

# Researchers Keep Rejecting Grandmother Cells after Running the Wrong Experiments: The Issue Is How Familiar Stimuli Are Identified

Jeffrey S. Bowers,\* Nicolas D. Martin, and Ella M. Gale

There is widespread agreement in neuroscience and psychology that the visual system identifies objects and faces based on a pattern of activation over many neurons, each neuron being involved in representing many different categories. The hypothesis that the visual system includes finely tuned neurons for specific objects or faces for the sake of identification, so-called “grandmother cells”, is widely rejected. Here it is argued that the rejection of grandmother cells is premature. Grandmother cells constitute a hypothesis of how familiar visual categories are identified, but the primary evidence against this hypothesis comes from studies that have failed to observe neurons that selectively respond to unfamiliar stimuli. These findings are reviewed and it is shown that they are irrelevant. Neuroscientists need to better understand existing models of face and object identification that include grandmother cells and then compare the selectivity of these units with single neurons responding to stimuli that can be identified.

## 1. Introduction

The hypothesis that single neurons mediate the identification of familiar visual categories (e.g., a familiar person) is often dismissed in neuroscience. Indeed, researchers tend to use the somewhat pejorative phrase “grandmother cell” to describe this form of neural coding. Instead, it is widely claimed that visual categories are identified on the basis of a pattern of activation over many neurons, each neuron being involved in representing many different categories, so-called “distributed” or “ensemble” coding. The distributed approach is thought to be supported by theoretical considerations and empirical findings. Indeed, single-cell recording studies are often taken to falsify the grandmother cell hypothesis.<sup>[1]</sup>

In this article, we detail why the widespread dismissal of grandmother cells is misguided. The core problem is that researchers tend to consider straw-man versions of the


hypothesis, and as a consequence, do not carry out experiments relevant to testing more plausible versions of the theory. We focus on one of the most common theoretical confusions that has led to a consistent failure to carry out relevant experiments, namely, the claim that grandmother cells constitute a theory of how both familiar and unfamiliar categories are perceived. Researchers then reject grandmother cells when they do not find neurons that respond selectively to unfamiliar stimuli (e.g., failing to find a neuron that selectively responds to an unfamiliar face). The problem with this, however, is that grandmother cells constitute a theory of how familiar categories are identified (e.g., how you identify your grandmother and how I recognize mine). Accordingly, the relevant question is how

neurons respond to familiar faces, objects, and words when they are identified. As we detail below, relatively few studies assess the selectivity of neurons under these conditions, but there is evidence that neural selectivity is greater for familiar than unfamiliar categories. Indeed, given the high level of selectivity that has been reported with familiar visual categories when they are identified, the grandmother cell hypothesis should be considered a serious hypothesis rather than dismissed out of hand.

This article is organized as follows: first, we describe in some detail what we mean by the term grandmother cell and then briefly describe a number of computational models of face, word, and object identification in psychology and neuroscience that implement grandmother cells. Second, we briefly review studies that assessed the response selectivity of single neurons in the monkey’s visual cortex to unfamiliar stimuli (e.g., unfamiliar faces or objects) or familiar stimuli that the monkey was never trained to identify (novel stimuli that the monkey repeatedly viewed without any need to categorize). Some of these studies are reviewed in the text; others are just noted in **Table 1**. Although these studies rarely report neurons that responded selectively to one stimulus category (e.g., a specific face), these studies are irrelevant to the grandmother cell hypothesis because grandmother cells constitute a theory of how familiar visual images are identified as members of specific categories. Third, we review several studies that do report highly selective responses in visual neurons when

---

Prof. J. S. Bowers, N. D. Martin, Dr. E. M. Gale  
School of Psychological Science  
University of Bristol  
Bristol BS8 1TU, UK  
E-mail: j.bowers@bristol.ac.uk

 The ORCID identification number(s) for the author(s) of this article can be found under <https://doi.org/10.1002/bies.201800248>.

DOI: 10.1002/bies.201800248

**Table 1.** Additional experiments that assessed response selectivity in visual cortex.

Year	Paper	Stimulus	Trained to identify image?	Conclusions regarding neural coding
2015	Dubois et al. <sup>[2]</sup>	F, O, B	No	No mention of GMC (propose DC)
2015	Meyers et al. <sup>[3]</sup>	F, O, B	No	No mention of GMC (propose DC)
2015	Taubert et al. <sup>[4]</sup>	F	No	Did not discuss coding
2012	Issa and DiCarlo <sup>[5]</sup>	F	No	Did not discuss coding
2010	Freiwald and Tsao <sup>[6]</sup>	F, O, B	No	No mention of GMC (propose DC)
2008	Meyers et al. <sup>[7]</sup>	A	No	No mention of GMC (propose DC)
2007	Franco et al. <sup>[8]</sup>	F, O, Sc	No	Did not discuss coding
2006	Tsao et al. <sup>[9]</sup>	F, O, B	No	No mention of GMC (propose DC)
2005	Freedman et al. <sup>[10]</sup>	A	No	Did not discuss coding
2001	Keyesers et al. <sup>[11]</sup>	F, O, Sc	No	No mention of GMC (propose DC)
2001	Sheinberg and Logothetis <sup>[12]</sup>	O	No	Did not discuss coding
2001	Tamura and Tanaka <sup>[13]</sup>	O, Sh	No	Did not discuss coding
1999	Sugase et al. <sup>[14]</sup>	F, Sh	No	No mention of GMC
1998	Booth and Rolls <sup>[15]</sup>	O	No	Reject GMC (propose DC)
1997	Rolls et al. <sup>[16]</sup>	F	No	Reject GMC (propose DC)
1996	Higuchi and Miyashita <sup>[17]</sup>	P	No	Did not discuss coding
1995	Rolls and Tovee <sup>[18]</sup>	F	No	Reject GMC (propose DC)
1992	Young and Yamane <sup>[19]</sup>	F	No	Reject GMC (propose DC)
1985	Perrett et al. <sup>[20]</sup>	F	No	Agnostic
1984	Rolls <sup>[21]</sup>	F	No	Reject GMC (propose DC)
1984	Desimone et al. <sup>[22]</sup>	F, B	No	Reject GMC (propose DC)
1982	Perrett et al. <sup>[23]</sup>	F	No	No mention of GMC (propose DC)

