

Spikes not slots: noise in neural populations limits working memory

Paul M. Bays

UCL Institute of Neurology, Queen Square, London, WC1N 3BG, UK

This opinion article argues that noise (randomness) in neural activity is the limiting factor in visual working memory (WM), determining how accurately we can maintain stable internal representations of external stimuli. Sharing of a fixed amount of neural activity between items in memory explains why WM can be successfully described as a continuous resource. This contrasts with the popular conception of WM as comprising a limited number of memory slots, each holding a representation of one stimulus – I argue that this view is challenged by computational theory and the latest neurophysiological evidence.

Deterministic and stochastic views of WM

WM refers to the ability of the nervous system to actively internally maintain information over brief intervals [1–4]. It is considered an essential component of most complex behaviours and is closely linked to general intelligence. Critically, WM is strongly limited in its ability to hold multiple representations simultaneously, constraining the complexity of mental operations. In vision, two classes of explanation have dominated attempts to account for this limit.

The first class proposes that the limit is due to a deterministic mechanism: there is a fixed maximum number of representations that can be held in memory at one time; once this limit is reached, further items cannot enter memory [5–7]. These models draw support from theoretical studies of synchronised neural activity [8,9] and imaging studies claiming to show plateaus in the activation function at a particular number [10,11] (see [3] for a critical review). The classic deterministic account of visual WM is the ‘slot’ model, which posits three or four independent memory slots, each holding a detailed representation of one visual object [6].

The second class asserts that the limit has a stochastic mechanism: representations in memory become increasingly variable as their number increases, until they can no longer be distinguished from random noise [12–19]. These models find support in the inherent variability of neural activity [20,21], and evidence from single-neuron recording and fMRI decoding studies [22–24]. In vision, the stochastic (see [Glossary](#)) view is exemplified by the ‘resource’

model, which proposes that a limited supply of a representational medium is continuously distributed between visual objects; items that receive more resource are stored with less noise [16–19]. Several ‘hybrid’ models have also been suggested, combining elements of both slot and resource accounts [25–27].

Crucial evidence favouring stochastic models has come from examining the distribution of errors in analogue recall tasks. A typical task, testing memory for orientation, and the observed distributions of error are shown in [Figure 1A,B](#) (data from [28]; similar distributions have been obtained for other visual features, e.g., colour and motion direction [15,16,18,29]). The first important observation is that variability (the width of the distribution) increases steadily with the number of items in the memory array ([Figure 1B,C](#)). The relationship between variance and set size is well described by a power law (appearing

Glossary

Analogue recall task: an experimental task in which an observer is required to select the value of a remembered stimulus feature from a continuous space of possibilities.

Delay period: the interval in a WM task between presentation and test, during which items must be maintained in memory.

Feature tuning: tendency of a visual neuron to fire maximally in response to stimuli with a particular visual feature, for example, a particular orientation. The full relationship between firing rate and feature value is described by tuning function of the neuron.

Gamma oscillation: high-frequency (30–100 Hz) rhythmic neural activity.

Kurtosis: a measure of the shape of a probability distribution, specifically ‘peakedness’ of the distribution in comparison to the normal distribution.

Lateral inhibition: the ability of an excited neuron to reduce the activation of other functionally similar neurons within its local environment.

Misbinding: an incorrect recombination of visual features belonging to two or more different objects. For example, recalling a red tractor when you actually saw a red car and a green tractor.

Neural oscillation: large-scale rhythmic activity in the nervous system, resulting from synchronisation of neural firing patterns.

Normal distribution: a bell-shaped continuous probability distribution. Sums or averages of random variables drawn from independent distributions are typically normally distributed. The normal distribution in a circular space is called the von Mises distribution.

Normalisation: rescaling a set of values to have a particular common magnitude.

Poisson process: a process that generates events independently at random with a particular mean rate. The spiking activity of a neuron can be approximated by a Poisson process.

Power law: a relationship in which one quantity varies as a power of another, $y = ax^n$. A power law with exponent $n = 1$ describes a linear relationship.

Set size: the number of distinct items presented for memorisation in a WM task.

Signal-to-noise: the ratio of meaningful signal to background noise in a source of information.

Spatial selectivity: tendency of a neuron to fire only in response to stimuli within a particular region of space (its receptive field).

Stochastic: partly or wholly unpredictable/random.

Theta oscillation: low-frequency (4–7 Hz) rhythmic neural activity.

Corresponding author: Bays, P.M. (p.bays@ucl.ac.uk).

Keywords: working memory; population code; normalisation; attractor network; diffusion to bound; binding.

1364-6613/

© 2015 Elsevier Ltd. All rights reserved. <http://dx.doi.org/10.1016/j.tics.2015.06.004>

