Daniel L. K. Yamins, PhD
617.230.4765
yamins@stanford.edu
Jordan Hall, Building 420
Stanford, CA 94305

September 5, 2017

Francis S. Collins, M.D., Ph.D.
Director, NIH
National Institutes of Health
9000 Rockville Pike, Bldg. 1
Bethesda, MD 20892

Dear Director Collins,

As a computational neuroscientist, my work involves building neural networks models of human behavior and brain function in response to sensory cues, such as visual images or auditory sounds. For this reason, a core activity in my lab involves actually measuring behavioral patterns and brain functioning when humans are presented with such cues.

I am planning to apply to the NIH for funding to support my work in an upcoming cycle. I've never considered myself to be a "clinical" researcher, but given recent concern about new NIH rules regarding the definition of clinical trials[1], I am writing to inquire whether any of the typical types of experiments performed in my lab will now be considered "clinical." I am also writing to register my concerns about the rule and its interpretation, and to make a potential suggestion for how to improve the rule going forward. I am aware of many other letters that have been written to you addressing this topic[2], so here I attempt only to cover only my own specific situation.

Three typical experimental scenarios in my laboratory.

To make my questions concrete, let me first describe three typical experimental scenarios for studies we might perform in my lab. These include:

1. One hundred anonymous healthy adult human participants are recruited via an online platform such as Amazon Mechanical Turk (AMT). Each participant is very briefly shown a series of visual images containing one or more everyday object (an animal, a car, a boat, and so on), under a variety of visual presentation circumstances. For each image, the subject is asked to choose what category the object in the image belongs to, given a menu of possible choices. Subjects typically answer correctly, but sometimes they don't (e.g. they might reliably confuse a car with an animal under certain circumstances). The experiment aggregates data across subjects, measuring the average error rate for each image. It is not hypothesized that participating the experiment will have a measurable long-term change on the behavior of the participants, nor will the participants be brought outside their normal physiological and mental ranges during the course of experiment: the experiment merely seeks to measure normal adult visual behavior. These data are then compared to error rates for neural network computational models of the human visual system: a candidate computational model is judged to be "good" when it correctly predicts the pattern of human error rates across the images.

2. The same experiment as above, but with thirty healthy adult human participants recruited from the local community coming physically into the lab, and being scanned in an fMRI device while performing the categorization task described above. The purpose of this experiment is to obtain brain data that can be matched against the internal states of the computational model, with the model being judged as "good" if it correctly predicts the measured neural response patterns.

3. The same type of task as 1. above, but with adults brought into the lab for a sequence of sessions over the course of several months, and with two defined conditions: (A) one in which images are shown in random order, and (B) a second condition in which images are shown in continuous sequences (e.g. videos) generated by natural motion. No neural data is being collected in this experiment. Unlike experiments 1 and 2, this experiment hypothesizes that the

[1] http://fabbs.org/2017/07/27/basic-scientists-remain-concerned-about-nih-clinical-trials-definition/
[2] http://fabbs.org/wp-content/uploads/2017/07/Letter-to-NIH-Collins-Re-Clinical-Trials-FINAL.pdf

continuous exposure condition may lead to a measurable (if small) improvement in some particular visual recognition capacity that is cumulative and lasts beyond the course of the experiment. The goal is see whether the improvement of the human in condition (B) vs condition (A), should it occur at all, is the same as observed in the computational model when the model is trained with condition (B) vs condition (A). The model is judged as "good" if its learning curve mimics the human learning curve, whatever that curve might be.

The first experiment is purely behavioral, the second is gathers a combination of behavioral and neural data, and the third involves a potential for long-term learning.

Applying the new rules to my experiments.

I have been attempting to understand whether any of all of these studies trigger the four criteria recently established by the NIH for determining if a trial is "clinical". As you no doubt are aware, the four criteria are: (1) whether the studies involve human participants, (2) whether the participants are prospectively assigned to receive an intervention, (3) whether the experiment seeks to measure the effect of the intervention on the participants, and (4) whether the effect being evaluated is a health-related biomedical or behavioral outcome. If all four criteria are met, the experiment is deemed a "clinical trial."

Specifically with regard to my experiments, then:

- **Criterion #1:** All three studies involve human subjects, so they clearly meet criterion #1.
- **Criterion #2:** I am confused, however, about the application of criterion #2 involving interventions. Certainly, all three studies prospectively assign participants to receive certain visual stimuli. Showing someone an image is definitely a causal manipulation of their sensory system, since the presence of a real-world object in their visual field will cause their retinas and visual brain areas to react (for about 200 milliseconds) in certain ways that are quite different than if they had been shown a gray monochrome background, or if they had just kept their eyes closed.

