Deep Learning
Prof. Mitesh M. Khapra
Department of Computer Science and Engineering
Indian Institute of Technology, Madras

Lecture - 01
Chapter 9: (Need for) Sanity

So, lot of fields have adopted deep learning now and lot of State of the Art systems are based on deep neural networks, but now what is needed is after all this madness were deep learning has taken over a lot of research areas. Can we now bring in some sanity to the proceeding, right. So, this is really a need for sanity.

(Refer Slide Time: 00:32)

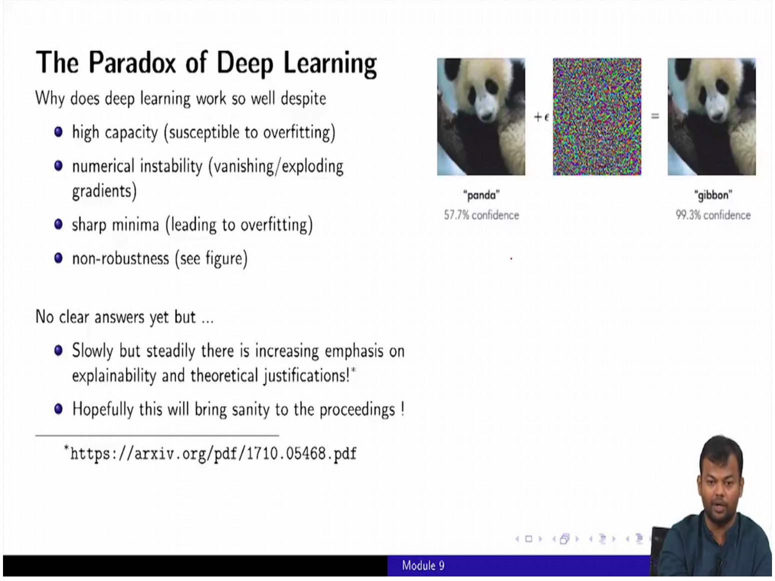

The Paradox of Deep Learning

Why does deep learning work so well despite

- high capacity (susceptible to overfitting)
- numerical instability (vanishing/exploding gradients)
- sharp minima (leading to overfitting)
- non-robustness (see figure)

No clear answers yet but ...

- Slowly but steadily there is increasing emphasis on explainability and theoretical justifications!*
- Hopefully this will bring sanity to the proceedings !

*https://arxiv.org/pdf/1710.05468.pdf

Why I say that is that because there is this paradox of deep learning. So, there is this interesting question that why does deep learning works so well despite having a high capacity.

So, the deep neural networks have a very high capacity which means that susceptible to over fitting. So, most of you would have done some course on machine learning. So, there you know that over fitting is bad because you are just memorizing the training data and then, you might not be able to do so well and at tested and over fitting happens when your model has a high capacity. So, even though deep neural networks have high capacity, why are they doing so well? We will focus on this high capacity, but when we talk about the universal approximation theorem and give some arguments for why deep neural networks have such a high capacity.

The other thing is they have this numerical instability, right. So, we spoke about these vanishing and exploding gradients and again, we will talk about this later on in the course. So, despite this training difficulties why is it that deep neural networks performs so well and of course, they have this sharp minima which is again it could lead to over fitting. So, if you look at there is an optimization problem, it is not a need convex optimization problem. So, it is a non convex optimization problem. So, why does it still do so well?

So, it is also not very robust. So, here is an example on the right hand side the figure that you see. So, the first figure is actually of a panda and the machine is able to detect this Panda with some 57 percent confidence, right. We have trained a machine for a lot of animal images. We have shown it a lot of animal images at test time. We show at this image. The first image that you see on the right hand side and is able to classify this is a Panda with 57 percent confidence, but now what I do is I add some very random noise. So, that second image that you see with some very random pixels if I add it to this image, I will get a new image, right.

