



Presidential Commission
for the Study of Bioethical Issues

TRANSCRIPT

Bernard Lo, M.D.

Professor of Medicine
Director, Program in Medical Ethics
University of California, San Francisco

Anthony Wagner, Ph.D.

Professor of Psychology and Neuroscience
Stanford University

Meeting 8, Session 7
February 3, 2012
San Francisco, CA

DR. GUTMANN: Yes, we are going to quickly switch speakers. So in this discussion we will focus on current clinical and research experience with advancing neuroscience, and I'm really delighted to have Dr. Bernie Lo with us who will introduce us to ethical issues in clinical care and Dr. Anthony Wagner who will be with us in a moment who will bring us up to date on recent scientific advances in neuroscience.

So let me introduce Professor and Dr. Bernie Lo. Dr. Lo is Professor Medicine and Director of the program in medical ethics here at UCSF. Thank you for having us at your institution. He's also the national program director for the Greenwall Faculty Scholars Program in Bioethics.

He's a member of the Institute of Medicine and serves on the IOM Council. He's been involved in a large number of studies on ethical issues and human participants research carried out by the IOM and the National Academy of Science.

He formerly chaired an IOM panel on confidentiality in health services research. He's a member of Ethics Working Group of the NIH-sponsored HIV Prevention Trials Network which carries out clinical trials in developing countries.

He is Co-Director of the Policy and Ethics Corps of the Center for AIDS Prevention Studies at UCSF which provides technical advice and consultation to researchers carrying out clinical research including research in resource-poor countries. Welcome, Dr. Lo.

DR. LO: Thanks very much. It's a pleasure for me to be here, and I hope you all enjoy the weather in San Francisco that Dr. Hauser arranged for you.

I'm going to talk about some clinical issues that followed from your discussion of self and selves, and I'm going to sort of drill down to actual clinical decisions and try to be a little more specific about what kinds of dilemmas patients and their families and doctors face.

I'm going to start with an example. This on the right is my favorite aunt as a young woman. That's my mom on the left. My aunt was a wonderful person. She was a chemistry professor, had a really inquiring mind, liked to talk about all kinds of things, and when she found out I was interested in bioethics, she was one of the few people in the family who had actually asked me questions and engaged me in discussions.

She unfortunately developed dementia, multi-

infarct dementia, got progressively physically and mentally impaired. This is my mom helping to take care of her, and there's a real change in her.

When she was younger we talked a lot and she repeatedly said she would never want to be kept alive if she had severe dementia. She talked about friends of hers, people she had visited, former colleagues, and I think for those of you in the law, her discussions would have risen to the level of clear and convincing evidence in terms of specificity and so forth.

When she became quite demented, I still visited her and I still interacted with her. She remembered who I was and remembered I had a son. Didn't remember much, would ask me about my son. I would pull out my picture book and show her pictures.

We have a nice conversation. She would smile. Five minutes later she would ask me about my son, but she was smiling, and she clearly was enjoying - as best as I could tell enjoying the interaction, so had she then adjusted to her limitations, found some sort of contentment in a situation that previously she had said would be intolerable and not worth extending, and her family later faced the

question of when she developed pneumonia should they follow her prior statements about refusing antibiotics, hospitalization, let alone a ventilator, or should they say her life still seemed to be valuable in a limited way, and it would be better for her to continue that existence.

There are lots of other dilemmas that perplex patients' families and physicians in dementia. Refusals of everyday medical interventions, when we all go to the doctor getting a blood drawing, going for an imaging study like an MRI or CT, taking pills are inconvenienced but not really burdensome, not really frightening.

When people have severe dementia, getting blood drawn, being put in a machine can be quite terrifying, and most much more important they don't understand what the benefit is and why it's important to do.

So again, if you think about people who have said previously if I develop Alzheimer's I really would want to be in a research study to help other people learn better to get better treatments, but when the times comes and she's eligible, she starts refusing all of the study procedures. What do we do?

Similarly people do give directives for clinical

care and sometimes they give maximum directives. They say I want my life extended. That's my religious belief that it's valuable even though I'm suffering, even though it may seem limited. I want you to do everything you can to extend that, but then if they're not - if they're balking at blood draws and taking pills, how do you actually do that.

There are other brain conditions where this comes up, and just recently in the past month two of my patients went into major depression. One of them recognized that she said I'm not myself. That's not who I am. I don't have the initiative, the energy, the upbeat hopefulness. That's who I am. The depression has taken that away from me.

