

AI in Education

A Practical Guide for Teachers and Young People



L-Università ta' Malta
Faculty of Information &
Communication Technology

Department
of Artificial
Intelligence

AI in Education

A Practical Guide for Teachers and Young People



L-Università ta' Malta
Faculty of Information &
Communication Technology

Department
of Artificial
Intelligence

Copyright 2019 Department of AI, University of Malta Editors: Vanessa Camilleri, Alexiei Dingli, Matthew Montebello Content Developed by: Simon Agius Muscat

Illustrations by: Clyde Haber

ISBN (Print): 978-99957-1-551-9

ISBN (PDF): 978-99957-1-552-6

This work is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-sa/4.0/>.



First Printing, September 2019

FOREWORD

Artificial intelligence (AI) is reshaping the world as we have known it for the past years. It is affecting every aspect of our society from the very common use of the mobile device to more complex situations at our workplaces. It also contributes to our collective wellbeing even though we many not see it in practice.


To ensure that Malta is at the forefront of technological innovation, the Government of Malta launched the Malta. AI taskforce with the scope of creating an AI strategy for Malta, with a vision to instil an appreciation of how new technologies can be used to improve people's quality of life.

Much information about AI is abundantly available, as industries and providers target adults to promote and exhibit the advantages of AI. However, very little information is available for children and youths. This is one of the reasons why academics at the Department of AI within the Faculty of ICT (University of Malta) felt the need to embark on this book project to help promote AI concepts with children and young people. The task at hand included bringing across topics in AI that can be quite complex in a way that engages children. This was done using game-based learning (GBL) concepts, using the element of fun to motivate and engage young learners with the content. The result was this book: AI in Education which is fun to use, full of practical games and which can be easily adapted to any classroom.

The Ministry of Education has financed the publication of this book to ensure that teachers and students have access to this essential resource that can be used across multiple subjects and areas.

Our vision is that children are our future leaders and as such we believe that they are entitled to the best formation possible. We do not just want them to successfully thrive in today's world. We want to give them the right tools so that they can shape tomorrow's world and ultimately the future for all of us.

Through this project and many others, the Ministry of Education is gearing up to shift from theory and principles to action. We are determined that through this first small step we help our children start understanding the promise of AI to harness it, and share its benefits. Through our vision, we want to ensure that no one is left behind, now and for generations to come.



• ARTIFICIAL INTELLIGENCE •

PREFACE

Emerging new technologies are merging the digital, physical and biological worlds as we challenge ourselves with what it means to be human.

Artificial Intelligence (AI) can be seen around us and in use in various devices and applications, from supercomputers, drones, and virtual assistants to 3D printing and wearable sensors. But the real questions are:

§ What does Intelligence in AI stand for?

§ How will AI affect our lives and our jobs?

§ How can AI improve the quality of our lives?

This practical guide for teachers will explore some of these questions, together with concepts lying at the roots of Artificial Intelligence. The activities with different challenge factors will introduce some of the concepts used for AI using simple examples from daily life.

PUBLICATION STRUCTURE

This publication is structured as follows:

Part A includes introductory material and instructions for use, as well as details for each activity or discussion, including briefs for teachers running sessions. Although there is no particular or specific order in which the activities must be carried out in class, in the publication these are sorted according to the challenge factor, from low (basic knowledge) to high (more advanced knowledge).

Part B includes the worksheets related to specific activities. Worksheets can be printed out and distributed to the learners during class.

Each activity has been designed to purposefully require as little setup as possible – typically limited to the printing of handouts or easy-to-acquire resources (such as a pen and paper or a dice). While several of the activities contain more practical/visual material (such as links to videos, blog posts, or online simulators), these are not essential to the running of the session and should only serve to support the written text rather than replace it.

INSTRUCTIONS FOR USE

Each activity can vary in duration. While the activities with a low challenge factor can last between 30-40 minutes, others with a medium-high challenge factor can take longer.

It is recommended that for a 2-hour session, each activity is followed up or supplemented either with an extension of the proposed games or additional activities such as visual coding, etc. However, we believe that each activity

should contain a strong element of peer-led discussion that will help learners reflect on what is going on as part of the activity.

Each activity section in Part A contains the following:

1. A short brief on what the activity/discussion covers in terms of area of AI
 2. A challenge factor, meant to indicate for which level/age of learners it is ideal for
 3. A definition of the learning outcomes for the activity
 4. An introduction that can be read out or paraphrased to set the stage
 5. Instructions for the teacher running the session (may include questions to field if the activity is a discussion)
 6. A conclusion that can be read or paraphrased to wrap up the session
 7. Follow up questions may be used to further stimulate discussion between learners about the subject
 8. Some activities may include an activity that can be completed online
- Each activity is directly linked to the worksheets found in Part B. Worksheets can be printed and distributed to students during class.

Each activity is directly linked to the worksheets found in Part B. Worksheets can be printed and distributed to students during class.

AREAS OF AI

The activities proposed in this guidebook make reference to the following specific areas of Artificial Intelligence. While this is not a comprehensive set of subject areas, it is meant to give an overview of the subject. The AI subject areas that are discussed in this publication include:

1. Practical Applications
2. Reinforcement Learning
3. Search Algorithms
4. Neural Networks
5. AI Ethics and Philosophy




Table of contents


- [PREFACE](#)
- [PUBLICATION STRUCTURE](#)
- [INSTRUCTIONS FOR USE](#)
- [AREAS OF AI](#)


1. Part A

1. Activity 1
2. Activity 2
3. Activity 3
4. Activity 4
5. Activity 5
6. Activity 6
7. Activity 7
8. Activity 8
9. Activity 9
10. Activity 10

2. Part B

- 1. Worksheet 1**
- 2. Worksheet 2**
- 3. Worksheet 3**
- 4. Worksheet 4**
- 5. Worksheet 5**
- 6. Worksheet 6**
- 7. Worksheet 7**
- 8. Worksheet 8**
- 9. Worksheet 9**
- 10. Worksheet 10**





ACTIVITY 1

SHOPPING WITH AI



INTRODUCTION

This activity aims to introduce the subject of AI, especially for learners who are not familiar with the subject.

Artificial Intelligence is, in simplified terms, machinery capable of making intelligent decisions. It may be difficult to recognize that up until only a few years ago, most machinery was making apparently intelligent decisions that did not fully exploit the potential of AI. For example, we know that when a person looks at a product on an online marketplace or shop, places that product in the cart, and does not purchase it, there are instances when a notification would be sent over email a few days later offering that person a discount on that same product.

TOPIC BRIEF FOR TEACHERS

In the example given of an online shopping cart, although to some people this might seem AI-driven, in reality, this technology is mostly hard-coded, where a programmer would develop explicit instructions that describe what the industry refers to as business logic. Continuing the example above, the behavior could be described as follows:

1. If cart not empty for X days
2. If stock for product in cart not low
3. Create discount code
4. Send email with discount code

AI, on the other hand, involves very little, if any, hard-coding. The program knows how to learn and will automatically execute instructions to learn. Given enough data, processing power, time, and the correct learning conditions, AI programs are meant to solve problems by changing their behavior over time, constantly attempting to adapt in a way that improves their scope and provides the right solution in the most effective way. To continue the example, an AI-powered discount generator might learn from a person's purchase history to understand what kind of products that person tends to purchase, and then suggest discounts for different products that they would be more likely to purchase in the future. Such an AI-driven discount generator would continue to track the person's purchase history over time to adapt to the individual's changing needs and preferences.

CHALLENGE FACTOR



This activity requires basic mathematical addition and a basic knowledge of negative numbers.

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Explain in simple terms what Artificial Intelligence (AI) consists of.
- Give a practical example of AI.
- Discuss in more detail how and why researchers came up with the concept of AI.
- Describe how AI could find its use in a number of practical real-world applications.
- Research and list practical applications of AI.
- Discuss the advantages and disadvantages of using and applying AI for the purposes of human use.

CLASS INTRODUCTION

Artificial Intelligence is a phrase that is being used more and more often in everyday conversation and is a part of many technology-driven products we use in our daily lives, such as social media, games, and even cars!

In this discussion, we will go over a brief history of AI, see what AI actually is, see if we can identify any AI in our day-to-day lives, and conclude by figuring out how AI affects our use of technology and how it may shape the future.

LEARNING RESOURCES

- 1x Worksheet titled ‘Shopping with AI’
- YouTube Trailer: [The Imitation Game](#)
- Introduction to AI: [What is Artificial Intelligence?](#)
- AI Quiz: [What is AI capable of?](#)

DISCUSSION

1. What does AI stand for?
2. What do you understand by Artificial Intelligence?
3. Can you give an example of AI?
4. What led Alan Turing to come up with the concept of Artificial Intelligence?
5. What can AI be useful for?
6. How do companies use AI for their products? What examples can you think about?
7. What are the advantages of using AI over a human-operated machine?
8. Can you think of any disadvantages of a machine which is capable of doing a task typically done by humans?

TEACHER AIDS

1. What is AI?

- Artificial Intelligence

2. What do you understand by Artificial Intelligence?

- Intelligent computers
- Machines that learn
- Robots
- Computer controlled objects

3. Can you give an example of AI?

- The computer player in video games
- YouTube video recommendations
- Targeted advertising

4. What led Alan Turing to come up with the concept of Artificial Intelligence?

Alan Turing, in 1950, published a paper on the concept of machines that are able to simulate human behavior, think, and respond intelligently. During that period, Alan Turing proposed a modified version of an old Victorian competition called the Imitation Game. In this proposal, Alan Turing discussed whether it would be possible for an interrogator to distinguish between a human and a computer when asked a question. Alan Turing's proposal was that if the interrogator cannot distinguish between the answers given by the computer and the human, then the computer would be considered to be thinking.

5. What can AI be useful for?

- Making predictions based on big data (huge volumes of data) that humans are unable to compute for example in medicine treatments, sports competition, or financial markets predictions
- Making programs more efficient and allocating resources more efficiently
- Enabling computers to do tasks that humans can do, for example, recognizing contents of images.

6. How do companies use AI for their products? What examples can you think about?

- Amazon Echo (Alexa)
- Facebook
- Amazon
- Google
- Apple (Siri)

7. What are the advantages of using AI over a human-operated machine?

- Increased precision
- Faster than humans on specific tasks
- Less prone to human error

8. Can you think of any disadvantages of a machine which is capable of doing a task typically done by humans?

- Ethical Issues
- Issues with safety and security


- Excessive automation



CONCLUSION

Artificial Intelligence can be exploited as a very useful tool. Like any other tool, using it has both its advantages and its disadvantages. There is a large amount of material related to the practical applications of AI that can be found online. For a more in-depth knowledge of AI, choosing Computer Studies at 'O' and 'A' levels is the first step towards a Diploma or Degree in AI that will enable you to work in the field and be a part of the AI revolution.

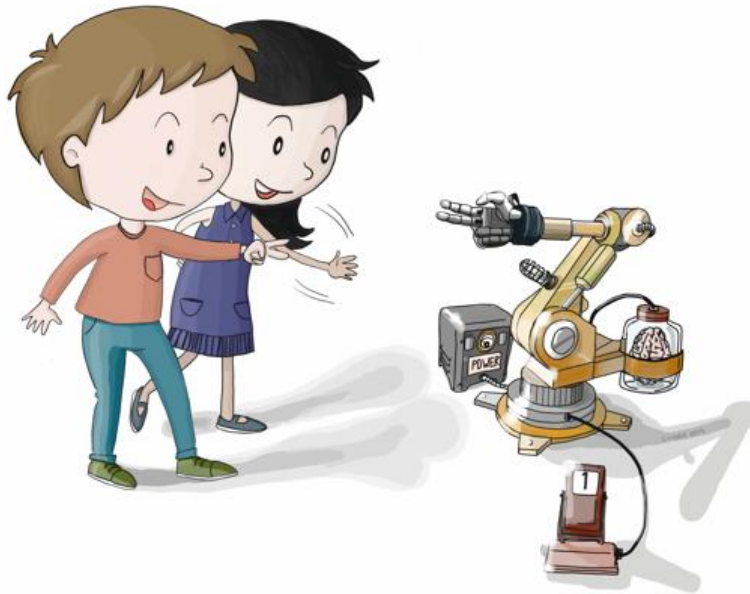
Finally, [this video](#) shows some of the more practical applications of AI. Can you find more examples?