So, every pixel in this image is added to this new noise image and I get the image which is see on the third. The third image that you see right to you and me or to any average human, this still looks like a Panda. There is hardly any difference between this image and the original image, but now if you pass this to the machine, all of a sudden instead of recognizing this is a Panda, it starts to recognize it as a Gibbon and that too with 99 percent confidence, right. So, why is it that they are not very robust and despite this not being very robust, why are deep neural networks so successful, right. So, people are interested in these questions and people have started asking these questions.

There are no clear answers yet, but slowly and steadily there is an increasing emphasis on explainability and theoretical justifications, right. So, it is not enough to say that your deep neural network works and gives you 99 percent accuracy. It is also good to have an explanation for why that happens is it that some components of the networks are really able to discriminate between certain patterns and so on. So, what is going on inside the network which is actually making it work so well, right and hopefully this will bring in some sanity to the proceedings.

So, instead of just saying that I apply deep learning to problem x and got 90 percent success, we will also make some kind of more scene arguments just to why this works and what is the further promise of this and thinks like that. So, that is roughly a quick historical recap of where deep learning started and where it is today; starting all the way back from advances in biology in 1871 to recent advances till 2017 and so on deep learning, right and here are few URL.

(Refer Slide Time: 03:55)

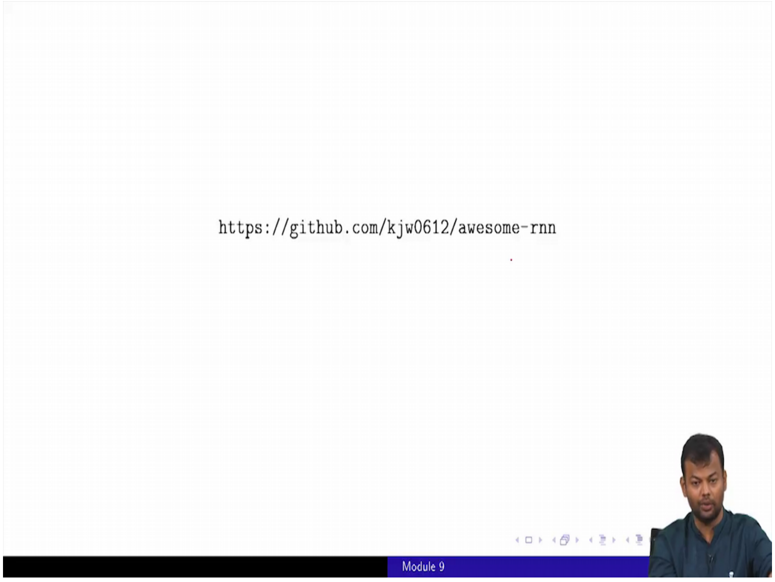

https://github.com/kjw0612/awesome-rnn

So, you could take a look at this for a lot of interesting applications of recurrent neural networks.

(Refer Slide Time: 04:00)

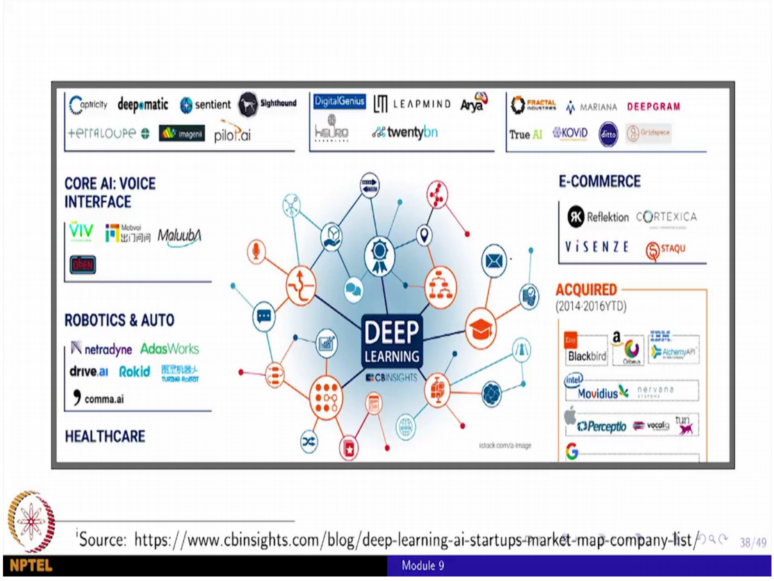

Source: https://www.cbinsights.com/blog/deep-learning-ai-startups-market-map-company-list/

