
CS 333: NLP

Fall 2023

Prof. Carolyn Anderson
Wellesley College

New policy: earn extra late days

You can earn bonus late days by attending a CS research talk. To be eligible:

- The talk must be about CS research or AI research in a related field
- The talk must be in-person (so that you have the ability to ask questions)
- You must write a paragraph about the talk and what you learned and email it to me.

Research

- *Manipulation-robust citizens' assembly selection* studies how to reduce incentives for people to try to increase their chances of being chosen for the assembly by misreporting their features.
- *The distortion of public-spirited participatory budgeting* studies the welfare of participatory budgeting outcomes in a beyond-worst-case model of voter behavior: instead of considering only their own interests, voters also weigh the interests of others. This model is motivated by the potential for this behavior to be cultivated in practice, via democratic deliberation.



Bailey Flanagan

October 25th

Computer Science Colloquium Series | Fall 2023

Supporting Responsible AI Practices in Public Sector Contexts

Anna is a third year PhD student at Carnegie Mellon's Human-Computer Interaction Institute. Her research focuses on improving the design, evaluation, and governance of AI technologies used to inform complex, consequential decisions in real-world organizations. In addition to her research, she will share prior experiences forming collaborations with public sector agencies, doing research internships with industry groups, travelling to conferences, and mentoring undergraduate students. The session will end with an open Q/A discussion on applying to and doing a PhD in Computer Science / Human-Computer Interaction and other topics.

