RESEARCHER INSIGHTS ISSUE NO.2

AI in Neuroscience

A report produced by Researcher

Welcome to this issue of 'Researcher Insights' – Researcher's special collection of articles, highlighting innovative ideas and leading voices from the research community.

Created with or by researchers, our carefully curated content will help you stay on top of new discoveries and developments in the field. This quarter, we are focusing on one of the hottest topics in neuroscience - artificial intelligence - and we are showcasing an even richer selection of contributors, with insights for clinicians, researchers, and beyond.

From our customised feeds, to our Live events and interviews, Researcher has always helped facilitate the conversation - but now we want to start it.

AI in Neuroscience

Editor:
Dr Kristine Lennie, Managing Editor - Researcher

Contributors:
Dr Karli Montague-Cardoso - Deputy Editor, Communications Biology
Roser Sanchez-Todo - Neurotwin
Dr Vijaya Kolachalama - Boston University Chobanian & Avedisian School of Medicine
Dr Benjamin Cowley - Cold Spring Harbor Laboratory
Dr Amy Kuceyeski - Cornell University
Dr Elishai Ezra Tsur - Open University of Israel
Dr Marco Capogrosso - University of Pittsburgh
Dr Weinan Sun - Janelia Research Campus

Design:
Alice Wilson

Contents

The Rise of AI - How technology is about to change the world of neuroscience

Dr Kristine Lennie

Over the last decade, artificial intelligence (AI) has shaped the way we access information, operate our homes, travel, enjoy entertainment, and even prove our identity.

We rely on AI technology more than we realise - when we asked just over 400 people (majority of which aged under 45) about their everyday habits, an astonishing (but unsurprising) 97.42% said they use their smartphone daily. 90.37% are on social media 'daily' or 'several times a week', while the percentage for digital voice assistants like Alexa, Siri and Google is 35.45%. Many also use face (50.47%) and gesture recognition (38.5%), photo filters (31.46%) and spellcheck (71.12%) - at the same frequency.

So, what if AI can also change the way we receive healthcare? Arguably, there are already some precedents for this. Of our 400+ respondents, 45.3% utilise fitness trackers 'daily' or 'several times a week', and the same is true for health tracking/monitoring apps (51.4%), smart scales (30.05%), smart rings for blood pressure (23.71%), or even smart posture correctors (19.25%). Another possible use is also on the horizon - with increase in life expectancy, the need for better treatments of age-related disorders is growing. When we asked our survey participants about their immediate family (grandparents, parents, siblings and spouses), 35.44% reported at least one of those relative was suffering from a neurological or neurodegenerative disease such as Alzheimer's; 36.39% from stroke, 26.29% 'some degree of paralysis', and nearly two-thirds (63.38%) an affective/mood disorder.

In this issue of Researcher Insights, we spotlight AI in neuroscience - from medical applications and tools, to understanding the brain. AI is already prevalent in research - when we asked 30 neuroscientists about their work, more than three quarters (76.6%) reported that they were either using AI (frequently or sometimes), or collaborated with groups that do (only 3.3% had no intention of using AI in their work). Deep neural networks (DNNs), the cornerstone of machine learning, are widely adopted to model complicated cognitive processes. These networks are trained on large amounts of data to make predictions, and in this issue of Researcher Insights, you will read about the extraordinary achievements of such algorithms - in some cases, matching the accuracy of humans to make decisions. You will also learn about the new directions that AI is heading in - the rise of spiking neural networks (SNN) which more closely mimic the function of the brain - as well as how AI can help us study brain function. You will hear from journal editor Dr Karli Montague-Cardoso on her perspectives of the field, and from Roser Sanchez-Todo - who talks about the ground-breaking Neurotwins project and its promise for patient care.

Among our 30 surveyed neuroscientists, 90% 'agreed' or 'strongly agreed' that the role of AI in healthcare will increase in the upcoming years. The top three directions where they anticipated growth were diagnosis (selected by 60% of the respondents), mass screening and disease prevention (56.7%), and public health monitoring (40%). As you'll see below, the current trends in neuroscience support these predictions as we head towards a new era of smarter, and faster, healthcare.

*The full data from both surveys can be found on pages 14-15.

Dr Kristine Lennie

Editor

AI in Neuro: Removing the fear of doctors and giving the gift of speech

by Dr Karli Montague-Cardoso

What if you had technology to generate speech simply by imagining the words? Or if you could take your child to the doctor without actually taking them to the doctor? Dr Karli Montague-Cardoso, Deputy Editor of Communications Biology and Consulting Editor for Communications Medicine shares her thoughts on cutting-edge science.

In recent years, the field of AI in neuro and brain–computer interfaces has developed rapidly, enabling the emergence of clinically valuable applications. As an editor, it is always exciting to receive a paper that makes a very clear and meaningful contribution to improving healthcare and people's quality of life. Something we see a lot nowadays is AI for tracking patient data in real-life settings. This is especially important for kids – since they tend to behave slightly differently when they're at the doctor's.

Earlier this year, Communications Medicine published a paper by Professor Sampsa Vanhatalo's lab (University of Helsinki) introducing a multi-sensor infant wearable device, which utilises deep learning for accurate quantification of postures and movement during spontaneous play. Essentially, this demonstrates that a less distressing 'out-of-hospital' assessment of infant motor ability and neurodevelopment is entirely plausible and can be employed for early intervention and individualised care strategies.

But, AI can be used for much more than observation. Towards the end of 2021, Communications Biology released an article that really excited me as both an editor and a neuroscientist. Although there are many different devices that aid the independence of patients with paralysis (for instance, robotic arms), what has been lacking is approaches which enable natural, verbal communication. Angrick et al. (2022) (Christian Herff's lab at Maastricht University) demonstrated real-time synthesis of imagined words in a 20-year-old female patient with severe epilepsy. The patient either whispered or imagined words while implanted with electrode shafts in various regions of the brain (predominantly the left frontal and lateral areas). A real-time decoding model was then used to analyse data from her intracranial recordings and generate audible speech.

This proof-of-principle case study makes an exciting step towards AI-aided speech neuroprosthesis. Aside from the obvious feat – converting imagined words into speech - the authors showed the feasibility of long-term, low-risk electrode implantation (without craniotomy). This makes the method more clinically viable even though, at present, the underlying model requires training with audible speech. For instance, in amyotrophic lateral sclerosis (ALS), where patients progressively lose their ability to speak, model training can be accomplished at the earlier stages of the disease, whilst the patient is still able to talk.

These papers, among many others published in our journals, show that the use of AI in neuroscience continues heading for something huge – I can't wait to see what that is!

Dr Karli Montague-Cardoso

Deputy Editor, Communications Biology
Consulting Editor, Communications Medicine

Key papers for this article:

Airaksinen, M., Gallen, A., Kivi, A., Vijayakrishnan, P., Häyrinen, T., Ilén, E., Räsänen, O., Haataja, L. & Vanhatalo, S. (2022). Intelligent wearable allows out-of-the-lab tracking of developing motor abilities in infants. *Communications Medicine*, 2(1), 1-14. https://doi.org/10.1038/s43856-022-00131-6

Angrick, M., Ottenhoff, M. C., Diener, L., Ivucic, D., Ivucic, G., Goulis, S., Saal, J., Colon, A. J., Wagner, L., Krusienski, D. J., Pieter, K. L., Schultz. T. & Herff, C. (2021). Real-time synthesis of imagined speech processes from minimally invasive recordings of neural activity. *Communications Biology*, 4(1), 1-10. https://doi.org/10.1038/s42003-021-02578-0

Nature Communications Medicine:

https://www.nature.com/commsmed

Nature Communications Biology:

https://www.nature.com/commsbio/

Neurotwins. Subject-specific data is used to personalize the computational models of the brain to provide optimised treatment using tES. We also use the data from mice under tES stimulation to validate our computational models.