Another one of my patients got better after anti-depressants and recognized that he was less depressed but also said I don't want to keep taking these pills. They make me a different person. He said there was like a fog going over his mind. He just wasn't thinking clearly, and he didn't want to continue in that situation even though he knew his depression was better.

So there are situations where you can sort of go back and forth depending on the relapses and remissions between one sort of self perhaps and another.

Let me just point out for psychiatric patients, people can sometimes recognize that their narrative if you want to say that involves lots of relapses to the point where they then have to be involuntarily hospitalized because of the extreme danger to themselves or others.

They may say when they're well if that happens to me, don't let me just keep getting worse and worse and make me involuntarily hospitalized. Do something before. Give me those medicines to prevent that precipitous line.

Do we honor that when the time comes and he's now refusing, screaming, cursing, kicking, and if we decide to honor the previous directive, how do you actually administer directives through refusing patients -- to administer medications?

I just want to point out other conditions aren't necessarily the brain involved. We want to then reformulate some of this in more ethics terms. First, an issue is what do you do when the prior directives of a person contradict what apparently is their current best interest as judged by others?

Do a person's advance directives, their previous values, no longer apply when they're "a different person,"

whatever that means, and pushing forward, it's disrespectful to override a competent - not competent, an adult patient who can refuse, and to what extent is it permissible to use undue influence, misrepresentation, or actually flat out coercion, so that's the ethical underside of honoring what you think is the patient's current best interest.

Secondly, there are also all kinds of other implicit promises that you make to a person and commitments made by that person, and again, in clinical medicine all the time, many people ask their families never to put them in a nursing home. That would be a fate worse than death they say, but later on as they physically and mentally decline, first of all it may not be possible to keep them at home, but they may actually seem content.

They're no longer repulsed by the smell, the appearance, of the nursing home, and have they changed? Should you still honor their previous refusal? I just want to put in the idea that behavioral scientists say that as people we are terrible at anticipating how we will adapt to new situations, so this whole notion of affective forecasting where we underestimate our ability to adapt and how do we factor that in.

I'm going to skip some of this. I want to highlight the idea that we need to think about what are the obligations of physicians and society or the health care system to facilitate continuity of self to the extent that these neurological psychiatric illness cause these changes in really fundamental aspects of people.

What kind of resources should we be committing to try and maintain the continuity of the narrative that the previous panel was talking about? Time and patience, and that's resources and that's money; ancillary and in-home services to try and keep people in a familiar environment as they become demented; support for family care givers. All those things I think are some pretty fair evidence and show that it prolongs, enhances a person's sense of who they are as their mental condition deteriorates.

Next I would suggest that we really need professional and legal standards. There was mention in a previous panel this notion of narrative and this idea that you judge a patient's best interest by sort of the arc of their narrative.

That's a very sort of appealing idea these days, and, Dan, I think you've actually written on that and talked

about it, but that's on the level sort of a philosophical maxim.

What about the level professional standards, legal standards, because that notion actually runs afoul of state laws and decisions in some jurisdictions, and then I think we need to ask a little bit harder what do we mean to say the authenticity of narrative. There are lots of potential endings to life stories. How do we choose what's authentic and what's not?

So I'm going to leave with some suggestions for you. First as a Commission, I would really urge you to take advantage of your unique role as a Commission to reach different audiences that wouldn't be reached by other people discussing these topics.

Secondly, I really encourage you to make recommendations to guide front-line decision makers. Your predecessor commission, the President's Council, addressed this issue to some extent in their report "Taking Care," but I really - if you want to tackle it, encourage you to sort of push further and ask under what circumstances is it permissible or even required to override previous directives that contradict a patient's current best interest and provide

guidance that really helps families and doctors making these decisions. Thanks very much.

DR. GUTMANN: Thank you very much, Bernie. We'd like now to welcome Dr. Anthony Wagner to update the Commission on the state of neuroscience and neuroimaging.

Anthony Wagner is Professor of Psychology and Neuroscience at Stanford where he directs the Stanford Memory Lab.

He is also Co-Director of Stanford Center for Cognitive and Neurobiological Imaging.

His research uses functional MRI and transcranial magnetic stimulation to examine the cognitive and neural underpinnings of learning, memory, and executive function.

He has published over 100 articles and chapters, and his research has been recognized through many awards including the American Psychological Association's distinguished scientific award for early career contribution and Alfred Sloan Research Fellowship and the Cognitive Neuroscience Society's Young Investigator Award. Congratulations and welcome, Anthony.

DR. WAGNER: Thank you. It's a pleasure and an honor to be here.

Having a bit of sense of self that I've built up over the years, I realize that I can run long, so I'm going to read prepared remarks to try to stay on point here or at least on time or close to it.

