Exercise
1. Give one example of application that use AI in Rwanda. 10 marks. 5 minutes

RWASIS
IREMBO
pindo.ai
gpt
deepseek
gemni

RWANDA POLYTECHNIC
HUYE COLLEGE
ICT DEPARTMENT
ITLPA701: PYTHON AND FUNDAMENTALS OF AI
15 LEARNING HOURS

NTAMBARA Etienne, Assistant Lecturer

February 9, 2025

Artificial Intelligence (AI) refers to the simulation of human intelligence in machines that are programmed to think, learn, solving problems and make decisions.

NTAMBARA Etienne 0783716761 [1 2 3] February 9, 2025

1. LEARNING UNIT 5–DEVELOP AI BASED APPLICATIONS

1.1. Learning Outcome 5.1: Introduce AI

1.1.1. Definitions of Key Terms

1. What is AI?

In today's world, technology is growing very fast, and we are getting in touch with different new technologies day by day. One of the booming technologies of computer science is **Artificial Intelligence (AI)**, which is ready to create a new revolution in the world by making intelligent machines. AI is now all around us, currently working with a variety of subfields, ranging from general to specific, such as self-driving cars, playing chess, proving theorems, playing music, painting, etc (Russell & Norvig, 2010).

AI is one of the fascinating and universal fields of computer science with great scope in the future. AI holds a tendency to cause a machine to work as a human (McCarthy, 2007). Artificial Intelligence is composed of two words: **Artificial** and **Intelligence**, where Artificial defines "man-made," and intelligence defines "thinking power"; hence AI means "a man-made thinking power."

Definition: AI is a branch of computer science by which we can create intelligent machines that can behave like a human, think like humans, and make decisions (Russell & Norvig, 2010). AI exists when a machine can have human-based skills such as learning, reasoning, and solving problems (Goodfellow et al., 2016). With AI, you do not need to preprogram a machine to do some work; instead, you can create a machine with programmed algorithms that can work with its intelligence. This is the essence of AI (McCarthy, 2007).

2. History of AI

The idea of *a machine that thinks* dates back to ancient Greece. However, the advent of electronic computing brought significant milestones in AI development:

- **1950:** Alan Turing publishes *Computing Machinery and Intelligence*, introducing the Turing Test (Turing, 1950).
- **1956:** John McCarthy coins the term *Artificial Intelligence* at Dartmouth College (McCarthy et al., 1956).
- **1967:** Frank Rosenblatt builds the Mark 1 Perceptron, the first computer-based neural network (Rosenblatt, 1958).
- **1980s:** Neural networks using backpropagation become widely used in AI applications (Rumelhart et al., 1986).
- **1997:** IBM's Deep Blue defeats world chess champion Garry Kasparov (Hammond, 1997).

- **2011:** IBM Watson wins against Jeopardy! champions Ken Jennings and Brad Rutter (Ferrucci, 2012).
- **2015:** Baidu's Minwa supercomputer surpasses human image recognition accuracy (Schmidhuber, 2015).
- **2016:** DeepMind's AlphaGo defeats world champion Go player Lee Sodol (Silver et al., 2016).

3. Benefits of AI

1. **Automation**: AI automates repetitive tasks, increasing efficiency and productivity (Russell & Norvig, 2020).
2. **Smart Decision Making**: AI analyzes trends, delivers data, and helps in informed decision-making (Goodfellow et al., 2016).
3. **Enhanced Customer Experience**: AI chatbots and NLP solutions improve customer support services (Hirschberg & Manning, 2015).
4. **Medical Advances**: AI helps in remote patient monitoring and disease prediction (Topol, 2019).
5. **Innovation**: AI enables new products, services, and advancements in robotics and NLP (Furst, 2018).
   NLP: Natural Language Processing
6. **Accessibility**: AI aids disabled individuals by enhancing mobility and communication (Koch & Marschollek, 2018).
7. **Business Continuity**: AI assists in business forecasting and risk management (**?**).
8. **Managing Repetitive Tasks**: AI-driven RPA tools reduce manual labor in business processes (van der Aalst & Song, 2020).
9. **Minimizing Errors**: AI reduces human errors in data processing and entry (Marcus, 2020).
10. **Increased Business Efficiency**: AI ensures 24/7 service availability and performance consistency (Brynjolfsson & Rock, 2018).

1.1.2. Types of AI

1. Weak AI

Narrow AI, also known as *Weak AI*, is designed to perform specific tasks within a limited domain. Examples include:

- Virtual assistants (e.g., Siri, Alexa) (Russell & Norvig, 2020).
- Recommendation systems (e.g., Netflix, Amazon) (Ricci & Rokach, 2011).
- Image recognition software (Yann LeCun & Hinton, 2015).

2 Strong AI

Strong AI consists of Artificial General Intelligence (AGI) and Artificial Super Intelligence (ASI). AGI refers to AI systems possessing human-like intelligence, capable of solving problems, learning, and planning for the future (Goertzel & Pennachin, 2014). ASI surpasses human intelligence and remains a theoretical concept (Bostrom, 2014).

1.1.3. Artificial Intelligence (AI) in Real Life

Artificial Intelligence (AI) refers to the simulation of human intelligence in machines that are programmed to think, learn, and make decisions. Below are some **real-life examples of AI**:

1. Self-Driving Cars

Self-driving cars, such as those developed by **Tesla** and **Waymo**, use AI algorithms to perceive their environment, make decisions, and navigate without human intervention. These cars rely on sensors, cameras, and machine learning models to detect obstacles, interpret traffic signals, and plan routes.

2. Navigation Systems

Navigation systems like **Google Maps** and **Waze** use AI to analyze real-time traffic data, predict congestion, and suggest the fastest routes. Machine learning algorithms process vast amounts of data from users and sensors to optimize travel time.

3. Chatbots

AI-powered chatbots, such as **ChatGPT** and **Google Assistant**, use Natural Language Processing (NLP) to understand and respond to user queries. These systems are widely used in customer service, healthcare, and education to provide instant support and information.

4. Human vs. Computer Games

AI has demonstrated its capabilities in strategic games like **Chess** and **Go**. For example, **Deep Blue** (IBM) defeated world chess champion Garry Kasparov in 1997, and **AlphaGo** (DeepMind) beat world Go champion Lee Sedol in 2016. These systems use advanced algorithms and reinforcement learning to master complex games.

5. Sophia Robot

Sophia, developed by Hanson Robotics, is a humanoid robot that uses AI to simulate human-like expressions, recognize faces, and hold conversations. Sophia represents advancements in robotics and AI integration, showcasing how machines can mimic human behavior.

