

# What is Next After Deep Learning



The first, and most important thing, to realize about deep learning is that it is not a “deep” subject, meaning that it is a very “shallow” topic with almost no theory underlying it. There are no guarantees of convergence (since we are after all talking about nonlinear optimization in high-dimensional spaces), and no performance guarantees of any kind (say, compared to what you get with other areas of machine learning, like kernel methods, sparse linear models etc.). It’s essentially like woodworking without physics. If you mix this type of polish with that kind of wood, you get this sort of

effect. The reason that there invariably has to be a future beyond deep learning is that one cannot build a solid engineering science of machine learning with bricks built out of hay. As Vladimir Vapnik once said, “The most practical thing in the world is good theory”, and that’s currently not available in deep learning. If deep learning is the best solution the machine learning community can do, as a card carrying member of this research community for over 30+ years, I’d have to say we are in serious trouble!

Let’s just take one example, the current rage over

generative adversarial models or GANs. There are close to 500+ papers on this topic, and almost 3 dozen variants of GANs with more appearing every week. However, there are barely any papers that show 1) whether GANs will converge reliably when trained (the original GANs do not!) 2) what the sample complexity of GANs are (no one knows) 3) what GANs can and cannot do. There's as far as I know 1-2 papers that attempt to give a theory of GANs, a particularly nice paper by Sanjeev Arora and colleagues, which is largely a negative result. It shows that the original GAN model does not converge, but that a modified multiple generator/multiple discriminator model might converge, in a very weak sense. Yet, this has not dampened any of the excitement about this model, far from it. There's also a collective sense of loss of reality when folks get excited about models like GANs. These models taken thousands and thousands of iterations to converge (when they do, and often, they don't), and each iteration requires many many passes through the data. At the end of the day, you burn through millions of CPU cycles, and you have to wonder whether after burning all that energy: is the game worth the candle? Where's all this energy getting us? is it leading us to a solid scientifically based theory of how to build a theory

of unsupervised learning? The vast majority of GAN papers are largely empirical, showing cute pictures of what a variant of GAN can do, but the metrics are often either non-existent or somewhat artificial.

So, many of us in the field indeed do look forward to a life beyond deep learning, where we can not only build impressive empirically substantiated learning systems, but also have a solid theory underlying it.

If you want an example of a truly "deep" science, look no further than this year's Nobel prize for the design of LIGO detectors, capping a 100 years of effort to detect gravitational waves from Einstein's general relativity theory. We can now detect collisions among black holes 2 billion light years away releasing more energy in one collision than all the energy from all the stars in the observable universe. And there's a very substantial amount of nontrivial mathematics that went into the building of the LIGO detectors and in advances in general relativity theory.

That's what a true "deep" learning theory should look like. I am confident that one day, machine learning will get there, but it will take many years of effort, and physicists provide us with an inspiration of what can be achieved