Before turning to my comments, I'd like to say just a little bit more about my background and my understanding of my charge here.

Again, I'm a cognitive neuroscientist at Stanford. My lab focuses on trying to understand the large-scale neurosystems that support learning and memory.

One of the systems that we focus on is the medial temporal lobe which is the circuit that's often impacted first in Alzheimer's patients, hence, their memory deficit.

Since 1994 I've been using functional brain imaging, principally fMRI which I'll say a little bit about in a moment.

Again, I direct the CNI, and kind of the mission statement of the CNI, this brain imaging facility which is sited on the main campus of Stanford, is to try to support the spread of neuroscience data across many disciplines, and our sort of tag line is neuroscience for society, and we take as sort of our charge to try to wrestle with sort of big

societal issues or try to at least generate data that might be relevant to major societal issues.

So turning to my introductory remarks, I took my charge today to seed a discussion about how advances in neuroscience may raise bioethical issues. Given my background, focus on advances stemming from modern neuroimaging techniques, I'll briefly at the outset just make or highlight the distinction between functional and structural neuroimaging.

Subsequently I'll introduce three approaches to using neuroimaging that have become increasingly prevalent in the field and that may have bioethical implications. Because I'm not an expert in bioethics or even in the more focused area of neuroethics, the goal of my comments is to draw the Committee's attention to these neuroimaging approaches.

So first I'll discuss neuroimaging efforts to try to detect mental states kind of related to this idea of can we actually from the outside know when John is about to cough even around the same time that he knows he's about to cough.

Here I'll briefly talk about fMRI lie detection and fMRI memory detection.

Second I'll introduce recent findings that raise the possibility that functional neuroimaging can be used to detect conscious awareness in patients that have been diagnosed as being in a vegetative state.

Finally I'll introduce how neuroimaging is being used to generate neuro predictors for different outcomes of interest that have implications for health, education, security, and employment screening among other areas.

So turning to neuroimaging methods, I didn't know the expertise on the panel, so forgive me if this is highly redundant. For human neuroimaging, there's two sort of gross classes. One is structural which are measures as you might suggest of brain anatomy, tissue composition, researchers are using structural MRI not only to detect frank neuro insults but also to get very subtle or to measure subtle individual differences in brain structure such as individual differences in cortical thickness that may predict a transition in a patient from going from mild cognitive impairment to Alzheimer's Disease.

Others are using structural imaging like diffusion tensor imaging to get measures of white matter integrity and try to relate these to other deficits such as dyslexia and

other reading disorders.

On the functional imaging side, there's been a revolution in the field over the last say 20 years or so. These are direct and indirect measures of brain activity as individuals are engaged in cognitive and behavioral acts, the dominant approach, up until the late '80s was scalp EEG. These are direct measures of electrical activity at the scalp.

You have very fine timing, but it's difficult often to know exactly where the signals are coming from. There are a number of other measures that have been developed in the '80s and into the '90s. The most important dysfunctional MRI with its advent in the early 1990s, it's revolutionized many fields. It's transformed psychological science, it's transformed neuroscience. We now can ask fundamental and basic neuroscience questions in the healthy human brain. We don't need to wait, excuse me, for patients of particular types.

I think as we'll see fMRI data are generating a number of possible sort of collisions between neuroscience and ethics.

So turning to the three topics, the first topic is

the use of brain imaging to detect mental states such as lies or memories.

One place clearly where neuroimaging and ethics may intersect relates to efforts used of fMRI to detect lies or to confirm an individual - to confirm that an individual is responding honestly.

Basic science on this topic has been conducted for about a decade. Based on this science and associated efforts, two companies now sell commercially fMRI lie detection services or truth verification services.

Over the past few years there's been an increasing effort to introduce these data as evidence in courts, and last year saw the first Daubert hearing in Federal Court in the Western District of Tennessee, and there the Judge, Judge Tu Pham, determined that the approach does not yet meet Daubert standards because its accuracy or error rate is outside of artificial lab settings. It's currently not known. Standards of use are not agreed upon, and it's not - the technique is not generally accepted within the scientific community.

My own assessment or analysis of the fMRI literature using these sort of artificial lab-based

approaches to assessing fMRI lie detection is perhaps even less generous or in that - my read of the literature is there's just no data right now that don't suffer from fundamental compounds that bear on this issue of sensitivity and specificity. I'd be happy to sort of talk about that if that's of interest.

This doesn't mean that I don't want to stress this. It's important to stress that it's quite possible with the right studies that the method may well be demonstrated to have pretty - above chance sensitivity and specificity.

