



Classification of Machine Learning Problems


Discrete Value and Probability Distribution: Representations for smile latent attribute:



Smile (discrete value)



Smile (probability distribution)



vs.

Vectors, Matrices, Tensors, Color Images

Vectors, matrices, tensors.

vector



$$\mathbf{v} \in \mathbb{R}^{64}$$

matrix



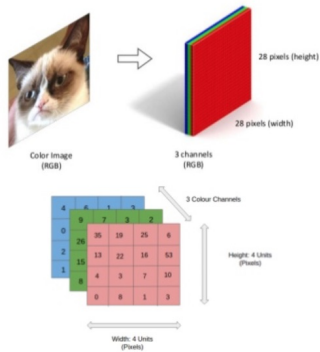
$$\mathbf{X} \in \mathbb{R}^{8 \times 8}$$

tensor



$$\mathcal{X} \in \mathbb{R}^{4 \times 4 \times 4}$$

Color images are three-dimensional tensors.




2D Convolutions

2D Convolutions


A 2D convolution is a mathematical operation that applies a filter to an image, producing a filtered output (also called a feature map).

Simple Example: Sharpen filter ...



$$5 \times 0 + 9 \times -1 + 13 \times 0 + 6 \times -1 + 10 \times 5 + 14 \times -1 + 7 \times 0 + 11 \times -1 + 15 \times 0 + 5 = 15$$

2D Convolutions



Convolution Neural Networks (CNNs)

CNN Layers:

- **Convolution Layers:** Filters or kernels that detect features such as edges or textures.
- **ReLU Activations:** Adds nonlinearity to the model, helping to learn complex complex patterns.
- **Pooling Layers:** Reduce the dimensionality of the image (making the network more efficient while preserving important features).
- **Fully Connection Layers:** After feature extraction, these layers make the final prediction based on the detected patterns.
- **Softmax Output:** Converts the network output into probabilities, showing the likelihood of each class.

Convolution Operations/Layers

Dropout Layer

A **dropout layer** randomly select a fraction of input units to zero during training, thereby working to **reduce overfitting** and ultimately **improving generalization**.




Figure 2-16. A dropout layer

CNN Summary


Advantages of CNNs:

- Automatic feature extraction.
- High accuracy in vision tasks.
- Translation invariance.


Disadvantages of CNNs:

- Resource intensive training (taking days or weeks).
- Large data requirement.
- Overfitting risk – model performs well on trained data, poorly on new data.
- Traditional CNN architectures require input images to be of a fixed size.

Probability Estimation (fig03)




Probability Estimation (fig06)



Probability Estimation (fig07)




Probability Estimation (fig09)



Probability Estimation (fig10)

Original Formulation




Deterministic Node



Random Node

Reparameterized Formulation



Feedforward



Differentiation

Example 2: Generation of Fashion Shapes

Original Images: (28 × 28) pixels, greyscale.



Reconstructions: (28 × 28) pixels, greyscale.



Fashion MNIST Dataset: A collection of greyscale images of clothing items, each (28 × 28) pixels, expanded to (32 × 32) pixels for processing. Individual pixels take a value 0 through 255.

Example 3: Generation of Cartoon Faces

Decoder:

Model: "Decoder"

Layer (type)	Output Shape	Param #
input_layer (InputLayer)	(None, 200)	0
dense_1 (Dense)	(None, 4096)	823,296
Reshape (Reshape)	(None, 8, 8, 64)	0
conv_transpose_1 (Conv2DTranspose)	(None, 16, 16, 64)	36,928
bn_1 (BatchNormalization)	(None, 16, 16, 64)	256
lrelu_1 (LeakyReLU)	(None, 16, 16, 64)	0
conv_transpose_2 (Conv2DTranspose)	(None, 32, 32, 64)	36,928
bn_2 (BatchNormalization)	(None, 32, 32, 64)	256
lrelu_2 (LeakyReLU)	(None, 32, 32, 64)	0
conv_transpose_3 (Conv2DTranspose)	(None, 64, 64, 64)	36,928

Example 3: Generation of Cartoon Faces

Encoder: 1.77 million parameters.