Bunch of start-ups which have come up in this space is working on very varied an interesting problems and here are all the references that I have used for this particular presentation.

(Refer Slide Time: 04:06)

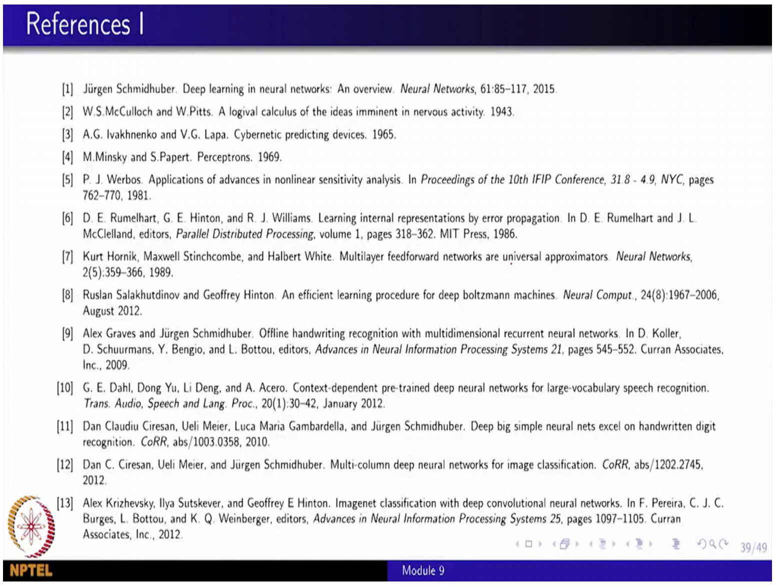
NPTEL
Module 9
39/49

References I

[1] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.

[2] W.S.McCulloch and W.Pitts. A logival calculus of the ideas imminent in nervous activity. 1943.

[3] A.G. Ivakhnenko and V.G. Lapa. Cybernetic predicting devices. 1965.

[4] M.Minsky and S.Papert. Perceptrons. 1969.

[5] P. J. Werbos. Applications of advances in nonlinear sensitivity analysis. In *Proceedings of the 10th IFIP Conference, 31.8 - 4.9, NYC*, pages 762–770, 1981.

[6] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In D. E. Rumelhart and J. L. McClelland, editors, *Parallel Distributed Processing*, volume 1, pages 318–362. MIT Press, 1986.

[7] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.

[8] Ruslan Salakhutdinov and Geoffrey Hinton. An efficient learning procedure for deep boltzmann machines. *Neural Comput.*, 24(8):1967–2006, August 2012.

[9] Alex Graves and Jürgen Schmidhuber. Offline handwriting recognition with multidimensional recurrent neural networks. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 545–552. Curran Associates, Inc., 2009.

[10] G. E. Dahl, Dong Yu, Li Deng, and A. Acero. Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition. *Trans. Audio, Speech and Lang. Proc.*, 20(1):30–42, January 2012.

[11] Dan Claudiu Ciresan, Ueli Meier, Luca Maria Gambardella, and Jürgen Schmidhuber. Deep big simple neural nets excel on handwritten digit recognition. *CoRR*, abs/1003.0358, 2010.

[12] Dan C. Ciresan, Ueli Meier, and Jürgen Schmidhuber. Multi-column deep neural networks for image classification. *CoRR*, abs/1202.2745, 2012.

[13] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.

So, that is where we end lecture 1 and I will see you again soon for lecture 2.

Thank you.