The persistent and unreasonable confusion regarding grandmother cells: Comment on Chang and Tsao (2017)

Jeffrey Bowers
School of Psychological Science,
University of Bristol
12a Priory Road, Bristol,
BS8-1TU
j.bowers@bristol.ac.uk

Abstract

Chang and Tsao (2017) failed to find any cells in monkey inferior temporal (IT) cortex that selectively respond to synthesized human faces and conclude that the visual system does not code for faces using “grandmother cells”. Here I show that this conclusion is not justified because grandmother cells are about how the visual system identifies familiar categories, such as images of faces that you know (e.g., your grandmother).

For many years I have argued that the “grandmother cell” hypothesis should be taken seriously (Bowers, 2002, 2009). On this view, single neurons code for familiar categories, such as well-known persons (Jennifer Aniston), objects, and words. The alternative view is that each neuron is involved in representing many different categories, and that a pattern of activation over many neurons codes for a specific category; so-called distributed coding. Unfortunately, it has been difficult to have a constructive debate on this issue because critics typically reject straw-man versions of the grandmother cell hypothesis. The most recent high-profile example of this was published a few months ago by Chang and Tsao (2017) in a paper entitled: “The Code for Facial Identity in the Primate Brain”. A key claim of the authors is that their findings rule out the hypothesis that single neurons code for single faces. As they write in the “In Brief” section of their article:

“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.**” [bold added]

This claim was also highlighted by Quian Quiroga in a “Leading Edge Previews” article in the same issue entitled “How Do We Recognize a Face?” 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.**” [bold added]

The first thing to note is that almost no one apart from a few oddballs (e.g., Simon Thorpe, Mike Page, and myself) are strongly arguing that the grandmother cell hypothesis should be taken seriously. Far from challenging “the long-standing assumption that single face cells encode individual faces”, the authors are defending the status-quo. But the more important point is that the Chang and Tsao paper provide no evidence against grandmother cells. The problem is that the authors have a basic misunderstanding of the hypothesis, and as a result, they have not carried out a relevant experiment.

What was the key finding that the authors took to refute grandmother cells? They failed to find any cells in monkey IT cortex that selectively responded to unfamiliar human faces. Instead, they found neurons that responded to single dimensions within a 50-dimensional “face space”, with a pattern of activation over ~ 200 of these neurons coding for these faces. That is, these faces were coded by distributed rather than grandmother cells.

The observation that these neurons represent part of a face space is an interesting claim. But the important point for present purposes is that a grandmother cell theory is a theory about how we recognize *familiar* categories – people like your grandmother, not unfamiliar people. No theory (grandmother or otherwise) predicts that single cells in monkey brains should selectively respond to unfamiliar human faces. In the same way, no theory should expect to find neurons in human visual cortex that selectively responds to an unfamiliar monkey face. Indeed, when Quian Quiroga and colleagues were looking for highly-selective cells in the hippocampus of humans they deliberately presented participants with images of highly familiar people, objects, and scenes.

But Chang and Tsao (2017) do not understand that grandmother cells are a theory about how we recognize familiar categories. This is made explicit in the PaperClip interview with Doris Tsao that is linked with the *Cell* paper. In the interview she says:

“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 you have enough cells to represent all the people that you possible 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” [bold added]

This does raise the question why they even bothered testing their version of grandmother theory given that it was ruled out *a priori* on the basis that there are not enough specialized neurons in a monkey brain.

But there is an interesting alternative grandmother cell hypothesis that a few researchers do entertain, namely, the view that some neurons selectively represent familiar categories, things like familiar faces, words, and objects. On this view, single neurons respond most strongly to inputs from one specific category, although the neuron may respond to a lesser degree to inputs from similar categories. The category is identified when the neuron fires beyond some threshold. This requires considerably less than 6 billion grandmother cells devoted to faces. Nevertheless, critics of grandmother cells keep rejecting the straw-man version, and don’t seem to even appreciate that there is an alternative hypothesis that they should consider.

It is not that this more plausible version of the grandmother cell hypothesis is obscure, or somehow a post-hoc attempt to maintain a discredited view. It is the hypothesis that all researchers sympathetic to grandmother cells have endorsed from the start, in multiple high-profile outlets (e.g., Gross, 2002; Konorski, 1967; Page, 2000; Thorpe, 1995). For example, in an article entitled “On the biological plausibility of grandmother cells: Implications for neural network theories in psychology and neuroscience”, I wrote:

“...the grandmother cell theory is a theory about how familiar items are coded (no one ever claimed there were individual neurons for unfamiliar words, novel faces, or novel thoughts).” (Bowers, 2009, 244-245)

Subsequently there was an exchange between myself (Bowers, 2010) and Quian Quiroga and Kreiman (2010) as well as Plaut and McClelland (2010). Just as Chang and Tsao (2017), Plaut and McClelland rejected grandmother cells because there just not enough neurons to code for all experiences. They note it is not possible to have a neuron devoted to a specific tulip on McClelland’s dining room table. Again, I gave the same response:

