CS601 Machine Learning

Unit 3: Convolutional Neural Networks (CNN)

RGPV Exam-Cracking Handwritten-Style Notes

What this PDF covers

CNN, flattening, subsampling, padding, stride, convolution layer, pooling layer, loss layer, dense layer, 1x1 convolution, inception network, input channels, transfer learning, one-shot learning, dimension reduction, and CNN implementation using TensorFlow/Keras.

How to use these notes

1. First read the diagrams and memory tricks.
2. Then revise definitions and PYQ boxes.
3. Practice 7-mark and 14-mark answers from the important questions section.
4. One day before exam, revise the final 24-hour revision page.

Unit 3 Official Syllabus

Convolutional neural network, flattening, subsampling, padding, stride, convolution layer, pooling layer, loss layer, dense layer, 1x1 convolution, inception network, input channels, transfer learning, one-shot learning, dimension reduction, implementation of CNN using TensorFlow and Keras.

PYQ Trend Alert

CNN architecture, convolution layer, pooling layer, padding, stride, transfer learning and implementation using TensorFlow/Keras are highly important. CNN has appeared repeatedly in RGPV papers and is a very high probability 7-mark or 14-mark question.

1. Convolutional Neural Network (CNN)

Why this topic exists

CNN is a deep learning model mainly used for image and visual data. It automatically extracts features such as edges, corners, shapes and objects from images.

Beginner Explanation

Imagine you are identifying a face. You first notice edges, then eyes and nose, then the complete face. CNN also learns in layers: low-level features first and complex features later.

7 Marks Definition

A Convolutional Neural Network is a deep neural network designed to process grid-like data such as images by using convolution, pooling and fully connected layers to extract features and perform classification.

Working / Step-by-Step

1. Input image is given to CNN.
2. Convolution layers extract features using filters.
3. Activation function adds non-linearity.
4. Pooling layer reduces size.
5. Flatten layer converts feature maps into vector.
6. Dense layer performs classification.

Visual Diagram

Input Image -> Convolution -> ReLU -> Pooling -> Flatten -> Dense -> Output

Applications

- Image classification
- Face recognition
- Medical image analysis
- Self-driving car vision
- Object detection

PYQ Intelligence

PYQ Frequency: ★★★★★
Most expected question: Explain CNN architecture with neat diagram.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Convolutional Neural Network (CNN): remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

2. Convolution Layer

Why this topic exists

Convolution layer is the main building block of CNN. It applies filters on input image to detect features.

Beginner Explanation

Think of filter as a small window that scans the image and detects a pattern like edge or curve.

7 Marks Definition

Convolution layer applies learnable filters over input data to generate feature maps that represent important visual patterns.

Working / Step-by-Step

7. A small filter/kernel is selected.
8. Filter slides over image.
9. Element-wise multiplication is performed.
10. Results are summed.
11. Feature map is generated.

Visual Diagram

Image + Filter/Kernel -> Sliding Operation -> Feature Map

Applications

- Edge detection

- Corner detection
- Texture detection
- Image feature extraction

PYQ Intelligence

PYQ Frequency: ★★★★★
Most expected question: Explain convolution layer and its working.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Convolution Layer: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

3. Filters / Kernels

Why this topic exists

A filter or kernel is a small matrix used in convolution to extract features from an image.

Beginner Explanation

Imagine moving a magnifying glass over a photo to find edges. Filter works like that magnifying glass.

7 Marks Definition

A kernel is a small matrix of weights that scans input data to produce feature maps in convolution operation.

Working / Step-by-Step

12. Filter size may be 3x3 or 5x5.
13. It moves over image.
14. It detects one specific feature.
15. Multiple filters detect multiple features.

Visual Diagram

3x3 Filter slides on Image -> Feature Detected

Applications

- Detecting vertical edges
- Detecting horizontal edges
- Finding curves
- Detecting textures

PYQ Intelligence

PYQ Frequency: ★★★★
Most expected question: What are filters/kernels in CNN?
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Filters / Kernels: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

4. Padding

Why this topic exists

Padding means adding extra pixels around image border before convolution.

