

The Mouse and the Cheese Game

A hands-on Reinforcement Learning experiment!

Goal of the Game

Help the mouse learn — all by itself — the best path to reach the cheese.

Each time the mouse succeeds, its “intelligence” improves... and the game becomes even more interesting!

Materials Needed (per group)

- A **3x3 grid board** (can be drawn or printed)
 - Two pieces:
 -  **The Mouse**
 -  **The Cheese**
 - **9 cups** (one for each square of the board)
 - **Direction cards:**
 - Up 
 - Down 
 - Left 
 - Right 
-

Setup

1. Place the board on the table.
 2. Put the mouse in one corner and the cheese in the opposite corner.
 3. Arrange the nine cups beside the board in three rows. Each cup is directly associated with one square of the grid.
 4. In each cup (one per square), place **one direction card for each possible move** from that square.
 - Example: if you're in a corner, there will only be two valid directions.
-

How to Play

1. Place the mouse in its starting position.
2. Take the cup corresponding to that square.
3. Draw one card at random and move the mouse in that direction.
4. Place the drawn card on the square from which the mouse moved.
 - If the mouse returns to a square it has already visited, return to the cups all the cards that form the closed loop leading the mouse back to that square.
5. Continue repeating:
 - Keep moving the mouse until it reaches the cheese.
 - Each turn represents a **decision made by the mouse** based on its “cards.”
6. The turn ends when the mouse reaches the cheese. 

💡 Reinforcement (Learning)

When the mouse manages to reach the cheese:

1. Trace back through all the squares the mouse has passed.
 2. For each cup along that path, return the card that is on the board **and add five identical cards** to the corresponding cup.
 - In other words, **reward** the decisions that led to success.
 3. Return the mouse to its starting position... and begin a new round!
-

🧠 What's Happening Behind the Scenes

- Each **square** represents a **state** of the environment.
 - Each **card** represents a **possible action**.
 - Every time the mouse reaches the cheese, it receives a **reward**.
 - Over time, the cups will fill with “smart” cards — the most successful directions.
 - Thus, the mouse **learns the optimal route** without being explicitly taught.
-

🏁 End of the Game

The game can be repeated as many times as you wish.

The goal is not to win, but to **observe how the system learns**:

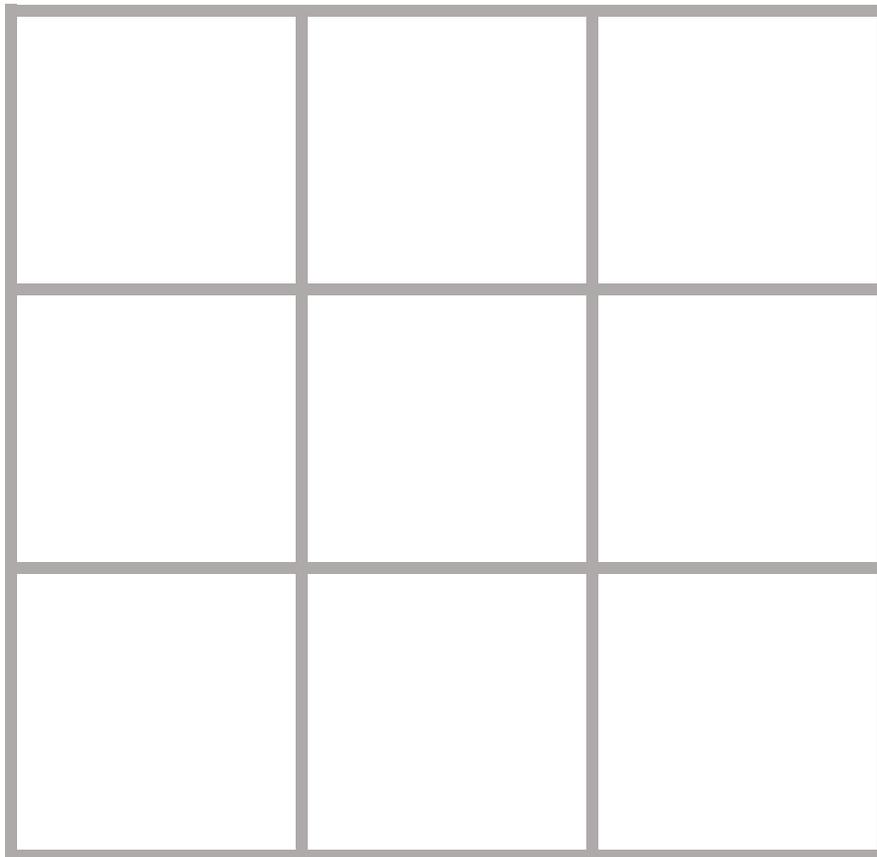
at first, the mouse wanders aimlessly, but soon it finds the path to the cheese with precision.