Turning to a highly related area, it's the area of trying to detect knowledge or guilty knowledge if it's within a criminal setting, but knowledge more broadly.

Can we read out using brain imaging methods whether somebody has a memory for some past experience or some past event.

For neuroscientists, well this approach with EEG has been around for about 15 years or so, this sort of brain fingerprinting approach. For neuroscientists, this really came to fore with this Mumbai, India, case in 2008 where a woman was convicted of murder based on what appears to be largely scalp EEG data from a forensic science lab that the

State of Mumbai, India, runs, and these data have been introduced in tens of cases.

As far as I know, the method may have one peer review publication at best, so it's unclear exactly what the method is, but responding to a call from leading journals including Nature Neuroscience that neuroscientists need to step up and assess sort of these methods. My lab supported by the MacArthur Law and Neuroscience Project over the last few years has taken on the task of trying to ask well, scalp EEG is not likely going to be the method, the most sensitive method.

How well does - do fMRI methods combined with the use of machine learning algorithms to treat fMRI data. How well do they do in terms of detecting memory states?

So we've been conducting lab-based studies looking at memories of single stimuli encountered in the lab plus memories of rich autobiographical events using a camera-based technique in assessing how well we can sort of read out people's memory states.

In our first study we observed as indicated here that somewhat surprisingly you can do quite well in these highly controlled situations.

On average you can be 75 to 80 percent in discriminating between whether a person is currently recognizing the stimulus in front of them as having been encountered in their past versus perceiving that stimulus as being novel.

If you use the approach more conservatively, you can ramp up performance up to towards ceiling.

The information about memories is distributed in many - it's present in many parts of the brain. While this initial study might suggest that this technique might be viable, there are many caveats including caveats that emerge from some of our data.

It's very - well we can discriminate or detect whether somebody is remembering versus whether or not they're perceiving something as novel. It's much harder to know whether when they're remembering it's a true memory or a false memory.

There are a number of other issues such as knowing the source of the memory countermeasures and a whole host of other complicated issues.

Turning to the next topic, detecting consciousness in patients that have been classified as being in a vegetative

state, this isn't an area of expertise for me in terms of many of you around the room are M.D.s and so you'll forgive me.

My understanding is a vegetative state often sort of reflects partial recovery from a coma due to a traumatic brain injury, that a vegetative state, these patients often demonstrate sleep/wake cycles but they don't seem to respond volitionally to external prompts, external cues, as well as they don't seem to generate self-initiated action.

In a very high profile and perhaps somewhat controversial paper, initial paper from Adrian Owen and colleagues published in 2006 in Science, they brought a single patient who's five months post accident, five months in a vegetative state, and brought her into an MRI scanner, scanned her brain while she was listening to sentences relative to control stimuli and observed that parts of the brain that seemed to represent her process language were active during this sentence presentation.

They presented her sentences that contained ambiguous words such as creek and creak and observed frontal lobe activation and past work in the literature suggested that these frontal lobe responses reflect controlled efforts to resolve uncertainty or conflict during sentence comprehension.

Perhaps most strikingly however, I will be only two or three more minutes.

DR. GUTMANN: Take a few more minutes.

DR. WAGNER: Thank you. In the most striking sort of part of this paper is they instructed the patient to engage in one or two imagery tasks for a brief - for 30-second chunks of time alternating back and forth between the two tasks.

One task was imagine you're playing tennis, hitting a ball against a wall over and over again or imagine that you're walking around your house, and in controlled subjects as you can see here, these two different imagery tasks are associated with different patterns of cortical activation and what they observed is in this patient who had been classified as being in a vegetative state who didn't appear to demonstrate any overt behavioral responses, they observed patterns of activation that looked like she was responding to the cue and was engaging in the tennis imagery versus the spatial navigation imagery.

There's a whole host - there's a number of tricky issues here. Subsequent work I think has kind of partly answered some of those issues. They've gone on to show that in fact you can use - you can instruct these - some of these

patients, a minority of them, but some of them you can instruct them to use these imagery tasks or imagery approaches to answer yes or no questions about their lives - is your father named Alexander, and a subset of the patients they actually can accurately read out from these patterns presumably the patient's responses.

There still remain a host of sort tricky issues here, but it seems like it gets at some of these fundamental issues of self awareness, et cetera.

Okay, last topic, and this will be rather brief, the next approach that is emerging in the basic science literature as well as applied literature is trying to use either structural imaging or functional imaging markers as predictors of future health outcomes or other outcomes, and I'll just sort of illustrate this in one, I think, particularly striking case.

