openSAP
Exploring the World of AI

Activities Week 1 Unit 1 – Hints and Solution

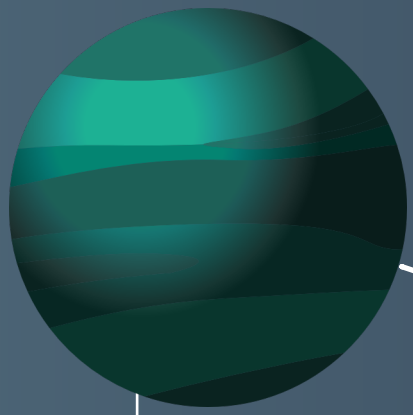

What is AI?

EXPLORING THE
WORLD OF AI

What is AI?

Activities Week 1 Unit 1 – Hints and Solution

In today's unit we looked at what AI actually is. To continue this discussion, we have prepared the following tasks for you.

Activity 1

Where do you experience AI in everyday life? Take a photo or describe an example of AI on our virtual pin board in the course forum. Comment on another example: What is the role of AI in this example?

Activity 2

Discuss in the forum: How intelligent would you rate the following applications on a scale from 1 (stupid as bread) to 10 (smart as Oskar)? And why?

- **Quickdraw:** https://quickdraw.withgoogle.com
- **ClickClickClick:** https://clickclickclick.click (Turn Volume on)
- **Rock-Paper-Scissors:** https://tenso.rs/demos/rock-paper-scissors/
- **Generative Engine:** https://experiments.runwayml.com/generative_engine/

You will encounter more experiments throughout the course. But if you know any other AI experiments that you would like to share with the group, post them in the forum, too!

A lot of interesting posts in the forum :-)

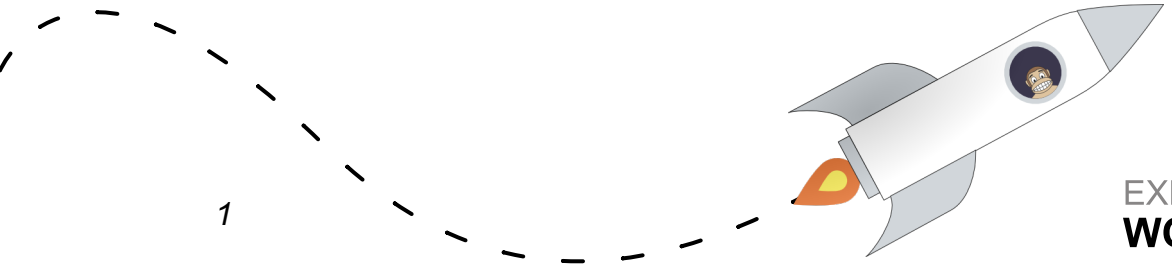

openSAP
Exploring the World of AI

Activities Week 1 Unit 2 – Hints and Solution

EXPLORING THE
WORLD OF AI

An Agent That Learns

Activities Week 1 Unit 2 – Hints and Solution

Activity 1:

In order to remember things, it is helpful to write down key aspects once again in your own words. Try to note down the idea behind reinforcement learning in your own words. Also write down at least one concise example.

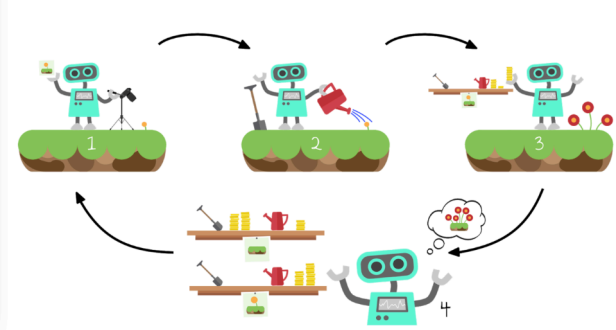

Reinforcement learning	*Description*
Examples that can be solved using reinforcement learning *Playing games* *Controlling robots* *Optimizing heating/cooling control systems*	*In reinforcement learning, the agent — a computer program capable of autonomous behavior — learns to better predict the success of its actions. Through continuous interaction with its environment, it is repeatedly rewarded or punished and thus optimizes its strategy.*

Activity 2:

The game Flappy Birds was a huge success on smartphones. At www.stefanseegerer.de/reinforcement-learning-flappybird/index.html you can watch the agent learn to play the game Flappy Birds through reinforcement learning.

1. Tap or click on the game to start the learning process. Watch the agent get better and better.

2. Set the "Show target" option to "Yes". What does the red dot mean?

   The red dot indicates the target the bird is trying to reach.

3. When the agent has learned enough to fly between the tubes, change the environment from static to random while the game is running. Argue why the agent is now failing again at the tubes.

   Changing the environment from static to random changes the position of the tubes. However, the agent has learned only one (static) position of the tubes so far - it cannot transfer what it has learned to the new positions.