Accessibility and Disability Resources:
accessibility@wellesley.edu



Anna Kawakami '21

Nov 2nd, 12:45-2:00 | SCI H401

Lunch will be served

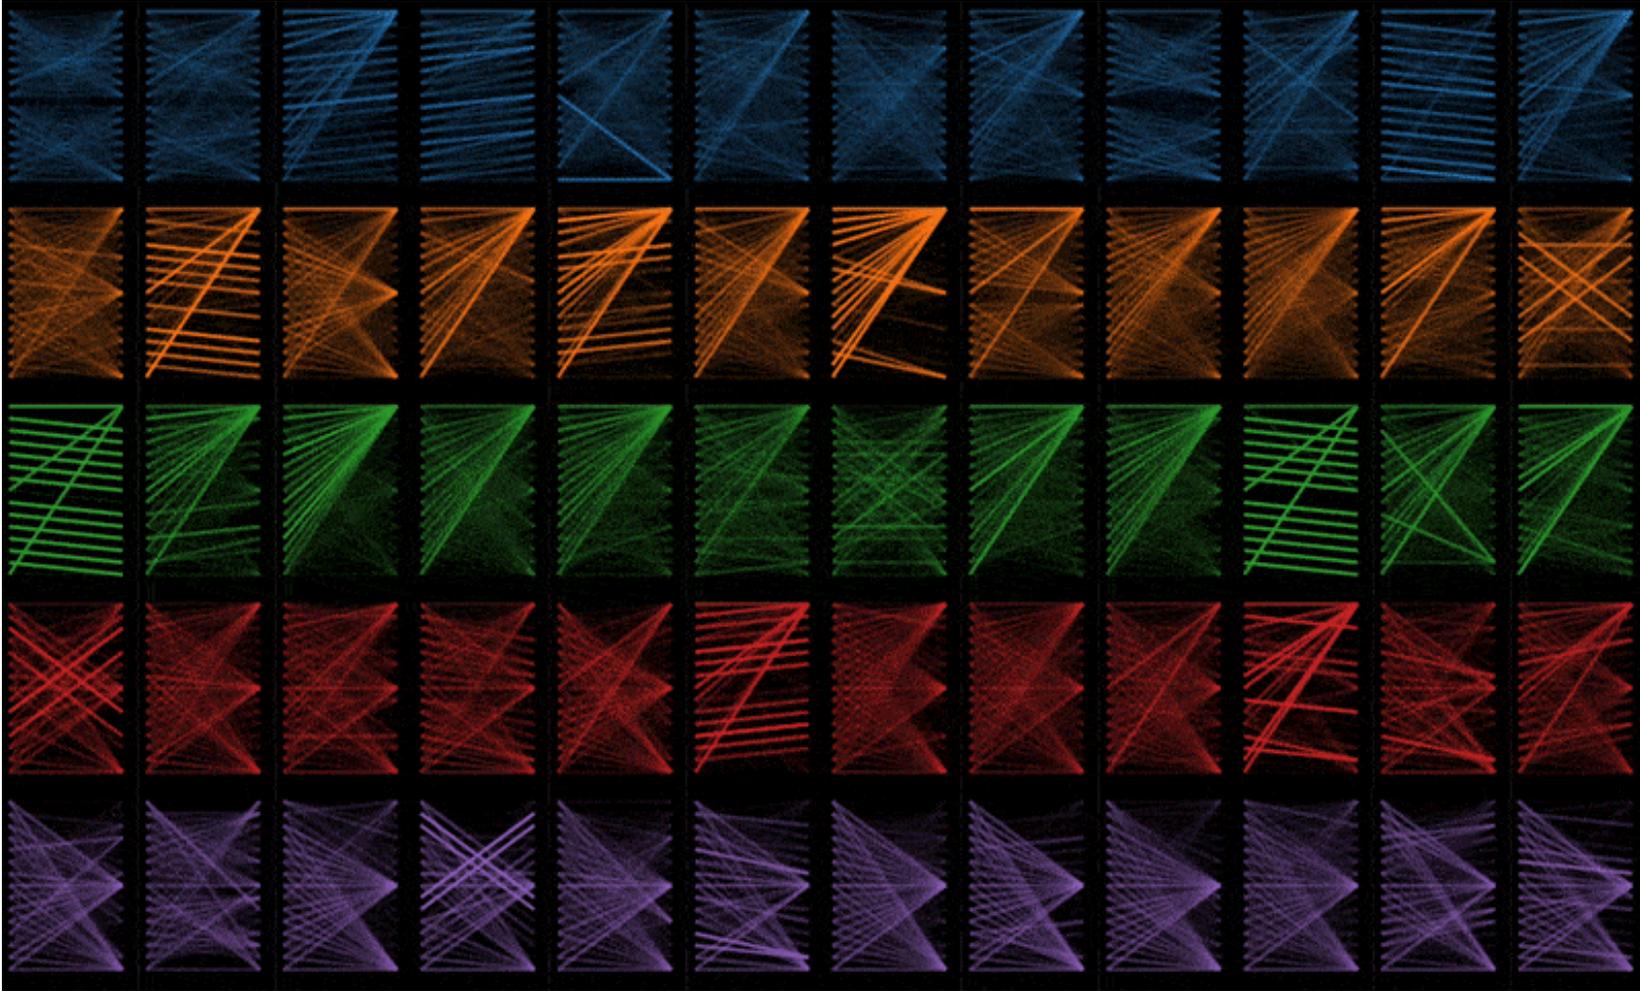
Questions??? eni.mustafara@wellesley.edu

November 2nd

Recap: regression model

Neural Networks

This is not your brain



<https://github.com/jessevig/bertviz>

It's a large language model (neural network)

