



How can computers see?

A look at Convolutional Neural Networks


What number does this look like?





What number does this look like?





What number does this look like?




Are these samples of the same character?



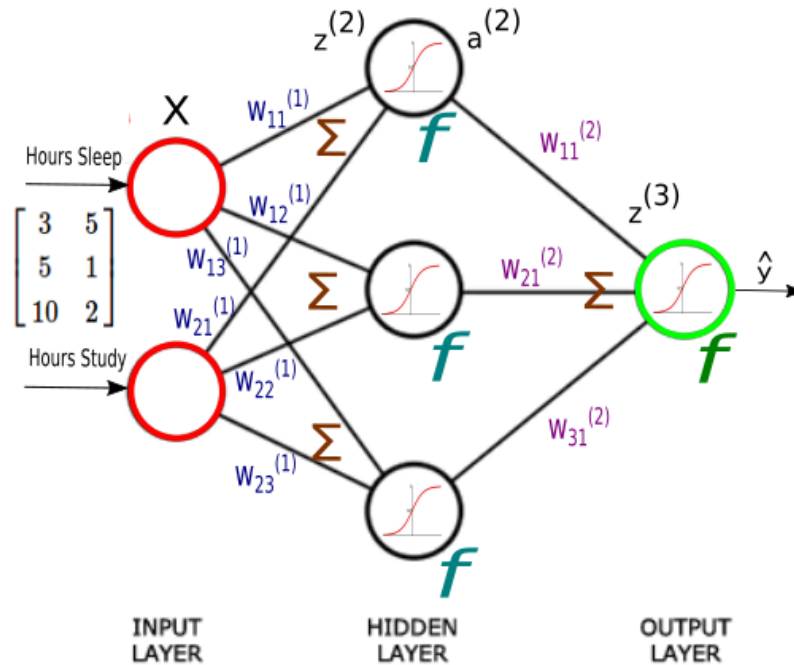
What is a neural network?



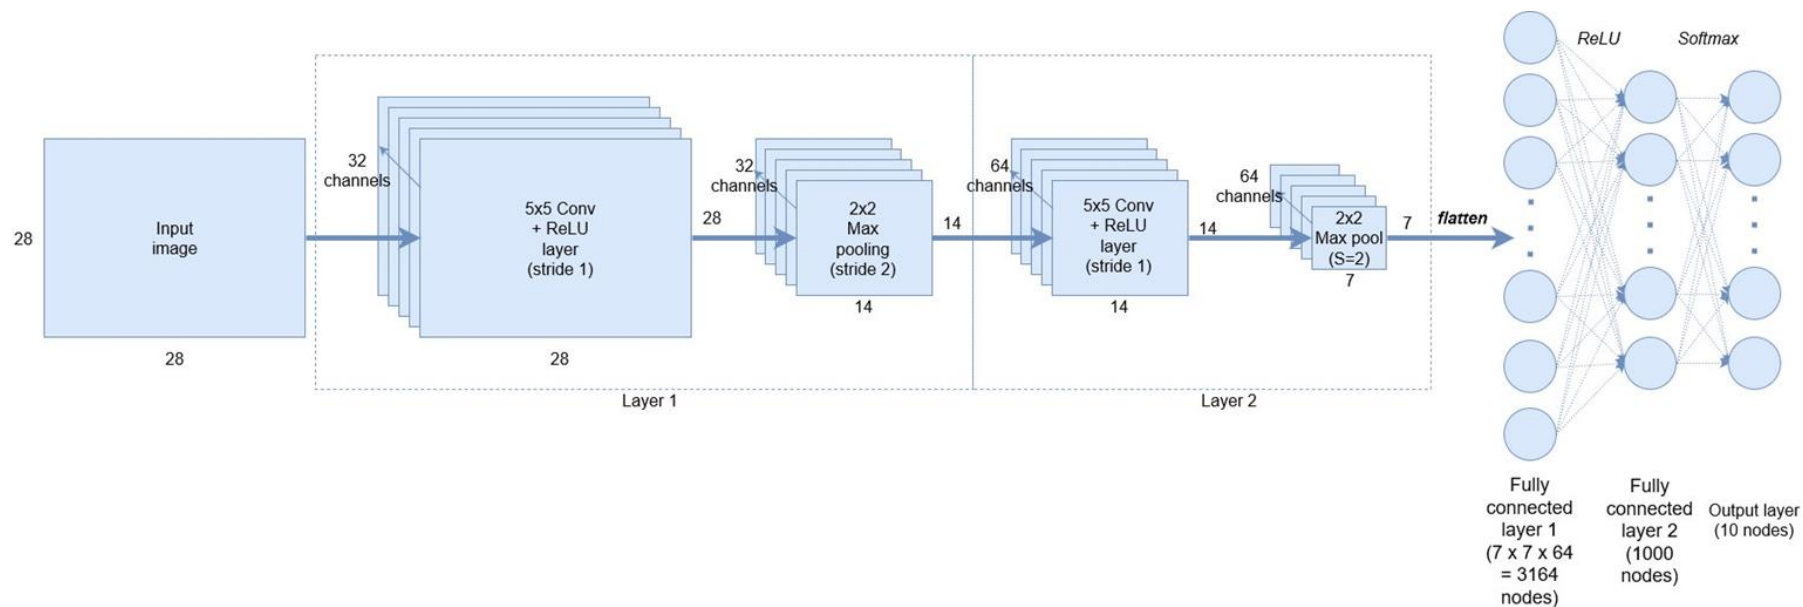
What is a neural network?