A, animals; B, body parts; DC, distributed coding; F, faces; GMC, grandmother cell coding; O, objects; P, patterns; Sc, scenes; Sh, shapes .

monkeys were tested on familiar stimuli they could identify—exactly the conditions in which grandmother cell theories predict highly selective responding. Finally, we briefly mention computational studies that show there are computational advantages with grandmother cell representations. Together, these empirical and computational studies show that it is premature to reject grandmother cells.

### 1.1. What Is the Grandmother Cell Hypothesis?

A grandmother cell is a neuron that represents a single familiar visual category, either a basic-level category (e.g., a cell that represents the visual category “dog” or “bicycle”) or a subordinate-level category (e.g., a cell that represents the visual category “my bicycle” or a specific person such as “my grandmother”). On this hypothesis, an object is identified at a basic level or a subordinate level when the corresponding cell(s) fires beyond some threshold.<sup>[24]</sup> Grandmother cells are often compared to highly sparse representations in which a visual category is represented by a small number of neurons, each neuron only contributing to the representation of a few different visual categories. However, the grandmother cell hypothesis is at the extreme end of the sparseness continuum: a single neuron represents a single visual category.<sup>[25]</sup> It should be noted that grandmother cell theories are not committed to

the claim that all neurons are grandmother cells, and indeed, some combination of grandmother cells and distributed representations may work hand-in-hand in some tasks. By contrast, critics of grandmother cell theories reject the hypothesis that there are any grandmother cells. Within psychology, the equivalent to a grandmother cell is a localist representation in which single units in models encode specific visual categories at a basic or a subordinate level. In this situation, the model identifies inputs by activating single units above some threshold. Unlike grandmother cells in neuroscience, localist representations in psychology are not couched in neuroscience terms but, functionally, grandmother cells and localist representations are equivalent.<sup>[26–28]</sup> Grandmother cells and localist representations constitute a hypothesis about how all basic and subordinate visual categories are identified, and not only about how we identify faces at a subordinate level.

### 1.2. Understanding Grandmother Cells through Modeling

In order to better understand the grandmother cell hypothesis, it is worth briefly reviewing a few models in computer science, neuroscience, and psychology that include grandmother cells (or their equivalent localist coding) in order to identify familiar visual categories at a basic or a subordinate level. The initial

motivation for many of these models was the observation that the visual system is hierarchical: “simple” cells in the visual cortex code for line orientations at a given retinal location and neurons in subsequent layers of the hierarchy code for more abstract (e.g., more spatially invariant) and complex stimuli.<sup>[25]</sup> According to the grandmother cell hypothesis, this process of abstraction continues over multiple layers, the neurons at the top of the visual hierarchy coding for complete objects or persons in such a way that a single neuron responds more strongly to images of this category compared to others. Hubel himself considered whether grandmother cells are the natural extension of his theory of early vision, writing:

“What happens beyond the primary visual area, and how is the information on orientation exploited at later stages? Is one to imagine ultimately finding a cell that responds specifically to some very particular item (Usually one’s grandmother is selected as the particular item, for reasons that escape us). Our answer is that we doubt there is such a cell, but we have no good alternative to offer.”<sup>[29]</sup>

### 1.3. Grandmother Cells in Computer Science and Neuroscience Models

Although Hubel doubts that there are cells that code for visual categories at the top of the visual hierarchy, it is important to note that there are many computational models that have just this property. An early example in computer science was the “neocognitron” model<sup>[30]</sup> that was directly inspired by Hubel and Wiesel’s model<sup>[25]</sup> of simple and complex cells in the primary visual cortex (V1). The neocognitron model roughly simulated the early processes in V1 and continued the hierarchy of processing steps to a top layer that coded each familiar category (written numbers between zero and nine) with a single unit that selectively responds to one number. Indeed, the authors used the term “grandmother cell” to describe these selective units in their model. The neocognitron was in turn one of the inspirations for computational models in neuroscience that include units that selectively represent familiar categories, such as the face of a specific person.<sup>[31]</sup>

More recently, Thorpe and colleagues<sup>[32,33]</sup> have been developing hierarchical artificial neural networks that incorporate the biologically plausible learning rule spike-time-dependent-plasticity (STDP). STDP is an unsupervised form of


learning that adjusts the connection strengths based on the relative timing of a particular neuron’s output and input action potentials. Importantly, for present purposes, models of object identification that employ STDP often learn highly selective units, as illustrated in **Figure 1**. In multiple papers, Thorpe highlights the computational advantages of highly selective representations and explicitly relates these selective representations to grandmother cells. For example, Thorpe wrote:

“I will discuss how a combination of STDP and temporal coding can allow highly selective responses to develop to frequently encountered stimuli. Finally, I will argue that “grandmother cell” coding has some specific advantages not shared by conventional distributed codes.”<sup>[34]</sup>

### 1.4. Localist Models in Psychology

Within psychology, the interactive activation (IA) model of visual word identification is an example of an early and influential theory that coded for familiar categories (words) with single units.<sup>[35]</sup> The IA model was again hierarchically organized with localist letter features at the input layer, localist letter detectors at a second layer, and localist word detectors at the output layer. In this scenario, a specific word is recognized when its localist word unit is activated beyond some threshold. Localist representations are commonplace in theories in psychology, including in models of visual word identification,<sup>[36]</sup> models of spoken word identification,<sup>[37]</sup> spoken word production,<sup>[38]</sup> face perception,<sup>[39]</sup> and many other domains of theorizing.