If you restart the game from scratch, you'll notice that each time it **learns a slightly different route**.

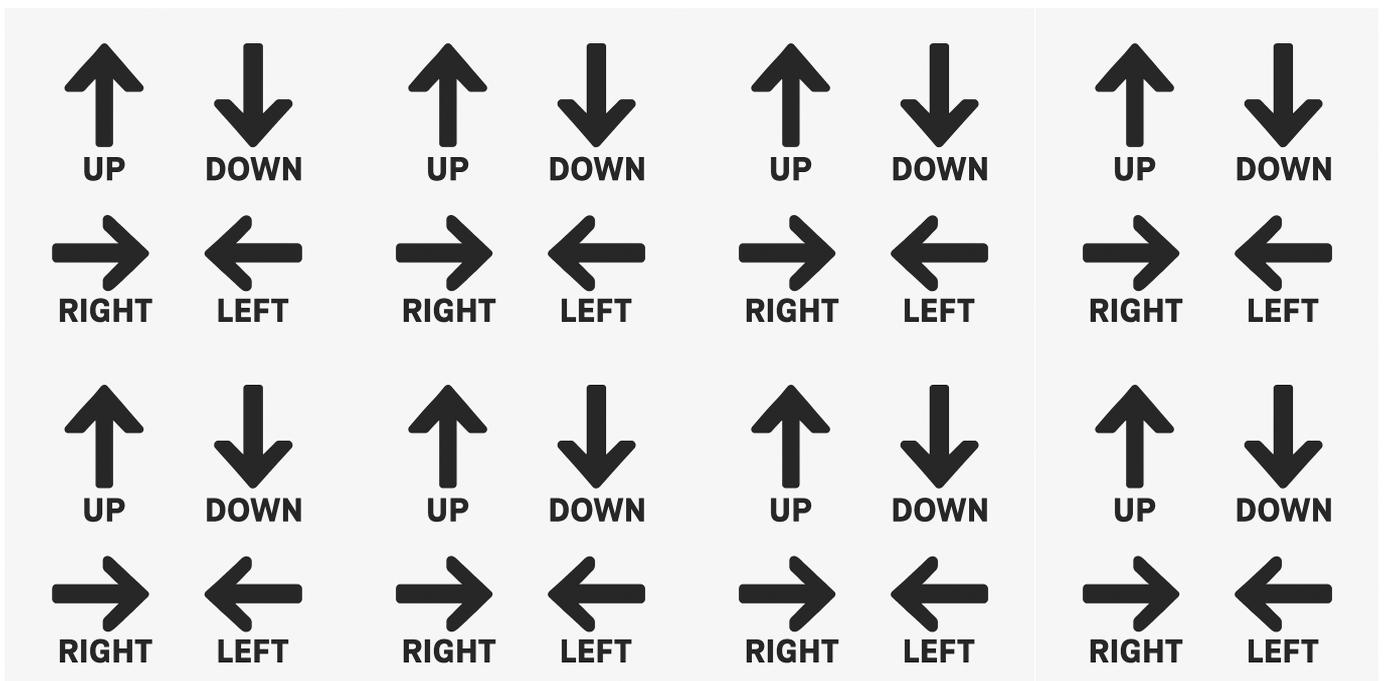


Mouse and Cheese Tokens

