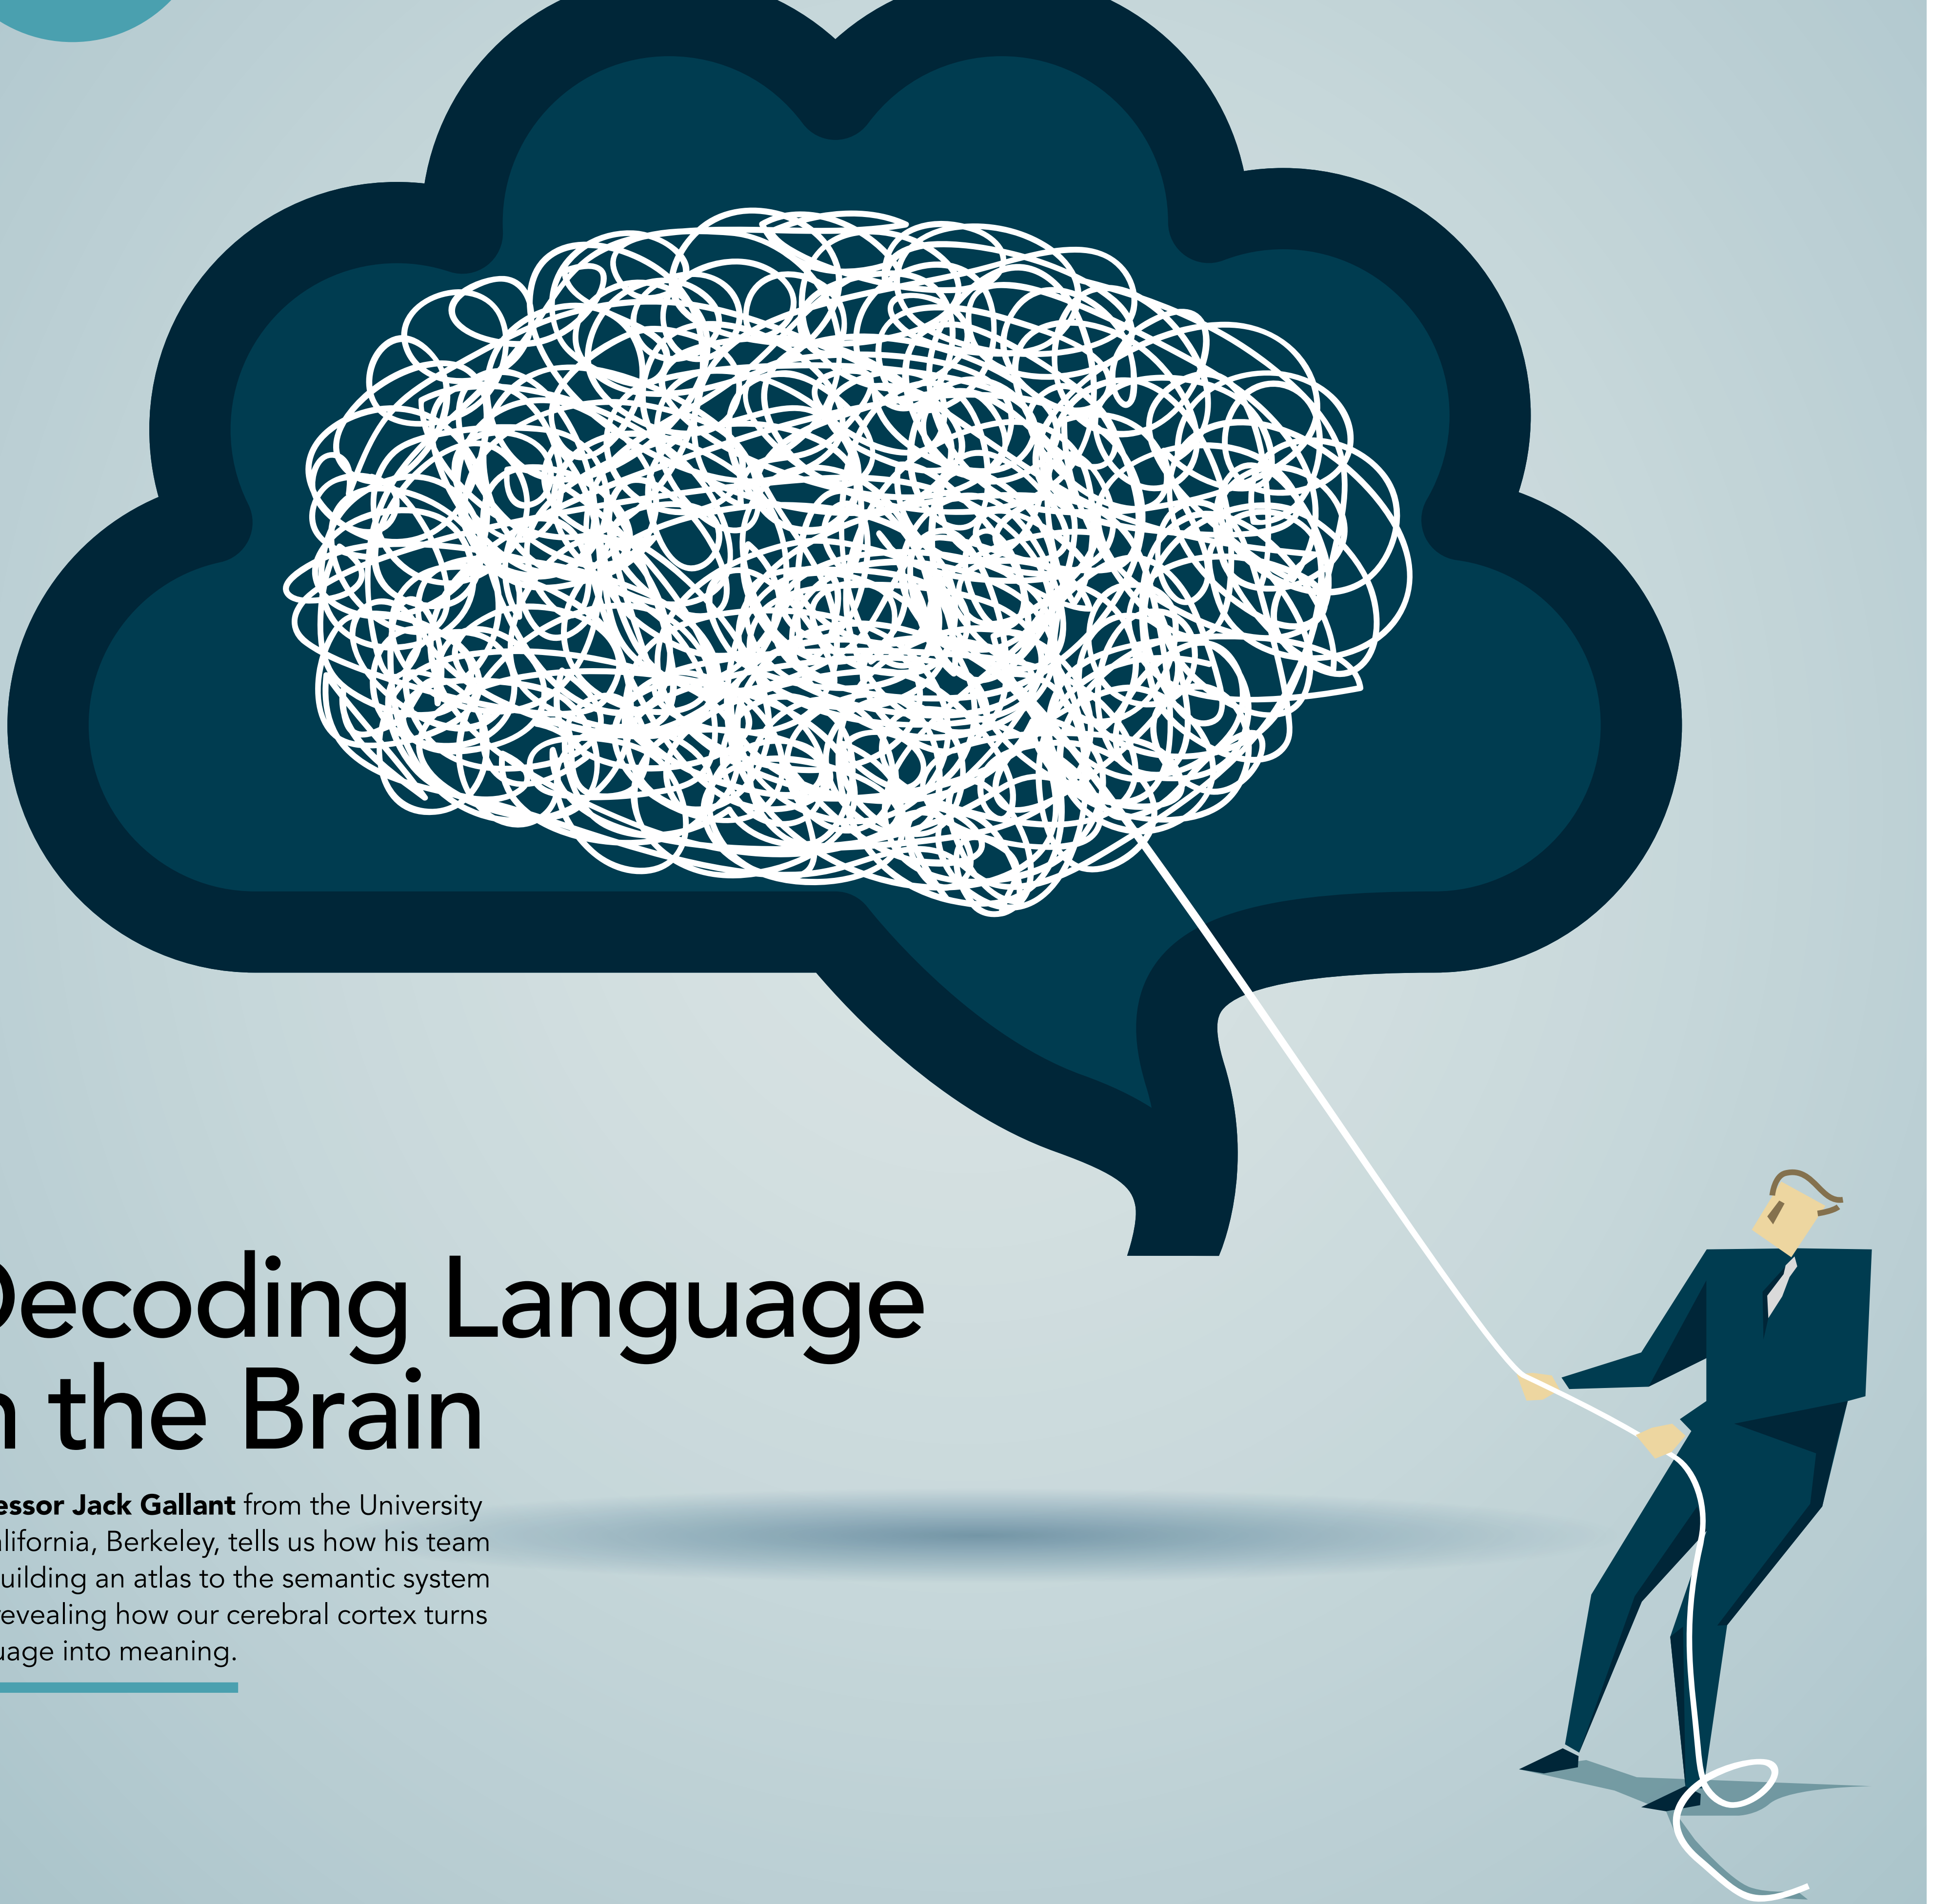

Decoding Language in the Brain

Professor Jack Gallant from the University of California, Berkeley, tells us how his team are building an atlas to the semantic system and revealing how our cerebral cortex turns language into meaning.

The human brain is one of the most complex biological machines on Earth. Despite many previous attempts to understand its workings, we may still have a long way to go. However, Professor Jack Gallant's team at the Gallant Lab, University of California, Berkeley, have taken a significant step forward into how the brain turns language into meaning. Developing new techniques in the process, the research team are attempting to map the semantic system, the regions of the brain that handle the complicated stream of data we call language.