Despite the existence of the above models, it is important to emphasize that most theorists in psychology and neuroscience strongly reject grandmother cells. For example, within psychology, the localist word representations in the IA model are often rejected in favor of distributed representations within the parallel distributed processing (PDP) models of word identifications.<sup>[40]</sup> Even researchers who endorse localist codes in psychological models are rarely committed to the claim that their model is implemented with grandmother cells; instead, the models are considered theories of cognition with no commitments to neurophysiology. Furthermore, in computational neuroscience where the models are intended to make claims about neural coding, it is very much a minority position to endorse grandmother cells. The far more common claim is



**Figure 1.** The network was trained to recognize objects (as depicted in the top row), and neurons in the hidden layer of the network learned to respond preferentially to image patterns that look like the objects (as depicted in the bottom row). That is, these neurons learned to represent information much like grandmother cells. Reproduced with permission.<sup>[30]</sup> Copyright 1982, Springer.

that single neurons do not code for specific categories.<sup>[41,42]</sup> As Averbeck, Latham, and Pouget put it:

“As in any good democracy, individual neurons count for little; it is the population of activity that matters.”<sup>[43]</sup>

Similarly, researchers who carry out single-cell recording studies overwhelmingly reject grandmother cells. But as we show later, this conclusion is compromised by the common assumption that grandmother cells should code for unfamiliar visual categories.

To summarize, then, a grandmother cell theory is committed to the claim that a subset of neurons selectively represents familiar visual categories at the basic level and at the subordinate level and that different grandmother cells encode basic and subordinate visual categories (one cell coding for the basic visual category “bicycle” and another cell coding for the subordinate “my bicycle”). These selective representations are hypothesized to support the identification of these familiar stimuli when their activation goes beyond some threshold. The grandmother cell hypothesis does not reject distributed representations (both forms of representations may play a role in vision), whereas critics of the grandmother cell hypothesis claim that grandmother cells do not exist.

### 1.5. The Theoretical Confusion that Leads to the Premature Rejection of Grandmother Cell Theory

Bowers has highlighted a number of confusions regarding grandmother cells that have led researchers to reject only straw-man versions of the hypotheses,<sup>[27]</sup> such as the claim that there is one grandmother neuron per visual category (such that a loss of one neuron can lead to a failure to recognize your grandmother)<sup>[44]</sup> and the claim that grandmother cells respond to one category of input and remain entirely silent to all other categories (contrary to the common observation that a neuron responds above baseline to images from multiple categories).<sup>[45,46]</sup> But there is no reason to adopt these characterizations of grandmother cell theory. Indeed, the grandmother/localist representations in the various models described above do not have the property that they only respond to one category of input and are silent to everything else. Rather they have the property that they respond most strongly to one category and less strongly to visually similar inputs from different categories. Of course, visually similar inputs should not activate a grandmother/localist unit beyond the threshold for identification, and if they do, it would result in a misidentification of the input. For a more detailed discussion and debate on this issue, see refs. <sup>[27,47–49]</sup> Similarly, multiple redundant grandmother cells are consistent with grandmother cell theory,<sup>[26,27,46]</sup> and indeed, redundant localist codes can emerge in artificial neural network models after training,<sup>[50]</sup> making the model more robust to the removal of a single unit.

But here we focus on the false claim that grandmother cell theories should explain how we perceive and identify unfamiliar things. This leads researchers to falsely reject grandmother cells on the basis that there are just not enough neurons in the brain to code for all possible unfamiliar stimuli,<sup>[48,51]</sup> and most critical for present purposes leads researchers to falsely reject

grandmother cells on the basis of not finding neurons that respond selectively to unfamiliar stimuli.

## 2. What Does the Research Show?

In order to review as many studies as possible, we describe some of the most relevant papers in the text and list other key studies in Table 1. In many or most studies reviewed below the monkeys were repeatedly presented with the stimuli used to drive visual neurons, either before the recording session or during the recording session itself. For this reason, the authors of these studies often report that the stimuli tested were familiar. However, in most of these studies, the monkeys simply saw the images passively without making any response, or they performed some task that did not require them to identify the images at the basic or at the subordinate level. Accordingly, there is no reason to assume that the monkeys were able to identify the images at the subordinate or basic levels. Our main point is that if animals cannot identify the stimuli, there is no reason to reject grandmother cells when single neurons fail to respond selectively to these stimuli.

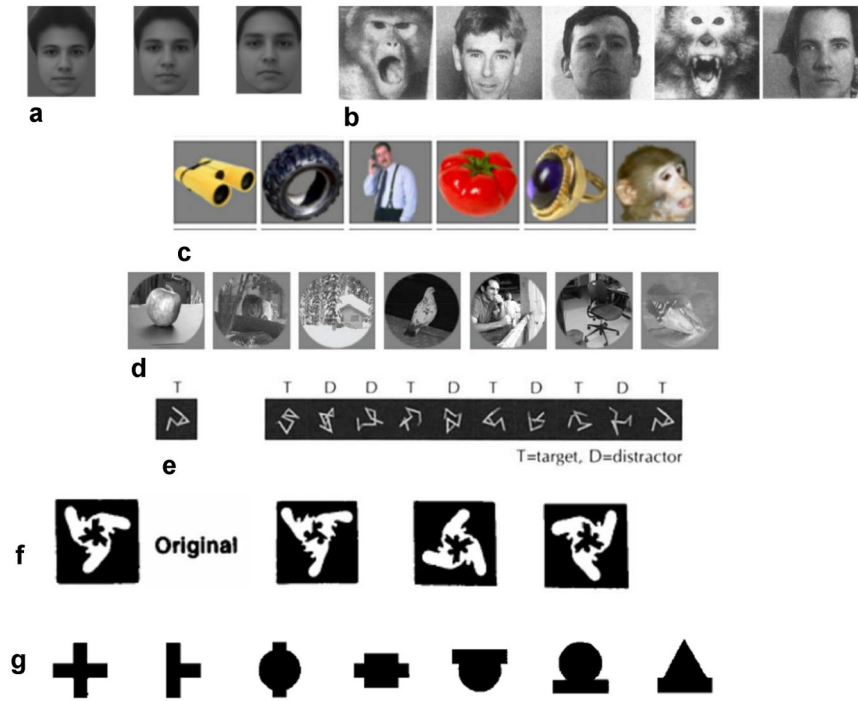
### 2.1. Chang and Tsao (2017)