Credit: © Neuroelectrics

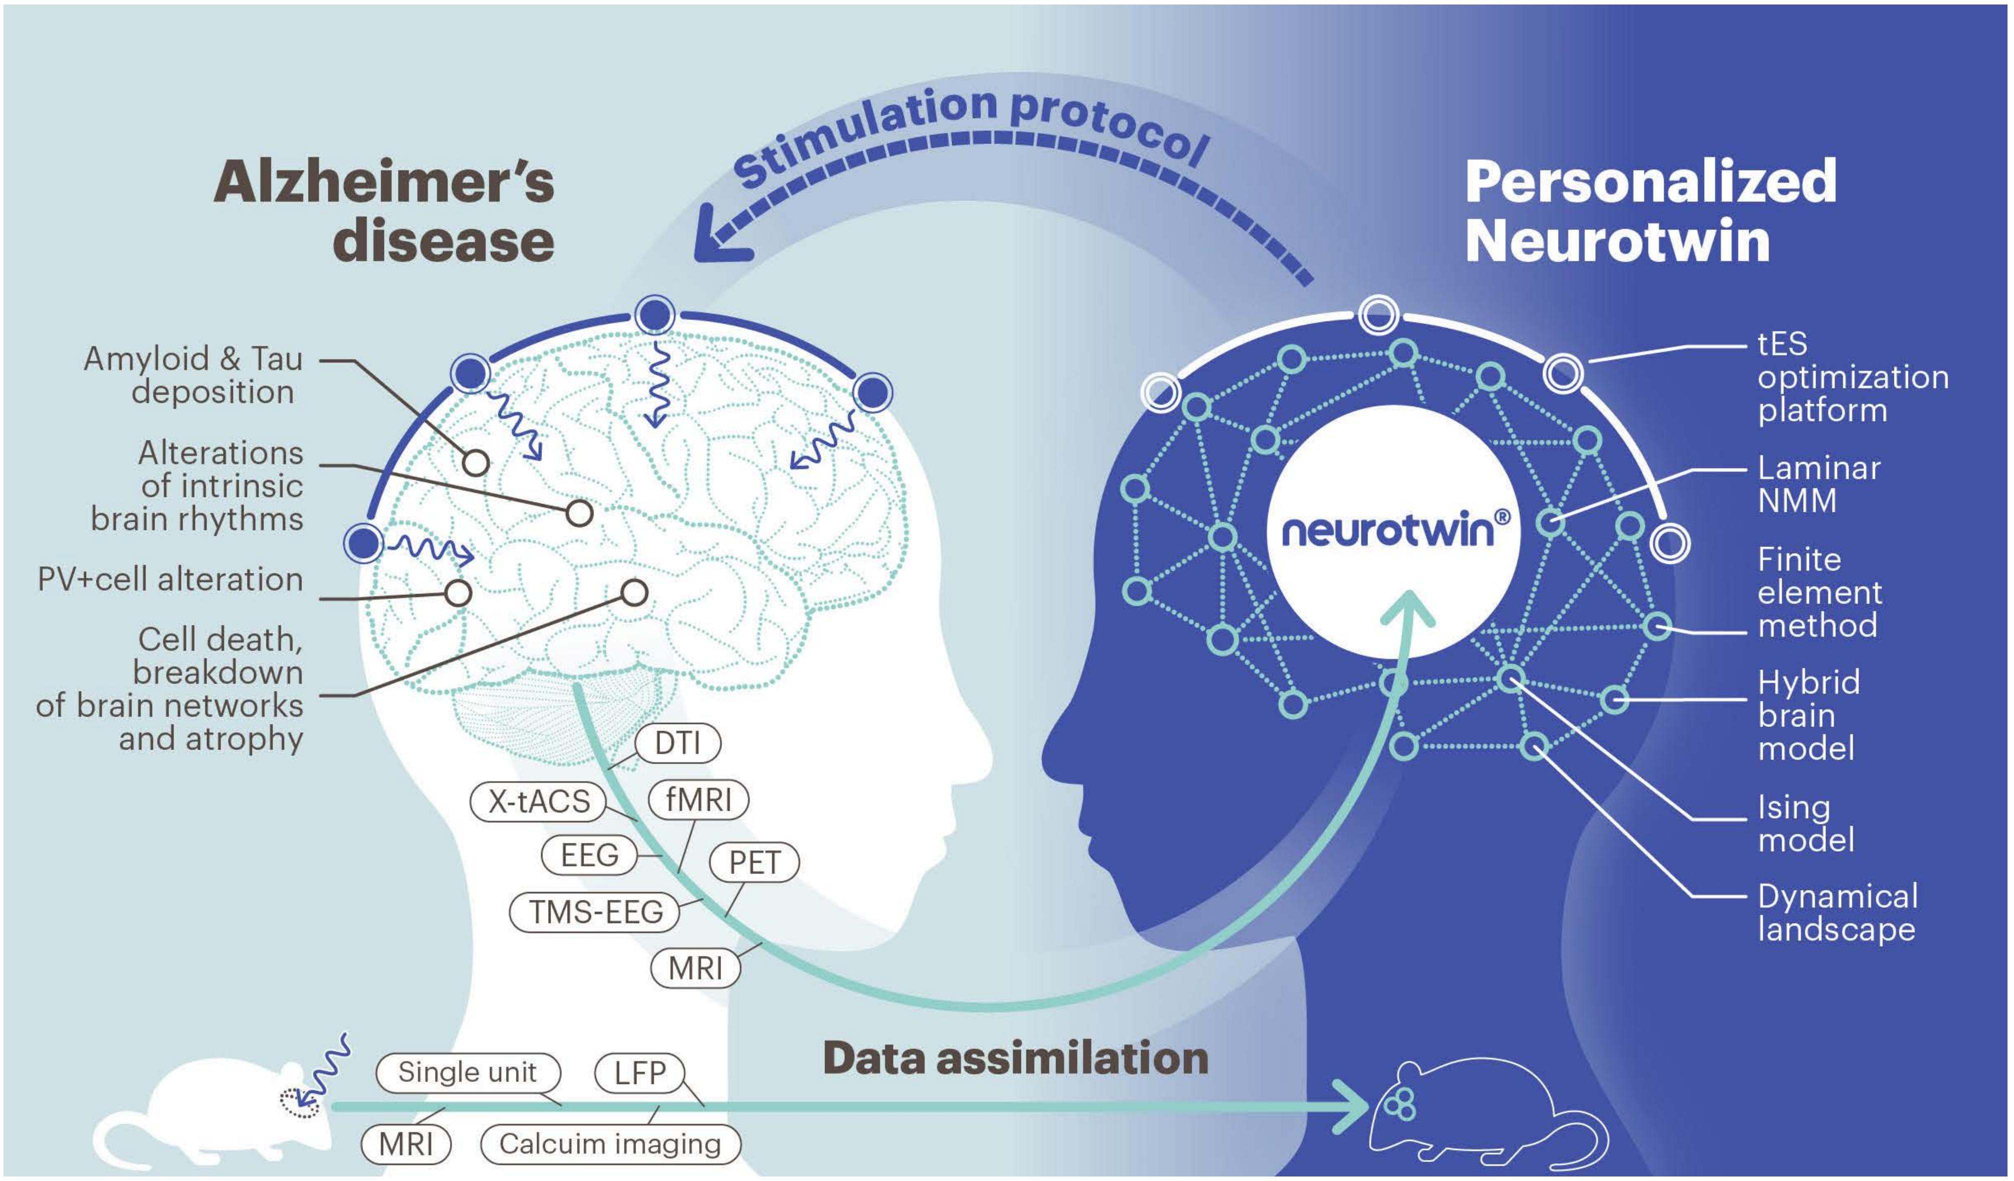

Neurotwins - Digital twins for non-invasive brain stimulation

Roser Sanchez-Todo, Neurotwin

During the last decade, an almost futuristic idea has been gaining momentum in translational medicine: using digital replicas of the human body as realistic proxies to 'beta test' diagnoses and treatment. A huge motivation for this has been the increasing prevalence of neurological diseases - the leading global cause of 'disability-adjusted life years', a measure of disease burden stemming from time lost to premature mortality and/or impaired health. Presently, there are no available cures for most neurological conditions, and challenges to existing therapies are amplified by individual differences between cases.

Neurotwin is a European-funded project set out to develop computational brain models that characterise personal pathology across different scales in order to predict the effects of non-invasive brain stimulation on patients. These personalised models - neurotwins - could be used to design optimal treatment protocols for Alzheimer's Disease (AD) - the most common type of dementia.

The goal of brain stimulation is to prompt neuronal populations to recover their healthy dynamics. Recent findings by Neurotwin's team and others suggest that techniques such as transcranial electrical current stimulation (tES) - specifically transcranial altering current stimulation (tACS) or transcranial magnetic stimulation (TMS) - may offer a valuable neurotherapeutic option for AD. However, substantial clinical leaps can only be made through a deeper understanding of the mechanisms underlying patient-specific brain networks. Anatomy, physiology, and pathology are known to be major sources of inter-individual variability in tES treatment response - something neurotwins could address. Our vision is that neurotwins can capture new insights into fundamental neuroscience, reduce uncertainty in diagnosis, and provide foundation for therapeutic breakthroughs. Our technology is built upon physiological models, such as the Neural Mass Models and anatomical models of the brain, e.g. Finite Element models. To personalise a neurotwin, we use patient-specific data: functional data (EEG, fMRI, TMS-EEG, PET) and anatomical data (MRI, dMRI) – but also tES stimulation data collected from healthy participants and non-human animal experiments. Neurotwin's team is intent on clinical translation – ensuring the project's successes reach patients – and to investigate the potential of Neurotwin-driven tES therapy in a controlled clinical pilot with AD patients.

Our dream is that, in the not-too-distant future, all patients will benefit from better health treatment stemming from the creation of their own neurotwins. Neuroimaging and model creation supported by cloud services will become the first step in diagnosis, optimising therapies in neurology and psychiatry. We anticipate, too, that the use of neurotwins will expand into medical practice well beyond the context of brain stimulation to include a wide range of conditions.

Roser Sanchez-Todo

Photo credit: Nacho Moragues

https://www.neurotwin.eu/

> Upward of 60 to 70% of all dementia cases are attributed to Alzheimer's disease – so addressing Alzheimer's will relieve much of the clinical and societal pressure.

AI for Dementia: A novel tool integrating multiple modalities to assess cognitive state

with Dr Vijaya Kolachalama, Boston University Chobanian & Avedisian School of Medicine

