neural activity to occur. We can take Maley's summary of Zacks' review findings at face value. The problem is that we still do not know what it is about these representations' use of magnitudes to represent magnitudes that predicts the response times.

5. Deep Neural Networks: An Alternative View of Retinotopic Organization

Even if the phenomenon of retinotopic organization cannot be connected with the mental rotation and mental scanning timing results in a way that supports the thesis that mental images are analog, some may see the observation of retinotopic cortical activation as sufficient evidence by itself for the thesis the mental images are analog (as in, e.g., Nanay [2023], p. 48)). After all, why would the brain bother to generate image-like patterns of activation while representing patterns in the world if it were not in fact making use of imagistic representations?

However, there are many reasons the brain might generate image-like patterns of activation even if it were not making use of analog, imagistic representations. An especially compelling reason has been obscured by the fact that the imagery debate's participants often assume a false choice: that mental imagery must either be *iconic* or *language-like* in its representational format. These, for many of the key participants, were the only formats worth serious consideration. That is surprising, given that connectionist alternatives—which invoke neither iconic nor language-like representations—were well-known and already much-discussed in the 1980's and 1990's, when the imagery debate was in full swing. The explanation, I suspect, is in part technological and in part sociological. At the time, there were no connectionist AI systems remotely capable of the kinds of image generation tasks we associate with imagery; while, simultaneously, the most influential researchers on both sides of the debate had in common a dim view of connectionism. Pylyshyn, in particular, is famous for his independent attack on connectionism as a theory of thought (Fodor & Pylyshyn [1988]).

Things are very different now. Today's sophisticated image classification and image generation AI systems all rely upon deep neural network (DNN) connectionist architectures. There are, by contrast, no AI systems of remotely comparable capacities that make use of iconic or language-like representations (*modulo* the assumption—which cannot be fully defended here—that the DNNs in question do not implement iconic or language-like representations). Thus, the possibility that DNN-like connectionist architectures underlie human vision and mental imagery deserves very close consideration.

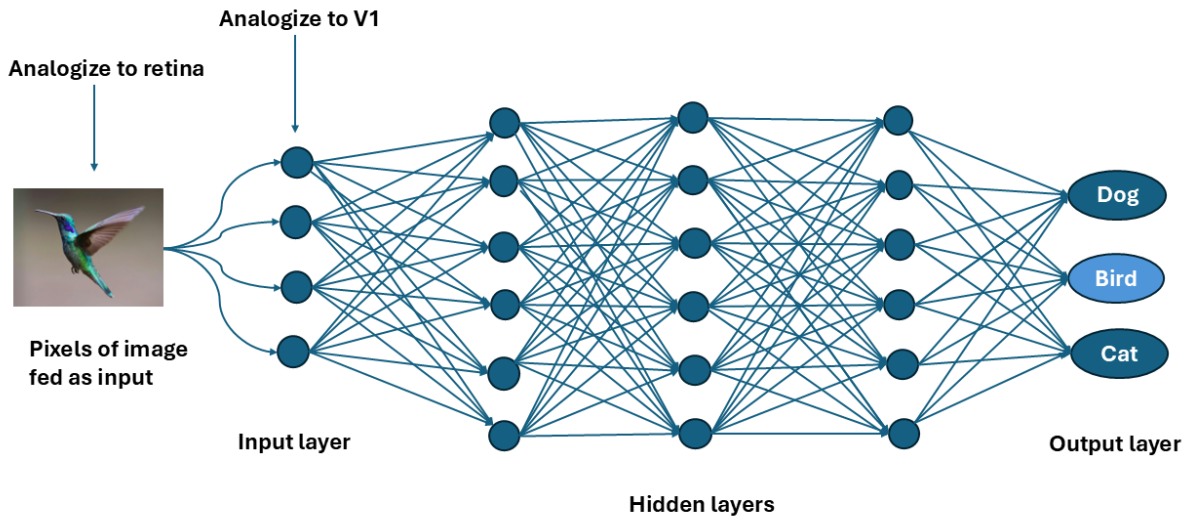

Figure 2: A representation of a convolutional deep neural network showing a characteristic lack of connection among nodes within any single layer.

A central selling point of connectionist architectures, considered as models of human cognition, is that the connections among layers of nodes in an artificial neural network bear clear structural similarities to the connections among neurons in the human brain. Accordingly, the important feature of image-processing DNNs for our purposes is that they offer an account of why we might see retinotopic activation in the brain during imagery (and visual perception) tasks, even if such activation did not constitute the tokening of iconic mental representations. Consider the feed-forward DNN shown in Figure 2, where an image of a hummingbird activates an input layer of four nodes, whose activations feed forward to nodes in three hidden layers, resulting in the activation of three nodes in the output layer. When the central node in the output layer is activated to a sufficient threshold, the network counts as indicating that a bird was in the input image. We can analogize this processing to human neural activation by letting the hummingbird image—considered as a 2D grid of pixel activations—stand for activation at the retina, and the input layer (also structured as a 2D arrangement of nodes) analogized to early visual areas that show retinotopic activation. If these analogies hold even roughly, it would be unsurprising to see structurally corresponding activations at the retina and at the input (and other early) cortical layers, simply because cells in the retina are activating subsequent neurons in early visual areas in ways corresponding to their own spatial outlay. Importantly, however,

when we conceive of what such activations are accomplishing computationally in connectionist terms, we see that there are no iconic (or depictive, or even analog) representations at work.

There are several interrelated reasons for this. First, it is only thanks to the cascade of activation across multiple layers—eventuating in the output layer's middle node activating above a threshold—that (say) a *bird* is represented. The cognitive system cannot be said to represent—by distinguishing from other stimuli—the sort of thing causing activation at the input layer until the nodes at other layers have played their role in shaping activation at the output layer. Further, the processing in connectionist networks is *parallel* in the specific sense that each node within a layer is processing inputs from each node in the prior level and sending a signal forward to each node in the subsequent layer *completely independently* of what else is going on with other nodes in the same layer. Critically, *there are no connections between nodes in the same layer* (as shown in Fig. 1). Because there are no connections among nodes within a particular layer, it makes no difference how far apart those nodes are from each other. It would not change the network's functioning at all to radically alter the space between nodes in a layer—making it entirely non-uniform, with any apparent retinotopic image disappearing—so long as each node maintained its connections to prior and subsequent nodes. In contrast, if spatial magnitudes were really being used to represent spatial magnitudes—if there really were iconic representations present within the networks—such changes should make *all the difference* to what is being represented.

In sum, the striking appearance of image-like activation in retinotopically organized cortex gives us no reason to suppose that the brain is making use of iconic representations. A brain whose processing mirrored the principles at work in our most sophisticated image processing and image generating artificial intelligence would also show such activations, without their being iconic representations. This is just one salient example of likely many for why we might see retinotopic activations in a system that is not using iconic (or analog) representations.

6. A Format-Agnostic Account of the Timing Results

What, then, explains the response times we see in the rotation studies? I end with a proposal that, while schematic, suggests that questions of format will not be at the fore. While I