independence means that at each position $\mathbf{x}$

$$P(n_1, n_2|\mathbf{x}) = P(n_1|\mathbf{x})P(n_2|\mathbf{x}), \tag{39}$$

where n_1 and n_2 are the number of spikes collected from the two cells within an arbitrary time window. For the data in Figure 1, the majority of the pairs of place cells were approximately independent simply because most place fields did not overlap. In other words, most cell pairs seldom fired together. Thus Eq. (39) may hold for the trivial reason that both sides are zero. A related consequence is that multiplying place fields in Bayesian reconstruction typically leads an extremely sharp distribution, with zero value for most pixels (Figures 2 and 13). The independence is no longer exactly true for cells with highly overlapping place fields [cf. Eichenbaum, Wiener, Shapiro & Cohen (1989)]. In fact, the correlations in the firing of these cells could be strengthened rapidly by Hebbian-type learning process (Wilson & McNaughton, 1994; Skaggs & McNaughton, 1996; Blum & Abbott, 1996; Mehta, Barnes & McNaughton, 1997). In such a situation, the independence assumption is only a simplifying approximation for theoretical convenience. There is additional information in correlated firing that could be exploited to improve the accuracy of reconstruction.

4.4 Continuity constraint: using two time steps

The trajectories reconstructed by all methods suffer from occasional erratic jumps, which are often caused by low instantaneous firing rates of the recorded place cells, especially when the animal stops running (Figure 3). Clearly, if all recorded cells stop firing, there is not enough information for accurate reconstruction.

Introducing a continuity constraint can improve reconstruction accuracy by reducing erratic jumps in the reconstructed trajectory. The possibility that some of the jumps may reflect real biological processes will be discussed later. To implement the continuity constraint, the reconstructed position from the preceding time step is used as a guide for estimating the current position. Speed information can be used to limit the change in position allowed within a single time step.

The continuity constraint can be incorporated easily into both the Bayesian and the basis function frameworks. Here we demonstrate the constraint only for the Bayesian method. The one-step Bayesian method presented previously computes the conditional probability $P(\mathbf{x}_t|\mathbf{n}_t)$ for the animal's position $\mathbf{x}_t$ at time step t, given the numbers of spikes $\mathbf{n}_t$ within a time window centered also at time t. Here subscripts denote the time step. When the reconstructed position $\mathbf{x}_{t-1}$ at the preceding time step is used as an additional piece of information, we can compute the conditional probability $P(\mathbf{x}_t|\mathbf{n}_t, \mathbf{x}_{t-1})$ from the general Bayes formula $P(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{n}_t)P(\mathbf{x}_t, \mathbf{n}_t) = P(\mathbf{x}_t|\mathbf{n}_t, \mathbf{x}_{t-1})P(\mathbf{n}_t, \mathbf{x}_{t-1})$ and the assumption that

$$P(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{n}_t) = P(\mathbf{x}_{t-1}|\mathbf{x}_t). \tag{40}$$

In the sense of probability, this assumption means that the activity $\mathbf{n}_t$ at the current time step cannot directly affect the position $\mathbf{x}_{t-1}$ at the preceding time step. With the help of simple equalities like $P(\mathbf{x}_t, \mathbf{n}_t) = P(\mathbf{x}_t|\mathbf{n}_t)P(\mathbf{n}_t)$ and $P(\mathbf{n}_t, \mathbf{x}_{t-1}) = P(\mathbf{x}_{t-1}|\mathbf{n}_t)P(\mathbf{n}_t)$, we immediately obtain the final equation:

$$P(\mathbf{x}_t|\mathbf{n}_t, \mathbf{x}_{t-1}) = kP(\mathbf{x}_t|\mathbf{n}_t)P(\mathbf{x}_{t-1}|\mathbf{x}_t), \tag{41}$$

where $k = 1/P(\mathbf{x}_{t-1}|\mathbf{n}_t)$ is a scaling factor that does not depend on $\mathbf{x}_t$ and can be determined by normalizing $P(\mathbf{x}_t|\mathbf{n}_t, \mathbf{x}_{t-1})$ over $\mathbf{x}_t$.

Now the reconstructed position at time t is

$$\hat{\mathbf{x}}_t = \arg\max_{\mathbf{x}_t} P(\mathbf{x}_t|\mathbf{n}_t, \mathbf{x}_{t-1}). \tag{42}$$

Compared with the single step Eq. (36), the only new factor in Eq. (41) is the conditional probability $P(\mathbf{x}_{t-1}|\mathbf{x}_t)$. Intuitively, the continuity constraint should require that given the current location $\mathbf{x}_t$, the preceding location $\mathbf{x}_{t-1}$ cannot be too far away. For a precise formulation, we assume a two-dimensional Gaussian distribution

$$P(\mathbf{x}_{t-1}|\mathbf{x}_t) = C \exp\left(\frac{-\|\mathbf{x}_{t-1} - \mathbf{x}_t\|^2}{2\sigma_t^2}\right), \tag{43}$$

where the standard deviation σ_t is related to the speed v_t at time t by the formula

$$\sigma_t = K\left(\frac{v_t}{V}\right)^d, \tag{44}$$

where both K and V are constant, and V is a speed so that the ratio in the parentheses is dimensionless. We expect the exponent $d = 1/2$ for random walks, and $d = 1$ for linear movements. The speed v_t can be estimated by using the approximate linear relationship between the instantaneous mean firing rate and running speed (Figure 8) or by reading the stored average speed at each position as computed from previous data. We used the latter method and set $v_t = U(\mathbf{x}_t)$, where function $U(\mathbf{x})$ stores the average speed at each position $\mathbf{x}$.

In general, if the standard deviation σ_t of the Gaussian prior is too large, the continuity constraint has little effect. On the other hand, if σ_t is too small, the constraint may become too restrictive and the reconstructed position might get stuck in the same position if the real position has moved away by a distance much larger than σ_t. In our analysis, we simply used $d = 1$ for linear movements, and confined σ_t to between 20 cm and 60 cm (good empirical values). The scaling in Eq. (44) was chosen so that σ_t would just reach the allowed maximum of 60 cm at top running speed. For generality, we did not include any bias for the direction of movement. As shown in Figures 3 and 5, the continuity constraint is effective in reducing reconstruction errors caused by erratic jumps.

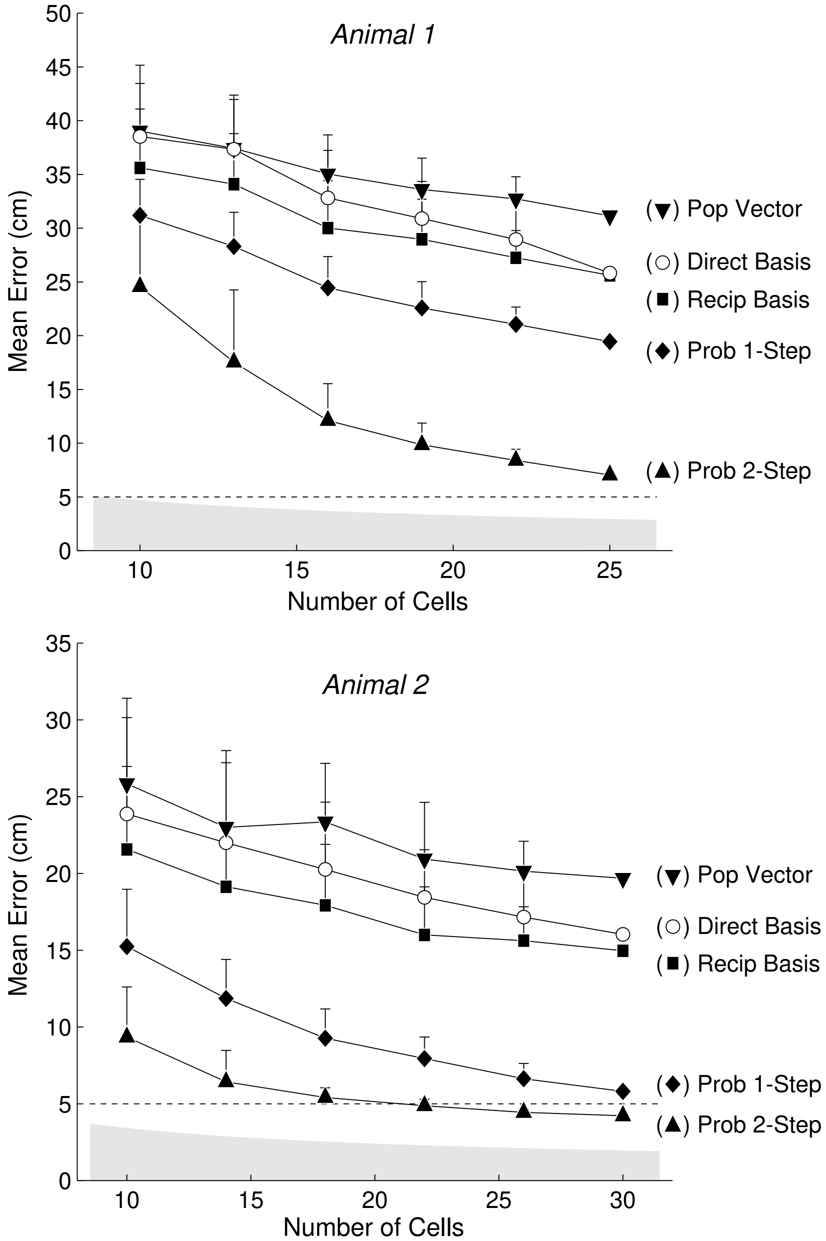

Figure 5: Mean errors of different reconstruction methods applied to data from two different rats. The errors all decreased as more cells were used, and the lowest errors were comparable with the intrinsic error of the tracking system, which was about 5 cm (horizontal lines). The shaded region represents reconstruction errors prohibited by the lower bound derived by Fisher information and the Cramér-Rao inequality [Eq. (53)]. Each data point and error bar represent the average and the standard deviation of the mean errors in 40 repetitive trials in which a subset of cells were drawn randomly from the whole sample. Animal 1 ran on the figure-8 maze, and 7 minutes of data were used for sampling, and the subsequent 7 minutes of data were used for reconstruction. Animal 2 run on the rectangular maze (upper loop of the figure-8 maze), and both the sampling and the reconstructing time intervals were 10 minutes. The time window was 1 second.

4.5 Incorporating firing rate modulation

Under the assumption of multiplicative modulation of firing rate such as in Eq. (6), the Bayesian Eq. (36) should be modified by replacing the fixed tuning function $f_i(\mathbf{x})$ of each cell i by

$$f_i(\mathbf{x}) \to f_i(\mathbf{x})m(t), \tag{45}$$

where, for simplicity, the modulation factor $m(t)$ is supposed to be the same for all cells. Here $m(t)$ is written as a function of the time step. This general formulation includes more specific variables such as modulation by speed as a special case. Note that the modulation factor can be estimated directly from spike data using the approximation

$$m(t) \approx F(t)/F_{\text{mean}} \,, \tag{46}$$

where $F(t)$ is the instantaneous firing rate averaged over all the cells and F_{mean} is the average of $F(t)$ over time. The modified Bayesian formula reads

$$P(\mathbf{x}|\mathbf{n}) = C'P(\mathbf{x}) \left(\prod_{i=1}^{N} f_i(\mathbf{x})^{n_i} \right) \exp\left(-m(t)\tau \sum_{i=1}^{N} f_i(\mathbf{x}) \right), \tag{47}$$

where factors independent of the position $\mathbf{x}$ have been absorbed in the constant C'.

4.6 Relationship to basis functions

General equivalence

In general, the performance of all basis function methods can be mimicked by the Bayesian method if appropriate template functions are used. To implement the sum $\sum_i n_i \phi(\mathbf{x})$ with the basis functions $\phi_i(\mathbf{x})$, choose the tuning function

$$f_i(\mathbf{x}) = \exp(\phi_i(\mathbf{x})) \tag{48}$$

and the prior probability for spatial occupancy

$$P(\mathbf{x}) = \exp\left(\tau \sum_i f_i(\mathbf{x}) \right), \tag{49}$$

where τ is time window. Then the Bayesian Eq. (36) becomes

$$P(\mathbf{x}|\mathbf{n}) = C \exp\left(\sum_i n_i \phi_i(\mathbf{x}) \right). \tag{50}$$

Its peak position is identical to that of $\sum_i n_i \phi(\mathbf{x})$ because the exponential function is monotonic.

Conversely, the Bayesian method can also be implemented in a basis function framework, provided that a constant bias is allowed. This relationship will be considered in detail in our later discussion on network implementations.

Relationship with population vector with scaling

Under special assumptions, the Bayesian method can lead to the population vector method with scaling by total activity. This scaling is needed for reconstruction of a positional vector. Suppose the tuning function for each cell i is a Gaussian $f_i(\mathbf{x}) = A_i \exp(-\|\mathbf{x} - \mathbf{x}_i\|^2/2\sigma_i^2)$, where $\mathbf{x}_i$ is the

Table 1: Relative mean errors of different reconstruction methods

Method	Animal 1	Animal 2	Key Equations
Population vector	10.61	9.41	(32)
Direct basis	8.80	7.66	(9) and (20)
Reciprocal basis	8.72	7.15	(9) and (21)
Modified basis (noiseless)	8.83	6.10	(80) and (83)
Modified basis (noise included)	8.76	4.84	(80) and (84)
Probabilistic 1-step	6.62	2.78	(36) and (37)
Probabilistic 1-step (modulated)	6.37	2.62	(36), (37) and (47)
Probabilistic 2-step	2.39	2.02	(41) and (42)
Probabilistic 2-step (modulated)	3.08	1.97	(41), (42), and (47)
Theoretical minimal error	1	1	(53)

Notes: The actual values of the theoretical minimal errors for animals 1 and 2 were 2.94 cm and 2.09 cm, respectively. The movement trajectory from the second half session was reconstructed using data from the first half session (7 minutes for animal 1 and 10 minutes for animal 2). The time window was 1 second, and the spatial grid was 64×64.

center of the place field of cell i and σ_i characterizes its width. Suppose the centers $\mathbf{x}_i$ and the peak firing rates A_i are uniformly distributed in space, then the sum $\sum_i f_i(\mathbf{x})$ is approximately constant in space. In this case, the optimal position $\hat{\mathbf{x}}$ in Bayesian reconstruction is given by

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}} \prod_i f_i(\mathbf{x})^{n_i} = \frac{\sum_i \left(n_i/\sigma_i^2\right) \mathbf{x}_i}{\sum_i n_i/\sigma_i^2}, \tag{51}$$

where the second equality can be obtained by setting the derivative of the product with respect to $\mathbf{x}$ to zero. This is a generalized population vector, which reduces to the form in Eq. (32) when all place fields have identical width; that is, when all $\sigma_i = \sigma$. This justifies scaling the positional population vector by the total activity $\sum_i n_i$ in Eq. (32).

5 Results of Place Cell Reconstruction and Discussion

5.1 Comparison of different methods

Reconstruction accuracy

The performance of different methods can be compared qualitatively in Figure 3 and quantitatively in Table 1 and Figure 5. The effects of various parameters are discussed later.

The reconstruction error of the direct basis method for a given number of cells is similar to that reported by Wilson & McNaughton (1993). This result makes sense because of the essential equivalence between the matching method and the direct basis method as discussed before.

The linear method with a reciprocal basis only improved slightly the performance of the direct basis because the place cells were mostly orthogonal and the two sets of basis functions were quite similar to each other (Figure 1). Further modifications of the reciprocal basis sometimes brought additional improvement (Table 1).

Probabilistic methods based on the Bayesian approach were more accurate. The best method overall was the Bayesian method using two time steps to enforce a continuity constraint on the reconstructed position. Its errors were in the range of the intrinsic error of the system for tracking the animal's position, which was estimated to be around 5 cm (Wilson & McNaughton, 1993; Skaggs, McNaughton, Wilson & Barnes, 1996).

Because reconstruction errors tend to be much larger during immobility (see below), excluding slow running periods from the analysis can improve reconstruction accuracy (Gerrard, Ginzburg, Barnes, Sejnowski & McNaughton, 1995). However data exclusion requires additional criteria. For all methods in this paper, we reconstructed the whole trajectory without introducing additional parameters.

