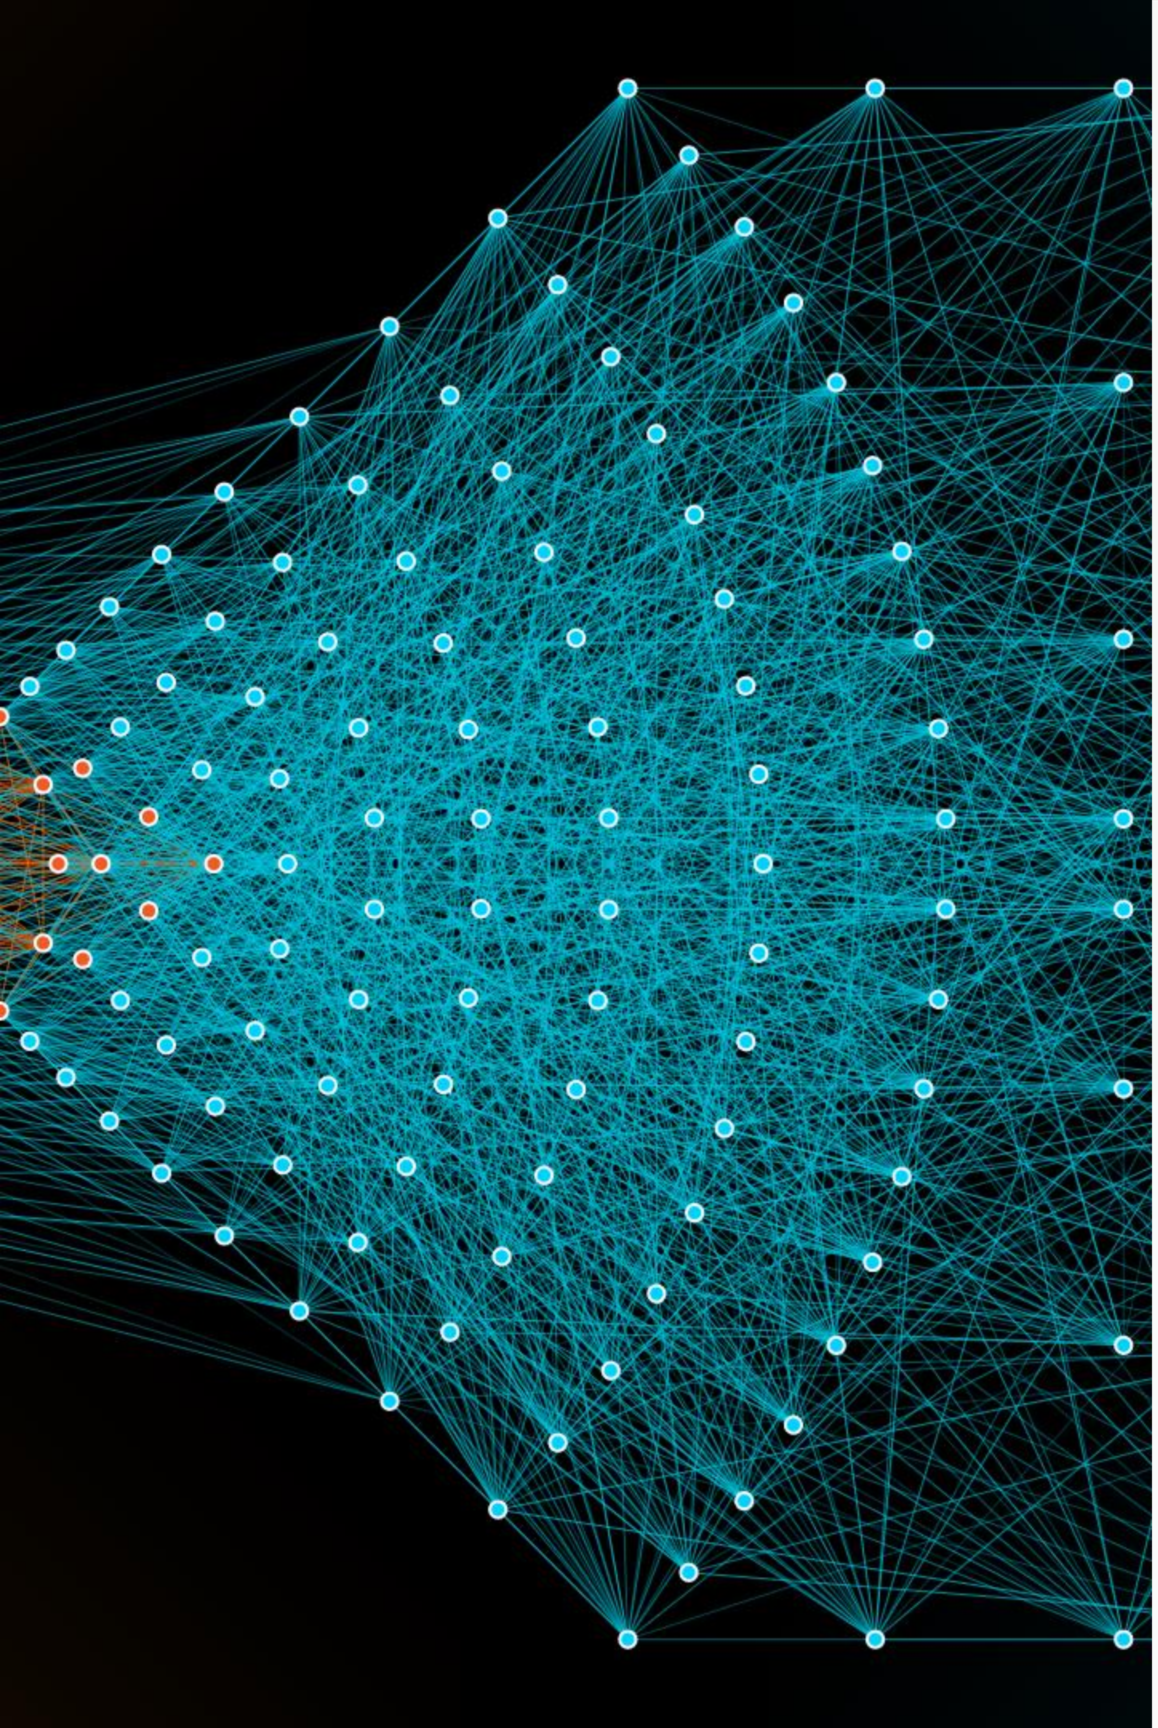

Convolution Neural Networks

In this presentation, we'll explore the fascinating world of Convolution Neural Networks and how they are transforming the field of Artificial Intelligence.

by Vineetha Venugopal

The Basic Structure of a CNN

Convolutional Layers

The first layer in a CNN that applies filters to the input image to create a feature map.

Pooling Layers

Downsamples the image by taking the max or average of a group of pixels.

Fully Connected Layers

