

NICK CHENEY: I'm Nick Cheney. I'm finishing my PhD in Computational Biology at Cornell University. Gabriel Kreiman and I are interested in seeing how deep networks respond to neuroplasticity. So in the brain, we know that the brain is constantly in flux. And neurons are growing and dying. Weights are changing in response to stimuli. But most of the time, machine learning, what we do is pre-train a network on some training set. But then we want to use it for real-- we freeze it and keep it in some static form.

There's been a lot more emphasis lately on more online learning so that you learn as you're working on a data set. And in those kind of environments, we think that the network will be changing quite a bit. So we're looking at how that network could be robust to those kind of changes much like the brain is to every day stimuli and actions it sees. To start out just doing a very simple test looking at perturbations of the network.

So just throwing random changes at the weights that make up this network and seeing how that affects its ability to classify images. After that, we're looking at how different parts of the network respond differently to these kind of changes. And then, ideally, we'd like to have some kind of rule that doesn't affect the performance of the network that much and it's able to maintain its ability to classify throughout seeing a number of stimuli.

So we know that the brain has certain learning rules like Hebb's rule, in which neurons that fire one after another end up strengthening their connections or conversely weakening their reactions. Then we're soon going to see if rules like that end up providing stable perturbations where the network can easily recover and maintain what it's doing or unstable ones where we're going to go down on track.

Or we know that deep networks act similar to how the brain works. Jim De Carlo gave a great talk about how the features we see in deep networks are similar to the features of the brain. And we know that the brain is constantly undergoing these kind of changes. So we're curious scientifically to see how these computer models respond in understanding how these two systems are the same or different. But also from an engineering context online learning where the network is changing while it's learning is going to be, I think, a much larger part of the use of machine learning going forward.

So understanding how stable these things are to constantly changing parameters, I think, will

be quite informative for those kind of studies. Being able to explore new types of materials and learn a lot about both computer vision and neuroscience has been a lot of fun. And, certainly, informative deep learning is a very hot topic right now. So being able to dive hands in a little bit and get some experience working with these models and some of the latest software packages, I think, will be useful going forwards too.