A problem with vector methods

The population vector method was the least accurate in reconstructing position in space for our dataset. The major problem is that it sometimes yielded implausible results. To see why, imagine that only two cells were active within the time window; then the weighted average of their positions would lie somewhere on the line connecting the two centers of the place fields. The estimated position may be in the center of the maze, which was never visited by the rat. This is a general problem for all vector methods including the reciprocal vector or optimal linear estimator. As long as the shape of the environment is not convex, a vector method may suffer from this problem. The existence of systematic bias in vector methods has been pointed out by Muller, Kubie & Ranck (1987). In contrast, both the basis function method and the probabilistic method take into account the entire two-dimensional distribution functions, and the final peak always occurs in a plausible position. See also the discussion following Eq. (33).

Remark on silent time steps

For a fair comparison of different methods, in all of our simulations we have used the heuristic that if all the recorded cells were silent within a time window, the current reconstructed position was set to be the reconstructed position at the preceding time step. We needed this rule because when all recorded cells are silent within the time window, all basis function methods fail, whereas the Bayesian methods are robust. With a time window of 0.5 second, about 15% of the time steps contained no spike in animal 1, and the percentage was only around 0.2% in animal 2. For a time window of 1 second, the percentages dropped to 4.3% and 0% for animals 1 and 2, respectively. There were fewer silent time steps in animal 2 because its maze was smaller but there were more cells so that the overall average firing rates were higher than that of animal 1.

5.2 Effects of parameters

Number of cells

When more cells were used in reconstruction, the errors tended to decrease (Figure 5). Different methods maintained their relative accuracy for different numbers of cells. In theory, the error of various reconstruction methods should be inversely proportional to the square root of the number of cells [Eq. (56)]. However, the inverse square root law is valid only when a sufficiently large number of cells is used, as shown by simulation (Figure 12). Here the number of cells was too small to apply the law.[1]

Time window

A time window with a duration of around 1 second was a good choice for the Bayesian 2-step method for these data (Figure 6A). The results of other methods were similar (not shown), although the errors for large time windows sometimes leveled off or only increased slightly. In theory, the error is expected to be inversely proportional to the square root of the time window [Eq. (56)]. This can account for the general tendency for the reconstruction to improve with a longer time window. However this theoretical consideration no longer holds when the time window exceeds the optimal size.

Time shift

In the Bayesian 2-step method, the center of the time window was shifted relative to the current time to see what shift produced the best accuracy [cf. Muller & Kubie (1989); Blair & Sharp (1995); Taube & Muller (1995); Blair, Lipscomb & Sharp (1997)]. Alignment of the time window to within 100 msec of the current time gave the best results. Figure 6B compares the reconstruction errors from the first half of the session to that from the second half. As expected, the errors for reconstructing the second half were larger than that of self-reconstruction, since all parameters for the reconstruction algorithm were determined using data from only the first half session.

Upon closer examination of Figure 6B, it would appear that the minimum of the curve for reconstructing the second half session (dark symbols) drifted towards the *past* with respect to the minimum of the self-reconstruction curve (light symbols). To quantify the drift, each curve was fitted with a 5th-order polynomial and the minimum of the smooth curve was found. The drift was computed as the difference between the two minima. For the data shown in Figure 6B, the value of the drift was -66 msec [$= -34$ msec $- 32$ msec] for animal 1, and -68 msec [$= -93$ msec $-$ $(-25$ msec$)$] for animal 2. Conversely, when the second half of data was used to reconstruct the trajectory of the first half, the minimum appeared to drift towards the *future* with respect to the minimum of self-reconstruction (not shown). Using the same smoothing method, the value of the drift was 118 msec [$= 107$ msec $- (-11$ msec$)$] for animal 1, and 93 msec [$= 17$ msec $- (-76$ msec$)$] for animal 2.

The accuracy for the estimation of the drift was limited by noise and the flatness of the curves. Nevertheless, the direction and the magnitude of the drift were consistent with the phenomenon of the back shift of place field centers during repetitive running as reported by Mehta, Barnes & McNaughton (1997), who directly computed the center of mass of place fields and found an average back shift of a few centimeters during a running period of the order of several minutes. On the other hand, we compared the spatial distribution of the reconstruction errors as a vector field for different time blocks (see below), but could not see an obvious systematic shift of the error direction in correlation with the running direction.

Time blocks: slow trends

[1] For comparison, we tested the square root law on the data taken from Figure 3 in the paper by Georgopoulos, Kettner & Schwartz (1988), where the error of population vector was computed using up to 475 motor cortical cells. The errors were given in terms of 95% confidence cone, which was expected to be proportional to the mean angular error. The square root law held quite well: on log-log scale, a linear fit yielded a slope of -0.519 with a correlation coefficient of 0.9984, as compared with the theoretical slope of $-1/2$.

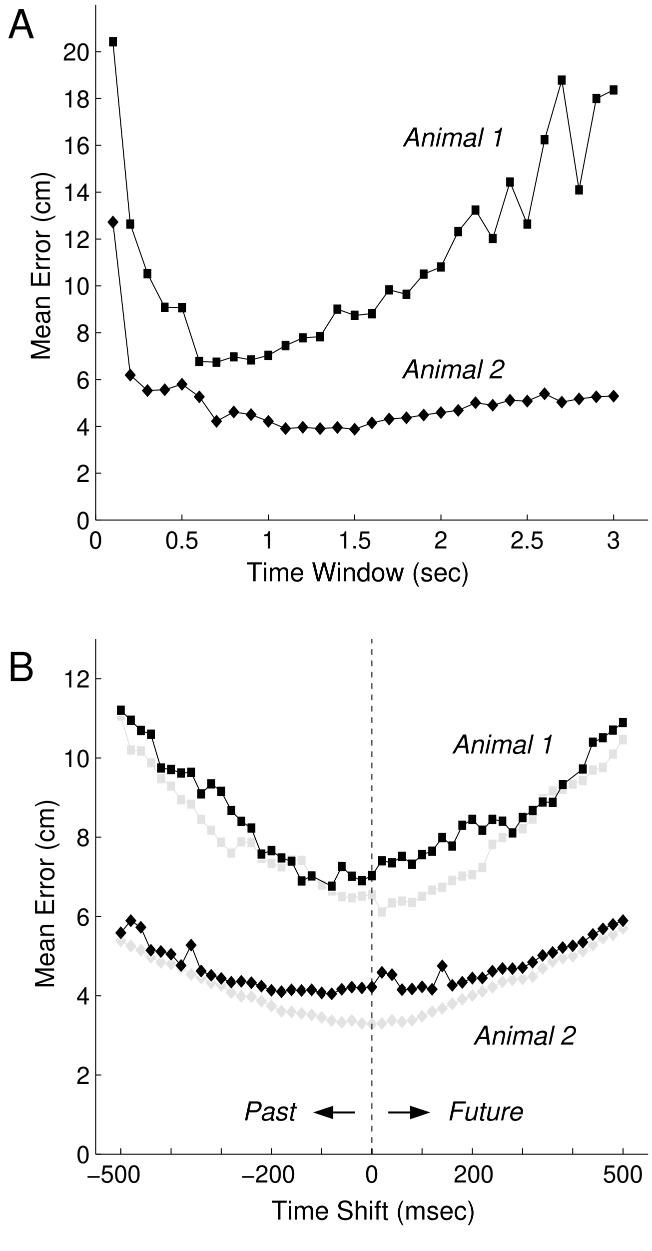

Figure 6: The mean error of the Bayesian 2-step reconstruction method depends on the time window and the shift of its center. (A) The length of the time window was altered systematically whereas all other parameters were identical to those in Figure 5, using all the cells. (B) The center of the time window (length fixed at 1 second) was shifted relative to the current instant of time. Only data from the first half session were used for sampling. Reconstruction was performed on data from both the first half session (light gray curves and symbols) and the second half session (black curves and symbols, with two outlier data points removed). For animals 1, each half session lasted 7 minutes, and for animal 2, each lasted 10 minutes. For both animals, the reconstruction for the second half session was most accurate when the center of the time window was slightly shifted towards the past.

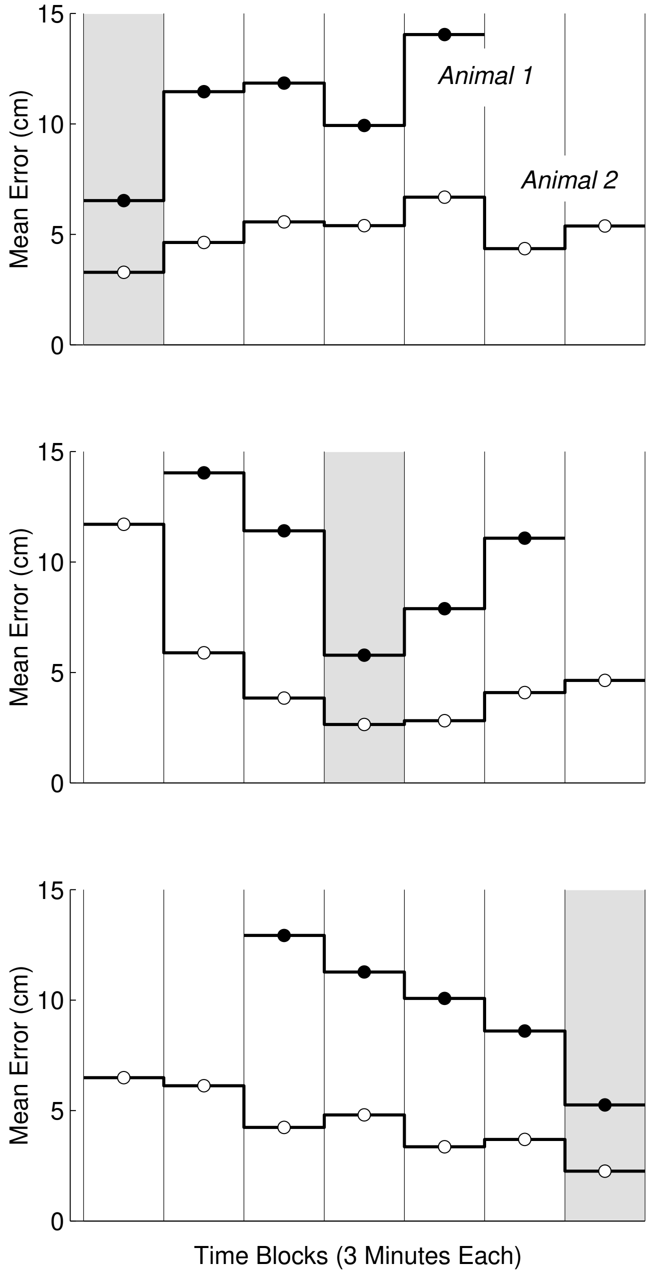

Figure 7: Mean reconstruction errors of the Bayesian 2-step method in consecutive and non-overlapping time blocks, each lasting 3 minutes. Only a single block (shaded) was used as the data block for sampling, while reconstruction was performed on all blocks. The graphs for the two animals are aligned along the data blocks. The average reconstruction errors tended to change gradually even though there was no overlap between contiguous time blocks.

As shown in Figure 7, when a single 3-minute time block of data was used to reconstruct the movement trajectory of other time blocks, the average reconstruction error tended to drift systematically. That is, reconstruction tended to get worse when going further into either the past or the future. Because there was no overlap between the time blocks, the tendency for gradual drift of error must be caused by gradual changes of either behavioral or neural origin.

During repetitive running, the mean running speed of both animals slowly decreased. For animal 2 (on the rectangular maze), the mean speed dropped from 9.0 cm/sec in the first 3-minute time block to 5.9 cm/sec in the last time block, with an overall average of 8.2 cm/sec across all seven time blocks. Similarly, the speed for animal 1 (on the figure-8 maze) dropped from 21.8 cm/sec to 16.6 cm/sec, for an average of 18.9 cm/sec. On the other hand, the mean firing rate increased almost monotonically for animal 2 (from 0.86 Hz to 1.34 Hz, for an average of 1.23 Hz) but showed less clear trend for animal 1 (from 0.94 Hz to 0.87 Hz, for an average of 0.92 Hz). Therefore, slow changes both in behaviors and in neural properties could contribute to the gradual error drift in Figure 7.

Firing rate modulation

Including firing rate modulation in reconstruction had only a small effect on the accuracy, even though the average firing rate was clearly modulated by running speed (Figure 8), as reported by McNaughton, Barnes & O'Keefe (1983).

Including firing rate modulation has no effect on the performance of basis function methods and the population vector, provided that an identical modulation factor is used for all cells as in Eq. (45). But under the same conditions, the performance of the modulated Bayesian method [Eq. (46) and Eq. (47)] may slightly improve (Table 1).

Other parameters

The reconstruction results were not sensitive to parameters like the exact width of the spatial grid and the exact radius of the blurring kernel, as long as they were kept below the intrinsic error of the tracking system. In all reconstruction data presented in this paper, we used 64×64 grid, corresponding to $111\,\text{cm} \times 111\,\text{cm}$ in real space. Adding a small constant background value to the firing rate maps sometimes slightly improved the Bayesian reconstruction.

Finally, in the Bayesian 2-step method, the standard deviation σ_t of allowable jumps in Eq. (44) varied between 20 cm and 60 cm, which gave good performance for the data of animal 1. Smaller values slightly improved the reconstruction for animal 2, which did not have many erratic jumps in the first place. Changing these values affects the reconstruction error only slightly, as long as they are not too small or too large.

5.3 Discontinuity in reconstructed trajectory

As shown in Figures 3 and 9, reconstruction errors were relatively small most of the time but were punctuated by large jumps. Most of these jumps occurred when the animal stopped running (Figure 9). This can also be seen in Figure 3, where the jumps tend to occur between the corners, the places where the animal stopped for food reward. The results from animal 2 were qualitatively similar, but the jumps in reconstruction errors were less frequent because more cells were used in reconstruction and their average firing rates were higher.

The occurrence of these jumps is correlated with the several confounding factors, all of which are apparent in Figure 9:

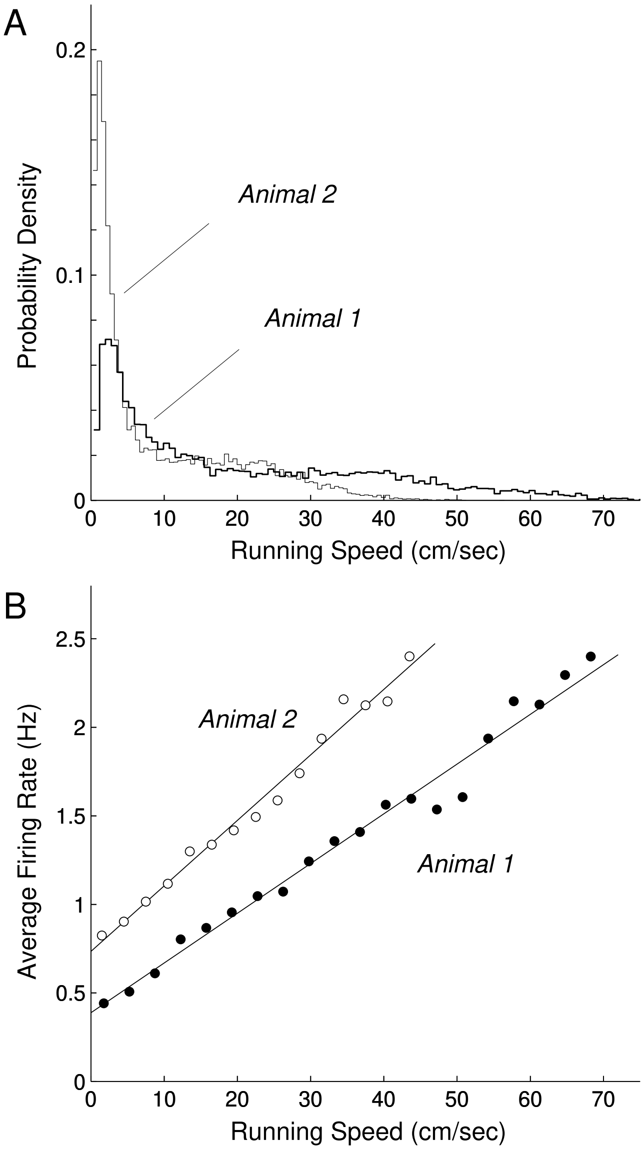

Figure 8: Modulation of place cell firing rate by running speed was nearly linear. (A) Histogram of speed distribution. The speed was computed after convolving the position trajectory with a Gaussian kernel. (B) The average firing rate across all cells as a function of the binned running speed. For animal 1, the mean speed was 19.6 cm/sec, the mean firing rate was 0.92 Hz, and the linear fit was $0.028v + 0.39$ with correlation coefficient $r = 0.991$. For animal 2, the corresponding parameters were 9.3 cm/sec, 1.1 Hz, and $0.037v + 0.73$ with $r = 0.992$. Only the first half of data are shown (7 minutes for animal 1 and 10 minutes for animal 2). (For the second half of data, the parameters for animal 1 were 18.0 cm/sec, 0.91 Hz, and $0.033v + 0.31$ with $r = 0.991$; for animal 2, they were 7.0 cm/sec, 1.38 Hz, and $0.023v + 1.25$ with $r = 0.956$.)