Beginner Explanation

If image border is not protected, corner information may be lost. Padding adds a protective boundary.

7 Marks Definition

Padding is the process of adding extra rows and columns around input data to control output size and preserve border information.

Working / Step-by-Step

16. Zeros are added around image.
17. Filter can cover border pixels.
18. Output size can be controlled.
19. Useful information near boundary is preserved.

Visual Diagram

Image -> Add Border/Zeros -> Apply Filter -> Better Feature Map

Applications

- Preserving image size
- Protecting edge information
- Deep CNN architectures

PYQ Intelligence

PYQ Frequency: ★★★★
Most expected question: Define padding. How does padding work in CNN?
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Padding: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

5. Stride

Why this topic exists

Stride is the number of pixels by which a filter moves over an image.

Beginner Explanation

If you read a book line by line, stride is 1. If you skip lines, stride is larger.

7 Marks Definition

Stride is the step size with which a convolution filter moves across the input image.

Working / Step-by-Step

20. Filter starts from top-left.
21. Moves according to stride value.
22. Larger stride reduces output size.
23. Smaller stride captures more detail.

Visual Diagram

Stride 1 = small steps; Stride 2 = skip one position

Applications

- Reducing feature map size
- Controlling computation
- Fast processing

PYQ Intelligence

PYQ Frequency: ★★★★
Most expected question: Explain stride in CNN.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Stride: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

6. Pooling Layer and Subsampling

Why this topic exists

Pooling reduces the spatial size of feature maps while keeping important information.

Beginner Explanation

Suppose you summarize a long chapter into key points. Pooling summarizes feature maps.

7 Marks Definition

Pooling layer performs down-sampling of feature maps to reduce dimensionality, computation and overfitting.

Working / Step-by-Step

24. Feature map is divided into small regions.
25. Max or average value is selected.
26. Output becomes smaller.
27. Important features are retained.

Visual Diagram

Feature Map -> Pooling Window -> Smaller Feature Map

Applications

- Max pooling
- Average pooling
- Reducing computation
- Making model robust

PYQ Intelligence

PYQ Frequency: ★★★★★
Most expected question: How do pooling layers reduce spatial dimensions?
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Pooling Layer and Subsampling: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

7. Flattening

Why this topic exists

Flattening converts 2D feature maps into a 1D vector before sending them to dense layers.

Beginner Explanation

Think of a table being converted into a single list.

7 Marks Definition

Flattening is the process of converting multidimensional feature maps into a one-dimensional vector for classification.

Working / Step-by-Step

28. CNN extracts feature maps.
29. Pooling reduces their size.

30. Flatten layer converts them into a vector.
31. Vector goes to dense layer.

Visual Diagram

2D Feature Maps -> Flatten -> 1D Vector -> Dense Layer

Applications

- Connecting CNN to classifier
- Image classification
- Neural network output stage

PYQ Intelligence

PYQ Frequency: ★★★★
Most expected question: Explain flattening in CNN.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Flattening: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

8. Dense Layer / Fully Connected Layer

Why this topic exists

Dense layer receives flattened features and performs final classification.

Beginner Explanation

It is like final decision maker after all features have been collected.

7 Marks Definition

Dense layer is a fully connected neural network layer where every neuron is connected to every neuron of the previous layer.

Working / Step-by-Step

32. Receives feature vector.

33. Applies weights and bias.
34. Uses activation function.
35. Produces final class scores.

Visual Diagram

Flatten Vector -> Dense Layer -> Class Prediction

Applications

- Cat/dog classification
- Digit recognition
- Medical image classification

PYQ Intelligence

PYQ Frequency: ★★★★
Most expected question: Explain dense layer in CNN.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Dense Layer / Fully Connected Layer: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

9. Loss Layer

Why this topic exists

Loss layer measures difference between predicted output and actual output.

