Electroacupuncture (EA) and the EEG:
An unfinished personal journey, 2001-2022.
From simple hypothesis to artificial intelligence (AI)

[Slide 1]

In this talk, I want to tell you about how acupuncture – in particular transcutaneous electroacupuncture, or TEAS – can affect the brain. I also want to tell the meandering story of how a traditional acupuncturist became involved in neuroscience, which you may find even more instructive.

You can find a list of sources, acknowledgements and credits at the end of the presentation.

[Slide 2]

When I left school, I was fascinated by the work of Abdus Salam on Lie groups and the symmetries of subatomic particles, but when I reached Cambridge soon found I just wasn't up to the maths involved, and changed subjects. I certainly had little love for computers and stats. It was thus as an art history undergraduate in 1970 that I was privileged to meet the neurophysiologist Grey Walter, a pioneer in both EEG research and robotics. His book 'The Living Brain' – still worth reading today – was an inspiration to me then, and led indirectly to my own later investigations into the effects of acupuncture on the EEG, electroencephalography. The *global* and *local* patterns implicit in the EEG are the nearest thing I have found to the satisfying symmetries of the physics that both fascinated and eluded me earlier.

After a first career in the world of artists' books and performance art, I became increasingly interested in what later became known as 'energy medicine', and ended up translating a German paper on 'orgone acupuncture' by Bernd Senf for the American Journal of Acupuncture, as well as excerpts from Romanian books on 'modern scientific acupuncture' and electronography by Ioan Dumitrescu. These were published in a little audio cassette magazine that I was co-editing at the time with Mike Weaver, my former American Arts supervisor from Exeter University. Another contributor to the magazine was a friend of Mike's, Roger Hill, who later became the chairman of the Traditional Acupuncture Society, one of the precursors of the British Acupuncture Council (BAcC). Only after that did I train in five-element acupuncture, in Leamington Spa. And only in 1988, after six years of practice, did I feel confident enough to incorporate electroacupuncture into my clinical work, coincidentally returning to a childhood interest in electricity.

At that time, there was little guidance available on how electroacupuncture should be used, other than 'zap it and see', so I gravitated to the methods clearly taught by the Society of Biophysical Medicine, later known as the Equinox Group (which included Gordon Gadsby, Mike Flowerdew and Rodney Robinson, among others). Part of their rationale was that stimulation applied at particular frequencies – whether through needles, TENS pads or hand-held probes – would elicit electrical activity in the brain at those *same* frequencies, as also suggested by the creator of the Canadian Codetron device, Norman Salansky. Thus, low frequencies (2.5 Hz, for example) should relax, medium frequencies (around 10 Hz) should normalise or stabilise, and high frequencies (80 Hz, or even 160 Hz) would excite. That made a lot of sense to me, as it echoed to some extent what Grey Walter had found with the ‘frequency-following response’ (or FFR) that occurs in the brain with photic stimulation, i.e. pulsing light.

[]

In 1996, I started to teach electroacupuncture, and took my own first faltering steps in the field of electroacupuncture research when I was invited by the publisher Elsevier to put together a textbook on the subject. I am afraid I am a man for detail, and this turned out to be an 11-year project, with contributions from around 20 experts in the field, including members of the AACP, BMAS, BAcC and the Equinox group. While the publishers wanted a practical introduction to the subject, I had always stubbornly wanted to write more of a research resource. In the end, I think the book more or less satisfied both of us, with two Forewords – one by Angela and John Hicks of the College of Integrated Chinese Medicine in Reading, the other by Zang-Hee Cho, a pioneer in the development of PET scanning (that’s positron emission tomography) and a Professor at the University of California.

[Slide 3]

The electroacupuncture book was finally published in 2007, and is still in print, associated with an online database of material garnered from some 8,000 clinical studies. Of course, 2007 is a long time ago now, and much has changed. In particular, the amount of research published on EA has increased exponentially over the years – particularly in the 1990s.