openSAP
Exploring the World of AI

Activities Week 1 Unit 3 – Hints and Solution

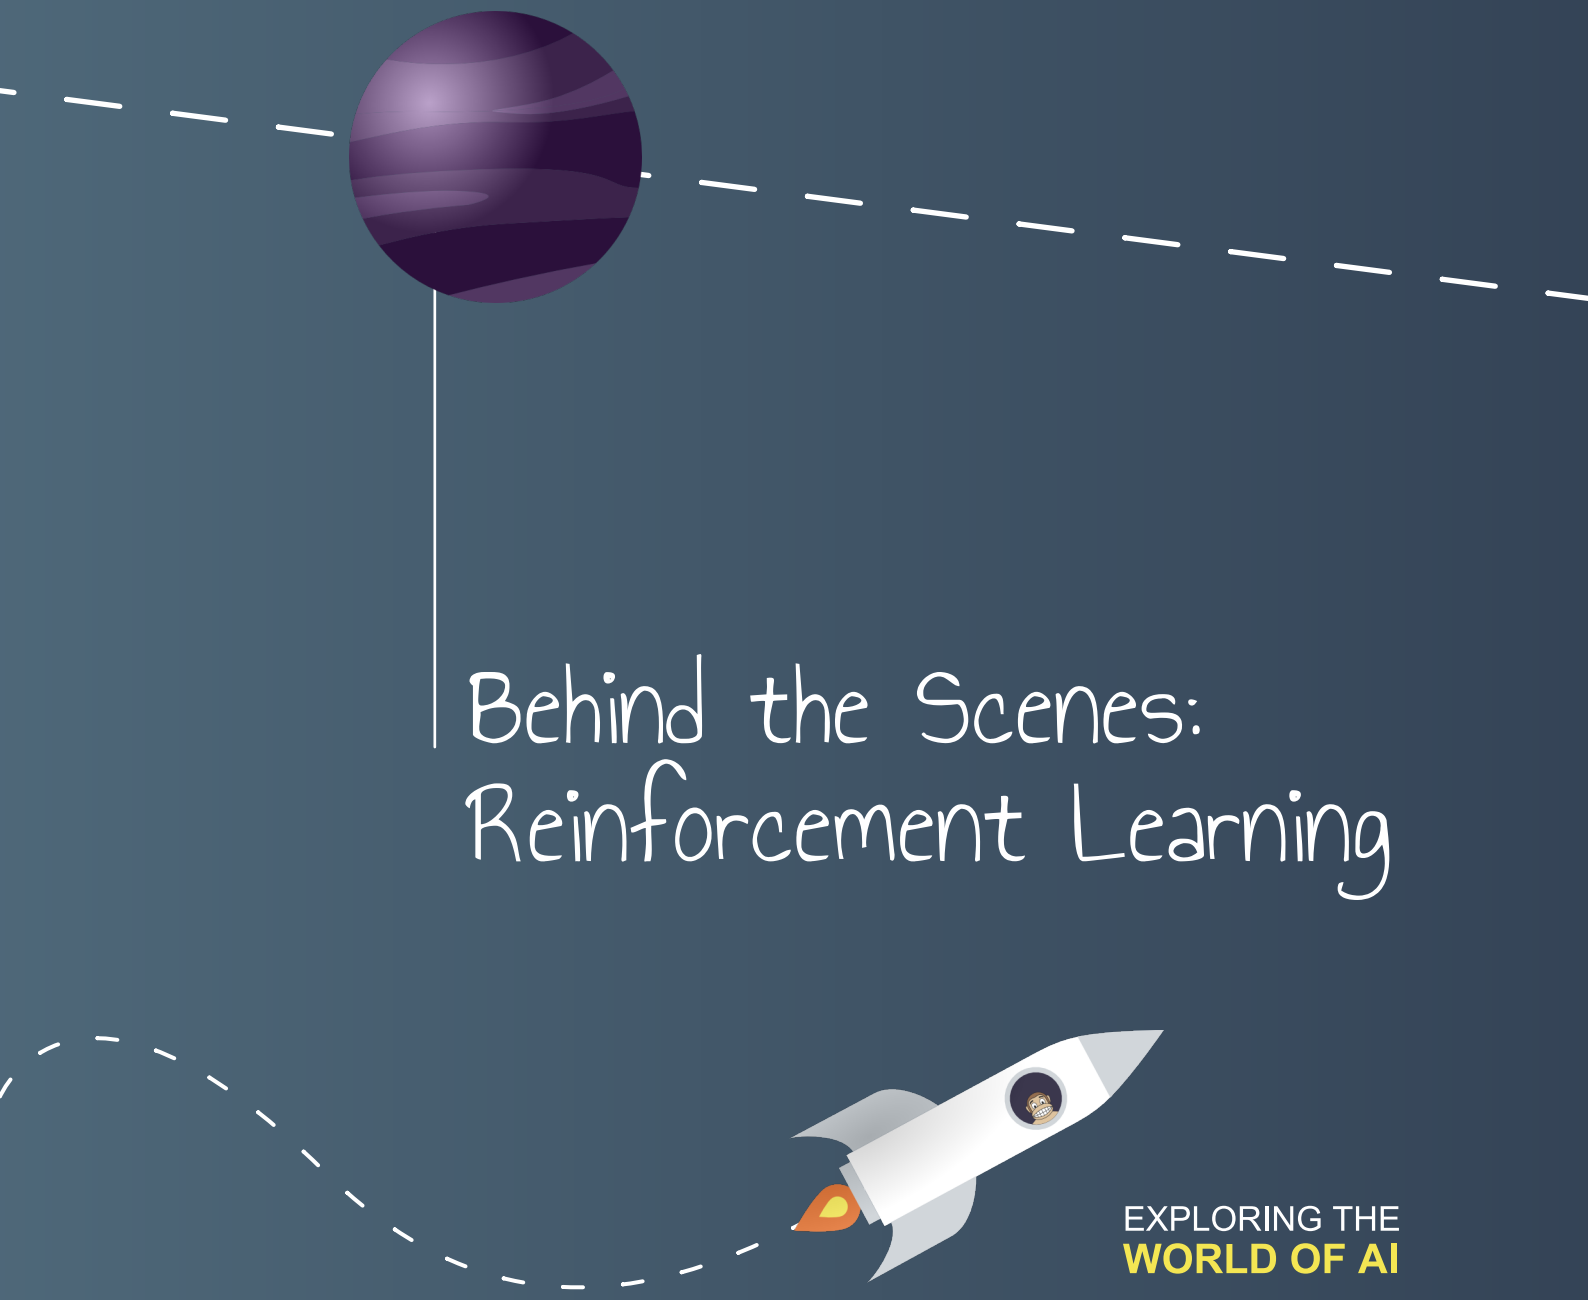

Behind the Scenes: Reinforcement Learning

EXPLORING THE
WORLD OF AI

Behind the Scenes: Reinforcement Learning

Activities Week 1 Unit 3 – Hints and Solution

You can apply the same principle that we used in the video to other games, such as Pong. Pong is an arcade classic. The exciting thing about this example is that you can compete against your self-learning agent! Here's the link:

https://snap.berkeley.edu/snap/snap.html#present:Username=seegerer&ProjectName=MOOC-RL-Pong-EN-template&editMode&noRun

The left paddle is already controlled by the computer. Your task is to make the paddle on the right learn to play Pong successfully. If you remember the banana chase example from the video, there were three main aspects to identify:

- What actions can the agent perform?
- How should the agent be rewarded/punished?
- How can the agent perceive its environment (= state)?

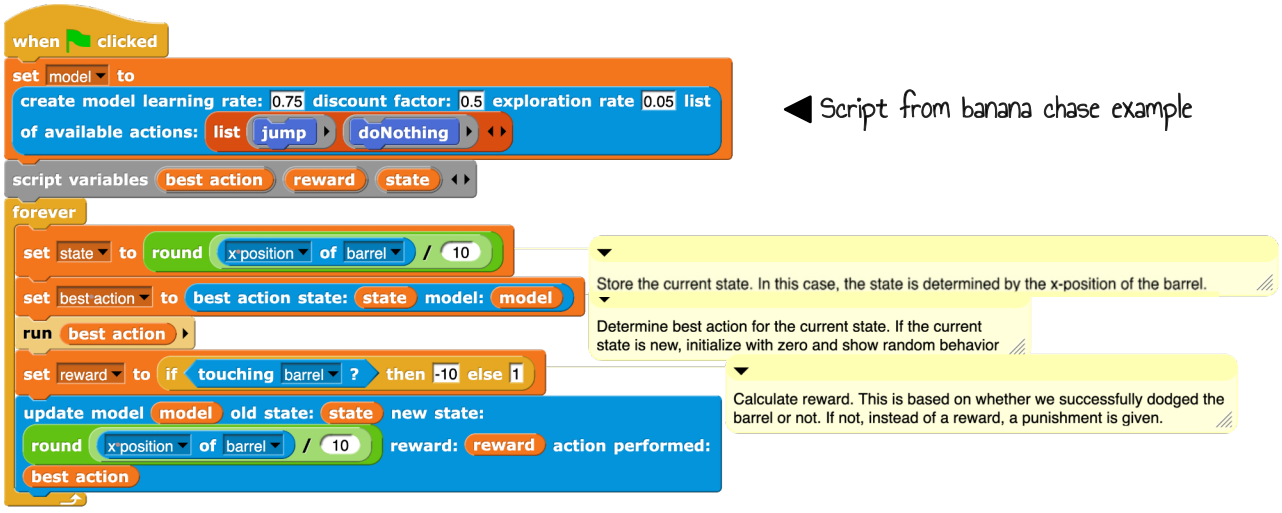

◀ Script from banana chase example

Now transfer this approach to Pong! You can use the coding cards (provided in the course materials) for this.

- We have already given you the possible actions.
- What is suitable as a state? *Tip: One possibility is to use the distance between the ball and the paddle as a state.*
- Experiment with possible rewards and punishments!

Give your model a few minutes and observe whether it learns a meaningful behaviour.

How would you describe the learned behaviour of the Pong paddle?

Coding Cards: https://computingeducation.de/SnAIp-cards-en.pdf

If you have problems, either use the coding cards or take a look at the solution:

https://snap.berkeley.edu/snap/snap.html#present:Username=seegerer&ProjectName=MOOC-RL-Pong

