

Localist representation can improve efficiency for detection and counting

Horace Barlow^a and Anthony Gardner-Medwin^b

^aPhysiological Laboratory, Cambridge CB2 3EG, England: hbb10@cam.ac.uk

^bDepartment of Physiology, University College London, London WC1E 6BT, England:

a.gardner-medwin@ucl.ac.uk

Abstract: Almost all representations have both distributed and localist aspects, depending upon what properties of the data are being considered. With noisy data, features represented in a localist way can be detected very efficiently, and in binary representations they can be counted more efficiently than those represented in a distributed way. Brains operate in noisy environments, so the localist representation of behaviourally important events is advantageous, and fits what has been found experimentally. Distributed representations require more neurons to perform as efficiently, but they do have greater versatility.

In addition to the merits Page argues for, localist representations have quantitative advantages that he does not bring out. The brain operates in an uncertain world where important signals are always liable to be contaminated and masked by unwanted ones, so it is important to consider how external noise from the environment affects the reliability and effectiveness of different forms of representation. In what follows, we shall adopt Page's definitions of localist and distributed representation, according to which almost any scheme or model has both components. In a scheme emphasising localist representation, the elements can nonetheless be used in combinations, and such combinations represent in a distributed way whatever input events cause them to occur. Similarly in a scheme emphasising distributed representation, each particular element is activated by a particular subset of the possible input patterns, and it represents this subset in a localist way; for example, a single bit in the ASCII code is a localist representation of the somewhat arbitrary collection of ASCII characters for which it is ON. We shall also assume for simplicity that the brain represents noisy data, but does not necessarily add noise; of course this is a simplification, but it is the appropriate starting point for the present problem.

Localist representations and matched filters. The principle of a matched filter is to collect *all* the signal caused by the target that is to be detected, and *only* this signal, excluding as much as possible signals caused by other stimuli. In this way the response from the target is maximised while pollution by noise from nontarget stimuli is minimised, yielding the best possible signal/noise ratio. Localist representations of features or patterns in the input data can be close approximations to matched filters. If the representation's elements are linear and use continuous variables, their outputs will be the weighted sums of their different inputs. If each weight is proportional to the ratio of signal amplitude to noise variance for that part of the input when the desired target is presented, the element will be a matched filter for that target. Some neurons in sensory areas of the cortex follow this prescription well, and it makes good sense to regard them as members of a vast array of matched filters, each with slightly different parameters for its trigger feature or optimal stimulus. In V5 or MT (an area specialising in coherent motion over small regions of the visual field) the receptive fields of the neurons differ from each other in position, size, direction, and velocity of their preferred motion (Felleman & Kaass 1984; Maunsell & Van Essen 1983a; Raiguel et al. 1995), and preferred depth or disparity of the stimulus (DeAngelis et al. 1998; Maunsell & Van Essen 1983b). It has been shown that many individual neurons can detect coherent motion with as great sensitivity as the entire conscious monkey (Britten et al. 1992; Newsome et al. 1989). Furthermore, human performance in similar tasks varies with stimulus parameters (area, duration, dot density, etc.) as if it was limited by the noise or uncertainty inherent in the stochastic stimuli that are used, so external noise appears to be an important limit (Barlow & Tripathy 1997). On re-examination it also turns out that external noise is important in monkey MT neurons (Mogi & Barlow 1998). For neurons to perform as well as they do, they must have properties close to those of optimum matched filters, which suggests that the whole visual cortex is a localist representation of the visual field using numerous different arrays of filters matched to different classes of feature. This insight may well apply to all sensory areas of the cortex and even to nonsensory parts, in which case the cortex would be a strongly localist representation throughout. Can efficient detection at higher levels always be done by the weighted combination of inputs from the elements of distributed representations at lower levels? This would require graded signals between the levels, and it is doubtful if those passing between cortical neurons have sufficient dynamic range. With binary signals, and a

task of counting occurrences rather than extracting signals from noise, there is an analogous problem of diminishing the effects of overlap in distributed representations.