Decoder: 974,000 parameters.

Model: "Encoder"

Layer (type)	Output Shape	Param #	Connected to
input_layer (InputLayer)	(None, 256, 256, 3)	0	-
conv_1 (Conv2D)	(None, 128, 128, 32)	896	input_layer[0][0]
bn_1 (BatchNormalization)	(None, 128, 128, 32)	128	conv_1[0][0]
lrelu_1 (LeakyReLU)	(None, 128, 128, 32)	0	bn_1[0][0]
conv_2 (Conv2D)	(None, 64, 64, 64)	18,496	lrelu_1[0][0]
bn_2 (BatchNormalization)	(None, 64, 64, 64)	256	conv_2[0][0]
lrelu_2 (LeakyReLU)	(None, 64, 64, 64)	0	bn_2[0][0]
conv_3 (Conv2D)	(None, 32, 32, 64)	36,928	lrelu_2[0][0]
bn_3 (BatchNormalization)	(None, 32, 32, 64)	256	conv_3[0][0]
lrelu_3 (LeakyReLU)	(None, 32, 32, 64)	0	bn_3[0][0]
conv_4 (Conv2D)	(None, 16, 16, 64)	36,928	lrelu_3[0][0]
bn_4 (BatchNormalization)	(None, 16, 16, 64)	256	conv_4[0][0]
lrelu_4 (LeakyReLU)	(None, 16, 16, 64)	0	bn_4[0][0]
conv_5 (Conv2D)	(None, 8, 8, 64)	36,928	lrelu_4[0][0]
bn_5 (BatchNormalization)	(None, 8, 8, 64)	256	conv_5[0][0]
lrelu_5 (LeakyReLU)	(None, 8, 8, 64)	0	bn_5[0][0]
Flatten (Flatten)	(None, 4096)	0	lrelu_5[0][0]
mean (Dense)	(None, 200)	819,400	Flatten[0][0]
log_var (Dense)	(None, 200)	819,400	Flatten[0][0]

Total params: 1,770,128 (6.75 MB)
Trainable params: 1,769,582 (6.75 MB)
Non-trainable params: 576 (2.25 KB)

Model: "Decoder"

Layer (type)	Output Shape	Param #
input_layer (InputLayer)	(None, 200)	0
dense_1 (Dense)	(None, 4096)	823,296
Reshape (Reshape)	(None, 8, 8, 64)	0
conv_transpose_1 (Conv2DTranspose)	(None, 16, 16, 64)	36,928
bn_1 (BatchNormalization)	(None, 16, 16, 64)	256
lrelu_1 (LeakyReLU)	(None, 16, 16, 64)	0
conv_transpose_2 (Conv2DTranspose)	(None, 32, 32, 64)	36,928
bn_2 (BatchNormalization)	(None, 32, 32, 64)	256
lrelu_2 (LeakyReLU)	(None, 32, 32, 64)	0
conv_transpose_3 (Conv2DTranspose)	(None, 64, 64, 64)	36,928
bn_3 (BatchNormalization)	(None, 64, 64, 64)	256
lrelu_3 (LeakyReLU)	(None, 64, 64, 64)	0
conv_transpose_4 (Conv2DTranspose)	(None, 128, 128, 32)	18,464
bn_4 (BatchNormalization)	(None, 128, 128, 32)	128
lrelu_4 (LeakyReLU)	(None, 128, 128, 32)	0
conv_transpose_5 (Conv2DTranspose)	(None, 256, 256, 3)	807

Total params: 954,387 (3.64 MB)
Trainable params: 952,859 (3.64 MB)
Non-trainable params: 448 (1.76 KB)

Example 3: Generation of Cartoon Faces

Iteration: 30



