

5 SUPERVISED LEARNING

PJD's opening remarks

Our conversation about rule-based AI frequently referenced neural networks and many of you were asking for more on that topic. Today's lecture will give you plenty to chew on and hopefully will stimulate your interest in learning more.

There are many tasks that humans can do well but no one had been able to get a rule-based machine to do well. Recognizing what is in an image is a common example. Can a machine distinguish a photo of a dog from a photo of a cat? Tell you which photos in your album include your mother? Spot a potential shooter carrying a weapon?

Neural networks were conceived in the 1940s as a possible way to organize a computer. A biological neuron is modeled with an electronic neuron. The original proposers speculated that a computer organized like a brain might be able to behave like a brain. Unfortunately, computers organized this way were too slow. This idea did not make the cut for the first electronic computers.

Researchers continued to pursue this idea and found some circuits of electronic neurons that would recognize rudimentary images. Those networks were called perceptrons. Perceptrons initially received bad press, but researchers pressed on and found that multi-layer networks – perceptrons stacked one after the next, if you will -- could recognize realistic images. They found ways to train these deep networks efficiently.

They needed vast computing power to do the training at the scale of typical images, which consist of megapixels or more. That computing power became available about a decade ago and the whole field exploded into public view with some astonishing feats.

Today, Professor Marko Orescanin is here to tell you about this. He will show you how neural networks operate. He will also show you a dark side – some neural networks are susceptible to attacks by people who want to cause them to malfunction. Marko just joined our faculty, with almost a decade of commercial experience where he most recently helped establish and has led AI and data capabilities in consumer electronics at Bose Corporation. He has developed deep learning models and Machine Learning algorithms that are now deployed in millions of devices. He holds several patents. He has published and developed products in Radars, Satellites, Cameras, Acoustics, Speech and other sensor technologies. Marko received his PhD at the University of Illinois at Urbana-Champaign.