It's an undoubted accomplishment of science that today we live longer than ever before. Improved healthcare, sanitation and immunisation promise the average person far more than half a century on this planet. But what is the cost of our boosted life expectancy?

According to the World Health Organisation (WHO), between 2020 and 2050, the number of people over 60 will more than double. With this comes an increase in the prevalence of dementia - the grouping term for the progressive loss of cognitive function most associated with advanced age. "The burden on the healthcare system due to neurodegenerative diseases is huge and set to keep rising", says Dr Vijaya Kolachalama, Associate Professor at Boston University. "Within this, upward of 60 to 70% of all dementia cases are attributed to Alzheimer's disease – so addressing Alzheimer's will relieve much of the clinical and societal pressure. At present, early detection is our best available management option. So, this is what our lab is focused on: building a new diagnostic tool for clinicians." Dr Kolachalama shares that his group spends a significant amount of time talking to medical professionals to learn how they work. "Neurologists examine different types of data - MRI scans, patient history, demographic information, neuropsychological or functional assessment, to determine cognitive status. We wanted to build a machine learning framework that could combine information from different modalities, similarly to a clinician, and determine if the patient has normal cognition, mild cognitive impairment, Alzheimer's or dementia due to other reasons."

Dr Kolachalama explains that to mimic the analysis of a medical expert (primary, secondary and differential diagnosis), his group's algorithm needed to include multiple layers of evaluation. "We developed a fusion model consisting of a convolutional neural network (CNN) for assessing image data, and a traditional classifier, to integrate different types of information: 3D, volumetric image data from MRI scans, together with non-imaging data like demographics, history, functional assessments and cognitive assessment test results." It took Dr Kolachalama's group about two years to train and validate the model on data from 8 diverse study cohorts. "Then, we wanted to see if our model could match the clinical standard outside of the training dataset. For this, we performed three independent experiments.

"First, we recruited 17 neurologists and gave them about 100 cases each to assess. Next, we asked a group of 7 neuroradiologists to examine 50 MRI scans each and identify sites of neurodegenerative changes. Finally, we collected post-mortem data from around 100 brains to examine their patterns of neuropathological lesions. We ran the model on these same sets of data, essentially generating a head-to-head comparison between the diagnostic performance based on the reference standard, and our model."

The results of the experiments were promising. The model performed on par with the medical professionals, both in terms of cognitive assessments, and ability to identify patterns of disease-related damage in the brain. This indicates a possibility to build a tool for faster diagnosis, that could reduce the burden on the healthcare system and improve patient's quality of life.

"We can't give the model to a clinician yet," says Dr Kolachalama. "We need to package it, make the software secure. Tests need to be run in a clinical setting, in real time. These are our next steps now, and we are very excited."