6. Turing Machine

The **Turing Machine**, proposed by Alan Turing, is a theoretical model of computation that laid the foundation for modern computers and AI. While not a physical machine, it inspired the development of algorithms and computational models that power AI systems today.

7. Many More Applications

AI is also used in:

- **Healthcare**: Diagnosing diseases (e.g., IBM Watson for Oncology).
- **Finance**: Fraud detection and algorithmic trading.
- **E-commerce**: Personalized recommendations (e.g., Amazon, Netflix).
- **Manufacturing**: Predictive maintenance and automation.

1.1.4. The Future of Artificial Intelligence (AI)

Artificial Intelligence (AI) is rapidly evolving and is expected to revolutionize various fields in the future. Below are some key areas where AI is likely to have a significant impact:

1. Military Bots

AI-powered military bots, such as autonomous drones and robotic soldiers, are being developed for surveillance, reconnaissance, and combat. These systems can reduce human casualties and enhance strategic decision-making. For example, the U.S. Department of Defense is investing in projects like **Project Maven** to integrate AI into military operations (of Defense, 2020).

2. The Perfect Lawyer

AI is poised to transform the legal industry by automating tasks like document review, legal research, and contract analysis. Tools like **ROSS Intelligence** (Intelligence, 2021) and **DoNotPay** (DoNotPay, 2022) already assist lawyers and individuals in navigating legal complexities. In the future, AI could even predict case outcomes based on historical data, making it the ”perfect lawyer.”

3. Music

AI is revolutionizing the music industry by composing, producing, and performing music. Tools like **OpenAI’s Jukedeck** (OpenAI, 2020) and **Amper Music** (Music, 2021) enable users to create original compositions using AI algorithms. In the future, AI could collaborate with human artists to push the boundaries of creativity and produce personalized music for listeners.

4. Business

AI is transforming businesses by optimizing operations, enhancing customer experiences, and enabling data-driven decision-making. For example:

- **Supply Chain Management**: AI predicts demand and optimizes logistics (Smith, 2022).
- **Customer Service**: AI-powered chatbots provide 24/7 support (Lee, 2021).
- **Marketing**: AI analyzes consumer behavior to deliver personalized campaigns (Brown, 2020).

In the future, AI could automate entire business processes, making organizations more efficient and competitive.

5. Healthcare

AI is set to revolutionize healthcare by improving diagnostics, treatment, and patient care. Examples include:

- **Diagnostics**: AI systems like **IBM Watson Health** analyze medical data to detect diseases early (IBM, 2021).
- **Personalized Medicine**: AI tailors treatments based on genetic and clinical data (Johnson, 2022).
- **Robotic Surgery**: AI-assisted robots perform precise and minimally invasive surgeries (Williams, 2021).

In the future, AI could enable predictive healthcare, where diseases are prevented before they occur.

1.2. Learning Outcome 5.2: Implement Machine Learning

1.2.1. Definition of Key Terms

1. Data Splitting

Definition: Dividing a dataset into subsets (e.g., training, validation, and test sets) to evaluate the performance of a machine learning model.

Python Example:

```python
from sklearn.datasets import load_iris
from sklearn.model_selection import
    train_test_split

# Load the Iris dataset
data = load_iris()
X, y = data.data, data.target

# Split into training (70%) and test (30%)
    sets
X_train, X_test, y_train, y_test =
    train_test_split(X, y, test_size=0.3,
    random_state=42)

print("Training set size:", X_train.shape
    [0])
print("Test set size:", X_test.shape[0])
```

2. Numerical Data

Definition: Data represented in numbers (e.g., age, temperature, salary).

Python Example:

```python
import pandas as pd

# Example of numerical data
data = {'Age': [25, 30, 35, 40], 'Salary':
    [50000, 60000, 70000, 80000]}
df = pd.DataFrame(data)
print(df)
```

3. Categorical Data

Definition: Data that represents categories or groups (e.g., gender, color, country).

Python Example:

```python
import pandas as pd

# Example of categorical data
data = {'Gender': ['Male', 'Female', '
    Female', 'Male'], 'Country': ['USA', '
    Canada', 'UK', 'USA']}
df = pd.DataFrame(data)
print(df)
```

4. Ordinal Data

Definition: Categorical data with a specific order or ranking (e.g., education level: high school, bachelor's, master's).

Python Example:

```python
import pandas as pd

# Example of ordinal data
data = {'Education': ['High School', '
    Bachelors', 'Masters', 'PhD']}
df = pd.DataFrame(data)
print(df)
```

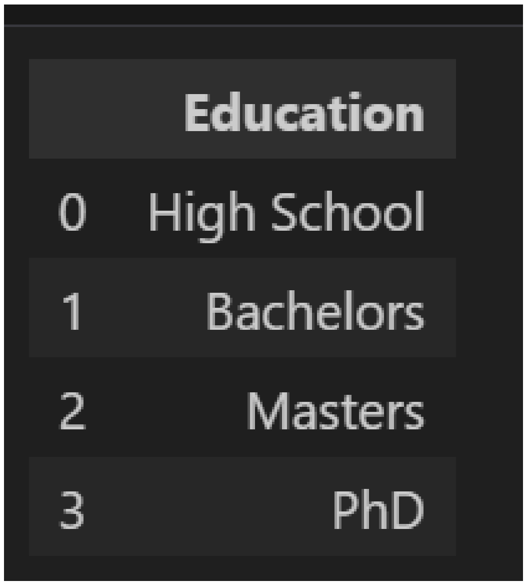

	Education
0	High School
1	Bachelors
2	Masters
3	PhD

Figure 1. Output.

5. Percentile

Definition: A measure indicating the value below which a given percentage of observations in a dataset fall.

Python Example:

```python
import numpy as np

# Example of calculating percentiles
data = [10, 20, 30, 40, 50, 60, 70, 80, 90,
     100]
print("25th Percentile:", np.percentile(
    data, 25))
print("50th Percentile (Median):", np.
    percentile(data, 50))
print("75th Percentile:", np.percentile(
    data, 75))
```

6. Data Distribution

Definition: The way data points are spread or distributed across a range (e.g., normal distribution, skewed distribution).

Python Example:

```python
import matplotlib.pyplot as plt
import seaborn as sns
import numpy as np

# Generate random data with a normal
    distribution
data = np.random.normal(0, 1, 1000)

# Plot the distribution
sns.histplot(data, kde=True)
plt.title("Normal Distribution")
plt.show()
```

7. Algorithm

Definition: A set of rules or steps used to solve a problem or perform a task (e.g., linear regression, decision trees).