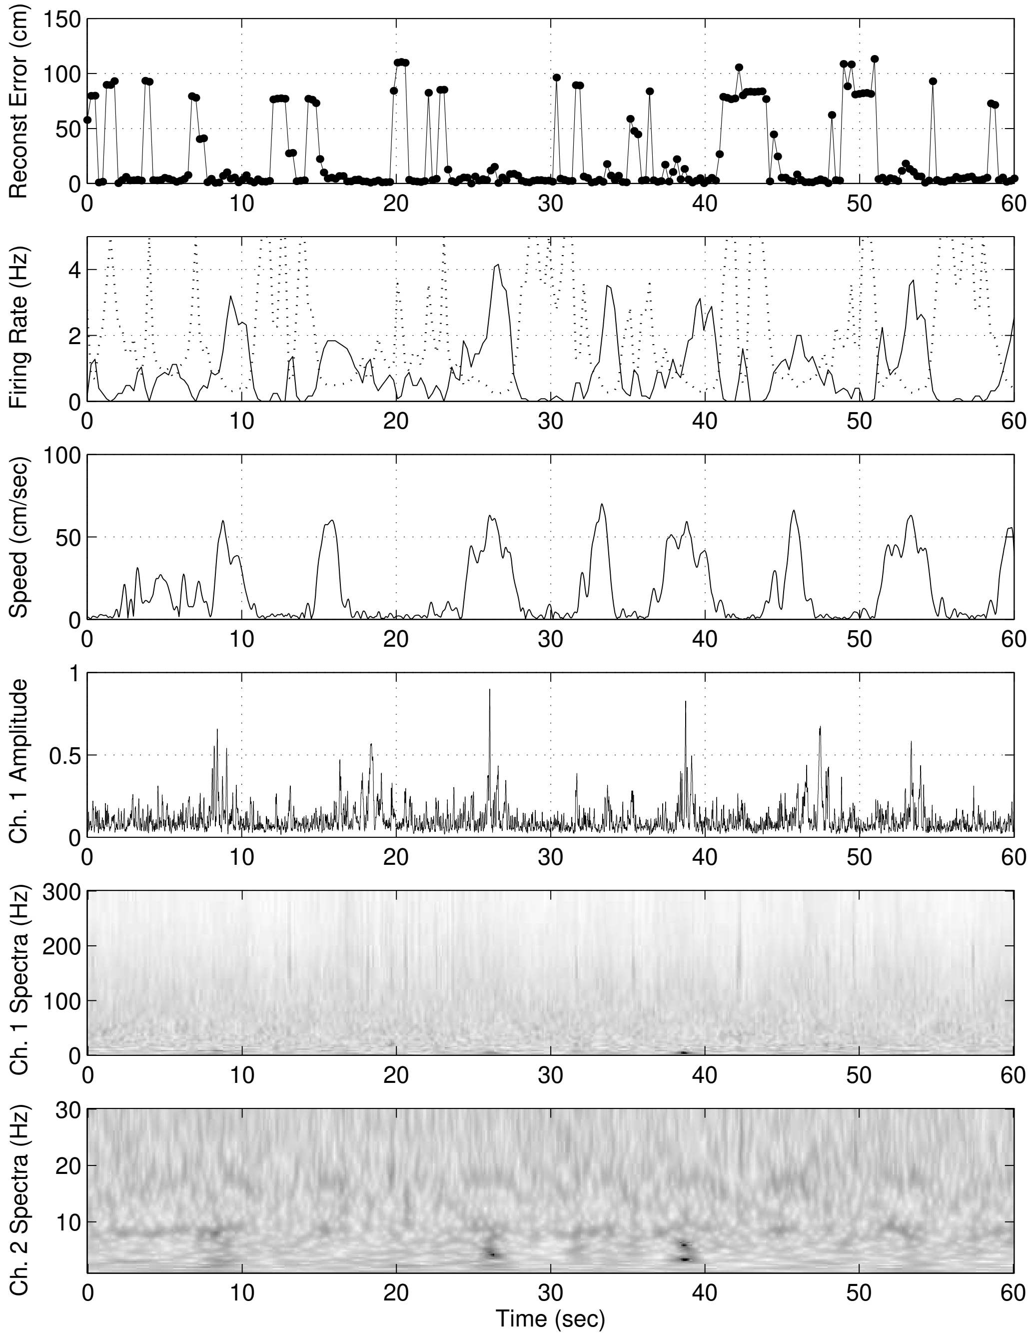

Figure 9: Sudden jumps of reconstruction error (here showing Bayesian 1-step method for animal 1) often occurred when the instantaneous firing rate averaged over all cells was low, which in turn was correlated with low running speed. In the firing rate diagram, the dotted curve shows the inverses of the average firing rates; this serves only as a visual guide to emphasize the moments of low firing rates. Running speed and the reconstruction error also had correlates with the hippocampal local field potentials, because the sharp waves and the theta rhythm are known to be more prominent during immobility and motion, respectively. The amplitude diagram shows the absolute value of the local field potential signal of channel 1, which was slightly blurred and shown in arbitrary units. Continuous Fourier power spectra were obtained by sliding a Hanning window with a length 10 times the period of each frequency. Channel 1 was sampled at 1 kHz to capture sharp waves, which can be seen in the Fourier power spectra as thin vertical strips around 150–200 Hz. Channel 2 was sampled at 200 Hz to capture theta rhythm, which can be seen in the Fourier power spectra as horizontal dark bands around 8 Hz with 2nd harmonics around 16 Hz.

1. Lower running speed or immobility, which itself is known to be correlated with:

2. Lower firing rates of place cells [McNaughton, Barnes & O'Keefe (1983); see also Figure 8];

3. Occurrence of large amplitude irregular activity (LIA) and more frequent sharp waves in local field potentials (Vanderwolf, Kramis, Gillespie & Bland, 1975; Buzsáki, 1986) (In Figure 9 the sharp waves can be seen as the thin vertical dark bands around 150–200 Hz in the Fourier power spectra of channel 1 of the local field potentials);

4. The disappearance of the theta rhythm (Vanderwolf, Kramis, Gillespie & Bland, 1975) (In Figure 9, theta modulation can be seen as the fuzzy horizontal dark bands around 8 Hz in the Fourier power spectra of channel 2 of the local field potentials).

Because of the correlated factors listed above, reconstruction tended to be more accurate during the theta rhythm when the animal was running (Gerrard, Ginzburg, Barnes, Sejnowski & McNaughton, 1995). This is consistent with the report that place cell firing represented an animal's position somewhat more faithfully during the theta rhythm (Kubie, Fox & Muller, 1984).

Another related factor is the sharp waves, which are believed to originate in the CA3 region perhaps because of the lateral connectivity in this region (Buzsáki, 1986, 1989). Sharp waves are more frequent during immobility. It is possible that the activity during sharp waves might be involved with other functions such as memory consolidation and reactivation of learned correlations rather than merely reflect the current spatial position of the animal. Reconstruction based on sharp wave activity may yield an incorrect position. However, because the jumps did not always coincide with the occurrence of sharp waves, we can only conclude that sharp waves are one possible contributing factor for discontinuity in reconstructed path.

A further consideration concerns the hippocampal view cells found in monkey (Ono, Nakamura, Nishijo & Eifuku, 1993; Rolls, Robertson & Georges-Francois, 1995). One cannot completely rule out the possibility that the activity of some cells in rat hippocampus might be similarly modulated by the animal's view rather than by the current position alone. If this is true, the reconstructed position may jump when a rat looks around or shifts attention.

In summary, most erratic error jumps occurred while the animal was not running. Many of these jumps can be accounted for by a momentary drop of instantaneous firing rates of recorded cells because fewer spikes imply less information for reconstruction. Some other jumps might reflect real biological activity in the hippocampus when the animal stopped running to eat, look around, and plan the next move.

6 Reconstruction Accuracy: Theoretical Limits and Biological Implications

6.1 Fisher information

The accuracy of different reconstruction methods varied greatly. What is the maximum accuracy that can possibly be achieved? The problem is to estimate the maximum amount of information inherent in the neuronal spikes that can be extracted for reconstruction.

Fisher information is particularly suitable for addressing this question (Paradiso, 1988; Seung & Sompolinsky, 1993; Snippe, 1996). It differs from Shannon information, which has been more commonly used in estimating information contained in spikes [for a recent review, see Rieke, Warland, de Ruyter van Steveninck & Bialek (1997)]. Although these two information measures are related, their relationship is not straightforward (Cover & Thomas, 1991). Fisher information depends on the shape and the slope of the tuning function (see Appendix B). In the place cell example, if the spatial bins are scrambled, the amount of Fisher information would be different. By contrast, Shannon information would be the same regardless of the ordering because it is not directly related to the slope of a distribution function. Shannon information offers a useful measure on the distribution of place fields (Skaggs, McNaughton, Gothard & Markus, 1993), but no existing theory has linked it to reconstruction accuracy. See also Treves, Barnes & Rolls (1996) for related discussion.

A key property of Fisher information is that its inverse is a lower bound, called the Cramér-Rao bound, on the variance of all unbiased estimators (Cramér, 1946; Rao, 1965; Zacks, 1971; Scharf, 1991; Kay, 1993). This directly yields an estimate of the minimal reconstruction error that can possibly be achieved.

6.2 Minimal error under Gaussian tuning in two-dimensional space

To estimate Fisher information in place cell data, several properties of the place fields must be provided: the average shapes of place fields, the spiking statistics of the cells, the density of the place fields per unit area, and the statistics of the tuning parameters. Instead of using noisy real data for direct numerical evaluation, we use analytical models based on simplifying assumptions. Although these assumptions are not rigorously true, they are reasonably good approximations that allow useful analytical estimates to be derived. First, assume that the spatial tuning function is a two-dimensional Gaussian:

$$f(\mathbf{x}) = f_{\max} \exp\left(-\frac{\|\mathbf{x}-\mathbf{x}_0\|^2}{2\sigma^2}\right), \tag{52}$$

where $f_{\max}$ is the maximal firing rate of a cell and σ is the tuning width. The second assumption is that spiking statistics are Poisson and different cells fire spike trains that are statistically independent. The third assumption is that in the spatial region of interest there are an adequate number of place cells, and their centers are distributed uniformly (Muller, Kubie & Ranck, 1987; Wilson & McNaughton, 1993). Because the tuning width σ and the peak firing rate $f_{\max}$ vary from cell to cell, our final assumption is that the distributions of the two parameters are statistically independent of each other (Muller, Kubie & Ranck, 1987).

Here we only state the final results of the detailed derivation provided in Appendix B. The simplest result is that reconstruction error cannot be smaller than

$$\text{Minimal Error} \approx \mathcal{F}_2\sqrt{\frac{2\left\langle\sigma^2\right\rangle}{N_{\text{spikes}}}}, \tag{53}$$

where $\mathcal{F}_2 = \sqrt{\pi}/2$ is a constant correction factor, $\langle\sigma^2\rangle$ is the average of the square of the tuning width for different place fields, and N_{spikes} is total number of spikes collected from all the cells within the time window, namely,

$$N_{\text{spikes}} = \tau N f_{\text{mean}}. \tag{54}$$

Here N is the total number of cells recorded, f_{mean} is their mean firing rate, and τ is the length of the time window. The mean firing rate is related to the parameters of the Gaussian place fields by the formula:

$$f_{\text{mean}} = \frac{2\pi}{A}\left\langle f_{\max}\right\rangle\left\langle\sigma^2\right\rangle. \tag{55}$$

where A is the area of interest, in which the N place fields are uniformly distributed, and $\langle f_{\max}\rangle$ is the average peak firing rate of all place cells. Here the assumption of the independence of peak firing rate and the tuning width has been used.

Another equivalent error formula is

$$\text{Minimal Error} \approx \mathcal{F}_2 \sqrt{\frac{A}{\pi \tau N \langle f_{\max} \rangle}}. \tag{56}$$

The equivalence of the two Eqs. (56) and (53) follows immediately from Eqs. (54) and (55).

Equations (53) and (56) are expected to be more reliable when a larger number of cells is used. This is because when the sample size is too small, a reconstruction method may suffer from systematic bias; that is, for a given true position, the average reconstructed position in repetitive trials does not approach the true position. In addition, in our derivation of the error formulas, the procedure of replacing a sum by an integral makes sense only when a large number of cells is used (Appendix B).

Unbiased reconstruction is implicitly assumed when the Cramér-Rao bound for variance is used to estimate the minimal reconstruction error. Otherwise, it would be trivial for a reconstruction method to achieve zero variance simply by giving a constant output regardless of the input. Because we used no more than 30 place cells in the reconstruction, systematic bias did exist for all reconstruction methods. This factor may reduce the accuracy of the error formulas (cf. Fig. 12).

In deriving these formulas, we have implicitly assumed that the duration of the time window is small compared with the running speed so that the rat does not move too far during that period of time. So if the time window is too large compared the running speed, the formulas are no longer reliable.

6.3 Minimal error under Gaussian tuning in arbitrary dimensional space

To generalize the results in the preceding section, we can derive error formulas for an arbitrary spatial dimension D by assuming Gaussian tuning function, Poisson spike statistics, uniform distribution of the centers of the Gaussians in space, and statistical independence of the firing of different cells as well as their peak rate $f_{\max}$ and tuning width σ. As the counterpart of Eq. (53), the minimal reconstruction error based on Fisher information is

$$\text{Minimal Error} \approx \mathcal{F}_D \sqrt{\frac{D \langle \sigma^D \rangle / \langle \sigma^{D-2} \rangle}{N_{\text{spikes}}}}, \tag{57}$$

where $\langle\ \rangle$ means average over all the cells. The counterpart of Eq. (56) is

$$\text{Minimal Error} \approx \frac{C_D}{\sqrt{\eta \tau \langle f_{\max} \rangle \langle \sigma^{D-2} \rangle}}, \tag{58}$$

where constant coefficient

$$C_D = (2\pi)^{-D/4} \sqrt{D} \mathcal{F}_D \tag{59}$$

depends only on the dimension D. Here, as before, $\mathcal{F}_D$ is the correction factor for dimension D [see Table 3 and Eq. (117)]; σ is the width and $f_{\max}$ is the peak firing rate for the Gaussian tuning function (52) in D-dimensional space; and N_{spikes} is the total number of spikes collected from N cells within the time window τ:

$$N_{\text{spikes}} = \tau N f_{\text{mean}}. \tag{60}$$

The mean firing rate f_{mean} averaged over all cells is related to the tuning parameters by

$$f_{\text{mean}} = (2\pi)^{D/2} \frac{\eta}{N} \langle f_{\max} \rangle \left\langle \sigma^D \right\rangle. \tag{61}$$

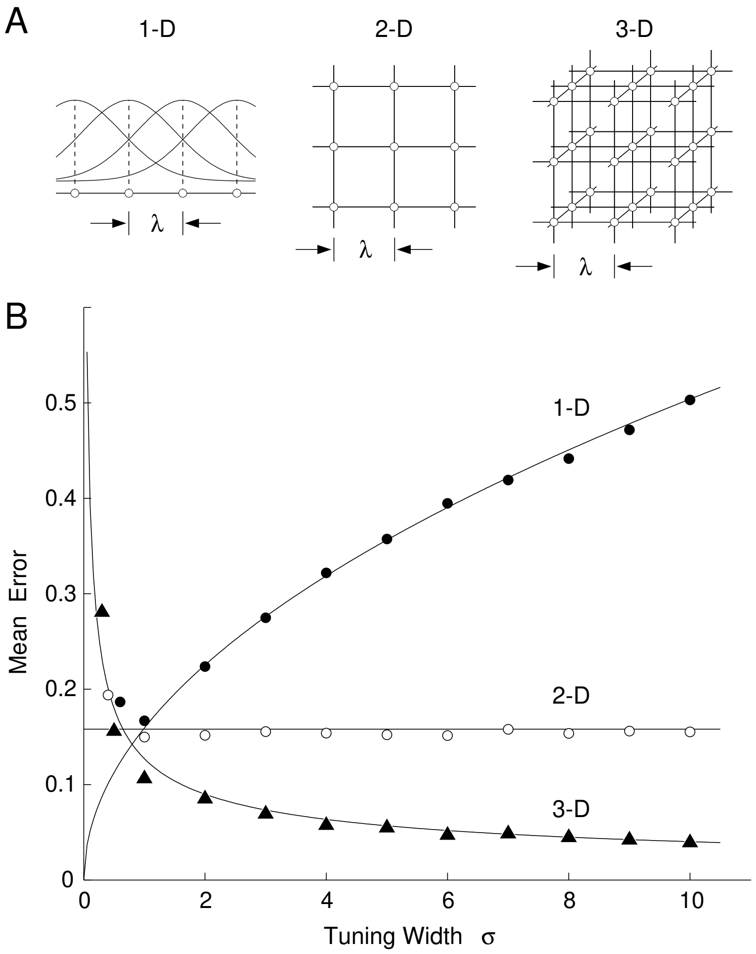

Figure 10: (A) Distribution of the centers of Gaussian tuning functions in one-, two-, and three-dimensional spaces. Only regular grids are shown here. In actual simulation, the centers were scattered randomly with a uniform distribution. The average length, area and volume per cell are given by λ, λ^2 and λ^3 for the one-, two- and three-dimensional cases, respectively. (B) The minimal achievable reconstruction error could increase, decrease or stay the same as the width of the Gaussian tuning function increases, depending on the dimension of space. The theoretical curves [Eq. (58)] are compared with numerical simulation (symbols) of Bayesian reconstruction assuming model cells with Poisson spike statistics and identical Gaussian tuning function. Each data point is the average of 10^4, 10^3, and 10^2 repetitive trials for the one-, two- and three-dimensional cases, respectively. Without loss of generality, we set $\lambda = 1$ in all simulation. The time window $\tau = 1$ second, and the peak firing rate $f_{\max} = 10$ Hz.

