

INTERVIEW WITH ROBERT KNIGHT AND BRIAN PASLEY:

BRAIN-MACHINE INTERFACES: NEURAL PROSTHETICS AND PATIENT CARE

Kapil Gururangan, Jingyan Wang, Prashant Bhat, Sushrita Neogi, Jared Rosen

Professor Robert Knight served as the Director of the Helen Wills Neuroscience Institute from 2001 to 2011 and has run a lab in human neuroscience within that institute since 1998. Using electrophysiological techniques, the Knight lab studies the role of the prefrontal cortex in human cognition. After graduating with a BS in Physics from the Illinois Institute of Technology, he earned his MD from Northwestern University, completed his neurology training at UC San Diego, and finished his post-doctoral work at the Salk Institute for Biological Studies. In 2010, Professor Knight helped found the Center for Neural Engineering and Prosthesis, a joint program between UC Berkeley and UCSF focused on brain-machine interfaces.

Brian Pasley, a post-doctoral student in Dr. Knight's lab since 2010, conceived the lab's latest publication, which announced the possibility for scientists to reconstruct human speech from recorded brain activity. Both Pasley and Knight hope the research will lead to advances in neural prosthetics for patients with neurological disorders.

BSJ: What has been the focus of your research in the past years? What does the Knight lab predominantly focus on?

Knight: My lab is mainly interested in the physiology of human behavior and over the years, we've looked at how the human frontal cortex works to regulate and control behavior.

BSJ: The frontal cortex is the region of the brain that is associated with decision-making or executive control mechanisms?

Knight: Evolutionarily, this region is heavily developed in humans – occupying about 30-35% of the cortical mantle. It has tremendous bidirectional connectivity to pretty much every region of the brain, so it's perfectly suited to receive input and deliver input back to many brain regions.

BSJ: We'd like to focus on the paper your lab has just published about reconstructing speech from the human auditory cortex. To start, how does this paper fit into your lab's general structure?

Knight: About 7 or 8 years ago, the lab started to do intracranial recordings, which became a big part of the lab. For these recordings, we actually have electrodes placed right on the cortical surface during an operation, to make sure no motor or language areas are damaged during the surgery, and for localization of epileptic activity. Naturally, we became interested in how the cortex works and how it supports behavior. It became fairly clear to me



about 4 or 5 years ago that because our electrodes were placed right on the cortex, we also had an opportunity to do brain-machine interface (BMI) research. I had been involved in recruiting Jose Carmena, who is a card-carrying world-expert BMI scientist in the Department of

Electrical Engineering and Computer Science (EECS), and we began to do work in motor brain-machine research. This wasn't because it was the main interest of the lab, but more so because we could do it and I felt that it would be important to have something in the lab that might have a more immediate translational application, a device that could actually help patients.

That's how it started. Brian joined the lab as a post-doctoral student two years ago and decided he wanted to tackle a more complicated BMI problem: understanding speech representation, both perceptually and imagined. That's how we got there.

BSJ: We were thinking that this had a lot of relevance to research the Gallant lab has been pursuing with encoding models and predicting what the brain is doing based on its activity. Did any of that work play a role in the conceptualization of this project?

Knight: I think it did. Brain would know more because he came up with this idea to tackle speech representation. I've been following the field of multivariate pattern-recognition for some time and was aware of Jack Gallant's work. The idea to use this approach for language really came from Brian. Brian was influenced by the notion that if you could do it in vision with fMRI, maybe you could do it in audition with direct brain recording.

Pasley: Both Jack Gallant and Frederic Theunissen have done a lot of pioneering work in this area. In fact, our coauthor in Maryland, Stephen David, was previously a graduate student in the Gallant lab. So, given this earlier work, we were interested to see if a similar approach would work with ECoG (electrocorticography).

BSJ: Can you briefly describe the experiment you guys did? Walk us through what a patient would have experienced in the procedure.

Knight: This is a very simple experiment. In the simplest condition, the patient is sitting there, awake, with these implanted electrodes over auditory areas and they hear either individual words or maybe sentences. Simple as that – they don't have to do anything, there's no task required. Let's say they heard 100 words. Brian would take half the words and use those words to try and fit models to see if you could reconstruct the word.

"I felt that it would be important to have something in the lab that might have a more immediate translational application, a device that could actually help patients."