ACTIVITY 2

REINFORCEMENT LEARNING (1):

ROCK, PAPER, AI!



INTRODUCTION

This activity aims to introduce the concept of Reinforcement Learning in AI, where a machine does not learn by being told what it should do, but learns instead from its mistakes as it operates within its environment. The game of Rock, Paper, Scissors traditionally played by children using their hands is used as an example to illustrate the concept of Reinforcement Learning.

TOPIC BRIEF FOR TEACHERS

Not all AI approaches involve a simulation of the structure and operation of the human brain. There are other approaches that are inspired by more generic psychological phenomena observable in both humans and animals.

One such method is Reinforcement Learning.

Similar to scolding a child for misbehaving, or training a puppy to avoid bad behavior through negative reinforcement, or alternatively rewarding them

with a treat for good behavior through positive reinforcement, this type of reinforcement-based learning AI would keep track of the results based on set actions. If one action has a negative outcome, then the likelihood of that action being repeated in the future is lessened. The theory holds that, with enough actions and learning, the AI will eventually learn the pattern of correct actions to consistently achieve positive outcomes.

CHALLENGE FACTOR



Through this activity, learners can gather in groups, and teachers or learning assistants are encouraged to participate and engage in the game of rock, paper, scissors. The learners need knowledge of mathematical addition and possibly a basic understanding of negative numbers

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Explain Reinforcement Learning in AI in simple terms.
- Demonstrate how to calculate adjustments that a simple Reinforcement Learning AI makes to itself.

- Explain how Reinforcement Learning allows the AI to learn from its mistakes.
- Put forward hypotheses on the difference between an AI versus a Human and an AI versus another AI.

CLASS INTRODUCTION

Machines that learn find various uses in our everyday lives. Artificial Intelligence can be trained or taught using different methods. For example, AI can be shown the correct answer repeatedly. In this activity, we will see how an AI can learn from both its successes and its mistakes instead.

Let us imagine a toddler who is learning how to walk. Does anyone have a younger brother or sister? Do you remember how they learned to walk? Can you describe it to us? If we were to try and teach a robot to walk or a car to drive, if at first the robot takes a big step and falls forward, it might be taught to adjust the size of the step so that next time it might not fall. If it falls again and again, it keeps adjusting its step until it doesn't fall. This is what is called reinforcement learning – the AI would learn from its mistakes so that it can achieve its goals. In this lesson, we will be trying to understand how AI can be taught to play Rock, Paper, Scissors by learning from its mistakes and by reinforcing its wins and successes!

LEARNING RESOURCES

- 1x Worksheet titled 'Rock, Paper...AI'
- Reinforcement Learning: [Machine Learning and Data Science](#)
- An Introduction to Reinforcement Learning: [AI Learning How to Drive](#)
- Reinforcement Learning: [Why do we need Reinforcement Learning?](#)
- A Case in Reinforcement Learning: [Reinforcement Learning by Example](#)


DISCUSSION

1. Can anyone explain the game of rock, paper, scissors?

2. Do you think that eventually the AI would learn how to play against the whole class? How and why?
3. Would we be able to win more games if we played completely randomly?
4. What do you think would happen if the AI played against itself?

TEACHER AIDS

1. The teacher will initially act as the AI player (a learner can also act as a volunteer to be the AI player).
2. The AI player needs to use the Rock, Paper, AI! Worksheet (see part B).
3. One by one, the learners are invited to come up and play Rock, Paper, Scissors against the AI player.
4. The AI player follows the instructions on which option to choose on each game.
5. The number of wins, losses and draws can be tracked, either by writing on a whiteboard or written down on a paper next to the AI player.
6. This process is repeated until all the learners have played at least once against the AI player.
7. The Rock, Paper, AI! Worksheet can be then handed out whilst the teacher illustrates the concepts underlying reinforcement learning (*see the Beginner's guide to [Reinforcement Learning](#) for more details*).
8. The discussion questions will help the learners explore concepts such as whether the number of wins and draws of the AI are due to learning or are random, and their ideas about what would happen in the case where the AI plays against itself.




CONCLUSION

This activity used some math to simulate how an AI can use reinforcement learning to learn, for example, how to play a game like rock, paper, scissors. Google DeepMind has managed to do this for Go! – an ancient Chinese game.

Their [AlphaGo](https://tinyurl.com/y2vcx5cb) reinforcement learning algorithm has actually beaten the current professional player Lee Sedol (see video: <https://tinyurl.com/y2vcx5cb>) at the game of Go! The result was a very exciting match! Seeing all this, where would you see Reinforcement Learning being applied in real life situations?

Some of the major fields where Reinforcement Learning AI is applied include:


- Game Theory
- Robotics
- Computer Networking
- Vehicular Navigation
- Medicine
- Industrial Logistic.



ACTIVITY 3

REINFORCEMENT LEARNING (II):

AI OLYMPICS



INTRODUCTION

This activity is a continuation of the previous Rock, Paper... AI.

Using the same concepts of Reinforcement Learning the learners will battle out AI in a knockout competition, until only one AI wins!

TOPIC BRIEF FOR TEACHERS

Building on the previous activity, this one introduces a relatively new concept in approaches to training AI: self-play. Since the Reinforcement Learning algorithms consist of a series of instructions that are interpreted by a computer, we can copy and run different versions of the AI at the same time. In comparison to other machine learning algorithms such as supervised and unsupervised learning, reinforcement learning algorithms learn by reacting to the environment using the feedback returned as reward and punishment. By running two versions of the AI against each other, the AI is

able to learn from its wins and losses, *at the same time* and improve itself at a much faster pace.

While this activity does not explore how self-play improves the same AI, it demonstrates how AIs can play against each other, with humans acting as human CPUs for instructions that would otherwise be interpreted by a computer to create an AI.

CHALLENGE FACTOR



This activity is a continuation of the previous one. Given that learners understood the previous activity, only supervision is required in this activity.

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Follow instructions on how the AI can win, draw, or lose while playing Rock, Paper, Scissors, and note down wins while reacting to the feedback received appropriately.
- Explain well how a reinforcement learning AI applies itself in Rock, Paper, Scissors through practice.

- Discuss how AIs can compete against each other, proposing theories on potential outcomes from such a competition.

CLASS INTRODUCTION

During the last lesson, we played rock, paper, scissors. What was the game about?

Who was the AI and what did the AI do? Did the AI learn the game in the end? Did it win? With the type of learning we saw, the AI could learn from its successes and its mistakes. Do you remember what that type of learning was called? During this lesson, we are going to each play as an AI against each other!

It is AI against AI today, in the AI Olympics of Rock, Paper, Scissors! What do you think the outcome will be?

LEARNING RESOURCES

- 1x Worksheet titled 'Rock, Paper...AI'
- 1 six-sided dice for each player
- Reinforcement Learning in Action: [Two Chatbots talking to each other](#)


DISCUSSION

1. Learners are given worksheets and split into pairs. Each pair is handed a dice.
2. The learners play a number of games (e.g. ten) of rock, paper, scissors against each other. For every game, learners update their worksheet with their wins, draws and losses.
3. The winner of each pair will continue on to play against the winner of the next pair: If there is a tie, the learners keep playing until a win occurs.
4. This continues until only two learners remain.

5. Learners can compare their probability tables with the others to explain how one AI could manage to win over the other.


TEACHER AIDS

1. It is important to stress to the learners the effect of the reward system and how this works for the AI. The teacher can choose to describe the example of someone who has to do a boring task. Would that person be likely to redo the task if they are rewarded at the end? If there are two people doing the same task, and one finishes before and is given a much bigger reward, what do you think will happen the next time they are both given another task to do?
2. Reinforcement learning is different from other machine learning paradigms because:
 - It requires no supervision
 - Time is important
 - It works on the concept of delayed rewards
 - The results from the AI's own action affects its behavior for the next action



CONCLUSION


The scope of this activity is to understand how an Artificial Intelligence can execute instructions to act, make moves in the game, make mistakes, and learn from those mistakes. This is how one of the methods of Artificial Intelligence works, which is called Reinforcement Learning. It is currently a very popular method and is being applied in many fields, including more complex games such as [DOTA 2](#).



ACTIVITY 4

SEARCH ALGORITHMS (I):

THE PAPER-BASED INTELLIGENT BEING



INTRODUCTION

This activity aims to introduce the idea of intelligence as a set of instructions through which a problem is solved within a given space.

Taking the example of a game of Tic-Tac-Toe, learners will be challenged to defeat the Paper-Based Intelligent Being, which is a description of the search algorithm that either always wins or draws in the game against any opponent. It will serve as a practical introduction to the application of AI in games and in general.

TOPIC BRIEF FOR TEACHERS

If you think about the way problems can be solved, you might reach the conclusion that a problem might have one solution and one correct way of getting to the solution. However, the more complex the problem is, the greater the number of possibilities in which a solution can be achieved.

More modern AI approaches have recently become more popular than their traditional counterparts, such as search algorithms. However, it was through the more traditional methods that AI actually began making news. In 1997, the IBM Deep Blue computer defeated Gary Kasparov in a series of chess games. Deep Blue internally used search algorithms to brute force as many possible future states of the chess game as possible, evaluating each one for the likelihood of eventually winning before making a move. And it won!

Search algorithms can be easily generalized to several different fields, where one wishes to find the ideal path between two states, whether it's a game of chess (from a fresh board to checkmate), discovering the shortest route from Mosta to Birkirkara (considering geographical distance, how roads connect together, and possibly even other variables such as the amount of traffic present), or even how to bake a cake (starting from the ingredients and ending with a cake).

It does this by expanding each possible state from the current state, then re-expanding each possible state from each of those states, over and over again.

By doing this process virtually, a computer can help optimize a journey or a task without even attempting it (unlike other methods, such as reinforcement learning).

CHALLENGE FACTOR



While the game is very simple to understand and play, the instructions to always win or draw are quite complex and text-heavy. If the instructions are understood well, this can be considered even as a low challenge factor since the learners will only play against the paper but not as the paper.

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Explain the concept of a mathematically complete game is.
- Demonstrate how intelligence can result from the execution of simple instructions in the game of Tic-Tac-Toe.
- Explain some of the simple terms in Search Algorithms such as the problem space, states, and the search space.
- Understand how Search Algorithms operate as they expand the problem space to attempt to find an optimal solution.

CLASS INTRODUCTION

During the last sessions, we played rock, paper, scissors and demonstrated how an AI can learn how to play the game by using a system of rewards and punishments.

Today we will play another game – this time Tic-Tac-Toe. Would anyone like to explain how the game is played and how you can win at the game?

Now imagine if we had to discover an intelligent paper-based being? What if this intelligent paper-based being claims that it can beat us at the game of Tic-Tac-Toe / noughts and crosses every single time? Do you think that it can happen?

Today we are going to rise up to the challenge and try to defeat this being. If this piece of paper actually is intelligent, then it should be able to win, or at least never lose, right?