The equivalence of Eqs. (57) and (58) follows from Eqs. (60) and (61). The new parameter η is the cell density. For example, suppose N cells are randomly distributed within a D-dimensional cube with edge length L, we have

$$\eta \equiv \frac{1}{\lambda^D} = \frac{N}{L^D}. \tag{62}$$

In the one-dimensional case, $\eta = N/L$ is the number of cells per unit length; in the two-dimensional case, $\eta = N/L^2$ is the number of cells per unit area; and in the three-dimensional case, $\eta = N/L^3$ is the number of cells per unit volume. Parameter λ is a characteristic length so that λ^D is the average length, area, or volume per cell when $D = 1$, 2, or 3, respectively (Figure 10A). When $D = 2$, the new Eqs. (57) and (58) reduce to the old Eqs. (53) and (56) for place fields.

The minimal reconstruction error can be reached by the Bayesian method (cf. Appendix C). In Figure 10, the theoretical curves are compared with the simulation results of Bayesian reconstruction, assuming model cells with Poisson spike statistics and Gaussian tuning functions with randomly distributed centers but identical tuning width and peak firing rate. Here the error of the Bayesian method has practically reached the theoretical lower bound as long as the tuning width σ was not too small compared with the characteristic length $\lambda = 1$. In this example, it follows from Eq. (58) that

$$\text{Minimal Error} \propto \sigma^{1-D/2}. \tag{63}$$

In other words, as the tuning width σ increases, the error increases with $\sqrt{\sigma}$ in the one-dimensional case, and keeps constant in the two-dimensional case, and decreases with $1/\sqrt{\sigma}$ in the three-dimensional case.

6.4 Minimal error under cosine tuning and comparison with population vector

Cosine tuning is widely used to model motor cortical cells (Georgopoulos, Kettner & Schwartz, 1988). We will show that under cosine tuning, the Bayesian method can achieve the best possible accuracy by reaching the theoretical lower bound, whereas the performance of the population vector is well above the bound.

However, the population vector can be more accurate than the Bayesian method in special situations. An example is given in Figure 11, where two cells had orthogonal preferred directions, perfect cosine tuning, and Poisson spike statistics. The population vector outperformed the Bayesian method in a small region where the test direction was the furthest away from the two preferred directions.

Consider the one-dimensional situation that corresponds to a two-dimensional work space for reaching. We have N idealized cells with identical cosine tuning curves so that the mean firing rate of cell i is described by

$$f_i(\theta) = A\cos(\theta - \phi_i) + B, \tag{64}$$

where θ is the angle for the test direction and ϕ_i is the preferred direction, which is drawn randomly from a uniform distribution around the circle, and

$$A = \frac{f_{\max} - f_{\min}}{2}, \quad B = \frac{f_{\max} + f_{\min}}{2} \tag{65}$$

are constants. Suppose the spikes have Poisson statistics and different cells are independent. The average reconstruction error $\varepsilon_{\text{Bayes}}$ for the Bayesian method Eq. (36) is expected to approach the minimal achievable error defined by Fisher information and Cramér-Rao bound, provided that a

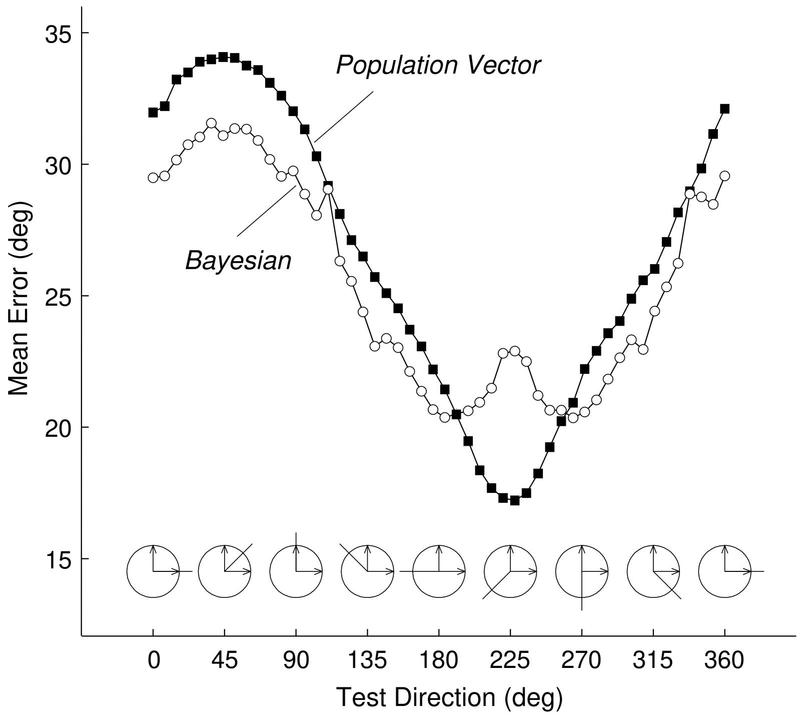

Figure 11: Mean angular errors for different test directions in a simple system with only two fixed orthogonal basis vectors, indicated by the arrows in the polar plots. A perfect cosine tuning curve and Poisson spike statistics were assumed. Each data point is an average over 5×10^4 repetitive trials. In this example, the population vector was more accurate than the Bayesian method when the test direction was close to 225°, but its average error across all directions (26.47°) was still larger than that of the Bayesian method (25.43°). Parameters: $f_{\max} = 10$ Hz, $f_{\min} = 1$ Hz, and the time window $\tau = 1$ second.

large number of cells is used (Appendix C). The minimal angular error (arc) derived by Fisher information is

$$\varepsilon_{\text{Bayes}} \approx \mathcal{F}_1 \sqrt{\frac{J_1^{-1}}{N}}, \tag{66}$$

where $\mathcal{F}_1$ is the one-dimensional correction factor (Table 3), N is the total number of cells, and

$$J_1 = \tau \left(\frac{f_{\max} + f_{\min}}{2} - \sqrt{f_{\max} f_{\min}} \right) \tag{67}$$

is the average Fisher information for a single cell over a time window τ. Simulation confirmed that this lower bound was indeed reached by the Bayesian method when the total number of cells was large enough (Figure 12).

In the same situation, the average angular error of the population vector method can also be estimated directly (Appendix D). The result is

$$\varepsilon_{\text{vector}} \approx \mathcal{F}_1 \sqrt{\frac{Q_1}{N}}, \tag{68}$$

where

$$Q_1 = \frac{1}{2} + \frac{2B}{\tau A^2}, \tag{69}$$

with A and B given by Eq. (65). In the two-dimensional case, corresponding to three-dimensional work space for reaching, Eq. (68) needs to be modified by replacing $\mathcal{F}_1$ by $\mathcal{F}_2$ and Q_1 by

$$Q_2 = \frac{6}{5} + \frac{6B}{\tau A^2}. \tag{70}$$

Because the error for the population vector in Eq. (68) and the error for the Bayesian method in Eq. (66) are both inversely proportional to the square root of the number of cells (see also footnote 1), the ratio

$$\frac{\varepsilon_{\text{vector}}}{\varepsilon_{\text{Bayes}}} \approx \sqrt{J_1 Q_1} \tag{71}$$

is independent of the number of cells. It can be verified analytically that $J_1 Q_1 > 1$ regardless of the values of $f_{\min}$, $f_{\max}$ and the time window τ. This means that here the Bayesian method is always more accurate than the population vector regardless of the choice of parameters.

6.5 Place cell reconstruction error compared with theoretical limits

Comparison with lower bound

As shown in Table 1 and Figure 5, our best reconstruction error by the Bayesian 2-step method using all the cells was larger than the theoretical limit in Eq. (53) by a factor of about two. To compute the theoretical lower bound, the width σ of the place fields was estimated using the first half of data by the method described in Appendix B. For animal 1, the average width of place fields was $\langle\sigma\rangle = 10.1$ cm, and $\sqrt{\langle\sigma^2\rangle} = 11.2$ cm. For animal 2, $\langle\sigma\rangle = 8.6$ cm, and $\sqrt{\langle\sigma^2\rangle} = 9.6$ cm. The mean firing rates of the 25 cells from animal 1 was $f_{\text{mean}} = 0.92$ Hz; and the average of the 30 cells from animal 2 was $f_{\text{mean}} = 1.09$ Hz.

Since the Bayesian 2-step method uses information from two time steps, is it fair to compare its performance with the Fisher information limit evaluated with the time window of a single time step? This may still be a reasonable comparison because the information from the previous time

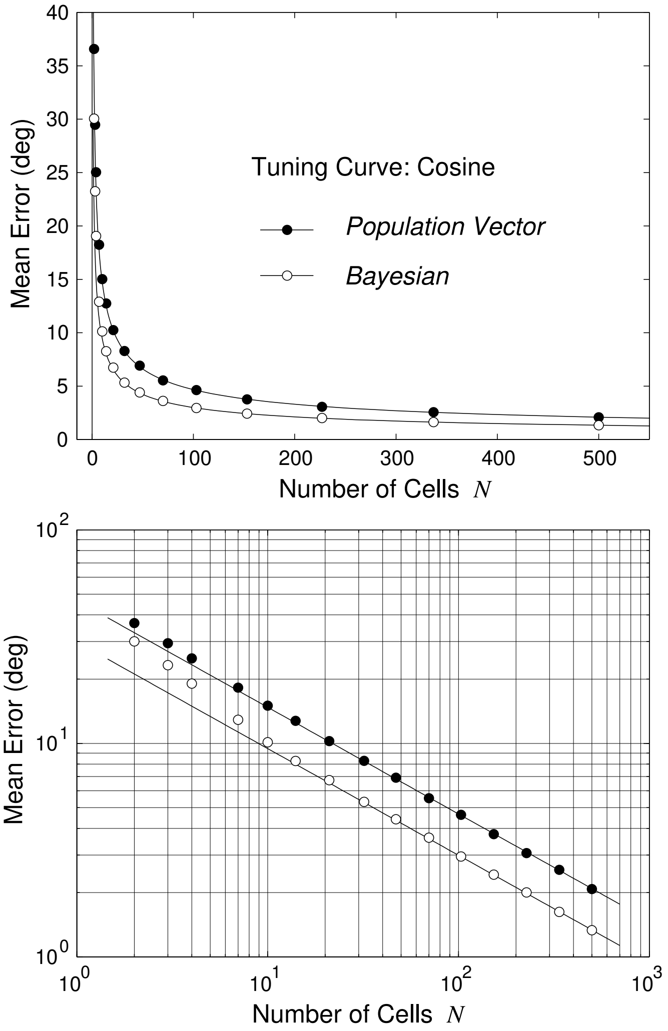

Figure 12: Comparison of the errors of the population vector and the Bayesian methods for two-dimensional reaching problem assuming cosine tuning, Poisson spike statistics, and a uniform probability distribution of the preferred directions around a circle. The simulation results (symbols) agreed with the theoretical results [curves and the straight lines in the log-log plot, given by Eqs. (66) and (68)] when the number of cells N was large so that the errors of both methods dropped inversely proportional to $\sqrt{N}$. In particular, the Cramér-Rao lower bound (the lower theoretical curve or line) was reached asymptotically by the Bayesian method when there were more than 20 cells. Parameters: $f_{\max} = 10$ Hz, $f_{\min} = 1$ Hz, and each data point is the average of 10^4 trials with randomly chosen preferred directions.

step is very crude: in actual reconstruction the width of the Gaussian prior distribution varied from 20 cm to 60 cm. It improved reconstruction accuracy mainly by prohibiting erratic jumps caused by low firing rates, rather than by providing precise location information. For example, it can be seen in Figure 3 that the Bayesian 1-step method, while suffering from the erratic jumps, already provides accurate information about the fine details of the movement trajectory. In animal 2, whose average firing rate was higher, the continuity constraint from the preceding time step actually did not contribute much to the accuracy (Figure 5).

In idealized situations, the performance of the Bayesian method should reach the Cramér-Rao lower bound (Figures 10 and 12, and Appendix C). For real place cell data, the basic assumptions of Poisson spike statistics and independence for the Bayesian formulas are not exactly true, which implies suboptimal performance. Reconstruction bias caused by the relatively small number of cells is another source of error (cf. Figure 12). Finally, the estimate of Cramér-Rao bound itself relies on some simplifying assumptions such as Gaussian tuning, which may lead to additional error.

Correction factor and spatial dimension

The correction factor $\mathcal{F}_D$ is the ratio of the mean error over the square root of the mean square error. Its theoretical value depends only on the dimension D of the work space, assuming Gaussian error distribution (Appendix B). As a consequence, its empirical value could be used to estimate the dimension D. For reconstruction errors obtained for all time steps by the Bayesian 2-step method, the empirical value was 0.8705 for animal 1, and 0.8855 for animal 2, surprisingly close to the value of the two-dimensional correction factor $\mathcal{F}_2 = 0.8862$ but quite different from $\mathcal{F}_1$ and $\mathcal{F}_3$ (Table 3). However, the actual error distribution differed from the theoretical probability density $\rho_2(r)$ (Table 3 and Figure 16) because it peaked at a smaller value than the standard deviation ($r = s$) and had a longer tail for large errors. This might be related to the fact that the rats run in two-dimensional space but within the confinement of a track so that the work space was neither purely two-dimensional nor purely one-dimensional.

Counterintuitive effect of place field size

Imagine we have a fixed number of place cells all having place fields within a given region. Suppose we can change the size of the place fields without altering the positions of the place field centers and the peak firing rates. How would this change affect the accuracy of spatial information encoded in the population? The same question in a different form is how the minimal achievable reconstruction error depends on the place field size. This question is related to how the accuracy of spatial representation is affected by the broader place fields found in ventral hippocampus of rats (Jung, Wiener & McNaughton, 1994) and in the hippocampus of the genetically engineered mouse (see below). The size of place fields also tends to increase during repetitive running on a track (Mehta, Barnes & McNaughton, 1997).

When the peak firing rate is fixed, changing the size of place fields should have no effect on the minimal reconstruction error or the accuracy of spatial representation. This is because Eq. (56) does not contain the width σ of the place fields. Here we have assumed that the place fields are of Gaussian shape, the spikes are Poisson, and different cells are independent.

The explanation is that larger place fields imply more overlaps and therefore higher overall average firing rate of the whole population, which will reduce reconstruction error. This effect turns out to precisely compensate for the less accurate spatial information provided by a single enlarged place field. A confirmation by simulation is shown in Figure 10. The fact that larger place field sizes imply higher mean firing rates is clearly seen in Eq. (55). This effect precisely cancels out the effect of the tuning width σ in Eq. (53) so that the width σ is absent in the equivalent Eq. (56).

