

Paul Werbos: The National Science Foundation and AI

<https://mindmatters.ai/podcast/ep138>

Robert J. Marks:

How does the NSF, the National Science Foundation, steer research in artificial intelligence? That's the topic today on Mind Matters News.

Robert J. Marks:

Greetings. I'm your ample host, Robert J. Marks. The National Science Foundation or NSF is a powerful force in guiding research in academia in the United States. A grant from the NSF is very prestigious. If a new professor is awarded an early career award from NSF, they got tenure almost certainly. So what's going on at NSF behind the scenes? Dr. Paul Werbos, our guest today, is the inventor of error backpropagation, which is the most commonly used technique used to train artificial neural networks today. He served for over 30 years as a program director at NSF. And during this time, Dr. Werbos steered NSF's funding in much of its machine intelligence research. And he's our guest today on Mind Matters News. Paul, welcome.

Paul Werbos:

Hi Robert.

Robert J. Marks:

Hey, what year did you retire from NSF? You said you're retired.

Paul Werbos:

Yes. So it's a long story. I started there in 1988 and my retirement - I came home to work for a new boss, my wife, on Valentine's Day 2015.

Robert J. Marks:

2015, so about six years. So you spent before that over 30 years as a program director at NSF that largely steered NSF's interest and research into artificial neural networks. During your tenure at NSF, what did you see the major turning points in machine intelligence? You mentioned last podcast, the genesis of deep learning as one of them.

Paul Werbos:

Yeah. It became obvious to me on day one. My understanding of how a brain works and what real intelligence is requires a lot of steps forward. And I wrote down these steps. And the question is, how long will it take for our science culture to catch up to even half of it? And I do have papers out there with roadmaps. One of them goes a hundred years into the future. And I know how to get there, it's just people have to do the work. There's just a lot of steps. And so I would say my whole tenure, my 30 years at NSF, I went through one step after another, after another. It's like every four years, it was like a different world. And the steps are not all done yet. In the very beginning people did not believe in neural networks almost anywhere.

Robert J. Marks:

And why was that? I think that that was because of the terrible fall of interest in neural networks. And I kind of time it to the publication of the book *Perceptrons* by Minsky and Papert that dried up the research. And so people were disillusioned.

Paul Werbos:

Yes.

Robert J. Marks:

And so neural networks had to be reproved using the algorithm that you created in order to show that many of the objections raised by Minsky and Papert were wrong. And so that was what you were up against, right?

Paul Werbos:

That is pretty much the story, but there were a whole lot of other little stories too. When I had that little talk with Minsky I mentioned.

Robert J. Marks:

Yes, on the last podcast. Yes.

Paul Werbos:

Yeah, yeah. I was talking about why it was that none of those neural networks they came up with were working. And I basically said there are a couple of things you got to change, but number one is it's a bad model of the neuron. You're assuming the neuron is a square wave generator. It's like you think it's a digital thing. Everybody believed you're scientific only if you're digital. You output ones and zeros, you're scientific, anything else is not scientific. And our brain being intelligent must be scientific. So it must be based on ones and zeros.

Robert J. Marks:

And that's right. Your error backpropagation algorithm, although highly digital, is based on continuous partial derivatives in calculus.

Paul Werbos:

So in front of Marvin Minsky, when he repeated, "Oh yeah, but everybody knows the brain is one zero." I said, "You know what? I took this course in neuroscience at Harvard. Let me show you the textbook. Let me show you the time series of the outputs of these neurons. Do these things look like square waves to you? Look at these things. Let's build a model that fits the data." And the data shows that it's basically frequency coding. Above a certain level there's a limit to how much output you have. There's a lower limit and upper limit. And you can pretend it's linear in the intermediate range, but the point is it's continuous. It's not one zero. You have continuous variables. And if you do it that way, then I said, "I have a new way to calculate derivatives. And I can prove that it'll work." So we have a provable way to do the derivatives. We have a provable way to train it, and it fits the empirical data.

Robert J. Marks:

But Minsky didn't like your argument.