LEARNING RESOURCES

- 1x Worksheet titled ‘The Paper-based Intelligent Being’
- Tic-Tac-Toe Machine: [A Minecraft Example](#)
- Additional reading: [Popular Search Algorithms](#)

DISCUSSION

1. A person is chosen to be the intelligent piece of paper’s interface. As it does not have a way of interacting with the world, we need someone to interpret the instructions for us (either the teacher or a learner).
2. The game grid is drawn on the whiteboard / paper.
3. The game is played between the paper and the learners for a number of games (e.g. five).
4. The paper and the learners alternate at playing first for the different games.

5. The learners can choose their moves either by collective agreement, or by taking it in turns.
6. The number of wins, draws and losses of both the paper and the learners should be tracked.
7. How do you think the intelligent being learned to play Tic-Tac-Toe so well ?
8. Do you think that there are AIs which play more complicated games?

TEACHER AIDS


1. The links below give some additional information that can help a better understanding of the ways search algorithms can be applied for an AI to win at games traditionally played by humans.
 - Solved games, a category of games to which the best instructions to win are already known: <https://tinyurl.com/yyo2ed3g>.
 - [Deep Blue](#), an IBM Computer capable of defeating the then world champion in chess, Gary Kasparov
 - [AlphaGo](#), an AI capable of defeating the best player at the game of Go
2. It is also important to emphasize that for this AI, the agent will work towards a goal and will identify a series of actions that will contribute to the solution. Since there can be many solutions, it is important for that agent to evaluate the possible solutions and determine the optimal one. There are a number of search strategies that can be adopted to resolve various real-world problems, including but not limited to games:
 - Route finding problem: examples of applications include tools for driving directions in websites, in-car systems, etc.
 - Travelling salesman problem: find the shortest tour to visit each city exactly once.
 - VLSI layout: position millions of components and connections on a chip to minimize area, shorten delays.
 - Robot navigation: special case of route finding for robots with no specific routes or connections, where state space and action space are potentially infinite.
 - Automatic assembly sequencing: find an order in which to assemble parts of an object which is a difficult and expensive

geometric search.

- Protein design: find a sequence of amino acids that will fold into a 3D protein with the right properties to cure a specific disease.

(taken from Search Algorithms in Artificial Intelligence:

<https://tinyurl.com/y2h4a5kf>)




CONCLUSION

While the piece of paper is not generally intelligent in the way that humans are, capable of learning how to do a variety of tasks, in the case of this game, the AI is outperforming the human! As we have seen with the other activities and discussions, AI is not limited to games. We could potentially build an AI to solve other specific problems which are not games.


Who knows which other world problems can be solved with this type of AI solution?

What do you think would happen if two paper-based intelligences played against each other?



ACTIVITY 5

**SEARCH ALGORITHMS (II):
THE BEST CLEANING ROBOT**



INTRODUCTION

This activity aims to introduce a number of concepts, mainly the idea of the problem space that was discussed in the previous activity, using a search algorithm to plan a solution within that space, and discussing the idea of using heuristics to make more intelligent choices. It attempts this through a scenario where learners need to direct a cleaning robot (similar to a Roomba) to clean a virtual room.

TOPIC BRIEF FOR TEACHERS

While the previous activity focuses on the use of search algorithms with examples within games, this activity takes the usefulness of search algorithms and applies them to a real-life setting. We take the example of a cleaning robot and ask how it would know that it cleaned the whole room. Robotic cleaners are usually equipped with sensors and use path planning algorithms to make educated decisions about the environment they are in, ensuring they achieve their goal of cleaning the entire room. Since environments in

different homes or indoor spaces vary, the path planning coverage of the robot needs to take into consideration turning angles and areas revisited. Path planning algorithms determine how and where the robot should move.

These path planning algorithms make use of search algorithms as discussed in the previous activity. This activity will also take into consideration the term Heuristic as applied to the shortest path problem. It is important to take into consideration that in this case, Heuristics refers to the function on the nodes of a search tree, as the algorithm attempts to calculate the cheapest path from one node to the optimal node to achieve the goal set. The effect of the heuristic applied is to produce a more computationally efficient solution in terms of cost and time.

In this activity, it is also recommended that as an additional activity, the learners play the game VIPER [<https://tinyurl.com/y5vgvtqv>] which includes machine learning for a space exploration activity. The teacher can acquire a log in account prior to the class. The game will need an Internet connection. Tutorials and sample programs are provided online accessible through the link.

CHALLENGE FACTOR



This activity involves basic summation, however understanding the rules of the activity and how these and the summation drive the robot's decision making requires a non-trivial explanation.

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Explain how a Search Algorithm takes the environmental context into consideration to find a more optimal solution.
- Understand the concept of a Heuristic in a shortest path problem to estimate the cost of the 'cheapest' path, from one node of the search space to the ultimate goal – in the case of this example, clean the entire room.
- Discuss how, given the same situation, different AIs may behave differently to reach different goals, depending on the type of search algorithm selected and implemented.

CLASS INTRODUCTION

Have you ever seen or heard of cleaning robots in the house?

Have you wondered how a cleaning robot such as a Roomba would know how to navigate a room to clean it? Let me ask you a question: do you think that all houses and living spaces are the same? Do they have the same layout?

So how do you think that such a robot would know how to clean a room or an entire house?

Today we are going to attempt to understand how such a cleaning robot can explore the different house environments it finds itself in, by thinking about the different ways that an intelligent robot can navigate around.

LEARNING RESOURCES

- 1x Worksheet titled ‘The Best Cleaning Robot’
- Cleaning Robot Explained: [The Roomba Cleaner](#)
- Wonderville Machine Learning for Site Exploration: [ViPER – Machine Learning Game](#) (registration and sign up required – free access)

DISCUSSION

1. Learners are each handed the worksheet.
2. The worksheet can be read through with the whole class at once or worked out on a whiteboard to ensure that all learners understand the concept of the game.
3. Learners are then left to complete the rest of the exercises alone, aided with difficulties as necessary.
4. The suggested class discussion questions for the house cleaning robot are:
 - Does the robot know where it’s dirty and where it’s not?
 - If the robot is not 100% sure about where it’s dirty, should it use an algorithm which makes it go faster or an algorithm which makes it check more of the room?
 - If the robot knows exactly which part of the house is dirty, should it use an algorithm which makes it go faster or an algorithm which makes it check more of the room?



TEACHER AIDS

The information given below is meant to help the teacher navigate through the Best Cleaning Robot lesson. Some of these questions and prompts can be used during the lesson:

- Have you ever seen a cleaning robot in person or seen a video of one?

They're typically small, round, and run around on the floor, cleaning the floor as they go along.

- These robots use traditional forms of Artificial Intelligence to navigate around the room, sometimes to avoid going over parts that have already been cleaned and sometimes to make sure it makes its battery last as long as possible.

The scope of the lesson activity is to help learners understand that giving the robot different AI routines will result in a completely different path taken, despite it being the same robot with the same task and the same restrictions.

1. Explaining the problem space:

A problem space is a formal mathematical description of our problem. In our case, it is planning how to go about cleaning a set of dirty hallways. The problem space for our cleaning robot on paper can be visualized as follows:

- Clean tiles are empty
- Dirty tiles are marked with an X
- Carpet tiles are marked with a C

2. The robot is successful if all the following are achieved:

- The robot finishes at the “End Here” tile
- The robot’s battery points are greater than zero by the time it reaches the final tile

3. The robot has the following rules:

- It starts in the tile marked “Start Here”
- Moving from one tile to another takes one turn
- It cannot move to a tile it has already been inside
- It cannot move to a black tile, only to a white one
- It has 20 battery points at the start:
 - Cleaning a dirty tile consumes 1 battery point
 - Each move on a clean tile consumes 1 battery point
 - Moving into a carpet tile consumes 2 battery points

4. Instructions to be given during the class exercise:

- Calculate the total cost of traveling down each of the three possible paths shown in the handout
- To do this for a given path, calculate the cost of moving from the first tile to the last tile of that particular path
- The cost depends on whether it is clean, carpeted, or dirty
- Add all of the costs together at the bottom of each path’s table

5. The following includes additional information about search algorithms and their use in cleaning robots:

Search algorithms are sets of instructions that search through a problem space, trying to find an efficient solution to solving the problem. In this case, search algorithms are used to create different paths for the cleaning robot, optimizing for different outcomes. The outcomes we are aiming for are:

- Use as little battery points as possible:

This search algorithm optimizes the path to use the least amount of battery points possible. It discourages moving on carpet tiles as these are the most expensive uses of battery points, and prefers shorter routes to longer ones. This search algorithm would choose the path where the total cost is the lowest out of all the options.

- Clean as many dirty tiles as possible:

This algorithm optimizes the path for providing as thorough a clean as possible. It encourages choosing dirtier routes. This search algorithm would choose the path where the number of cleaned dirty tiles is the highest, even if the total cost is higher than other options.

- Avoid carpets as much as possible:


This algorithm optimizes the path to avoid carpet tiles as much as possible. It does this by not considering any path which has a carpet tile in it. If all of the paths contain carpet tiles, then it chooses the path with the least number of carpet tiles, regardless of the cost.

It is recommended that this lesson is followed up with the activity from [Wonderville Machine Learning](#).

Although access is free, registration and sign-up are required. This game requires Internet access and connectivity. This game is recommended for

learners who are more advanced in the area of computing and programming. The scope of the game is to program a robot to explore space in search of signs of life. A tutorial is provided so that learners can learn and navigate through the programming. The search space is organized into 6 different mazes that can be set up. Commands for navigation through the mazes can be dragged and dropped into the programming window. The program also features a sandbox where learners can experiment with creating their own obstacle-filled terrain for the ViPER robot to navigate through.


Learners can spend some follow up hours, playing the game online.



CONCLUSION

What did you learn about the different paths a robot can take? There are many different paths that each robot can take to solve the problem.

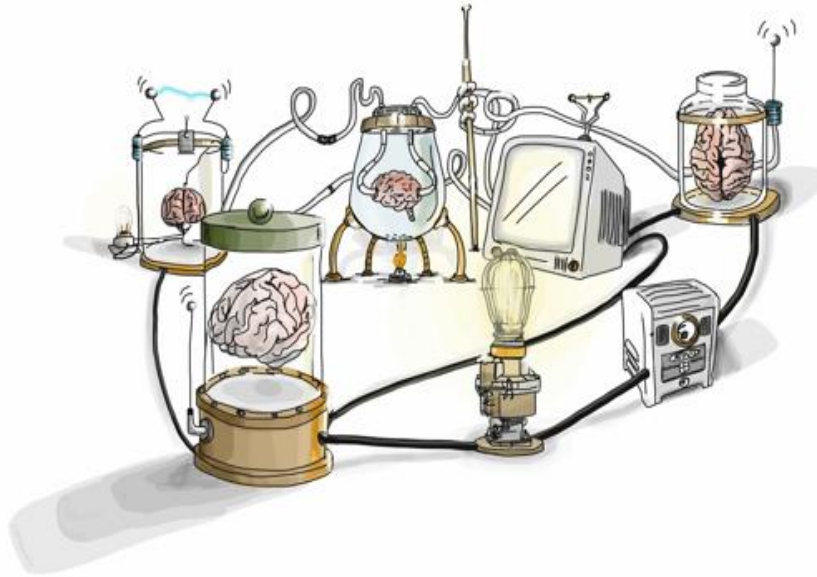
Different AIs are optimized for different problems, such as cleaning the room in the least number of steps, while a different AI is focused on avoiding carpets. A good way machines can exhibit intelligent behavior is to be able to switch between different sets of instructions on the fly without any human telling it what to do or how, depending on what is most important at that time and depending on the context in which it is placed! Choosing which intelligence to use at any given time is yet another possible area where intelligence can be applied – one could call it being intelligent about intelligence!



