

The face detective

Doris Tsao unlocked the brain's code for recognizing faces. Now she wants to work out how we see everything else.

BY ALISON ABBOTT

Doris Tsao launched her career deciphering faces — but for a few weeks in September, she struggled to control the expression on her own. Tsao had just won a MacArthur Foundation ‘genius’ award, an honour that comes with more than half a million dollars to use however the recipient wants. But she was sworn to secrecy — even when the foundation sent a film crew to her laboratory at the California Institute of Technology (Caltech) in Pasadena. Thrilled and embarrassed at the same time, she had to invent an explanation, all while keeping her face in check.

It was her work on faces that won Tsao awards and acclaim. Last year, she cracked the code that the brain uses to recognize faces from a multitude of minuscule differences in shapes, distances between features, tones and textures. The simplicity of the coding surprised and impressed the neuroscience community.

“Her work has been transformative,” says Tom Mrsic-Flogel, director of the Sainsbury Wellcome Centre for Neural Circuits and Behaviour at University College London.

But Tsao doesn't want to be remembered just as the scientist who discovered the face code. It is a means to an end, she says, a good tool for approaching the question that really interests her: how does the brain build up a complete, coherent model of the world by filling in gaps in perception? “This idea has an elegant mathematical formulation,” she says, but it has been notoriously hard to put to the test. Tsao now has an idea of how to begin.

Her ambitions for unlocking some of the most recalcitrant mysteries of the mind are no surprise to neuroscientist Margaret Livingstone, who advised Tsao throughout her PhD at Harvard Medical School in Boston, Massachusetts. “Doris never got sidetracked,” she recalls. “She was quiet and focused, and always went for the big questions.”

Tsao grew up in a household filled with science. Her mother worked as a computer

programmer and her father was a machine-vision researcher. They emigrated to the United States from Changzhou, China, when Tsao was just four, “for a better life with more opportunities”, she says.

“My father is probably the key reason why I study vision, though I try to deny it,” Tsao says. Back when she was in high school, they discussed mathematical theories for how the brain might process aspects of vision. She found them “incredibly beautiful”, she says. “He helped plant in my head the idea that vision requires a profound explanation.”

She graduated in mathematics and biology at Caltech before joining Livingstone's team in 1996, where she initially studied the way the brain perceives depth of vision.

THE FACE CODE

Livingstone's lab works with macaques, which have a similar visual system and brain organization to those of humans. The view of the world through any primate's eyes is funnelled from the retina into the visual cortex, the various layers of which do the initial processing of incoming information. At first, it's little more than pixels of dark or bright colours, but within 100 milliseconds the information zaps through a network of brain areas for further processing to generate a consciously recognized, 3D landscape with numerous objects moving around in it.

During most of her PhD, Tsao was focused on the outermost layers of the visual cortex, where information from the retina first arrives. She learnt how to insert tiny electrodes — sensitive enough to record the firing of single brain cells — into this area of the monkeys' brains. But to help her probe deeper into the visual cortex, she decided to add brain imaging to her repertoire. The broader maps of brain activation provided by functional magnetic resonance imaging (fMRI) could help guide the


more-precise single-cell recording techniques. Few labs at the time were imaging the brains of animals, but Wim Vanduffel, a pioneer of monkey fMRI at the Catholic University of Leuven, Belgium, helped Tsao to establish the infrastructure needed to do the work in Boston.

While learning about the technique, she became aware of a surprising fMRI discovery made by neuroscientist Nancy Kanwisher from the nearby Massachusetts Institute of Technology. Kanwisher had identified a small area of the brain in humans that lights up whenever a person is shown a picture of a face, but not when they are shown pictures of other objects such as a house or a spoon.

Tsao reasoned that if the same face-recognition system existed in monkeys, she could use her sensitive electrodes to probe the neurons involved and work out how they function.