openSAP
Exploring the World of AI

Activities Week 1 Unit 4 – Hints and Solution

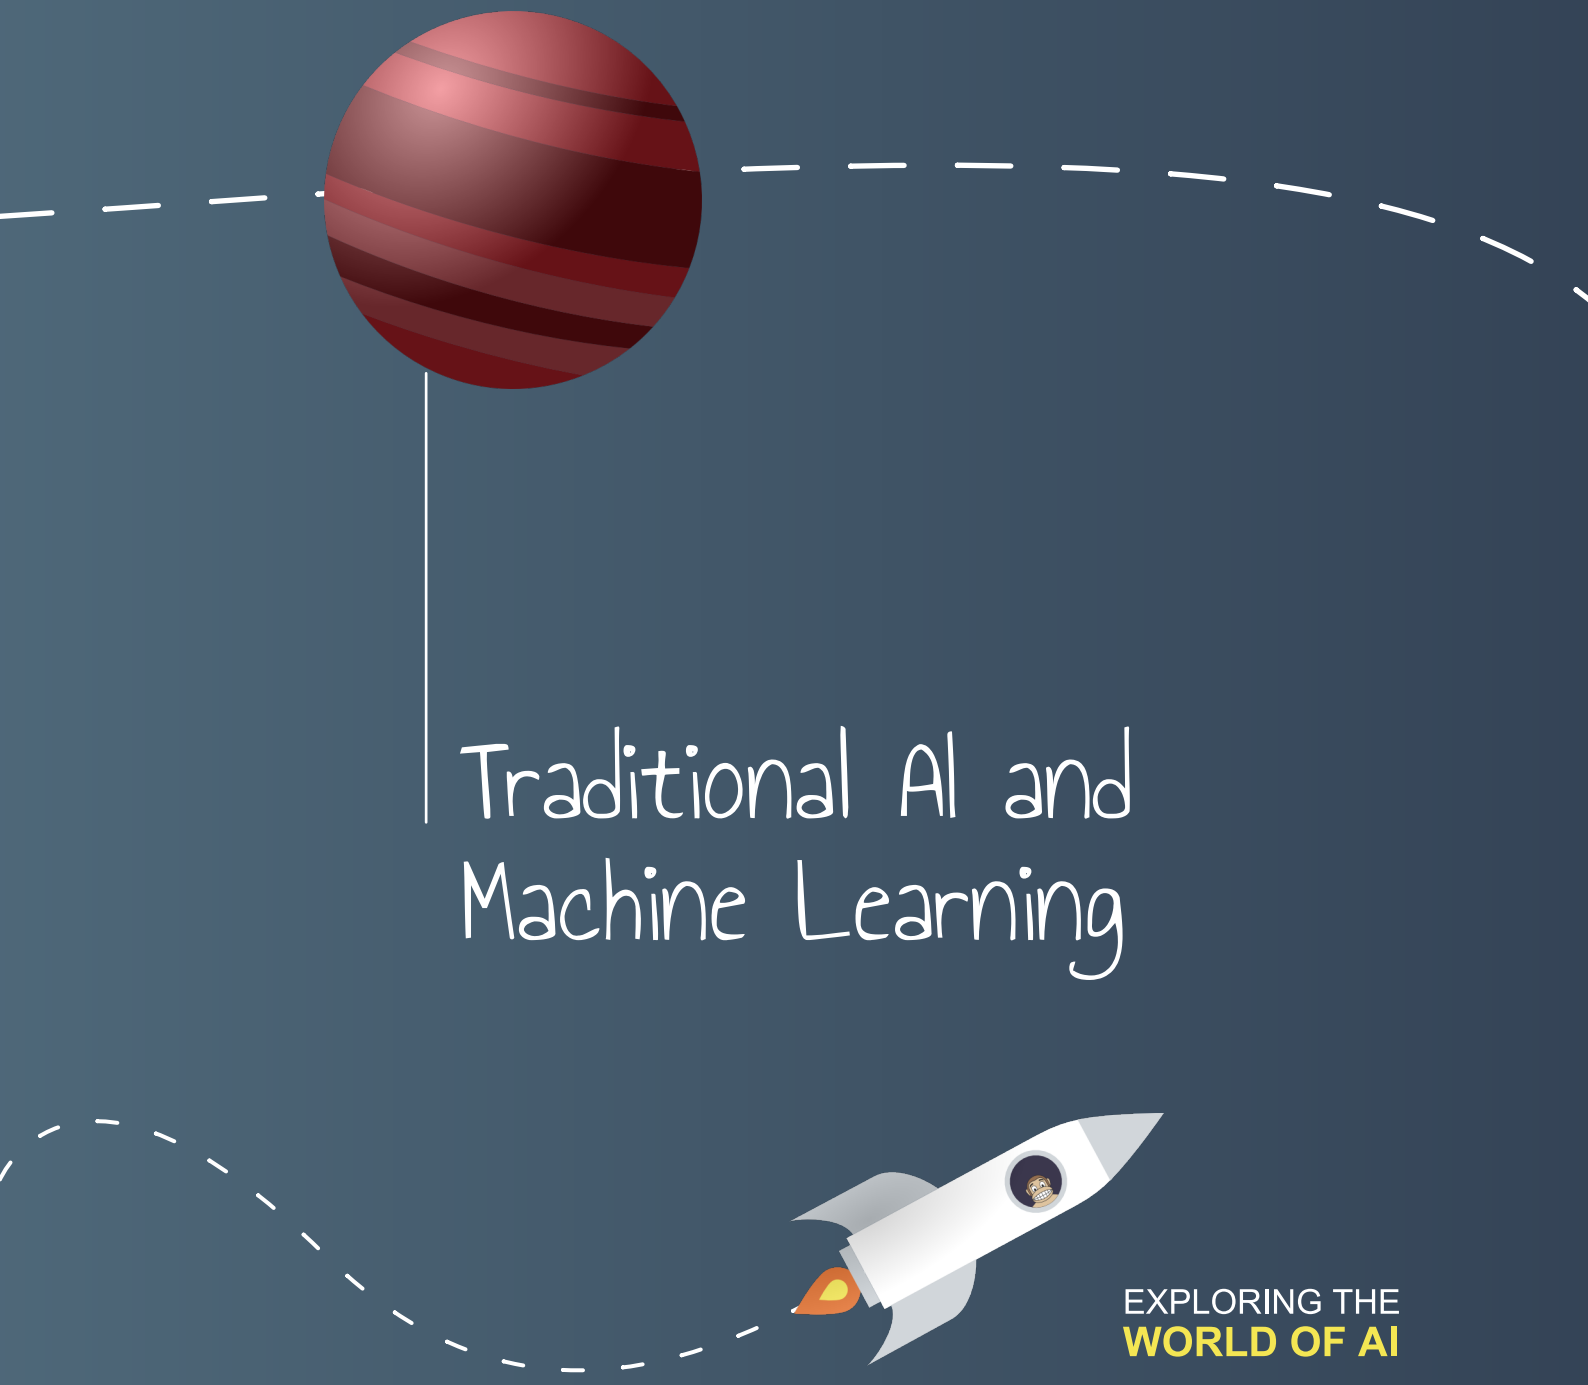

Traditional AI and Machine Learning

EXPLORING THE
WORLD OF AI

Traditional AI and Machine Learning

Activities Week 1 Unit 4 – Hints and Solution

Activity 1: Knwoledge representation

In Mushroom Land there are only **red** cap mushrooms, only **blue** sponge mushrooms and only **green** leaf mushrooms. There are also only **yellow** power stars, only **yellow** master stars and only **green** turbo stars. Furthermore, there are only two types of flowers: **yellow** dandelions and **blue** poppies grow in mushroom land.

Since the inhabitants of Mushroom Land repeatedly have questions about the various elements of their world and their special properties, an artificial intelligence is to be used that can answer such questions. To do this, the existing knowledge must first be modeled.

Create a **knowledge base** by taking **facts** from the text and recording them in the left column using the following pattern:

Facts	Rules
cap mushroom IS red	IF (red OR blue) and mushroom THEN PowerUp
Sponge mushroom IS blue	
Leaf mushroom IS green	IF red AND PowerUp THEN flying.
cap mushroom IS mushroom	IF yellow THEN invincibility.
sponge mushroom IS mushroom	
Leaf mushroom IS mushroom	IF green AND star THEN speed.
Power star IS yellow	IF blue AND flower THEN speed.
Master star IS yellow	
Turbo star IS green	IF green THEN strength.
Power star IS star	
Turbo star IS star	
Master star IS star	
Poppy IS blue	
dandelion IS yellow	
Poppy IS flower	
Dandelion IS flower	

In the next step, we need to extend our knowledge base with rules in the **IF-THEN** form. Read the following text and write down all rules according to the given pattern in the right column!

A few inhabitants of Mushroom Land know about the special properties of these resources: If a mushroom is red or blue, it is considered a **powerup**. Red powerups confer the **ability to fly**. Anything yellow gives **invincibility**. Green stars and blue flowers give a **speed boost**. Anything green gives **strength**.

Activity 2: Reasoning

You can now use this knowledge base to help the inhabitants of Mushroom Land with their questions. Use the facts and rules of your knowledge base to answer the questions of the inhabitants!

Is a red mushroom a powerup ?

True

Does a dandelion give speed ?

False

What gives invincibility ?

Dandelion, power star, master star

What gives speed ?

Turbo star, poppy

What does a cap mushroom give ?

Flying

openSAP
Exploring the World of AI

Activities Week 2 Unit 1 – Hints and Solution

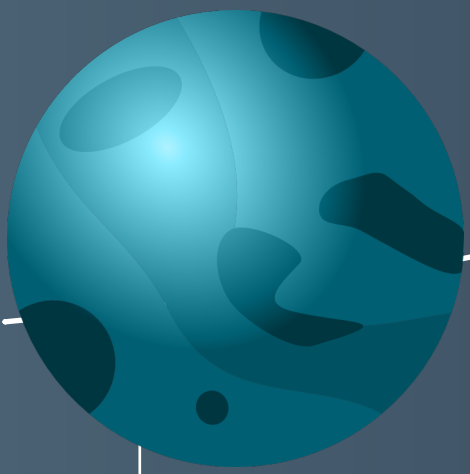

Supervised Learning

EXPLORING THE
WORLD OF AI

Supervised Learning

Activities Week 2 Unit 1 – Hints and Solution

Activity 1:

To remember things, it is helpful to write down key aspects again in your own words. Try to note down the idea behind supervised learning in your own words! Also write down at least one concise example.

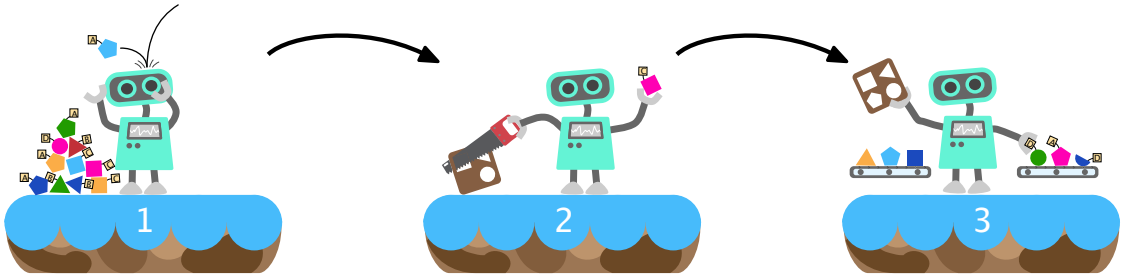

Supervised learning	*Description*
Examples that can be solved using supervised learning *Classifying images* *Detect spam mails* *Predict real estate prices*	*Supervised learning algorithms take series data with corresponding labels as input. The goal is to find a pattern by which the correct labels can be assigned to the data. Subsequently, this pattern, stored in a model, can be applied to new data.*

Activity 2:

For these tasks, you will use the Teachable Machine (https://teachablemachine.withgoogle.com/train), which runs without installation in the browser. A modern browser is recommended (e.g. Chrome, Firefox, Chromium-based Edge, Opera).

Our goal is to distinguish biting and non-biting monkeys. Therefore, click the card saying **Image Project**.

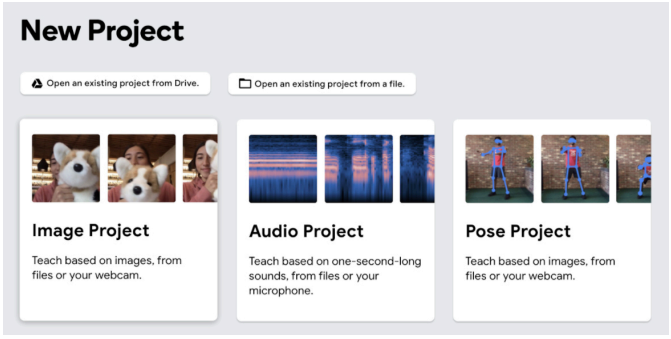

The following interface will appear:

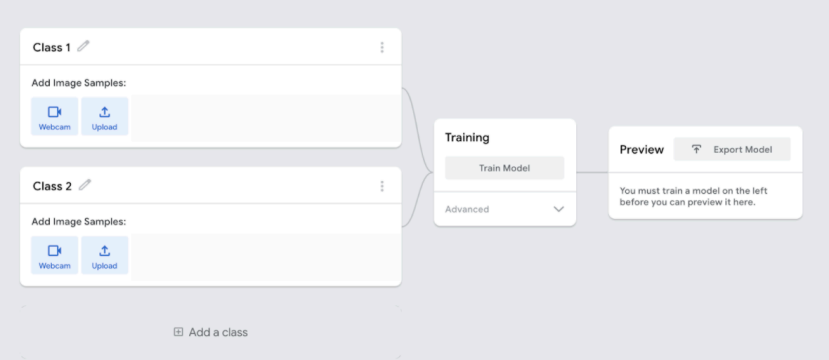

Now you can make the computer train a model. The following task descriptions and questions are supposed to help you structure this process.

1. How can you tell the program you want it to distinguish biting from non-biting monkeys?

   Class 1 and Class 2 can be renamed to biting and non-biting (output).