This is a study conducted by Talma Hendler and colleagues at the University of Tel Aviv and what they did is they brought in individuals, 18-year-olds, who were going into IDF doing their required service. They were going to be trained to be paramedics in IDF and spend two years in the IDF.

They scanned these individuals prior to going into the IDF and they asked these subjects as well as controlled subjects to give self report measures of stress, how stressed they feel, behavioral symptoms of stress, and then they went in and they were scanned 18 months - after 18 months of paramedic service in the IDF, and what they observed - and they collected these behavioral measures again, and what you see here is that the paramedics 18 months later there was a mean increase in self-reported behavioral symptoms of stress.

Controls didn't see that over the 18-month period. There was also the wild or marked variability across individuals in terms of this magnitude of change in reported stressful symptoms, and what they wanted to know is could you predict this magnitude of change in stress induced presumably by having spent 18 months and having these stressful experiences in the military. Could you predict them from the brain imaging data acquired prior to going into the IDF?

What they observed - the upper panel here and this is my last slide, the upper panel shows that post-service the amygdala activation, the amygdala structure important for a sort of salient signaling and affective processing and arousal as well as the hippocampus important for learning and memory

that the magnitude of activation in these structures post service were correlated with the change in self-reported stress, but more importantly on the lower panel what they observed is that amygdala activation 18 months over accounted for over 40 percent of the variants in reported change in experienced stress following the service, and so this is just one of many instances in which neuroscience data may be able to provide or may sort of be used to provide some form of prediction.

It clearly has some benefits such as possible early intervention, but one could imagine a number all sort of challenges. Thank you.

DR. GUTMANN: Thank you very much. We're open for questions from any members of the Commission. Nita.

DR. FARAHANY: Thank you for both of those presentations which were provocative and interesting. I have two questions which are somewhat unrelated.

The first one is tying together what both of you were speaking about, so, Bernie, thinking about how - whether or not current selves combine future selves and directive and then the recent consciousness studies and thinking about the ability to access individuals who may have had future

directives but may now in fact be responsive if only by something like an elaborate technical process, so we know that the persons that don't wish to be in a vegetative state are kept alive in a vegetative state. They're now in a complete vegetative state but were, in fact, able to have them be responsive to the fMRI assuming that we can stretch the interpretation of those studies to say that they're true. They are being responsive which, you know, I think we can given the difference between walking through rooms of your house versus playing tennis and the fact that that seems to be pretty responsive.

So how do we think about those two things together, right? So now we have a self which is a very diminished self locked into a body that is unable to respond except via neural activation that can be detected via fMRI, and how do we disaggregate and decide whether the previous self bound the current self or if even despite those clear directives we now have some kind of communication via fMRI instead? So that's the first question which I will put out there, and then if there's time, I'll go with the second one.

DR. LO: Well first I think it would be very interesting to try some of these new neuroimaging techniques,

and I take it now there's modified EEG that is our gauge to similar things and to pose questions and see what kind of answers you get. This you have to be very exploratory.

Terms of directives from a former self being binding on a current self I think we have to ask what kind of direction - directives, and what are the inherent limitations of directives.

So if I give a directive say to my wife, if I'm ever in a vegetative state pull the feeding tube, don't treat infections, but maybe I wasn't aware of this new data and especially if it turns out to be prognostically significant that if I show these there may be some evidence that I'm more likely to recover to a certain extent.

So if my directives cover didn't - when I made the directive, if I wasn't aware of pertinent information about my current state, to what extent should my family, my surrogate, my physician, literally follow those directives or use them as sort of a guide or override them because they might conclude Bernie didn't understand his current situation.

What he said then, doesn't really apply here.

DR. GUTMANN: Bernie, can I just ask just to make sure I understand the full thing. Do you feel - is this

symmetrical, so for example if I give a prior directive that I want to be maintained alive and then there's some evidence in from the fMRI studies that I want to die, do you feel the same way that my earlier directive should be possibly discounted?

DR. LO: Well we need to have something.

DR. GUTMANN: So I want to make sure we're not loading it on the direction that this example is taking which is someone wanting now to live. What about someone who say very clearly I want to be maintained alive and then in the fMRI state with some probability of not knowing the one-to-one correspondence between what lights up in the brain and what the, you know, the actual statement would be changes his or her mind?

DR. LO: The simple answer is yes. It should be symmetric.

DR. GUTMANN: Keep going. I just wanted to make sure that we've had that directional --

DR. LO: It seems to me that similar consideration so there are many patients who say for lots of reasons including very deeply held religious beliefs that they want their life prolonged no matter the quality of life, no matter what kind of suffering, but we've had cases where that patient

gets to a medical situation where just the daily care of changing their clothes, changing their linen, is really painful and it's not relieved by the things we usually do to relieve pain.