The last layer in a CNN that processes the output from the convolution and pooling layers to make predictions.

Activation Function

Transforms the pooled data to give a nonlinear output that can be used for deep learning.

Convolutional Layer and Feature Maps

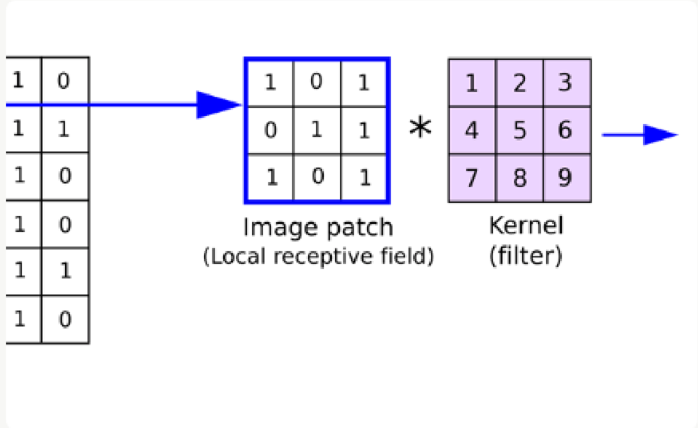

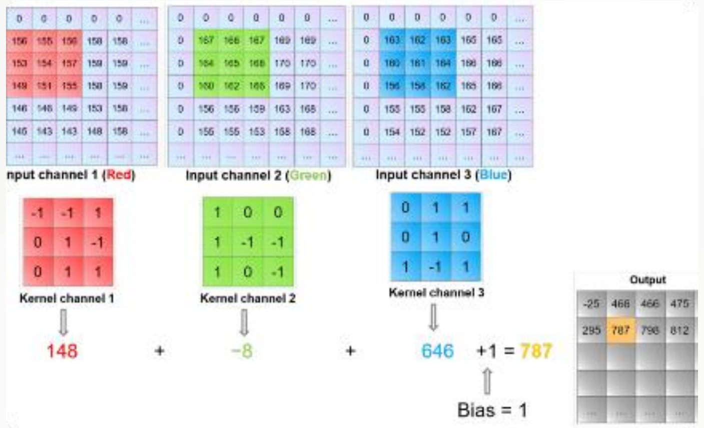

Convolutional Layer

The layer that performs a mathematical operation (convolution) on the input image and filters to produce a feature map.