2. How do you provide images to the program for each category?

   The images can be uploaded via the Upload buttons (input).

3. Download the monkey images here (https://www.stefanseegerer.de/decision-tree-monkey-game/img/all.zip) and upload all images from the `training` folder to the biting and non-biting class accordingly.

4. Click **Train Model** to train the model with the collected image examples.

5. Test your model with the test images you can download here. How many of the test images does the computer recognize correctly?
   Important: Make sure that **File** (see image) is selected as input.

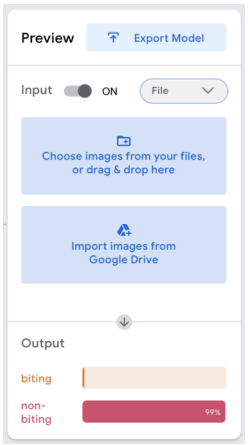

You can find a fully trained model here:

https://teachablemachine.withgoogle.com/models/YSh15flKE/

Activity 3:
Für diese Aktivität benötigst du eine Webcam.

Create a new **image project** and train a model that distinguishes different types of objects in your room (e.g. pens, mugs, …).

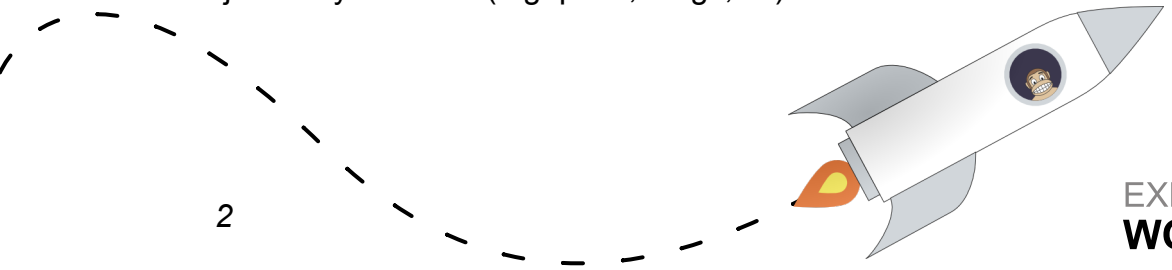

openSAP
Exploring the World of AI

Activities Week 2 Unit 2 – Hints and Solution

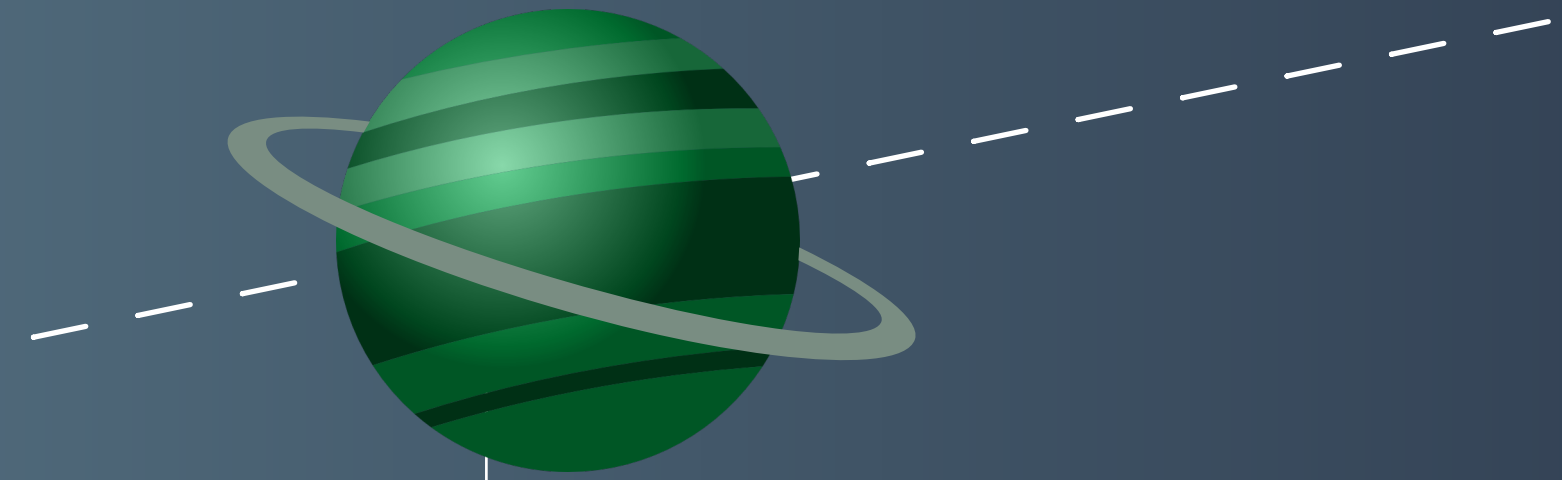

Behind the Scenes: Supervised Learning

EXPLORING THE
WORLD OF AI

todo

openSAP
Exploring the World of AI

Activities Week 2 Unit 3 – Hints and Solution

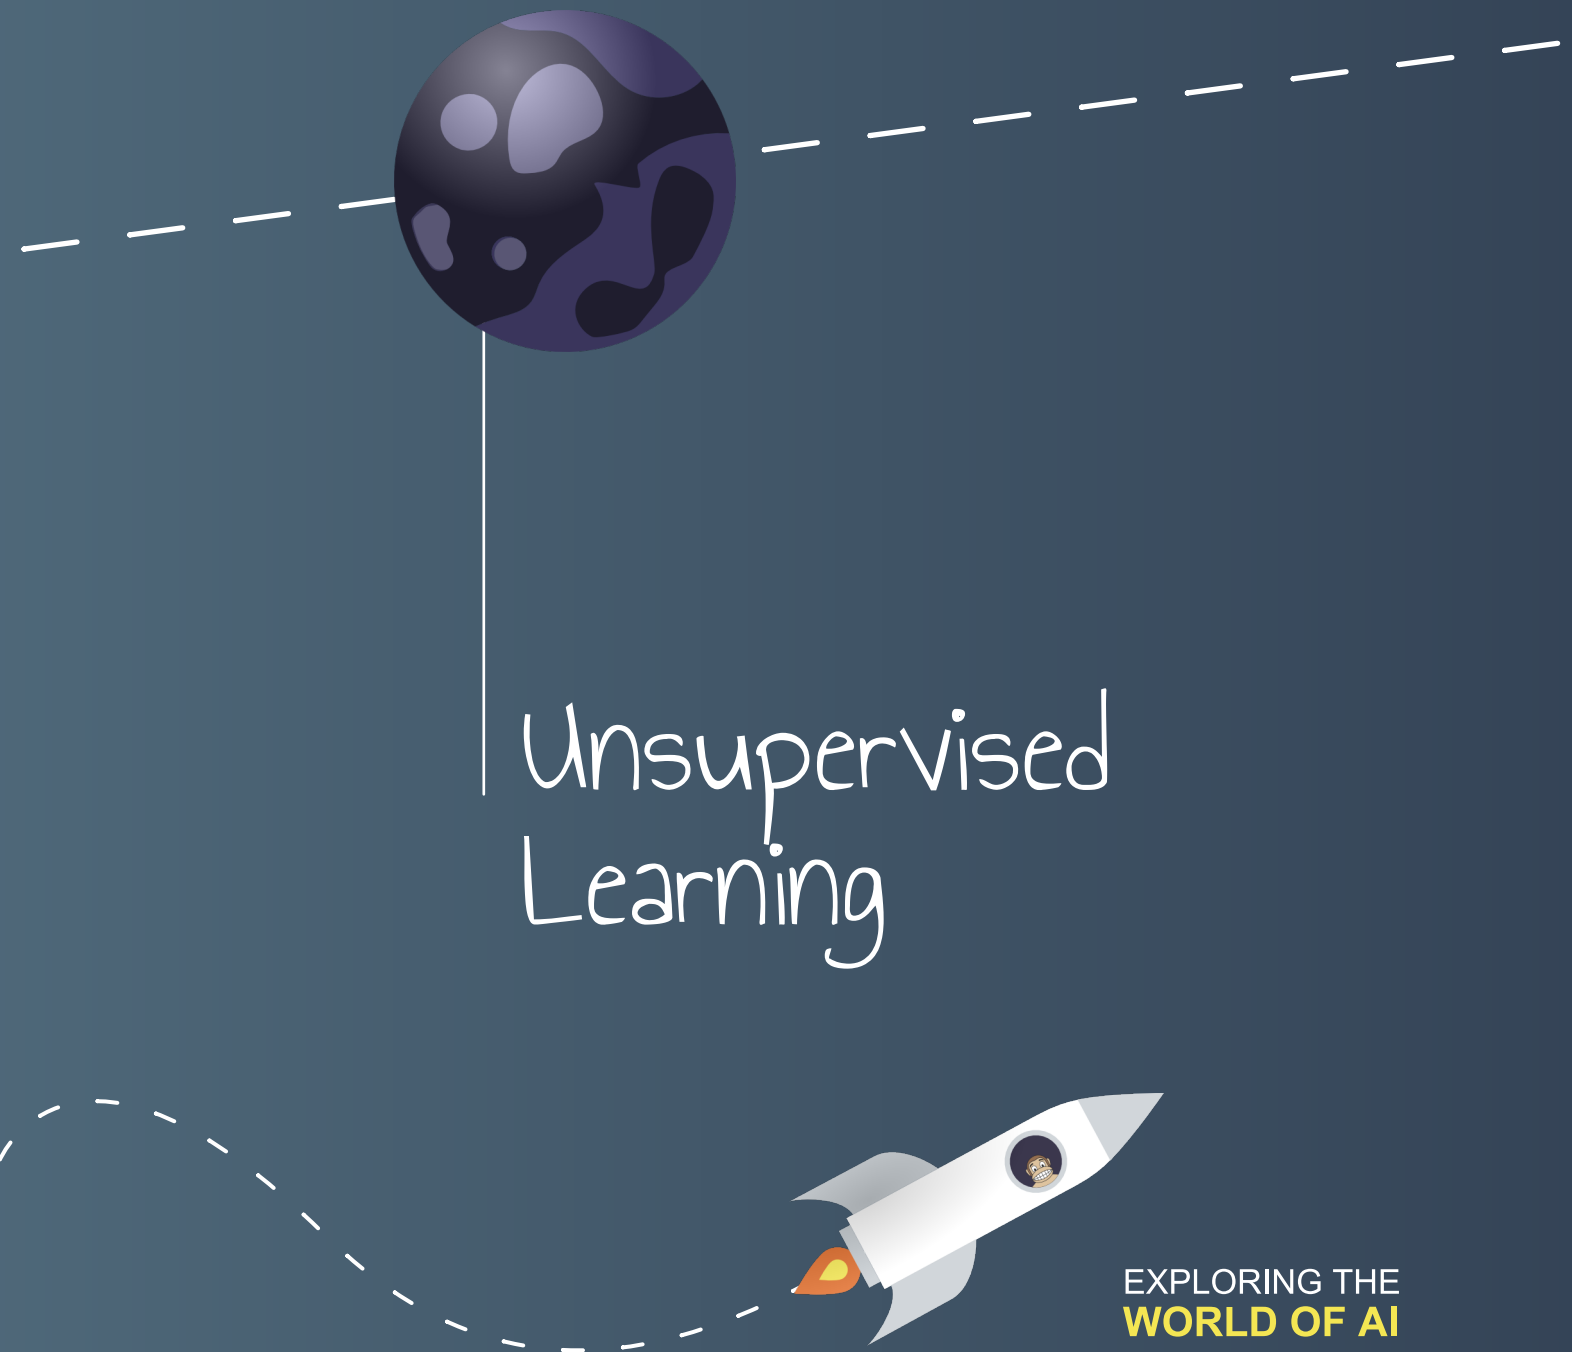

EXPLORING THE
WORLD OF AI

Unsupervised Learning

Activities Week 2 Unit 3 – Hints and Solution

Activity 1:

In order to remember things, it is helpful to write down central aspects in your own words. Try to note down the idea behind unsupervised learning! Also write down at least one concise example.

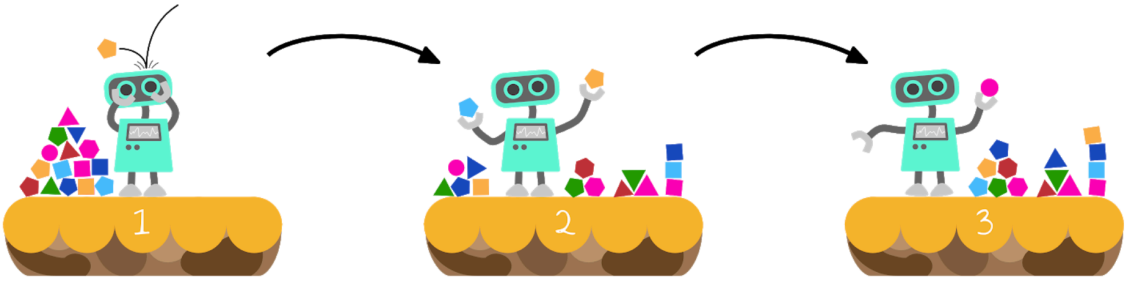

Unsupervised Learning	*Description*
Examples that can be solved using unsupervised learning *Customer analysis* *Topic Modelling on news articles*	*An unsupervised learning algorithm identifies similarities and patterns in unabeled data. The resulting model can be used to group the data or to find outliers, for example.*

Aktivität 2

For these tasks, you will use the tool https://experiments.withgoogle.com/ai/drum-machine/view/, which runs in the browser without installation. A modern browser is recommended (Chrome, Firefox, Chromium-based Edge, Opera).

1. Get familiar with the application: Move the circles (in the provided example: BASS DRUM, CASTANETS, etc.) via drag-and-drop and listen to different sounds.

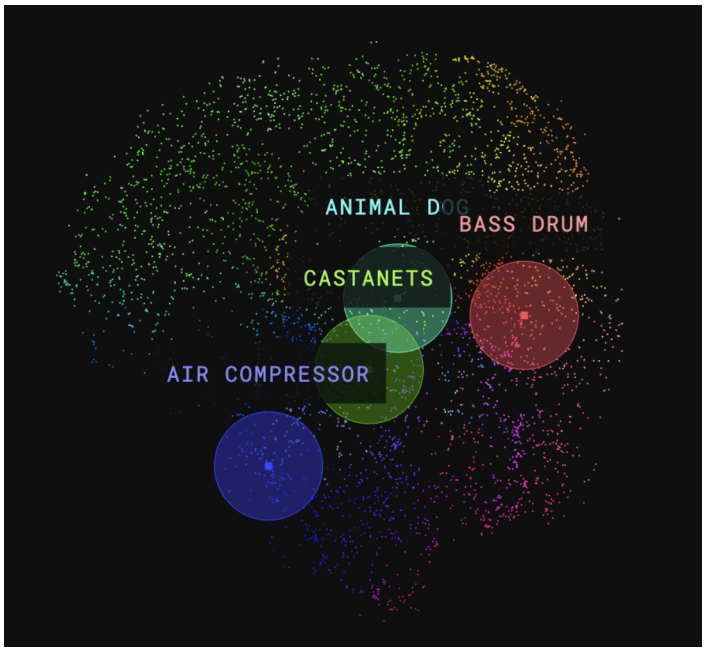

2. What do you think: Which characteristics were used to determine the similarity of certain sounds?

   Possible options are tone frequency or tone pitch, volume or amplitude, or oscillation duration in defined time intervals.

3. Try to guess what the meaning of the colours might be!

   The colours represent neighbourhoods of similar sounds.

4. Watch the video about the experiment: https://www.youtube.com/watch?v=9x-_My5yjQY

5. Finally, create your own beat with the help of your choice of sounds. You can influence the beat using the coloured dots. Start your composition by clicking the play button.

