

A brief history of deep learning

Alice Gao

Lecture 10

Readings: [this article](#) by Andrew Kurenkov

Based on work by K. Leyton-Brown, K. Larson, and P. van Beek

Outline

Learning Goals

A brief history of deep learning

Lessons learned summarized by Geoff Hinton

Revisiting the Learning goals

Learning Goals

By the end of the lecture, you should be able to

- ▶ Describe the causes and the resolutions of the two AI winters in the past.
- ▶ Describe the four lessons learned summarized by Geoff Hinton.

A brief history of deep learning

A brief history of deep learning based on [this article](#) by Andrew Kurenkov.

The birth of machine learning

(Frank Rosenblatt 1957)

- ▶ A perceptron can be used to represent and learn logic operators (AND, OR, and NOT).
- ▶ It was widely believed that AI is solved if computers can perform formal logical reasoning.

First AI winter

(Marvin Minsky 1969)

- ▶ A rigorous analysis of the limitations of perceptrons.
- ▶ (1) We need to use multi-layer perceptrons to represent simple non-linear functions such as XOR.
- ▶ (2) No one knows a good way to train multi-layer perceptrons.
- ▶ Led to the first AI winter from the 70s to the 80s.

The back-propagation algorithm

(Rumelhart, Hinton and Williams 1986)

- ▶ In their Nature article, they precisely and concisely explained how we can use the back-propagation algorithm to train multi-layer perceptrons.

Second AI Winter

- ▶ Multi-layer neural networks trained with back-propagation don't work well, especially compared to simpler models (e.g. support vector machines and/or random forests).
- ▶ Led to the second AI winter in the mid 90s.

The rise of deep learning

- ▶ In 2006, Hinton , Osindero and Teh showed that back-propagation works if the initial weights are chosen in a smart way.
- ▶ In 2012, Hinton entered the ImageNet competition using deep convolutional neural networks and did far better than the next closest entry (an error rate 15.3% rather than 26.2% for the second best entry).
- ▶ The deep learning tsunami continues today.

Lessons learned summarized by Geoff Hinton

- ▶ Our labeled data sets were thousands of times too small.
- ▶ Our computers were millions of times too slow.
- ▶ We initialized the weights in a stupid way.
- ▶ We used the wrong type of non-linearity.

Revisiting the Learning Goals

By the end of the lecture, you should be able to

- ▶ Describe the causes and the resolutions of the two AI winters in the past.
- ▶ Describe the four lessons learned summarized by Geoff Hinton.