A related issue is the recent reports of increased place field size for CA1 cells in the hippocampus of genetically altered mice with impairment of long-term potentiation (LTP) restricted to CA1 (McHugh, Blum, Tsien, Tonegawa & Wilson, 1996; Rotenberg, Mayford, Hawkins, Kandel & Muller, 1996). A decease of the peak firing rate of place cells was explicitly reported by Rotenberg, Mayford, Hawkins, Kandel & Muller (1996). According to Eq. (55), such a decrease was also implied in the observation that the mean firing rate of the genetically altered mice were the same as the normal (McHugh, Blum, Tsien, Tonegawa & Wilson, 1996). The increase of reconstruction error for these mice (McHugh, Blum, Tsien, Tonegawa & Wilson, 1996) may be attributed to a decrease in peak firing rate rather than to an increase in place field size per se, because Eq. (56) suggests that reconstruction error should be independent of place field size. Equivalently, if the mean firing rate is fixed, it is equally valid to apply Eq. (53), which suggests that in this case the error should be larger for larger place fields. Instability of place fields is another possible source of reconstruction error (cf. Rotenberg, Mayford, Hawkins, Kandel & Muller (1996)). A more detailed analysis of the data would be required for quantitative comparison.

The accuracy of spatial representation can be improved by increasing the number of cells or their firing rates, but not by sharpening the place fields. For example, after repetitive running on a track, both the size of the place fields and the average firing rate tended to increase (Mehta, Barnes & McNaughton, 1997). The increased firing rate would imply better accuracy of spatial representation, regardless of the place field size.

6.6 General biological discussion

The minimal achievable reconstruction considered in the preceding sections is actually a general measure of how accurately a physical variable is encoded by a population of neurons, regardless of the decoding method used. As a consequence, the theoretical limit is valid not only for the computational problem of reconstruction, but also for all neurophysiological and behavioral measures of acuity or accuracy in the sense of mean square error, as long as the information is derived solely from the same neuronal population.

We have derived explicit formulas for cells with Gaussian and cosine tuning functions because of their common usage, although the methods can be readily adapted for more general situations. From these formulas, we know precisely how the reconstruction accuracy depends on the parameters of individual cells and the cell number density. For example, suppose we have spatial region of 1 meter $\times$ 1 meter. How many place cells must be used to achieve a spatial acuity of $d = 1$ cm everywhere in this region? Using Eq. (56), we have

$$N = \frac{A}{4d^2 f_{\max}\tau}. \tag{72}$$

Assuming a mean peak firing rate $f_{\max} \sim 15$ Hz and a biologically plausible time window $\tau \sim 200$ msec, we get $N \sim 10^3$. That is, a minimum of about one thousand cells would be required. As another example, using up all the place cells in a rat, say, $N \sim 10^5$ (Amaral & Witter, 1995), what is the maximum area that can be covered under the acuity of 1 cm? Using the same parameters considered above, we find the total area is about $A \sim 10^6$ cm^2, the same as a 10 meter $\times$ 10 meter square.

The insensitivity of place field size on reconstruction accuracy is a peculiarity of two dimensions. As shown in Figure 10, more information is encoded by a sharper tuning curve in one-dimensional

Table 2: Unified formulation of reconstruction methods

Methods	Basis $\phi_i(\mathbf{x})$	Bias $A(\mathbf{x})$
Vector methods[a]	$\mathbf{x}_i \cdot \mathbf{x}$	0
Direct basis[b]	$f_i(\mathbf{x})P(\mathbf{x})$	0
Reciprocal basis[b]	$g_i(\mathbf{x})P(\mathbf{x})$	0
Probabilistic (1-step)	$\ln f_i(\mathbf{x})$	$\ln P(\mathbf{x}) - \tau \sum_i f_i(\mathbf{x})$
Probabilistic (2-step)[c]	$\ln f_i(\mathbf{x})$	'above' $- k \left\|\mathbf{x} - \mathbf{x}_{\text{pre}}\right\|^2 / u(\mathbf{x})^d$

Notes: A summary of reconstruction methods using Eq. (74). [*a*] Examples of vector methods include population vector and optimal linear estimator. Here only a directional vector is shown. For a positional vector like two-dimensional position, scaling is needed as in Eq. (32). [*b*] f_i and g_i are related by pseudoinverse. [*c*] In the probabilistic 2-step method (Bayesian), $k > 0$ is a constant, $\mathbf{x}_{\text{pre}}$ is the reconstructed position from the preceding time step, $u(\mathbf{x})$ is the dimensionless average speed at position $\mathbf{x}$, and d is a constant ($1 \le d \le 2$).

space. Examples of one-dimensional tuning curves include the orientation tuning in visual cortex and directional tuning of head-direction cells (Taube, Muller & Ranck, 1990). In contrast, in three-dimensional space or space of higher dimensions, a sharper tuning width implies worse information encoding. So the two-dimensional example of place cell is a critical case where the tuning width does not matter. To increase representation accuracy, one should use neurons with broader tuning functions to represent a variable higher than two dimensions, assuming that the peak firing rate is fixed. However, broader tuning width implies higher energy consumption because more neurons are activated at any given time. See Hinton, McClelland & Rumelhart (1986), Baldi & Heiligenberg (1988), Zohary (1992), Gerstner, Kempter, van Hemmen & Wagner (1996) for related discussions on the consequences of broad tuning.

7 Biological Plausibility of Reconstruction Methods

7.1 Unified formulation

The main result in this section is that all the reconstruction methods discussed in this paper can be implemented as linear feedforward neural networks. This result is in sharp contrast with the general belief that reconstruction methods such as the Bayesian method or the template matching method are merely mathematical techniques without any biological relevance.

The reconstructed position $\hat{\mathbf{x}}$ in all reconstruction methods considered above can be expressed as

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}} \psi(\mathbf{x}), \tag{73}$$

and function $\phi(\mathbf{x})$ can be written in the form

$$\psi(\mathbf{x}) = \sum_{i=1}^{N} n_i \phi_i(\mathbf{x}) + A(\mathbf{x}), \tag{74}$$

where as before n_i is the number of spikes that cell i fired in a certain time window, $\phi_i(\mathbf{x})$ is a basis function associated with cell i. The function $A(\mathbf{x})$ is an additive bias, which is independent of the

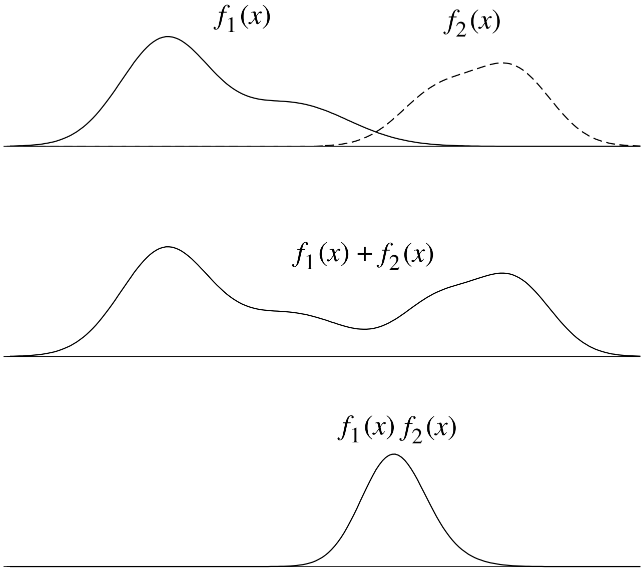

Figure 13: An example of addition and multiplication of two basis functions. While addition yields a distribution with two peaks, multiplication yields only a single peak. The key operation of the Bayesian reconstruction method is multiplication, and its output distribution is often sharper and less likely to contain multiple peaks than the output distribution of the basis function method, whose key operation is addition. See Figure 2 for an actual example of comparison.

activity of the cells and is reminiscent of a regularization term (Tikhonov & Arsenin, 1977; Poggio, Torre & Koch, 1985). See also Zemel, Dayan & Pouget (1997) for related discussion. The bias is needed only for implementing Bayesian methods. For all basis function methods, there is no bias term; that is, $A(\mathbf{x}) \equiv 0$. See Table 2 for a summary.

Here the Bayesian methods are implemented by taking the logarithm of the posterior probability in Eqs. (36) and (41). Since logarithm is monotonic and does not change the peak position, maximizing the product $\prod_i f_i(\mathbf{x})^{n_i}$ is equivalent to maximizing the sum $\sum_i n_i \phi_i(\mathbf{x})$, where

$$\phi_i(\mathbf{x}) = \ln f_i(\mathbf{x}). \tag{75}$$

To avoid taking the logarithm of zero, an arbitrary small positive number may be added to the tuning function f_i. The other factors in the Bayesian Eq. (36) becomes the additive bias

$$A(\mathbf{x}) = \ln P(\mathbf{x}) - \tau \sum_{i=1}^{N} f_i(\mathbf{x}), \tag{76}$$

which does not depend on activity n_i. An intuitive interpretation of the bias is that it favors positions that animal visits frequently, that is, places with high $P(\mathbf{x})$, and at the same time disfavors places that are over-represented by a large number of cells or cells with excessively high firing rates. The bias term becomes a constant and can be ignored if the animal is equally likely to visit all positions and the average firing rates of all cells are uniform across space.

The difference between addition and multiplication of basis functions is illustrated in Figure 13, although taking logarithm can transform a product into a sum. For the direct basis method, the computation is the sum:

$$n_1 f_1(\mathbf{x}) + n_2 f_2(\mathbf{x}) + \cdots = \sum_i n_i f_i(\mathbf{x}). \tag{77}$$

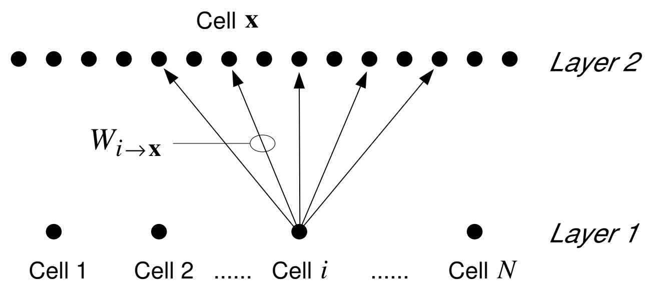

Figure 14: A unified network implementation for various reconstruction methods. The first layer contains N cells tuned arbitrarily to an variable $\mathbf{x}$. The cells in the second layer represent the value of $\mathbf{x}$ by their locations in the layer. The activity distribution in the second layer represents how likely it is for each possible value of $\mathbf{x}$ to be the true value.

In Bayesian method, the crucial part of the main Eq. (36) is the product:

$$f_1(\mathbf{x})^{n_1} f_2(\mathbf{x})^{n_2} \cdots = \prod_i f_i(\mathbf{x})^{n_i}. \tag{78}$$

The product is like a logical AND operation and the final peaks correspond to those positions where all contributing template functions in the product are not close to zero. By contrast, the sum behaves like a logical OR operation and all peaks from contributing template functions are preserved. This difference explains why output distributions of basis function methods are usually much broader and more likely to contain multiple peaks than that of the Bayesian methods (Figure 2).

7.2 Biological implications

The unified formulation of reconstruction methods in the preceding section can be implemented readily by a biologically plausible neural network structure. The first computational step in Eq. (74) is linear combination of fixed basis functions. Here the coefficients are the numbers of spikes, which are proportional to the instantaneous firing rates within the given time window, which is plausibly a fraction of a second. The linear combination only requires a feedforward network in which the basis functions are implemented by the connection weights (Figure 14). More precisely, the sum $\sum_{i=1}^{N} n_i \phi_i(\mathbf{x})$ is computed as

$$\sum_{i=1}^{N} n_i \phi_i(\mathbf{x}) = \sum_{i=1}^{N} \left(\frac{n_i}{\tau} \right) W_{i \to \mathbf{x}} \ , \tag{79}$$

where n_i/τ is the instantaneous firing rate of cell i in the first layer, and $W_{i \to \mathbf{x}}$ is the weight from cell i in the first layer to cell $\mathbf{x}$ in the second layer. The weight pattern is defined to be proportional to the template function: $W_{i \to \mathbf{x}} \equiv \tau \phi_i(\mathbf{x})$.

Indeed, linear combination of basis functions is biologically plausible computational strategy that has been used by many authors, such as in three-dimensional object representations (Poggio & Girosi, 1990; Ullman & Basri, 1991), parietal spatial representation (Zipser & Andersen, 1988; Pouget & Sejnowski, 1994; Pouget & Sejnowski, 1997), modeling place cells (Zipser, 1985; Redish & Touretzky, 1996; O'Keefe & Burgess, 1996), and or as a general strategy for neural representation (Poggio, 1990; Anderson & Van Essen, 1994; Zemel, Dayan & Pouget, 1997).

The maximization in the second computational step in Eq. (73) can be implemented by any circuits that can approximate a winner-take-all operation on the activity distribution in the second layer after the first feedforward step. An exact winner-take-all implementation might not be needed in the biological system because there are different ways of reading the distributed information after the first step is done. In general, the full activity distribution after the first feedforward step contains more information about the probability distribution of the position, and it may be useful to maintain this information for the purpose of further computations that need to take into account the variance of the estimate as well as the mean of the position (Nowlan & Sejnowski, 1995).

The bias term for the 1-step Bayesian method can be implemented simply as constant inputs to the second layer. To implement the additional bias for the 2-step Bayesian method, a type of short-term memory is needed for the winner-take-all operation. The quadratic bias in Table 2 comes from the Gaussian distribution in Eq. (43), but the precise functional form is not critical. In network implementation, any facilitation mechanism that favors units close to the previously selected winner unit and prohibits units far away should suffice.

Therefore, it is feasible for a biological system to implement all the reconstruction methods considered in this paper. Moreover, a simple Hebb rule should be enough for learning the weights used in the Bayesian method and the direct basis method because the desired template function for each cell depends only the tuning function of the same cell. However, the reciprocal basis method may need a nonlocal learning rule.

A related observation is that attractor dynamics in a recurrent network may perform operations similar to winner-take-all, noise cleaning or maximum likelihood estimation (Pouget & Zhang, 1997). This approach has some biological plausibility because attractor dynamics might actually be used by the brain, for example, in the hippocampal system (McNaughton, Barnes, Gerrard, Gothard, Jung, Knierim, Kudrimoti, Qin, Skaggs, Suster & Weaver, 1996; Samsonovich & McNaughton, 1997; Tsodyks & Sejnowski, 1995a; Zhang, 1996), the primary visual cortex (Ben-Yishai, Bar-Or & Sompolinsky, 1995; Somers, Nelson & Sur, 1995; Tsodyks & Sejnowski, 1995b), and motor cortex (Lukashin, Amirikian, Mozhaev, Wilcox & Georgopoulos, 1996).

One may ask which reconstruction algorithm is actually used by the biological system, for example, for spatial navigation using the hippocampus place cells. There are a few general predictable effects. According to the unified network implementation in Eq. (74), different methods are distinguished mainly by the template functions that are implemented as synaptic weights. For example, the direct basis method would imply a weight pattern directly proportional to the mean firing rate given by the tuning function, whereas the Bayesian method would imply a weight pattern proportional to the mean firing rate on a logarithm scale. Whether the brain uses the most efficient construction method might also be tested by spatial acuity at neural or behavioral levels. For our dataset, such information is not available. What we have shown here is that the biological system does have the computational resources to implement the most efficient algorithms we considered, which can reach the best possible accuracy defined by Cramér-Rao lower bound under idealized situations.

8 Conclusions

The goal of this paper has been in part to compare different reconstruction methods, to improve the accuracy of these methods, and to assess their performance against the theoretical lower bound on reconstruction accuracy for all possible methods. For trajectory reconstruction based on spike trains from simultaneously recorded hippocampal place cells, probabilistic methods based on Bayesian formula were exceptionally accurate, especially when information about the reconstructed position from the previous time step was used to discourage discontinuities in the trajectory. In our best reconstruction using an average of up to about 30 spikes within the time window, the mean errors were in the range of the 5 cm intrinsic error of the position tracking system.

This reconstruction study has also revealed some interesting properties of place cells, such as erratic jumps of reconstruction errors and their correlation with movement and the local field potentials, the slow systematic drift of reconstruction error, as well as a possible shift of the optimal center of the reconstruction time window towards the past during continuous periodic running.

There are two major implications for our consideration of the theoretical lower bound on reconstruction accuracy derived by Fisher information and Cramér-Rao inequality. First, this lower bound determines precisely how accurately a physical variable is encoded in the neuronal population in the sense of mean square error, regardless of which method is used for reading out the information. That is, this bound quantifies the amount of information encoded in the neuronal population regardless of decoding method. As a consequence, the best achievable accuracy is valid not only for the reconstruction problem, but also for all possible neural or behavioral measures of acuity or accuracy, as long as the information is derived solely from the neuronal population in question. Second, we have identified that the Bayesian reconstruction methods can reach the theoretical lower bound when the conditions of Poisson spike statistics and independence of different cells are satisfied. In such situations, optimal reconstruction has been achieved. For our place cell dataset, the best reconstruction errors were above the estimated theoretical limit by perhaps a factor of two or so. The remaining discrepancy may be caused by various approximations involved.