Paul Werbos:

And Minsky said, "But all the computational neuroscientists know the brain is a square wave generator. And they all know it's ones and zeros." That's what McCullough told us. And if I disagreed with the neuroscientists, they would laugh at me because I'm not from the right field. I'm not from their field. You've got to be a neuroscientist to get away with stuff like that. And I said, "But what if there's data?" "Even data won't change them. If you're not one of the right people, they won't listen to you." That's what he said.

Robert J. Marks:

Oh my goodness. Okay. So what were some of the other turning points you saw? Could you elaborate on the deep learning? I think the deep learning is largely credited to Geoffrey Hinton who came up with convolutional neural networks. What's that?

Paul Werbos:

Oy vey. I'm not Jewish, but there are times when you got to say "oy vey." So there is so much false history out there it's just pathetic.

Robert J. Marks:

Well, I've seen even an article where Geoffrey Hinton was credited because of his involvement with the PDP book as being the originator of error backpropagation.

Paul Werbos:

Yeah. There's a whole lot of stuff.

Robert J. Marks:

Which is not true. You are the one that did that.

Paul Werbos:

No. And not only that, it wasn't independently rediscovered either. That's also false. I don't want to get into the whole history of who did what with what motive. It's not very constructive.

Robert J. Marks:

Well, we talk about the book Talking Nets, where you get into that.

Paul Werbos:

Yeah, some of that, right. There's a lot more. They thought that was racy, they didn't even see the half of what's happened. But let me come back.

Robert J. Marks:

Okay.

Paul Werbos:

The PDP books do have a role in history. There's no question that the history of the field was strongly influenced by the PDP books. And I think that that old boss of mine was involved in funding them

somehow. I would talk to you about that too. But bottom line is these books were influential, no question about it. And the cognitive science part was really very insightful in cognitive science. But the main machine that Hinton was pushing was something called a Boltzmann machine. And the quick summary is it basically didn't work. It was based on not understanding the math.

Robert J. Marks:

I remember Boltzmann machines, and we're going to talk about the kneeling later on, which I think he used in the Boltzmann machine. But let's not go there now.

Paul Werbos:

The bottom line is you got into Hinton because of something else.

Robert J. Marks:

I got into Hinton because convolutional neural networks, which I think by the way, convolutional neural networks, I don't know who came up with the idea, but they're brilliant.

Paul Werbos:

Okay. Convolutional neural nets became popular when I was at NSF. And I have to admit the first that I heard of the standard convolutional neural network, I thought of it as a kluge. So see it works. And it was something I wasn't backing. Where did it come from? It came from AT&T Bell Labs.

Robert J. Marks:

Really?

Paul Werbos:

And if I had to guess, I would say, I think it came from a French woman, to be honest, either that or an Israeli woman. They had a whole group. The Bell Labs group was really amazing. And certainly Yann LeCun was a key part of that group. There's no doubt. Yann LeCun was a lot of the success of that group, but there were a lot of other really important and creative people there whose names should come to my head. But the bottom line is they were the ones who took this idea, I believe it started in France. I think it was the group that was working together in a small company in Paris. And then they moved to Bell Labs where they became routinized. And I think Yann LeCun has some personal affection for Jeff Hinton who was kind of bringing him to the world of cognitive science, which is an important world, connecting him to a community. But to be honest, I really don't think the convolutional network came from Hinton. I'm pretty sure it came from this group in France.

Paul Werbos:

And then with LeCun. And for many, many years, I certainly cited this group. In fact, I remember there were 15 authors of the seminal paper on the convolutional network. And there were early tests by the post office on who can recognize zip code characters. There were two groups that just beat the world in the competition that the post office set up.

Robert J. Marks:

Oh, that would be a perfect application for convolutional neural networks. Wouldn't it?