We start with the Chang and Tsao study<sup>[52]</sup> because it has received a lot of attention, makes strong claims in support of distributed over grandmother cell coding for face identification, and it is the most recent of a series of papers by Tsao and colleagues on the neural coding of faces (see Table 1 for other related studies from this group).

The authors showed images of human faces to two monkeys while recording from neurons from the middle lateral (ML)/middle fundus (MF), and anterior medial (AM) regions of the inferotemporal (IT) cortex that have previously been shown to be involved in different aspects of face processing. The faces were synthetic human faces that had been transformed in high-dimensional space in order to assess the impact of various transformations on neural firing (see **Figure 2a** for example faces). The faces were presented for 150 ms, interleaved with 150 ms of a gray screen, each image being presented between three and five times. The stimuli included 2000 images of parametrized frontal faces, 2000 images of parameterized profile faces, and the task of the monkeys was simply to fixate on a fixation point for a juice reward, and accordingly, the faces were just viewed passively. They reported that a pattern of activation over approximately 200 neurons coded for specific faces and took the findings to support distributed rather than grandmother cell coding. In the “In Brief” summary of their article, they wrote:

“Facial identity is encoded via a remarkably simple neural code that relies on the ability of neurons to distinguish facial features along specific axes in face space, disavowing the long-standing assumption that single face cells encode individual faces.”<sup>[52]</sup>

But no theory (grandmother or otherwise) predicts that single cells in monkey brains should selectively respond to



**Figure 2.** Stimuli used across multiple experiments: a) Chang and Tsao, Reproduced with permission.<sup>[52]</sup> Copyright 2017, Elsevier. b) Rolls, Reproduced with permission.<sup>[59]</sup> Copyright 1995, The American Physiological Society. c) Lehky et al., Reproduced with permission.<sup>[53]</sup> Copyright 2011, The American Physiological Society. d) Rust and DiCarlo, Reproduced with permission.<sup>[54]</sup> Copyright 2012, The Authors, published by the Society for Neuroscience. e) Logothetis et al., Reproduced with permission.<sup>[55]</sup> Copyright 1995, Elsevier. f) Sakai et al., Reproduced with permission.<sup>[56]</sup> Copyright 1994, Cold Spring Harbor Laboratory Press. and g) Kobatake et al., Reproduced with permission.<sup>[57]</sup> Copyright 1998, The American Physiological Society.

unfamiliar human faces. In the same way, no theory should expect to find neurons in the human visual cortex that selectively responds to an unfamiliar monkey face.

The conceptual confusion of the authors was highlighted more strongly in the interview with Doris Tsao, who said the following in an interview linked with the Chang and Tsao (2017) paper:

“Before this work that is described in this paper itself... people thought at the highest levels of the brain’s face recognition system there are cells that are selective for specific individuals, all the people that you know and recognize there are cells encoding them.

And obviously this raised a question, which is how can one have enough cells to represent all the people that you possibly could recognize. There are 6 billion people on this earth, and obviously you do not have 6 billion cells specialized for face recognition in your brain. So it was a mystery how it is ultimately done.”<sup>[58]</sup>

The conclusion Tsao draws is that the brain must rely on distributed coding rather than grandmother cells. But given this conception of grandmother cells, it raises the question why the authors even bothered using their data to argue against grandmother cells, given that it was ruled out a priori on the basis that there are not enough specialized neurons in a monkey brain. Nevertheless, this same conclusion was drawn by Quian Quiroga in a “Leading Edge Previews” article in the same issue entitled “How Do We Recognize a Face?”<sup>[1]</sup>

He writes:

“As the authors argue, their results imply that there are no detectors for face identity at the single neuron level in the face patch system and, consequently, this may put an end to the long-standing dispute about the existence of grandmother cells in visual cortex.”<sup>[1]</sup>

## 2.2. Rolls (2017)

Rolls, in a special issue of a journal devoted to the topic of grandmother cells, reviewed a number of findings from his laboratory that he takes to falsify the grandmother cells.<sup>[42]</sup> As he writes in the abstract: “The encoding of information in the primate inferior temporal visual cortex, hippocampus, orbito-frontal cortex, and insula is described. All these areas have sparse distributed graded firing rate representations.” However, once again, this conclusion comes from studies in which animals were presented with unfamiliar stimuli, or where the familiarity of the stimuli was not discussed, and where the animal was not needed to distinguish images one from another.