Key papers for this article:

Qiu, S., Joshi, P. S., Miller, M. I., Xue, C., Zhou, X., Karjadi, C., Chang, G. H., Joshi, A. S., Dwyer, B., Zhu, S., Kaku, M., Zhou, Y., Alderazi, Y. J., Swaminathan, A., Kedar, S., Saint-Hilaire, M., Auerbach, S. H., Yuan, J., Sartor, E. A., Au, R. & Kolachalama, V. B. (2020). Development and validation of an interpretable deep learning framework for Alzheimer's disease classification. *Brain*, 143(6), 1920-1933. https://doi.org/10.1093/brain/awaa137

Qiu, S., Miller, M. I., Joshi, P. S., Lee, J. C., Xue, C., Ni, Y., Wang, Y., De Anda-Duran, I., Hwang, P. H., Cramer, J. A., Dwyer, B. C., Hao, H., Kaku, M. C., Kedar, S., Lee, P. H., Mian, A. Z., Murman, D. L., O'Shea, S., Paul, A. B., Saint-Hilaire, M., Sartor, A. E., Saxena, A. R., Shih, L. C., Small, J. E., Smith, M. J., Swaminathan, A., Takahashi, C. E., Taraschenko, O., You, H., Yuan, J., Zhou, Y., Zhu, S., Alosco, M. L., Mez, J., Stein, T. D., Poston, K. L., Au, R. & Kolachalama, V. B. (2022). Multimodal deep learning for Alzheimer's disease dementia assessment. *Nature Communications*, 13(1), 1-17. https://doi.org/10.1038/s41467-022-31037-5

Dr Vijaya Kolachalama

Photo credit: Boston University Chobanian & Avedisian School of Medicine

Many of us rely on vision to navigate our everyday lives. The visual stimuli we receive from the environment are interpreted in the brain through electrical signals that photoreceptors in the retina generate in response to light. The resulting brain activity patterns can, in part, explain how the brain 'sees', but the exact transformations that allow us to recognise faces and objects, and move through space, are still elusive. Moreover, though we can record brain responses while viewing images, our ability to modify these images to achieve a specific pattern of activity is still limited.

Could AI hold the answer?

Gaudy Images: A faster way to train DNNs for the visual cortex

with Dr Benjamin Cowley, Cold Spring Harbor Laboratory

Individual neuronal responses (spikes) in the visual system can be represented by computational frameworks that map selected images to activity in the brain. Some of the first-generation models designed for this were Generalised Linear Models (GLMs) which made assessments by testing images for the presence of a single specific feature. The visual system, however, likely operates using more complex, but poorly defined hierarchical processing algorithms. This makes deep neural networks (DNNs) a better candidate for modelling – but one that requires large swathes of training data.

"There are some highly predictive models for neuronal spikes already available for the primary visual cortex (V1)," says Dr Benjamin Cowley, faculty at Cold Spring Harbour Laboratory. "But V1 is the simplest and best-studied region of the visual cortex. For more sophisticated visual areas, such as V4 (a mid-level visual cortical area) and the inferior temporal cortex (IT), we don't have enough data. A workaround is to rely on DNNs pre-trained on object recognition, where we can linearly map the activity of a middle layer of the DNN to the neural responses. However, because this DNN was never trained on real neural data, we may be missing important features. For best prediction, we would ideally train an end-to-end model whose sole purpose is predicting neural responses."

The multiple layers of a DNN can be programmed to detect different aspects of vision, such as textures, colours, shapes, and spatial orientation. Following training, a highly accurate model could then be used to investigate the underlying 'decision-making' procedure: the model's layers can act as 'proxies' for real neurons, and the framework can be interrogated to provide insight into the step-by-step computations of the brain. "So, how do we build such a model without much data?" Dr Cowley asks. "One way is a procedure called 'closed-loop active learning'."

Active learning is based on iteratively allowing the model to select its own training data. For a visual model, a small collection of images would first be shown to subjects and the neural brain responses recorded in a laboratory setting. These data would be fed to the model, and once the training is finished, a subsequent large set of new images (without the responses) would be presented for assessment. The model would select a set of images from this new dataset that maximise uncertainty, and those would be sent out to the lab to again collect responses. This is repeated until a sufficiently accurate model is obtained.

"So, how can we improve this process?" asks Dr Cowley. "We tested many different active learning algorithms and all of them preferred the same type of images: high-contrast edges, garish colours, bold geometric patterns. We termed this class of images 'gaudy images' for their over-the-top colours and contrasts – and they turned out to train DNN models with even less data than active learning algorithms."

Dr Cowley's gaudy images are natural images where the pixel intensity of each colour channel (red, green or blue) is pushed to its maximum or minimal value: i.e. 0 if it's below and 255 if it's above the mean pixel intensity of the image.

He tested training with gaudy images against active learning in a GLM and a DNN. "In both, gaudy images outperform close-loop active learning,' says Dr Cowley. "Our results hold for different activation functions, across different DNN architectures and layer complexity, as well as against DNNs that are pre-trained on object recognition."

With these gaudy images, Dr Cowley was able to create a model of visual cortical activity – while providing a new strategy for faster and more successful training of DNNs. "We can now examine questions we couldn't a decade ago: what image features do higher-order visual cortical neurons prefer? How large do our models need to be? What computations are necessary and sufficient to predict neural responses? This will move us closer to understanding the step-by-step computations of the brain. But it'll also be exciting to find 'gaudy images' for other sensory systems. Are there gaudy sounds? Smells? Tastes?"

Dr Benjamin Cowley

Key paper for this article:

Cowley, B., & Pillow, J. W. (2020). High-contrast "gaudy" images improve the training of deep neural network models of visual cortex. *Advances in Neural Information Processing Systems*, 33, 21591-21603.

> NeuroGen could potentially allow clinicians to stimulate connections in patient's brains to change the underlying connectivity and possibly aid recovery from disease or injury.

NeuroGen: A new platform for discovery neuroscience

with Dr Amy Kuceyeski, Cornell University

"I went to a talk by Dr Leila Wehbe a few years ago about a model that could predict brain activity patterns from a presentation of words," thoughtfully recalls Dr Amy Kuceyeski, Cornell University. "I turned to my colleague, and I said, wouldn't it be great if we could find a way to create images that can stimulate any brain region of our choosing?"

Unlike Dr Cowley, Dr Kuceyeski's focus is on macro-scale brain regions rather than individual neurons, in *in vivo* human imaging instead of animal models. Recently, she has been looking at neuromodulation – technology aimed at altering neural activity patterns.

"There's been a rise in the number of publications examining how synthetic images – meaning, artificial images created by a computer generator – can be used to affect brain responses. Researchers have shown that in monkeys, certain synthetic images can achieve higher neuron firing rates than natural images." Dr Kuceyeski references a 2019 paper by Bashidan et al. and another by Ponce et al. from the same year (see references page). "But some electrophysiological measures can't be taken in humans like they can in monkeys – for humans, the main tool is non-invasive, fMRI. So, to our knowledge, we were the first to work with synthetic images to manipulate human brain responses."

Typical stimulus-response data from experiments links images shown to an individual in an MRI scanner to that person's corresponding brain responses. Producing a neural network that predicts the output (the brain activity measured using functional MRI) from an image (the input) is known as an encoding model. To achieve the opposite and produce synthetic images able to activate specific brain regions, Dr Kuceyeski's lab first needed to create such an encoding model –and then 'reverse-engineer' the process.

Several scientific resources were used for this project. The first one was the Natural Scenes Dataset (NSD), a collection of over 20,000 image-fMRI scan pairs obtained from 8 individuals. The second was a pre-existing encoding model by St-Yvest and Naselaris (2018) known as a 'deepnet feature-weighted receptive field encoding model' (fwRF) that can predict the image features that produce a given brain response pattern. Next was AlexNet, an image classifier trained on the ImageNet database of 14mil annotated images organised in over 1,000 categories, and finally BigGAN-deep – a conditional image generator also trained on ImageNet and able to create images based on a pre-defined category from ImageNet's list.