A word comes in, it has a spectrogram – it has a sound profile and different frequencies depending on the speaker and the structure of the word. He wanted to know whether he could fit a model from the electrical activity in each individual electrode and turn that activity back into the frequency spectrum of the word that was perceived. Now, if you have unlimited parameters, you can get anything to fit anything. The question really is whether it is predictive. He tried different models and each one had different success rates and sensitivities to the type of word. I don't think that necessarily is so important to go into in detail now, since his approaches are nicely laid out in his 2012 PLOS Biology paper. The important part is that once he got models that worked, he presented to the models a word it had never heard before because it was not part of the training set. The key question now is whether the model can predict what it's heard.

The test presented two words, "guitar" and "orangutan", and we wanted to see at what degree of reliability the model would say one word was "orangutan" and not "guitar". That is where Brian was able to predict with roughly 90% accuracy what word had been heard, which is an amazing level of accuracy. In the sound files, I think you can actually hear the word presented to the patient, if you dream and fantasize a little bit.

BSJ: Intracranial EEG (electroencephalography) takes direct surface recordings from the cortex. Why is that advantageous compared to an fMRI?

Knight: You're right on the cortex. The electrodes are placed directly on the cortex, so you're right on to the neural source. You have two things that are quite nice: perfect timing and exquisite spatial information. The electric fields of the neuronal populations are changing in real time and we're able to capture that in real time. It's not like an fMRI BOLD response, which measures the changes in blood oxygen levels. There's a delay of roughly 5-6 seconds, whereas EEG can notice a sound that arrives at your auditory cortex in 10-15 milliseconds after entering the ear. Each technique has its strengths and weakness and BOLD is a very powerful technique and a wonderful addition to human research because it had whole brain coverage. We only have ECoG access to the area where our surgeons put the electrodes. We don't tell them where to put the ECoG electrodes because it's a clinical decision, so where the electrodes are is what we're limited to for our recordings. We rarely get any depth information, so we

wouldn't know what's going on in the basal ganglia or the cerebellum, where the fMRI could get at these areas easily. It's a trade off. I like being right on the cortex. It appeals to me as a neurologist.

BSJ: There are multiple layers to the cortex and there's a depth to that gray matter layer. How much coverage does the EEG give you of the auditory cortex?

Knight: The auditory cortex is infolded, so it's actually partially inside on the superior temporal plane. We actually didn't record from the primary auditory cortex, called A1. When the signals come from the thalamus, it ends up at A1 in about 15 milliseconds. We're recording in secondary auditory regions referred to as auditory association cortex (AII) on the lateral part of the hemisphere. The response onsets in AII around 50 milliseconds rather than 15 milliseconds, so there's already some processing going on before we get the information. The electrodes have recording surfaces of 2.3 millimeters, so that's a lot of cortex that we're recording from in each electrode. A cortical column is half a millimeter at most; so in a 2.3-millimeter diameter surface, you have a whole bunch of columns and millions of cells that are contributing to the electrode field we are recording. We're really looking at field potentials, not individual cells, and oscillatory activity rather than single-cell spiking. Now, the nice thing is that it's becoming clear and accepted that the high-frequency oscillations we record with ECoG on the cortical surface, which occur between 60 and 200 Hertz, are surrogates for measures of cortical excitability. There have been some interesting papers on animal neurophysiology that have come out and made that claim.

Brian specifically looked at the most informative frequency in the electrocorticogram, which was in the high frequency band between 60 and 200 Hertz, and what he observed was that certain electrodes had stronger correlation with certain frequencies or modulation spectrums of words. Brian came up with, I think, the cleverest way of making a simple analogy for his analysis approach. Beethoven, when he was deaf, could still produce music, and in his head, he could still hear the notes as he played them. If Beethoven watched you play the piano and hit the keys, he couldn't hear you but because he knew which keys were being pressed, he could, in his mind, know what piece you were playing. That's the same thing Brian is doing. He's looking at the ECoG electrodes, and they're all like piano keys. He's able to figure out which one goes with which note or which chord in the music analogy and then reconstruct the word.

BSJ: Is it the same when you think about word and listen to the word?

Knight: That's the \$64K question determining whether this work is going to become something useful as a prosthetic device. Right now, it's interesting and it's a very beautiful piece of science. We can reconstruct that you heard the word "hungry," that's interesting and that's what we did. The real question is if you can't speak because of a neurological problem and you thought the word "hungry," could we decode that? If you could do that, now you're really talking. You have the potential for the development of a true speech prosthetic device.