The semantic system is distributed across much of the cerebral cortex. It is vital to the lives of modern human's, allowing us to communicate and understand our diverse thoughts, opinions and emotions. The term semantics is derived from the Ancient Greek word for 'significant', and studies the meaning we draw from incoming sensory information. The challenge for the researchers at the Gallant Lab is to understand the relationship between signifiers, such as words and phrases, and their denotation - what the words actually mean to us.

Q&A

Is it likely that this research will enable decoding of language in the minds of individuals with communication disorders such as ALS anytime soon?
The quality of any brain decoding device is inevitably limited by three factors: the quality of brain measurements, the quality of the model, and computer power. Of these three, the main limitation on brain decoding is measurement technology. Although functional MRI is powerful, it only recovers a very small fraction of the information that is potentially available in the brain at any given time. Currently, this isn't really sufficient to enable language decoding with high fidelity. However, when better, more portable brain measurement technologies become available, it will be possible to decode language from brain activity.

Are you encouraging other researchers to adopt your method of visualising the data on the surface of a 'flattened cortex'?
Although visualising data on flattened cortical maps is unusual in the fMRI community, this method has been long used in neuroanatomy. Flat maps provide a simple way to visualise the entire cortex simultaneously, though at first it can be difficult to think about the relationship between the 3D brain and the flat map.

How do you expect the semantic map to differ in someone who is multilingual?
This is an open question, and one that we are currently pursuing. Some have proposed that true bilinguals should have coincident maps of the two languages, while late second language learners will have different maps for the two languages. However, no study thus far has had sufficient spatial resolution to resolve this matter.

What are the limits of using fMRI analysis and how can brain activity measurement be improved for a finer resolution in the future?
fMRI measures changes in blood oxygenation, flow and volume that are caused (partly) by neural activity. For this reason, fMRI measurements provide only a very blurry view of brain activity. The next generation of fMRI should provide roughly a 100 fold improvement in resolution, but even then we will still be far below the resolution that we would like to have.

Do you think that this research could help lead to the ability to upload language to the brain one day?
In principle, brain writing is much more difficult than brain reading (i.e. decoding). A brain decoder does not act on the brain directly, so it does not affect the brain if the decoder is inaccurate. However, a brain writer would affect the brain, and any inaccuracies in the writing process would be uninterpretable to the brain. For this reason it will take much longer to develop the technology necessary to write into the brain than it will to read the brain.