[Slide 4]

Thus, while more than 1700 studies on EA were listed in PubMed up to the end of 2005, around 4,700 have been indexed since – with many more, of course, listed in Chinese databases like CNKI, the China National Knowledge Infrastructure. It is impossible to keep up! And so the book and accompanying online database on which I spent so much time, energy and money (paying translators), much to my wife’s despair, is now well out of date – although still useful (I have to say that, of course!). Nonetheless, the whole process taught me a lot – about electroacupuncture, research, and myself, about my own ambition and greed for knowledge, and also about the inexorable effects of time. Fortunately, you now have experienced and inspiring teachers like Stephen Lee, Lynn Pearce and Kevin Young who can impart the practical skills of electroacupuncture, especially for treating pain, acute or chronic.

Please note that what follows in this presentation is about *experimental*, not clinical research.

[Slide 5]

A major inspiration for me while preparing the electroacupuncture book was the work of professor Han Jisheng, a pioneer in the research of basic mechanisms of acupuncture since 1965, and currently Director of the Neuroscience Research Institute at Beijing University.

You are probably already familiar with the conclusions from the many Chinese animal studies, especially his, as well as Western research on TENS, that different frequencies of stimulation have different neurochemical effects in the spinal cord and brain.

[Slide 6]

As the Han institute website states:

"The most important determinant of acupuncture effect is the *frequency* of impulses transmitted along the nerve fibres from periphery to central nervous system. Signals of different frequencies can induce the release of different kinds of chemical mediators." This chimes well with Grey Walter's statement that "Signals due to rhythmic stimulation ... appear to reach parts of the central nervous system which are inaccessible to impulses set up by non-rhythmic stimuli, however intense".

On the other hand, *varying* the stimulation in some way may be beneficial, avoiding habituation and a dulling down of response. This is part of the justification for the 'dense-disperse' *alternation* of low and high frequencies (such as 2 Hz and 100 Hz) advocated by Han, and also of the Codetron device developed by Pomeranz and Salansky in Canada in the 1980s, where electrodes were energised in random order.

In 2001, Zang-Hee Cho and I both gave presentations at the annual AACP Conference (held that year in March at the Latimer Conference Centre near Chesham), he on the functional magnetic resonance imaging (fMRI) effects of acupuncture at the point Bladder 67 and other findings, and I on central nervous system resonances to peripheral stimulation, the principle underlying the methods taught by the Society of Biophysical Medicine. Mine was quite a muddled talk, so I won't bother you with the details now. You can always find it in the AACP journal if you really want to.

Suffice it to say that I mentioned the 'homoeodynamic' nature of neuroplasticity, and that signals in the healthy brain are not regular or fixed, but continually changing, even 'chaotic', with greater EEG complexity or variability usually a concomitant of both mental and neurological health. At the end of the presentation, I outlined a simple research proposal, using electroacupuncture with needles or pads (i.e. TEAS) at major points such as LI4 and ST36, in unilateral and bilateral configurations, stressing the need to record the EEG both during <u>and after</u> stimulation.

[]

However, it's all very well having ideas about experiments that it would be interesting to perform, quite another thing to get around to doing them, especially when you are a busy practitioner with a family but without funding. I have to say, though, that the AACP has always been generous – making me an honorary member in around 2003, and even providing a small research grant for our most recent study, which I will describe shortly. My *own* professional organisation, the BAcC, on the other hand, did not. I'm probably not traditional enough.

While preparing the electroacupuncture book, I'd visited Tim Watson, the king of electrotherapy in the UK, while he was at Brunel University. In 2008, after his move to the University of Hertfordshire, he encouraged me to become a research associate there, but it was not until two years later that fortuitously, through a Russian artist, Evgenia Emets, who had an interest in visualising and sonifying the electromagnetic data detected from sensors on the body, that I met a fellow maverick and EEG biofeedback practitioner, New Zealander Tony Steffert, who was willing to collaborate without charging the earth for his time.

Suddenly, everything became possible. With Tim's support and supervision, Tony and I embarked on our journey together, undertaking a series of five small pilot studies with 40 participants in total, mostly in my own clinic, and then a larger study at the University of Hertfordshire in 2015-2016, with 66 participants, mostly from the University. Our aim was to try and answer research questions about how needling, EA and TEAS at different frequencies might affect physiology. We have investigated not just the EEG response, but – with Tony's expertise – also changes in heart rate variability (HRV), respiration, blood flow, temperature and even head movement. In 2014, realising we couldn't really do everything ourselves after all, a Computer Science graduate was taken on for a self-funded PhD on our fifth pilot study. Ronak Bhavsar both recorded and processed the data, and analysed results. We presented posters side by side at the ARRC Symposium in London, and, as she said later, 'we had a lot fun'. However, quite rightly, her thesis took a different direction from the one I'd hoped, and did little in the end to help answer my research question about the effects of stimulation frequency. Having a baby also meant that it took her a lot longer than planned to complete.