Is that possible? I think it is probably for a few reasons. We know that if you move your hand and then you simply think about moving your hand in the same way, the same areas are activated in your brain. It's the core principle of brain-machine interface. That's in the motor domain. On the perceptual domain, if you look at an orange, and then close your eyes and image the orange, the same visual areas are activated. Imagery activates the same brain areas. It's been shown for vision and olfaction. Hearing a word and thinking of the word in fMRI studies have also shown the same areas of the brain being activated. What we want to know is, can we reconstruct an imagined word? Brian is doing that research right now. We have some preliminary evidence that we might be able to do it. If that's true, then it's going to be a clear track to developing an implantable prosthetic device for a person who can think, who can formulate ideas, but cannot speak, due to a stroke or motor impairment, like Stephen Hawking's ALS. That would be the goal. We're already working with our Engineering colleagues. In fact, Jan Rabaey's group in the Berkeley Wireless Research Center is working on a sixty-four channel implantable, wireless, externally chargeable device, which is what you'd want to make this a real thing that a patient could use.

BSJ: Are there a lot of labs working with BMI's and the possibility of either using machines to control a brain or using the brain to control a machine.

Knight: I wouldn't say there's a lot, but the number is growing rapidly internationally. Jose Carmena is the lynchpin of the operation in Berkeley. I'd say my lab is growing into the area. There are other people who do work, but really it's the engineers who are getting interested in the devices that makes this the most exciting to me. For instance, Michel Maharbiz in EECS is getting involved with controlling beetles using neural circuitry and is a world expert in small recording devices. In terms of neuroscientists, there's not a lot of people doing BMI research explicitly at Berkeley. There are people whose work is directly relevant, but they wouldn't list themselves as doing BMI research. There are sensory physiologists, like Daniel Feldman, who's doing work

that is very salient to this problem, as well as Richard Ivry, who is a superb motor control neuroscientist in the psychology department. We now have a Center for Neural Engineering and Prostheses between UC Berkeley and UCSF to capture the talents of all these faculty with a common goal to make an implantable prosthetic device to restore motor and language function.

BSJ: There is a barrier, an economic barrier, between the research and discovery salient to this problem and presenting the research to the market. Is your group comfortable making that leap?

Knight: No, I'm not. We want to deliver the best kind of signal quality to the engineers and the development people. We may eventually want an industry partner. I could help with that, but I couldn't do the engineering and the devices. It's a classic multidisciplinary problem. You've got to know the neuroscience, the engineering. You want to be able to go from bench to bedside. This is a classic problem – you have to get the signal, make the device, and put it in the patient and we have the team now to do that. We're excited about that and now we just have to deliver the goods!

BSJ: What are your goals now after that paper?

Knight: Our internal goal is to have an implantable device for a speech and motor prosthesis within 5 years. It sounds like a long time, but in scientific terms, 5 years isn't that long and it would be amazing if we could pull it off.

BSJ: Our journal's topic for this issue is how science can be both beneficial and detrimental in terms of the results that it can produce. Can you speak to the benefits of your research – what has this research and similar projects created that has been beneficial to the human experience?

Knight: Well right now, the whole field of BMI is in its infancy. There have been devices that help people spell. I was involved in the initial development of the P-3 Speller about 15 years ago. Based on brain waves, patients can actually select letters and spell words.

We have had some success with EEG-controlled wheelchairs and computer screens. That work is being driven by John Wolpaw's group in Albany and José del Millán's group in Lausanne, Switzerland. There's a nice movie of one of his subjects driving his wheelchair with an EEG cap available online.

I would say the actual number of patients benefited by these innovations is just a handful at the moment – maybe 50 patients, maybe 100, maybe a little more. It's certainly not made a huge impact yet, but the potential number of patients that could benefit are staggering. If you just think of people who've had a stroke and who can't move their hand, it's not easy for them to open up a coke can with one hand or take off a bottle top. A simple device that would allow them to open and close their paralyzed hand would allow them to hold the bottle and open it. Those are major things to someone who's had their fluid interaction with the environment taken away. So, the total number of patients who could benefit is quite large. In the US alone, the number likely exceeds 200,000 people.