What is a neural network?




What is a Convolutional Neural Network (CNN)?



What is the benefit of a CNN?

When working with images, CNNs improve on NNs with:




- Feature detection
- Improved noise filtering
- Convolution layers
- Pooling layers



How are CNNs trained?

Once a stable structure is found, it is best to then change:


- Activation function
- Learning rate
- Window size (“kernel size”)
- Window border size (“padding”)
- Scanning speed (“stride”)
- Window dilation (“dilation”)



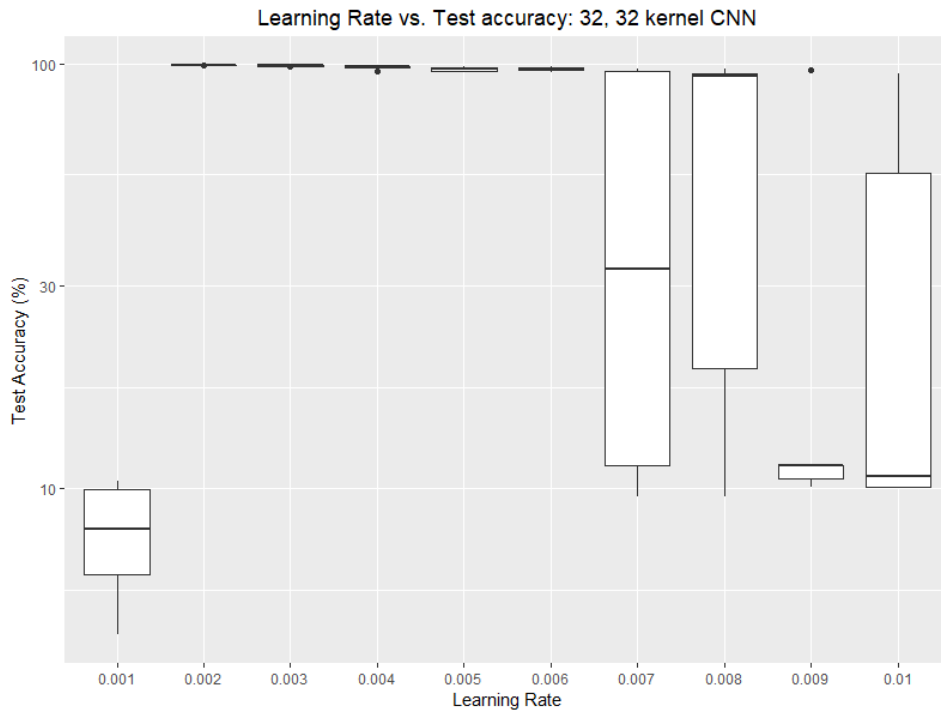
MNIST



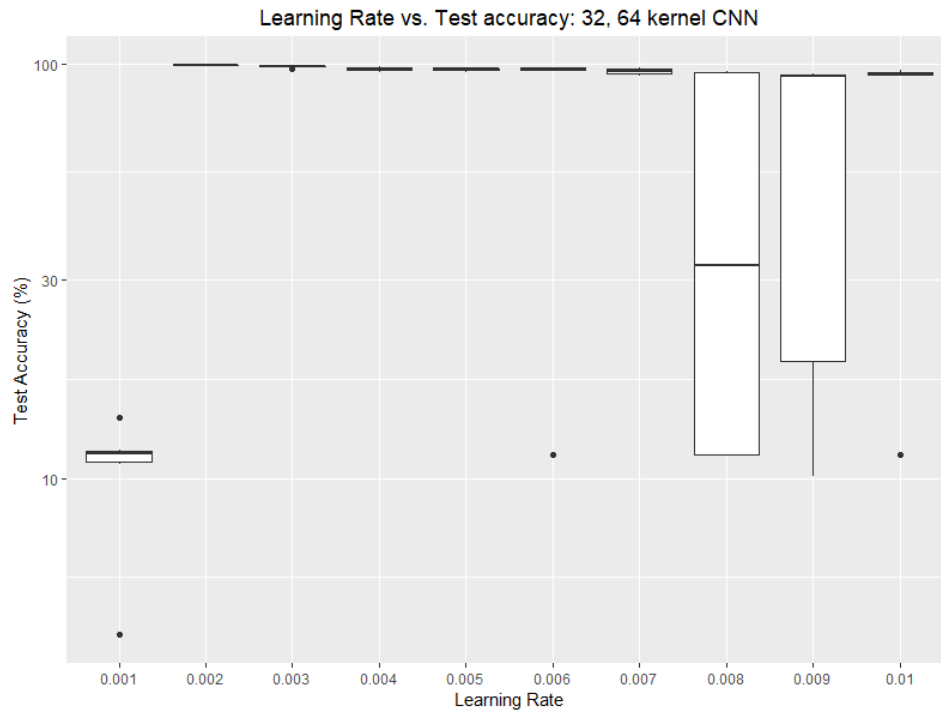
How well do CNNs perform?



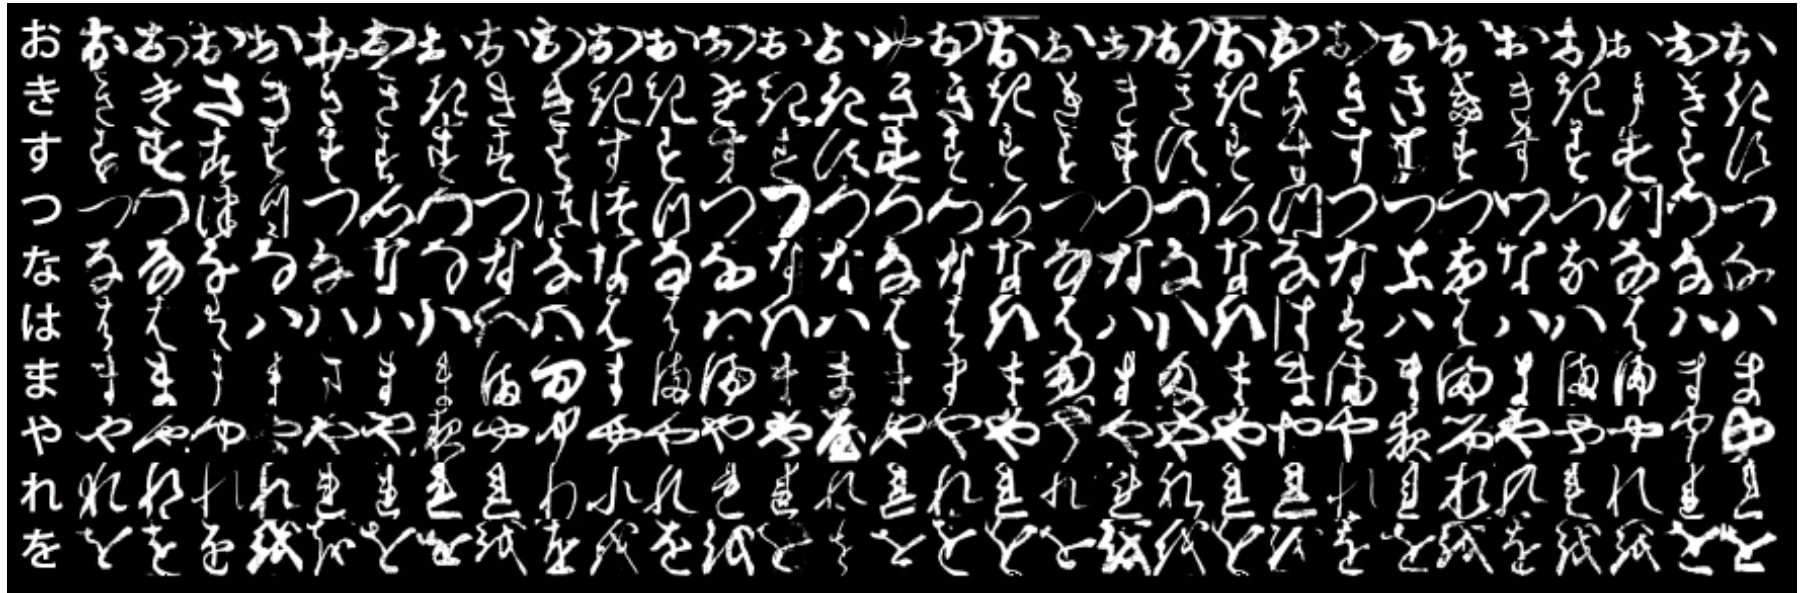
How well do CNNs perform?