A little diversion on entropy

[Slide 7]

Entropy, the amount of uncertainty and unpredictability in a system, or in a signal such as the EEG or ECG, was one of the main topics in Ronak's PhD. Closely related to concepts such as chaos and complexity, research into entropy is another area of exponential growth, far more so, in fact, than acupuncture – as you can see from the graph. Given that Ronak had not, after all, used entropy measures to investigate the differential effects of stimulation frequency on the EEG and HRV, I started to think about how we could do that ourselves. Ronak's principal supervisor, Na Helian, came up with another PhD student, Deepak Panday, who turned out to be a whizz in the programming language MATLAB.

[Slide 8]

Thanks to Deepak and his wife Harikala, we were finally able to develop CEPS, a software tool with a 'graphical user interface' (or GUI) that clinicians without computer science skills could use *themselves* to analyse the complexity and entropy of signals. It currently includes around 120 measures, and more than 20 researchers from around the world have actively contributed to its development. This may have been a major diversion from my primary research question, but has proved both fascinating and extremely fruitful. If any of you are interested in using this, feel free to get in touch. Don't worry – it's not as complicated as it looks!

[Slide 9]

In our most recent study, 66 reasonably healthy participants attended for four sessions, a week or more apart (except for four who dropped out after only one session, and another who only completed three). Sessions were conducted in the University's Physiotherapy Lab, but despite our best efforts this could not be completely soundproofed or temperature-controlled, and was quite brightly lit. Data were collected by Tony, Aistė Noreikaitė or our other research assistant, Lidia Zaleczna. I administered questionnaires, positioned and removed electrodes and ran around trying to make sure everything ran smoothly.

Participants were seated upright in a comfortable chair, eyes open, with both forearms supported. Informed consent was obtained, questionnaires completed, and the participants were then prepared for the session. This preparation, which took around 15 minutes, involved fitting an EEG cap with head movement sensors attached, and affixing ECG electrodes to the forearms, as well as other sensors to the fingers of both hands. These were all worn for the remainder of the session (usually around 60 to 90 minutes).

[Slide 10]

Following an initial 5-minute baseline recording, TEAS was applied for 20 minutes to both hands, with a short pause halfway to allow a brief rest. EEG recording continued during stimulation, which was between the acupuncture point LI4 (*hegu*) and the ulnar border of the *same* hand. In other words, current only passed between the electrodes on each hand, and did *not* flow through the arms and torso, so that, in principle, it shouldn't affect the heart – or brain - directly.

After stimulation, recording was continued for a further 15 minutes to assess post-stimulation changes.

[Slide 11]

A charge-balanced Equinox stimulator was used in all four sessions, and in each session was set at one of four different frequencies – 2.5, 10, 80 or 160 Hz.

(Strictly speaking, as stimulation consisted of alternating monophasic pulses, frequency should be in units of 'pulses per second', but I'll use the more familiar 'Hertz' here.)

For the three lower frequencies, output amplitude was set to provide a 'strong but comfortable' sensation for that particular participant. In contrast, 160 Hz was applied as a 'sham' treatment, with the device switched on (and a flashing light visible), but the output amplitude remaining at zero throughout – although a *pretence* was made of turning it up out of sight of the participants. Some participants were in fact aware of this supposedly subthreshold, clinically irrelevant current. The different stimulation frequencies for each participant were applied in a semi-randomised balanced order.

[Slide 12]

The so-called 10/20 system of EEG electrode location was used (19 electrodes, 'montaged' with linked ears as reference and ground anterior to Fz). Data collection followed standard EEG procedures, with electrode caps selected by size for maximum comfort.

[]

After spending more than a year in the Lab, we finally had our data! Again, though, while it may be relatively simple – if laborious and challenging – to write Ethics applications and gather data, once you have all your recordings, what do you do with them? Physiological data are rarely neat and tidy, can be corrupted by noise or other artefacts (as shown in a previous slide), and have to be in the right format for the analysis you want to undertake.