She teamed up with Winrich Freiwald, who was then a postdoc in Kanwisher's lab, and began a series of experiments combining fMRI with single-cell recording techniques to probe the inferior temporal (IT) cortex, the brain region that Kanwisher had identified. Over the next eight years or so, Freiwald, Tsao and their collaborators made a number of important discoveries¹⁻³. Passing picture after picture in front of the macaques, they mapped out the individual cells that fired in response to a human or monkey face. This allowed them to identify six patches on each side of the brain, distributed along the IT cortex. If the researchers electrically stimulated any one of the patches, the others lit up. Seeing those face patches working together in a network for the first time “was a joyful moment”, says Freiwald, who is now at Rockefeller University in New York City.

Freiwald and Tsao also discovered that the patches tended to be specialized. By showing monkeys a series of cartoon faces with various details such as hair, a nose or irises missing, they could determine which cells fire in response to specific facial features. A cell's rate of firing



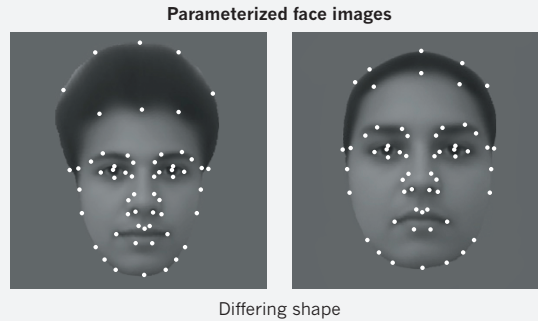
Doris Tsao has revealed fundamental aspects of visual systems by showing hundreds of human faces to monkeys.

DECODING THE FACE

To decode facial recognition in monkeys, Doris Tsao and Steven Le Chang recorded signals from neurons in brain areas called face patches, while showing the animals hundreds of pictures of human faces. Face patches are located in the inferior temporal (IT) cortex, which is involved in visual processing of objects.

SHAPE DIFFERENCES

Specific cells in a patch at the top of the IT cortex respond to a given set of shape features, such as the space between the eyes, the width of a mouth or the shape of the hairline. These cells tend to respond to faces that are in one orientation.




APPEARANCE DIFFERENCES

Cells in a patch found deeper in the IT cortex respond to appearance features, such as skin tone and texture. Some respond to the same faces regardless of orientation.



PROJECTION TO PREDICTION

The degree of difference between facial features dictates the firing rate of neurons that respond. When researchers put together the neural activity about face shape and appearance from just 205 neurons, they could predict the features of a face the monkey was looking at. New work suggests that this type of patch organization, together with firing-rate responses, could underlie the mechanism by which other objects are recognized in the IT cortex.



would ramp up according to how extreme the feature is, a property known as ramp-shaped tuning that turned out to be fundamental for face coding. A cell responding to the distance between two eyes, for example, might fire slowly in response to close-set eyes, but rapidly to ones set farther apart. When they showed monkeys real faces that were looking in different directions, the researchers discovered that cells in the patches closest to the visual cortex tended to fire in response to specific orientations of any face, whereas those in the deepest patch responded to a few individual faces, no matter what their orientations.

To investigate how the IT cortex might be encoding full faces from this information, Tsao realized that every face could be created by mixing the most important dimensions of 'faceness', such as how pointy a nose is, how eyes are set or complexion. She and her postdoc Steven Le

Chang identified the 50 dimensions that varied most across faces — 25 for shape and 25 for appearance — and created a set of 2,000 face images in which the value of all 50 dimensions was known⁴. They flashed these images in front of the monkeys while measuring responses from 205 neurons in two face patches. The code started to reveal itself.

Cells in the more superficial patch tended to be tuned to shape dimensions, whereas many of those located deeper in the IT cortex responded to appearance dimensions. This made sense because the deeper cells might have to account for distorted shape dimensions when a head is turned. Tsao and Chang could predict how the neurons would fire on the basis of the dimensions of any face, and they could even reconstruct a face just from the firing patterns of these cells (see 'Decoding the face').