"We trained our encoding model on the Natural Scenes Dataset," explains Dr Kuceyeski. "We took the layers (the feature extractor part) of the pre-trained classifier AlexNet and fed the Natural Scenes Dataset images through that. The extracted features were then run through the St-Yvest and Naselaris (2018)'s regression model, whose aim was to predict as output one of 24 visual regions in the brain. Once we validated our encoding model against fMRI data, we paired it with the BigGAN-deep synthetic image generator, which produces a synthetic image based on a noise vector and a specified category."

Upon selecting a target brain region to be activated, BigGAN-deep produces 100 images for each of ImageNet's 1,000+ categories. Each of those sets of 100 images is fed to the encoding model to predict 100 corresponding brain responses and calculate an average (with respect to the target region). The top 10 categories with the highest average are then selected, and the noise vector optimised again to return the best image in each. "We showed that NeuroGen can produce images that outperform natural images in the intensity of predicted responses," adds Dr Kuceyeski. "We then had a discovery tool for neuroscientists that could create images for a desired brain response. So, we next focused on personalisation. We collected image-fMRI scan data in prospective individuals, then fine-tuned our encoding model for each person. We showed that NeuroGen can be customised for a specific individual."

These results prove that Dr Kuceyeski's framework can be used to enrich training datasets for fMRI experiments in a controlled way. But beyond that, in the future, NeuroGen could potentially allow clinicians to stimulate connections in patient's brains to change the underlying connectivity and possibly aid recovery from disease or injury.

"We are looking at some decoding work now, how to reconstruct an image a person is looking at from brain activity patterns," concludes Dr Kuceyeski. "Work at the intersection of biological and artificial neural networks is getting very exciting."

Key paper for this article:

Gu, Z., Jamison, K. W., Khosla, M., Allen, E. J., Wu, Y., Naselaris, T., Kay, K., Sabuncu, M. R., & Kuceyeski, A. (2022). NeuroGen: activation optimized image synthesis for discovery neuroscience. *NeuroImage*, 247, 118812. https://doi.org/10.1016/j.neuroimage.2021.118812

Dr Amy Kuceyeski

Photo credit: Weill Cornell Medicine

The Role of AI Today

To what extent is AI already integral to our everyday lives – be it for work, leisure or staying healthy?

This quarter, we collected answers from 426 respondents to find out about our audience's habits and relationship with AI (and we also asked about their health).

HOW OLD ARE YOU?

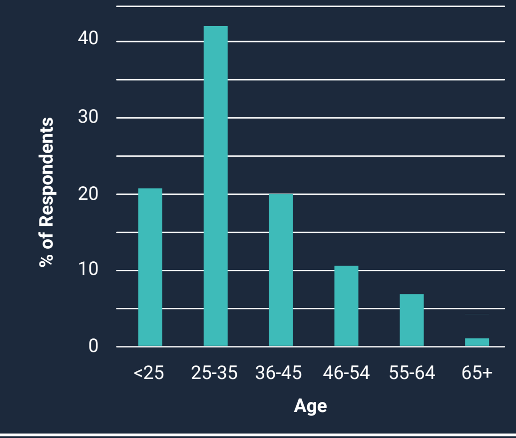

HOW OFTEN DO YOU USE THE FOLLOWING?

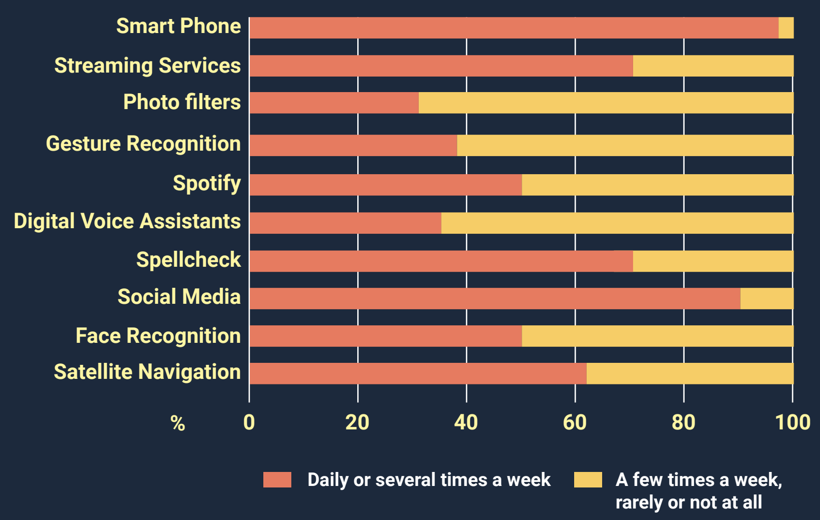

HOW OFTEN DO YOU USE THE FOLLOWING?

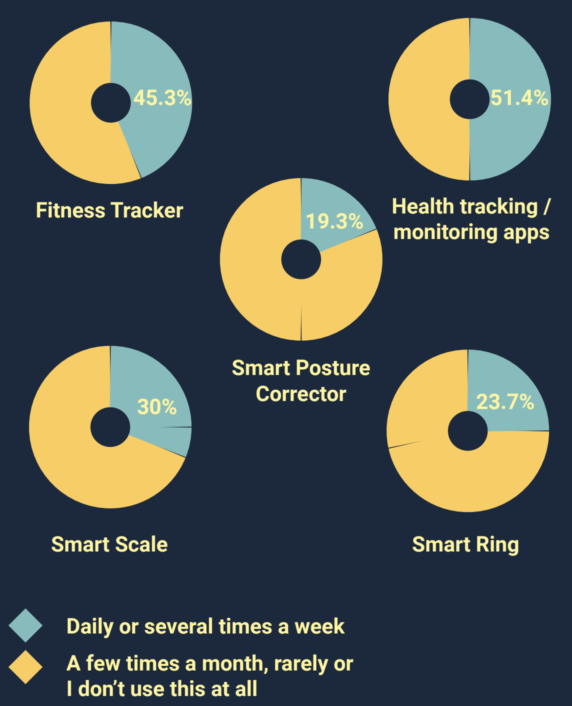

PERCENTAGE OF RESPONDENTS WHERE AT LEAST ONE FAMILY MEMBER WITHIN IMMEDIATE FAMILY (GRANDPARENTS, PARENTS, SIBLINGS, SPOUSE AND CHILDREN) SUFFER FROM:

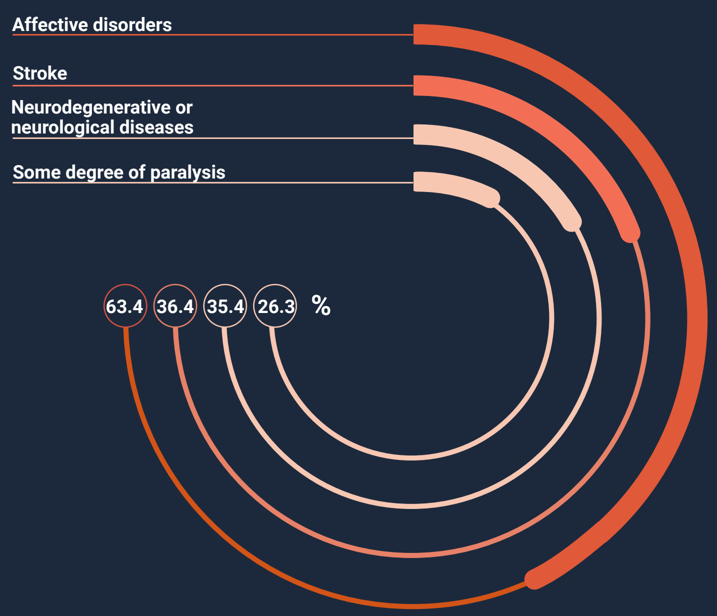

Neuroscience is already embracing AI as an essential tool for studying the brain and its function.