Paul Werbos:

Yeah. And that's sort of where it started, this huge zip code recognition competition funded by the post office. And I later worked with the guy, I believe he was a Chinese American who moved to UTC who led that competition. There were so many false papers about that competition. So many people said, I won this competition with X. I won it with Y. And so I went to the guy who actually ran the competition. I said, "Hey, you got the scores. Who won it?" He said, "Two groups won it. And they both had convolutional neural nets from independent sources." One German, one French originally. And I believe that LeCun's group was one of the two that really won the competition. It was a huge breakthrough. They did much better than anybody ever did on that problem before.

Paul Werbos:

But it was just the zip code recognizer, because the original convolutional network had certain limitations. And we had ways to go beyond that. In fact, we have networks today that are much more powerful than the convolutional networks. And in the test, the new networks are much more powerful, but most of the people graduating computer science don't know about them.

Robert J. Marks:

You mentioned deep learning as one of the things that took off during your tenure at NSF. Could you differentiate a convolutional neural networks, which are a specific case of deep learning from the field in general?

Paul Werbos:

One bad thing that happens when you make it to Oprah and everybody reinvents you, they come up with 20 definitions of the same word. You used to think you had a word, and then all of a sudden, there are 20 different definitions.

Robert J. Marks:

The one that I remember on that, Paul, is artificial intelligence. When I was a boy in artificial intelligence, artificial intelligence meant the stuff that Minsky did, which was a rule-based sort of thing. If this, then this, if this, then this. And we didn't like this in connectionist. I think you would probably agree. That's the reason that IEEE started the word computational intelligence to differentiate from artificial intelligence.

Paul Werbos:

Well, I think yeah, you're the guy who came up with the computational intelligence.

Robert J. Marks:

Well, yeah, I was on a team of three people. I think Jim Bezdek and Pat Simpson and I came up with it in an email exchange. But the point is, is that to your point, the definitions change. Today, artificial intelligence subsumes all of this. It subsumes neural networks. It subsumes expert systems and everything else. So there is, there's an evolution.

Paul Werbos:

It assumes a hundred cultures that use different definitions. The same thing with the word consciousness, and even the word love means different things to different people. Certainly the word God means different things to different people.

Robert J. Marks:

Well, yes, okay.

Paul Werbos:

Yes, indeed. And then the question is, which of these things are the real ones in which are the made up definitions? And believe it or not, we have that problem even in the world of neural net. So to me convolutional neural nets and deep learning, I think they're well-defined words, but most people are hearing different things these days. To me, what the deep learning originally meant is just you have many, many layers typically with a feed forward network. And we were doing many, many layers long before these people were using the term deep learning. I was funding, supporting - and actually I wrote the original mathematics, because with one or two layers everybody could see just by eyeball.

Paul Werbos:

If you wanted to do it with many layers, you got to know the math. So I funded a guy named Krishna Kumar, working with NASA who applied many, many layers. And he was doing this long before any computer scientists had even heard of the term deep learning. And he didn't even call it that. I didn't call it that. I think my paper in ICNN in 1988 was the first reasonably open publication of what you can do with many, many, many layers and how to do it. But there's a problem with that. Okay? Yes, we can use many, many layers. The mathematics of many, many layers is good, but the way the brain does it is different. The brain has a more powerful way to do the same thing using the same math, but more general. The kind of recurrence that we have in the brain can do anything a deep network can do and basically better.

Paul Werbos:

So understanding how to really use recurrence is much more powerful than what they're doing with the deep learning. But having many layers can be useful and there's good mathematics, and it's the same mathematics. So that's deep learning. The term deep learning became popular when a bunch of computer scientists learned 20 years later stuff we were doing before, but they applied it and it was useful. They have every reason to be proud that they brought it to application. But convolutional networks are totally different. Convolutional networks are something I don't think I really understood until the late 1990s.

Robert J. Marks:

Well, you beat me. I only understood them about five years ago. So, okay.

Paul Werbos:

I think I got something like seven patents about just the late 1990s, and those patents already address many of the technologies of the future that they still haven't fully caught up to because they're just a whole lot of pieces.

Robert J. Marks:

Oh, how many patents do you have?

Paul Werbos:

They're probably expired by now, but probably about, I don't know, seven, eight, something like that.