More diversions

while wrestling with the intricacies of how to process the mass of EEG data we had gathered, we decided we had to learn more about how to do this ourselves. As a result, Tony and I ended up at the Medical Research Council Cognition and Brain Sciences Unit in Cambridge, organising a memorable course, on ‘Brainstorm’, a collaborative, open-source application for the analysis of brain recordings that had been developed in Montreal.

[Slide 13]

On the second evening, after a full-on day’s training, we relaxed on the banks of the river Cam, listening to EEG sonification created by Lithuanian artist Aistė Noreikaitė, one of the two research assistants who had helped us record and process our EEG and HRV data. Quite something – and a little puzzling to some of the neuroscientists, although the event was also an experiment in ‘hyperscanning’, sonifying a whole orchestra of brains!

[Slide 14]

Later, we formed Cephalopedia, ‘a loose-knit international group of experimental artists and performers ... with a neuroscience edge’. We performed in London, Paris, Lithuania and, in 2018, the University of Hertfordshire. But enough of that.

Now down to the nitty-gritty.

Cleaning up the data

One of those who came to our University performance was Thea Radüntz, from the German Federal Institute of Occupational Safety and Health. She had developed an automated EEG artifact elimination method using machine learning, and generously applied that to our EEG data for free, although the results were still difficult for us to understand. As we were to find later, findings using ‘artificial intelligence’ methods – ‘machine learning’, and more especially so-called ‘deep learning’ - are notoriously difficult to interpret.

In 2019, through a community of researchers using EEGLab, the most widely used software package for processing EEG (or MEG) data, we eventually found someone willing to help us clean up our EEG data and prepare it for analysis in a way that we, as clinicians rather than computer scientists, could interpret. Paul Steinfath, then at the Max Planck Institute for Human Cognitive and Brain Sciences in Leipzig, helped us through to the stage where Tony and I could do some basic number crunching and statistical analysis. Stats, of course, are essential when designing and interpreting study results, and Neil Spencer, the University of Hertfordshire’s Professor of Practical Statistics, has supported us on several madcap projects, from analysing the effects of paced breathing on HRV to exploring the effects of space and terrestrial weather on the EEG [a study in progress with Ciar*án Beg*gan of the British Geological Survey].

[Slide 15]

I don’t want to get too technical, but in our TEAS study, data were recorded initially in WinEEG software, saved out in EDF, the more commonly used ‘European Data Format’, and then each separate session file was cut into eight separate MATLAB files, one for each 5-minute ‘slot’. Data were filtered between 0.5 and 45 Hz, to include the standard EEG ranges (delta to gamma). A 50-Hz ‘notch filter’ was also used (49-51 Hz) to remove mains interference (particularly important as we also intended to analyse the EEG in higher-frequency bands). Data were also *re*-montaged using an average reference.

We ended up with six sets of cleaned EEG data, in addition to Thea's. Under Paul's guidance, Tony used two versions of what is called 'independent component analysis' (ICA) to sort the real EEG from noise (such as eye blinks, ECG artefact or muscle movement): first, 'extended Infomax' (based on maximising the nonlinear measure of 'mutual information' between input and output) and also 'AMICA' (Adaptive Mixture ICA). Other EEGLab 'plug-ins' were used as well, and data were then 're-montaged' yet again, this time with the Laplacian form of local average reference. Various methods of time-frequency analysis were then used to extract different sub-bands from the raw EEG signal: the Welch periodogram, complex Morlet wavelet and Thomson's multitaper. Results were compared, and later I will describe some of those obtained with the Infomax/multitaper combination.

At this stage, data from a number of participants were excluded, either because of inadvertent differences in sampling rate (in four sessions), missing data or because recordings were cut short for one reason or another (e.g., discomfort from wearing the cap – or even having to take a trip to the bathroom). 1,536 files for 48 participants remained, out of a possible 2,112 if the data were complete.

[Slide 16]

Now, at least, we had some idea of the EEG power in different bands, whether these were the conventional bands – delta, theta, … etc. – or the tighter bands, 1-Hz or 3-Hz 'bins', that we used to try and detect resonance or frequency-following effects. And we could also calculate other measures based on these. As you can see, there are many, many possible ways of quantifying changes in response to treatment.