Python Example:

```python
from sklearn.linear_model import
    LinearRegression
import numpy as np

# Example of a linear regression algorithm
X = np.array([[1], [2], [3], [4]])
y = np.array([2, 4, 6, 8])

model = LinearRegression()
model.fit(X, y)

print("Prediction for X=5:", model.predict
    ([[5]]))
```

Classification Algo:
1. Logistic regression
2. Naive Bayes

Regression Algo
1. Linear Regression

Others Works on Both
1. Random Forest
2.Decision tree
3. SVM
4. XGBOOST
5. LightGBM
6. CAT BOOST

8. Model

Definition: The output of a machine learning algorithm after training on data, used to make predictions.

Python Example:

```python
from sklearn.datasets import load_iris
from sklearn.tree import
    DecisionTreeClassifier
from sklearn.model_selection import
    train_test_split

# Load the Iris dataset
data = load_iris()
X, y = data.data, data.target

# Split into training and test sets
X_train, X_test, y_train, y_test =
    train_test_split(X, y, test_size=0.3,
    random_state=42)

# Train a decision tree model
model = DecisionTreeClassifier()
```

STOP HERE

```python
model.fit(X_train, y_train)

# Make predictions
y_pred = model.predict(X_test)
print("Predictions:", y_pred)
```

9. Scatter Plot

Definition: A graphical representation of data points on a two-dimensional plane, used to visualize relationships between variables.

Python Example:

```python
import matplotlib.pyplot as plt
import numpy as np

# Generate random data
X = np.random.rand(50)
y = 2 * X + np.random.randn(50)

# Create a scatter plot
plt.scatter(X, y)
plt.title("Scatter Plot Example")
plt.xlabel("X")
plt.ylabel("y")
plt.show()
```

10. Importance of Machine Learning

Definition: Machine learning enables automation, prediction, and decision-making based on data, revolutionizing industries like healthcare, finance, and technology.

Python Example:

```python
from sklearn.datasets import load_iris
from sklearn.model_selection import
    train_test_split
from sklearn.ensemble import
    RandomForestClassifier
from sklearn.metrics import accuracy_score

# Load the Iris dataset
data = load_iris()
X, y = data.data, data.target

# Split into training and test sets
X_train, X_test, y_train, y_test =
    train_test_split(X, y, test_size=0.3,
    random_state=42)

# Train a Random Forest model
model = RandomForestClassifier()
model.fit(X_train, y_train)

# Make predictions
y_pred = model.predict(X_test)

# Evaluate the model
accuracy = accuracy_score(y_test, y_pred)
print("Model Accuracy:", accuracy)
```

11. Data Splicing

Definition: A technique to split or segment data for analysis or processing.

Python Example:

```python
import pandas as pd

# Example of data splicing
data = {'Values': [10, 20, 30, 40, 50, 60,
    70, 80, 90, 100]}
df = pd.DataFrame(data)

# Splice the first 5 rows
spliced_data = df[:5]
print("Spliced Data:\n", spliced_data)
```

1.2.2. Define Variables and Data

Working with bellow data-frame as an example

```python
# Python Code for Defining Variables
import pandas as pd
from sklearn.model_selection import
    train_test_split

data = {
    'Size': [1200, 1500, 1800, 2000, 2200],
    'Price': [150000, 200000, 250000,
    300000, 350000]
}
df = pd.DataFrame(data)

X = df[['Size']]
y = df['Price']

X_train, X_test, y_train, y_test =
    train_test_split(X, y, test_size=0.2,
    random_state=42)
```

1. Predictor Variable (Independent Variable)

The predictor variable is the input feature used to predict the output. In this case, the predictor variable is the **Size of the house**, which will be used to predict the house price.

```python
# Predictor variable (independent variable)
X = df[['Size']]  # Size is the predictor variable
```

2. Response Variable (Dependent Variable)

The response variable is the output we aim to predict. In this case, the response variable is the **Price of the house**, which we will predict based on the size.

```python
# Response variable (dependent variable)
y = df['Price']  # Price is the response variable
```

3. Training Data

Training data is the dataset used to train the ML model. It includes both predictor and response variables. In this case, we use the dataset to train the model to learn the relationship between house size and price.

```python
# Split the dataset into training and testing data
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)
```

4. Testing Data

Testing data is used to evaluate the performance of the trained model. It is separate from the training data. In this case, after training the model, we use the testing data to check how well the model performs on unseen data.

```python
# The testing data is used for evaluation
# We will use the testing data (X_test) to predict and evaluate the model
```

5. Data Scraping

Data scraping involves extracting data from websites or other sources. In this example, we scrape house sizes and prices from a website for use in the model.

```python
# Data Scraping Example
url = "https://example.com/house-prices"
response = requests.get(url)
soup = BeautifulSoup(response.content, 'html.parser')

# Extract house sizes and prices from the website
house_sizes = []
house_prices = []

for item in soup.find_all('div', class_='house-item'):
    size = item.find('span', class_='size').text
    price = item.find('span', class_='price').text
    house_sizes.append(int(size))
    house_prices.append(int(price.replace('$', '').replace(',', '')))

# Create a DataFrame from scraped data
scraped_data = pd.DataFrame({'Size': house_sizes, 'Price': house_prices})
```

1.2.3. Machine Learning Processes

The ML process (life-cycle) includes defining the objective, data gathering, preparing data, data exploration, building the model, model evaluation, and making predictions.

1. Define Objective

The first step is to define the objective of the ML project.

```python
# Define the objective
objective = "Predict house prices based on features like size, location, and number of bedrooms."
print("Objective:", objective)
```

2. Data Gathering

Data gathering involves collecting the necessary data for the ML model.

```python
import pandas as pd

# Load dataset from a CSV file
data = pd.read_csv("house_prices.csv")
print("Data Head:")
print(data.head())
```

3. Preparing Data

Data preparation involves cleaning and preprocessing the data.

```python
# Handle missing values
data = data.dropna()

# Encode categorical variables
data = pd.get_dummies(data, columns=['
    Location'], drop_first=True)

print("Prepared Data Head:")
print(data.head())
```

4. Data Exploration

Data exploration involves analyzing the dataset to understand its structure.

```python
import matplotlib.pyplot as plt

# Scatter plot of Size vs Price
plt.scatter(data['Size'], data['Price'])
plt.xlabel("Size (sq ft)")
plt.ylabel("Price ($)")
plt.title("House Size vs Price")
plt.show()
```

5. Building Model

Building the model involves selecting an algorithm and training it.

```python
from sklearn.linear_model import
    LinearRegression
from sklearn.model_selection import
    train_test_split

# Define predictor (X) and response (y)
    variables
X = data[['Size', 'Bedrooms', '
    Location_Urban']]
y = data['Price']

# Split data into training and testing sets
X_train, X_test, y_train, y_test =
    train_test_split(X, y, test_size=0.2,
    random_state=42)

# Build and train the model
model = LinearRegression()
model.fit(X_train, y_train)

print("Model trained successfully!")
```