For example, the first study he summarizes was carried out by Rolls and Tovee<sup>[59]</sup> that assessed the neural coding of face and object identity in the temporal visual cortex. This highly cited paper assessed the firing selectivity of 14 neurons to a set of 68 stimuli. None of the neurons responded selectively to one of the stimuli, and Rolls takes the findings as inconsistent with the grandmother cell hypothesis.<sup>[42]</sup> This is a strong conclusion

to reach on the basis of recording from 19 neurons, but more relevant for present purposes, the stimuli were composed of unfamiliar human and monkey faces and unfamiliar nonface stimuli that included images of woodland, countryside, and foods: some of these images are reproduced in Figure 2b. The monkey's task was simply to look at the photos without responding (the monkeys received a reward by licking a tube between image trials), and accordingly, there was no pressure to learn to identify the images. Given that the monkey was not trained to identify the stimuli, the experiment does not test the grandmother cell hypothesis.

A search of Google Scholar using the search terms "grandmother" and "Rolls" (as author) identifies 87 papers and, as far as we can tell, there is only one paper where there was a brief note that selectivity might be greater for familiar categories,<sup>[21]</sup> but this was limited to a single paragraph where they simply note that their findings "raise the possibility" (p. 213) that visual experience may impact on the tuning of single neurons.

### 2.3. Lehky, Kiani, Esteky, and Tanaka (2011)

Lehky et al. recorded from IT cells of 674 monkeys, each stimulated by 806 object photographs.<sup>[53]</sup> Although they found some neurons to respond highly selectively, they failed to obtain evidence for grandmother cells, writing:

"We believe that the data presented here do not support "grandmother cell" coding.... On average, the second largest response was almost the same size as the largest response (89.3% for single-neuron responses, 79.5% for population responses). This is not a characteristic of "grandmother cell" coding."<sup>[53]</sup>

But again, the conclusion is not justified, given that the stimuli were not familiar and the monkeys were not trained to discriminate between the stimuli (see Figure 2c for some example images from this experiment). Interestingly, the authors do leave the door open for grandmother cells for familiar, meaningful stimuli, writing: "Nevertheless, it is still possible that grandmother cell encoding could occur for a small number of special objects to which the observer was highly exposed, and which also had strong behavioral significance." Why such selective coding should be restricted to a few categories with strong behavioral significance as opposed to all visual categories that can be identified is unclear.

### 2.4. Rust and DiCarlo (2012)

Rust and DiCarlo<sup>[54]</sup> designed a study to measure V4 and IT neuronal sparseness using a set of 300 natural images that included an object in its natural context, each object being distinct (it would be called by a different name). The authors wrote that: "Images included a wide variety of content, including objects familiar to the animal, other (unfamiliar) animals, man-made objects, other monkeys, and people," but provided no analyses of whether the results were impacted on familiarity. But the stimuli reported in their figures suggest that

the vast majority of images would be unfamiliar and a few example images are depicted here in Figure 2d. Indeed, the authors appear to consider the familiarity of the images a problem, writing: "To guard against possible nonstationary effects (e.g., familiarity with the images), recordings were altered between V4 and IT" (p. 10173). Again, the monkeys were not required to discriminate between any of these images. Based on the fact that they did not find any neurons that selectively responded to the test images, the authors argue that the visual system codes information in a distributed manner.

But again, no grandmother cell theory would predict selective responding to the unfamiliar images, and the authors did not specifically discuss or analyze the results for the familiar images in the experiment. As summarized in Table 1, other studies have also failed to observe highly selective responses of IT neurons in response to stimuli that a monkey was not trained to identify.

There is yet another issue with some of the above studies that further weakens the conclusions that have been drawn regarding grandmother cells. That is, there is growing evidence that monkeys are quite poor at recognizing faces and that they recognize faces in a qualitatively different way than humans.<sup>[60]</sup> The findings include the observation that monkeys require extensive exposure and training in order to achieve only moderate performance in identifying specific faces in a laboratory setting, performance on matching tasks is worse for faces than for other objects, and monkeys are no better at matching familiar compared to unfamiliar faces (unlike humans). In addition, Macaques show no inversion effects, have no distinct ventral face-specific pathway, and no right hemisphere specialization for faces. This further undermines the common rejection of grandmother cells based on studies that did not even train monkeys to identify (unfamiliar) faces.

## 3. Highly Selective Representations Are More Often Found when Familiar Stimuli Are Tested

A number of studies<sup>[57,61–63]</sup> have reported that the selectivity of neural firing increases as the stimuli become more familiar to the animal and this alone suggests that the above results provide a poor basis for characterizing the representations that support object and face identification. But the important question for present purposes is whether stimuli that can be identified at a basic or at a subordinate level evoke levels of selectivity in visual neurons that are consistent with grandmother cells.

The most high-profile set of experiments that have reported highly selective neural firing in response to familiar stimuli were recorded from the hippocampus and related structures in humans (e.g., the "Jennifer Aniston" neurons<sup>[64]</sup>). It is important to note, however, that these neurons are part of the memory rather than the visual system, and accordingly, the results are not directly relevant to the classic grandmother cell hypothesis that is concerned with vision. Nevertheless, it raises the question of whether similar levels of selectivity can be obtained with familiar stimuli within the visual system.

In fact, as reviewed below, a few studies have reported highly selective responses in IT cells when monkeys are presented with stimuli they can identify, and, in one case, the level of selectivity is as high as any selectivity observed in the

hippocampus. Although many of the studies are quite old, they are highly relevant today. Unfortunately, the central message that selectivity can be extremely high for familiar objects in the visual cortex seems to have been lost on the many recent experimental single-cell recording studies reviewed above as well as recent theoretical discussions of these findings;<sup>[1,65,66]</sup> but see Guntupalli and Gobbin<sup>[67]</sup> for a counterexample.