Feature Maps

The maps created by convolving the input image with the filters to extract specific features.

Convolution Operation

The process of sliding a filter over the input image, computing element-wise multiplications and sums, and producing a convolved feature map.

Pooling Layer and Downsampling

1

Max Pooling

Retains the maximum value of a group of pixels in the feature map, resulting in a reduced output size and translation invariance.

2

Average Pooling

Computes the average value of a group of pixels in the feature map, resulting in a reduced output size and smoothing the image.

3

Global Pooling

Takes the average or maximum value of the entire feature map, resulting in a single scalar output.

Activation Functions and Backpropagation

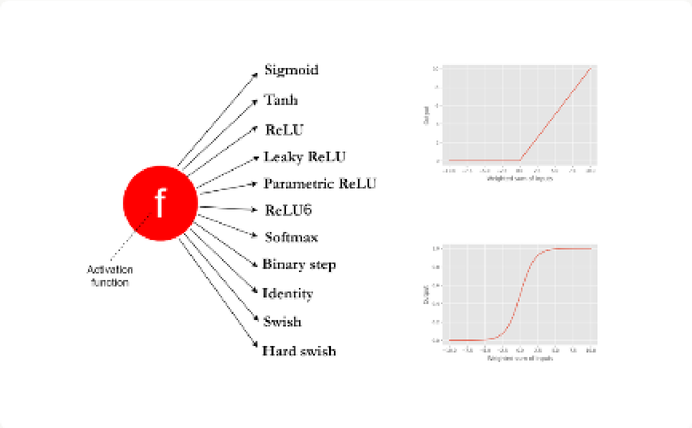

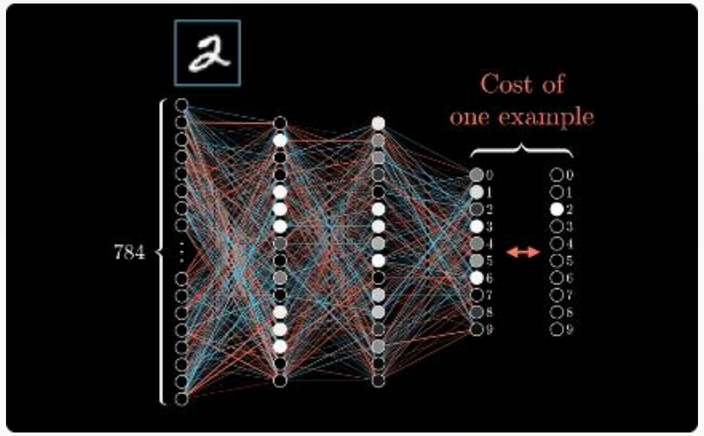

Sigmoid Function

Can map any input value to a value between 0 and 1, suitable for binary classification tasks.

ReLU Function

Returns 0 for all negative inputs and the input itself for all positive inputs, useful for avoiding the vanishing gradient problem.

Backpropagation

The process of updating the weights in the neural network to minimize the difference between predicted and actual output, using the gradient of the loss function.

Applications of CNNs in Image and Video Video Processing

1

2

3

Object Detection

Identifying multiple objects in an image and drawing bounding boxes around them using Region-based CNNs.

Image Classification

Identifying an object or scene in an image and assigning it to a specific category using CNNs trained on large datasets.

Video Processing

Extracting features from consecutive frames in a video and analyzing motion and action using CNNs with Long Short-Term Memory (LSTM) networks.

Limitations and Future Advancements

1 Overfitting

When the neural network performs well on the training data but poorly on the test data due to memorizing rather than learning the features.

2 Interpretability

Understanding how the neural network makes its predictions and determining which features are important.

3 Continual Learning

Training neural networks on multiple tasks and domains while retaining previously learned knowledge and adapting to new environments.

In Conclusion

Convolution Neural Networks have revolutionized the field of computer vision and are making significant contributions to many other domains. With further research and development, we can unlock their full potential and improve our lives in countless ways.