One counterintuitive result emerging from our theoretical consideration of accuracy is that making the size of all place fields smaller (or larger) without changing their peak firing rate would have no effect on the accuracy of coding the animal's position in the whole population. However, for one-dimensional work space, a narrower tuning curve gives more accurate coding. By contrast, for three and higher dimensions, a broader tuning function gives more accurate coding. Thus two dimensions are a critical case in which the coding accuracy is independent of the width of the tuning function.

Despite popular belief to the contrary, we have shown that all reconstruction methods considered in this paper including template matching and the Bayesian method can be implemented as simple feedforward neural networks. This result suggests that biological systems are capable of implementing all of these reconstruction methods. The reduction of the nonlinear Bayesian method to the linear basis function framework relies on the assumptions of Poisson spiking statistics and independence of different cells. For the place cell data, they are reasonable approximations, although they are not exactly true.

Whenever neuronal activity is correlated with a measurable physical variable, reconstruction of this variable from population activity is a relevant problem both as a research tool and as a hint to how the brain might solve the same problem. There are no intrinsic constraints on the type of physical variables that can be reconstructed, or on the type of tuning functions that the cells can encode. For example, instead of the position in a continuous two-dimensional space, the variable could as well be discrete and disjoint, which might be more suitable for representing distinct classes of objects or categories.

Acknowledgments

We thank Hemant Kudrimoti and William E. Skaggs for their valuable help in data processing and for many discussions, Jie Wang for cluster cutting, Michael S. Lewicki for discussions on Bayesian methods, Alexandre Pouget for discussions on population coding and Fisher information, Andrew B. Schwartz and Robert U. Muller for helpful suggestions on the manuscript, and two anonymous reviewers for useful comments. Data acquisition by Jason L. Gerrard and Carol A.

Barnes. Supported by HHMI and NS20331.

Appendices

Appendix A. Reciprocal basis with modifications

In this section we modify the reciprocal basis by taking into account a nonuniform probability distribution for visiting different positions as well as the randomness in firing rate. As with the original reciprocal basis, the modified basis functions are optimal with respect to mean square error. However, the reciprocal relationship is lost. For our dataset, the performances of these modified methods were sometimes better than the original reciprocal basis, as shown in Table 1.

In general, consider a basis function method

$$\hat{\mathbf{x}}_{\text{basis}} = \arg\max_{\mathbf{x}} \sum_{i=1}^{N} \frac{n_i}{\tau} g_i(\mathbf{x}), \tag{80}$$

where N is the total number of cells, $\mathbf{n} = (n_1, \ldots, n_N)$ are the numbers of spikes fired by these cells within the time window τ, and $(g_1, \ldots, g_N)$ are the basis functions to be determined. Here τ is included in the formula for the convenience of scaling; it does not change the result. The discrete version of the basis function $g_i(\mathbf{x})$ is the column vector $\mathbf{g}_i$, which is defined by first digitizing the position $\mathbf{x}$ and then concatenating the pixels. The vectors are collected into a single matrix $\mathbf{G} = [\mathbf{g}_1, \ldots, \mathbf{g}_N]$. As shown in Eq. (15), the simple reciprocal basis is related to the firing rate matrix $\mathbf{F}$ by

$$\mathbf{G} = \mathbf{F}^{\dagger\mathrm{T}} = \mathbf{F}(\mathbf{F}^{\mathrm{T}}\mathbf{F})^{-1}, \tag{81}$$

where the second equality is valid only when the matrix inverse is nonsingular.

The reciprocal basis functions obtained by Eq. (81) are appropriate only if the animal visits each spatial position equally frequently. When the visiting probability is nonuniform, there are several ways to modify the matrix $\mathbf{G}$ or the basis functions. First, the simple Eq. (21) is the same as the prescription

$$\mathbf{G} = \mathbf{P}\mathbf{F}^{\dagger\mathrm{T}} = \mathbf{P}\mathbf{F}(\mathbf{F}^{\mathrm{T}}\mathbf{F})^{-1}. \tag{82}$$

This is the formula used in our systematic comparison of different reconstruction methods in Figure 5. Alternatively, we may choose

$$\mathbf{G} = \sqrt{\mathbf{P}}\left(\sqrt{\mathbf{P}}\mathbf{F}\right)^{\dagger\mathrm{T}} = \mathbf{P}\mathbf{F}(\mathbf{F}^{\mathrm{T}}\mathbf{P}\mathbf{F})^{-1}. \tag{83}$$

A third possible formulation is

$$\mathbf{G} = \mathbf{P}\mathbf{F}\left(\mathbf{F}^{\mathrm{T}}\mathbf{P}\mathbf{F} + \frac{1}{\tau}\mathbf{D}\right)^{-1}. \tag{84}$$

Here three diagonal matrices have been used, which are defined as

$$\mathbf{P} = \begin{pmatrix} P_1 & & & \\ & P_2 & & \\ & & \ddots & \\ & & & P_K \end{pmatrix}, \tag{85}$$

$$\sqrt{\mathbf{P}} = \begin{pmatrix} \sqrt{P_1} & & & \\ & \sqrt{P_2} & & \\ & & \ddots & \\ & & & \sqrt{P_K} \end{pmatrix}, \tag{86}$$

and

$$\mathbf{D} = \begin{pmatrix} \bar{f}_1 & & & \\ & \bar{f}_2 & & \\ & & \ddots & \\ & & & \bar{f}_N \end{pmatrix}, \tag{87}$$

where K is the total number of position pixels, P_k is the probability for visiting position k, and $\bar{f}_i$ is the average firing rate of cell i.

These formulas are related. In Eq. (84), the $\mathbf{D}$ term is related to firing rate variability (see below). For large time window τ, the $\mathbf{D}$ term vanishes, and Eq. (84) is reduced to Eq. (83). When the visiting probability P_k is the same for all positions, Eq. (83) is reduced to the original reciprocal basis Eq. (81).

Equation (84) has been derived by minimizing the error function

$$E = \sum_{k,\mathbf{n}} \left\| \mathbf{I}_k - \mathbf{G}\left(\frac{1}{\tau}\mathbf{n}\right) \right\|^2 P_k P_{\mathbf{n}|k}, \tag{88}$$

where $P_{\mathbf{n}|k}$ is the probability for the spike vector $\mathbf{n}$ to occur at position k, and the position k is represented by the vector $\mathbf{I}_k$, the k-th column of the $K \times K$ identity matrix. Assuming Poisson spike statistics, we can expand Eq. (88) and evaluate each term using

$$\langle n_i n_j \rangle = \begin{cases} \langle n_i \rangle^2 + \langle n_i \rangle, & i = j \\ \langle n_i \rangle \langle n_j \rangle, & i \neq j \end{cases} \tag{89}$$

where $\langle\ \rangle$ means average over spikes, after setting derivative of the error function with respect to matrix $\mathbf{G}$ to zero. Thus we obtain Eq. (84). The diagonal matrix $\mathbf{D}$ appears here because the self-correlation $\langle n_i^2 \rangle = \langle n_i \rangle^2 + \langle n_i \rangle$ contains the additional term $\langle n_i \rangle$ due to firing rate variability [cf. Salinas & Abbott (1994)].

When firing rate variability is ignored, the equation

$$\frac{1}{\tau}\mathbf{n} = \mathbf{F}^{\mathrm{T}}\mathbf{I}_k \tag{90}$$

holds exactly. This is a reasonable approximation when the firing rates are high or the time window is large. Dropping out the term $P_{\mathbf{n}|k}$ in Eq. (88), we only need to minimize

$$E = \sum_{k=1}^{K} \left\| \mathbf{I}_k - \mathbf{G}\mathbf{F}^{\mathrm{T}}\mathbf{I}_k \right\|^2 P_k = \left\| \left(\mathbf{I} - \mathbf{G}\mathbf{F}^{\mathrm{T}}\mathbf{I}\right)\sqrt{\mathbf{P}} \right\|^2. \tag{91}$$

The result is Eq. (83).

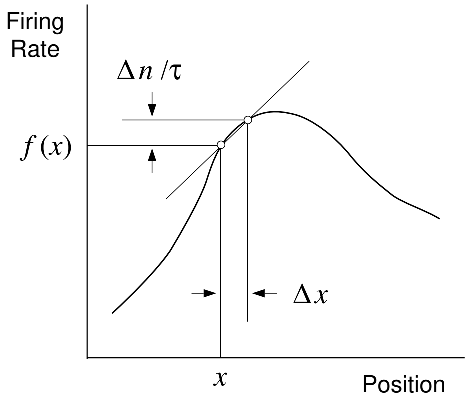

Figure 15: Intuitive argument for Fisher information in estimating the position x from the instantaneous firing rate and a known tuning curve $f(x)$. Fluctuation in the firing rate leads to an error Δx in the estimated position, which decreases with the slope of the tuning curve.

Appendix B. Minimal Reconstruction Error Based on Fisher Information

Basic approach

We provide an intuitive argument for using Fisher information and then sketch our method for estimating a lower bound for reconstruction error using the Cramér-Rao inequality.

As an intuitive example, consider a one-dimensional situation with an idealized place cell whose firing rate $f(x)$ depends only on position x (Figure 15). Suppose the cell fired n spikes within the time window τ. Because of random variation in the exact number of spikes, the formula

$$n/\tau = f(x) \tag{92}$$

is true only on average but not for each individual trial. When we use this single cell to estimate the position x from the actual firing rate n/τ and the shape of the tuning curve $f(x)$, the fluctuation of the number of spikes would result in an error Δx that obeys

$$\Delta n/\tau \approx f'(x)\Delta x. \tag{93}$$

Thus the mean squared error is

$$\left\langle \Delta x^2 \right\rangle \approx \frac{\langle \Delta n^2 \rangle}{\tau^2 f'(x)^2} = \frac{f(x)}{\tau f'(x)^2}. \tag{94}$$

In the last step, Poisson spiking statistics was assumed; that is, the variance $\langle \Delta n^2 \rangle$ is equal to the mean $\langle n \rangle = \tau f(x)$. Here the reconstruction error is expressed in terms of the tuning function $f(x)$ and its slope $f'(x)$. It makes sense that a larger slope implies smaller error because a small change of position would imply a large and noticeable change in the firing rate (Figure 15).

An identical result is obtained by using the formal definition of Fisher information (Rao, 1965; Zacks, 1971; Cover & Thomas, 1991; Scharf, 1991; Kay, 1993; Seung & Sompolinsky, 1993):

$$J(x) = \left\langle \left| \frac{\partial}{\partial x} \ln P(n|x) \right|^2 \right\rangle = \frac{\tau f'(x)^2}{f(x)}, \tag{95}$$

where the last equality follows from straightforward calculation under the assumption that the probability $P(n|x)$ of firing n spikes at position x has Poisson distribution with the mean firing rate $\tau f(x)$.

The squared error for any unbiased estimator of position x cannot exceed the Cramér-Rao bound, namely,

$$\left\langle \Delta x^2 \right\rangle \geq \frac{1}{J(x)}, \tag{96}$$

whose validity originates from the universal Cauchy-Schwartz inequality. This result agrees with the intuitive Eq. (94).

When more than one cell is used in reconstruction, as long as they are independent, the total Fisher information is the sum

$$J(x) = \sum_{i=1}^{N} J_i(x). \tag{97}$$

This leads to an effective composition rule of minimal errors:

$$\frac{1}{\varepsilon^2} = \sum_{i=1}^{N} \frac{1}{\varepsilon_i^2}, \tag{98}$$

where ε^2 is the minimum achievable value of $\langle \Delta x^2 \rangle$ by using all the cells from 1 to N, and ε_i^2 is defined as the same error using only cell i.

So far, we have only considered a one-dimensional problem. Extension to two-dimensional estimates is simple if we assume that the error statistics in x and y directions are the same. Then the squared error of reconstruction is

$$\left\langle \Delta x^2 \right\rangle + \left\langle \Delta y^2 \right\rangle = 2 \left\langle \Delta x^2 \right\rangle \geq \frac{2}{J}, \tag{99}$$

where J is the Fisher information for variable x alone. The minimal reconstruction error is estimated as $\mathcal{F}_2 \sqrt{2/J}$. The final result is given by Eq. (56) or Eq. (53).

Derivation of error formula in arbitrary dimensions

Following the same approach considered in the preceding section, we list the keys steps in deriving the general formula for minimal reconstruction error in arbitrary spatial dimension D. The mean firing rate of each cell is assumed to obey a Gaussian tuning function in a D-dimensional space:

$$f(\mathbf{x}) = f_{\max} \exp\left(-\frac{\|\mathbf{x} - \mathbf{c}\|^2}{2\sigma^2} \right), \tag{100}$$

where $\mathbf{x} = (x_1, \ldots, x_D)$ and $\|\mathbf{x}\|^2 = x_1^2 + \cdots + x_D^2$ and $\mathbf{c}$ is the center of the Gaussian. The centers for different cells are assumed to be uniformly distributed in space. The peak firing rate $f_{\max}$ and the tuning width σ may differ from cell to cell. For each cell, at position $\mathbf{x}$, the probability for obtaining n spikes within a time window τ is assumed to obey a Poisson distribution:

$$P(n|\mathbf{x}) = \frac{(\tau f(\mathbf{x}))^n}{n!} \exp(-\tau f(\mathbf{x})). \tag{101}$$

The Fisher information with respect to a single dimension, say x_1, is

$$J^{[1]}(\mathbf{x}) = \left\langle \left| \frac{\partial}{\partial x_1} \ln P(n|\mathbf{x}) \right|^2 \right\rangle = \frac{\tau}{f(\mathbf{x})} \left(\frac{\partial f(\mathbf{x})}{\partial x_1} \right)^2 = \frac{\tau x_1^2}{\sigma^4} f(\mathbf{x}), \tag{102}$$

where the first equality is a definition, the second equality follows from the Poisson statistics Eq. (101), and the last equality follows from the Gaussian tuning Eq. (100).

First consider the simple case where all cells have identical peak firing rate and tuning width. Now the total Fisher information with respect to x_1 for all N cells should be

$$J_N^{[1]} = \sum_{i=1}^{N} J^{[1]}(\mathbf{x} - \mathbf{c}_i) \approx \eta \underbrace{\int \cdots \int}_{D \text{ times}} J^{[1]}(\mathbf{x})\, dx_1 \cdots dx_D, \tag{103}$$

where $\mathbf{c}_i$ is the center of the Gaussian tuning function for cell i, and η is the number of cells per unit volume in D-dimensional space [see Eq. (62)]. The sum is replaced by the integral under the assumption that the number of cells N is large and the centers $\mathbf{c}_1, \ldots, \mathbf{c}_N$ are uniformly distributed in space. Because of the symmetry with respect to different dimensions, we can use the same argument in Eq. (99) to obtain the Fisher information for all N cells with all dimensions included:

$$J_N = D^{-1} J_N^{[1]} = (2\pi)^{D/2} D^{-1} \eta \tau f_{\max} \sigma^{D-2}. \tag{104}$$

If different cells have different peak firing rate $f_{\max}$ and tuning width σ, one can easily see that Eq. (104) should be modified to

$$J_N = (2\pi)^{D/2} D^{-1} \eta \tau \left\langle f_{\max} \sigma^{D-2} \right\rangle, \tag{105}$$

where the average $\langle\ \rangle$ is with respect to the populations of cells with different peak firing rate and tuning width. Assuming independence of the tuning parameters, we have

$$\left\langle f_{\max} \sigma^{D-2} \right\rangle = \langle f_{\max} \rangle \left\langle \sigma^{D-2} \right\rangle. \tag{106}$$

Finally,

$$\text{Minimal Error} \approx \mathcal{F}_D / \sqrt{J_N}, \tag{107}$$

where the correction factor $\mathcal{F}_D$ is given by Eq. (117). The result is Eq. (58). In addition, it is easy to derive Eq. (61), which immediately leads to the equivalent Eq. (57).