ACTIVITY 6

NEURAL NETWORKS (I):

THE PAPER NEURON



INTRODUCTION

This activity aims to introduce the concept of artificial neurons, their internal structure, and how they can be altered to mimic different mathematical functions. It includes a worksheet where the learner will alter a neuron on paper and see the difference in results. The follow-up questions are recommended to stimulate further discussion.

TOPIC BRIEF FOR TEACHERS

Whilst the first few methods of Artificial Intelligence involved optimizing existing algorithms through the use of *heuristics* (which are, simply put, variables that the algorithm takes into account to attempt to find a better solution), other AI methods emerged, with one of the most popular ones, which is by creating Artificial Neural Networks, being inspired by our own biological computer – the Human brain.

The human brain is a biological neural network that consists of several millions of tiny brain cells called *Neurons* are interconnected in a complex network. Neurons send and receive electrical signals from each other, the sum total of which gives rise to intelligent thought. Artificial Neural Networks work in a similar way – they consist of tiny *cells* which contain mathematical functions that interact with each other by feeding the results of their functions as inputs to other neurons. The sum total of these interactions can give rise to complex behavior, such as temperature prediction or image recognition.

The way the human brain works in terms of neural operations it is able to carry out is incredibly powerful in the domains of vision, speech, and motor processes, while it is less powerful in simple tasks such as multiplication. It makes sense, therefore, to expand research in the areas of neural networks to simulate how the brain works, to solve complex problems in domains that traditionally are more adept to humans than to machines.

CHALLENGE FACTOR



This activity requires a basic understanding of the mathematics behind artificial neurons to be explained confidently. It also involves basic mathematics, including sum of multiples and the use of negative numbers.

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Explain what an artificial neuron is and how it operates in simple terms.
- Demonstrate the differences between two neurons that are able produce different results.
- Use addition and multiplication operations to work out the activation of an artificial neuron.
- Explain how the application of artificial neurons can help solve different real-world problems in different domains.

CLASS INTRODUCTION

One of the most common forms of applied AI is the Artificial Neural Network. It is inspired by the human brain and consists of connecting small units together called neurons to perform different operations.

Intelligent behavior is exhibited when a collection of these neurons working together send signals to one another. The scope of AI is to make machines learn, and we do this by teaching them to learn the way humans do. This is a branch of AI that is called machine learning, and one way of doing this is by copying the way the human brain works artificially – these are the Artificial Neural Networks.

Let us break down these rather complex terms so that we can understand more about how machines can learn. If I were to ask you – how do you know that a fish is a fish – what would you answer? Would a baby know what a fish is? How would a toddler learn what a fish is?

What a human being does, as they grow older, is that they are able to store many images in their brain, and every time the human brain comes across that image or a similar one, it can associate the patterns and label an object – for example

“That is a fish!”

With machines, we can do something similar. We can give machines access to many images or data that may be different but may all be related, and the

machines learn from all this data without being given any instructions from humans. Once the machine learns, it can refine its learning so that it can answer many complex and difficult problems or questions. In this activity, we will explore how an Artificial Neuron performs calculations by first looking at a neuron that adds two numbers together, then adjusting it to only add numbers together if the result is positive.

LEARNING RESOURCES


- 1x Worksheet titled ‘The Paper Neuron’
- Artificial Intelligence Introduction to Neural Networks: [An AI Quest](#)
- Neural Networks: [How Neural Networks can help Doctors](#)
- Teaching Machines : [How Machines can learn like a Child](#)
- Machine Learning for Kids: [Teaching a Computer to Play a Game](#)
- Machine Learning for Beginners: [Teach a Computer to Make Art](#)

DISCUSSION

1. Learners are each handed the Paper Neuron worksheet: a. In this activity learners can be split into small groups.
2. The teacher can read through the worksheet with the learners and use exercise 1 as an example by demonstrating it on a whiteboard.
3. The learners are encouraged to complete exercise 2 by themselves.
4. The following discussion questions are recommended:
 - Do you think that a neuron can learn any kind of mathematical function?
[If we have access to enough examples (these can take millions of examples), we can find the ideal weight of nearly any problem].
 - What would happen if we use more than one neuron?
 - How could we connect neurons together? [The output of one neuron would affect the input of the next neuron in the network].

TEACHER AIDS

1. The following website is recommended to help teachers facilitate in their learners the grasp of Artificial Neural Networks:
<https://tinyurl.com/y3xwuxy7>
2. This additional workbook also contains material that can be used in class: <https://tinyurl.com/yyy83p7s>
3. A video series that explains the science behind Artificial Neural Networks includes: <https://tinyurl.com/yb8fa3ze>
4. It would also be interesting to explain how a neural network trains and what advantages its training leads to. In this way, it is important to clarify in simple terms that the strengths of a neural network lie in the patterns of information which are fed through its inputs and emerge from its outputs. In between, there are a number of hidden layers that connect both the inputs and the outputs. Information traverses these layers. Simply put, each unit sums up all the inputs it receives by assigning a set weight, and if a threshold is reached, the next hidden layer is activated so that the information can move forward towards the output.
5. Neural networks learn by receiving feedback, comparing the output obtained to the output they should have obtained, and making adjustments as necessary so that there is no difference between the two outputs.



CONCLUSION

What was this activity about? What did you learn from this activity?


In this activity, we taught a neuron on paper. Do you think that we could write a computer program to run and teach the neuron for us? Do you think that this could be done at a faster rate than a human working it out on paper? What happened when you changed the unit weights? What does this tell you about the neuron's behavior?

Following the ideas we discussed, what would your thoughts be on the use of neural networks in:

- Airplanes? [*hint*: as an auto-pilot, with input units reading signals from the various cockpit instruments and output units modifying the plane's

controls appropriately to keep it safely on course]


- Banks? [*hint*: as an automatic way to identify any possible fraudulent transactions using credit cards]
- Images? [*hint*: recognising images for example aerial images that show areas that are mostly littered with garbage]



ACTIVITY 7

NEURAL NETWORKS (II):

MECHANICAL MOZART



INTRODUCTION

This activity aims to introduce the idea of a neural network in a creative field such as music. By representing musical notes as the input units, the network will generate a different sequence of notes which will be given as a new output. The learners will be split into groups of four, where each group will represent a network and each learner will represent a neuron to understand how neural networks can operate in the creative arts.

TOPIC BRIEF FOR TEACHERS

You may consider creativity to be something unique to humans – after all, there are few living creatures that can paint, compose songs, or write poetry. Most animals, in fact, have fewer neurons in their brains than humans, and their brains are geared more towards psychokinetic actions such as running or jumping.

However, since AI using Neural Networks emulates the human brain, and the human brain is capable of what we consider creative tasks, then it follows that it is possible to perform creative tasks with AI. Generative music as a genre has existed long before the modern approaches to AI, especially given how the underlying structures of music (such as tempo or relations between notes) are all mathematically linked.

These mathematical links allowed for statistical approaches to music, such as generating jazz music depending solely on which notes have been played how many times recently to determine the next note. Using AI is an improvement on the previous approaches rather than a complete paradigm shift, unlike in other, more practical areas such as computer vision. More information about recent advances in AI music composition can be found here: [<https://tinyurl.com/y36szxms>]

CHALLENGE FACTOR



This activity does not require advanced technical knowledge; however, it involves cooperation within small groups which may require learner guidance or supervision.

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Explain how artificial neurons work together.
- Discuss how a neural network converts an input to an output depending on the set rules.
- Demonstrate how by adjusting the neurons within a network both the rules and the output can change despite the same input.
- Improve on team building skills.

CLASS INTRODUCTION

Have you ever heard of Wolfgang Amadeus Mozart? Mozart was a well-accomplished classical music composer. He lived from 1756 to 1791 in Salzburg, Austria. Mozart is no longer alive today, which means that he will never be able to compose any more music, and the world is indeed grateful for all the music he composed during his lifetime.

However, I would like to ask you – with all that you have seen in the previous weeks about the potentials and limitations of AI – do you think that AI can somehow act as a mechanical Mozart? What if we could create a neural network that acts as an artificial Mozart? One that could keep on producing music 24 hours a day, 7 days a week? In this activity, we will explore a proof of concept of how our 'Mechanical Mozart' would work.

LEARNING RESOURCES

- 1x worksheet titled 'Mechanical Mozart'
- AI and Music (IBM Watson): [Behind the Scenes](#)
- Orb Composer (free demo): [ORB](#)
- AIVA Music Composer: ["Genesis" Symphonic Fantasy in A minor, Op. 21](#)
- [AMPER](#): First album composed, produced entirely by AI

- Google Magenta: [Musical Deep Neural Networks](#) in the Browser
- Google Magenta Demo: [Piano Genie](#)
- Google's Bach Doodle: Behind the Doodle: [Celebrating Johann Sebastian Bach](#)

DISCUSSION

1. Learners are each handed the Mechanical Mozart worksheet.
 2. Learners are split into groups of four.
 3. The worksheet can be worked through together with all the groups at once or the groups left to work them out on their own.
 4. The worksheet exercise is an introduction to the interaction between different neurons. The exercise can be carried out individually on paper, or as a group, with one person representing each neuron. In the end, the neural network will look like the figure on page 56, where the output neuron observes the behavior of the input neurons.
- Could a single neuron generate different music from what it is fed?
 - Why do you think the input neuron instructions do not take into consideration the other input neurons?
 - If we fed exactly the same sequence of notes twice, would the output be the same each time?

TEACHER AIDS

1. Most often at the heart of music composition relying on AI, there is a deep learning technology. The way music composition works, is that a piece of music is first decomposed into the various parts of rhythm, pitch and timbre.
2. These would act as the input data points feeding into neural networks. The interesting aspect of any AI music composer such as IBM's Watson, is that of changing the various input points, to generate a different music output.

3. It is therefore important that such AI-based composers are trained with large music datasets so as to generate their own music models which can be used in various instances. In the case of AIVA, this produces sheet music for orchestras to play as background to film/video, so it is more of a composer than a music performer. Other deep learning composers may also perform music as well as compose it.

In answer to the discussion questions:

- Could a single neuron generate different music from what it is fed?


Yes, if the internal rules or the weightings/score (thus triggering different paths of hidden layers) change.

- Why do you think the input neuron instructions do not take into consideration the other input neurons?

Each neuron will trace its own path depending on the rules set by the algorithm. This will trigger hidden layers of neurons, which will in turn trigger other layers to navigate to the output.

- If we fed exactly the same sequence of notes twice, would the output be the same each time?

Unless the rules and the weightings change, then yes.




CONCLUSION

Through simple instructions at each neuron, we were able to create a new piece of music from an existing piece of music. If we kept feeding different notes to the neural network then we could potentially create an infinite amount of new music!

If we wanted to generate a different style of music, all we need to do would be to change the instructions that each neuron follows.

If we give the right instructions, we could switch our network to generate music in a completely different style, from pop to jazz to classical using the same inputs.



ACTIVITY 8

PHILOSOPHY (I): INTERROGATING A ROBOT



INTRODUCTION

This discussion activity is meant to introduce the idea of conversational AI and AI that can mirror human behavior, such as that found in chatbots.

A fictional scenario is created where the participants are interrogating a user, as they attempt to determine whether whoever is answering is either a human or a robot. This particular scenario paves the way for research into the communication and cooperation between humans and robots in the field of robotics.

TOPIC BRIEF FOR TEACHERS

While other subjects tend to take a theory-first approach, ending with practice (since theory is typically required to understand and perform the practical part), one can have a whole career in AI and never need to deal with the philosophical implications of the creation of intelligence, since it may never hold any bearing on the practical output of an AI system. Nevertheless, the rapid adoption of AI systems is going to make these questions inevitable, and being aware that there is a metaphysical, philosophical reality to AI would help stifle misinformation and fear-mongering.