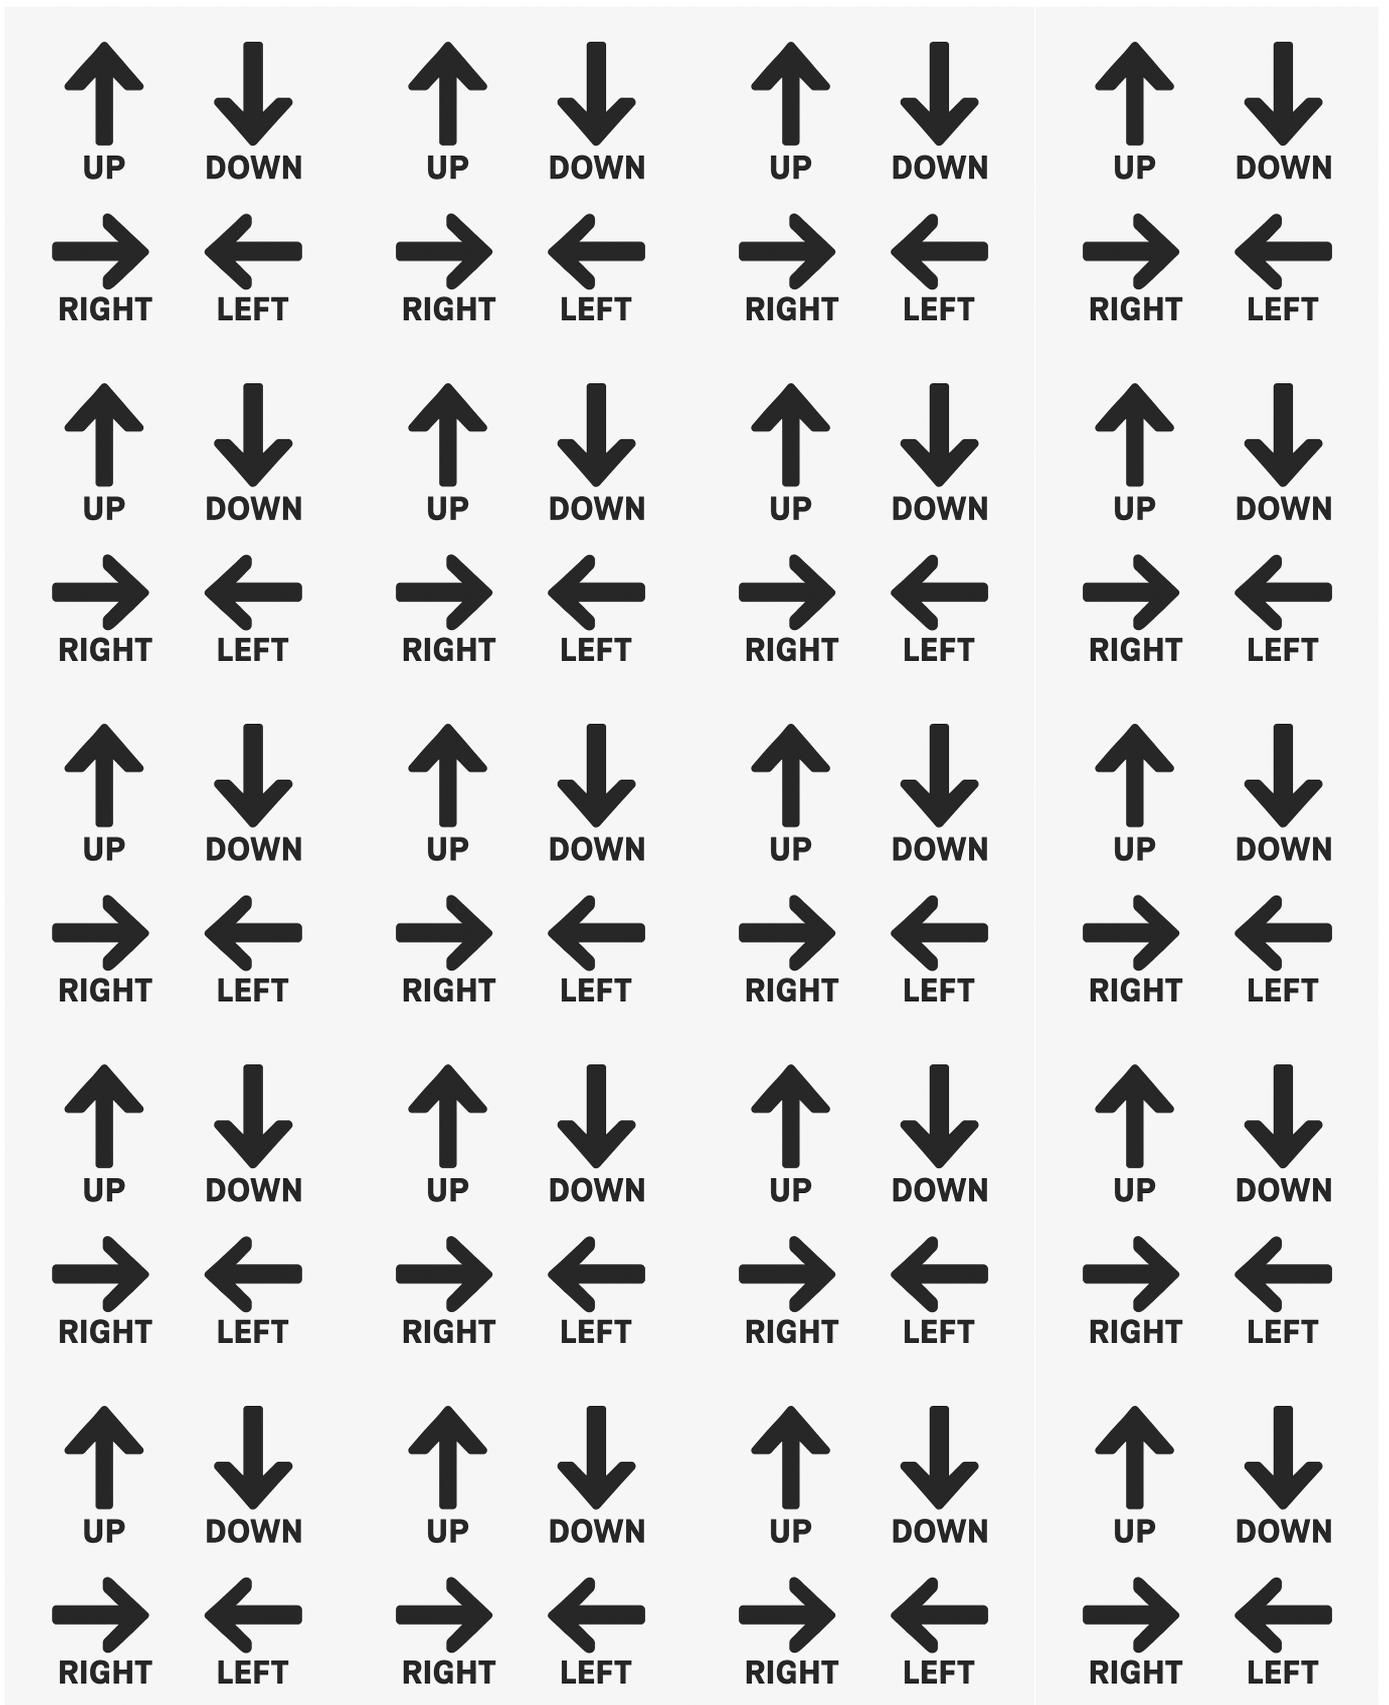
[Artificial brains - How machines learn to think](#)