The research seemed to point to a mechanism

by which individual cells in the cortex interpret increasingly complex visual information, until, at the deepest points, individual cells code for particular people.

That idea made intuitive sense. In 2005, Rodrigo Quian Quiroga, then a postdoc at Caltech, had identified what became known as Jennifer Aniston cells. By working with people who had had electrodes implanted in their brains to treat epileptic seizures, Quian Quiroga found signals from single neurons that responded to pictures of familiar or famous people. The cells also responded to any concept of that person. For example, one neuron fired in response to a photograph of the actor Jennifer Aniston, but also to her written name or even the title of a film she had starred in. These 'concept' cells resided in the hippocampus, which lies a little deeper in the brain than does the IT cortex⁵.

Tsao met Quian Quiroga, now at the University of Leicester, UK, in 2015 at a small meeting in Ascona, Switzerland, where she was presenting her latest results. Over dinner, he asked her how she thought her face cells related to his concept cells. "They are probably their precursors," she told him. But she fretted about her answer throughout the night. One thing had always bothered her. The deep IT cortical cells that she had been working on often fired in response to several individual faces — those that didn't look like each other at all.

Unable to sleep that night, she thought through the mathematical analysis that she and Chang had been applying to their data. Then a moment of insight struck. She had gone over the maths that so neatly described the ramp-shaped tuning responses of cells a million times. But in the dark, silent hotel room, she realized that it was the same as a mathematical operation that describes a type of projection. Projection explains, for example, how the Sun might cast the same shadow for two different objects depending on how they are positioned. If the cells are simply projecting combined dimensions from a multidimensional 'face space', she says, "it would explain why lots of different faces could elicit the same response in a face cell". The IT cortex is not homing in on one particular person at all; that transformation must happen at a point even deeper in the brain.

A CATEGORICAL CHANGE

At breakfast, she told Quian Quiroga about her new hunch and found that he had been thinking the same thing. So, she made an unusual wager: she bet him a bottle of expensive wine that it would be wrong, "because if it were true, I would be happy without wine".

Rushing back to the lab, she and Chang embarked on additional experiments that lost her the bottle, but culminated in the publication⁴ of the facial-recognition code in 2017.

The code was thrillingly — perhaps just a touch disappointingly — simple, says Tsao. That realization "was one of the happiest moments for me", she says.

There is a good chance that the same

simple code might apply over the whole of the IT cortex. Scientists have discovered other networks similar to the face-patch network that respond to other things, including bodies⁶, scenes⁷ and coloured objects⁸. But most of the IT cortex is uncharted territory. At a neuroscience meeting in Berlin this summer, Tsao presented some details of her current work. With her postdoc Pinglei Bao, she electrically stimulated cells in what she calls the no-man's land of the IT cortex, while scanning the monkey's brain. Two patches lit up, indicating another network — but this time she had no idea of its function.

To find out, she targeted the patches with her recording electrodes and monitored neuron activity as a monkey viewed pictures of 50 randomly chosen objects — from animals and vehicles to vegetables and houses — each from 24 different angles. The neurons did not respond to faces, but neither did the pattern of firing activity suggest that any other specific category of objects was associated with the network. Instead, the neurons seem to encode general properties of different objects. They seem to register, for example, whether something is spiky like a camera tripod or stubby like a USB stick; animate like a cat or inanimate like a house.

The way that this network processes information has remarkable parallels to how the face-patch network processes faces. Individual cells respond to elements of shape or character, with ramp-shaped tuning. A cell tuned to an object's animacy, for example, might fire slowly for a washing machine and rapidly for a cat. Cells in the more superficial patch tended to respond to similar categories of objects of similar orientation, whereas those in patches deepest in the IT cortex tended to respond to a handful of specific objects, whatever the angle. And Tsao and Bao were able to correctly predict the appearance of any object by looking at firing patterns from just 400 or so neurons.