Estimating tuning width

To evaluate the Fisher information, the width σ in the Gaussian tuning model needs to be estimated. Given the spatial distribution of the average firing rate $f(x, y)$ measured experimentally for a single cell, we want to estimate the tuning width σ assuming a Gaussian model:

$$f(x, y) = f_{\max} \exp\left(-\frac{(x - x_0)^2 + (y - y_0)^2}{2\sigma^2} \right). \tag{108}$$

There are several methods for estimating σ. The first method uses entropy or Shannon information:

$$\sigma^2 = -\frac{1}{2\pi} \left\langle \frac{f}{f_{\max}} \ln \frac{f}{f_{\max}} \right\rangle. \tag{109}$$

This method requires an estimate of the peak rate $f_{\max}$, which is sensitive to noise and smoothing.

The second method uses the variance:

$$\sigma^2 = \frac{1}{4\pi} \frac{\langle f \rangle^2}{\langle f^2 \rangle}. \tag{110}$$

Table 3: Correction factor for estimating minimal reconstruction error

Spatial Dimension D	Correction Factor $\mathcal{F}_D$	Error Probability Density $\rho_D(r)$
$D=1$	$\mathcal{F}_1=\sqrt{\frac{2}{\pi}}\approx 0.7979$	$\rho_1(r)=\sqrt{\frac{2}{\pi}}\frac{1}{s}\exp\left(-\frac{r^2}{2s^2}\right)$ (half Gaussian)
$D=2$	$\mathcal{F}_2=\frac{\sqrt{\pi}}{2}\approx 0.8862$	$\rho_2(r)=\frac{r}{s^2}\exp\left(-\frac{r^2}{2s^2}\right)$ (Rayleigh)
$D=3$	$\mathcal{F}_3=\sqrt{\frac{8}{3\pi}}\approx 0.9213$	$\rho_3(r)=\sqrt{\frac{2}{\pi}}\frac{r^2}{s^3}\exp\left(-\frac{r^2}{2s^2}\right)$ (Maxwell)

In the above, all averages $\langle\ \rangle$ should be interpreted as follows. Suppose each unit square in the grid has the area h^2, and $f(i,j)$ is the value of average firing rate at grid point (i,j), then $\langle f\rangle = h^2\sum_{i,j} f(i,j)$, and so on. Equations (109) and (110) are closely related to the information content measure and sparsity measure considered by Skaggs, McNaughton, Wilson & Barnes (1996). One difference is that here the probability for spatial occupancy should not be included.

The third method exploits the linear relationship between the area A of a place field above a given cutoff firing rate f_{cutoff} and the logarithm of f_{cutoff} (Muller, Kubie & Ranck, 1987). This linearity is consistent with the Gaussian model in Eq. (108), which gives

$$A = 4\pi\sigma^2(\ln f_{\max} - \ln f_{\text{cutoff}})\,. \tag{111}$$

Because A depends linearly on $\ln f_{\text{cutoff}}$, both σ and $f_{\max}$ can be obtained by linear regression.

In our dataset, the movement of the animal was restricted on a narrow track. Simply assuming that each observed place field is a circular symmetric two-dimensional Gaussian cut by the track, we can estimate the width σ of the uncut Gaussian by a modified variance method. First, concatenate the pixels in the two-dimensional distribution $f(i,j)$ to obtain a one-dimensional distribution $f(i)$, and then assume a one-dimensional Gaussian distribution to obtain the one-dimensional tuning width

$$\sigma_1 = \frac{1}{2\sqrt{\pi}}\frac{\langle f\rangle^2}{\langle f^2\rangle}. \tag{112}$$

Here $\langle f\rangle = h\sum_i f(i)$, and so on, with h being the interval in the one-dimensional grid. The two-dimensional tuning width σ should equal approximately to σ_1 divided by the number of grid points spanned by the width of the track. This is because when a two-dimensional Gaussian is cut by an arbitrary straight line, the resultant one-dimensional distribution is a Gaussian of the same tuning width. The above procedure is valid if different one-dimensional slices cut in parallel with the track have approximately the same peak rate.

Distribution of reconstruction error and correction factor

A correction factor is needed because the Cramér-Rao bound only gives an estimate of the variance of unbiased reconstruction, that is, the average square of the reconstruction error $\langle r^2\rangle$. Its square root $\sqrt{\langle r^2\rangle}$ differs slightly from the direct average of reconstruction error $\langle r\rangle$. The task of this section is to determine the ratio $\mathcal{F}_D = \langle r\rangle/\sqrt{\langle r^2\rangle}$ analytically.

For reconstruction in D-dimensional space, it is natural to assume that the distribution of the

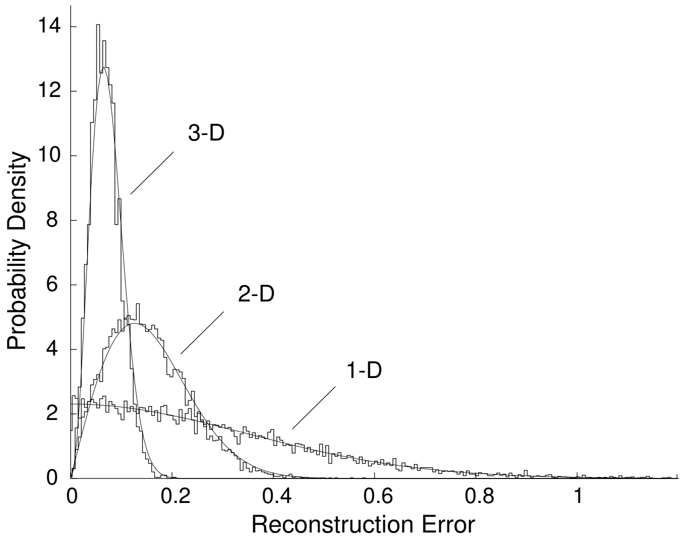

Figure 16: Distributions of reconstruction errors in one-, two-, and three-dimensional spaces. The continuous curves are the theoretical distributions in Table 3, which agree nicely with the histograms obtained by simulations of Bayesian reconstructions repeated for 5000 to 10000 trials, assuming Gaussian tuning function and Poisson spike statistics. The tuning width was $\sigma = 3$. All other parameters were identical to those in Figure 10.

error vector obeys a D-dimensional Gaussian distribution:

$$\int \left(\sqrt{2\pi}s\right)^{-D} \exp\left(-\frac{r^2}{2s^2}\right) dV = \int_0^\infty \rho_D(r)\, dr = 1, \tag{113}$$

where $r = \sqrt{x_1^2 + \cdots + x_D^2}$ is the error of reconstruction in D-dimensional space, with x_i being the error in the i-th dimension. Here

$$dV = dx_1 \cdots dx_D = \left[\frac{2\pi^{D/2}}{\Gamma(D/2)}\right] r^{D-1}\, dr \tag{114}$$

is the volume element in the D-dimensional space, and the factor in the brackets is the surface area of the $(D-1)$-dimensional unit sphere, and Γ is Euler's gamma function.

The probability density $\rho_D(r)$ for reconstruction error can thus be obtained from Eqs. (113) and (114):

$$\rho_D(r) = \frac{r^{D-1}}{2^{D/2-1}\Gamma(D/2)s^D} \exp\left(-\frac{r^2}{2s^2}\right), \tag{115}$$

from which it is possible to compute the moment

$$\langle r^m \rangle = \int_0^\infty r^m \rho_D(r)\, dr = 2^{m/2} \frac{\Gamma\left((D+m)/2\right)}{\Gamma\left(D/2\right)} s^m. \tag{116}$$

For example, when $m = 2$ the formula has a particularly simple form $\langle r^2 \rangle = Ds^2$, and $(r/s)^2$ obeys the chi square distribution with D degrees of freedom.

The correction factor, defined as $\langle r \rangle / \sqrt{\langle r^2 \rangle}$, is thus:

$$\mathcal{F}_D \equiv \frac{\langle r \rangle}{\sqrt{\langle r^2 \rangle}} = \sqrt{\frac{2}{D}} \frac{\Gamma\left((D+1)/2\right)}{\Gamma\left(D/2\right)}. \tag{117}$$

Note that $\mathcal{F}_D$ depends only on the spatial dimension D. The the width parameter s for the error distribution has disappeared. Although $\mathcal{F}_D < 1$ for all dimension D, it approaches 1 monotonically as $D \to \infty$.

The results for one-, two- and three-dimensional spaces are listed in Table 3. The theoretical curves can accurately describe the empirical error distribution of Bayesian reconstruction on synthetic data (Figure 16). This result is consistent with the assumption of Gaussian error distribution in the original D-dimensional space. In fact, for a locally uniform prior probability distribution, Bayesian reconstruction method is equivalent to maximum likelihood estimation, the error distribution of which is known to be asymptotically Gaussian in the limit of large sample (Cramér, 1946).

Appendix C. Optimality of Bayesian Reconstruction

As shown in Figures 10 and 12, Bayesian reconstruction in Eq. (36) can achieve the best possible accuracy in the sense of mean square error defined by Fisher information and the Cramér-Rao bound, provided that the conditions of Poisson spike statistics and independence of different cells are satisfied. In such a situation, the optimality of Bayesian reconstruction method can be verified directly as follows. This analysis can provide an explicit estimate of the reconstruction error in each individual trial. In general, for a uniform prior probability distribution, Bayesian reconstruction method is equivalent to maximum likelihood estimation, which can reach the Cramér-Rao bound in the limit of large sample (Cramér, 1946).

Consider the Fisher information for a one-dimensional reconstruction under a uniform prior probability:

$$J = \tau \sum_i \frac{f_i'(x)^2}{f_i(x)}, \tag{118}$$

where $f_i(x)$ is the tuning curve for cell i and τ is the time window (see Appendix B). Let Δx be the reconstruction error of the Bayesian method, which seeks the position x that maximizes

$$L = \sum_i \left[\frac{n_i}{\tau} \ln f_i(x) - f_i(x) \right], \tag{119}$$

or satisfies

$$\frac{\partial L}{\partial x} = \sum_i \left(\frac{n_i}{\tau} \frac{f_i'(x)}{f_i(x)} - f_i'(x) \right) = 0. \tag{120}$$

Cramér-Rao inequality requires that the average square error $\langle \Delta x^2 \rangle \geq 1/J$. Here we want to show that

$$\left\langle \Delta x^2 \right\rangle \approx 1/J \tag{121}$$

when a large number of cells is used. That is, the Cramér-Rao lower bound is achieved.

Consider how fluctuation in the instantaneous firing rate n_i/τ affects the estimated x since without the fluctuation the reconstruction would always be perfect. By definition, the average of n_i/τ is $f_i(x)$. Calling the difference ϵ_i, we write

$$n_i/\tau = f_i(x) + \epsilon_i. \tag{122}$$

Replace x by $x + \Delta x$ in Eq. (120) and then make Taylor expansion, keeping only the first-order term of Δx. After straightforward algebra we find

$$\Delta x \approx \frac{\tau}{J} \sum_i \epsilon_i \frac{f_i'(x)}{f_i(x)}. \tag{123}$$

This shows how the estimated position depends on the firing rate fluctuation in an individual trial. The above approximation procedure is valid if a large number of cells is used so that the error is small. Since by definition $\langle \epsilon_i \rangle = 0$, it follows from Eq. (123) that the estimated position must be unbiased; that is $\langle \Delta x \rangle = 0$. Squaring Eq. (123) and averaging yield

$$\left\langle \Delta x^2 \right\rangle \approx \frac{\tau^2}{J^2} \sum_i \left\langle \epsilon_i^2 \right\rangle \frac{f_i'(x)^2}{f_i(x)^2}, \tag{124}$$

where all cross terms $\langle \epsilon_i \epsilon_j \rangle$ with $i \neq j$ have vanished because of the independence of different cells. Using the definition Eq. (118) and Poisson spike statistics $\langle \epsilon_i^2 \rangle = f_i(x)/\tau$, we obtain Eq. (121) from Eq. (124) as desired.

Appendix D. Error of Population Vector under Cosine Tuning

Two-dimensional work space

Consider N cells with cosine tuning functions and preferred directions scattered randomly around the circle with a uniform distribution. Let the spikes of different cells occur independently of each other and obey a Poisson distribution. The task is to estimate the average angular error of the population vector.

Note that by *average over trials*, denoted by $\langle\ \rangle_{\mathrm{T}}$, we mean that the preferred directions of all cells are fixed, and the average is performed over repetitive trials where the only randomness is the Poisson statistics of the spikes; by *average over configurations*, denoted by $\langle\ \rangle_{\mathrm{C}}$, we mean averaging over random choices of the preferred directions. The error of population vector comes from both the randomness of the spikes and the randomness of the preferred directions.

Because of the rotation symmetry, we need to consider only a single test direction along the x-axis, $\mathbf{u} = (1, 0)$. The average firing rate of cell i is

$$f_i = A\mathbf{u} \cdot \mathbf{u}_i + B = Au_{ix} + B, \tag{125}$$

where A and B are constants [cf. Eq. (64)], and vector $\mathbf{u}_i = (u_{ix}, u_{iy})$ is of unit length and represents the preferred direction of cell i. The population vector is defined as

$$\mathbf{U} = \sum_{i=1}^{N} c_i \mathbf{u}_i \tag{126}$$

with the coefficients given by

$$c_i = n_i - \tau B, \tag{127}$$

where n_i is the number of spikes for cell i within the time window τ, and the background firing rate B is subtracted to improve reconstruction accuracy (Georgopoulos, Kettner & Schwartz, 1988). The x component of the population vector is estimated as follows:

$$\begin{aligned} U_x &= \sum_{i=1}^{N} c_i u_{ix} \approx \sum_{i=1}^{N} \langle c_i \rangle_{\mathrm{T}}\, u_{ix} \\ &= \tau A \sum_{i=1}^{N} u_{ix}^2 \approx \tau A N \left\langle u_{ix}^2 \right\rangle_{\mathrm{C}} \\ &= \frac{1}{2} \tau A N, \end{aligned} \tag{128}$$

in which the two approximation steps are valid if the number of cells N is large. In the above derivation, we have evaluated several averages:

$$\langle c_i \rangle_{\mathrm{T}} = \tau A u_{ix} \tag{129}$$

follows directly from the definition Eq. (127) and the fact that $\langle n_i \rangle_{\mathrm{T}} = \tau f_i$; the equations

$$\left\langle u_{ix}^2 \right\rangle_{\mathrm{C}} = \left\langle u_{iy}^2 \right\rangle_{\mathrm{C}} = \frac{1}{2}, \quad \left\langle u_{ix}^2 u_{iy}^2 \right\rangle_{\mathrm{C}} = \frac{1}{8} \tag{130}$$

are a consequence of the uniform distribution of preferred direction around the circle.

Because the test direction is assumed to be along the x-axis, population vector $\mathbf{U}$ is also close to the x-axis. Thus, $|\mathbf{U}| \approx U_x$, and the angular error of the population vector can be approximately as

$$\Delta\theta \approx \frac{U_y}{|\mathbf{U}|} \approx \frac{U_y}{U_x}, \tag{131}$$

where U_x is a constant given by Eq. (128) and $U_y = \sum_{i=1}^{N} c_i u_{iy}$ with c_i given by Eq. (127). It follows immediately that

$$\left\langle \Delta\theta^2 \right\rangle_{\mathrm{T}} \approx \frac{\left\langle U_y^2 \right\rangle_{\mathrm{T}}}{U_x^2}. \tag{132}$$

After expanding $\left\langle U_y^2 \right\rangle_{\mathrm{T}}$, evaluate each term using

$$\langle n_i n_j \rangle_{\mathrm{T}} = \begin{cases} \langle n_i \rangle_{\mathrm{T}}^2 + \langle n_i \rangle_{\mathrm{T}}, & i = j \\ \langle n_i \rangle_{\mathrm{T}} \langle n_j \rangle_{\mathrm{T}}, & i \neq j \end{cases} \tag{133}$$

which follows from Poisson spike statistics and the independence of different cells. Next, substitute $\langle n_i \rangle_{\mathrm{T}} = \tau f_i = \tau(A u_{ix} + B)$ into the expansion and note that all terms containing odd powers such as $u_{ix} u_{iy}^2$ and $u_{ix} u_{iy} u_{jx} u_{jy}$ will vanish after averaging over configurations. Thus, we need only keep terms with even powers and obtain

$$\begin{aligned} \left\langle U_y^2 \right\rangle_{\mathrm{T}} &\approx \tau B \sum_{i=1}^{N} u_{iy}^2 + \tau^2 A^2 \sum_{i=1}^{N} u_{ix}^2 u_{iy}^2 \\ &\approx \tau B N \left\langle u_{iy}^2 \right\rangle_{\mathrm{C}} + \tau^2 A^2 N \left\langle u_{ix}^2 u_{iy}^2 \right\rangle_{\mathrm{C}}, \end{aligned} \tag{134}$$

which can be evaluated by applying Eq. (130) again. The final result of the average angular error is $\mathcal{F}_1 \sqrt{\langle \Delta\theta^2 \rangle_{\mathrm{T}}} = \mathcal{F}_1 \sqrt{Q_1/N}$ with Q_1 given by Eq. (69).