[Slide 17]

One such measure that interested us was 'cordance', a measure developed by Andrew Leuchter and his colleagues from UCLA. Cordance, used primarily in studies on depression, is derived from both absolute and relative power at the EEG electrodes, relative power being absolute power in *each* frequency band as a percentage of the absolute power summed over *all* frequency bands.

Cordance is supposedly better correlated with regional blood flow in the brain than either absolute or relative EEG power individually, although not everyone agrees with Leuchter's conclusions on this. Unfortunately, by the time we were ready to use cordance, the online calculator I'd originally found was no longer available, and despite Leuchter's personal assurances, code in a form that Tony could use (in the programming language MATLAB) was not forthcoming. Hunting around, I noted that another research group, this time in Turkey, had published a number of cordance studies. This time we struck lucky, and in 2017 Türker Tekin Ergüzel, of Üsküdar University in Istanbul, was willing to share his method of calculating cordance in Excel. But unfortunately again, by 2020, when finally we had cleaned data on which we could use his method, he was becoming very impatient at what he saw as our unnecessary delays. Finally though, last year Tony managed to translate Türker's code into MATLAB and we could proceed.

[Slide 18]

Brief literature review

Before I share our own findings on TEAS and the EEG, I'd like to review very briefly what other acupuncture researchers have found. There are over 300 studies mentioning acupuncture and the EEG indexed in PubMed, and almost 50 on TENS and the EEG. Of these, 147, published between 1986 and 2022, were easily retrieved and could be examined in depth. 56 (around 38%) were from China, 15 from Korea, 13 from the US, 11 from Japan and 10 from Taiwan. Other countries were represented by fewer than 10 studies each. Of the Chinese studies, almost half (25, so more than 44%), were from the prestigious Hebei University of Technology, aka Tianjin University.

Most of these EEG studies were on manual acupuncture and, of course, you have to be careful about extrapolating from one modality to another. As stated in a 2007 review by Dhond *et al.*: 'TEAS is different from insertive electroacupuncture in many ways, and the results from these studies may not apply to acupuncture'.

Studies were mostly small (median *N* = 14, IQR 10 to 25). Twenty were animal (mostly rat) studies; eight of these did not provide information on how many animals were involved – as if the individual animals didn't matter.

Most of the studies located on EEG and acupuncture or TENS variants were experimental rather than clinical. Of the remainder, fifteen concerned pain (either experimental or clinical).

In the acupuncture-related studies, the points most commonly used - as also noted in a 2018 systematic review of 19 EEG acupuncture studies - were ST36, LI4 and P6.

Stimulation and EEG duration:

Studies were very heterogeneous, some being on manual acupuncture (MA), some on EA, some on TEAS or TENS, but durations of 20-30 minutes were those most commonly used for the latter – as you'd expect from clinical practice. EEG was often recorded before, during and after stimulation, and artefact-free segments from each period were then selected for further analysis. 2-second segments of EEG were most frequently used.

[Slide 19]

Methods of analysing changes in the EEG were varied. Measures based on EEG power occurred in similar numbers of studies published before and after 2013, the median year of publication for the 147 studies located, as did nonlinear entropy and complexity measures. Only one study on cordance was located. Functional connectivity measures based on what is called graph or network theory – i.e. quantifying *relationships* between the EEG at different electrodes - were found in only one study before 2013, but in 13 of the 72 studies published since then.

Of the Tianjin studies located, 14 were on manual acupuncture, 10 on non-invasive magnetic stimulation (TMS), one on moxibustion and one on 100 Hz microcurrent TEAS. Half the Tianjin acupuncture studies, published between 2010 and 2021, investigated the effects of different frequencies of needle-twirling at ST36 on the EEG. Participants were lying down, eyes closed, in a darkened room. Three different frequencies were used in the same session, with between 4- and 10-minute rests between them, depending on the study.

In contrast, only one upper limb TENS study, a BSc thesis from Holland, investigated the effects of stimulation frequency on the EEG, and did not use low frequency stimulation.