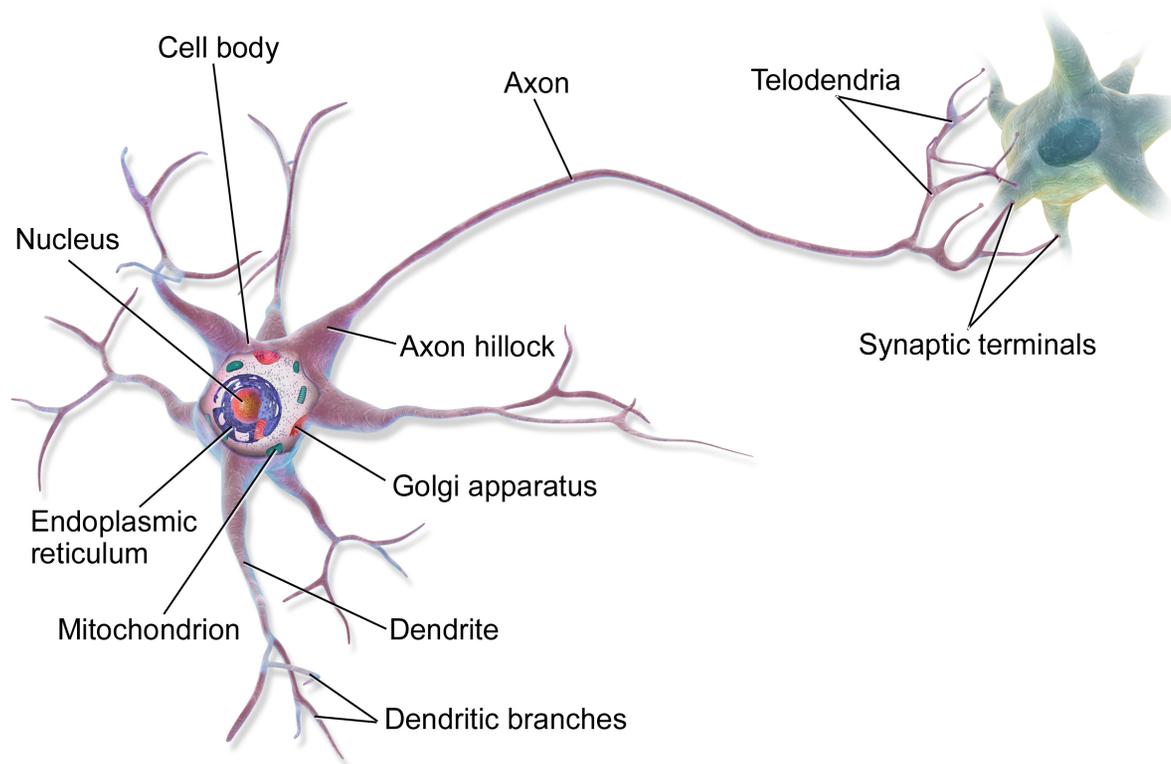
contemporary models have billions of parameters (GPT3: 175 billion)