6. Model Evaluation

Model evaluation involves assessing the model's performance.

```python
from sklearn.metrics import
    mean_squared_error, r2_score

# Make predictions on the testing data
y_pred = model.predict(X_test)

# Evaluate the model
mse = mean_squared_error(y_test, y_pred)
r2 = r2_score(y_test, y_pred)

print(f"Mean Squared Error: {mse}")
print(f"R-squared: {r2}")
```

7. Predictions

The trained model can be used to make predictions on new data.

```python
# Predict the price of a new house
new_house = [[1500, 3, 1]]  # Size: 1500 sq
     ft, Bedrooms: 3, Location: Urban
predicted_price = model.predict(new_house)

print(f"Predicted Price: ${predicted_price
    [0]:.2f}")
```

1.2.4. Types of Machine Learning

Machine learning can be supervised, unsupervised, or reinforcement learning.

```python
# Python Code for Supervised and
    Unsupervised Learning
from sklearn.linear_model import
    LinearRegression
from sklearn.cluster import KMeans

# Supervised Learning
model = LinearRegression()
model.fit(X_train, y_train)

# Unsupervised Learning
kmeans = KMeans(n_clusters=2)
kmeans.fit(X)
```

Machine learning (ML) can be categorized into three primary types based on the type of data and the method of learning: Supervised Learning, Unsupervised Learning, and Reinforcement Learning.

1. Supervised Learning

Supervised learning is a type of machine learning where the model is trained on labeled data. In other words, the input data (X) and the corresponding output (Y) are provided during training. The goal is to learn a mapping from inputs to outputs, and once the model is trained, it can predict the output for new, unseen data.

Example: Linear Regression

In supervised learning, the model is trained to predict the response variable (dependent variable) based on the predictor variables (independent variables). For example, predicting the price of a house based on its size and number of rooms is a supervised learning problem.

```python
# Example of Supervised Learning: Linear Regression

from sklearn.linear_model import LinearRegression
from sklearn.model_selection import train_test_split
from sklearn.metrics import mean_squared_error
import pandas as pd

# Create a dataset
data = {'Size': [1200, 1500, 1800, 2000, 2200, 2400, 2600, 2800, 3000, 3200],
        'Price': [150000, 200000, 250000, 300000, 350000, 400000, 450000, 500000, 550000, 600000]}

df = pd.DataFrame(data)

# Define predictor and response variables
X = df[['Size']]  # Predictor variable
y = df['Price']   # Response variable

# Split the data into training and testing sets
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)

# Train the model using Linear Regression
model = LinearRegression()
model.fit(X_train, y_train)

# Make predictions
y_pred = model.predict(X_test)

# Evaluate the model
mse = mean_squared_error(y_test, y_pred)
print(f'Mean Squared Error: {mse}')
```

In this example, the model learns the relationship between the house size and the price, and it can predict the price of a house based on its size.

2. Unsupervised Learning

Unsupervised learning is used when the model is provided with data that has no labels (i.e., there is no output variable to predict). The goal is to find hidden patterns or intrinsic structures in the data. Common tasks in unsupervised learning include clustering and dimensionality reduction.

Example: K-Means Clustering

In unsupervised learning, one common task is clustering, where the model groups similar data points together. For instance, grouping houses into clusters based on similar characteristics like size, price, and location.

```python
# Example of Unsupervised Learning: K-Means Clustering

from sklearn.cluster import KMeans
import numpy as np

# Example data: House sizes and prices
data = np.array([[1200, 150000],
                 [1500, 200000],
                 [1800, 250000],
                 [2000, 300000],
                 [2200, 350000],
                 [2400, 400000],
                 [2600, 450000],
                 [2800, 500000],
                 [3000, 550000],
                 [3200, 600000]])

# Create the KMeans model
```

```python
kmeans = KMeans(n_clusters=2, random_state
    =42)

# Fit the model to the data
kmeans.fit(data)

# Print the cluster centers
print("Cluster Centers:", kmeans.
    cluster_centers_)

# Predict which cluster each data point
    belongs to
clusters = kmeans.predict(data)
print("Cluster Assignments:", clusters)
```

In this example, the K-Means algorithm groups houses into two clusters based on size and price. The model identifies patterns within the data and groups similar houses together.

3. Reinforcement Learning

Reinforcement learning is a type of machine learning where an agent learns by interacting with an environment. The agent takes actions, receives feedback in the form of rewards or penalties, and adjusts its strategy accordingly to maximize cumulative rewards. This type of learning is used in scenarios where the agent must make decisions sequentially, such as playing a game or controlling a robot.

Example: Q-Learning

In reinforcement learning, the agent learns a policy to take the best action given a state in order to maximize its reward. For example, a robot learns to navigate a maze by receiving positive rewards for correct moves and negative penalties for wrong ones.

```python
# Example of Reinforcement Learning: Q-
    Learning (simplified)

import numpy as np

# Define the environment (grid)
grid = np.array([[0, 0, 0, 0],
                 [0, -1, 0, 0],
                 [0, 0, 0, 10]])

# Define Q-table
Q = np.zeros_like(grid, dtype=float)

# Define parameters
learning_rate = 0.1
discount_factor = 0.9
epsilon = 0.1
iterations = 1000

# Q-learning algorithm
for _ in range(iterations):
    state = (2, 0)  # Starting position
    while state != (0, 3):  # Goal state
        # Explore or exploit
        if np.random.rand() < epsilon:
            action = np.random.choice([0,
    1, 2, 3])  # Random action (explore)
        else:
            action = np.argmax(Q[state])  #
     Best action (exploit)

        # Take action and observe the next
    state
        if action == 0:
            next_state = (max(0, state[0] -
     1), state[1])  # Move up
        elif action == 1:
            next_state = (min(2, state[0] +
     1), state[1])  # Move down
        elif action == 2:
            next_state = (state[0], max(0,
    state[1] - 1))  # Move left
        else:
            next_state = (state[0], min(3,
    state[1] + 1))  # Move right

        # Reward is based on the
    environment grid
        reward = grid[next_state]

        # Update Q-table
        Q[state] = Q[state] + learning_rate
     * (reward + discount_factor * np.max(Q
    [next_state]) - Q[state])

        state = next_state

# Output the learned Q-table
print("Q-table:", Q)
```

In this example, the Q-learning algorithm helps the agent (robot) learn how to navigate a grid to reach the goal while maximizing the total reward. It explores different actions and exploits the best actions based on the learned Q-values.