The core problem with this analysis is that a grandmother cell theory is only committed to the claim that single neurons code for an equivalent class of familiar things. Accordingly, it is only necessary to devote a single unit to a specific tulip on McClelland’s dining room table if McClelland can identify it (as opposed

to other tulips). Barring this, it is possible that there is a unit for tulips (in general) at the top of his visual processing hierarchy, and a tulip is identified when the tulip cell is activated beyond some threshold. The perceptual vividness of each tulip might be due to the specific set of coactive neurons across all the levels of the visual hierarchy. (Bowers, 2010, p. 303)

Nevertheless, despite this exchange, Quian Quiroga (2017) claims that failing to find selective cells for unfamiliar human faces in the visual cortex of monkeys might finally put a rest to the grandmother cell hypothesis.

I think there is one point we all can agree on, namely, there are no grandmother cells of the sort that Chang and Tsao (2017), Quian Quiroga (2017) and Plaut and McClelland (2010) are rejecting. They are correct: There really are not enough neurons in the human brain to code for all possible monkey and human faces, let alone all possible scenes, sentences, and thoughts, etc., using grandmother cells. The same goes for monkey brains. It is also a view that no one ever held.

At the same time, it is important not reject a serious hypothesis based on the Chang and Tsao study, namely, the view that some neurons selectively represent *familiar* categories. This version of grandmother cells may also turn out be wrong, but few people even consider this hypothesis because they think rejecting an absurd theory of a grandmother cell rules out all versions. In fact, there is an abundance of neuroscientific and computational evidence published in high-profile journals that highlight the biological plausibility and computational advantages of grandmother cells for coding familiar categories. Indeed, grandmother cells have are learned in artificial PDP and deep neural networks, and these findings may help provide an explanation as to why some neurons in cortex respond in such highly-selective manner (for review see my recent TICS paper; Bowers, 2017). It would be great if researchers who reject grandmother cells would address these points.

To get a range of perspectives on this issue, I would suggest having a look at the special issue of the journal *Language, Cognition, and Neuroscience*, 2017, issue 32, entitled “Cognitive and neurophysiological evidence for and against localist ‘grandmother cell’ representations”.

If you would like to comment on this piece, please go to the following link:

<https://jeffbowers.blogs.bristol.ac.uk/blog/grandmother-cells/>

References:

- Bowers, J.S. (2002). Challenging the widespread assumption that connectionism and distributed representations go hand-in-hand. *Cognitive Psychology*, 45, 413-445.
- Bowers, J.S. (2009). On the biological plausibility of grandmother cells: Implications for neural network theories in psychology and neuroscience. *Psychological Review*, 116, 220-251.
- Bowers, J.S. (2010). More on grandmother cells and the biological implausibility of PDP models of cognition: A reply to Plaut and McClelland (2010) and Quian Quiroga and Kreiman (2010). *Psychological Review*, 117, 300-306. doi:10.1037/a0018047
- Bowers, J.S. (2017). Parallel Distributed Processing theory in the age of deep networks. *Trends in Cognitive Science*. <https://doi.org/10.1016/j.tics.2017.09.013>
- Chang, L., & Tsao, D. Y. (2017). The Code for Facial Identity in the Primate Brain. *Cell*, 169(6), 1013-1028.
- Gross, C. G. (2002). Genealogy of the “grandmother cell”. *The Neuroscientist*, 8(5), 512-518.
- Konorski, J. (1967). *Integrative activity of the brain: an interdisciplinary approach*. Chicago: University of Chicago Press.
- Page, M. P. A. (2000). Connectionist modeling in psychology: A localist manifesto. *Behavioral and Brain Sciences*, 23, 443-512.
- Page, M. P. A (2017). Localist models are compatible with information measures, sparseness indices, and complementary-learning systems in the brain. *Language, Cognition and Neuroscience*, 32(3), 366-379.
- Plaut, D.C., and McClelland, J.L. (2010). Locating object knowledge in the brain: Comment on Bower’s (2009) attempt to revive the grandmother cell hypothesis. *Psychological Review*, 117, 284-288. doi:10.1037/a0017101
- Quian Quiroga, R., and Kreiman, G. (2010). Measuring sparseness in the brain: Comment on Bowers (2009). *Psychological Review*, 117, 291-297. doi:10.1037/a0016917
- Quian Quiroga in a “Leading Edge Previews” article in the same issue: Quiroga, R. Q. (2017). How Do We Recognize a Face? *Cell*, 169(6), 975-977.
- Rossion, B., & Taubert, J. (2017). Commentary: The Code for Facial Identity in the Primate Brain. *Frontiers in Human Neuroscience*, 11, 550.
- Thorpe, S. (1995). Localized versus distributed representations. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks*. Cambridge, MA: MIT Press