This is someone's brain



estimated to have 100 trillion synapses connecting 100 billion neurons

This is in your brain

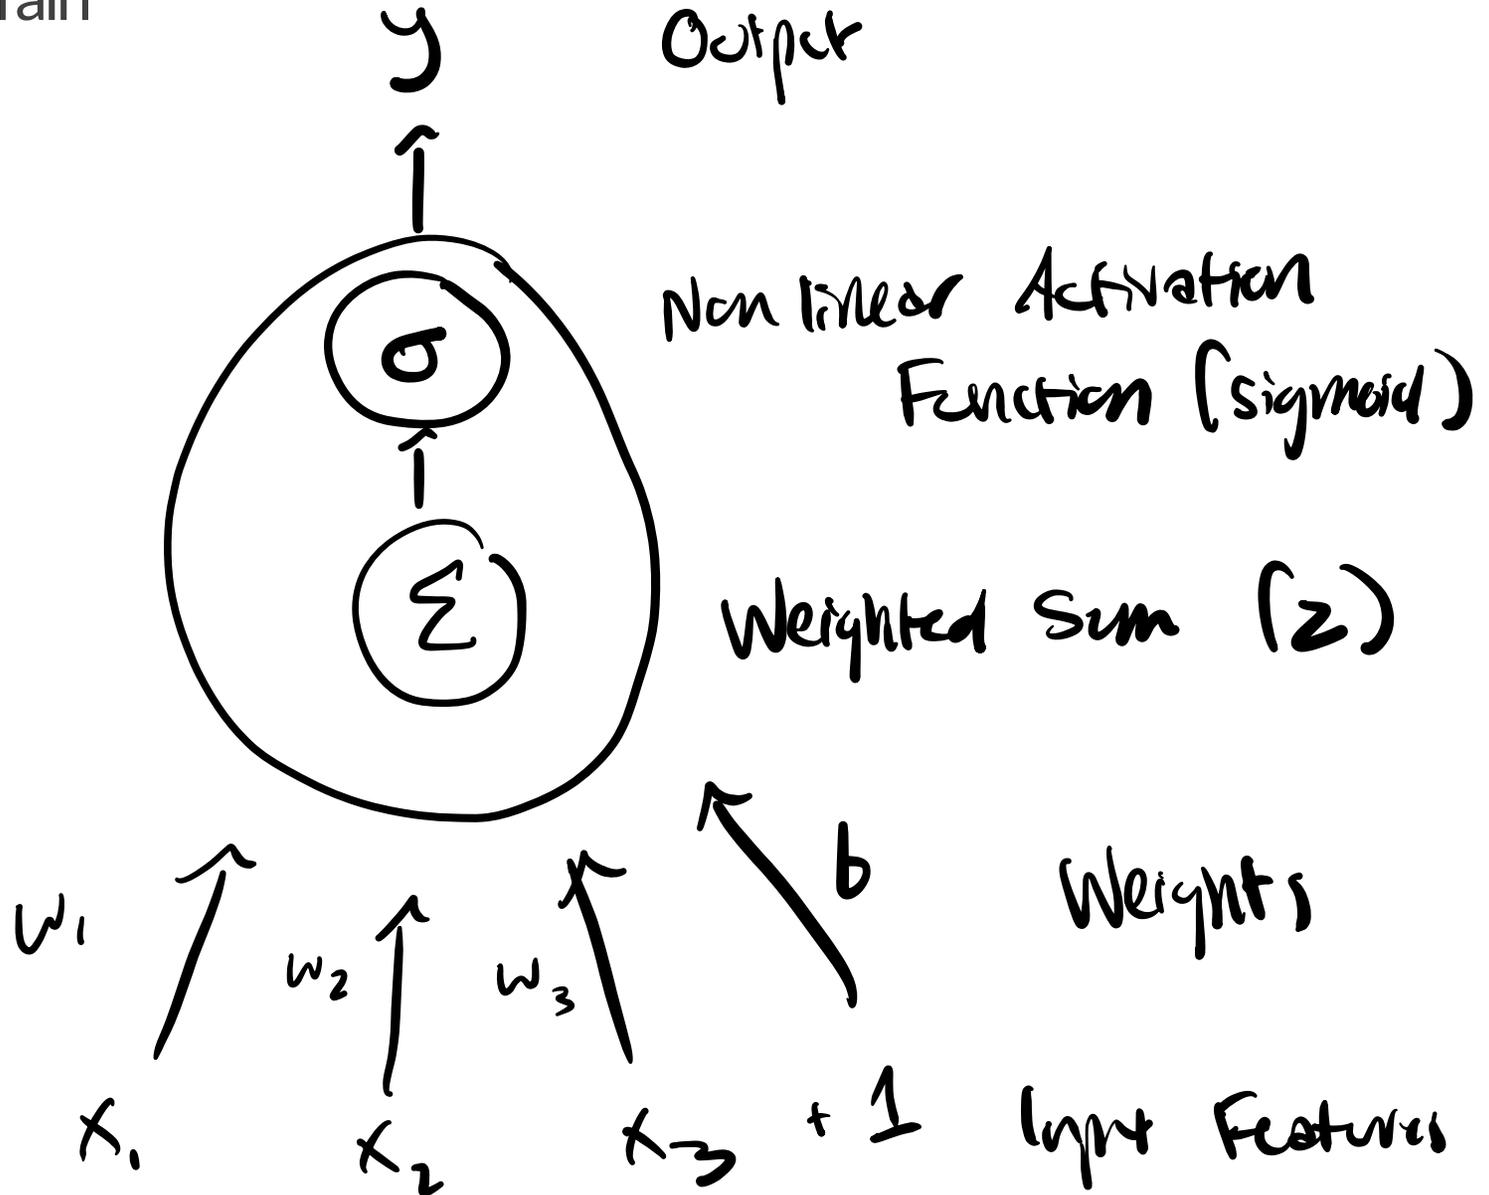


By BruceBlaus - Own work, CC BY 3.0, <https://commons.wikimedia.org/w/index.php?curid=28761830>