Here, we showcase survey responses from 30 academics from the field, highlighting the current status quo, emerging trends, and future directions.

DO YOU USE AI IN YOUR RESEARCH?

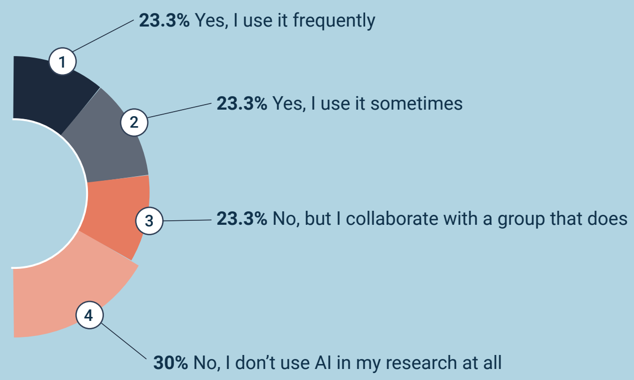

WHAT IS YOUR OCCUPATIONAL LEVEL?

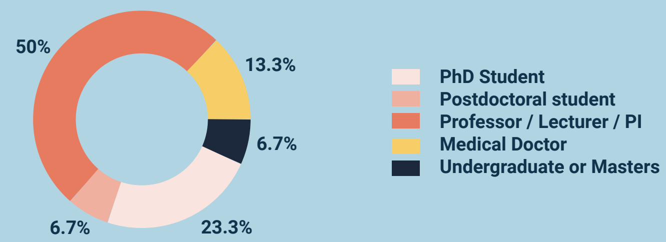

DO YOU INTEND TO USE AI IN RESEARCH IN THE NEAR FUTURE?

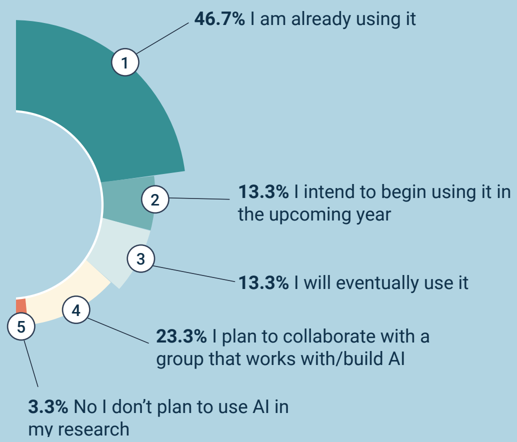

TO WHAT EXTENT DO YOU AGREE WITH THE FOLLOWING STATEMENT:

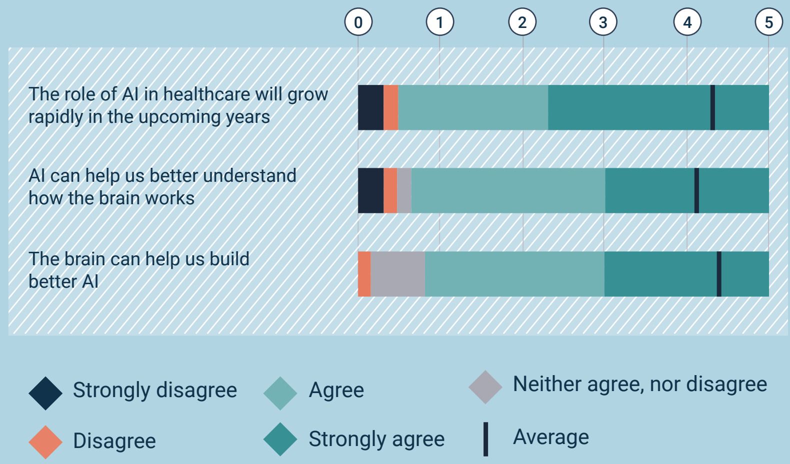

WHAT ARE THE TOP 5 DIRECTIONS IN AI IN HEALTHCARE THAT WILL SEE GROWTH IN THE UPCOMING DECADE?

1. Diagnosis
2. AI for mass screening & disease prevention
3. AI for public health monitoring
4. AI for patient data tracking
5. AI for speech comprehension

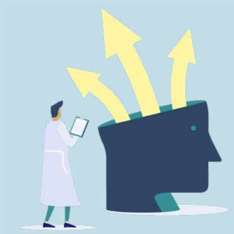

WHICH FORM OF NEW INITIATIVES DO YOU CONSIDER MOST HELPFUL FOR SPEEDING UP PROGRESS IN AI IN NEUROSCIENCE RESEARCH?

1. New inter-departmental collaborations (between computational/engineering and neuroscience)
2. More governmental funding for AI in neuroscience projects
3. Workshop/training programmes introducing AI to neuroscientists or vice versa
4. More institutes and programmes focused specifically on AI in neuroscience

Developing a Brain-inspired Robotic Arm for Wheelchair Use

with Dr Elishai Ezra Tsur, Open University of Israel

"I often visit a children's rehabilitation hospital, where many of the kids have mobility impairments," says Dr Elishai Ezra Tsur, head of the Neuro-Biomorphic Engineering Lab (NBEL-lab.com) at the Open University of Israel. "One day I asked a 6-year-old girl, 'What would you like me to build for you to make your life easier?'. She said, 'My dream is to be able to comb my hair'." Dr Tsur mimics the action with his hand, emphasising the wrist rotation. "If you think about it, the task of being able to comb your hair is much more complicated than programming a robot to, say, manufacture a car."

Creating adaptive movement is still one of the most challenging tasks in robotics – and with current technology, an energetically costly one to resolve. Dr

> The task of being able to comb your hair is much more complicated than programming a robot to, say, manufacture a car.

Tsur believes the answer is neuromorphic engineering – a point of convergence between biology, AI, and machine building, which utilises ideas from neuroscience to enhance technological advancements.

"Robots run on microprocessors that have a von Neumann architecture – they fetch programming instructions and execute them using register-based memory and arithmetic logic units. A typical robot has those microprocessors, continuously multiplying large matrices to calculate movement trajectories," Dr Tsur gestures towards his head for comparison. "The brain does not resolve systems of linear equations to create motion. Instead, it uses spiking neural networks, comprising of neurons that 'communicate' with each other with electrical impulses."

Neuronal spikes (as mentioned in Dr Cowley's article above) are the communication system of the neurons. A spike occurs when the membrane potential (the difference in electrical charge between the interior and exterior of the cell) reaches a certain threshold causing the neuron to 'fire'. It was this type of behavior that Dr Tsur and his colleagues wanted to replicate with their tech.