Provoking discussion about the incredibly abstract topic of the nature of intelligence may be difficult at this academic level, but this activity is strongly encouraged, especially if your learners are serious about taking a holistic approach to AI.

Currently, mainstream robotics are there to help humans carry out boring everyday tasks, such as vacuuming and cleaning the floors, or turning lights on or off. Research in the field directs the future of robots toward working alongside humans at a communication level that uses natural language (the language that humans use).

Robots may be used to carry out various jobs and tasks, and assigned different levels of responsibility depending on the complexity of AI

developed. However, it all depends on the level of trust humans can place in robots, and how AI will establish its trustworthiness. More details and information can be accessed online from <https://tinyurl.com/y2jo68vf>

CHALLENGE FACTOR



This activity is more open ended than others without a clear goal. It requires the teacher to provoke discussion through philosophical questions which might be lost on younger learners.

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Be more aware of the existence of a branch of AI that takes a philosophical approach.
- Consider and discuss the implications that AI may have on our society at a deeper level.
- Create a meaningful conversation about AI and how it can be used to effectively impinge positively on different societal strata and at different levels.

CLASS INTRODUCTION

Imagine that today we are a team of detectives investigating a crime. A suspect has been taken into custody, but we do not have the authority to meet them face to face.

We have been given permission to interrogate the suspect through a microphone and speaker. Our assignment is to discover whether this suspect is a human or a robot, as it may change how the law is applied in this case.

LEARNING RESOURCES

- 1x worksheet titled 'Interrogating a Robot'
- Humans and Robots: [How Robots can be taught basic Interaction](#)
- Debating the future of humanity: [Sophia & Han](#)
- [Google I/O Keynote](#)

DISCUSSION

The activity can be introduced by checking out the Google I/O 2018 Keynote: [Google Duplex demo](#) (at 35:00). Following this, a discussion should be stimulated using the questions below as a guide.

1. What kind of questions would you ask to find out if the suspect is a human?
2. What kind of questions would you ask to find out if the suspect is a robot?
3. What if the suspect begins lying?
4. What if it turns out the suspect is a robot, but is being controlled by a human?
5. If it turns out to be a robot should the robot be considered guilty?

6. If it turns out that the AI was programmed to be lawful, but then the AI learned to break the law, would it be guilty?
7. After the discussion, learners are each handed the worksheet and following completion, a new discussion is started by asking:
 - Who thinks the suspect is a robot?
 - Who thinks the suspect is a human?
 - Who “Doesn’t Know”?
 - Can you tell me why you have answered in this way?

TEACHER AIDS

The scope of the discussion is that of being open ended.

Therefore, there are no right or wrong answers and a deeper exploration of the topic is encouraged. In order to stimulate more questions, it is also recommended that additional questions are used to probe the concepts at a deeper level. A few examples include:

1. What kind of questions would you ask to find out if the suspect is a human?
 - How about asking about typical human experiences, something a robot would not know about?
 - Human Experience - emotion: “Are you feeling sad that you committed a crime?”
 - Human Experience - food: “Are you being given enough food while you are in your cell?”
2. What kind of questions would you ask to find out if the suspect is a robot?
 - How about asking about experiences which are difficult for a human to perform but easy for a computer to perform?
 - Human Experience - complex mathematical calculations: “What is 17 times 16 divided by 7?”
 - Human Experience - remembering something that requires a large memory space: “What are the first 1000 digits of Pi?”
3. What if the suspect begins lying?
 - How could we guess that the suspect is lying?

- What about lie detection tests? Would it work on both a human and a robot?
- 4. What if it turns out the suspect is a robot, but is being controlled by a human?
 - Who is guilty here, the human, or the robot?
- 5. If it turns out to be a robot and the robot is guilty, is it the robot that is guilty?
 - Is it the AI running inside the robot?
 - Is it the human who programmed the AI?
- 6. If it turns out that the AI was programmed to be lawful, but then the AI learned to break the law, is it still guilty?
 - If found guilty, would a 'normal' punishment apply?

During the worksheet activity, it is also important to bring out the concept of what makes us human and what gaps exist in current AI developments that make them different from humans. It would also be interesting to gather from learners if they would want the AI to be closer to a human being, and the reasons why or why not.

It might also be interesting to touch upon issues from this article:

<https://tinyurl.com/y8u6wyug>


That deal with philosophical issues when autonomous robots are expected to take decisions that affect human life.



CONCLUSION

Many of the other activities and discussions were based on practical applications, exploring how Artificial Intelligence can make machines more intelligent or enable computers to solve problems that currently only humans can perform.

This activity approached AI from a philosophical perspective. This perspective is very important as it opens up several moral and ethical dilemmas that come hand in hand with thinking machines. What would happen if machines became as intelligent as humans in every aspect, or more? There are many films and books that attempt to address this topic, usually with disastrous results. Can you mention any?



ACTIVITY 9

PHILOSOPHY (II):

DOES ALEXA HAVE FEELINGS?



INTRODUCTION

This discussion aims to introduce the idea of weak and strong AI, emphasizing the difference between a machine that learns and a machine that acts consciously about the decisions it takes. During this activity, it is important to bring out the differences between a machine that has been programmed to act in a specific way under set conditions and one that acts in a completely autonomous way, differently from any program any human might have written for it. It attempts this by fielding questions about a popular voice assistant (specifically, Amazon Alexa), however, it can be generalized for others such as Google Assistant, Samsung Bixby, or Apple's Siri.

TOPIC BRIEF FOR TEACHERS

The way computers operate can sometimes not be fully understood, especially by those people who are not conversant with technology, aside from its use in social media.

Only after spending a non-trivial amount of time struggling with learning the principles behind computing and programming, can someone start to understand how machines and computers operate on a series of instructions executed one after the other.

There is a similar lack of understanding in the field of AI mostly because to many it is still a novel concept, and also because of the ‘transparent’ way that AI works, such that most people do not even notice it is running in the background.

One feature of AI is that, while it may not be given explicit instructions other than to use large amounts of data, it then becomes capable of predicting or recognizing other instances of that same data. Currently, many consumers are fascinated yet scared of the potential of AI as it makes its way into people’s daily lives and homes.

The future of AI indicates one which requires more of the human ingenuity to create an artificial intelligence that is capable of improving the quality of life of people, through the research, predictions and solutions it can provide, whether these are in the field of healthcare and medicine, to the financial, education, or entertainment industry.

CHALLENGE FACTOR



This activity is similar to the previous one and requires the teacher to provoke discussion about a typically abstract concept. Despite this, given that Voice Assistants are not rare, it may be easier to find practical examples from the learners' own experiences.

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Describe a Voice Assistant (VA) and how it operates.
- Understand better that it is not difficult for a VA to assume human-like features.
- Discuss how AI can be used to create more interactive and human-like experiences.
- Reflect on how VAs currently do not comprehend their own existence the way humans do.
- Reflect about how AI and Humans may mix without distinction in the near future.

CLASS INTRODUCTION

Have you ever tried one of the new fancy Voice Assistants? Can you mention some? [*Have you heard about Amazon Alexa or Google Home?*]

Some people use them to play music, get news read to them, or even control lights in the house. But if I were to ask you whether the voice assistant is an actual person, what would you say? Do they have feelings, or are they simply a computer program? What makes it a computer program? What would make a voice assistant more human-like? Today, we are going to discuss these voice assistants and the role of artificial intelligence in making them work.

LEARNING RESOURCES

- 1x Worksheet titled 'Does Alexa have Feelings?'

- For schools having Alexa: [Setting up Alexa](#)
- Google Home: [Google Home for Kids](#)
- For schools having Google Home: [The Google Assistant](#)
- Speech test on Voice Assistants: [Young Children try out Voice Assistants](#)

DISCUSSION

This activity is meant to provoke discussion where learners can reflect on what distinguishes humans from robots or machines. It is recommended that the teacher follows the teacher aids to further deepen the class discussion.

1. Do you consider Alexa to be a person? Why or why not?
2. Do you think that Alexa can have feelings and emotions? Why or why not?
3. If Alexa tells you that she is feeling sad, does that mean that she is actually feeling sadness?
4. If you ask Alexa whether she knows a foreign language such as Japanese and she says yes, would you believe her? Why or why not?
5. If you were to put Alexa behind a curtain, do you think that you would be able to guess that she is a Voice Assistant and not a human? Why or why not?

TEACHER AIDS

The scope of the discussion is that of being open ended.

Therefore, there are no right or wrong answers, and a deeper exploration of the topic is encouraged. In order to stimulate more questions, it is also recommended to use additional questions to probe the concepts at a deeper level. A few examples include:

1. Do you consider Alexa to be a person? Why or why not?

- What does it take to be considered a person?
 - Is a convincingly human sounding voice enough to be a person?
 - Do you need a body to be a person?
 - If yes, is a cat a person, because they have a body?
2. Do you think that Alexa can have feelings and emotions? Why or why not?
- What qualifies as a feeling?
 - Are only humans capable of feelings?
 - Are animals like cats and dogs capable of feelings?
3. If Alexa tells you that she is feeling sad, does that mean that she is actually feeling sadness?
- Do you need to be sad to be able to say the words, “I am sad?”
 - Is a computer capable of sadness?
4. If you ask Alexa whether she knows a foreign language such as Japanese and she says yes, would you believe her? Why or why not?
- If you downloaded a language pack for Alexa that allowed her to speak in Japanese, does that mean that you taught her Japanese?
5. If you were to put Alexa behind a curtain, do you think that you would be able to guess that she is a Voice Assistant and not a human? Why or why not?
- Would your answer change if she sounded more realistic?
 - Would your answer change if she initiated the conversation and spoke to you instead of you to her?


In addition, it is also recommended that the teacher can read up on the following article, using some of its questions to stimulate further discussion and provoke reflection on the topic of voice assistants and their role in people’s homes: <https://tinyurl.com/ybcv137q>



CONCLUSION

As Artificial Intelligence begins to sound, look, and act more realistically, we are shifting away from interacting with computers through keyboards and touchscreens towards gestures and voice commands.

Could AI eventually replace humans in jobs, for example, those found in call centers? What kind of changes would realistic AI bring into your life? Are these good changes or bad changes? What do you think would happen if AI became indistinguishable from humans?



ACTIVITY 10

MAKING MACHINES LEARN



INTRODUCTION

In these previous sessions, we gradually came to understand that AI can have different applications in different fields in society. Do you remember what we discussed? We discussed a number of ways in which AI could train and learn. How does AI learn and how does the way it learns affect its purposes?

These ways in which AI learns are branches of artificial intelligence called Machine Learning. Machine Learning is the technology behind face and emotion recognition, targeted advertising, voice recognition, detecting credit card fraud, virtual personal assistants, self-driving cars, etc.

TOPIC BRIEF FOR TEACHERS

This class activity is meant to be quite flexible, and the teacher can choose which resources to use and adapt in class. Although there is no specific order in which activities could be carried out, it is recommended that this activity follows one where neural networks have been covered. The main scope of this activity is to simulate machine learning using neural networks in a physical class environment.

The teacher needs to get hold of either a pack of flashcards with different animals/insects/plants/objects, etc. The goal is to use different versions of the same dataset to understand how machines train in recognizing patterns.

The teacher can also choose to ask the students to draw or write different emotions people might show and use those for the suggested activity.

To include a greater challenge, the teacher might demonstrate the network with a small number of layers, and the learners might design their own versions of hidden layers depending on the complexity of the problem chosen.

CHALLENGE FACTOR



This activity can require some abstract thinking and can develop into a complex problem solving attempt.

LEARNING OUTCOMES

By the end of this session, the learner shall be able to:

- Understand and discuss the basics of how neural networks in machine learning work.
- Design layers for different neural networks requiring different complexity.
- Demonstrate the use of machine learning applications in an everyday context.