Neural Network Unit

This is not in your brain

$$y = \sigma(wx + b)$$



Units in Neural Networks

Neural unit

$$z = b + \sum_i W_i x_i$$

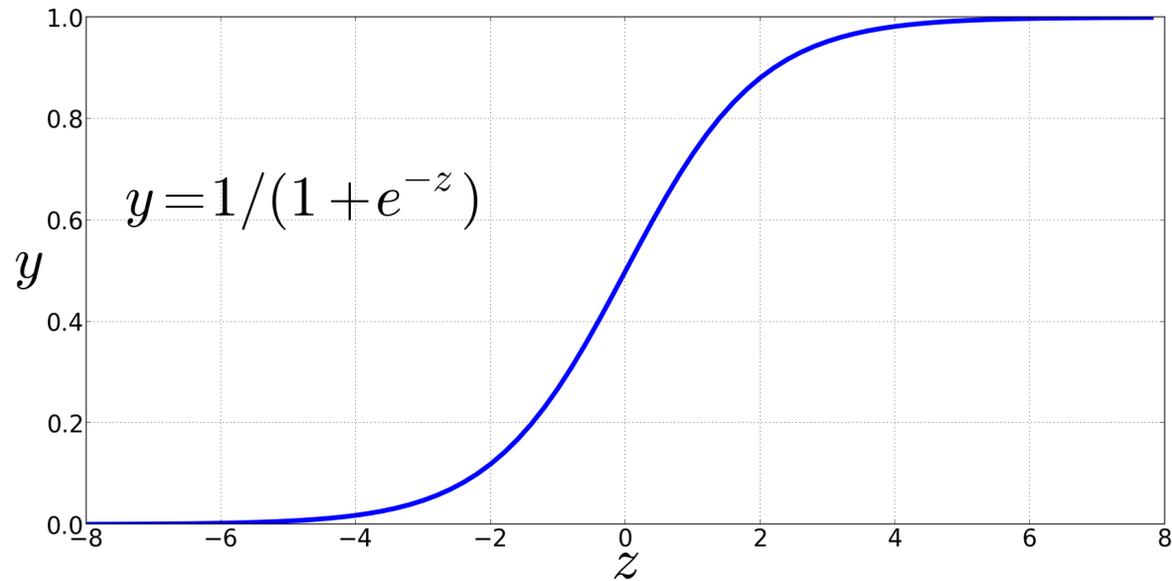
$$z = wx + b$$

$$y = \sigma(z)$$

Non-Linear Activation Functions

We've already seen the sigmoid for logistic regression:

Sigmoid

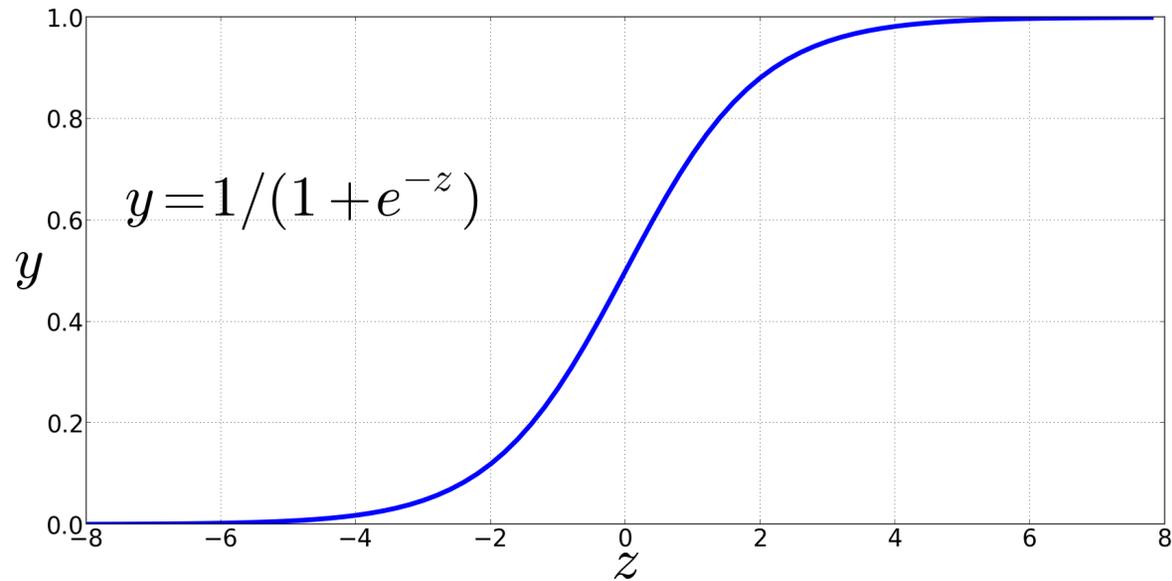


Non-Linear Activation Functions

We've already seen the sigmoid for logistic regression:

Sigmoid

$$y = s(z) = \frac{1}{1 + e^{-z}}$$



Final function of a unit:

$$y = \sigma(wx+b) = \frac{1}{1 + \exp(-(wx+b))}$$

Spot the differences

Neural Network Unit

$$z = b + \sum_i w_i x_i$$

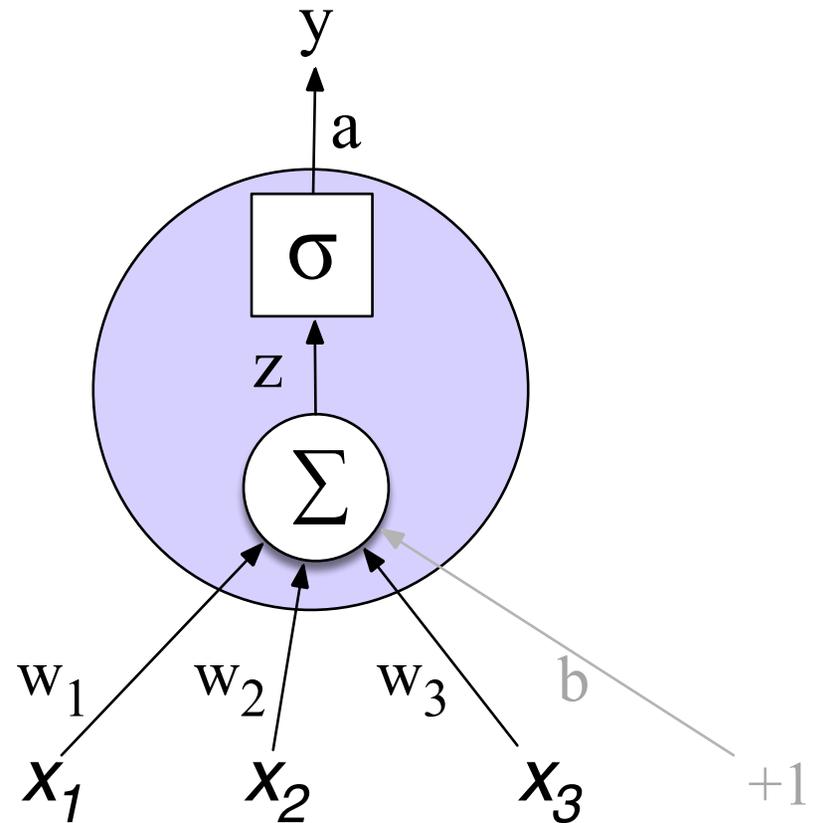
$$y = \sigma(wx + b)$$

Logistic Regression

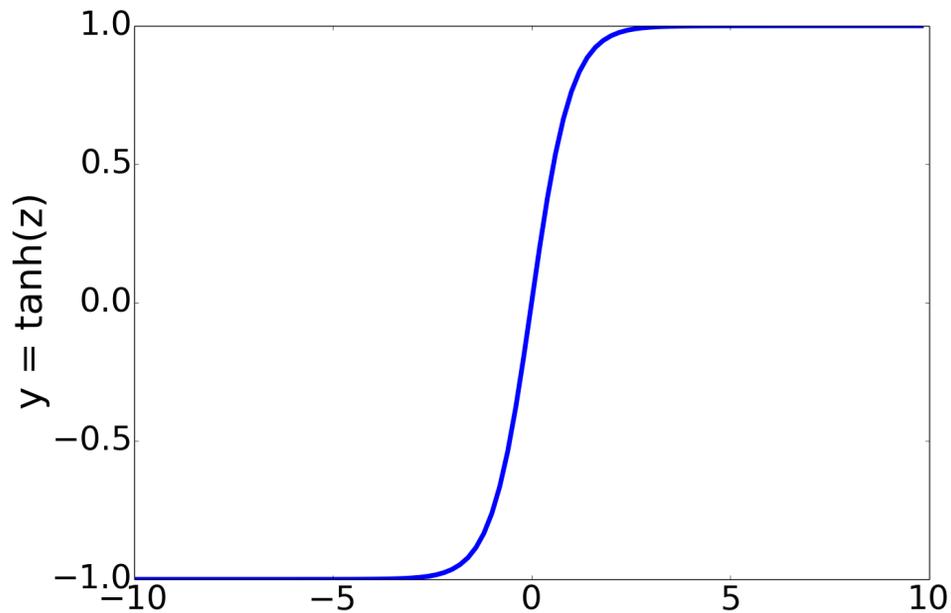
$$z = \left(\sum_i w_i x_i \right) + b$$

$$p(y=1) = \sigma(wx + b)$$

Final unit again



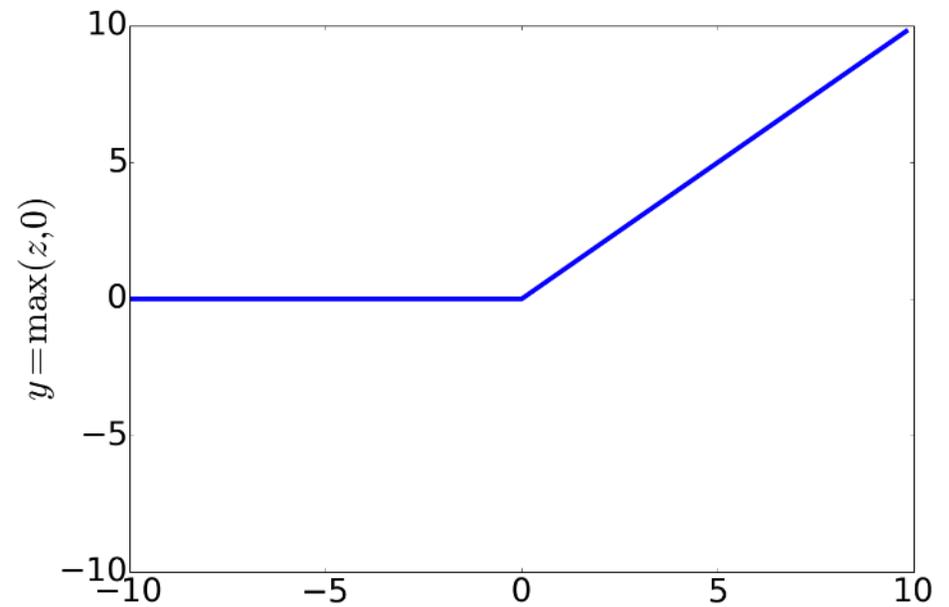
Non-Linear Activation Functions besides sigmoid



tanh

$$y = \frac{e^z - e^{-z}}{e^z + e^{-z}}$$

Most Common:



ReLU

Rectified Linear Unit

$$y = \max(z, 0)$$