1.2.5. Machine Learning Algorithms

Common ML algorithms include linear regression, logistic regression, decision trees, random forests, KNN, and SVM.

```python
# Python Code for Various ML Algorithms
from sklearn.linear_model import LinearRegression, LogisticRegression
from sklearn.tree import DecisionTreeClassifier
from sklearn.ensemble import RandomForestClassifier
from sklearn.neighbors import KNeighborsClassifier
from sklearn.svm import SVC

# Linear Regression
model = LinearRegression()
model.fit(X_train, y_train)

# Logistic Regression
model = LogisticRegression()
model.fit(X_train, y_train)

# Decision Tree
model = DecisionTreeClassifier()
model.fit(X_train, y_train)

# Random Forest
model = RandomForestClassifier()
model.fit(X_train, y_train)

# K Nearest Neighbour
model = KNeighborsClassifier(n_neighbors=3)
model.fit(X_train, y_train)

# Support Vector Machine
model = SVC()
model.fit(X_train, y_train)
```

In this section, we discuss some of the popular machine learning algorithms, including both supervised and unsupervised methods, and provide Python code examples for each.

1. Linear Regression

Linear regression is a supervised learning algorithm used to predict a continuous target variable based on one or more input features. It finds the linear relationship between the input features and the target variable (Smith, 2018).

Example: Predicting House Price Based on Size

```python
# Linear Regression Example

from sklearn.linear_model import LinearRegression
from sklearn.model_selection import train_test_split
import pandas as pd

# Create a dataset
data = {'Size': [1200, 1500, 1800, 2000, 2200, 2400, 2600, 2800, 3000, 3200],
        'Price': [150000, 200000, 250000, 300000, 350000, 400000, 450000, 500000, 550000, 600000]}

df = pd.DataFrame(data)

# Define predictor and response variables
X = df[['Size']]  # Predictor variable
y = df['Price']   # Response variable

# Split the data into training and testing sets
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)

# Train the model using Linear Regression
model = LinearRegression()
model.fit(X_train, y_train)

# Make predictions
y_pred = model.predict(X_test)
```

2. Logistic Regression

Logistic regression is used for binary classification problems (Doe, 2019).

Example: Predicting the Probability of Default on a Loan

```python
# Logistic Regression Example

from sklearn.linear_model import LogisticRegression
from sklearn.model_selection import train_test_split
from sklearn.metrics import accuracy_score

```

```python
# Sample dataset for binary classification
data = {'Age': [25, 30, 35, 40, 45, 50, 55, 60, 65, 70],
        'Income': [35000, 40000, 45000, 50000, 55000, 60000, 65000, 70000, 75000, 80000],
        'Default': [0, 0, 0, 1, 1, 1, 1, 0, 0, 0]}  # 0: No Default, 1: Default

df = pd.DataFrame(data)

# Define predictor and response variables
X = df[['Age', 'Income']]
y = df['Default']

# Split the data into training and testing sets
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.3, random_state=42)

# Train the model using Logistic Regression
model = LogisticRegression()
model.fit(X_train, y_train)

# Make predictions
y_pred = model.predict(X_test)

# Evaluate the model
accuracy = accuracy_score(y_test, y_pred)
print(f'Accuracy: {accuracy}')
```

3. Decision Tree

A decision tree is a supervised learning algorithm (Johnson, 2017a).

Example: Classifying Animals Based on Features

```python
# Decision Tree Example

from sklearn.tree import DecisionTreeClassifier
from sklearn.model_selection import train_test_split
from sklearn.metrics import accuracy_score

# Sample dataset for classification
data = {'Weight': [20, 10, 50, 100, 200],
        'Height': [0.5, 0.3, 1.0, 1.5, 2.0],
        'Animal': ['Dog', 'Cat', 'Horse', 'Elephant', 'Giraffe']}

df = pd.DataFrame(data)

# Define predictor and response variables
X = df[['Weight', 'Height']]
y = df['Animal']

# Split the data into training and testing sets
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.3, random_state=42)

# Train the model using Decision Tree
model = DecisionTreeClassifier()
model.fit(X_train, y_train)

# Make predictions
y_pred = model.predict(X_test)

# Evaluate the model
accuracy = accuracy_score(y_test, y_pred)
print(f'Accuracy: {accuracy}')
```

4. Random Forest

Random Forest is an ensemble learning method (Lee, 2020).

Example: Predicting Customer Churn

```python
# Random Forest Example

from sklearn.ensemble import RandomForestClassifier
from sklearn.model_selection import train_test_split
from sklearn.metrics import accuracy_score

# Sample dataset for classification
data = {'Age': [25, 30, 35, 40, 45, 50, 55, 60, 65, 70],
        'Income': [35000, 40000, 45000, 50000, 55000, 60000, 65000, 70000, 75000, 80000],
        'Churn': [0, 0, 1, 1, 0, 1, 1, 0, 0, 1]}  # 0: No Churn, 1: Churn

df = pd.DataFrame(data)
```

```

# Define predictor and response variables
X = df[['Age', 'Income']]
y = df['Churn']

# Split the data into training and testing
    sets
X_train, X_test, y_train, y_test =
    train_test_split(X, y, test_size=0.3,
    random_state=42)

# Train the model using Random Forest
model = RandomForestClassifier(n_estimators
    =100, random_state=42)
model.fit(X_train, y_train)

# Make predictions
y_pred = model.predict(X_test)

# Evaluate the model
accuracy = accuracy_score(y_test, y_pred)
print(f'Accuracy: {accuracy}')
```

5. K-Nearest Neighbors (KNN)

K-Nearest Neighbors is a simple, instance-based learning algorithm (Brown, 2021).

Example: Classifying Flowers Based on Petal Length and Width

```
# K-Nearest Neighbors Example

from sklearn.neighbors import
    KNeighborsClassifier
from sklearn.model_selection import
    train_test_split
from sklearn.metrics import accuracy_score

# Sample dataset for classification
data = {'Petal Length': [1.4, 1.3, 1.5,
    1.4, 1.7],
        'Petal Width': [0.2, 0.2, 0.3, 0.3,
     0.4],
        'Species': ['Setosa', 'Setosa', '
    Versicolor', 'Versicolor', 'Virginica'
    ]}

df = pd.DataFrame(data)

# Define predictor and response variables
X = df[['Petal Length', 'Petal Width']]
y = df['Species']

# Split the data into training and testing
    sets
X_train, X_test, y_train, y_test =
    train_test_split(X, y, test_size=0.3,
    random_state=42)

# Train the model using K-Nearest Neighbors
model = KNeighborsClassifier(n_neighbors=3)
model.fit(X_train, y_train)

# Make predictions
y_pred = model.predict(X_test)

# Evaluate the model
accuracy = accuracy_score(y_test, y_pred)
print(f'Accuracy: {accuracy}')
```

6. Support Vector Machine (SVM)

Support Vector Machine is a supervised learning algorithm (Davis, 2016).