"Our neuromorphic computers don't have a CPU or centralised register-based memory chips," he remarks. "Our computing hardware is comprised of highly connected electrical neurons that don't operate by simultaneous timesteps, but asynchronously send spikes to each other - just like the brain does. This is called a spiking neural network (SNN). We and other groups have shown, that modelling adaptive control systems (i.e. systems where the parameters are constantly changing in response to the environment) with SNNs, conserves energy and improves performance."

Wheelchairs are one type of equipment where battery life is essential. Wheelchair-mounted robotic arms, which also rely on power, are manufactured by companies but are incredibly expensive to produce and therefore inaccessible to patients. "So, using our technology, we set off to build a state-of-the-art, energy-efficient robotic arm from affordable components," says Dr Tsur. "And we succeeded."

The robotic arm's adaptive motor control uses accelerometers – non-expensive sensors that track changes in direction and speed (e.g. shaking, rotating, tilting) – an abundant technology that features in many devices, including smartphones. The acceleration readings are integrated using a SNN so that the spatial location of the arm is being measured in real time for faster, energy-efficient error detection and minimisation.

"Instead of using huge amounts of training data, feedback from the environment reinforms the machine continuously. If the arm is trying to move distance x to a target but only moves distance x-5 due to, for example, heavy load or a weak engine, the accelerometers will detect that the target was missed, the information will be fed into the computational framework and adjustments will be made to reach the target," clarifies Dr Tsur.

Computations are conducted by Intel's neuromorphic chip Loihi (up to 1000 times more energy efficient and 100 times faster than a CPU), encoded with the corresponding Nengo library. For the interested reader, Nengo is an open-source Python package designed for building neural networks, including SNNs. It's based on the Neural Engineering Framework (NEF) theoretical approach which builds, from single neurons with predefined properties and inter-neuronal interactions, large-scale cognitive models.

The current version of the robotic arm can help users drink from a cup, eat, pick up and move items. "To our knowledge, this is the first time that someone has been able to show clinically relevant application to neuromorphic engineering," concludes Dr Tsur, whose interest in the field extends beyond this project as evidenced by his recent book 'Neuromorphic Engineering: The Scientist's, Algorithms Designer's and Computer Architect's Perspectives on Brain-Inspired Computing' – a must read for enthusiasts in robotics and computer architectures.

> To our knowledge, this is the first time that someone has been able to show clinically relevant application to neuromorphic engineering.

Dr Elishai Ezra Tsur

Elishai Ezra Tsur's book: Tsur, E. E. (2021). Neuromorphic Engineering: The Scientist's, Algorithm Designer's, and Computer Architect's Perspectives on Brain-Inspired Computing. CRC Press.

Key paper for this article:

Ehrlich, M., Zaidel, Y., Weiss, P. L., Yekel, A. M., Gefen, N., Supic, L., & Tsur, E. E. (2022). Adaptive control of a wheelchair mounted robotic arm with neuromorphically integrated velocity readings and online-learning. *Frontiers in Neuroscience*, 16, 1007736-1007736. https://doi.org/10.3389/fnins.2022.1007736

Spinal Cord Stimulation as a New Tool for Managing Paralysis

with Dr Marco Capogrosso,
University of Pittsburgh

Long-lasting motor deficit is a growing clinical and economic concern. According to a 2013 study, over 5 million people in the USA alone live with some form of paralysis, with 33% of those cases induced by stroke - the leading cause of the condition. Only around 15% of hospitalised stroke patients recover completely from paralysis while a recent review from 2021 (spanning the period 2002-2017), estimated that the collective cost of stroke in the USA exceeds $100 billion per year, of which 2/3rds attributed to indirect burdens such as unemployment (see references page). In all this, the landscape of treatment and assistance options for patients with paralysis is limited– fortunately, this might be about to change.

A lab at the University of Pittsburgh has been working on ways to artificially elicit and fine-tune neuroactivity in the spinal cord to restore movement. Dr Marco Capogrosso, director of the Spinal Cord Stimulation Laboratory and member of the Rehab and Neural Engineering Labs is the head of the initiative and a co-founder of Reach Neuro, a start-up for helping stroke-caused motor deficits.

"In paralysis, there is a problem with the communication between the brain and the cells in the spinal cord that control muscles. In most cases, the physical infrastructure downstream is actually preserved – for example, in stroke, the damage is in the brain but the spinal cord, where the cells that control movement are located, is fully intact. Similarly, in spinal cord injury, everything below the lesion is healthy. So, can we restore this missing neural transmission?" Dr Capogrosso has been working in the neuroscience of movement field since 2009, trying to find a safe way to amplify residual capacity in the nervous system to re-establish muscle control. At the heart of his work is an electrical stimulation method called epidural spinal cord stimulation (SCS).

"An electrode is inserted between the spinal cord and vertebrae (the epidural cavity)," Dr Capogrosso explains. "The nervous system isn't pierced which minimises the risks, though the procedure still requires surgical intervention. SCS has been shown to improve motor function empirically since the 70s – but it is not known exactly how. Which sensory and motor fibers co-activate during stimulation, and how do we fine-tune specific movement?"

To answer these questions, Dr Capogrosso's lab developed a 3D geometrical model capable of replicating the effects of SCS therapy on motor and sensory fibers in the body. The model combined aspects of physics, such as the electric membrane properties of neurons and the electric potential that SCS generates inside the body, with anatomically realistic measures of the cervical spine (volume, localisation and properties of the various cells and tissues), to predict medical outcomes for a patient. The results of the simulations - including a range of different electrode placements and the corresponding neural trajectories - were validated with experiments in animals and humans.

"We wanted to know what the likely targets of such stimulation would be and estimate the activation thresholds of the nerves in order to predict how upper body mobility would be affected by the procedure," explains Dr Capogrosso. "We found that SCS can activate dormant motor circuits by injecting excitatory inputs through the sensory afferents (nerve fibers important for proprioception), i.e. by directly stimulating fibers that normally carry sensation of movement. These fibers, in turn, convey excitatory inputs to the circuits below the lesion, thereby inducing an alternative signal route to the one lost to disease or injury – and allowing movement."

> "SCS can be fine-tuned to selective engage muscles of the arm and hand to restore functional reaching and grasping in monkeys with spinal cord injury.

