

Neural Nets: “Recent” History

- ▶ 1989: George Cybenko proves that neural nets with sigmoids are

Neural Nets: “Recent” History

- ▶ 1989: George Cybenko proves that neural nets with sigmoidals are universal function approximators.

Neural Nets: “Recent” History

- ▶ 1989: George Cybenko proves that neural nets with sigmoidals are universal function approximators.
- ▶ 1989: Yann LeCun- convolutional neural networks to read handwritten digits.

Neural Nets: “Recent” History

- ▶ 1989: George Cybenko proves that neural nets with sigmoidals are universal function approximators.
- ▶ 1989: Yann LeCun- convolutional neural networks to read handwritten digits. Turin prize in 2018 for this work.

Neural Nets: “Recent” History

- ▶ 1989: George Cybenko proves that neural nets with sigmoidals are universal function approximators.
- ▶ 1989: Yann LeCun- convolutional neural networks to read handwritten digits. Turin prize in 2018 for this work.
- ▶ 1993: Support Vector Machines (SVM), C. Cortes and V. Vapnik.

Neural Nets: “Recent” History

- ▶ 1989: George Cybenko proves that neural nets with sigmoidals are universal function approximators.
- ▶ 1989: Yann LeCun- convolutional neural networks to read handwritten digits. Turin prize in 2018 for this work.
- ▶ 1993: Support Vector Machines (SVM), C. Cortes and V. Vapnik.
- ▶ 1998: Yann LeCun brings in Stochastic Gradient Descent.

Neural Nets: “Recent” History

- ▶ 2006 or so: Approximately the beginning of “deep learning”, as coined by Geoffrey Hinton (who was given the Turing award in 2018).

Neural Nets: “Recent” History

- ▶ 2006 or so: Approximately the beginning of “deep learning”, as coined by Geoffrey Hinton (who was given the Turing award in 2018).
- ▶ 2009: The launch of ImageNet (Fei-Fei Li at Stanford). More than 14 million **labeled** images.

Neural Nets: “Recent” History

- ▶ 2006 or so: Approximately the beginning of “deep learning”, as coined by Geoffrey Hinton (who was given the Turing award in 2018).
- ▶ 2009: The launch of ImageNet (Fei-Fei Li at Stanford). More than 14 million **labeled** images.
- ▶ 2011: IBM’s Watson computer wins *Jeopardy!*.

Neural Nets: “Recent” History

- ▶ 2006 or so: Approximately the beginning of “deep learning”, as coined by Geoffrey Hinton (who was given the Turing award in 2018).
- ▶ 2009: The launch of ImageNet (Fei-Fei Li at Stanford). More than 14 million **labeled** images.
- ▶ 2011: IBM’s Watson computer wins *Jeopardy!*.
- ▶ 2011: The use of the Rectified Linear Unit (or ReLU) as activation (transfer) function.

Neural Nets: “Recent” History

- ▶ 2012: Creation of AlexNet - kicked off a spike in researching convolutional neural nets (that continues to this day).

Neural Nets: “Recent” History

- ▶ 2012: Creation of AlexNet - kicked off a spike in researching convolutional neural nets (that continues to this day).
(5 conv. layers, some use max-pooling layers, last three layers are fully connected nets with ReLU)

Neural Nets: “Recent” History

- ▶ 2012: Creation of AlexNet - kicked off a spike in researching convolutional neural nets (that continues to this day).
(5 conv. layers, some use max-pooling layers, last three layers are fully connected nets with ReLU)
- ▶ 2014: Introduction of GANs.
Ian Goodfellow introduced them; realistic data generation, especially images.

Neural Nets: “Recent” History

- ▶ 2012: Creation of AlexNet - kicked off a spike in researching convolutional neural nets (that continues to this day).
(5 conv. layers, some use max-pooling layers, last three layers are fully connected nets with ReLU)
- ▶ 2014: Introduction of GANs.
Ian Goodfellow introduced them; realistic data generation, especially images.
- ▶ 2014: Adam: A Method for Stochastic Optimization (D.P. Kingma and J. Ba, the name is derived from *adaptive moment estimation*).

Neural Nets: “Recent” History

- ▶ 2012: Creation of AlexNet - kicked off a spike in researching convolutional neural nets (that continues to this day).
(5 conv. layers, some use max-pooling layers, last three layers are fully connected nets with ReLU)
- ▶ 2014: Introduction of GANs.
Ian Goodfellow introduced them; realistic data generation, especially images.
- ▶ 2014: Adam: A Method for Stochastic Optimization (D.P. Kingma and J. Ba, the name is derived from *adaptive moment estimation*).
- ▶ 2016 – AlphaGo Defeats Lee Sedol in Go using deep reinforcement learning.