Example: Classifying Text Data

```
# Support Vector Machine Example

from sklearn.svm import SVC
from sklearn.model_selection import
    train_test_split
from sklearn.metrics import accuracy_score

# Sample dataset for classification
data = {'Text Length': [200, 150, 300, 500,
     600],
        'Text Complexity': [1, 0, 2, 4, 5],
        'Category': ['Sports', 'Politics',
    'Sports', 'Technology', 'Politics']}

df = pd.DataFrame(data)

# Define predictor and response variables
X = df[['Text Length', 'Text Complexity']]
y = df['Category']

# Split the data into training and testing
    sets
X_train, X_test, y_train, y_test =
    train_test_split(X, y, test_size=0.3,
    random_state=42)
```

```

# Train the model using Support Vector
    Machine
model = SVC(kernel='linear')
model.fit(X_train, y_train)

# Make predictions
y_pred = model.predict(X_test)

# Evaluate the model
accuracy = accuracy_score(y_test, y_pred)
print(f'Accuracy: {accuracy}')
```

1.3. Learning Outcome 5.3: Building Data Models

This section focuses on building data models using various machine learning techniques, particularly Artificial Neural Networks (ANNs) and Deep Learning (DL).

1.3.1. Artificial Neural Networks (ANNs)

1. Definitions

Artificial Neural Networks (ANNs) are a class of machine learning models inspired by the human brain. They consist of layers of nodes, or "neurons", that are connected by weighted edges. These models can learn complex patterns from data, making them particularly useful in tasks like classification, regression, and even reinforcement learning (Johnson, 2017b).

2. Python Example: Artificial Neural Network (ANN)

Below is an example of how to implement a simple Artificial Neural Network using the Keras library for regression problems, like predicting house prices based on a feature (e.g., size of the house):

```
import numpy as np
from tensorflow.keras.models import
    Sequential
from tensorflow.keras.layers import Dense
from sklearn.model_selection import
    train_test_split
import pandas as pd

# Sample data (house size vs price)
data = {
    'Size': [1200, 1500, 1800, 2000, 2200,
    2400, 2600, 2800, 3000, 3200],
    'Price': [150000, 200000, 250000,
    300000, 350000, 400000, 450000, 500000,
     550000, 600000]
}
df = pd.DataFrame(data)

# Define predictor (X) and response (y)
X = df[['Size']]
y = df['Price']

# Split the data into training and testing
    sets
X_train, X_test, y_train, y_test =
    train_test_split(X, y, test_size=0.2,
    random_state=42)

# Build the ANN model
model = Sequential()
model.add(Dense(units=64, activation='relu'
    , input_dim=1))  # Input layer
model.add(Dense(units=64, activation='relu'
    ))  # Hidden layer
model.add(Dense(units=1))  # Output layer

# Compile the model
model.compile(optimizer='adam', loss='
    mean_squared_error')

# Train the model
model.fit(X_train, y_train, epochs=100,
    batch_size=10, verbose=1)

# Evaluate the model
loss = model.evaluate(X_test, y_test)
print(f'Mean Squared Error on test data: {
    loss}')
```

3. Use Case Implementation Steps

Building an artificial neural network involves several key steps:

- **Data Preprocessing:** Prepare the data by normalizing and cleaning it.
- **Model Architecture:** Choose the number of layers, neurons per layer, and activation functions.
- **Training the Model:** Use backpropagation to train the

model by adjusting weights based on errors.

- **Evaluation:** Test the model on unseen data to assess its accuracy.
- **Hyperparameter Tuning:** Experiment with different learning rates, activation functions, etc.

Neuro Networks Examples

- **Feedforward Neural Networks:** The simplest type of neural network where information moves in one direction, from input to output.
- **Convolutional Neural Networks (CNNs):** Specialized for processing grid-like data, such as images.
- **Recurrent Neural Networks (RNNs):** Used for sequential data, such as time series or language data.

1.3.2. Building a Model: Steps

1. Identify Business Problems

Before building a data model, it's crucial to define the business problem you are solving. For example, predicting customer churn in a telecommunications company or detecting fraudulent transactions in financial services.

2. Identify and Understand Data

Understanding the available data is essential for choosing the right model. Determine the type of data you have (e.g., numerical, categorical) and identify key features that might influence the model's prediction.

3. Collect and Prepare Data

Data preparation steps include:

- Handling missing values
- Encoding categorical variables
- Normalizing numerical values

Proper data preparation ensures that the model performs optimally.

4. Determine Models and Train Data

Once the data is ready, select appropriate models (e.g., neural networks, decision trees, etc.) and train them on the dataset using an appropriate training method.

5. Evaluate Models

After training, evaluate the model's performance using metrics such as accuracy, precision, recall, and F1-score.

6. Experiment and Adjust Model

Based on the evaluation, experiment with different algorithms, architectures, and hyperparameters to improve the model's performance.

1.3.3. Deep Learning

1. Definitions

Deep learning is a subset of machine learning that involves neural networks with many layers (hence the term "deep"). These models can automatically learn high-level representations of data, making them suitable for tasks like image recognition, natural language processing (NLP), and more (Goodfellow, 2016).

Python Example: Deep Learning Model (CNN for Image Classification)

Here is an example of a Convolutional Neural Network (CNN) implemented using Keras for image classification:

```python
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Conv2D, MaxPooling2D, Flatten, Dense
from tensorflow.keras.datasets import mnist
from tensorflow.keras.utils import to_categorical

# Load dataset (MNIST)
(x_train, y_train), (x_test, y_test) = mnist.load_data()

# Reshape and normalize the data
x_train = x_train.reshape(x_train.shape[0], 28, 28, 1).astype('float32') / 255
x_test = x_test.reshape(x_test.shape[0], 28, 28, 1).astype('float32') / 255
```

```

# One-hot encode the labels
y_train = to_categorical(y_train, 10)
y_test = to_categorical(y_test, 10)

# Build the CNN model
model = Sequential()
model.add(Conv2D(32, kernel_size=(3, 3),
    activation='relu', input_shape=(28, 28,
     1)))
model.add(MaxPooling2D(pool_size=(2, 2)))
model.add(Flatten())
model.add(Dense(128, activation='relu'))
model.add(Dense(10, activation='softmax'))

# Compile the model
model.compile(optimizer='adam', loss='
    categorical_crossentropy', metrics=['
    accuracy'])

# Train the model
model.fit(x_train, y_train, epochs=5,
    batch_size=200, verbose=1)

# Evaluate the model
score = model.evaluate(x_test, y_test,
    verbose=0)
print(f'Accuracy on test data: {score
    [1]*100}%')
```

2. NLP (Natural Language Processing)

NLP is a field of deep learning that focuses on enabling machines to understand and process human language. Common tasks in NLP include sentiment analysis, translation, and text summarization.