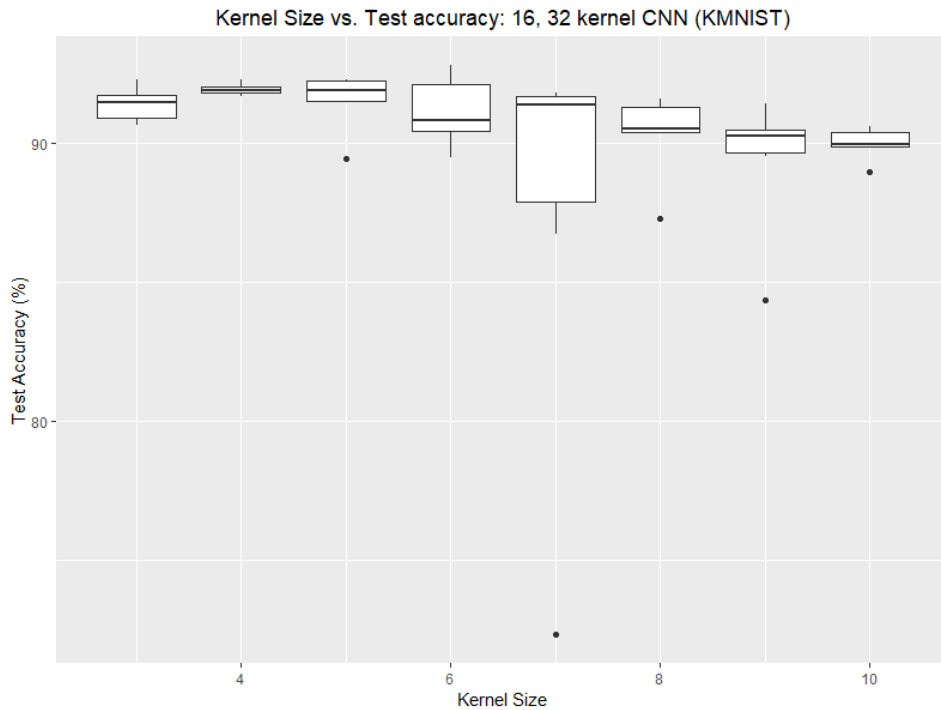
How well do CNNs perform?




KMNIST



How well do CNNs perform on characters?



How well do CNNs perform on characters?



Conclusions and future work

For MNIST, we know:

- Best learning rate
- Best architecture
- Kernel size
- Upward of 98% test accuracy

For KMNIST, we know:

- Best kernel size
- Upward of 92% test accuracy

What are some other things CNNs can do?

- General image classification
- Spatial data prediction
- Facial recognition
- On-board camera analysis




Image source: www.spiedigitallibrary.org