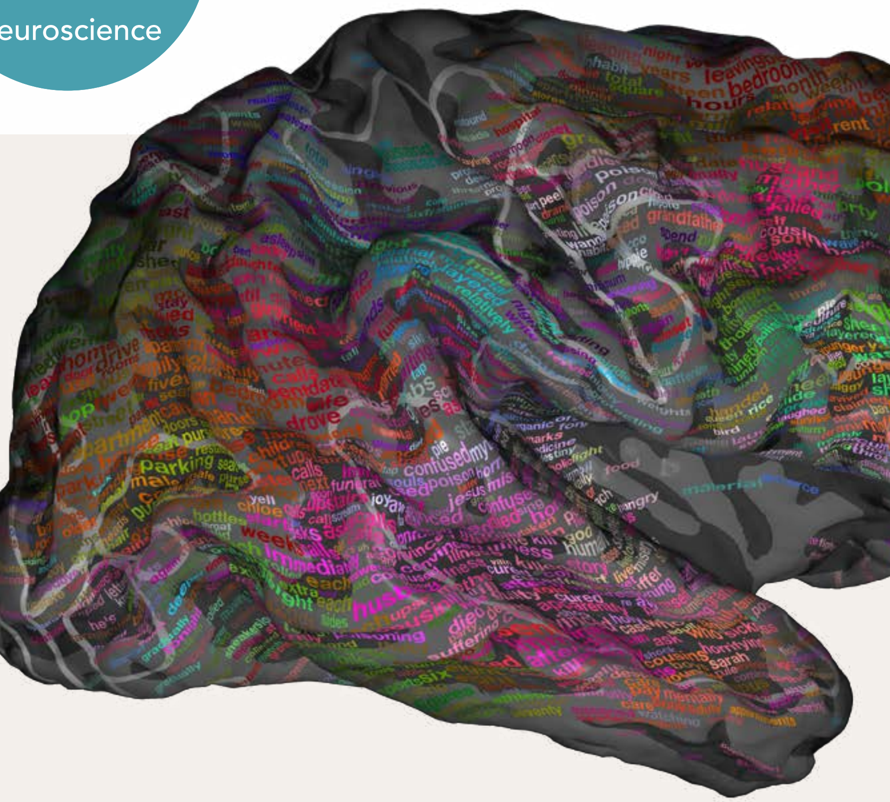

A lateral view of one person's right cerebral hemisphere. When heard in context, the overlaid words are predicted to evoke strong responses near the corresponding cortical location. The colour of each word indicates its semantic category. For example, green words are mostly visual and tactile concepts, while red words are mostly social concepts.

BUILDING A GIANT ATLAS

Until now, little has been discovered about the functional and anatomical organisation of the semantic system. In their recent study published in Nature, Alexander Huth, Jack Gallant and their colleagues set out to map the functional regions of the brain involved in semantic processing. "Our goal was to build a giant atlas that shows how one specific aspect of language is represented in the brain, in this case semantics, or the meanings of words," said Jack Gallant. This difficult task involved recording functional MRI (fMRI) data (see box 'What is fMRI?') of seven subjects' cerebral cortex while they listened to natural narrative speech. As the subjects listened to autobiographical stories from The Moth Radio Hour, a popular US radio show, the researchers recorded exactly which parts of the brain were activated by different words. This information was then used to build mathematical models that predict brain activity in response to the words heard. These models were shown to predict responses in the cerebral cortex relatively well by playing new stories to the subjects and comparing the predicted activity against the actual fMRI observations.

However, to visualise these complex models and identify the centres of different types of information was the next challenge. The 985-dimensional semantic models were reduced to the 3 most important dimensions, using principal component analysis to preserve as much information as possible.

The team then applied this information to a 3D model of the brain to generate the fascinating semantic map visualisations like the one shown above. To compare these results across subjects, the team developed a new analysis technique with the equally carefully crafted acronym PrAGMATiC. This Probabilistic And Generative Model of Areas Tiling the Cortex approach found similar functional areas in different subjects, allowing for anatomical variability in each subject.

The cerebral cortex is a sheet of neural tissue that wraps around the brain. Only a few millimeters thick, the cortex is highly folded so we can fit more in processing power. Unfortunately, these folds rather complicate the process of mapping the activity across the sheet. To overcome this problem the research team inflated the model cortex using a computational process, as shown opposite. Projecting the data onto this simplified brain model creates a much clearer visualisation. These flat maps are easier to interpret and the team suggest that the technique could be very useful for many other fMRI studies.

WHAT IS FMRI?

Neural processes are difficult to measure at high resolution. Functional Magnetic Resonance Imaging measures brain activity indirectly by detecting changes in blood flow, oxygenation and volume. The assumption is that fMRI signals are closely related to synaptic activity just upstream of the point of measurement.

Using fMRI technology has been validated in previous studies, which show that fMRI signals are closely correlated with local field potentials and local multi-unit activity - other measurements that are closely associated with activity of excitatory neurons. "That said, it is prudent to remain agnostic about the precise relationship between fMRI signals and neural signals" states Huth.

> **It's possible that this approach could be used to decode information about what words a person is hearing, reading, or possibly even thinking ”**

WORD ASSOCIATION

"Our study found that different individuals have remarkably similar semantic maps" comments Huth. For example, all seven of the subjects showed a particular area for words related to people. Interestingly, that area is surrounded by others selective for visual, tactile and number words. Groups of related concepts such as these suggest semantic associations are a basic organising principle for the brain and the meaning of language. A region that represents the word 'victim' also responds to 'killed', 'convicted' and 'murdered'. However, all of the subjects in this study are native English speakers. It remains to be discovered how these semantic maps change with different languages and cultures - one of the aims of the continuing research.

The colourfully tiled visualisations of these distinct areas for different information provide the first comprehensive view into how meaning is represented around the cerebral cortex. They show that language is processed across very broad regions of the brain, not limited to a few areas as previously thought. Furthermore, the images show that semantic activity is roughly as large and varied in both hemispheres of the brain. This may force some scientists to rethink their belief that language only involves the left hemisphere. However, this belief was inherited from studies of language production, not comprehension as studied here, leaving plenty of room for debate and further study.