Python Example: NLP (Sentiment Analysis with LSTM)

Here is an example of a Long Short-Term Memory (LSTM) network for sentiment analysis of text data:

```
from tensorflow.keras.models import
    Sequential
from tensorflow.keras.layers import LSTM,
    Dense, Embedding, SpatialDropout1D
from tensorflow.keras.preprocessing.
    sequence import pad_sequences
from tensorflow.keras.datasets import imdb

# Load and preprocess the IMDB dataset
(X_train, y_train), (X_test, y_test) = imdb
    .load_data(num_words=5000)
X_train = pad_sequences(X_train, maxlen
    =500)
X_test = pad_sequences(X_test, maxlen=500)

# Build the LSTM model
model = Sequential()
model.add(Embedding(5000, 128, input_length
    =500))
model.add(SpatialDropout1D(0.2))
model.add(LSTM(100, dropout=0.2,
    recurrent_dropout=0.2))
model.add(Dense(1, activation='sigmoid'))

# Compile the model
model.compile(loss='binary_crossentropy',
    optimizer='adam', metrics=['accuracy'])

# Train the model
model.fit(X_train, y_train, epochs=5,
    batch_size=64, verbose=1)

# Evaluate the model
score = model.evaluate(X_test, y_test,
    verbose=0)
print(f'Accuracy on test data: {score
    [1]*100}%')
```

1.3.4. 3.Image and Object Recognition

Image and object recognition is a computer vision task that uses deep learning techniques, particularly Convolutional Neural Networks (CNNs), to identify and classify objects in images. This task is widely used in applications like autonomous driving, facial recognition, and medical image analysis.

Python Example: Image Classification with CNN

Here is an example of a Convolutional Neural Network (CNN) used for image classification, specifically classifying images from the CIFAR-10 dataset:

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

```python
from tensorflow.keras.datasets import cifar10
from tensorflow.keras.utils import to_categorical

# Load the CIFAR-10 dataset
(x_train, y_train), (x_test, y_test) = cifar10.load_data()

# Normalize the data
x_train = x_train.astype('float32') / 255
x_test = x_test.astype('float32') / 255

# One-hot encode the labels
y_train = to_categorical(y_train, 10)
y_test = to_categorical(y_test, 10)

# Build the CNN model
model = Sequential()
model.add(Conv2D(32, kernel_size=(3, 3), activation='relu', input_shape=(32, 32, 3)))
model.add(MaxPooling2D(pool_size=(2, 2)))
model.add(Conv2D(64, kernel_size=(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'))

# Compile the model
model.compile(optimizer='adam', loss='categorical_crossentropy', metrics=['accuracy'])

# Train the model
model.fit(x_train, y_train, epochs=10, batch_size=64, verbose=1)

# Evaluate the model
score = model.evaluate(x_test, y_test, verbose=0)
print(f'Accuracy on test data: {score[1]*100}%')
```

This code demonstrates how to implement a CNN for image classification using the CIFAR-10 dataset, which contains images of 10 different classes (e.g., airplanes, cars, cats).

4. Prediction

Prediction refers to the use of trained models to make forecasts or decisions based on new, unseen data. In machine learning, prediction can involve regression (for continuous values) or classification (for discrete categories).

Python Example: House Price Prediction using Linear Regression

Here is an example of using linear regression to predict house prices based on features such as size and number of bedrooms:

```python
import numpy as np
from sklearn.linear_model import LinearRegression
from sklearn.model_selection import train_test_split
import pandas as pd

# Sample data (house size, number of bedrooms vs price)
data = {
    'Size': [1200, 1500, 1800, 2000, 2200, 2400, 2600, 2800, 3000, 3200],
    'Bedrooms': [2, 3, 3, 4, 4, 4, 5, 5, 6, 6],
    'Price': [150000, 200000, 250000, 300000, 350000, 400000, 450000, 500000, 550000, 600000]
}
df = pd.DataFrame(data)

# Define predictor (X) and response (y)
X = df[['Size', 'Bedrooms']]
y = df['Price']

# Split the data into training and testing sets
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)

# Train the Linear Regression model
model = LinearRegression()
model.fit(X_train, y_train)

# Make predictions on the test set
predictions = model.predict(X_test)

# Evaluate the model
```

```python
print(f'Predicted Prices: {predictions}')
print(f'Actual Prices: {y_test.values}')
```

In this code, a simple linear regression model predicts house prices based on the features ”Size” and ”Bedrooms.”

5. Deep Learning Layers

Deep learning models, particularly neural networks, are composed of multiple layers that perform different tasks. The primary types of layers in a neural network include:

- **Input Layer:** The first layer that receives the input data.
- **Hidden Layers:** Layers between the input and output layers where computations are performed.
- **Output Layer:** The final layer that produces the predictions or classifications.
- **Convolutional Layers:** Used in CNNs to apply filters and detect features in images.
- **Pooling Layers:** Reduce the spatial dimensions of the input (used in CNNs).
- **Fully Connected Layers:** Layers where each neuron is connected to all neurons in the previous layer.
- **Recurrent Layers:** Used in Recurrent Neural Networks (RNNs) to process sequential data.

Python Example: Neural Network with Different Layers

Here is an example showing how to create a deep neural network with multiple layers using Keras:

```python
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Dense, Dropout
import numpy as np
from sklearn.model_selection import train_test_split

# Generate random data for binary classification
X = np.random.rand(1000, 20)
y = np.random.randint(0, 2, size=(1000, 1))

# Split the data into training and testing sets
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)

# Build a deep neural network model
model = Sequential()
model.add(Dense(64, activation='relu', input_dim=20))  # Input layer
model.add(Dropout(0.5))  # Regularization to prevent overfitting
model.add(Dense(128, activation='relu'))  # Hidden layer
model.add(Dense(64, activation='relu'))  # Hidden layer
model.add(Dense(1, activation='sigmoid'))  # Output layer

# Compile the model
model.compile(optimizer='adam', loss='binary_crossentropy', metrics=['accuracy'])

# Train the model
model.fit(X_train, y_train, epochs=10, batch_size=32, verbose=1)

# Evaluate the model
score = model.evaluate(X_test, y_test, verbose=0)
print(f'Accuracy on test data: {score[1]*100}%')
```

This code creates a deep neural network with an input layer, hidden layers, dropout layers for regularization, and an output layer for binary classification.