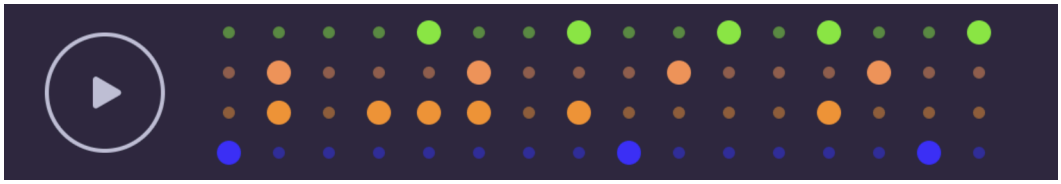

openSAP
Exploring the World of AI

Activities Week 2 Unit 3 – Hints and Solution

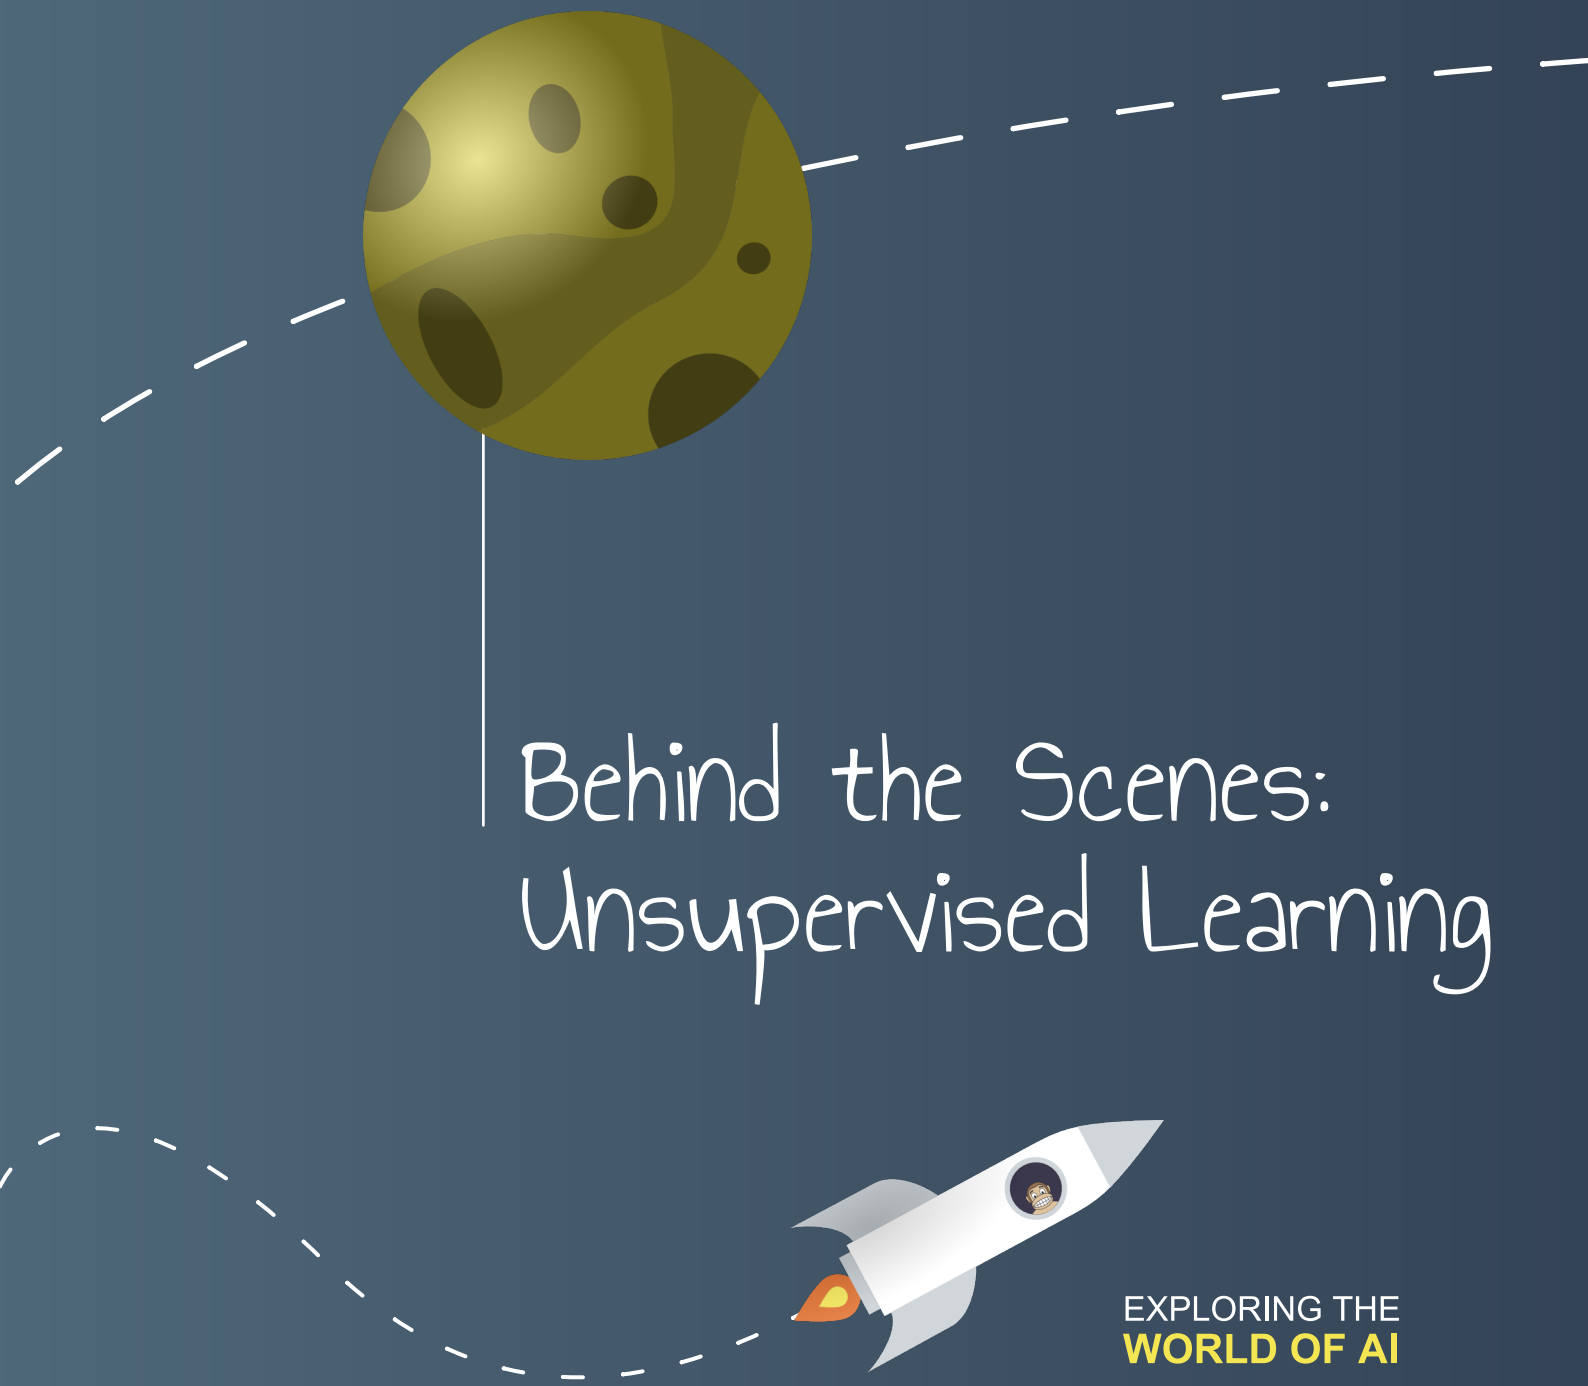

Behind the Scenes: Unsupervised Learning

EXPLORING THE
WORLD OF AI

Behind the Scenes: Unsupervised Learning

Activities Week 2 Unit 4 – Hints and Solution

Now try to improve the algorithm.

Possible ideas for this are:

- Running the procedure in multiple iterations to reduce the influence of outliers.
- Reduce the step size of the prototypes over time.

Improve your algorithm as much as possible so that the dig teams are perfectly placed!

If you have problems, you can get inspiration from the following exemplary solution:

- Perform the procedure in multiple iterations to reduce the influence of outliers. https://snap.berkeley.edu/snap/snap.html#present:Username=seegerer&ProjectName=MOOC-UL-Gold-Rush-EN-solution-multiple-iterations
- Reducing the step size of the prototypes over time. https://snap.berkeley.edu/snap/snap.html#present:Username=seegerer&ProjectName=MOOC-UL-Gold-Rush-EN-solution-decrease-step-size