### 3.1. Logothetis, Pauls, and Poggio (1995)

Logothetis et al.<sup>[55]</sup> trained two monkeys to identify over 100 novel computer-generated objects from various viewpoints over the course of months. After learning a substantial subset of stimuli, the monkeys performed a visual matching task in which they first fixated at a target stimulus from one viewpoint and then saw a series of test stimuli from various viewpoints that were from the same class or not, as illustrated in Figure 2e. The monkeys categorized the objects as matching or mismatching while the authors recorded from 796 neurons in the upper bank of the AM temporal sulcus. The neurons showed a range of selectivity, a few (3/796; 0.37%) responding selectively to only one object presented from any viewpoint, and a larger set (93/796; 11.6%) responding selectively to a subset of views of one of the known target objects but less frequently (or not at all) to highly similar objects. These are perhaps the most selective responses ever recorded in any part of the brain. Critically, no selective responses were encountered for views that the animal systematically failed to recognize. That is, highly selective codes were associated with the ability to identify the object. Furthermore, the authors found that the percentages of cells responding to objects from a given class correlated with the amount of training. This led the authors to write: "Thus, it seems that neurons in this area may develop complex configurational selectivity as the animal is trained to recognize specific objects."<sup>[55]</sup>

### 3.2. Sakai, Naya, and Miyashita (1994)

Sakai et al.<sup>[56]</sup> trained monkeys to recognize 12 pairs of computer-generated Fourier patterns while recording from 474 neurons in the Anterior Inferior Temporal cortex. One of the two stimuli from a pair (the "cue") was presented for 0.5 s and after a 5 s delay, the paired associate image or a foil image was presented side by side and the monkey had to select the associate pattern to receive an award. This required the monkeys to learn and discriminate these patterns. The authors identified 89 neurons that responded most strongly to one of the cue stimuli. The Fourier pattern cue stimuli were then systematically manipulated in order to assess the impact of varying the learned images on neural responses. In the vast majority of cases, the neurons responded more strongly to the trained pattern compared to the untrained transformed ones, and in no case did the neuron respond more strongly to the transformed pattern. This suggests that the neurons were tuned to the trained visual patterns. This is exactly as one might expect from a grandmother cell theory. See Figure 2f for an example of trained images and manipulated images.

### 3.3. Perrett, Smith, Potter, Mistlin, Head, Milner, and Jeeves (1985)

Perrett et al.<sup>[20]</sup> recorded from cells in the superior temporal sulcus of the macaque monkey in studies that were designed to distinguish between faces from nonface stimuli (which is the general category of the face rather than the identity of a specific face). But during testing, the authors noted that some cells responded to one familiar person more than others. For instance, one cell responded more to a wide range of views of one familiar person (Paul Smith, one of the experimenters) compared to another (David Perrett, another experimenter), such that the identity of the person (based on whole body, not just face) could be determined by the activation of this one neuron. Perrett and colleagues considered these findings to highlight the role of learning on neural selectivity and their findings were at least consistent with grandmother cell coding schemes, writing:

"The responses of these cells have many of the properties hypothesized for "gnostic units" and provide insight into the final stages of visual processing leading to the recognition of an object as a face and more specifically the identity of the fact."<sup>[68]</sup>

Note that the term "gnostic unit" is being used there the same way as the grandmother cell.

### 3.4. Kobatake, Wang, and Tanaka (1998)

Kobatake et al. trained two monkeys to discriminate between 28 moderately complex shapes and then recorded from 131 cells in IT cortex.<sup>[57]</sup> A single shape (for examples, see Figure 2) was displayed on a computer screen and it disappeared when it was touched. After a short delay, the stimulus was displayed along with four additional foils from the same set of 28 stimuli. The task was to select the repeated stimulus. The training was extensive and was completed when the monkeys performed 500 successful trials per day, with a success rate of over 75%. Neural recordings to these stimuli were performed while these two trained monkeys were under anesthesia, as well as 130 IT neurons in two additional monkeys who were not trained to identify and remember the stimuli.

For the trained monkeys, there were 28 cells that maximally responded to one of the trained stimuli and for these neurons, an average of three other stimuli from the trained set evoked responses >50% of individual cells' maximal response. Based on this the authors concluded that "The broad tuning may suggest that the discrimination depended on the activity of cell population." Unfortunately, however, the authors did not report whether any of the 26 critical failed to respond strongly to other stimuli (the overall average of three stimuli evoking strong response may well have included some neurons that were more selective). The failure to report the most selective responding neuron is unfortunate. Importantly, the authors also reported that the selectivity of the neurons was much greater for the trained compared to nontrained monkeys, highlighting the fact that selectivity is greater for stimuli that can be identified.

### 3.5. Khuvis, Yeagle, Norman, Grossman, Malach, and Mehta (2018)

To the best of our knowledge, Khuvis et al.<sup>[69]</sup> have provided the first test of whether individual neurons in the human visual cortex code for specific familiar faces. The authors recorded from 63 neurons in the ventral temporal cortex across eight patients. A set of images of ten famous faces as well as ten body parts, ten houses, ten tools, and ten abstract patterns were presented centrally and participants performed a one-back task during recording, in which they indicated whether the current stimulus matched the previous one. Twenty-six neurons were found to be selective for the general category faces and seven for nonface categories. However, within the face category, the authors did not report any neurons that selectively responded to one face. It will be important to see whether similar results are obtained in humans in follow-up studies, especially if it is possible to test a wider range of neurons and face images.

We want to be clear that we do not consider reports of extreme selectivity in visual neurons in response to trained stimuli as evidence in support of the grandmother cell hypothesis. There are too few relevant studies, the results are mixed, and it is perfectly plausible that the most selective neurons are not so selective that they are tuned to maximally respond to one visual category. Nevertheless, the most striking demonstrations of selectivity are consistent with the grandmother cell hypothesis, and in our view, it is premature to reject grandmother cells based on the current findings.