Beginner Explanation

It is like marks deducted for wrong answer. More wrong prediction means more loss.

7 Marks Definition

Loss layer calculates prediction error and helps update network weights during training.

Working / Step-by-Step

36. Model predicts output.

37. Actual label is compared.
38. Loss is calculated.
39. Backpropagation uses this loss.

Visual Diagram

Prediction + Actual Label -> Loss -> Weight Update

Applications

- Cross entropy loss
- MSE loss
- Model training

PYQ Intelligence

PYQ Frequency: ★★★★
Most expected question: Explain loss layer in CNN.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Loss Layer: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

10. Input Channels

Why this topic exists

Input channels represent different layers of input image data. RGB image has three channels: Red, Green and Blue.

Beginner Explanation

A color image is not one sheet; it is like three transparent sheets stacked together.

7 Marks Definition

Input channels are separate data channels provided to a CNN, such as RGB channels in color images.

Working / Step-by-Step

40. Grayscale image has 1 channel.
41. RGB image has 3 channels.
42. Filters process all channels.
43. Feature maps are produced.

Visual Diagram

RGB Image -> Red + Green + Blue Channels -> CNN

Applications

- Color image processing
- Medical image channels
- Video frames

PYQ Intelligence

PYQ Frequency: ★★★
Most expected question: Explain input channels in CNN.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Input Channels: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

11. 1x1 Convolution

Why this topic exists

1x1 convolution uses a 1 by 1 filter to combine channels and reduce or increase depth.

Beginner Explanation

It looks small but is powerful because it works across channels.

7 Marks Definition

1x1 convolution is a convolution operation using a 1x1 kernel to change channel depth and reduce computation.

Working / Step-by-Step

44. Takes all input channels at one pixel location.
45. Combines channel information.
46. Reduces or expands depth.
47. Used before expensive convolutions.

Visual Diagram

Input Channels -> 1x1 Convolution -> Reduced Channels

Applications

- Dimension reduction
- Inception network
- Computational efficiency

PYQ Intelligence

PYQ Frequency: ★★★
Most expected question: Explain 1x1 convolution and its use.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

1x1 Convolution: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

12. Inception Network

Why this topic exists

Inception network uses multiple filter sizes in parallel and combines their outputs.

Beginner Explanation

Instead of guessing one lens size, it uses multiple lenses together.

7 Marks Definition

Inception Network is a CNN architecture that applies multiple convolution and pooling operations in parallel to capture features at different scales.

Working / Step-by-Step

48. Apply 1x1, 3x3 and 5x5 convolutions in parallel.
49. Apply pooling in parallel.
50. Concatenate outputs.
51. Capture multi-scale features.

Visual Diagram

Input -> [1x1, 3x3, 5x5, Pooling] -> Concatenate -> Output

Applications

- Image classification
- GoogleNet
- Feature extraction

PYQ Intelligence

PYQ Frequency: ★★★
Most expected question: Explain Inception Network.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Inception Network: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

13. Transfer Learning

Why this topic exists

Transfer learning uses a pre-trained model and adapts it for a new related task.

Beginner Explanation

It is like learning driving a car and then quickly learning to drive an SUV.

7 Marks Definition

Transfer Learning is a technique where knowledge learned from one task is reused for another related task to reduce training time and data requirement.

Working / Step-by-Step

52. Choose pre-trained model.
53. Remove final layer.
54. Add new classifier.
55. Fine-tune on new dataset.
56. Deploy for new task.

Visual Diagram

Pre-trained CNN -> Replace Final Layer -> Train on New Data -> New Model

Applications

- Medical image classification
- Small dataset tasks
- Face recognition
- Object classification

PYQ Intelligence

PYQ Frequency: ★★★★★
Most expected question: Explain benefits of transfer learning and Inception architecture.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Transfer Learning: remember definition, diagram, working steps and applications. For exam, write a neat

diagram first, then explain working in points.