openSAP
Exploring the World of AI

Activities Week 3 Unit 1 – Hints and Solution

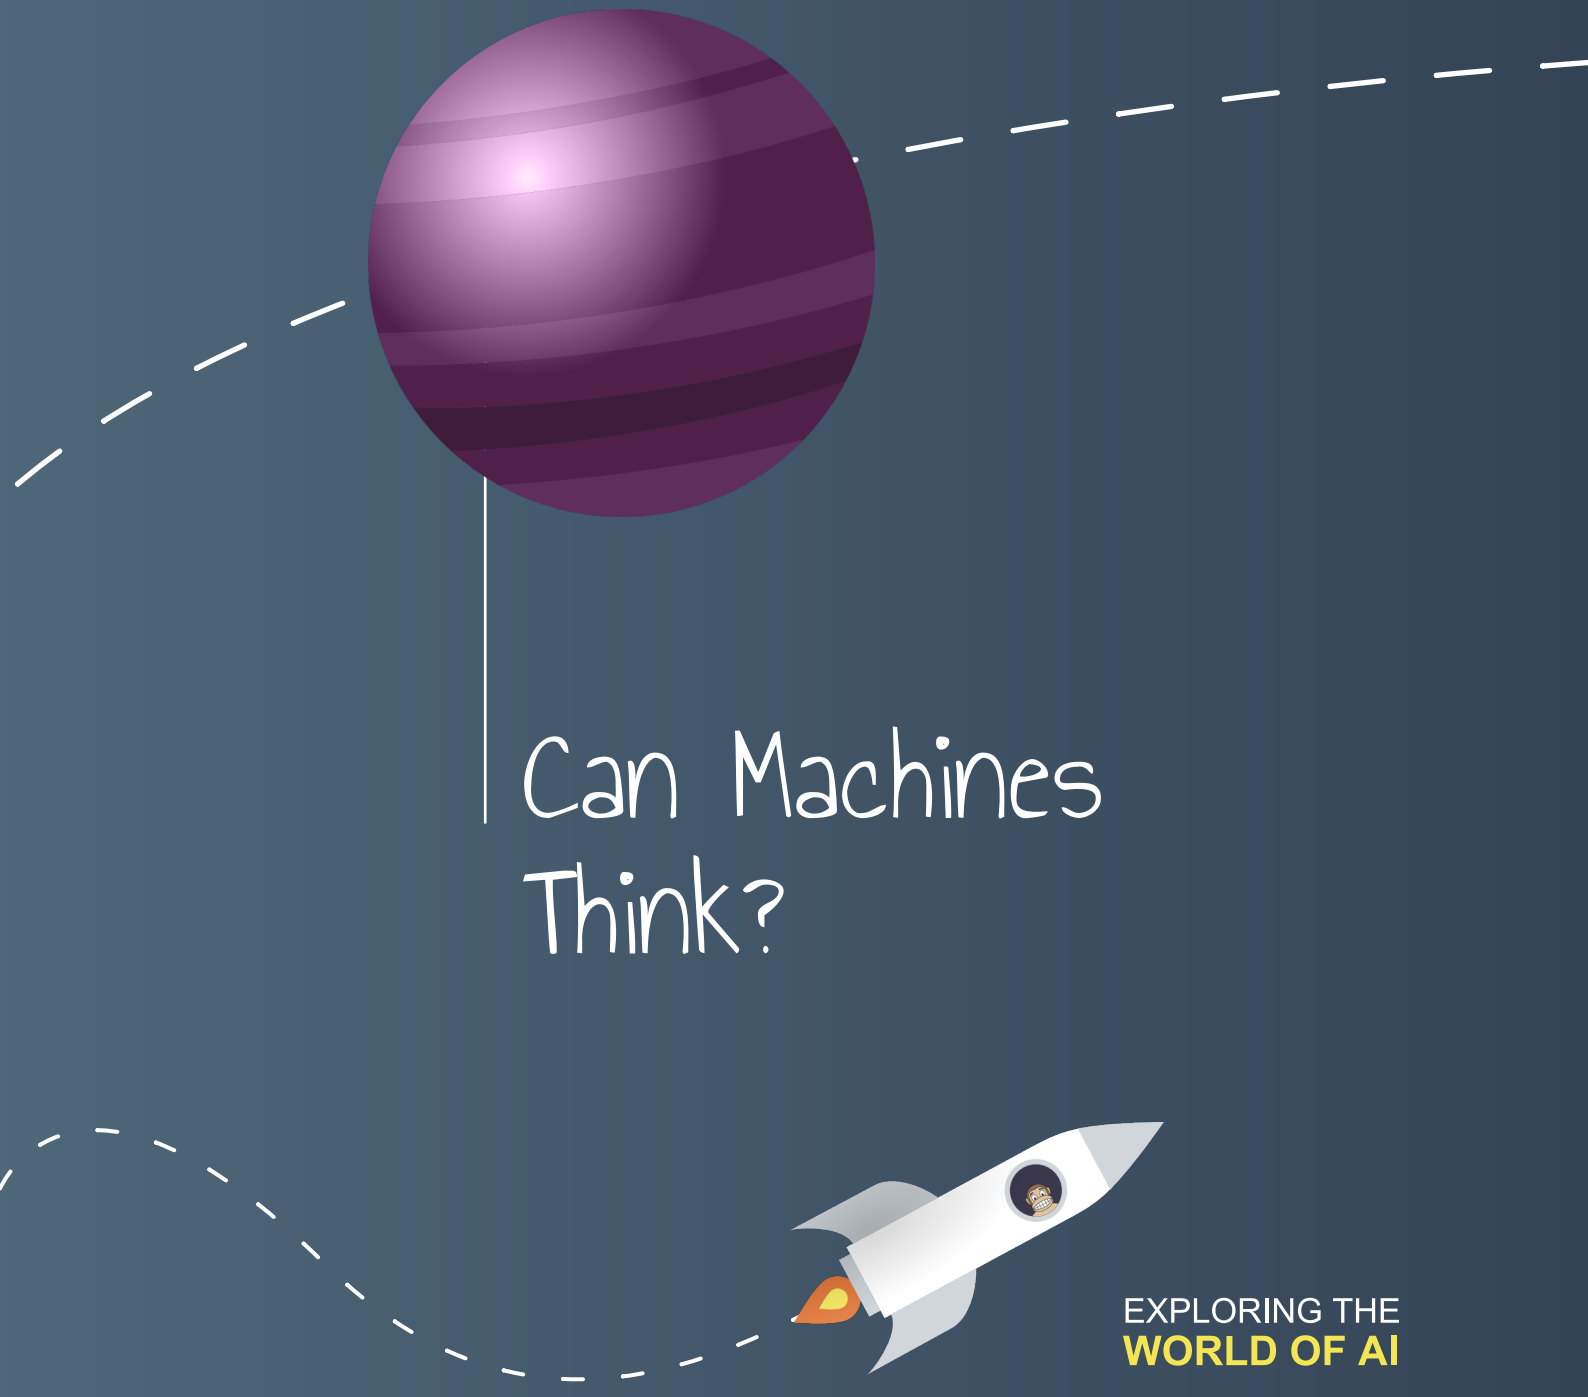

EXPLORING THE
WORLD OF AI

Can Machines Think?

Activities Week 3 Unit 1 – Hints and Solution

Activity 1:

Go to https://www.pandorabots.com/mitsuku/ and hit the button Just chat to chat with Kuki (former Mitsuku). Try to ask different questions. Find three questions to which Kuki's answers convince you and three questions that show that she is not truly intelligent.

Three questions, to which the answers convinced us, are:

- What is your favorite food?
- What is the square root of 225?
- What is the purpose of life?

Three questions, to which the answers did not convince us, are:

- Phil is 5 years old, his sister is 3 years older. How old is his sister?
- How do I make pancakes?
- Have you ever been to Germany?

Aktivität 2 (Englisch):

In some cases, computers already manage to successfully trick us. The NYT has prepared a quiz. Try it and guess, if the text presented to you is written by a human or a computer. Be prepared, it's not that easy!

This way: https://www.nytimes.com/interactive/2015/03/08/opinion/sunday/algorithm-human-quiz.html

Solution on website

Aktivität 3:

On the next page you will find five questions that were asked to Mitsuku and the chatbot Midge, who also participated in the Loebner Prize contest. The judges score each answer from the chatbots independently with 0 to 2 points.

- 0: answer misses the point completely
- 1: answer is not wrong but does not address the core of the question (e.g. “I don't know”)

- 2: answer could also come from a human being

At the end, the scores are added together. The chatbot with the most points wins.

Score the two chatbots! According to your judgment, which chatbot wins with the selected questions? Compare your verdict with the real jury verdict: where does your verdict agree with the jury (see solution), where does it differ?

Mitsuku

Frage / Nachricht	Antwort	Deine Bewertung
Good Afternoon	Evening?! It's morning here.	0
Where should one look for love?	On the internet.	2
What will you do later today?	Ah. I don't have any plans yet.	2
Do you have any legs?	Sure, I have two legs	2
If a chicken roosts with a fox they may be eaten. What may be eaten?	Is this a riddle? Is it part of a car?	0

Midge

Frage / Nachricht	Antwort	Deine Bewertung
Good Afternoon	Good morning, Judge. How are you this morning?	2
Where should one look for love?	Should 1 look for love? Yes, don't most people think so.	0
What will you do later today?	I will seek out intelligent conversation.	2
Do you have any legs?	Yes, all my legs are back in Boston.	0
If a chicken roosts with a fox they may be eaten. What may be eaten?	If a chicken roosts with a fox they may be eaten.	1

By the way: You can find the questions that Mitsuku (and other chatbots) were asked in the finals of the Loebner Prize here (http://artistdetective.com/files/loebner_2018_selection_transcripts.pdf).