Again, I don't want to oversell this, because this is a new area and there are things to consider. First, is it safe? It requires a surgical procedure. Will people want it? I think the answer to the last thing is: yes, if it works, there's no doubt they would. Initially, people think all these

“my ethical considerations are that if I didn't do this research, it would be unethical, because I wasn't paying attention to the needs of the patients.”

ideas about machine-human interface are crazy. I've been around long enough to know that people didn't think implantable cardiac pacemakers and defibrillators would work back in the 1970's. They are now standard medical practice. You might have a loved one who has a

cardiac induction block and if the doctor doesn't put in a pacemaker now, it may be malpractice. There's an example of a cardiac signal, not conceptually different than a brain signal, used in a device that improves a person's function. The next important example would be cochlear implants. People were skeptical: you have 25,000 hair cells or neurons, and they're going to put in a stimulator that's got 64 or 128 contact points. People said it would never work. It turns out that the initial sound coming in is indistinct, but within a few weeks to months, the brain makes sense of the sound and the person can understand speech. Something that was felt to be unlikely is now routine. One or two kids or young adults probably received cochlear implants at UCSF this week. I think we're on the upslope of BMI research. We don't know what the steady state will be.

BSJ: A lot of discussions about your research in news articles and press releases have contained a small section about the detrimental effects of this research. They have called it things like “mind-reading” and “a polygraph.” I want to know what you think about the claims of where this research could lead.

Knight: I can assure you that every journalist has asked about the potential downside of this research. “What are the ethical considerations?” I say: “my ethical considerations are that if I didn’t do this research, it would be unethical, because I wasn’t paying attention to the needs of the patients.” That’s where I’m coming from. Now, I’ll speak to how it could be used “badly.” First, the stuff we’re doing requires major surgery – you have to be right on the brain. Would that stop somebody if they really wanted to get the information? Probably not. The second thing is that we’re only measuring what you hear, not what you think. You have to get to the next stage, which we think we can get to. Then the question would be: have you made the person you’re interrogating think of what you want them to think of? I think it is interesting to talk about “bad use of BMI”, but I do not think it is likely to ever happen. I understand that people get very nervous about these things, but I can’t control that. I can only control basic science. I’m not trying to diminish the concerns. I understand the concerns, but I don’t think the research is close to being at that point.

BSJ: It’s a very fundamental question of whether your brain can mimic its own activity without actually producing a motor action.

Knight: Yes, well we know it’s true for motor. I think it’s true for audition too, and we have some evidence in the literature to support that. The question is more technical: knowing when it happens, when the action onsets, when you average the signals. With our research, we need a nice time code to do our analysis. If you’re just sitting there thinking and your mind goes to “hungry” how do we know when to try to decode the signal? It’s a little bit more difficult signal analysis problem. That’s good because not everybody will try to figure it out. However, I will be ecstatic if tomorrow someone announced they figured it all out and they were a month away from having an implantable device.

BSJ: Because of the benefit to the patient.

Knight: Absolutely. I would be tremendously pleased for sure. It’s not easy to do this intracranial recording. I mean, the patients are sick. They are in the operating room or the intensive care unit and it’s not easy. For every 10 people that are implanted, we do not get very good data for 3 because they have problems. They may start seizing or be experiencing post-operative pain and need narcotics. You have to have a high tolerance of frustration to do this research.

BSJ: So these patients are epileptic?

Knight: The intra-operative surgical patients are generally people getting tumors removed. The other patients are all epileptic patients being monitored to localize their seizure onset for palliative surgery. A paper just came out last week showing tremendous improvement in seizure control using early surgery for epilepsy. They took groups of patients with uncontrolled seizures and had one group on multiple medications and the other group had surgery. They followed up with them for a few years and they found that the ones who had no surgery still had multiple seizures whereas 70% of the people who had surgery did not have seizures. This data is really going to push people toward earlier surgery. Which is good, because the longer you have seizures, 1) you are more likely to get hurt during a seizure because something could happen physically, and 2) your personality gets a little distorted due to a combination of things: subtle social ostracism, plus the repetitive seizure discharges – it’s not the greatest thing for the developing or adult brain.

BSJ: Is the brain just at firing at random?

