

# INDEX

## A

accuracy, 58  
activation function, 62  
AlexNet, 34  
algorithm, 5  
alignment, 132  
Analytical Engine, 6, 27  
architecture, 69  
artificial intelligence, 5, 29

- algorithm, 8
- artificial general intelligence, 116
- artificial narrow intelligence, 116
- deep learning, 5, 7, 34
- machine learning, 5, 7
- model, 7
- narrow vs. general, 116
- supervised learning, 8

Asilomar Conference on Beneficial AI, 36  
Atkinson, Rowan, 25  
attention, 131  
automatic differentiation, 95  
AutoML, 94  
Ayers, John W., 147

## B

Babbage, Charles, 6, 27  
backpropagation, 32, 39, 75  
backward pass, 69  
bagging, 49–50  
batch normalization, 40  
Belloc, Hilaire, 19  
Bengio, Yoshua, 33  
bias, 63  
binary model, 9  
Boiko, Daniil A., 148  
Boole, George, 27  
Boolean algebra, 27  
Bostrom, Nick, 150

Bottou, Léon, 33  
bounding box, 93  
Box, George, 8  
Breiman, Leo, 33  
Bubeck, Sébastien, 148

## C

Carroll, Lewis, 147  
central processing unit (CPU), 38  
ChatGPT, 36, 115, 147  
Cheng, Eugenia, 151  
Choi, Yejin, 141  
CIFAR-10, 47, 82, 90  
classifier, 12  
class label, 9, 12–13  
computational graph, 95  
conditional GAN, 103  
confusion matrix, 15, 57

- accuracy, 58
- Matthews correlation coefficient, 58

connectionism, 26  
context encoding, 131  
context window, 132  
convolution, 85

- kernel, 54, 85

convolutional layer, 87  
convolutional neural network (CNN), 30, 33, 86

- AlexNet, 34
- effective receptive field, 90
- embedding, 92
- end-to-end learning, 81
- filter, 88
- LeNet, 87
- pooling layer, 87

correlation, 20  
Cortes, Corinna, 32  
Cover, Thomas, 30  
CPU (central processing unit), 38

curse of dimensionality, 46  
curve fitting, 17

## D

DALL-E 2, 110  
Dartmouth workshop, 29  
data  
    features, 10  
    feature space, 11  
    feature vector, 10  
    hard negatives, 21  
    labeled, 8  
    sample, 10  
    test set, 10  
    training set, 10  
data augmentation, 70  
dataset  
    CIFAR-10, 47, 82, 90  
    dinosaur footprints, 56  
    ImageNet, 34, 56  
    iris, 9, 63, 133  
    MNIST, 14, 33, 45, 82, 100  
decision tree, 12, 48  
decoder, 130  
Deep Blue, 32  
deep learning, 5, 7, 34  
    algorithm, 8  
    toolkits, 94  
dense layer, 86  
diffusion model, 107  
    conditional, 109  
    DALL-E 2, 110  
    errors, 111  
    fixed pseudorandom number  
        generator seed, 111  
    forward algorithm, 109  
    noise schedule, 107–108  
    reverse algorithm, 109  
    Stable Diffusion, 110  
dinosaur footprints, 56  
discriminator, 97  
distribution, 18  
dropout, 40  
dual numbers, 95  
Duda, Richard, 30

## E

effective receptive field, 90  
embedding, 92  
embedding space, 130  
emergence, 135  
end-to-end learning, 81  
ensembling, 49, 52  
entanglement, 105  
epoch, 76  
Euclid, 125  
evolutionary algorithm, 59  
expert system, 31  
explainable AI, 22  
extrapolation, 16, 70

## F

Falkingham, Peter L., 56  
false negative, 21, 57  
false positive, 21, 57  
features, 10  
    feature selection, 49–50  
    feature space, 11  
    feature vector, 10  
filter, 88  
forward pass, 69  
Fridman, Lex, 150  
Fukushima, Kunihiko, 30, 33

## G

generalization, 11  
generative adversarial network, 34, 97  
    conditional, 103  
    controllable, 104  
    discriminator, 97  
    entanglement, 105  
    generator, 97  
    mode collapse, 103  
    noise vector, 100  
    training algorithm, 98  
generative AI, 34  
generative pretrained transformer  
    (GPT), 130  
generator, 97  
genetic programming, 59  
global minimum, 72

Gomes, Gabe, 148  
Goodfellow, Ian, 34  
gradient descent, 39, 71  
    global vs. local minima, 72  
    learning rate, 75  
graphics processing unit (GPU), 38

## H

Haffner, Patrick, 33  
hallucination, 128, 140  
Hart, Peter, 30  
hidden layer, 64  
Hinton, Geoffrey, 32  
Hopfield, John, 31  
hyperparameters, 55

## I

ImageNet, 34, 56  
in-context learning, 132  
inference, 57–58  
initialization, 67  
interpolation, 16, 70  
iris, 9, 63, 133

## K

Kahneman, Daniel, 121  
Kasparov, Garry, 32  
kernel, 54, 85  
Khan, Sal, 146  
Kosinski, Michal, 119