14. One-Shot Learning

Why this topic exists

One-shot learning trains a model to recognize a class from only one or very few examples.

Beginner Explanation

Humans can identify a new person after seeing one photo. One-shot learning tries to do the same.

7 Marks Definition

One-Shot Learning is a learning approach where a model learns to recognize objects or classes using only one or very few training examples.

Working / Step-by-Step

57. One or few examples are provided.
58. Model learns similarity.
59. New input is compared with known example.
60. Closest match is predicted.

Visual Diagram

One Example -> Feature Encoding -> Similarity Check -> Prediction

Applications

- Face verification
- Signature verification
- Rare disease detection

PYQ Intelligence

PYQ Frequency: ★★★★
Most expected question: Explain one-shot learning.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

One-Shot Learning: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

15. Dimension Reduction

Why this topic exists

Dimension reduction reduces number of features while keeping important information.

Beginner Explanation

It is like reducing a long chapter into useful summary points.

7 Marks Definition

Dimension Reduction is a technique used to reduce the number of input features while preserving important information and improving model efficiency.

Working / Step-by-Step

61. Remove redundant features.
62. Compress feature space.
63. Reduce computation.
64. Improve visualization and training.

Visual Diagram

High Dimensional Data -> Dimension Reduction -> Compact Features

Applications

- PCA
- Autoencoders
- Feature selection
- Faster training

PYQ Intelligence

PYQ Frequency: ★★★★★
Most expected question: What do you mean by dimension reduction?
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

Dimension Reduction: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

16. CNN Implementation using TensorFlow and Keras

Why this topic exists

TensorFlow and Keras are popular libraries used to build CNN models quickly.

Beginner Explanation

Instead of writing everything from scratch, Keras gives ready-made layers like Conv2D, MaxPooling2D, Flatten and Dense.

7 Marks Definition

CNN implementation using TensorFlow and Keras involves defining convolution, pooling, flatten and dense layers, compiling the model and training it on image data.

Working / Step-by-Step

65. Import TensorFlow/Keras.
66. Create Sequential model.
67. Add Conv2D layer.
68. Add MaxPooling2D layer.
69. Add Flatten layer.
70. Add Dense output layer.
71. Compile and train model.

Visual Diagram

Input -> Conv2D -> MaxPooling2D -> Flatten -> Dense -> Output

Applications

- Image classification
- Digit recognition
- CNN practical implementation

PYQ Intelligence

PYQ Frequency: ★★★★★
Most expected question: Explain process of implementing CNN in TensorFlow.
Question pattern: Explain / diagram based / short note.

Examiner Keywords

Feature maps, filters, convolution, pooling, dimensionality, CNN architecture, classification, training, TensorFlow/Keras.

One-Minute Revision

CNN Implementation using TensorFlow and Keras: remember definition, diagram, working steps and applications. For exam, write a neat diagram first, then explain working in points.

CNN Architecture - Complete RGPV Answer Format

Neat Diagram to draw in exam

Input Image
↓
Convolution Layer + ReLU
↓
Pooling Layer
↓
Convolution Layer + ReLU
↓
Pooling Layer
↓
Flattening
↓
Dense / Fully Connected Layer
↓
Output Class

In a 14-mark answer, start with definition, draw this architecture, explain each layer and finish with applications. This format is enough to score good marks in RGPV exam.

TensorFlow / Keras CNN Code Skeleton

Code Skeleton

```
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Conv2D, MaxPooling2D, Flatten, Dense

model = Sequential()
```

```
model.add(Conv2D(32, (3,3), activation="relu", input_shape=(64,64,3)))
model.add(MaxPooling2D(pool_size=(2,2)))
model.add(Conv2D(64, (3,3), activation="relu"))
model.add(MaxPooling2D(pool_size=(2,2)))
model.add(Flatten())
model.add(Dense(128, activation="relu"))
model.add(Dense(10, activation="softmax"))
model.compile(optimizer="adam", loss="categorical_crossentropy", metrics=["accuracy"])
```

Exam me code exact yaad na ho to bhi layer names yaad rakho: Conv2D, MaxPooling2D, Flatten, Dense, Compile, Fit. Ye keywords examiner ko clearly show karte hain ki implementation concept samajh me hai.