Robert J. Marks:

Did you get these when you worked for NSF or are these after?

Paul Werbos:

Yes.

Robert J. Marks:

So it's assigned to the government, I suppose?

Paul Werbos:

You asked about how NSF is doing it.

Robert J. Marks:

Okay. Yes.

Paul Werbos:

And I'm glad I won't have to get into every gory detail because they got ups and downs like the whole world. But for many, many years, the Office of General Counsel, the number one lawyer for NSF, there was a guy named Charles Brown. And when I think of the greatest glories of NSF in its absolute peak period, Charles Brown, the general counsel, had a lot to do with it. He went to the Harvard Law school. He learned about freedom. He learned about free speech, intellectual truth. I mean a whole lot of very basic principles. He understood the spirit of the Constitution. And his way of dealing with regulations and rules, he respected the rights and the privacy of everybody and the integrity of the system. So at some point the idea came up, gee, I've got a few ideas. What should I do about patenting? And I talked to him.

Paul Werbos:

And he said, "We now have a clear policy." And in those days we did. He said, "The clear policy is if you come up with a patentable idea, we want you to get a patent. And it won't be an NSF patent. We will have to make a decision, whether the idea came from your work as an NSF employee, or did it come from outside of your job? And if it came from your work, the government has an interest, which is like a no free use of the thing. But even then you're allowed to patent and we encourage you to patent because if you don't, the Chinese will patent it two years after you invent it and we won't be allowed to use it," he said.

Robert J. Marks:

Is that right? Do we have a patent treaty with China?

Paul Werbos:

Well times have changed. Patent law has changed.

Robert J. Marks:

Okay.

Paul Werbos:

Patent law has changed. I'd probably have 10 more patents if they hadn't changed the patent law.

Robert J. Marks:

Okay.

Paul Werbos:

But they have substantially changed the patent law in ways that are not so good for small inventors in all fairness.

Robert J. Marks:

Yes. I acted as an expert witness in a number of cases involving kind of computational intelligence, artificial intelligence, and things have changed. I mean, it's almost impossible for the little guy anymore.

Paul Werbos:

Right. So the only thing the little guy can do is talk to the people in China who are willing to pay for that kind of stuff. And a lot of that is going on.

Robert J. Marks:

I have a friend who has a bunch of patents and he says, all the patent is good for is it gives you the right to sue people. And that's basically what it is. And it costs big bucks to sue people.

Paul Werbos:

So anyway, basically I had a misunderstanding when I was in early graduate school. Some of my first papers on reinforcement learning, brain-like intelligence, I believed that the kind of design I did for my PhD thesis with a few tweaks, parameters, basically could replicate what a mammal brain does. And then I learned a lot about mammal brains. I learned step by step by step mammal brains, do this, do that. There's some fundamental mathematical principles. And you can build a universal intelligence doing what I did in my thesis, but it's not a fast universal intelligence. You can build a faster, more powerful one by using some very fundamental principles. And I even did a paper in World Congress of Computational Intelligence 2014 in WICCI. I said, "Here is a four step procedure math one, math two, math three, that'll bring you up to at least the mouse level of universal general learning ability."

Paul Werbos:

And one of them, one of the key levels I called full spatial intelligence. And it's like a convolutional network, except it exploits more general spacial symmetries. And the more general spatial symmetries give you a power that a normal network doesn't have, a convolutional network doesn't have. You have to have this additional power in the spatial intelligence algorithm. And I was so proud giving that talk at WICCI and everybody seemed to listen, except the Chinese government really listened and started a

new program. And there was a US intelligence agency guy there who said, "No, no, we can't do this. We have to stop this work in the US."

Robert J. Marks:

Oh boy.

Paul Werbos:

So they stopped that line of research in the US and they expanded it in China.

Robert J. Marks:

Why? For what reason would they want to stop the US research? Were they going to take it under the umbrella of NSA or something? Probably you can only speculate.

Paul Werbos:

Well, I spoke to the guy, but he didn't tell me his total game plan.