## 4. Conclusions and Outlook

We hope that most readers will accept our argument that it is inappropriate to reject grandmother cells on the basis of testing stimuli that cannot be identified, and it is an important observation that neural selectivity is increased when familiar stimuli are tested. At the same time, we expect many will continue to dismiss the grandmother cell hypothesis. The notion that among the billions of neurons in a brain there is a subset of neurons tuned to code for specific categories might seem implausible. We understand this intuition, but it is worth emphasizing again that there are existing models of visual object recognition that are designed to be biologically plausible that have grandmother cells.<sup>[31]</sup> Indeed, there are biologically plausible models of visual object identification that learn grandmother cells.<sup>[70]</sup> Even artificial neural network models that are claimed to learn distributed codes in fact learn grandmother representations when trained to coactivate multiple words (or objects or faces) at the same time in short-term memory.<sup>[50,71]</sup> These findings suggest that there are computational advantages of grandmother cells in the context of short-term memory, much like there are computational advantages of learning highly selective representations in the hippocampus for the sake of episodic memory.<sup>[72]</sup> Grandmother cells have also been observed in artificial neural network models that learn to identify faces.<sup>[73]</sup>

In our view, these computational considerations, in combination with the empirical findings above, should lead researchers and theorists to take the grandmother cell hypothesis more

seriously. Going forward, researchers need a better understanding of existing models of face and object identification that include grandmother cells, the conditions in which artificial neural networks learn grandmother cells,<sup>[50,71]</sup> and then design experiments that compare the selectivity of these units to the selectivity of visual neurons responding to stimuli that can be identified.

## Conflict of Interest

The authors declare no conflict of interest.

## Keywords

distributed representations, grandmother cells, object identification

Received: December 4, 2018

Revised: May 1, 2019

Published online:

- [1] R. Quiñ Quiroga, *Cell*. **2017**, *169*, 975.
- [2] J. Dubois, A. O. de Berker, D. Y. Tsao, *J. Neurosci.* **2015**, *35*, 2791.
- [3] E. M. Meyers, M. Borzello, W. A. Freiwald, D. Tsao, *J. Neurosci.* **2015**, *35*, 7069.
- [4] J. Taubert, G. Van Belle, W. Vanduffel, B. Rossion, R. Vogels, *J. Neurosci.* **2015**, *35*, 9872.
- [5] E. B. Issa, J. J. DiCarlo, *J. Neurosci.* **2012**, *32*, 16666.
- [6] W. A. Freiwald, D. Y. Tsao, *Science* **2010**, *330*, 845.
- [7] E. M. Meyers, D. J. Freedman, G. Kreiman, E. K. Miller, T. Poggio, *J. Neurophysiol.* **2008**, *100*, 1407.
- [8] L. Franco, E. T. Rolls, N. C. Aggelopoulos, J. M. Jerez, *Biol. Cybern.* **2007**, *96*, 547.
- [9] D. Y. Tsao, W. A. Freiwald, R. B. H. Tootell, M. S. Livingstone, *Science* **2006**, *311*, 670.
- [10] D. J. Freedman, M. Riesenhuber, T. Poggio, E. K. Miller, *Cereb. Cortex*. **2005**, *16*, 1631.
- [11] C. Keyers, D. K. Xiao, P. Földiák, D. I. Perrett, *J. Cognitive Neurosci.* **2001**, *13*, 90.
- [12] D. L. Sheinberg, N. K. Logothetis, *J. Neurosci.* **2001**, *21*, 1340.
- [13] H. Tamura, K. Tanaka, *Cereb. Cortex* **2001**, *11*, 384.
- [14] Y. Sugase, S. Yamane, S. Ueno, K. Kawano, *Nature* **1999**, *400*, 869.
- [15] M. C. Booth, E. T. Rolls, *Cereb. Cortex* **1998**, *8*, 510.
- [16] E. T. Rolls, A. Treves, M. J. Tovee, *Exp. Brain Res.* **1997**, *114*, 149.
- [17] S. Higuchi, Y. Miyashita, *Proc. Natl. Acad. Sci. U. S. A.* **1996**, *93*, 739.
- [18] E. T. Rolls, M. J. Tovee, *J. Neurophysiol.* **1995**, *73*, 713.
- [19] M. Young, S. Yamane, *Science* **1992**, *256*, 1327.
- [20] D. I. Perrett, P. A. Smith, D. D. Potter, A. J. Mistlin, A. S. Head, A. D. Milner, M. A. Jeeves, *Proc. R. Soc. London* **1985**, *223*, 293.
- [21] E. T. Rolls, *Hum. Neurobiol.* **1984**, *3*, 209.
- [22] R. Desimone, T. Albright, C. Gross, C. Bruce, *J. Neurosci.* **1984**, *4*, 2051.
- [23] D. I. Perrett, E. T. Rolls, W. Caan, *Exp. Brain Res.* **1982**, *47*, 329.
- [24] J. S. Bowers, *Lang. Cogn. Neurosci.* **2017**, *32*, 257.
- [25] D. H. Hubel, T. N. Wiesel, *Am. J. Physiol.* **1959**, *148*, 574.
- [26] M. Page, *Behav. Brain Sci.* **2000**, *23*, 443.
- [27] J. S. Bowers, *Psychol. Rev.* **2009**, *116*, 220.
- [28] J. S. Bowers, *Trends Cognit. Sci.* **2017**, *21*, 950.
- [29] D. H. Hubel, T. N. Wiesel, *Sci. Am.* **1979**, *241*, 150.
- [30] K. Fukushima, S. Miyake, *Competition and Cooperation in Neural Nets*, Springer, Berlin **1982**, pp. 267–285.
- [31] M. Riesenhuber, T. Poggio, *Nat. Neurosci.* **1999**, *2*, 1019.
- [32] R. Guyonnet, R. VanRullen, S. J. Thorpe, *Neural Comput.* **2005**, *17*, 859.