Clearly, Tianjin is the place to be if you want to research the effects of acupuncture-related modalities on the EEG; and graph theory is a method now very much in vogue. As yet, there are only two studies in PubMed on artificial intelligence (whether machine learning or deep learning), the EEG and acupuncture or transcutaneous stimulation (TENS or TEAS). One of these is from Tianjin.

The plot shows how, while electroacupuncture, TENS and TEAS publication rates increase linearly, those for machine learning and deep learning are increasing exponentially.

I will now move on to describe some results from our most recent study.

[Slide 20]

To cut many long and tedious stories short, having pre-processed the EEG data using Infomax ICA, I tested our EEG data for normality of distribution using the Kolmogorov-Smirnov test and by checking for skewness and kurtosis. I then started our analysis using the maximum and median values of power in specified frequency bands defined by using the multitaper method for our 5-minute recordings.

[]

If there is a frequency-following neural transmission effect (FFR), maximum power would be expected to occur in bands centred on the stimulation frequency or its harmonics – although this could also occur if transmission from the hands to the brain occurs in muscle tissue – so-called 'volume conduction' (VC) - rather than in ascending nerve pathways. When stimulating LI4, the FFR is also likely to be marked over the index finger regions of the primary somatosensory cortex, beneath EEG electrodes C3 and C4. The corresponding regions for the little fingers lie between these and the vertex, Cz, but please note that the illustration on this slide is not intended to be accurate – it's more of an *aide mémoire*. There is not a simple 1:1 mapping of the underlying cortical areas to the surface electrodes, and as the signal spreads through the brain, the FFR may become less clear.

The results here are all for absolute EEG power. Those for relative EEG power have not yet been calculated.

The upper left Table shows counts of *maximum* power (or *global* peaks) across all 17 bins listed in Slide 16, for all 19 electrodes.

The upper right Table shows corresponding results for the *local* peaks in the data.

For comparison, in the lower Table local peaks in control bins *unrelated* to our stimulation frequencies are shown. The alpha peak, which is around 9-10 Hz, varying in different individuals, is still visible, and there is still some effect at 2.25 Hz in response to 2.5 Hz stimulation, but less than the purported 'resonance' effect at 2.5 Hz itself.

[Slide 21]

The signal from the Equinox stimulator is approximately a square wave. In a perfectly conducting medium, a proper square wave would be made up solely of odd harmonics, as already mentioned in slide 11. However, more local EEG peaks were found for the *even* than for the odd harmonics of the stimulation frequency, perhaps suggesting that they may not all be due to volume conduction. (Although, of course, what start out as square waves do not remain that shape as they are conducted through living tissue.)

The plots show the numbers of times each electrode shows peaks at odd and even harmonics of 2.5 Hz.

However, median peak *height* in these narrow bins was greater with *stronger* stimulation at some electrodes for both 2.5 and 10 Hz stimulation, which could suggest volume conduction.

So – although I'm biased - I think we have a draw here – Frequency Following Response (FFR) 1, Volume Conduction (VC) 1.

[Slides 22-24]

The EEG electrodes can be divided into anterior and posterior [slide], left and right [slide], and central and peripheral, or outer [slide].

[Slide 25]

Interestingly, during stimulation more peaks - and stimulation *harmonic* frequency peaks in particular – were found over the *back* of the brain rather than frontally, regardless of frequency, on the *Left* side (except during 10 Hz TEAS), and for 2.5 Hz TEAS *Centrally* rather than at the Outer electrodes. Bear in mind that this central effect might result from neural rather than volume conduction.

[**BACK** to Slide 24]

Conversely, during stimulation at 80 Hz or 160 Hz (Sham), the effect was more at the *Outer* electrodes - perhaps indicating *volume* conduction up the neck muscles (a strong frequency-following or resonance effect is less likely at these higher frequencies).

These results suggest neural conduction is occurring at 2.5 Hz, but more of a volume conduction effect at 80 Hz or 160 Hz (sham). 10 Hz stimulation may involve *both* mechanisms.

A similar pattern of higher central absolute *power* (rather than peak counts) for *lower* stimulation frequency and *lower* central power for higher frequencies (80 or 160 Hz) also looks likely, but this requires further investigation.

So we still have a draw: FFR 2, VC 2.