Three-dimensional work space

We need to consider only a single test direction along the x-axis: $\mathbf{u} = (1, 0, 0)$. The method is very similar to that in the case of two-dimensional work space, except that now

$$\left\langle u_{ix}^2 \right\rangle_{\mathrm{C}} = \frac{1}{3}, \quad \left\langle u_{ix}^2 u_{iy}^2 \right\rangle_{\mathrm{C}} = \frac{1}{15}, \quad \text{etc.} \tag{135}$$

In addition, the error now occurs in both y and z directions so that

$$\left\langle \Delta\theta^2 \right\rangle_{\mathrm{T}} \approx \frac{\left\langle U_y^2 + U_z^2 \right\rangle_{\mathrm{T}}}{U_x^2} = \frac{2 \left\langle U_y^2 \right\rangle_{\mathrm{T}}}{U_x^2}. \tag{136}$$

The final result of the average angular error is $\mathcal{F}_2 \sqrt{\langle \Delta\theta^2 \rangle_{\mathrm{T}}} = \mathcal{F}_2 \sqrt{Q_2/N}$ with Q_2 given by Eq. (70).

References

Abbott, L. F. (1994). Decoding neuronal firing and modelling neural networks. *Quarterly Review of Biophysics*, *27*, 291–331.

Abbott, L. F. & Blum, K. I. (1996). Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral Cortex*, *6*, 406–416.

Amaral, D. G. & Witter, M. P. (1995). Hippocampal formation. In G. Paxinos (Ed.), *The Rat Nervous System* (2nd ed.)., pages 443–493. San Diego: Academic Press.

Anderson, C. H. & Van Essen, D. C. (1994). Neurobiological computational systems. In J. M. Zurada, R. J. Marks II, & C. J. Robinson (Eds.), *Computational Intelligence: Imitating Life*, pages 213–222. New York: IEEE Press.

Baldi, P. & Heiligenberg, W. (1988). How sensory maps could enhance resolution through ordered arrangements of broadly tuned receivers. *Biological Cybernetics*, *59*, 313–318.

Barnes, C. A., Suster, M. S., Shen, J., & McNaughton, B. L. (1997). Multistability of cognitive maps in the hippocampus of old rats. *Nature*, *388*, 272–275.

Ben-Yishai, R., Bar-Or, R. L., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences USA*, *92*, 3844–3848.

Bialek, W., Rieke, F., de Ruyter van Steveninck, R. R., & Warland, D. (1991). Reading a neural code. *Science*, *252*, 1854–1857.

Blair, H. T., Lipscomb, B. W., & Sharp, P. E. (1997). Anticipatory time intervals of head-direction cells in the anterior thalamus of the rat: Implications for path integration in the head-direction circuit. *Journal of Neurophysiology*, *78*, 145–159.

Blair, H. T. & Sharp, P. E. (1995). Anticipatory head-direction cells in anterior thalamus: Evidence for a thalamocortical circuit that integrates angular head velocity to compute head direction. *Journal of Neuroscience*, *15*, 6260–6270.

Blum, K. I. & Abbott, L. F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, *8*, 85–93.

Brown, E. N., Frank, L., & Wilson, M. (1996). Statistical approaches to place field estimation and neuronal population decoding. *Society for Neuroscience Abstracts*, *22*, 910.

Burgess, N., Recce, M., & O'Keefe, J. (1994). A model of hippocampal function. *Neural Networks*, *7*, 1065–1081.

Buzsáki, G. (1986). Hippocampal sharp waves: Their origin and significance. *Brain Research*, *398*, 242–252.

Buzsáki, G. (1989). Two-stage model of memory trace formation: A role for "noisy" brain states. *Neuroscience*, *31*, 551–570.

Cover, T. M. & Thomas, J. A. (1991). *Elements of Information Theory*. New York: Wiley.

Cramér, H. (1946). *Mathematical Methods of Statistics*. Princeton, New Jersey: Princeton University Press.

Eichenbaum, H., Wiener, S. I., Shapiro, M. L., & Cohen, N. J. (1989). The organization of spatial coding in the hippocampus: A study of neural ensemble activity. *Journal of Neuroscience*, *9*, 2764–2775.

Fenton, A. A. & Muller, R. U. (1997). Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proceedings of the National Academy of Sciences USA*. In press.

Földiák, P. (1993). The 'ideal humunculus': Statistical inference from neural population responses. In F. Eeckman & J. Bower (Eds.), *Computation and Neural Systems 1992*. Norwell, Massachusetts: Kluwer Academic Publishers.

Georgopoulos, A. P., Kettner, R. E., & Schwartz, A. B. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience*, *8*, 2928–2937.

Georgopoulos, A. P., Lurito, J. T., Petrides, M., Schwartz, A. B., & Massey, J. T. (1989). Mental rotation of the neuronal population vector. *Science*, *243*, 234–236.

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*, 1416–1419.

Gerrard, J. L., Ginzburg, I., Barnes, C. A., Sejnowski, T. J., & McNaughton, B. L. (1995). Accuracy of hippocampal place fields in predicting location is enhanced in the presence of theta rhythm. *Society for Neuroscience Abstracts*, *21*, 942.

Gerstner, W., Kempter, R., van Hemmen, J. L., & Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature*, *383*, 76–78.

Golub, G. H. & Van Loan, C. F. (1996). *Matrix Computations* (3rd ed.). Baltimore: The Johns Hopkins University Press.

Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel Distributed Processing*, volume 1, pages 77–109. Cambridge, Massachusetts: MIT Press.

Jensen, O. & Lisman, J. E. (1996). Hippocampal CA3 region predicts memory sequences: Accounting for the phase precession of place cells. *Learning and Memory*, *3*, 279–287.

Jung, M. W., Wiener, S. I., & McNaughton, B. L. (1994). Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *Journal of Neuroscience*, *14*, 7347–7356.

Kay, S. M. (1993). *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, New Jersey: Prentice Hall.

Kubie, J. L., Fox, S. E., & Muller, R. U. (1984). Variations in place cell firing with the state of the hippocampal EEG. *Society for Neuroscience Abstracts*, *10*, 599.

Lehky, S. R. & Sejnowski, T. J. (1990). Neural model of stereoacuity and depth interpolation based on a distributed representation of stereo disparity. *Journal of Neuroscience*, *10*, 2281–2299.

Lukashin, A. V., Amirikian, B. R., Mozhaev, V. L., Wilcox, G. L., & Georgopoulos, A. P. (1996). Modeling motor cortical operations by an attractor network of stochastic neurons. *Biological Cybernetics*, *74*, 255–261.

McHugh, T. J., Blum, K. I., Tsien, J. Z., Tonegawa, S., & Wilson, M. A. (1996). Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. *Cell*, *87*, 1339–1349.

McNaughton, B. L., Barnes, C. A., Gerrard, J. L., Gothard, K., Jung, M. W., Knierim, J. J., Kudrimoti, H., Qin, Y., Skaggs, W. E., Suster, M., & Weaver, K. L. (1996). Deciphering the hippocampal polyglot: The hippocampus as a path integrating system. *Journal of Experimental Biology*, *199*, 173–185.

McNaughton, B. L., Barnes, C. A., & O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, *52*, 41–49.

Mehta, M. R., Barnes, C. A., & McNaughton, B. L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences USA*, *94*, 8918–8921.

Muller, R. U., Bostock, E., Taube, J. S., & Kubie, J. L. (1994). On the directional firing properties of hippocampal place cells. *Journal of Neuroscience*, *14*, 7235–7251.

Muller, R. U. & Kubie, J. L. (1989). The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience*, *9*, 4101–4110.

Muller, R. U., Kubie, J. L., & Ranck, Jr., J. B. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *Journal of Neuroscience*, *7*, 1935–1950.

Nowlan, S. & Sejnowski, T. J. (1995). A selection model for motion processing in area MT of primates. *Journal of Neuroscience*, *15*, 1195–1214.

O'Keefe, J. & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, *381*, 425–428.

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. & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, *3*, 317–330.

Olton, D. S., Branch, M., & Best, P. J. (1978). Spatial correlates of hippocampal unit activity. *Experimental Neurology*, *58*, 387–409.

Ono, T., Nakamura, K., Nishijo, H., & Eifuku, S. (1993). Monkey hippocampal neurons related to spatial and nonspatial functions. *Journal of Neurophysiology*, *70*, 1516–1529.

Optican, L. M. & Richmond, B. J. (1987). Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *Journal of Neurophysiology*, *57*, 162–178.

Paradiso, M. A. (1988). A theory for the use of visual orientation information which exploits the columnar structure of striate cortex. *Biological Cybernetics*, *58*, 35–49.

Pellionisz, A. & Llinas, R. (1985). Tensor network theory of the metaorganization of functional geometries in the central nervous system. *Neuroscience*, *16*, 245–273.

Pellionisz, A. J. (1984). Coordination: A vector-matrix description of transformation of overcomplete CNS coordinates and a tensorial solution using the Moore-Penrose generalized inverse. *Journal of Theoretical Biology*, *110*, 353–375.

Poggio, T. (1990). A theory of how the brain might work. In *Cold Spring Harbor Symposia on Quantitative Biology, Volume LV, The Brain*, pages 899–910. Plainview, New York: Cold Spring Harbor Laboratory Press.

Poggio, T. & Girosi, F. (1990). Regularization algorithms for learning that are equivalent to multilayer networks. *Science*, *247*, 978–982.

Poggio, T., Torre, V., & Koch, C. (1985). Computational vision and regularization theory. *Nature*, *317*, 314–319.

Pouget, A. & Sejnowski, T. J. (1994). A neural model of the cortical representation of egocentric distance. *Cerebral Cortex*, *4*, 314–329.

Pouget, A. & Sejnowski, T. J. (1997). Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience*, *9*, 222–237.

Pouget, A. & Zhang, K.-C. (1997). Statistically efficient estimations using cortical lateral connections. In M. C. Mozer, M. I. Jordan, & T. Petsche (Eds.), *Advances in Neural Information Processing Systems*, volume 9, pages 97–103. Cambridge, Massachusetts: The MIT Press.

Ranck, Jr., J. B. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats: Part I. Behavioral correlates and firing repertoires. *Experimental Neurology*, *41*, 461–555.

Rao, C. R. (1965). *Linear Statistical Inference and Its Applications.* New York: Wiley.

Redish, A. D. & Touretzky, D. S. (1996). Modeling interactions of the rat's place and head direction system. In D. S. Touretzky, M. C. Mozer, & M. E. Hasselmo (Eds.), *Advances in Neural Information Processing Systems*, volume 8, pages 61–67. Cambridge, Massachusetts: The MIT Press.

Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1997). *Spikes: Exploring the Neural Code.* Cambridge, Massachusetts: MIT Press.

Rolls, E. T., Robertson, R. G., & Georges-Francois, P. (1995). The representation of space in the primate hippocampus. *Society for Neuroscience Abstracts*, *21*, 1494.

Rotenberg, A., Mayford, M., Hawkins, R. D., Kandel, E. R., & Muller, R. U. (1996). Mice expressing activated CaMKII lack low frequency LTP and do not form stable place cells in the CA1 region of the hippocampus. *Cell*, *87*, 1351–1361.

Salinas, E. & Abbott, L. F. (1994). Vector reconstruction from firing rates. *Journal of Computational Neuroscience*, *1*, 89–107.

Samsonovich, A. & McNaughton, B. L. (1996). Attractor-map-based path integrator model of the hippocampus reproduces the phase precession phenmenon. *Society for Neuroscience Abstracts*, *22*, 1872.

Samsonovich, A. & McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, *17*, 5900–5920.

Sanger, T. D. (1994). Theoretical considerations for the analysis of population coding in motor cortex. *Neural Computation*, *6*, 12–21.

Sanger, T. D. (1996). Probability density estimation for the interpretation of neural population codes. *Journal of Neurophysiology*, *76*, 2790–2793.

Scharf, L. L. (1991). *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis.* Reading, Massachusetts: Addison-Wesley.

Schwartz, A. B. (1993). Motor cortical activity during drawing movements: Population representation during sinusoid tracing. *Journal of Neurophysiology*, *70*, 28–36.

Schwartz, A. B. (1994). Direct cortical representation of drawing. *Science*, *265*, 540–542.

Seung, H. S. & Sompolinsky, H. (1993). Simple models for reading neuronal population codes. *Proceedings of the National Academy of Sciences USA*, *90*, 10749–10753.

Skaggs, W. E. & McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, *271*, 1870–1873.

Skaggs, W. E., McNaughton, B. L., Gothard, K. M., & Markus, E. J. (1993). An information-theoretic approach to deciphering the hippocampal code. In S. J. Hanson, J. D. Cowan, & C. L. Giles (Eds.), *Advances in Neural Information Processing Systems*, volume 5, pages 1030–1037. Cambridge, Massachusetts: The MIT Press.

Skaggs, W. E., McNaughton, B. L., Wilson, M. A., & Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, *6*, 149–172.

Snippe, H. P. (1996). Parameter extraction from population codes: A critical assessment. *Neural Computation*, *8*, 511–529.

Somers, D. C., Nelson, S. B., & Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, *15*, 5448–5465.

Stewart, M. & Fox, S. E. (1990). Do septal neurons pace the hippocampal theta rhythm? *Trends in Neurosciences*, *13*, 163–168.

Taube, J. S. & Muller, R. U. (1995). Head direction cell activity in the anterior thalamic nuclei, but not the postsubiculum, predicts the animal's future directional heading. *Society for Neuroscience Abstracts*, *21*, 946.

Taube, J. S., Muller, R. U., & Ranck, Jr., J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, *10*, 420–435.

Tikhonov, A. N. & Arsenin, V. Y. (1977). *Solutions of Ill-Posed Problems.* New York: Wiley.

Treves, A., Barnes, C. A., & Rolls, E. T. (1996). Quantitative analysis of network models and of hippocampus data. In T. Ono, B. L. McNaughton, S. Molotchinkoff, E. T. Rolls, & H. Nishijo (Eds.), *Perception, Memory, and Emotion: Frontiers in Neuroscience,* pages 567–579. Oxford: Pergamon.

Tsodyks, M. V. & Sejnowski, T. J. (1995a). Associative memory and hippocampal place cells. *International Journal of Neural Systems, 6 (suppl.),* 81–86.

Tsodyks, M. V. & Sejnowski, T. J. (1995b). Rapid state switching in balanced cortical network model. *Network, 6,* 111–124.

Tsodyks, M. V., Skaggs, W. E., Sejnowski, T. J., & McNaughton, B. L. (1996). Population dynamics and theta rhythm phase precession of hippocampal place cell firing: A spiking neuron model. *Hippocampus, 6,* 271–180.

Tuckwell, H. C. (1988). *Introduction to Theoretical Neurobiology, Volume 2: Nonlinear and Stochastic Theories.* Cambridge, England: Cambridge University Press.

Ullman, S. & Basri, R. (1991). Recognition by linear combinations of models. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 13,* 992–1006.

Vanderwolf, C. H., Kramis, R., Gillespie, L. A., & Bland, B. H. (1975). Hippocampal rhythmical slow activity and neocortical low voltage fast activity: Relations to behavior. In R. L. Isaacson & K. H. Pribram (Eds.), *The Hippocampus, Volume 2: Neurophysiology and Behavior,* pages 101–128. New York: Plenum Press.

Wilson, M. A. & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science, 261,* 1055–1058. Corrections in vol. 264, p. 16.

Wilson, M. A. & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science, 265,* 676–679.

Zacks, S. (1971). *The Theory of Statistical Inference.* New York: Wiley.

Zemel, R. S., Dayan, P., & Pouget, A. (1997). Probabilistic interpretation of population codes. In M. C. Mozer, M. I. Jordan, & T. Petsche (Eds.), *Advances in Neural Information Processing Systems,* volume 9, pages 676–682. Cambridge, Massachusetts: The MIT Press.

Zhang, K.-C. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience, 16,* 2112–2126.

Zipser, D. (1985). A computational model of hippocampal place fields. *Behavioral Neuroscience, 99,* 1006–1018.

Zipser, D. & Andersen, R. A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature, 331,* 679–684.

Zohary, E. (1992). Population coding of visual stimuli by cortical neurons tuned to more than one dimension. *Biological Cybernetics, 66,* 265–272.

Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature, 370,* 140–143.