Neural Nets: Recent History

- ▶ 2017 – Transformers and “Attention Is All You Need”.

Neural Nets: Recent History

- ▶ 2017 – Transformers and “Attention Is All You Need”. Vaswani et al. proposed “transformers” replacing recurrent nets for NLP. Foundational for future language models.

Neural Nets: Recent History

- ▶ 2017 – Transformers and “Attention Is All You Need”. Vaswani et al. proposed “transformers” replacing recurrent nets for NLP. Foundational for future language models.
- ▶ 2018 – Google’s BERT (Bidirectional Encoder Representations from Transformers)

Neural Nets: Recent History

- ▶ 2017 – Transformers and “Attention Is All You Need”. Vaswani et al. proposed “transformers” replacing recurrent nets for NLP. Foundational for future language models.
- ▶ 2018 – Google’s BERT (Bidirectional Encoder Representations from Transformers)
BERT achieved state-of-the-art results and became a foundational model

Neural Nets: Recent History

- ▶ 2017 – Transformers and “Attention Is All You Need”. Vaswani et al. proposed “transformers” replacing recurrent nets for NLP. Foundational for future language models.
- ▶ 2018 – Google’s BERT (Bidirectional Encoder Representations from Transformers)
BERT achieved state-of-the-art results and became a foundational model
- ▶ 2020 – Release of GPT-3.

Neural Nets: Recent History

- ▶ 2017 – Transformers and “Attention Is All You Need”. Vaswani et al. proposed “transformers” replacing recurrent nets for NLP. Foundational for future language models.
- ▶ 2018 – Google’s BERT (Bidirectional Encoder Representations from Transformers) BERT achieved state-of-the-art results and became a foundational model
- ▶ 2020 – Release of GPT-3. OpenAI released GPT-3 with 175 billion parameters. Demonstrated unprecedented language generation ability.

Neural Nets: Recent History

- ▶ 2017 – Transformers and “Attention Is All You Need”. Vaswani et al. proposed “transformers” replacing recurrent nets for NLP. Foundational for future language models.
- ▶ 2018 – Google’s BERT (Bidirectional Encoder Representations from Transformers)
BERT achieved state-of-the-art results and became a foundational model
- ▶ 2020 – Release of GPT-3.
OpenAI released GPT-3 with 175 billion parameters. Demonstrated unprecedented language generation ability.
- ▶ 2021 – AlphaFold2 Solves Protein Folding.

Neural Nets: Recent History

- ▶ 2017 – Transformers and “Attention Is All You Need”.
Vaswani et al. proposed “transformers” replacing recurrent nets for NLP. Foundational for future language models.
- ▶ 2018 – Google’s BERT (Bidirectional Encoder Representations from Transformers)
BERT achieved state-of-the-art results and became a foundational model
- ▶ 2020 – Release of GPT-3.
OpenAI released GPT-3 with 175 billion parameters.
Demonstrated unprecedented language generation ability.
- ▶ 2021 – AlphaFold2 Solves Protein Folding.
DeepMind’s AlphaFold2 achieved near-experimental accuracy in protein structure prediction, revolutionizing computational biology.

Neural Nets: Recent History

- ▶ 2022 – Text-to-Image Models (e.g., Stable Diffusion).

Neural Nets: Recent History

- ▶ 2022 – Text-to-Image Models (e.g., Stable Diffusion).
Tools like DALL·E 2 and Stable Diffusion enabled high-quality image generation from text prompts, popularizing generative AI among artists and the general public.

Neural Nets: Recent History

- ▶ 2022 – Text-to-Image Models (e.g., Stable Diffusion).
Tools like DALL·E 2 and Stable Diffusion enabled high-quality image generation from text prompts, popularizing generative AI among artists and the general public.
- ▶ 2022 – Launch of ChatGPT.

Neural Nets: Recent History

- ▶ 2022 – Text-to-Image Models (e.g., Stable Diffusion).
Tools like DALL·E 2 and Stable Diffusion enabled high-quality image generation from text prompts, popularizing generative AI among artists and the general public.
- ▶ 2022 – Launch of ChatGPT.
- ▶ 2023 – GPT-4 and the Rise of Multimodal AI.

Neural Nets: Recent History

- ▶ 2022 – Text-to-Image Models (e.g., Stable Diffusion).
Tools like DALL·E 2 and Stable Diffusion enabled high-quality image generation from text prompts, popularizing generative AI among artists and the general public.
- ▶ 2022 – Launch of ChatGPT.
- ▶ 2023 – GPT-4 and the Rise of Multimodal AI.
Multimodal capabilities, allowing a single model to process text, images, and more.