[Slide 26]

Counting all local peaks at each electrode in narrow (3-Hz) bins, for the four different stimulation frequencies, we can see that during stimulation 2.5 Hz is the most active frequency (produces most peaks), followed by 10 Hz, but that the after-effect of 2.5 Hz is – surprisingly - less than that of 10 Hz TEAS.

I also explored other simple EEG measures, such as 'spectral edge frequency', and spectral and spatial centroids, but they did not produce particularly interesting findings.

Cordance, however, was a little better.

[Slide 27]

At Cz, the most central electrode, changes in cordance values from baseline differed with stimulation frequency, being greatest for 2.5 Hz in Delta (a resonance effect?), and with the steepest decrease over time in Alpha occurring for 10 Hz.

It is intriguing that the amplitude of *peripheral* blood flow, assessed using fingertip PPG (photoplethysmography), was strongly, and negatively, correlated with Alpha cordance at electrodes C3 *and* C4 during and following 10 Hz TEAS, but positively correlated with Beta cordance. The other stimulation frequencies did not show this strong effect. Remember that Leuchter found corresponding correlations between *cerebral* perfusion and cordance at these same electrodes during a motor task.

[Slide 28]

Although cordance is usually calculated for the four standard EEG bands, results when 11 1-Hz bins were used demonstrate the effect of stimulation more clearly. Note: these are *counts* of significant results, not cordance *values*.

Particularly on the right, you can see that differences in cordance with frequency were fewer *post*-TEAS, and that these differences were more evident in 1-Hz bins than standard bands.

[]

If volume conduction played a major role, cordance values at different electrodes might be expected to be strongly correlated *during* stimulation. However, numbers of positive correlations were higher during *non*-stimulation than during stimulation slots at all except six *electrodes*, and for only one of these (T4) was there a *significant* preponderance of positive correlations during stimulation as against non-stimulation ~~(p = 0.024)~~.

Similarly, numbers of positive correlations were higher during *non*-stimulation than during stimulation slots in all EEG *bands*, except for 80 Hz stimulation in Delta, Theta and Alpha. This is another clue that perhaps what we are seeing are not simply the results of volume conduction, except at higher frequencies.

Possibly FFR now has the advantage: FFR 3, VC 2.

I am still trying to figure out if the other results included on this slide are simply the result of ‘fishing’ for something meaningful, or something more.

[Slide 29]

In 1970, Bo Hjorth introduced three general parameters for single-channel data that do not depend on frequency analysis of the EEG. He called them Activity, Mobility and Complexity. To make life simpler still, these were then generalised further by others so that multi-channel EEG data could be summarised using just three linear global ‘state space’ descriptors: Global field strength (Σ), Global frequency of field changes (Φ) and spatial complexity (Ω).

An Erasmus student at the University of Hertfordshire, Firgan Feradov from Bulgaria, kindly provided code for the Hjorth parameters. Jiří Wackermann, although now retired, provided code for the global descriptors he had helped create. Tony implemented the matlab codes, and I conducted a simple analysis using SPSS and Excel.

Some intriguing results were found for both the Hjorth parameters, and for the Global descriptors, but to be able to interpret them will take further thought.

As shown Top Right, Global field strength (Σ) for the group increased during stimulation at all frequencies except 10 Hz, for which it *de*creased. Generalised frequency (Φ) and spatial complexity (Ω) also increased during stimulation at 2.5 Hz, but this was not the case for Ω at 10 Hz.

Whether at baseline, during or after stimulation, Σ tended to increase rather than decrease during each 5-minute slot.

Apart from using the software toolbox CEPS on our EEG data, we had now got pretty much to the end of where we could go without further assistance from 'real' neuroscientists with more advanced computer skills, when fate stepped in unexpectedly.

Our CEPS paper was listed on ResearchGate, and Çağlar Uyulan, a young Turkish researcher, started following our publications. It turned out that he was a former supervisee of Türker Tekin Ergüzel, and together they were now involved in research on the EEG using AI, notably 'deep learning' methods, some of which were based on entropy measures. I was hooked, and after some initial caution from Türker, still annoyed at my slowness in using his methods of computing cordance, we signed a data-sharing agreement and I sent gigabytes of it off to Istanbul.