## L

Lallensack, Jens N., 56  
large language model (LLM), 116  
    alignment, 132  
    attention, 131  
    context encoding, 131  
    context window, 132  
    decoder, 130  
    embedding space, 130  
    emergence, 135  
    hallucination, 128, 140  
    in-context learning, 132  
    pretraining, 130  
    reinforcement learning, 34  
        from human feedback  
        (RLHF), 132

    token, 130  
    transformer, 130  
leaky ReLU, 101  
learning rate, 75  
LeCun, Yann, 33–34  
Lee, Peter, 148  
Leibniz, Gottfried, 27  
LeNet, 87  
Lighthill, James, 30  
local minimum, 73  
loss, 70  
Lovelace, Ada, 6, 27, 36

## M

machine learning, 5, 7  
    algorithm, 8  
MacKnight, Robert, 148  
manifold, 47  
matrix, 9  
Matthews correlation coefficient, 58  
McCarthy, John, 5, 29  
McCorduck, Pamela, 26  
McCulloch, Warren, 28–29  
Mechanical Turk, 26  
metric, 56  
minibatch, 76  
Minsky, Marvin, 29–30  
MNIST, 14, 33, 45, 82, 100  
mode collapse, 103  
model, 7  
    binary, 9  
    classifier, 12  
    decision tree, 12, 48  
    expert system, 31  
    generalization, 11  
    hyperparameters, 55  
    multiclass, 9, 55  
    nearest neighbor, 30, 44  
    optimization, 11, 54, 72  
    parameters, 8  
    random forest, 33, 49  
    support vector machine (SVM),  
        32, 52  
    training, 8  
*Mona Lisa*, 110–111  
multilayer perceptron, 78

## N

Nash, John, 29  
nearest neighbor, 30, 44  
Neocognitron, 31  
neural network, 26

- activation function, 62
  - AlexNet, 34
- architecture, 39, 69
- bias, 63
- convolutional, 30, 33, 86
- deep learning, 34
- dense layer, 86
- generative adversarial, 34, 97
- generative AI, 34
- hidden layer, 64
- initialization, 67
- loss, 70
- neuron, 62
- node, 63
- overfitting, 23, 70
- preprocessing, 66
- recurrent, 130
- regularizer, 70
- reinforcement learning, 34
- training algorithm, 69
- unsupervised learning, 34
- weight, 63

neuron, 62  
Noah, Trevor, 148  
node, 63  
noise vector, 100  
normal distribution, 99  
normalization, 40

## O

Offray de La Mettrie, Julien, 27  
one-hot encoding, 103  
one-versus-one, 56  
one-versus-rest, 56  
optimization, 11, 54, 72  
overfitting, 23, 70

## P

Pantoliano, Joe, 149  
paper clip catastrophe, 150  
Papert, Seymour, 30  
parameters, 8  
Perceptron, 29, 32, 78

Petro, Joseph, 148  
Pitts, Walter, 28  
pooling layer, 87  
preprocessing, 66  
pretraining, 130  
primary visual cortex, 84  
prompt engineer, 111

## R

random forest, 33, 49

- bagging, 49–50
- ensembling, 49, 52
- feature selection, 49–50

rectified linear unit (ReLU), 39

- leaky ReLU, 101

recurrent neural network (RNN), 130  
regularizer, 70  
reinforcement learning, 34

- from human feedback (RLHF), 132

Romilio, Anthony, 56  
Rosenblatt, Frank, 29, 115  
Rumelhart, David, 32  
Russell, Stuart, 150

## S

sample, 10  
schedule, 107  
Sedol, Lee, 35  
Sejnowski, Terrence, 140  
semantic segmentation, 93  
Shannon, Claude, 29  
Shazeer, Noam, 130  
Shelley, Mary, 27  
sigmoid, 64  
Solomonoff, Ray, 29  
“Sparks of Artificial General Intelligence,” 116  
Stable Diffusion, 110  
stochastic gradient descent, 76  
Stork, David, 30  
supervised learning, 8  
support vector machine (SVM), 32, 52

- kernel, 54, 85
- optimization, 11, 54, 72
- support vector, 54

swarm intelligence, 59  
symbolic AI, 26

## **T**

Tattersall, Ian, 37  
Tegmark, Max, 150  
testing, 10  
    test set, 10  
token, 130  
training, 8  
    training set, 10  
transformer, 130  
    attention, 131  
    decoder, 130  
    encoder, 130  
true negative, 57  
true positive, 57  
Turing, Alan, 6, 28  
Turing machine, 28  
Turing test, 28  
Tversky, Amos, 121

## **U**

Uhr, Leonard, 30  
unsupervised learning, 34

## **V**

van Gogh, Vincent, 83  
Vapnik, Vladimir, 32  
Vaswani, Ashish, 130  
vector, 9  
Vossler, Charles, 30

## **W**

weight, 63  
Williams, Ronald, 32  
Wooldridge, Michael, 26  
Wright, Letitia, 148