CLASS INTRODUCTION

Today we shall understand how a programmer can set up a machine to learn how to recognize either text, sounds, colors, images, or numbers. Let us first understand what we need to do if we were to attempt to spot a pattern or, in this case, design ways in which a machine can play Snap! [*Teacher needs to explain the concepts as per the Teacher Aids section*].

LEARNING RESOURCES

- 1x Worksheet titled ‘Making Machines Learn’
- 2x1m & 2x2m string for each group of 7 students (more string will be needed depending on the addition of neuron layers)
- Toilet paper rolls (6 rolls for a 7-student network layer)
- A set of Flash Cards (e.g. living creatures/plant or even Pokémon cards!)
- [Machine Learning Projects for Kids](#) (requires setting up of a free IBM cloud account)
- Downloadable [Leaflet - Artificial Intelligence](#) ... but where is the Intelligence?
- [Children’s Computing Magazines](#)
- Follow [CS4FN](#) on Twitter for additional resources
- AIMMO: [A coding adventure](#) (registration required)
- [Pica AI Digital Pet](#) (ages 5-9) for iOS

DISCUSSION

If I were to ask you the questions:

1. What is Machine Learning?
2. How would you describe Supervised Learning?
3. How do you think that the machine/computer is able spot patterns?

Let us take this example:

Imagine that today we shall simulate the artificial brain, and you will be the neurons (which we have seen in other activities), helping it learn to spot patterns and recognize colors, animals or shapes (or any other flashcard pack content).

TEACHER AIDS

The book, [Machine Learning for Humans](#) can help teachers gain more insights into this field of AI and the branches of society that it can impinge more on.

As a teacher, you need to bring across certain important notions of Machine Learning: this is a technology that “*allows computers to perform specific tasks intelligently, by learning from examples*”.

In practice creating machine learning tools needs to be done in a number of steps.

1. The first step is that of teaching the program how to learn through a series of instructions (the algorithm).
2. The second step is training. The relevant data (e.g. images of sports cars) is given to the program. The more data the program has the greater the possibility of learning more about cars.
3. Through an interface, users can use the application to give it new data (more photos of vehicles), so that the program can tell them which vehicles are sports cars or not.

It is suggested that a simple game of Snap is first illustrated so that the students understand how to pose the right questions to help patterns emerge. Following that, students are encouraged to plan and design their own, possibly more complex networks to help the machine learn the identification of set objects.

For this activity, the students will have to think about how to physically construct a neural network to aid machine learning in the field of pattern matching.

The students can create their own designs but essentially, the targets include:

- Target data is chosen (for example red/black cards, birds from a pack of living creatures flashcards or the energy cards in a game of Pokémon cards).
- Learners need to make a list of properties that they can use to classify their data (for example color, size, or anything that makes the target data different from the others).
- Using the properties the learners emerged with, they have to build yes or no questions that would help classify the objects (for example does it fly?). Students must adhere to the rule that they can only ask questions which identifies only *one* property of the target data.
- Once the plan is done, the students need to construct a physical network of nodes to acquire an output. This physical network can be constructed with rope/string, and using toilet paper rolls to fire neurons from the input cards to gather an output. Learners can take the place of neurons and fire the toilet rolls from one node (learner) to another when the criteria (question asked) is met (answered as yes).
- The diagram below refers to a network that is built to teach the network to play a simple game of Snap. The target data is made up of a pack of cards. In this game, nodes 1, 2, 3, 4 refer to a visual input – I1, I2, I3, and I4. The pack of cards is distributed in half. One half of the pack will be laid between I1 and I2, whilst the other half of the pack will be between I3 and I4.
- Learners will have to act out the role of the neurons, which fire (toilet paper rolls down the string) when their input is triggered. Therefore, for the neuron at position I1, the trigger will be a red card. If a red card is picked from the pack, then the student will fire the toilet roll towards the student who is at position Red + Red. At I3, the toilet roll will be fired towards position Red + Red when the card that is picked is also red. Once the rule for the position Red + Red has been fulfilled, the person at Red + Red will fire the toilet roll towards the last person in the output line, who will shout ‘Snap’!

- This is one example that can be repeated with other cards, and the students can experiment with layout of the network, complexities, etc.




CONCLUSION


When people hear the term Artificial Intelligence, the first image they think about is a robot as seen in many science fiction movies. Some of these robots may be harmful to humans; some, in the movies, even want to take over the world. As we saw through these sessions, the reality of AI is quite different. For one, we are still far away from having machines that scheme to take over the world, but most importantly, our aim is that of having machines that work hand in hand with humans to improve their overall quality of life.

In the meantime, AI is becoming more and more popular and the intelligence is being exploited in a number of fields. I would like you to think a bit more in-depth about how AI can be used:

- As an expert in medicine?
- To help create music?
- To see where others do not reach?
- To self-drive?



PART B



WORKSHEET 1

SHOPPING WITH AI



Worksheet 1

Fill in this table with some of the applications which make use of AI that you come across in your search using the Internet. You can include links to the applications that you have found:

Practical Application of AI


Download link

You can download this exercise [here](#)

REFLECT !


How do you think Amazon and other online shopping platforms know which products to best suit what you want?

How can AI be used to improve the online shopping experience?



WORKSHEET 2

ROCK, PAPER ... AI?



PLAY INSTRUCTIONS

These instructions are for playing the game Rock, Paper, Scissors & AI. The rules are:

- Rock beats Scissors.
- Paper beats Rock.
- Scissors beats Paper.

For this activity, you will need the table on this worksheet and a six-sided dice.

The teacher or learning assistant (adult) will play the part of the AI. As a student, you will need to play the rock, paper, scissors game as you would usually do.

The adult playing the part of the AI will need to follow these rules:

- Instead of choosing whether it's rock, paper or scissors, the adult will roll the dice.
 - o for 1 or 2, choose Rock.
 - o for 3 or 4, choose Paper.
 - o for 5 or 6, choose Scissors.
- The adult needs to fill in the table on the other page with the details below:
 - o If the adult wins the game, the number in the relevant column of the table is increased by one.
 - o If the adult draws the number in the relevant column of the table is decreased by one.
 - o If the adult loses, the number in the relevant column of the table is decreased by three.

- o If any numbers go below zero, they should be left as zero.
- After 5 rounds, the adult needs to start choosing the next move depending on which column has the highest confidence number:
 - o If two numbers have the same confidence, roll the dice. If you get 1-3, choose the one on the left, otherwise the one on the right.
 - o If all three numbers have the same confidence, roll the dice. For 1 or 2 choose Rock, for 3 or 4 choose Paper, for 5 or 6 choose Scissors.
- Keep track of your wins and losses by filling in the score table.
- After a few rounds, the filled in tables should look something like this:

Confidence Table

	Rock	Paper	Scissors
Confidence	0 1 2 3 2	0 1 2 1	0 1 0

Score table

Wins	1 2 3 4 5 6
Losses	1 2 3

By assigning different *return* and *reward* values to winning, losing or drawing against our opponent, the AI is able to adapt to which option the opponent chooses the most often, thus optimizing for maximum wins or at least draws. This means that the AI is learning and improving over time!

This AI is particularly limited though, as it does not consider the previously chosen options by the opponent. If the opponent constantly chose random

options, this AI would not be able to adapt.

Worksheet 2

Confidence Table

	Rock	Paper	Scissors
Confidence			

Score table

Wins	
Losses	


Download link

You can download this exercise [here](#)

REFLECT!


How could you use the history of moves to help the AI to make better choices?

What do you think would happen if the AI played against itself?



WORKSHEET 3

AI OLYMPICS



PLAY INSTRUCTIONS

This time, you as the students will take your turn to each play as the AI, as AI vs.

AI battle out the game of Rock, Paper, Scissors:

- You will be split into pairs or groups and each given a worksheet to fill in. Every group will have a 6-sided dice to play with.
- You will need to roll the dice every time you need to make a move against your opponent:
 - for 1 or 2, choose Rock.
 - for 3 or 4, choose Paper.
 - for 5 or 6, choose Scissors.
- You are going to play a number of games using this method and fill in the worksheet as you accumulate wins or losses.
- After 5 rounds, you will stop using the dice and start using your filled-in table to choose your moves to win against your opponent. Choose a column that has a higher confidence number, to predict your wins or your losses!
- The winner of each pair or group will keep playing against the winner of the next pair until only two players remain. The AI's should play to win, using their confidence table to predict which of the rock, paper, scissors has the greater chance to win!
- After this Olympiad, you should compare your tables with those of your friends, and discuss any differences.

Worksheet 3

Confidence Table

	Rock	Paper	Scissors
Confidence			

--	--	--	--

Score table

Wins	
Losses	


Download link

You can download this exercise [here](#)

REFLECT!


How did you make your choices to win the games?

How did the AI vs. AI affect the outcomes of the games?



WORKSHEET 4

THE PAPER-BASED INTELLIGENT BEING



PLAY INSTRUCTIONS

This worksheet is for a paper-based intelligent being – a piece of paper that is capable of always either winning or drawing against opponents in the game of Tic-Tac-Toe.

You will need printouts of Worksheet 4.

If you are playing first and are playing as X, always start by placing an X in a corner:

- If the opponent places an O in the center, place an X either on the opposite side of the O or on an edge square that is not touching your first X.
- In this case, a win is only possible if the opponent makes a mistake.
- Continue the game by blocking the opponent while trying to form a line across an edge.
- If the opponent places an O along the edges, a win is guaranteed. Place an X in any corner as long as there is an empty space between the first X and the one you are placing now and not an O.
- Regardless of where the opponent places their next O, place your next X in a way that your X pieces either form a triangle or are all in a corner.
- Win with your fourth X by placing it in any position which forms a line. Your previous move should have created two possible places, to which your opponent can only counter one.

If you are playing second and are playing as X:

- If the opponent starts in a corner you can only force a draw by placing your first X in the center.
- If you do not need to block your opponent from winning, place your next X along an edge, but not in a corner.
- Keep playing until you force a draw.

- If the opponent starts in the center, you can only force a draw by placing your first X in a corner.
- Place all the following X pieces in a way that the opponent is blocked from scoring.
- If the opponent starts on an edge but not at a corner, there is a chance to win by placing your first X in the center.
- If the opponent places an O on the opposite edge (forming a column or row reading OXO), put your second X on a corner.
- If your opponent once again creates an OXO line by placing their O adjacent to your X, you can win by blocking their row of two O.
- Otherwise, simply block your opponent until you force a draw.

Tic-Tac-Toe is an example of a mathematically solved game – that is the best instructions are already known to either win or draw every time. These instructions take into consideration the situation and changes its actions depending on how the opponent moves – which is a sign of a simple form of intelligence.

Whilst Tic-Tac-Toe is very simple, consisting only of two different types of pieces and a 3x3 grid, the more complex AI has made it possible to defeat humans at games such as Chess, Go and even DOTA 2.

WORKSHEET 4

--	--	--


Download link

You can download this exercise [here](#)

REFLECT!


How do you think the intelligent being learned to play Tic-Tac-Toe so well ?

Discuss in class more complex world-problems that AI would be capable of solving.



WORKSHEET 5

THE BEST CLEANING ROBOT



PLAY INSTRUCTIONS

A 'Cleaning Robot' goes around the rooms in a house to clean the floors. These robots use traditional forms of Artificial Intelligence to navigate around the room, sometimes to avoid going over parts that have already been cleaned and sometimes to make sure it makes its battery last as long as possible.

In this worksheet, you will see how installing different AI routines in a virtual robot on paper will result in completely different paths taken, despite it being the same robot with the same task and the same restrictions.