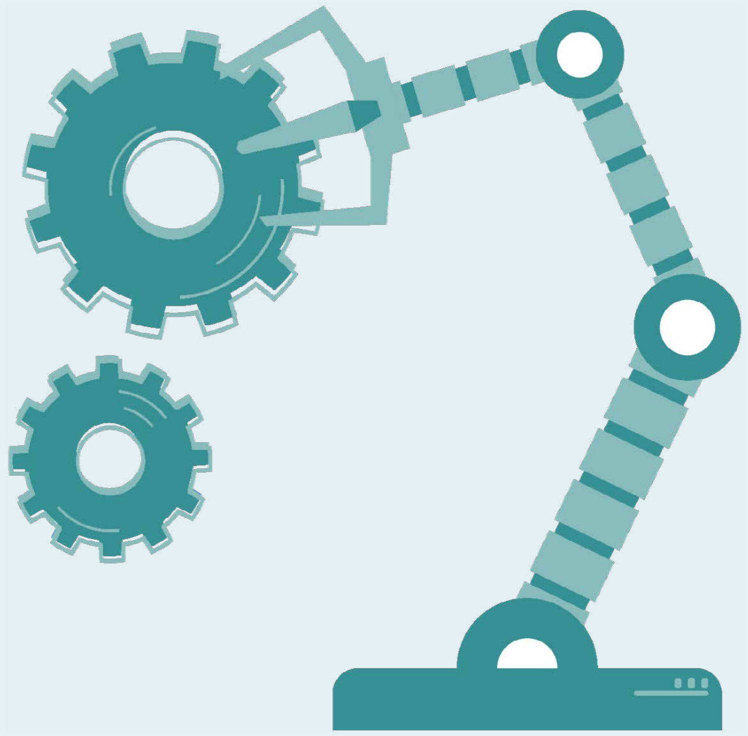

Dr Capogrosso's laboratory showed that SCS can be fine-tuned to selectively engage muscles of the arm and hand to restore functional reaching and grasping in monkeys with spinal cord injury. This work is important for translation to humans since monkeys are the most accurate animal model for control of the arm and hand.

"We have now started testing this technology in humans with stroke in an ongoing clinical trial, where we are hoping to see results similar to those in monkeys." Dr Capogrosso adds that one of the current challenges is optimising the stimulator for each patient (adjusting signal strength and amplitude) as this requires a lengthy trial and error process. "We have a collaboration with Prof Doug Weber, Carnegie Mellon University, where his team and he are trying to transform this into an automated machine intelligence-guided optimisation," Dr Capogrosso says. "With that, and our two upcoming clinical trials, we might be able to completely change the field."

Dr Marco Capogrosso

Key papers for this article:

Greiner, N., Barra, B., Schiavone, G., Lorach, H., James, N., Conti, S., Kaeser, M., Fallegger, F., Borgognon, S., Lacour, S., Bloch, J., Courtine, G., & Capogrosso, M. (2021). Recruitment of upper-limb motoneurons with epidural electrical stimulation of the cervical spinal cord. *Nature Communications*, 12(1), 1-19. https://doi.org/10.1038/s41467-020-20703-1

Capogrosso, M. Spinal Cord Stimulation for Restoration of Arm and Hand Function in People With Subcortical Stroke. https://clinicaltrials.gov/ct2/show/NCT04512690

Memorisation and Generalisation: An AI-inspired theory of memory processing

with Dr Weinan Sun, Janelia Research Campus

Many of the AI systems we know, have been heavily inspired by the brain: from the architecture of the neurons within it, to the way it processes and stores information. But can AI itself inspire ideas about the brain?

Investigators at the HHMI Janelia Research Campus and University College London have been studying memory formation through the lens of machine learning concepts. Dr Weinan Sun (HHMI Janelia) told us more about the project: "The hippocampus is this structure in the brain which acts as a sort of fast 'tape recorder,' capturing memories before gradually transferring them to the more stable, higher capacity neocortex." This hippocampus-coordinated process of information transfer is known as 'systems consolidation' and is thought to be key to integrating knowledge. Without the hippocampus, the ability to form new memories (more specifically, so-called 'episodic memories') is severely impaired.

"Researchers have been looking at humans and animals with hippocampal lesions or surgically removed hippocampus. For such cases, the classic theory has been that only memories formed recently before the hippocampus damage (i.e. unprocessed memories) would be affected. However, there are studies, where we see loss of both recent, and old memories." Dr Sun points to a paper by Ocampo et al. (2017). "We thought concepts from machine learning might explain why."

Neural networks are trained using data, for example images of animals such as cats and dogs. The network 'learns' from the dataset to make decisions on new data. Over time, the network becomes better at 'generalisation', meaning it can more accurately identify cats and dogs in new images. "But the training data given to the network isn't perfect. There may be a mislabelled animal, or a weird-looking cat or dog... If the network learns from the same pool of data for too long, it begins trying to find all the features it sees in the training set, including the noise. And so, the generalisation error begins to increase again – this is called 'overfitting'." Dr Sun argues that this is actually similar to people's experience of the world. Information, some useful and reliable, is captured - but so is noise. The hippocampus then starts to train the neocortex to learn about the world. "And in the past, it's been suggested that the hippocampus replays memories to optimise memorisation. We propose that, instead, the goal is to optimise generalisation. We built a model to test this theory."

Investigators Dr James Fitzgerald (HHMI Janelia) and Dr Andrew Saxe (University College London) suggested studying the problem using a novel 'teacher-student-notebook' framework, which consists of three neural networks. The 'teacher' represents the environment which generates "examples" to teach the 'student' - the neocortex – whereas the 'notebook' (the hippocampus) records the examples for replay. The 'student' and the 'teacher' are both feed-forward networks taking an input and producing an output. The teacher generates examples and passes them to the student. The goal of the student is to learn the teacher's transformation through an error optimisation algorithm called 'gradient descent'. Gradient descent aims to minimise the cost function (the model's output error) and updates the weights in the student's weight matrix accordingly after each 'learning example'. The notebook is a Hopfield network and iteratively replays the examples to the student. Dr Sun's model tests the student after every learning epoch on two things: new data (to assess 'generalisation') and old data (to assess 'memorisation'). If the teacher's data is perfect—i.e. it has no noise—the memorisation and the generalisation errors eventually reach zero. If the teacher's data is 'imperfect', memorisation error reaches zero, while 'generalisation' eventually suffers from overfitting.

"Through learning, predictable memories are ultimately stored in the neocortex. But idiosyncratic memories, which may harm generalisation, are stored in the hippocampus. This natural separation stops unpredictable information from causing overfitting. It also explains how damage to the hippocampus has different effects on reliable and 'accidental' memories." Dr Sun hopes the student-teacher-notebook framework inspires the experimental neuroscience community to run real-life tests of his theory. For example, mice can be exposed to either predictable or unpredictable experiences in the form of visuals and associate that with a reward. Afterwards, the mice's hippocampus will be lesioned, expecting the predictable task to be preserved while the unpredictable task – not.

"We've provided novel angles to resolve a multi-decade-old debate," Dr Sun concludes. "This theory is also medically relevant for many disorders in which the memory transformation process goes awry, such as PTSD, generalized anxiety disorder, and more."

Dr Sun's project was conducted under the guidance of James Fitzgerald (Janelia), Andrew Saxe (UCL), and Nelson Spruston (Janelia).

Left to right: Dr Andrew Saxe, Dr James Fitzgerald, Dr Weinan Sun

Photo credit: Dr Weinan Sun

References:

Ocampo, A. C., Squire, L. R., & Clark, R. E. (2017). Hippocampal area CA1 and remote memory in rats. *Learning & Memory*, 24(11), 563-568. https://doi.org/10.1101/lm.045781.117

Key paper for this article:

Sun, W., Advani, M., Spruston, N., Saxe, A., & Fitzgerald, J. E. (2021). Organizing memories for generalization in complementary learning systems. *BioRxiv*. https://doi.org/10.1101/2021.10.13.463791

References

Access all the references for this issue by scanning the QR code or clicking here.

Contact Us

Email: kristine.lennie@researcher-app.com
Website: https://www.researcher-app.com/