  Before reading the Case Studies recently published by the NIH[3], I would definitely *not* have guessed that the stimulus presentation activity in my studies 1 and 2 would be considered an "intervention" for clinical purposes. I've always thought of that term as being akin to medical treatment, e.g. something that, if improvement on some metric is established, might lead directly to a new recommendation for patients in a clinical setting. Experiments 1 and 2 are definitely not a "treatment" in that sense, since participating in the experiment is extremely unlikely to have a differential effect on the future physiology of the participant, nor is there a metric on which "improvement" would sensibly be measured. However, Case Study #18 indicates that *merely asking subjects to perform a working memory task counts as an intervention*. My task involves a working memory component, since the participants have to remember what object they saw until queried later about its category. Am I right in thus thinking that the task in my Studies 1 and 2 *does* count as an "intervention" for the NIH's purposes, even if the experiments will not have any kind of lasting effect on the participants after the duration of the experiment, nor were the participants brought outside the range of their normal physiological or behavioral range during the course of the experiment?

  Unlike study 1 and 2, the activity of my study 3 definitely does seem to me to count as an "intervention" in the usual sense of the term, because it very well might have a long-term behavioral consequence. (On the other hand, this intervention does not seems directly "clinically relevant" in the usual sense of those terms, a point which I will return to below.)
- **Criterion #3:** All three studies seek to measure the effect of the "intervention" on the participants, namely, what category choice they make on each image. Of course, in Studies 1 and 2, the "effect" is not long-term, since the changes in physiological state of the brain's visual system upon seeing an image and (perhaps) holding it in one's memory only applies during the course of a single trial in the experiment — that is, on the order of hundreds of milliseconds or at most seconds. Before having read the Case Studies, I would have though that "effect" would be intended to connote a long-term change in the participants physiology or behavior. However, Case Study #18 suggests otherwise, since in the working memory task (as in my visual psychophysics task), any "effect" disappears as soon as the experiment is over and the "intervention" ends. Again, with my study 3, the case is different, in that the experiment definitely seeks to measure a long-term effect that might outlast the experimental timeframe.
- **Criterion #4:** I am unsure whether the effects in any of the studies being evaluated are "health-related biomedical or behavioral outcomes". Of course, it is expected that in the medium or longer term my experiments will lead to increased understanding of the normal adult human visual system, which will in turn form the basis for future proposed treatments of visual impairment that *would* seem to naturally need "clinical trials" designation. However, I would have thought that none of the three basic-science studies described above measure "health-related outcomes" directly enough to be designated "clinical".

[3] https://grants.nih.gov/policy/clinical-trials/case-studies.htm

However, Case Study #18 appears to indicate that merely measuring "brain function," as is done in an MRI scanner, counts as a "health-related outcome." This seems odd to me, and quite mistaken even in the situation of Case Study #18 itself. Brain operation is in some extended sense "health-related", just as any physiological or behavioral state is "health-related".[4] However, there is no specific "outcome" being evaluated in Case Study #18 — and certainly not an outcome that might have direct clinical implications. Just because a brain measurement is involved doesn't somehow mean that a "biomedical" treatment outcome is being evaluated. Measuring "brain function" in an fMRI machine, as in the Case Study description, is simply a form of *monitoring* the normal state of affairs in the participants' brains during the course of experiment. The working memory experiment in the Case Study will (a) not have a long-term effect on the person's physiology or brain state, (b) does not bring the participant outside the normal range of their behavior or physiology at any time, and (c) there is no "treatment" vs "control" scenario, such that success on the "treatment" would lead to a direct clinical recommendation. I would think at least one of these would have to be true for an outcome to really be "clinical" in any meaningful sense, and the fact that brain monitoring is being performed doesn't change any of these considerations. I very much hope the NIH will reconsider the interpretation in Case Study #18, as it seems scientifically inaccurate.

As for my experiments, unfortunately I suspect that the same reasoning would facially appear to apply to my study 2 described above, simply because it involves using an fMRI machine. Is my suspicion accurate?

However, I am especially confused about, and would like guidance on, whether my Study 1, which is purely behavioral, would also count as having a "health-related outcome". There is no example in the Case Studies that directly addresses this situation. However, if merely monitoring brain operation with an MRI during an "intervention" counts as a "health-related outcome", why wouldn't monitoring behavior also count as such? In fact, behavior is in general *more* directly related to clinically-relevant outcomes than a given MRI scan. If the latter counts as "health-related outcome", then the former should as well. But wouldn't this have the consequence of classifying effectively *any* basic research study in sensory psychology a "clinical trial"? Of course my response to this is that the premise of counting either study 1 or study 2 as such is flawed.