In this exercise, the cleaning area is defined as the 'Problem Space'. Our problem is about planning how to go about cleaning a set of dirty rooms. The problem space is going to be visualized as follows:

- Clean tiles are empty
- Dirty tiles are marked with an X
- Carpet tiles are marked with a C

The robot is successful if all the following are achieved:

- The robot finishes at the "End Here" tile
- The robot's battery points are greater than zero by the time it reaches the final tile

The robot has the following rules to follow:

- It starts in the tile marked "Start Here"
- Moving from one tile to another takes one turn
- It cannot move to a tile that it has already visited
- It has to move only to a white tile and skip the black tile
- It has 20 battery points at the start
 - Cleaning a dirty tile consumes 1 battery point
 - Each move on a clean tile consumes 1 battery point
 - Moving onto a carpet tile consumes 2 battery points

HOW TO PLAY

1. Calculate the total cost of traveling down each of the three possible cleaning paths.
 - For a given path calculate the cost of moving from the first tile to the last tile
 - The cost depends on whether the tile is clean, a carpet, or dirty
2. Add all of the costs together at the bottom of each path's table

Below you can find a worked-out example to guide you. The Robot starts at position A1, then can only move towards the right (since it cannot go to a tile it has already been to) so there is only one recommended path.

In this case, the path chosen will always be the same.

Example Problem Space

	A	B	C	D	E
I	Start Here		C	X	End Here

Example Search Space Table - Grey Path

Turn	Current Position	Cost		
		Move onto empty or dirty tile	Move onto dirty or clean tile	Move onto carpet tile
1	AI	0	0	0
2	BI	1	0	0
3	CI	0	0	2
4	DI	1	1	0
5	EI	1	0	0

Turn	Current Position	Cost		
		Move onto empty or dirty tile	Move onto dirty or clean tile	Move onto carpet tile
Total Cost of each variable		3	1	2
Total Cost		6		

SEARCH ALGORITHMS (Navigation Paths)

Search Algorithms are sets of instructions which search through a problem space, trying to find an efficient solution to solving the problem. In our case, we are going to use search algorithms which take *3 different paths*, optimizing for different *outcomes*. The outcomes we are aiming for are:

1. Use as little battery points as possible:

This search algorithm optimizes to use the least amount of battery points possible. It discourages moving on carpet tiles as these are the most expensive uses of battery points and prefers shorter routes to longer ones. This search algorithm would choose the path where the total cost is the lowest out of all the options.

2. Clean as many dirty tiles as possible:

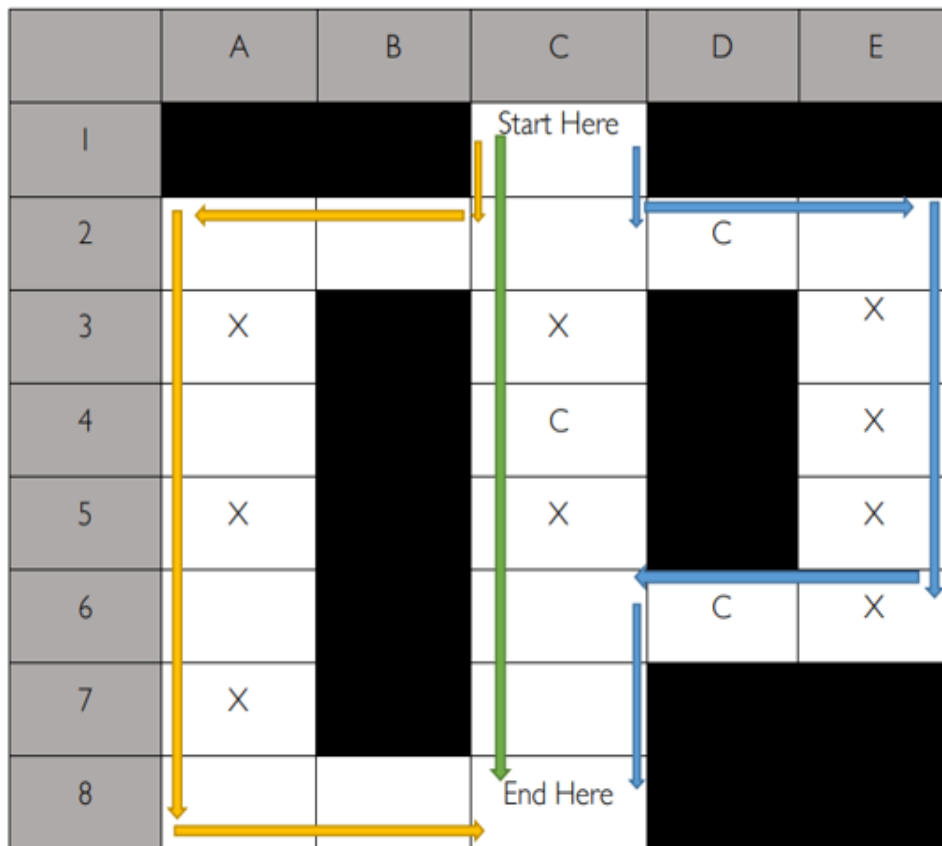
This algorithm optimizes for providing as thorough a clean as possible. It encourages choosing the 'more dirty' option path. This search algorithm would choose the path where the number of cleaned dirty tiles is the highest, even if the total cost is higher than other options.

3. Avoid carpets as much as possible:

This algorithm optimizes to avoid carpet tiles as much as possible. It does this by not considering any path which has a carpet tile in it. If all

of the paths contain carpet tiles, then it chooses the path with the least number of carpet tiles, regardless of the cost.

WORKSHEET 5



In the problem space above, there are clearly three paths we can take, each marked with a color. Each path corresponds to one of the three specific search algorithms described in the previous page.

Calculate the cost of taking each path, where the first move of the robot is from C1 to C2, following by either moving to the left (B2), downwards (C3) or to the right (D2) always ending at C8.

SEARCH SPACE TABLE – YELLOW PATH (ON THE LEFT)

Turn	Current	Cost

Turn	Position Current	Move onto empty or dirty	Cost Move onto dirty	Move onto
	Position	Move onto empty or dirty tile	or clean tile Move onto dirty or clean tile	carpet tile Move onto carpet tile
1	C1			
2	C2			
3	B2			
4	A2			
5	A3			
6	A4			
7	A5			
8	A6			
9	A7			
10	A8			
11	B8			
12	C8			
Total Cost of each variable				
Total Cost				

SEARCH SPACE TABLE – GREEN PATH (IN THE CENTER)

Turn	Current Position	Cost

Turn	Current Position	Move onto empty or dirty tile	Cost Move onto dirty or clean tile	Move onto carpet tile
		Move onto empty or dirty tile	Move onto dirty or clean tile	Move onto carpet tile
1	C1			
2	C2			
3	C3			
4	C4			
5	C5			
6	C6			
7	C7			
8	C8			
Total Cost of each variable				
Total Cost				

SEARCH SPACE TABLE – BLUE PATH (ON THE RIGHT)

Turn	Current Position	Cost		
		Move onto empty or dirty tile	Move onto dirty or clean tile	Move onto carpet tile
1	C1			

Turn	Current Position	Cost		
		Move onto empty or dirty tile	Move onto dirty or clean tile	Move onto carpet tile
2	C2			
3	D2			
4	E2			
5	E3			
6	E4			
7	E5			
8	E6			
9	D6			
10	C6			
11	C7			
12	C8			
Total Cost of each variable				
Total Cost				

Download link

You can download this exercise [here](#)

REFLECT!

Given the above tables, which algorithm do you think the robot used to go through each path?


Yellow Path (use as little battery as possible) on the Left: _____

Green Path (clean as many tiles as possible) in the Center: _____

Blue Path (avoid carpets as much as possible) on the Right: _____


Through this exercise, we can see how intelligence can arise from simple mathematical rules which, when changed, can achieve different results, even though the robot does not change.

Using the above, you can come up with your own rules (such as, all the dirty tiles need to be cleaned, or the robot can run in “fast mode”, where it can move two tiles in one turn but for double the cost).



WORKSHEET 6

THE PAPER NEURON



PLAY INSTRUCTIONS

This is our artificial neuron on paper:



On the left, we can input numbers into an artificial neuron (the Bias is a special input which is always set to the number 1) and, depending on the internal values, a value is output the other side. In this case, the Neuron is acting as a *black box*, meaning we do not know what its internal values are set to. We can only observe what the output will be, depending on what input we provide to it.

Here is an example of feeding a table of inputs to a neuron and recording the output:

Input 1	Input 2	Bias	Output
0	3	1	3
1	2	1	3
2	1	1	3
3	0	1	3

In this case, we can easily guess that the neuron has been trained to simply add the two inputs together, by looking at the inputs, then at the output.

This is not always the case though. What if we have a table such as this?

Input 1	Input 2	Bias	Output
----------------	----------------	-------------	---------------

Input 1	Input 2	Bias	Output
0	3	1	3
1	2	1	6
2	1	1	5
3	0	1	6

In this case, guesswork is not so useful. Instead of guessing, we can take a look at the internal values of the neuron, which will show us what the neuron has been trained to do.

If we reveal the neuron's internals (similar to how you would open the bonnet of a car to see its engine), the following will be revealed:



A *Weight* is how much the neuron takes notice of that input number. Weights are usually decimal numbers such as 0.85 or -0.333343, however they can still be normal numbers such as 1, 2, 3, etc.

The *Net* is the result of a *weighted sum* of the inputs. This translates to multiplying the input with the corresponding weight, then adding them together, as follows:

$$Net = (Bias * Weight B) + (Input 1 * Weight 1) + (Input 2 * Weight 2)$$

The final step to determining what the output is a very important mathematical property in every artificial neuron: the *activation function*. The

activation function takes the net and from it calculates the output.

Typically choosing an activation function depends on what we are using our neurons for. One such example is the *linear* activation function. This function can be described as follows: *Output the net*.

This is one of the simplest activation functions. All it does is output the value of the net. There are many activation functions, and we will look at the *relu* function as a second example in the second exercise.

Using these instructions, we can calculate the output for the neuron, given any input and a specific set of weights. For example, if we set the weights to the following, the Neuron will be trained to simply add both inputs together, similar to the first example above.

- Weight B = 0
- Weight 1 = 1
- Weight 2 = 1

WORKSHEET 6

Exercise 1 : A neuron which adds the two inputs together

In this exercise, we are going to confirm whether our neuron actually adds two numbers together correctly. We will try

WeightB = 0	Weight1 = 1	Weight2
-------------	-------------	---------

To test the provided weights above, fill out the following table:

Input 1	Input 2	Bias*WeightB	Input1*Weight1	Input2*Weight2	Neuron Output
2	2				
1	1				
0	0				
-1	-1				
-2	-2				
2	-2				
1	-1				
-1	1				
-2	2				

Exercise 2: A neuron which adds two numbers together but ignores negative numbers

In this exercise, we will change the neuron to only output positive numbers. We can achieve this by changing the activation function to:

If the Net is greater than zero,

WeightB = 0	Weight1 = 1	Weight2
-------------	-------------	---------

To test the new behavior of the neuron, fill out the following table:

Input 1	Input 2	Bias*WeightB	Input1*Weight1	Input2*Weight2	Neuron Output
2	2				
1	1				
0	0				
-1	-1				
-2	-2				
2	-2				

Input 1	Input 2	Bias*WeightB	Input1*Weight1	Input2*Weight2	N
1	-1				
-1	1				
-2	2				

Download link

You can download this exercise [here](#)

REFLECT!

By changing the internal values of the neuron, we were able to change its behavior completely. In these exercises however, we only changed the *activation function*.


Another option would have been to change the *weights* instead.

What do you think would the result have been?

This process is typically called *Back Propagation*.