You could make the same argument, when they made that statement were they aware that being kept alive according to their wishes would entail terrible suffering that wouldn't be relieved with all the things doctors - you have to allow people to change their minds and to sort of cancel something they previously said, but if they can't speak for themselves, who has that right to overrule?

DR. GUTMANN: Anthony, before you - DR.

FARAHANY: Let me just ask you to build into your answer. This was a very splashy week for neuroscience news for mind reading, so while there was another study that came out of Berkeley about being able to decode words and minds, then there was the UCSD study about the EEG helmet which enables decoding of thoughts, and so all over the news media this week there was mind reading is now possible, and given your research, I was hoping, you know, in conjunction with this fMRI consciousness study to be able to speak a little bit to the quality of actually being able to read somebody's, you

know, mind and thoughts in either a persistent vegetative state or ends occurrence and uncooperative or cooperative state.

DR. WAGNER: So my response to your first question was first to note that I think the data right now are very impoverished in terms of what we know about these patients and what they - where they would have ended up if you waited in terms of prognosis, what their - whether they were misclassified and they really shouldn't have been classified as being in a vegetative state or now, but it's also impoverished in that most of the studies up until there was mention of sort of scalp EEG, a recent study in the last - in the Lancet using scalp EEG.

Up until this most recent study, the method for assessing and trying to get some insights as to what their state of awareness is and how high-level sort of cognitive function sort of what was preserved and what wasn't, it's very limited, right. It's a handful. Most of them I think it's five, six, seven questions, so you know very little about really sort of how broadly aware they are, how well or not they're functioning.

The recent study where three out of 16 patients,

vegetative-state patients, with scalp EEG could actually using a sort of imagine moving your thumb versus imagine moving your foot, you could decode the responses.

That unlocks the ability of science now to begin to try to weigh in and try to at least assess on an individual-by-individual basis awareness because now you don't have to bring them into the scanner. You can wear these nets that's long term and get many, many probes.

So if it proves to be the case that high-level cognition is sort of reasonably preserved in some of these patients, I don't know that it's a special case. I don't think it's that different from locked-in syndrome or any other syndrome where you've got to struggle with somebody who can communicate with you, a desire now relative to some past desire.

Now with respect to mind reading, you'll forgive me. I haven't been paying too much attention to literature this week, kind of prepping for this because it's going to be moving me outside of my sort of comfort zone and knowledge space.

Let me take a broader - yes, sort of Jack's Lab and other labs. Let me just take a broader - mental decoding

efforts are still primitive. You can decode what somebody is looking at. You know, are they looking at an image of a face versus an image of a scene versus a common every-day object.

You can decode pretty well. Are they imagining a face, a scene, playing tennis, et cetera. You can decode what noun they might be thinking or what verb they might be thinking of.

Once you train a classifier off of other data, you can get some basic sort of thought decoding. You can - some were suggesting you can decode which of one or two particular memories the person might be retrieving or bringing back to mind, but this is a very well - highly controlled where they only learned two events and they overlearned them.

So I don't know - and then you can begin to reconstruct what images they're looking at, et cetera.

I don't know data right now that indicate that you can decode very complex, you can decode complex thought, a fully-formed rich nuance sort of. There will be efforts to go in this direction. Whether ultimately the data will support that, it's not clear. I think decoding efforts will continue to improve, what we can read out.

There's a lot of information in these signals.

Once you start combining fMRI signals with EEG signals, I think you're going to be doing - we're going to be doing better and better, but whether we'll get to the point where we can read out the kinds of things that might be sort of make us uncomfortable or create legal or ethical sort of challenges, it's not clear yet.

DR. GUTMANN: Jim.

DR. WAGNER: Actually you took her question pretty much to cover most of what I was going to ask about, but just a real quick tag to that.

How much of the next - of future progress is based on the measurement technology and how much do you imagine is based on understanding the data that we have?

For example, we can develop the optical microscope as far as we want and we'll never image an atom. That kind of a question.

DR. WAGNER: That's a great question. Advances will continue to come from hardware advances, getting better measurements. I'm certain of that, but I think Martha Farah a very - if you haven't read it, a great article on neuroethics in the Annual Review of Psychology that just came out here in 2012 kind of carved up the imaging sort of field into

different decades, and, you know, the decade of 1990 was really a hardware decade, getting better MR - structural MR with higher resolution, improvements in magnet strengths, pulse sequences, et cetera, as well as functional measure. That was a measurement decade, and there may - there's still room for improvement.