Comparison Tables for Quick Revision

Concept	Meaning	Exam Point
Convolution	Filter scans image	Feature extraction
Padding	Adds border pixels	Preserves size and edges
Stride	Filter movement step	Controls output size
Pooling	Down-sampling	Reduces computation
Flattening	2D to 1D conversion	Connects CNN to dense layer
Dense Layer	Fully connected classifier	Final prediction

Technique	Main Purpose	High Probability Question
Transfer Learning	Reuse pre-trained model	Explain benefits of transfer learning
One-Shot Learning	Learn from few examples	Explain one-shot learning
1x1 Convolution	Reduce channels	Use in Inception network
Inception Network	Multi-scale feature extraction	Explain architecture
Dimension Reduction	Reduce features	Why high-dimensional data is challenging

Top 20 Important Questions - Unit 3

72. Explain Convolutional Neural Network with neat architecture diagram.
73. Explain convolution layer and its working in CNN.
74. Define padding. How does padding work in CNN?
75. Explain stride and its effect on output size.
76. How do pooling layers reduce spatial dimensions of feature maps?
77. Explain flattening in CNN.
78. Explain dense layer and loss layer in CNN.
79. Explain input channels in CNN with RGB image example.
80. Explain 1x1 convolution and its importance.
81. Explain Inception Network architecture.
82. Explain transfer learning and its benefits.

83. Explain one-shot learning with example.
84. Explain dimension reduction in Machine Learning.
85. Explain CNN implementation using TensorFlow/Keras.
86. Differentiate convolution and pooling layers.
87. Differentiate padding and stride.
88. Explain CNN applications in computer vision.
89. Explain how CNN extracts hierarchical features.
90. Write short note on Max Pooling and Average Pooling.
91. Explain complete CNN workflow from image input to output class.

Most Expected Questions 2026

- 95% - Explain CNN architecture with neat diagram.
- 95% - Explain implementation of CNN in TensorFlow/Keras.
- 90% - Define padding. How does padding work in CNN?
- 90% - How do pooling layers reduce spatial dimensions?
- 90% - Explain transfer learning and Inception network.
- 85% - Explain one-shot learning and dimension reduction.
- 85% - Explain 1x1 convolution and input channels.

24-Hour Exam Revision Sheet

Read this one day before exam

CNN = Image learning deep network
Convolution = feature extraction
Filter/Kernel = small matrix that scans image
Padding = protects border and controls size
Stride = filter movement step
Pooling/Subsampling = reduces feature map size
Flattening = 2D to 1D vector
Dense Layer = final classifier
Loss Layer = measures error
1x1 Convolution = channel reduction
Inception = parallel filters of different sizes
Transfer Learning = reuse pre-trained model
One-Shot Learning = learn from very few examples
Dimension Reduction = reduce features
TensorFlow/Keras = Conv2D, MaxPooling2D, Flatten, Dense

RGPV 14-Mark Answer Template for CNN

92. Introduction: CNN is a deep learning model used for image data.
93. Definition: Write 7-mark definition.
94. Diagram: Draw CNN architecture.
95. Layers: Explain convolution, ReLU, pooling, flattening, dense and loss layer.
96. Working: Explain step-by-step feature extraction to classification.
97. Advantages: automatic feature extraction, high accuracy, less preprocessing.
98. Applications: face recognition, medical imaging, object detection, self-driving cars.
99. Conclusion: CNN is powerful for computer vision and image classification.

Final Quality Check

If a weak student studies only these notes one day before exam, they can confidently write 7-mark and 14-mark answers on CNN, padding, pooling, transfer learning, one-shot learning, dimension reduction and CNN implementation using TensorFlow/Keras.