6. TensorFlow

TensorFlow is a popular open-source machine learning framework developed by Google. It is widely used for deep learning tasks and supports both CPU and GPU computations. TensorFlow allows you to build and train machine learning models efficiently using high-level APIs like Keras.

Python Example: Simple Neural Network with TensorFlow

Here is an example of how to build a simple neural network for classification using TensorFlow's Keras API:

```python
import tensorflow as tf
from tensorflow.keras.models import
    Sequential
from tensorflow.keras.layers import Dense
from tensorflow.keras.datasets import mnist
from tensorflow.keras.utils import
    to_categorical

# Load dataset (MNIST)
(x_train, y_train), (x_test, y_test) =
    mnist.load_data()

# Reshape and normalize the data
x_train = x_train.reshape(x_train.shape[0],
     28, 28, 1).astype('float32') / 255
x_test = x_test.reshape(x_test.shape[0],
    28, 28, 1).astype('float32') / 255

# One-hot encode the labels
y_train = to_categorical(y_train, 10)
y_test = to_categorical(y_test, 10)

# Build the model
model = Sequential()
model.add(Dense(128, activation='relu',
    input_shape=(28, 28, 1)))
model.add(Dense(10, activation='softmax'))

# Compile the model
model.compile(optimizer='adam', loss='
    categorical_crossentropy', metrics=['
    accuracy'])

# Train the model
model.fit(x_train, y_train, epochs=5,
    batch_size=200, verbose=1)

# Evaluate the model
score = model.evaluate(x_test, y_test,
    verbose=0)
print(f'Accuracy on test data: {score
    [1]*100}%')
```

In this code, we use TensorFlow's Keras API to implement a simple neural network for classifying MNIST digits.

[1]RWANDA POLYTECHNIC - HUYE COLLEGE [2]Department of Information and Communication Technology [3]Assistant Lecturer. Correspondence to: Ntambara Etienne <ntambaraienne94@gmail.com>.

References

Bostrom, N. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, 2014.

Brown, A. Ai in marketing: Personalization and beyond. *Marketing Tech*, 2020.

Brown, R. K-nearest neighbors for classification. *Journal of Machine Learning*, 15:45–58, 2021.

Brynjolfsson, E. and Rock, D. Artificial intelligence and the modern economy. *Annual Review of Economics*, 10: 643–665, 2018.

Davis, E. *Support Vector Machines for Classification*. Wiley, 2016.

Doe, J. Logistic regression in classification problems. *Data Science Review*, 5:20–30, 2019.

DoNotPay. Donotpay: The ai lawyer. *Tech Innovations*, 2022.

Ferrucci, D. Introduction to ”this is watson”. *IBM Journal of Research and Development*, 56(3.4):1–15, 2012. doi: 10.1147/JRD.2012.2184356.

Furst, J. Ai and innovation in robotics. *International Journal of Robotics Research*, 37(4):275–290, 2018.

Goertzel, B. and Pennachin, C. *Artificial General Intelligence*. Springer, 2014.

Goodfellow, I. *Deep Learning: Methods and Applications*. MIT Press, 2016.

Goodfellow, I., Bengio, Y., and Courville, A. Deep learning. *MIT Press*, 2016.

Hammond, J. H. Deep blue's grand challenge. *IEEE Spectrum*, 34(6):38–43, 1997. doi: 10.1109/6.591666.

Hirschberg, J. and Manning, C. D. Advances in natural language processing. *Science*, 349(6245):261–266, 2015.

IBM. Ibm watson health: Ai in diagnostics. *Healthcare AI*, 2021.

Intelligence, R. Ross intelligence: Ai for legal research. *Legal Tech Review*, 2021.

Johnson, A. *Decision Trees and Their Applications*. Springer, 2017a.

Johnson, E. Ai in personalized medicine. *Medical AI Review*, 2022.

Johnson, M. Introduction to artificial neural networks. *Artificial Intelligence Journal*, 22:10–20, 2017b.

Koch, S. and Marschollek, M. Ai for health: Opportunities and challenges. *Journal of Medical Systems*, 42(5):89, 2018.

Lee, D. *Random Forests in Machine Learning*. Oxford University Press, 2020.

Lee, J. Ai chatbots in customer service. *Customer Experience Review*, 2021.

Marcus, G. Rebooting ai: Fixing artificial intelligence. *Artificial Intelligence Journal*, 2020.

McCarthy, J. What is artificial intelligence? *Stanford University*, 2007. URL http://jmc.stanford.edu/artificial-intelligence/what-is-ai.html.

McCarthy, J., Minsky, M., Rochester, N., and Shannon, C. A proposal for the dartmouth summer research project on artificial intelligence. In *Dartmouth Conference on Artificial Intelligence*, 1956. URL https://www.aaai.org/ojs/index.php/aimagazine/article/view/1904.

Music, A. Amper music: Ai for music production. *Creative AI*, 2021.

of Defense, D. Project maven: Ai in military operations. *Defense Tech Journal*, 2020.

OpenAI. Jukedeck: Ai music composition. *Music Tech*, 2020.

Ricci, F. and Rokach, L. Recommender systems handbook. *Springer*, 2011.

Rosenblatt, F. *The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain*, volume 65. Psychological Review, 1958. doi: 10.1037/h0042519.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986. doi: 10.1038/323533a0.

Russell, S. and Norvig, P. *Artificial Intelligence: A Modern Approach*. Pearson, 3rd edition, 2010.

Russell, S. and Norvig, P. *Artificial Intelligence: A Modern Approach*. Pearson, 2020.

Schmidhuber, J. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015. doi: 10.1016/j.neunet.2014.09.003.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Driessche, G. V. D., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–489, 2016. doi: 10.1038/nature16961.

Smith, J. An introduction to linear regression. *Machine Learning Journal*, 10:1–12, 2018.

Smith, J. Ai in supply chain management. *Business AI Journal*, 2022.

Topol, E. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books, 2019.

Turing, A. M. Computing machinery and intelligence. *Mind*, LIX(236):433–460, 1950. doi: 10.1093/mind/LIX.236.433.

van der Aalst, W. and Song, M. Robotic process automation and artificial intelligence. *Business Information Systems Engineering*, 62(4):287–293, 2020.

Williams, D. Ai-assisted robotic surgery. *Surgical Innovations*, 2021.

Yann LeCun, Y. B. and Hinton, G. Deep learning. *Nature*, 521:436–444, 2015.