- [33] T. Masquelier, S. J. Thorpe, *PLoS Comput. Biol.* **2007**, *3*, e31.
- [34] S. J. Thorpe, *Visual Population Codes: Toward a Common Multivariate Framework for Cell Recording and Functional Imaging*, MIT Press, Cambridge, MA **2011**, pp. 23–51.
- [35] J. L. McClelland, D. E. Rumelhart, *Psychol. Rev.* **1981**, *88*, 375.
- [36] C. J. Davis, *Psychol. Rev.* **2010**, *117*, 713.
- [37] T. Hannagan, J. S. Magnuson, J. Grainger, *Front. Psychol.* **2013**, *4*, 563.
- [38] G. S. Dell, *Psychol. Rev.* **1986**, *93*, 283.
- [39] A. M. Burton, V. Bruce, R. A. Johnston, *Br. J. Psychol.* **1990**, *81*, 361.
- [40] M. S. Seidenberg, J. L. McClelland, *Psychol. Rev.* **1989**, *96*, 523.
- [41] P. S. Churchland, T. J. Sejnowski, T. A. Poggio, *The Computational Brain*, MIT Press, Cambridge, MA **2016**.
- [42] E. T. Rolls, *Lang. Cogn. Neurosci.* **2017**, *32*, 316.
- [43] B. B. Averbeck, P. E. Latham, A. Pouget, *Nat. Rev. Neurosci.* **2006**, *7*, 358.
- [44] W. W. Lytton, *From Computer to Brain: Foundations of Computational Neuroscience*, Springer-Verlag, New York **2007**.
- [45] S. Waydo, A. Kraskov, R. Quian Quiroga, I. Fried, C. Koch, *J. Neurosci.* **2006**, *26*, 10232.
- [46] C. G. Gross, *Neuroscientist* **2002**, *8*, 512.
- [47] J. S. Bowers, *Psychol. Rev.* **2010**, *117*, 300.
- [48] D. C. Plaut, J. L. McClelland, *Psychol. Rev.* **2010**, *117*, 284.
- [49] R. Quian Quiroga, G. Kreiman, *Psychol. Rev.* **2010**, *117*, 291.
- [50] J. S. Bowers, I. I. Vankov, M. F. Damian, C. J. Davis, *Psychol. Rev.* **2014**, *121*, 248.
- [51] H. Eichenbaum, in *Dictionary of Biological Psychology* (Ed: P. Winn), Routledge, London and New York **2003**, pp. 558–559.
- [52] L. Chang, D. Y. Tsao, *Cell* **2017**, *169*, 1013.
- [53] S. R. Lehky, R. Kiani, H. Esteky, K. Tanaka, *J. Neurophysiol.* **2011**, *106*, 1097.
- [54] N. C. Rust, J. J. DiCarlo, *J. Neurosci.* **2012**, *32*, 10170.
- [55] N. K. Logothetis, J. Pauls, T. Poggio, *Curr. Biol.* **1995**, *5*, 552.
- [56] K. Sakai, Y. Naya, Y. Miyashita, *Learn. Mem.* **1994**, *1*, 83.
- [57] E. Kobatake, G. Wang, K. Tanaka, *J. Neurophysiol.* **1998**, *80*, 324.
- [58] D. Y. Tsao, PaperClip Interview with author Dr. Doris Tsao, *Cell*, **2017**, *169*, [https://www.cell.com/cell/fulltext/S0092-8674\(17\)30538-X](https://www.cell.com/cell/fulltext/S0092-8674(17)30538-X) (accessed: April 29, 2019).
- [59] E. T. Rolls, M. J. Tovee, *J. Neurophysiol.* **1995**, *73*, 713.
- [60] B. Rossion, J. Taubert, *Vision Res.* **2018**, *13*
- [61] C. A. Erickson, B. Jagadeesh, R. Desimone, *Nat. Neurosci.* **2000**, *3*, 1143.
- [62] N. Li, J. J. DiCarlo, *Science*. **2008**, *321*, 1502.
- [63] M. Riesenhuber, L. S. Glezer, *Lang. Cogn. Neurosci.* **2017**, *32*, 286.
- [64] R. Q. Quiroga, L. Reddy, G. Kreiman, C. Koch, I. Fried, *Nature*. **2005**, *435*, 1102.
- [65] W. J. Clark, M. Colombo, in *Encyclopedia of Animal Cognition and Behavior* (Eds: J. Vonk, T. Shackelford), Springer, Cham **2018**.
- [66] B. R. Conway, *Annu. Rev. Vis. Sci.* **2018**, *4*, 402.
- [67] J. S. Guntupalli, M. I. Gobbini, *Trends Cognit. Sci.* **2017**, *21*, 915.
- [68] D. I. Perrett, A. J. Mistlin, A. J. Chitty, *Trends Neurosci.* **1987**, *10*, 358.
- [69] S. Khuvis, E. M. Yeagle, Y. Norman, S. Grossman, R. Malach, A. D. Mehta, *bioRxiv* **2018**, 487686.
- [70] S. R. Kheradpisheh, M. Ganjtabesh, S. J. Thorpe, T. Masquelier, *Neural Networks* **2018**, *99*, 56.
- [71] J. S. Bowers, I. I. Vankov, M. F. Damian, C. J. Davis, *Cognition*. **2016**, *148*, 47.
- [72] D. Marr, D. Willshaw, B. McNaughton, *From the Retina to the Neocortex*, Springer, Berlin **1991**, p. 59.
- [73] I. I. Vankov, J. S. Bowers, *Lang. Cogn. Neurosci.* **2017**, *32*, 392.