The early 2000s, it's really a data treatment decade. There's so much information and so data rich, and most of this decoding work is realizing that one can take these machine-learning algorithms that have been used on many other kinds of data sets and apply them to these rich distributed patterns of data that we get, be it EEG or fMRI, and I think there's going to continue - there's going to be continued improvement, there's going to be I think integration of behavioral data into these algorithms with the neurodata, integration of multiple neuro - kind of neurodata, and I think there will be integration of predictions from math models or computational models of the way the system is working.

So I think that's one instance or one example of where improved analysis, treatment of data, and then the other I think is going to be the aggregating of data.

Right now most, not all, but the majority of

published studies are low-end studies. It's costly science. You've got 15, 20 subjects in your sample.

If it's a clinical study, you might have 50 or 100, but we're all accumulating these data sets from our studies. We have - there are thousands and thousands of volumes of structural MRI volumes sitting in research labs across the country, and there's beginning to be this effort that aggregate those data, apply information extractions sort of algorithms over those data, and I think that that's going to give rise to sort of important advances that we won't be able to achieve if we didn't have those large ends.

DR. GUTMANN: Thank you. Dan and then Christine, and then we'll take your responses and then we're going to take a break. Dan.

DR. SULMASY: First is just probably a comment that I suspect you'll agree with which is we've got to be very careful in these kinds of discussions we've just been having about using these techniques to get answers about decisions about whether to continue life-sustaining treatments on patients whose brains are already incredibly damaged for whom in the case of Alzheimer's Disease we already have all the evidence we need if the person does not have decision-making

capacity so no matter what they say on fMRI, we're not going to accept that either.

If someone is in the minimally-conscious state, then, in fact, using fMRI to try to detect whether they have the decision-making capacity to process the information so that we could then accept their answer is going to be an extraordinarily difficult judgment to make where it could work with somebody with a locked-in syndrome where clearly, you know, a spinal cord injury or something. We pretty well can judge that that's the case there.

So that's the comment. I suspect you'll agree there's a caution, and then the second is a question about our use of the term person here and different, particularly different person.

I think we have to be very careful about the metaphorical way in which we typically use that and then taking that metaphor too strongly, so, you know, John Donne's, shall we then say that we are not those persons which we were, so who's the reason we're breaking up, or we say it's this sort of primitive statement we use metaphorically to say that we're feeling sick, I don't feel myself today.

We don't really mean that we are a different human

being, different person, at that point. We feel there's a disruption perhaps in the narrative of ourself, and so when we take - do you agree that we take these things too seriously we really disrupt our ethical thinking, our philosophical thinking, and even our medical thinking about the care of patients who are severely brain damaged.

DR. GUTMANN: Bernie, what do you -

DR. LO: So with your first comment I totally agree that knowing that someone has cortical activity that seems volitional or intentional is very different than saying they have decision-making capacity to the extent we will respect their preferences.

I agree with your idea that we shouldn't take the everyday language that I'm a different person to have deep philosophical or moral impact.

I'm actually less concerned about is this the same person or not as the question to what extent should our decisions be guided by the patient's current values, preferences, directives versus what they said in the past.

I think, you know, it's really hard. People change over time. People can undergo changes in political philosophy, religious conversions, and yet somehow they're

still the same person. St. Paul is still the same person.

But how we would let their ideas guide the decisions they - we allow them to make or others make for them when they can't really decide for themselves.

Whether or not we call them a person - which values do we follow, and what are the standards by which we say we're going to follow these values, these decisions, versus those values.

DR. GUTMANN: Can I just ask a follow up because this is relevant to advance directives, right, but advanced directives specify a state that one needs to be in for the advance directive to take hold.

Are you suggesting that the state is mischaracterized because if I do an advance directive and I change my mind, I can change my mind. It's just that if I don't any longer have a mind that's in the state that the advance directive anticipates, if my mind is no longer there, then the advance directive takes over.

Are you suggesting it shouldn't?

DR. LO: Well let me sort of try and separate out your questions with a series of questions, so the first is when do we follow what the patient is currently saying versus

what the advance directive tells us to do, so that's a decision-making capacity or legal competency decision.

Then the question is if we say this person just can't make up their - can't make a meaningful decision, let's be guided by an advance directive.

Then the question is advance directives generally need to be interpreted so that first of all they don't always specify - sometimes they specify the surrogate, now ask my spouse to make the decision, and then that spouse has the power guided by in general what they think they would want, but often advance directives don't have both intervention and a situation combined so for me to say I wouldn't want any done if I'm in a vegetative state. Well what exactly did I mean by that? I don't want nursing care? I do or don't want a feeding tube. I do or don't want interventions for easily-treatable infection, and sometimes they say I wouldn't want the intervention, and as a physician we always say, well, you say you would never want to be in intensive care.