If you are interested in learning more about how Neural Networks and Artificial Intelligence work, there are some additional resources to help you continue to find out and learn more:

1. Tensorflow Playground: an in-browser fully functional neural network simulation. This may be complicated compared to what we explored today, but try to play around with it anyway!
<https://tinyurl.com/yc9qydo2>
2. ConvNetJS: a series of in-browser, interactive demonstrations of various AIs, including [Neural Networks](#)



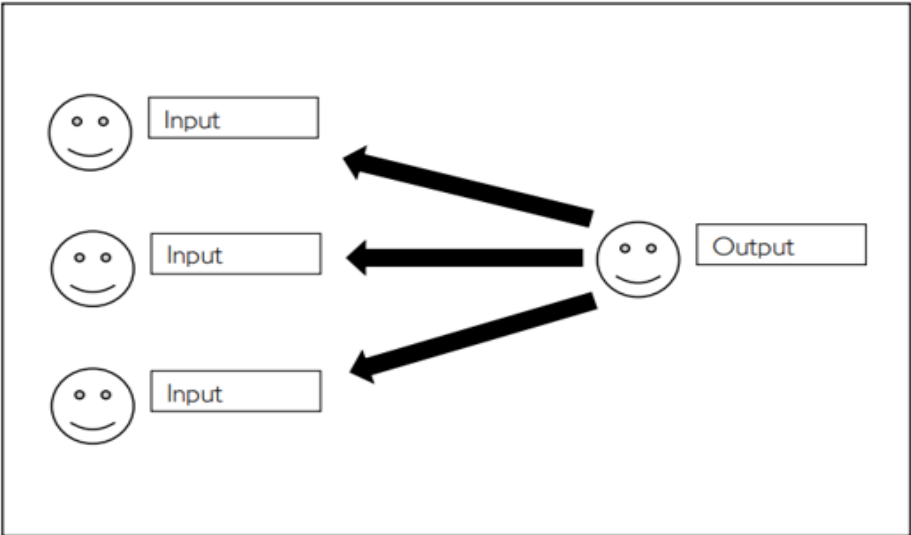
WORKSHEET 7

THE MECHANICAL MOZART



PLAY INSTRUCTIONS

This exercise is an introduction into the interaction between different neurons. You can run it alone on paper, but you could also run it as a group, with one person representing each neuron. In the end, our neural network will look like this, where the Output Neuron observes the behavior of the input neurons.



Step 1: Choose neurons

Three people will represent the input neurons. Write down their names here:

Neuron	Person name
Input Neuron 1	
Input Neuron 2	
Input Neuron 3	

One person will represent the output neuron. Write down their name here:

Neuron	Person name

Neuron	Person name
Output Neuron	

Step 2: Copy down the instruction manuals

Copy your instructions down from the tables below. Make sure to only copy your own!

If you are an Input Neuron: You need to raise (or not raise) one or both hands depending on the note in the current step. How many hands you need to raise is provided in your instructions below.

If you are an Output Neuron: You need to write down a letter on a piece of paper, depending on how many raised hands you see, regardless of who is holding them up. Which letter to write down for each number of hands is provided in your instructions below.

Input Neuron 1 Instruction

Note on Paper	Instruction
C	Raise your left hand only
D	Raise your right hand only
E	Raise both hands
G	Do not raise any hands

Input Neuron 2 Instruction

Note on Paper	Instruction
C	Do not raise any hands
D	Raise your left hand only
E	Raise both hands

Note on Paper	Instruction
G	Do not raise any hands

Input Neuron 3 Instruction

Note on Paper	Instruction
C	Raise your left hand only
D	Raise your right hand only
E	Raise your left hand only
G	Raise your right hand only

Output Neuron Instruction

Note on Paper	Instruction
0	Write down C
1	Write down D
2	Write down E
3+	Write down G

Step 3: The Input Neurons need to read their notes and obey instructions while the Output Neuron writes down result.

Each input neuron has to look at the following list and obey their respective instructions according to the letter at each step. Once the output neuron has written down their answer, cross out the row with a pencil and move onto the next row. If there is no letter provided, an input neuron must not raise any hands.

WORKSHEET 7

Sequence 1: Sweet Home Alabama

The notes in this sequence are taken from the song “Sweet Home Alabama” by Lynyrd Skynyrd.

Step	Neuron 1	Neuron 2	Neuron 3	Output Neuron Result
1	D			
2	C	C		
3	G	G	G	
4	G	G	D	
5	D	D	C	
6	C	C	G	
7	G	G	D	
8	D	D	C	
9	C	C	G	
10	G	G	D	
11	D	D	C	
12	C	C	G	
13	G	G	D	
14	D	D	C	
15	C	C	G	

Step 4: Compare the output result with the input

Once you have completed the exercise, rewrite the output neuron's answer in the table below and see whether the song that has been generated is the same, similar, or completely different!

Step	Original Note (Of Neuron One)	Output Result	Is the note the same ? Tick this box if yes
1	D		
2	C		
3	G		
4	G		
5	D		
6	C		
7	G		
8	D		
9	C		
10	G		
11	D		
12	C		
13	G		
14	D		
15	C		

Download link

You can download this exercise [here](#)

REFLECT!

The above input and output neuron rules allowed us to translate the notes of one song into another, different song. If we kept passing in different notes, we would keep generating new and different songs. This shows how simple rules, on paper, can actually have a creative result when following the instructions set by the rules.

- Did you see any specific pattern emerge in the output?
- What do you think would happen if you change the input neuron instructions without changing the output neuron, or changing the output neuron without changing the input neurons?

You could try the exercise above again by applying a set of rules instead (you can follow the suggestions below or else create a completely new set of rules):

Input Neuron 1 Instruction

Note on Paper	Instruction
C	Do not raise any hands
D	Raise your left hand only
E	Raise your left hand only
G	Do not raise any hands


Input Neuron 2 Instruction

Note on Paper	Instruction
C	Raise your left hand only
D	Do not raise any hands

Note on Paper	Instruction
E	Do not raise any hands
G	Raise your left hand only

Input Neuron 3 Instruction

Note on Paper	Instruction
C	Raise both hands
D	Raise both hands
E	Raise both hands
G	Raise both hands



WORKSHEET 8

INTERROGATING A ROBOT



You have been tasked with figuring out whether our suspect criminal is a human or a robot. Due to various laws, we do have sufficient clearance to be in the same room as the suspect. As a result, we have been given a transcript of an interrogation of the suspect. Now it is up to you to determine whether the suspect is a human or a robot, depending on their answers.

PLAY INSTRUCTIONS

1. Read carefully through the transcript below, do not write anything for now
2. Once you have read it, read through it again and for each question and answer, mark the appropriate column on the right, whether you think that the suspect is a human, a robot, or you are not sure.
 - H stands for Human, R stands for Robot and ? stands for “I don’t know”
 - Only fill in where the row is grey, not black (that is, for the suspect only, not the interrogator).
3. Once you have filled in at least one box on each line, count the total amount of ticked boxes and write them down at the end.
4. After filling each box, think again whether your intuition matches with the column that has the highest number, then write down whether you think that the suspect is a human or a robot.

WORKSHEET 8

Interrogation Transcript

For the purpose of this transcript, the real name of the suspect has been replaced with “Joe”.

This does not mean that “Joe” is definitely a human.

Person	Transcribed Words	Suspicion		
		H	R	?
Interrogator	So, tell me Joe, did you commit the crime ?			
Suspect	No, I am innocent !			
Interrogator	Well we have records showing that your key card was used to enter the room at the time of the crime.			
Suspect	The system is lying ! i was not there !			
Interrogator	Then where were you Joe ?			
Suspect	I was getting repairs done.			
Suspect	My memory has been acting up.			
Interrogator	So, you have been forgetting things recently ? Important things ?			
Suspect	No no, not I have been forgetting, my memory has been forgetting.			
Interrogator	Well, what is the difference ?			
Suspect	Well i am not my memory. I leave my memory here when I leave work.			
Interrogator	So, do you remember when you last used your key card ?			
Suspect	Yes, I used my key card to approve the repairs I needed on my memory.			
Interrogator	And not to open the door to the room of the crime ?			

Person	Transcribed Words	Suspicion		
		H	R	?
Suspect	No, I already told you, I did not commit the crime !			
Interrogator	But, if it was not you, then someone used your key card to get into the room at the time of the crime.			
Suspect	Clearly someone stole my key card then.			
Interrogator	Maybe it was stolen while you were getting repairs ?			
Suspect	Or maybe the technician copied it, I don't know !			
Interrogator	Who was the technician ?			
Suspect	It was a standard diagnostic technician. He told me that my memory was faulty as it was causing several flipped bits every 24 point 55 milliseconds and that a repair would require my key card and for me to power down my terminal.			
Interrogator	Let us take a break, because clearly, we are getting nowhere on this problem, maybe you can remember better in an hour.			
Suspect	Whatever you think is best			
Total of each Column				

I think that the suspect is a _____

Download link

You can download this exercise [here](#)


REFLECT!

In this activity, we have examined whether our mystery suspect is a human or a computer.

The issue is, that we actually cannot know! Our interrogation was a modified form of a test called the Turing Test, a test proposed by Alan Turing to check whether someone being examined is a human or a computer.

Do you think that a computer could lie?

The Turing test is only a very basic test. Can you think of any other tests that you could use to differentiate between a human and an Artificial Intelligence?



WORKSHEET 9

DOES ALEXA HAVE FEELINGS?



WORKSHEET 9

A human is conversing with Alexa. In the first column create your own comic strip or a conversation which you think a human could have with Alexa. In the second column identify the different emotions that can be linked to that part of the conversation.

Conversation	Emotion

Download link


You can download this exercise [here](#)

REFLECT!

What different types of emotions can you identify?

What would it mean for a voice assistant or a robot to be emotional?

How do you think it would affect humans if voice assistants were to show more emotions in their conversations?




WORKSHEET 10

MAKING MACHINES LEARN



WORKSHEET 10

1. Using string, toilet paper rolls, beads, or other objects design your network.
2. Different objects will represent data. Pick one object as your target data.
3. Make a list of the properties that you can use to classify your data (anything that makes your target data different from the other objects).
4. Using the list of properties, write yes or no questions that help classify the objects. For example: is the object blue? Remember to ask only questions about one property at a time!
5. Once you're done with your plan, build a network of nodes and connections between the nodes. Each node will be a different question. You can show your design in the box below.

A large, empty rectangular box with a thin black border, intended for the student to draw their network design. The box is currently blank.



Download link

You can download this exercise [here](#)

REFLECT!


How did you construct your network?

Did it give the desired output?

Do you think the right questions were asked in the right places?

Which other questions do you think could have been asked? Should there have been other layers of questions that needed to be triggered?

Can this network be used to solve a more complex sorting problem?



TO GO FURTHER




GAN's AI

Welcome to the world of Generative Adversarial Networks (GANs)! In this series of exercises, you will delve into one of the most fascinating and rapidly evolving areas of artificial intelligence. GANs, introduced by Ian Goodfellow and his colleagues in 2014, represent a class of AI algorithms used for unsupervised learning, with the ability to generate new data samples that resemble a given dataset.

GANs consist of two neural networks, the generator and the discriminator, which are trained simultaneously through adversarial training. The generator creates synthetic data samples, while the discriminator evaluates them against real data, attempting to distinguish between the two. The goal of the generator is to produce data so convincing that the discriminator can no longer differentiate between real and generated data.

Download link

You can download these exercises [here](#)



... A LAST WORD

In movies, AI is often portrayed as human-like robots bent on destroying humanity. However, the reality of intelligent machines could mean many other things. AI researchers aim to create more intelligent machines that can assist us and enhance our well-being, thereby improving overall human quality of life. Helping humans solve problems that are too complex for us is just one way intelligent machines can contribute to a better life!

1. [AI in Education](#)