[Slide 30]

Unfortunately, we soon discovered that collaboration at a distance, for people with very different skills, mindsets, languages and characters, was not going to be straightforward. I could not understand the language of deep learning, and Çağlar could not understand why I wanted to understand. Türker stood on the sidelines, getting impatient again.

Working towards producing a paper that will satisfy all of us continues to be a stressful business. I won't bore you with the details, as it's not yet published and I still don't understand them all, but the main conclusion is, very simply, that the EEG during TEAS is different from the EEG at baseline or afterwards, and that these differences vary with stimulation frequency.

However, we have already seen this from the simpler methods I have used – so have we learned anything really? Sadly – and this is an inherent problem with all deep learning methods – *which* particular EEG characteristics are significantly different remains unclear. The results are difficult to interpret.

So, despite weeks of Turkish computing time and continuing frustration on all sides, we hadn't got very far with AI, despite our initial optimism. A good lesson: look before you leap!

[Slide 31]

But all was not lost! At this point, I remembered Iosif Mporas, who'd supervised Firgan's PhD on the Hjorth parameters in 2016, and is now Reader in Signal Processing and Machine Learning at the University of Hertfordshire. In an online meeting with Tony and me, he immediately saw that my enthusiasm for the new had led me in a wrong direction, and is now helping us to complete our paper and explore m*achine* learning - rather than *deep* learning – as a way to analyse our EEG data.

Through Iosif and Deepak, we are now recruiting computer science MSc students to help us with this work, … as well – for instance - as trying to analyse the spontaneous head movements that occur during TEAS: are these dependent on frequency of stimulation, respiration rate or other factors? We are also starting to explore the heady realms of EEG connectivity and graph theory, which should finally help us sort out which effects are due to a frequency-following response and which to volume conduction. Stavros Dimitriadis from Cardiff University, an expert in this area, has offered to help with this part of our analysis. We will also be working with French mathematician Davide Faranda to apply his dynamical systems climate modelling metrics to our EEG data. All ground-breaking stuff!

In addition, I mentioned that in our first study on CEPS we had investigated the effects of paced breathing on HRV. Tony is now gathering data for a second study. Initial results indicate very different patterns of response to paced, self-paced and spontaneous breathing at different rates. You will hear more about breathing from Peter Deadman this afternoon.

In conclusion

We have not fully confirmed our hypothesis of a frequency-following response, but have not ruled it out either, particularly for low frequency (2.5 Hz) TEAS. It is also possible that volume conduction may play a greater role in *high* frequency stimulation. Another potentially fruitful avenue to explore would be to examine the effects of stimulation frequency on both central and peripheral blood flow as well as the EEG itself, at the same time.

Research can be fun as well as painstaking and obsessive, as simple or as complicated as you want to make it, but in my experience the road to results may not be straightforward, and the journey is unlikely to go exactly to plan. Indeed, as in life, sometimes it is the unpredictable that can tumble you into new, creative directions. You certainly don’t have to know everything when you start, but it’s good to know people who do know more than you and are willing to share or trade that knowledge. In fact, research of the kind I’ve been talking about is necessarily collaborative: it may not be *what* you know, but *who* you get to know on the journey that is important. If you’re shrewd, research also doesn’t necessarily have to be hugely expensive. And if you want to, you can get a qualification for what you do – although this is not something that’s been a major incentive for me, and I often hear grumbles that being part of an academic institution takes more energy than it sometimes provides.

[Slide 32]

In addition to the many researchers who have assisted me in my own journey, listed here, I would also like to thank my wife and family, for putting up with me when so often my head is full of whirling numbers instead of the business of cooking a meal, family dynamics, collecting grandchildren from school, or play.

[Slide 33]

Here, and on the next slide, you can see the references used in preparing this talk, and the image credits.

[Slide 34]

[Slide 35]

Finally, I would like to dedicate this lecture to the memory of Hugh MacPherson, whose contributions in the field of acupuncture research, including fMRI and magnetoencephalography, MEG, remain an inspiration to us all.

I hope you will consider using the EEG in your own research, and have found this talk interesting rather than overwhelming. Thrown into retirement by the pandemic, I have maybe had too much time to ruminate on the patterns and symmetries of life and research – although I hope this is not my swan song for the AACP.

Thank you very much.