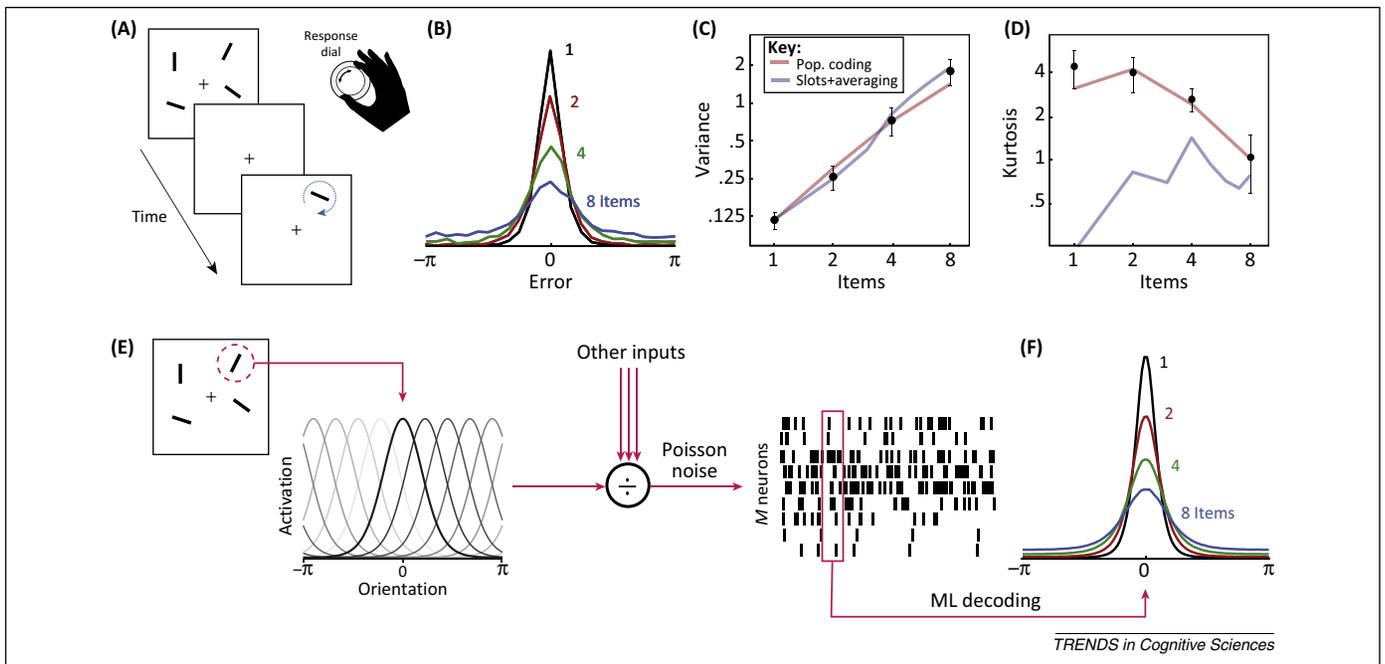


Figure 1. Errors in working memory (WM) reflect noise in neural coding. **(A)** An analogue report task, in which participants adjust a probe stimulus to match the orientation recalled at the same location in the memory array. **(B)** The distribution of responses around the true orientation changes as a function of the number of items in the memory array. **(C)** Error variability (black) increases monotonically with the number of array items. Curves show predictions of the population coding model (red; [28]) and slots + averaging model (blue; [27]) with best-fitting parameters. **(D)** Kurtosis (black) measures deviations from normality of error distributions. Curves as in (C). **(E)** The population coding model. Each stimulus is encoded by a set of orientation-selective neurons with bell-shaped (normal) tuning functions. Normalisation operates across the whole population, scaling summed activity to a fixed level. Neurons generate spikes according to a noisy Poisson process, and recall is modelled as maximum likelihood (ML) decoding of the spiking activity over a fixed time window. **(F)** Error distributions predicted by the population coding model with ML parameters (compare with B). Adapted and reprinted from [28], with permission from the Society for Neuroscience.

linear on log–log axes as in Figure 1C) and there is no evidence for any abrupt discontinuity as might be expected on reaching a deterministic limit. This observation is at the heart of stochastic models of WM and appears incompatible with the classic deterministic view, which predicts no changes in variability so long as all items are stored, that is, until the fixed limit is exceeded.

The second key observation is that WM error distributions deviate significantly from the familiar normal distribution; in particular, the distributions have excess kurtosis (Figure 1D) – they are strongly peaked with long tails. This observation is important because mathematical models in psychology and neuroscience typically assume that internal variables have noise that is normally distributed. On this assumption, the observed deviations from normality have been interpreted variously as evidence for random ‘guessing’ [27] or variability in the precision of memory representations [18,30].

A different perspective on these observations is obtained by considering how visual features are represented in the nervous system. Simple visual features, such as orientation, are encoded in neural activity by population codes (Box 1). A recent study [28] has shown that stimulus estimates obtained by optimal decoding of a noisy population code have non-normal distributions of the type observed in WM experiments (Figure 1E,F). According to this model, the stimulus features in the memory array are encoded in the firing rates of a population of neurons with spatial selectivity and feature tuning. The population activity is normalised, meaning that the total activity level (sum of firing rates) remains constant across changes in set size. The activity pattern persists once the memory array is

removed, with neurons firing stochastically according to a Poisson process.

Recall of a probed item is modelled as maximum likelihood decoding (Box 1) of the persistent spiking activity over a fixed time window. Because of the noise in spiking activity, this estimate exhibits variability around the true value; the fewer spikes available to decode, the greater the variability. Because the population as a whole is normalised, the activity coding for each memory item declines with increasing set size. This simple model accurately reproduces the power law relationship between variance and set size observed in experiments (red line in Figure 1C). Furthermore, estimates decoded from a tuned population are not in general normally distributed, and the specific deviations from normality expected fit those observed in experiments (red line in Figure 1D).

Slots and resources

In the population coding model [28], total spiking activity is limited due to normalisation, and is distributed (shared out) between visual items; in this way, the model provides a plausible biological basis for the concept of a limited WM resource. Neural activity can be selectively allocated to stimuli with the greatest behavioural priority, and this quantitatively reproduces empirical observations of enhanced WM precision for such items, as well as corresponding costs for non-prioritised stimuli [17,31,32].

Can slot-based models also account for WM error distributions? An influential study [27] proposed a hybrid model, the ‘slots + averaging’ model that claimed to fit WM errors. The slots + averaging model modifies the classic slot model to allow multiple slots to be allocated

Download English Version:

<https://daneshyari.com/en/article/141387>

Download Persian Version:

<https://daneshyari.com/article/141387>

[Daneshyari.com](https://daneshyari.com)