DECODING LANGUAGE IN THE BRAIN

"It's possible that this approach could be used to decode information about what words a person is hearing, reading, or possibly even thinking from fMRI data" says Huth. The team are some way off a general purpose language decoder, however in the future it could make a huge difference to those suffering from communication disorders such as ALS or locked-in syndrome. Future studies may help to understand how the semantic system differs with language disorders like dyslexia and autism. The work could also help develop tools to study how people recover from brain injuries and strokes.

Currently, they are working on applying the method to map other kinds of language information, such as phonemic information, syntactic information and narrative. More research is also planned to verify the semantic atlas they have developed, hoping to test subjects listening to different stimuli, reading instead of listening and listening in a different language. Semantic maps could vary between individuals from different cultures and who speak other languages. The team also want to look at the links between the semantic system and other senses, for example, "what is the relationship between the representation of a visual scene containing a dog, and the representation of a story about a dog?"

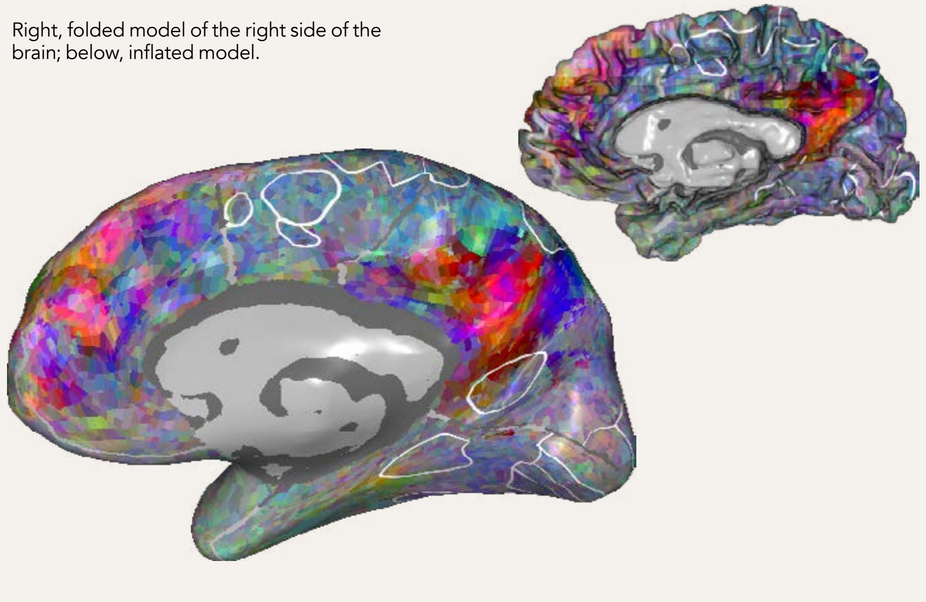

Right, folded model of the right side of the brain; below, inflated model.

3D view of one person's cerebral cortex. The colours indicate semantic selectivity, or which category of words it is selective for. For example, green areas mostly selective for visual and tactile concepts, while red areas are mostly selective social concepts. The inflated image demonstrates the approach to make the information easier to interpret by flattening the cerebral cortex.

Main illustration © iStock/ erhui1979

Detail

RESEARCH OBJECTIVES

- Use system identification to discover the computational processes of the brain
- Fit computational models from fMRI data to the brain to facilitate decoding
- Develop software optimised for visualising the results

FUNDING

National Science Foundation

PARTNERS

Caseforge Head Stabilisation

Pycortex Surface Visualisation

Neurotree

STRFLab

Neural Prediction Challenge

BIO

Professor Jack Gallant is the head of the cognitive, computational and systems neuroscience Gallant Lab at University of California, Berkeley. His team use functional MRI, computational modelling and machine learning to map perceptual, language and cognitive functions across the human brain.

CONTACT

Professor Jack Gallant
Department of Psychology
University of California, Berkeley
CA 94720-1650

E: gallant@berkeley.edu
T: (US) 510-642-5293
W: http://www.gallantlab.org/

@gallantlab
/UCBerkeley
/GallantLab