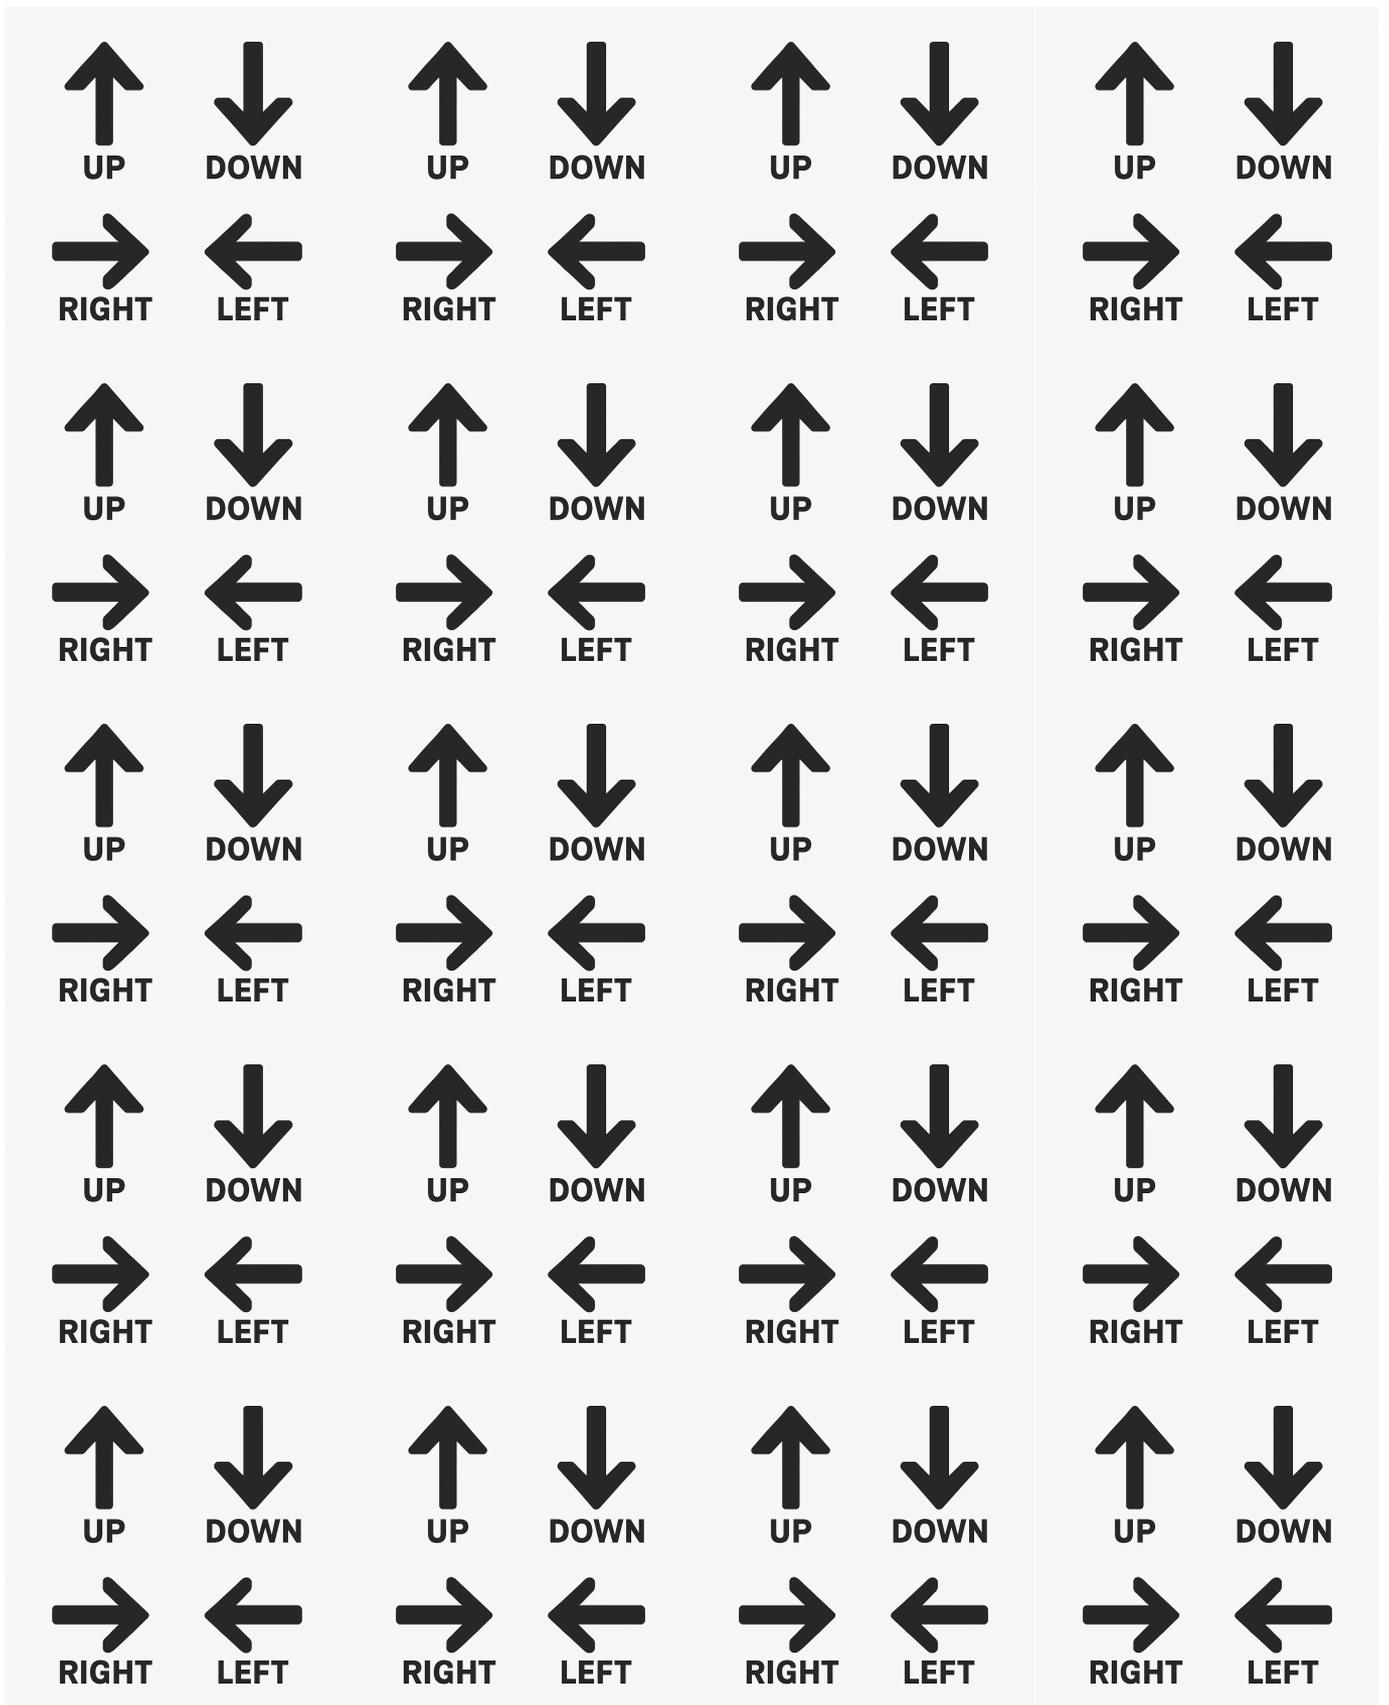
3x3 Board



[Artificial brains - How machines learn to think](#)



[Artificial brains - How machines learn to think](#)



[Artificial brains - How machines learn to think](#)