Let's suppose you broke your leg, you could have the hip repaired to good functional use, but you'd have to be in an intensive care unit overnight. Wouldn't you want that? Oh, yes, yes. That's not what I meant.

So advance directives, we have this hope or illusion that we don't have to make the decisions, the family, the doctor, does make the decision. The patient back there is telling us what to do, but language isn't ever that simple, rules are never that simple, you have to interpret it.

DR. GRADY: Thank you both very much. I think my question builds on actually both Amy's and Dan's question a little bit.

I found myself thinking, you know, in the clinical environment we very much value this assessment of someone's capacity to make decisions, and with capacity we give people choices, a wide range of choices, but I felt myself wondering in both of your talks if because of new thinking about the brain, number one, but also technologies that allow us to look at the brain in ways that go beyond what people can explain to us, are we facing a situation where those - the priority on capacitated choice is changing?

So one of the things I heard you say, Bernie, was maybe we should take into account contentment and values and preferences that are current even in an incapacitated individual as guidance for making decisions about that individual that might supersede capacity determinations with

the directives and all those other things as part of that story.

Then I wondered if I know maybe we're not there now, but with techniques of studying the brain and studying how people make decisions, are we even thinking - are we even - or should we be thinking about, you know, will we know more about how people make choices than they do and will we, therefore, use different criteria for deciding when somebody made a choice or not.

I mean I guess - I mean I'm sort of trying to think if we're re-thinking those - how we value-

DR. WAGNER: This seems like a very tricky issue because what you're wrestling with is the fact that much of the discussion today about the self has been really rooted in memory. Much of, I think, the earlier panel's presentation was making this point that the brain learns in many different ways We have different memory systems, some of which kind of operate much more automatically outside of awareness, and that can give rise to these automatic sort of responses that might make us content such as automatic affective responses triggered by the current situation is sort of this struggle of sort of willed action versus automatic or habitual acts -

habitual action as well.

So should we, I don't have an answer and shy away from given my background, weighing in on this, but should we take contentment that's derived from sort of more automatic memory processes that are not the kind of memory processes that give rise to rich sort of life narratives that are rooted in sort of gluing stuff together that were robbed in Alzheimer's disease.

How do we weigh that, and that was actually part of the - with respect to the vegetative state data that was part of the question. Are these automatic responses or are they richer things? How should we - so I don't have answer there, but I think it is an important sort of thing to appreciate which is we are - and I bring my biases here, we're not entirely but we're a lot our memories, but our memories aren't a singular thing, and so what we are are many different things, and if we loose one class of the things that we are, there still might be a lot of value or we might still way to weigh the other sort of things we are.

DR. GUTMANN: Bernie, do you want - have a quick respond to a very rich and - question that is not amenable to a quick response?

DR. LO: I just told Christine I would say that right now in clinical medicine, doctors all the time don't just literally accept a choice or a decision that sometimes we think that the decision-making process hasn't really been inclusive so very quickly sort tossed up a - mentioned colectomy, amputation, disfiguring surgery. Their patients say, no way, I don't want that. Let me out of here.

We tend to go back and say, well, did you realize that some people who initially say that later on are able to lead a very rich life although very different, so I think we always to intervene with not just accepting the first decision.

You also raised the question earlier in your very complex question about what moral weight do we put on the preferences or refusals by people who we don't think have decision-making capacity.

I think that's really an understudied and underanalyzed problem, so if you talk to emergency room nurses or psychiatric hospital nurses, all the time they're being asked to do something to a patient who is deemed to be subject to involuntary hospitalization because they no longer have the capacity and are dangerous, but their big and often strong

enough to resist, kick, scream, curse.

People who work with severely demented patients who don't take medicines they really need for their comfort, do they resort to hiding the pill and medicine and lying about what they're doing and things like that?

At what point do we say, you know, it would be so burdensome to sort of force this person to do this even though we have the legal authority that maybe we shouldn't do it on a very practical level that when you actually think on how you implement the care, it seems to be almost infinite.

DR. GUTMANN: I just want to thank you both very much for a very thoughtful and provocative discussion.

(Applause.)

DR. GUTMANN: We will take a 15-minute break or at least we will reconvene at 11:10, whatever that makes this break. We will reconvene at 11:10.

(Whereupon, the above-entitled matter went off the record at 10:57 a.m. and resumed at 11:10 a.m.)