Counting accuracy in localist and distributed representation. For many of the computations that are important in the brain, such as learning, or detecting that two stimuli are associated, it is necessary to count or estimate how often a specific type of event has occurred. It is easy to see that, because the elements active in the distributed representation of an event that is to be counted also respond to other events, the mean response rates of those elements will be greater than the mean responses due solely to the event to be counted. The average effect of this inflation can readily be allowed for, but in a noisy environment the variance as well as the mean will be increased, and this cannot be corrected. The only way to avoid this problem completely would be to have localist representations for the counted events, though as shown elsewhere (Gardner-Medwin & Barlow, submitted), distributed representations can be efficient at counting if they employ enough elements with sufficient redundancy. It may be suggested that brains often learn from a single experience and do not need to count accurately, but such an argument would be misleading. Efficient statistics are what an animal needs in order to make correct inferences with the minimum amount of data collection, and this is more, not less, important when the number of available trials is low. A system cannot use inefficient methods of representation if one-shot learning is to occur reliably when it is appropriate and not when it is not. The relative merits of localist and distributed representations are sometimes finely balanced and are discussed in greater detail elsewhere (Gardner-Medwin & Barlow, submitted). Localist representations have the edge in terms of efficiency, but one must know in advance what needs to be detected and counted, so they are mainly appropriate for frequent, regularly recurring features of the environment. In spite of the large numbers of neurons required, the ability of distributed representations to handle unexpected and unspecified events without ambiguity makes them better for handling novel experiences.

The principle of local computation. Finally it should be pointed out that the merit of localist representations stems from the fact that computation in the brain is done by local biophysical processes. Every element of a computation requires a locus in the brain where all the necessary factors are collected together so that they can take part in the biophysical process. As an example of the relevance of this principle, consider the Hebbian assumption about the locus of learning. Biophysical processes close to a synapse can readily be influenced by both pre- and postsynaptic activity, since the required information is present there in the way that the principle requires, but it would not be reasonable to assume that distributed patterns of synchronous activity in remote neurons could work in the same way. The implied ban on “action at a distance” may eventually need qualification through better understanding of neuromodulators and dendritic interactions, but localist representations have the advantage that they already collect at one element all the information required for detection and counting; this is what makes it possible for them to perform these jobs efficiently. Page ends his manifesto by saying “if the brain does not use localist representations then evolution has missed an excellent trick.” Plenty of neurophysiological evidence shows that it has not, in fact, missed this trick that is so valuable for achieving sensitive and reliable detection of weak signals in a noisy background, and for the frequency estimations needed for reliable and efficient learning. Doubtless evolution has also exploited the advantages that distributed representation can bring to the handling of the unexpected.

This commentary was published within:

BEHAVIORAL AND BRAIN SCIENCES (2000) 23, 443-512

Connectionist modelling in psychology: A localist manifesto

Mike Page Medical Research Council Cognition and Brain Sciences Unit, Cambridge, CB2 2EF, United Kingdom

Abstract: Over the last decade, fully distributed models have become dominant in connectionist psychological modelling, whereas the virtues of localist models have been underestimated. This target article illustrates some of the benefits of localist modelling. Localist models are characterized by the presence of localist representations rather than the absence of distributed representations. A generalized localist model is proposed that exhibits many of the properties of fully distributed models. It can be applied to a number of problems that are difficult for fully distributed models, and its applicability can be extended through comparisons with a number of classic mathematical models of behaviour. There are reasons why localist models have been underused, though these often misconstrue the localist position. In particular, many conclusions about connectionist representation, based on neuroscientific observation, can be called into question. There are still some problems inherent in the application of fully distributed systems and some inadequacies in proposed solutions to these problems. In the domain of psychological modelling, localist modelling is to be preferred.