"We think the entire IT cortex may be using the same organization into networks of connected patches, and the same code for all types of object recognition," says Tsao.

That's an idea that resonates with neuroscientist Georg Keller at the Friedrich Miescher Institute for Biomedical Research in Basel, Switzerland. "It gives hope that such a feature-based coding may operate widely in the brain," he says.

THE HALLUCINATING ENGINE

Now, however, Tsao wants to address the even bigger picture of how the brain captures the entirety of the world, rather than just how it decodes objects. This means understanding not just how visual and other sensory information flowing into the brain is processed, but also how high-level knowledge, which experience has embedded deep in the brain, affects perception. "Think about how we know that a blurry blob on a lake is likely to be a duck," she says.

The brain is not just a sequence of passive sieves fishing out faces, food or ducks, she says,



Combining brain imaging and electrophysiology has helped Doris Tsao to peer deep into the primate brain.

"but a hallucinating engine that is generating a version of reality based on the current best internal model of the world". Her ideas draw on Bayesian inference theory; only by combining perception with high-level knowledge can the brain arrive at the best possible understanding of reality, she says.

One possible mechanism is a long-debated theory called predictive processing, which is currently attracting interest among neuroscientists. Predictive processing holds that the brain operates by predicting how its immediate surroundings will change millisecond by

“Think about how we know that a blurry blob on a lake is likely to be a duck.”

millisecond, and comparing that prediction with the information it receives through the various senses. It uses any mismatch — ‘prediction error’ — to update its model of the world.

To find out what's going on, Tsao wants to learn how the hallucinating engine of the brain is wired. But unsure of which approach will work best, she's trying several simultaneously and recording from ever deeper parts of the brain.

One of her methods involves probing optical illusions, such as the famous face–vase picture. The brain automatically flips between the two perceptions after some seconds of staring at it. By recording single neurons as monkeys stare at the picture, Tsao is trying to identify where and how the flip happens in the brain,

and how it resets the internal representation of the world. Another method involves showing a monkey a picture of a familiar face, then morphing it into another familiar face, while recording in the brain. The primate's brain will automatically try to categorize a face as familiar, and at a precise point it will switch its perception of which of the two individuals it is seeing. "Ten years ago, no one would have known where to start investigating these phenomena because we didn't know where faces — or vases — were processed in the brain," says Tsao. Now that both location and code are known, "we can ask questions about exactly what changes as perception shifts".

The approach in non-human primates "has a lot of potential", says Keller, who studies predictive coding in the mouse visual cortex. Mice have a limited internal model of the world, he says, and it is unclear whether results from the mouse will be applicable to people. Although he and others can study predictive coding in the human brain using fMRI and electroencephalograms, such techniques will allow only a superficial inspection. "We won't be able to get at the mechanism, or how it is implemented, in the human like Doris will be able to."

Tsao continues to probe deeper into the brain in search of the sort of beautiful equations that her father inspired her with when she was young. She no longer has to hide her excitement, however. Now, it spreads across her entire face. ■

Alison Abbott is Nature's Senior European Correspondent.

1. Tsao, D. Y. *et al. Science* **311**, 670–674 (2006).
2. Moeller, S. *et al. Science* **320**, 1355–1359 (2008).
3. Freiwald, W. A. & Tsao, D. Y. *Science* **330**, 845–851 (2010).
4. Chang, L. & Tsao, D. Y. *Cell* **169**, 1013–1028 (2017).
5. Quiroga, R. *et al. Nature* **435**, 1102–1107 (2005).
6. Popivanov, I. D. *et al. J. Neurosci.* **34**, 95–111 (2014).
7. Kornblith, S. *et al. Neuron* **79**, 766–781 (2013).
8. Lafer-Sousa, R. & Conway, B. R. *Nature Neurosci.* **16**, 1870–1878 (2013).