

Here's an implementation of a basic machine learning approach using reinforcement learning to solve a labyrinth (or maze). We'll use **Q-learning**, which is a popular reinforcement learning algorithm. This algorithm is particularly suited for problems like maze-solving, where an agent learns to navigate an environment by receiving rewards and penalties.

Q-learning Overview:

- **Agent:** The entity trying to solve the labyrinth.
- **Environment:** The labyrinth (represented as a grid).
- **State:** Each position in the labyrinth.
- **Action:** Moving up, down, left, or right.
- **Reward:** A numerical value received for performing an action (e.g., reaching the goal gives a positive reward, hitting a wall gives a negative reward).

Here's a step-by-step implementation of Q-learning to solve a labyrinth:

Step 1: Import Libraries

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

Step 2: Create the Labyrinth Environment

We will represent the labyrinth as a grid, with:

- `0` indicating a free space,
- `1` indicating a wall,
- `2` indicating the starting point,
- `3` indicating the goal.

```
```python
Labyrinth definition: 0=free space, 1=wall, 2=start, 3=goal
labyrinth = np.array([
 [1, 1, 1, 1, 1, 1],
 [1, 2, 0, 0, 3, 1],
 [1, 0, 1, 0, 1, 1],
 [1, 0, 1, 0, 0, 1],
 [1, 0, 0, 1, 0, 1],
 [1, 1, 1, 1, 1, 1]
])

start_state = (1, 1)
goal_state = (1, 4)
```
```

Step 3: Define Hyperparameters and Initialize Q-table

```
```python
```

```

Parameters
alpha = 0.1 # Learning rate
gamma = 0.9 # Discount factor
epsilon = 1.0 # Exploration rate
epsilon_decay = 0.995 # Decay rate for exploration
min_epsilon = 0.01 # Minimum value for epsilon
episodes = 5000 # Number of episodes for training
max_steps = 100 # Max steps per episode

Actions: 0 = Up, 1 = Right, 2 = Down, 3 = Left
actions = ['up', 'right', 'down', 'left']

Q-table: a table where rows are states and columns are actions
q_table = np.zeros((labyrinth.shape[0], labyrinth.shape[1], len(actions)))

Reward function (labyrinth with penalties and rewards)
def get_reward(state):
 if state == goal_state:
 return 100 # Large reward for reaching the goal
 elif labyrinth[state] == 1:
 return -100 # Large penalty for hitting a wall
 else:
 return -1 # Small penalty for each step taken
...

Step 4: Define Functions for Agent Movement

```python
# Determine next state based on action
def next_state(state, action):
    x, y = state
    if action == 0: # Up
        return (x - 1, y)
    elif action == 1: # Right
        return (x, y + 1)
    elif action == 2: # Down
        return (x + 1, y)
    elif action == 3: # Left
        return (x, y - 1)

# Check if the state is valid
def is_valid_state(state):
    x, y = state
    if x < 0 or x >= labyrinth.shape[0] or y < 0 or y >= labyrinth.shape[1]:
        return False # Out of bounds
    if labyrinth[state] == 1:
        return False # Hitting a wall
    return True
...

```

Step 5: Q-learning Algorithm Implementation

```
```python
Training the Q-learning agent
for episode in range(episodes):
 state = start_state
 total_reward = 0
 for step in range(max_steps):
 # Choose action: Exploration vs Exploitation
 if random.uniform(0, 1) < epsilon:
 action = random.choice(range(len(actions))) # Exploration
 else:
 action = np.argmax(q_table[state]) # Exploitation

 # Get next state and reward
 next_s = next_state(state, action)

 if not is_valid_state(next_s):
 next_s = state # Stay in place if next state is invalid

 reward = get_reward(next_s)

 # Update Q-value (Bellman equation)
 best_next_action = np.argmax(q_table[next_s])
 q_table[state][action] = q_table[state][action] + alpha * (
 reward + gamma * q_table[next_s][best_next_action] - q_table[state][action])

 state = next_s
 total_reward += reward

 # If goal is reached, end episode
 if state == goal_state:
 break

 # Decay exploration rate
 epsilon = max(min_epsilon, epsilon * epsilon_decay)

Print final Q-table
print("Training complete!")
```
```

Step 6: Test the Agent

After training, we can test the agent to see if it can successfully solve the labyrinth.

```
```python
Test the trained agent
state = start_state
```

```

steps = [state]

for _ in range(max_steps):
 action = np.argmax(q_table[state]) # Choose best action from Q-table
 state = next_state(state, action)

 if not is_valid_state(state):
 break # Stop if agent tries to move into a wall

 steps.append(state)

 if state == goal_state:
 print("Goal reached!")
 break

Visualize the path taken by the agent
print("Steps taken by agent:", steps)

Plot the labyrinth and the path
labyrinth_copy = labyrinth.copy()
for step in steps:
 labyrinth_copy[step] = 4 # Mark the path

plt.imshow(labyrinth_copy, cmap='coolwarm')
plt.title('Path taken by Q-Learning Agent')
plt.show()
`

```

### ### Explanation:

1. **Q-learning Setup:**
  - The Q-table stores the expected rewards for every state-action pair. The agent updates the Q-values over time to find the optimal path.
2. **Agent Training:**
  - The agent starts at the beginning of the labyrinth and explores the environment while balancing between exploration (random actions) and exploitation (choosing actions based on learned Q-values).
3. **Movement and Rewards:**
  - The agent receives a small penalty for each step to encourage it to find the shortest path. A large positive reward is given for reaching the goal, while a large negative reward is assigned for hitting walls.
4. **Visualization:**
  - After training, the labyrinth and the path taken by the agent are visualized using `matplotlib`.

This implementation demonstrates how reinforcement learning (specifically Q-learning) can solve a labyrinth by training an agent to navigate towards the goal effectively.