Finally, I am also very reluctant to characterize my study 3 as having a "health-related outcome". That is because, even though the intervention might have a long-term effect on some particular behavioral capacity, there is no obvious sense in which that result would directly lead to any particular clinical recommendation. After all, seeing random image sequences might be less effective for learning than seeing continuous visual motion (a hypothesis of the experiment), but the potentially "better" (or at any rate, more learning-inducing) continuous motion scenario is *already* what every human gets anyhow during their normal course of life. In other words, the experiment is seeking to determine how normal learning occurs, for the purposes of basic science, rather than immediately making a direct clinically-relevant comparison. In such cases, the added scrutiny applied to clinical trials seems like massive overkill.

Consequences for funding opportunities.

I'm also confused about consequences of the new rule on funding opportunities for my work. I have read guidance from NIH that suggests that funding opportunities available for clinical trials will now be limited to Funding Opportunity Announcements (FOAs) targeted specifically for clinical trials.[5] Would this mean that I would not be able to submit any of my above studies as part of a standard RO1 Research Project Grant at the NIH National Eye Institute (NEI),[6] which has funded basic research studies in visual psychophysics, like the ones I describe above, for many years? Does this mean that unless the NIH plans to release clinical FOAs specifically targeted at visual psychophysics that I would no longer have an NEI funding opportunity applicable to my research program? How specific will the new FOAs be? Will there be "clinical trial" versions for each existing NEI granting mechanism, including the ones that previously covered basic research that is now being reclassified as "clinical"?

Potential alternatives.

In summary, I am uncertain about the status of my own experimental program. I am also more generally concerned that there are misunderstandings in the NIH's interpretation of its own rule — especially as exposed in Case Study #18 — that would lead to wholesale misclassification of psychology and cognitive neuroscience studies. The core problem is that neither the rule criteria nor the interpretations of it make any explicit reference to *direct clinical relevance*. **As a result, it appears that many well-motivated basic research studies, especially in brain and behavioral sciences, will meet**

[4] I would also think that the educational outcome measured in Case Study #22 would be health-related in this larger sense too, and find the outcome suggested by the NIH in that case a bit confusing (though it is not directly relevant to my situation).

[5] https://grants.nih.gov/grants/guide/notice-files/NOT-OD-17-043.html

[6] See https://nei.nih.gov/funding/neifm

the four criteria for being deemed "clinical trials" even when they have no direct clinical implications. This is a highly undesirable situation for both researchers and the NIH.

Simple ways to avoid this situation would be to either change the interpretation of some of the terms in the four criteria to implicitly test for direct clinical implications, or to add a criterion to the list making this point explicitly. For example:

1. The definition of "intervention" could be modified to require that it (a) have some effect beyond the duration of the experiment, or (b) be testing a condition that is outside the range of normal functioning of the population being tested.

2. The definition of "outcome" could be changed to require that the study seeks to contrast at least two experimental conditions, measuring differential results between the conditions on one or more metrics that would be relevant for some clinical practice or treatment.

3. A fifth criterion could simply ask: "Does the outcome being evaluated have any direct clinical implications?"

Adopting one or more of these changes would resolve what I believe is an apparent major inconsistency in the new NIH rule, while causing very little "collateral" damage. For example, few if any of the Case Studies (aside from #18) would change their outcome.[7] After all, essentially all of the other Case Studies that are deemed clinical trials already explicitly assume a type of clinical scenario in their construction that is missing (or mistakenly found) in Case Study #18.

* * *

One of the few clear rationales that I have come across for designating basic behavioral and cognitive neuroscience experiments as "clinical trials" is that doing so would require that the studies results be registered in the NIH Clinical Studies database, thereby increasing experimental rigor, transparency, and integration. This seems like a well-intentioned idea. I would welcome having a standard place where the results of my experiments could be posted and browsed by the public. However, shoehorning what are not really clinical trials into the clinical framework to achieve this end is both scientifically inaccurate and has unnecessarily burdensome bureaucratic side-effects. There are a number of scientist-led efforts on experimental registration databases[8] that would likely be significantly more cost effective solutions for achieving research integrity and transparency. The NIH should probably look into endorsing one of these (and perhaps requiring NIH-funded neuroscience or behavioral research to participate in one) before deciding to significantly dilute the existing Clinical Studies database with a lot of basic science studies that will not have immediate clinical implications.

* * *

I apologize for the length of this letter and for its sometimes unsatisfied tone. However, I'm sure you can appreciate the seriousness that this issue has for me, and for many other researchers in a similar situation. I look forward to hearing your response.

Yours sincerely,

Daniel L. K. Yamins, PhD
Assistant Professor of Psychology and Computer Science
Faculty Scholar, Stanford Neurosciences Institute
Stanford University

[7] With the possible exception of #22, since "educational results" might sensibly be considered outcomes of clinical relevance.
[8] `https://openfmri.org/`