Robert J. Marks:

Okay.

Paul Werbos:

And I have circumstantial evidence about what happened. I have discussed it with other people at NSF. There were a lot of people in NSF who were concerned about changes that occurred in many areas, and we debated whose fault are they and what's the reason. People once told me when I worked for the Department of Energy 10 years before NSF people used to tell a joke. The Department of Energy has already been privatized. It's a joint corporation owned 60% by the oil industry and 40% by the nukes. And it was like that my last few years there. When I went to NSF, I said, "Now I'm in a place where the powers that be are the American Physical Society and the deans." And the American Physical Society and the deans respect scientific values a lot more than those other guys did. And that was a big part of why I moved to NSF because the culture was totally the deans and the APS.

Paul Werbos:

And they're not perfect, but they have high values. But then came this Washington corruption business. My theory is that Washington corruption is what really hit more than anything else. But that's my speculation. I had friends who had other theories.

Robert J. Marks:

I wanted to ask you since your retirement I'm sure you've kept your fingers on the pulse of what's happening in our neural networks, artificial intelligence. What do you see some of the major advances in machine intelligence since your retirement?

Paul Werbos:

I have to confess one of the joys in retirement is I have a few other retired friends. And one is a woman named Frederica Darema who used to be head of the Air Force Office of Scientific Research. And every time people ask me, what's a quick summary of what is going on today with all of these algorithms and all of these different forms of AI, the best source I recommend to them, Frederica Darema organized a

conference a few months ago, global Zoom type conference, on something she calls DDDAS. And if you want to know the answer to your question, the starting point would be do a search on DDDAS video 2020 or 2021. And from her conference, I learned just how diverse the world is.

Paul Werbos:

There are so many different things they call AI. And there's so many experts who think they know things that are mutually contradictory. But I would say out of a hundred groups, there's a major percentage where they think the best way to solve the problems of AI is to get rid of it and go back to what we did 40 years ago in mathematics, before they even had AI, let alone neural networks. There are people getting funded to do that. And they persuade people this is the greatest clever thing you haven't heard of in the last 20 years. That's because it's a hundred years old.

Robert J. Marks:

Wait a minute, Paul. So they want to abandon AI? I think that that is especially stupid in terms of only national security.

Paul Werbos:

If you go to the DDDAS website, you'll see that there are groups who are trying to sell that approach. Out of maybe a hundred groups, the biggest number are sort of your routine, cut and dry, baby MLP, 20 years old at least. There were two groups that are just way the hell ahead of them. And in this particular conference, one of them was PWC, and one was RTX. RTX is the new merger of Raytheon and UTC where my old post office friend went. And PWC is a Pricewaterhouse, what's left over after they killed Arthur Henderson and they got a new competitor. And both of those organizations have people who are not doing the more advanced stuff we talked about, but they at least caught up with the most advanced stuff that we were doing. And so people will tell you, "Oh, you'll never build a Terminator robot. Don't worry about things like slaughter bots and Terminator weapons and autonomous weapons. That's all science fiction." And these groups would say that's science fiction, maybe in a hundred years we'll learn how to do it.

Paul Werbos:

And then another group is doing it. And they're building it. And nobody is stopping them from doing it. There are people stopping them from advertising it for obvious reasons. But the bottom line is that there are very advanced projects going on in the world that are decades ahead of almost all of their competition. There is mathematics, which is decades beyond them. So it's just incredible what the diversity is.

Robert J. Marks:

Wow. Very interesting. Thank you, Paul. We've been talking to Paul Werbos, inventor of the most commonly used technique to train artificial neural networks, an algorithm called error backpropagation. Thanks for listening. Until next time, be of good cheer.

Announcer:

This has been Mind Matters News with your host Robert J. Marks. Explore more at mindmatters.ai. That's mindmatters.ai. Mind Matters News is directed and edited by Austin Egbert. The opinions expressed on this program are solely those of the speakers. Mind Matters News is produced and copyrighted by the Walter Bradley Center for Natural and Artificial Intelligence at Discovery Institute.