Knight: Well, they are not at random. Actually brain cells get quite synchronized during

a seizure. The beauty of humans is that most of our neurons are random and we kind of put neural signals together into meaningful groups of activity. In seizures, everything gets activated at once and everything is firing away in synchrony. In a tonic-clonic seizure, the whole brain surges at once, which causes the synchronized convulsions.

BSJ: We had some questions concerning the basic science of the paper. Why are you putting the EEG on the association areas along the superior temporal gyrus rather than the auditory cortex itself? What was the reason for that?

Knight: In order to get to the primary cortex, you would have to cut the membrane and slide the electrodes in, which would be a risk for the patient and we wouldn’t do it. You generally do not get epileptic sources from primary auditory cortex, and it is not a region you have to map for language, so there is no reason for the surgeon to put them there. You know the phrase, “physician, do no harm.” We are not going to do anything just for the sake of discovery. If we had electrodes on the primary auditory cortex, we would have better, more precise reconstruction, but we

“The beauty of humans is that most of our neurons are random and we kind of put neural signals together into meaningful groups of activity.”

would never have electrodes down there in the first place. Maybe a rare case will need it, but I don't think it's going to happen.

BSJ: In using epileptic patients, is there a potential for their condition to affect the EEG recordings?

Knight: No, we are very careful to account for that in our analysis. You have to remember that people who have epilepsy look and behave normally in between seizures. Every so often the seizure discharge goes crazy and takes over the brain and that's when you have a clinical seizure, but the abnormal area starting the seizure is quite constricted. What we do in our analysis is to make sure not to include any area that is affected by epilepsy. We focus on the normal areas of the brain that are, by and large, 98% or more of the person's brain. You also want to make sure you're not going to analyze epochs of activity when an abnormal spike spreads all over the brain. When these things spread, they take over the whole brain, so we are very careful to remove them. To accomplish this, we go through every second of the EEG and make sure we clean it before we do any analysis. When I say we clean it, I don't mean that we are tampering it, but we are taking out sections that have epileptic activity. If you get a spike – a spike has a very sharp rise to it, which means if you do a power spectrum there is all kinds of energy in the high frequency band. The spike would be producing high frequency activity that has nothing to do with normal brain function and it would contaminate our recording of the critical high frequency signal needed for Brian's reconstruction.

BSJ: Where do you see this research going in the future, not only in terms of your research, but broadly speaking in terms of neurological research?

Knight: First a general statement. In the last ten to twenty years it has become apparent that we can get tremendous physiological insight into human behavior. I think the field is still in its infancy, even though people think it is very advanced. I think the big questions are yet to be poised and certainly not yet solved. That will have important implications for understanding normal behavior, how your child develops, how you learn, how you retain, how you age in a healthy manner. We will have a better physiological understanding of the normal development and aging life span, in addition to a behavioral one.

Second, these methods are going to have increasing influence in terms of neurological and psychiatric disorders. The dogma in psychiatry is that there is something wrong with your neurotransmitters. I think people are now beginning to realize that's probably a limited view. There

are more papers emerging on electrical dysfunction, oscillatory dysfunction, and impaired connectivity between brain areas. Think about depression, for instance. People can get really depressed, get treated and get over it, and then not be depressed for years, and then they get depressed again. Well, how do you explain that? The fact that it comes and goes doesn't fit very easily with "your chemicals are bad" or "your brain is deteriorating" because you're normal in between. It does fit with abnormalities of system biology, network connectivity, things being out of synch, so I think you are going to see real advances in both diagnosis and monitoring, perhaps even treatment. There are more papers emerging where people are putting stimulators in at certain frequency bands to actually improve neurological function. There was a paper a couple weeks ago in the New England Journal of Medicine on stimulation in the hippocampus in humans improving memory. I wouldn't be surprised if some of the research extends into that area. That is the broad answer.

The narrow answer is that the field of bio-machine interface is in its infancy. The field has been dominated by engineers, who are fantastic because they can do things and provide key devices. On the other hand, they don't necessarily care what the signal they use for control is. They care whether the signal differentiates. A signal from your nose and your ear would be okay if it allows you to go left or right, and that's a totally valid approach. In our view, to get the best BMI for motor control, you really have to understand the motor physiology of the motor cortex. We're focused on not just finding two signals to be orthogonal enough to be used to go left or go right. We're more interested in understanding how the human cortex supports motor and language behavior. We think in a long run, it will have a bigger payoff to the field of brain machine interface.