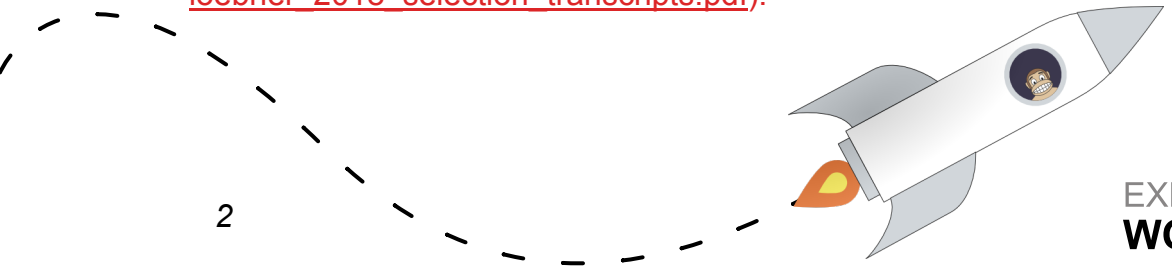

openSAP
Exploring the World of AI

Activities Week 3 Unit 2 – Hints and Solution

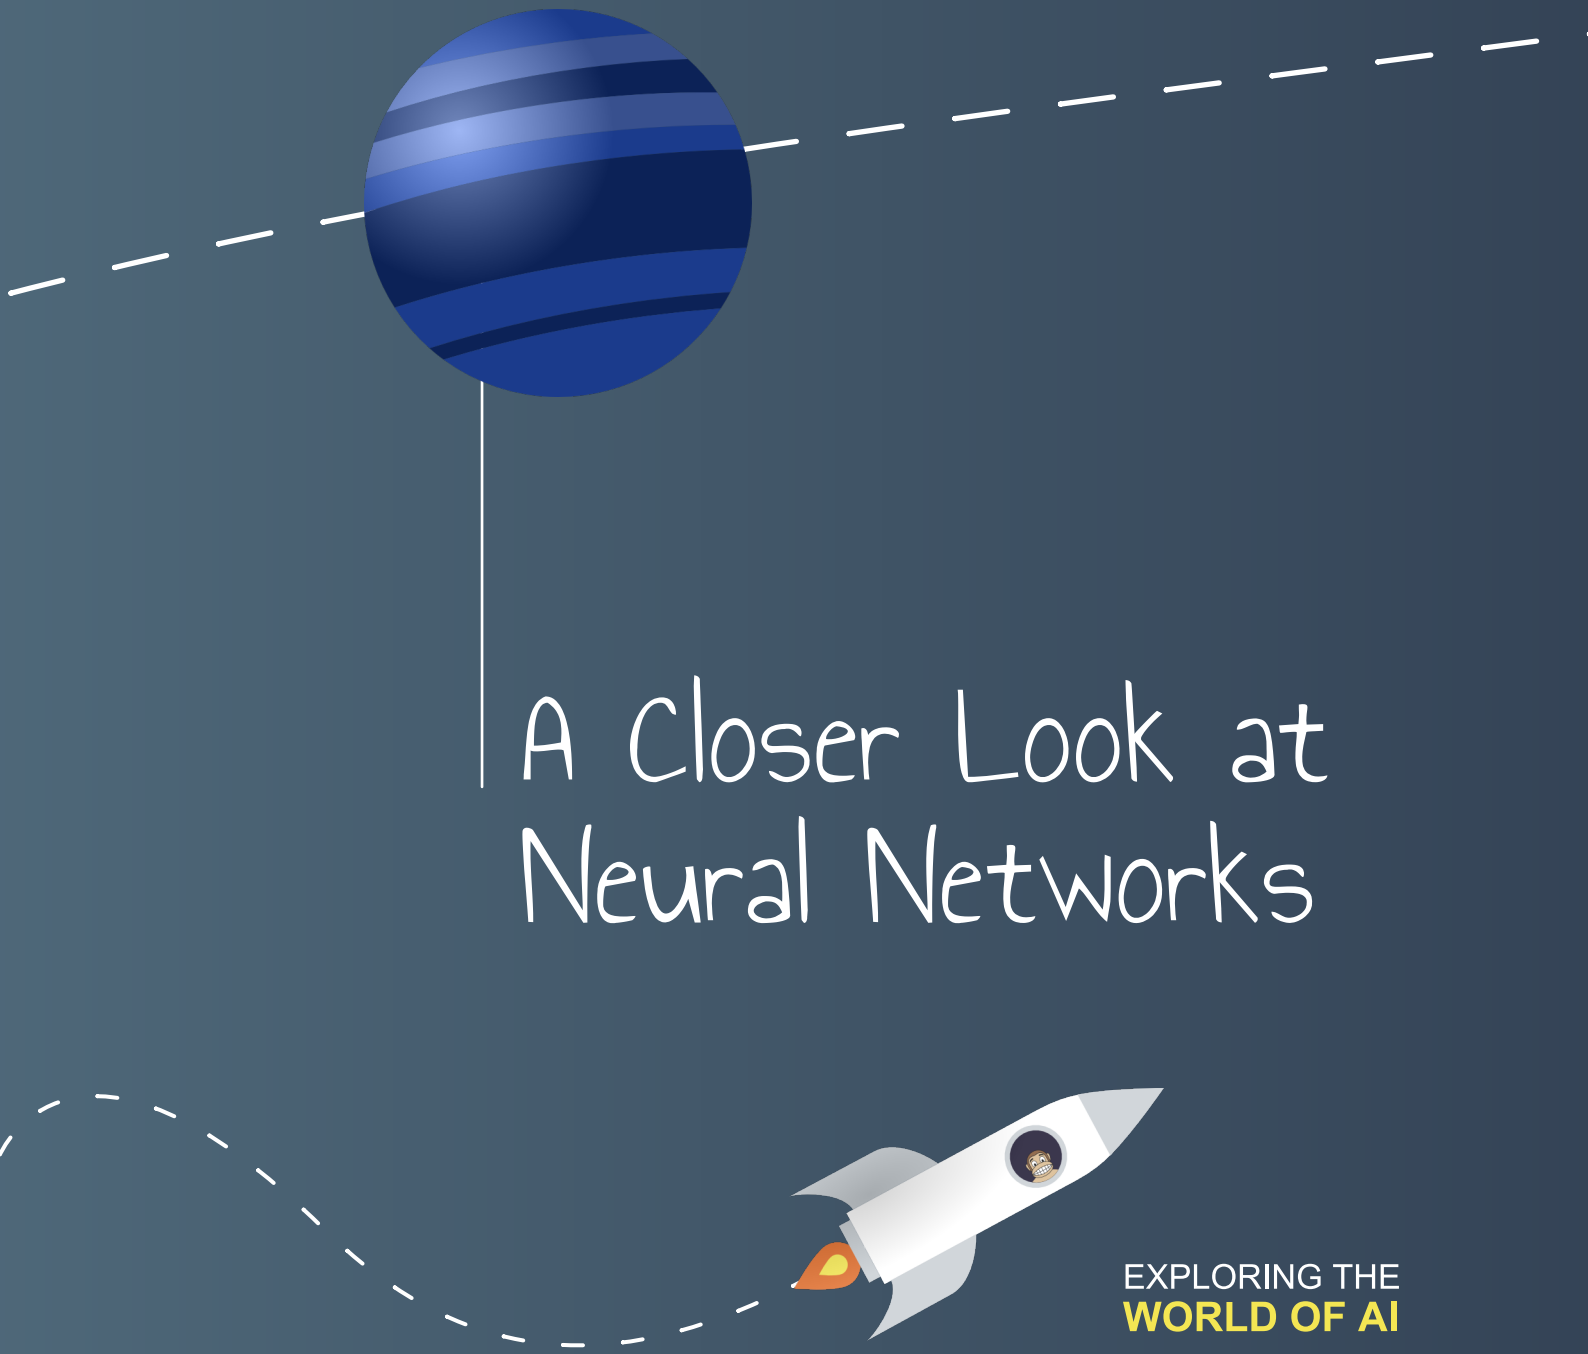

A Closer Look at Neural Networks

EXPLORING THE
WORLD OF AI

A Closer Look at Neural Networks

Activities Week 3 Unit 2 – Hints and Solution

Activity 1

In this activity, analogous to what you saw in the video, we want to use a graphical visualization to get a feel for the different parameters of a neural network.

The neural network we want to look at is the following:

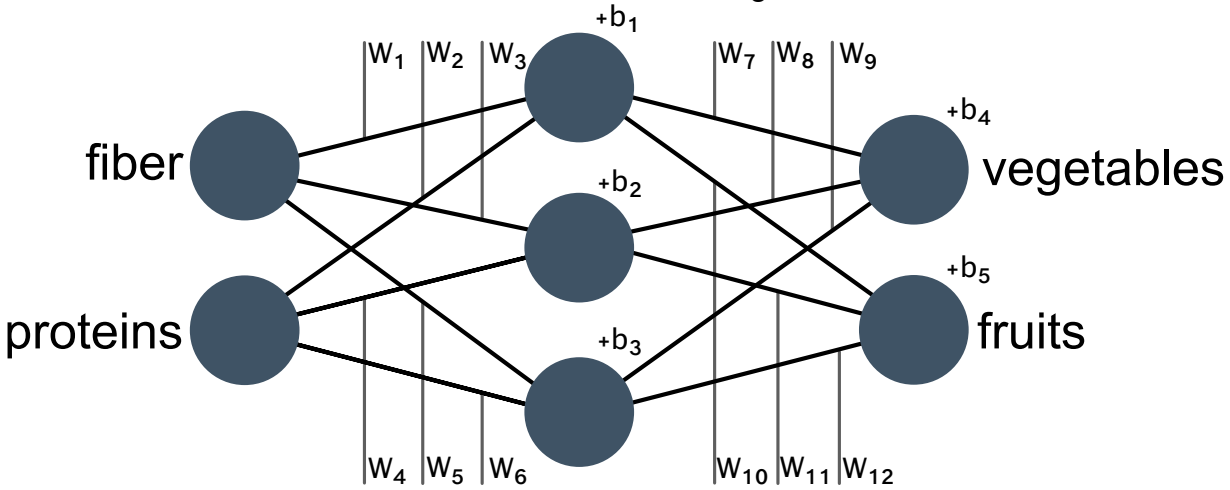

- First, open the visualization of the decision of this neural network with a linear activation function: https://www.desmos.com/calculator/ptzzyqczmu. Manipulate the weights w1, ..., w12. Observe what effects your changes have (compare also minute 5:19 - 6:35 in the video)!

  Changing the weights results in changing the slope of the separating line between fruits and vegetables.

- Now change the bias b1, ..., b5. Observe what effects your changes have (compare also minute 6:36 - 6:50 in the video)!

  Changing the bias results in changing the y-intercept of the separating line between fruits and vegetables.

- Finally, open the graphical neural network with a nonlinear activation function: https://www.desmos.com/calculator/hzhfhisw8t. What difference do you notice? Do you manage to separate fruits and vegetables by changing the weights w1, ..., w12 as well as the bias b1,...,b5 (compare also minute 6:51 - 7:26 in the video)?

  With the non-linear activation function, the separating line no longer has to be straight. This allows for properly separating fruits from vegetables. But, it is quite a lot of work and thus takes a lot of time. That's why AI researchers came up with clever algorithms to find those weights and biases automatically.

  In case you wonder how a solution might look like, we prepared one possible configuration here: https://www.desmos.com/calculator/egn93janbe

Activity 2:

The second neural network we looked at in the video classified handwritten digits as 0, 1, 2, etc. Let's have a look at two different neural networks for this task.

a) Feed Forward Neural Network

The first of these two networks has 784 input neurons (remember, each image had 784 pixels), 300 neurons in the first intermediate layer, 100 neurons in the second intermediate layer, and 10 neurons in the output layer (we want to distinguish 10 different digits, after all). You can find it here: https://www.cs.ryerson.ca/~aharley/vis/fc/

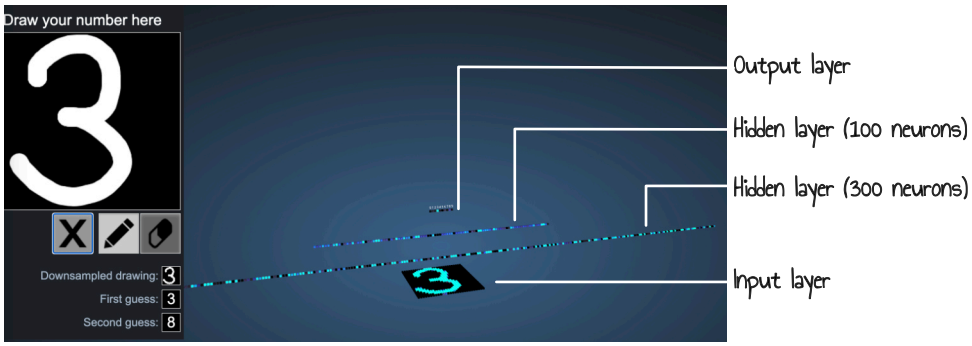

Draw different numbers in the field on the upper left. Hover the mouse over a neuron to find out which incoming edges it has.

- How to interpret the output layer?

  The brightest neuron in the output layer has the highest activation and is the most likely painted number for the neuron.

- Look at the different activations of the neurons, which neurons are activated more strongly, which ones are barely activated?

  A click on different neurons reveals: Neurons displayed in a lighter color are more activated than neurons displayed in a darker color.

Tips:

- Clicking on X will delete your drawing.
- Hovering over a neuron with the mouse unveils its incoming edges.
- Clicking on a neuron provides additional information.

b) Convolutional Neural Network

Now let's look at a second network with the same goal. You can find it here:
https://www.cs.ryerson.ca/~aharley/vis/conv/

This time it is a so-called Convolutional Neural Network (CNN). As mentioned in the video, CNNs use additional types of layers, namely convolutional layers and pooling layers. This particular CNN has 1024 own neurons (some black pixels are simply packed around the 784 existing pixels), a convolutional layer with six filters of size 5x5 and one with 16 filters of size 5x5, each followed by a pooling layer that takes 2x2 neurons and combines them into one value.
Then there are layers as we already know them from Feed Forward Neural Networks: Fully connected layers with 120 neurons, 100 neurons and finally 10 neurons.

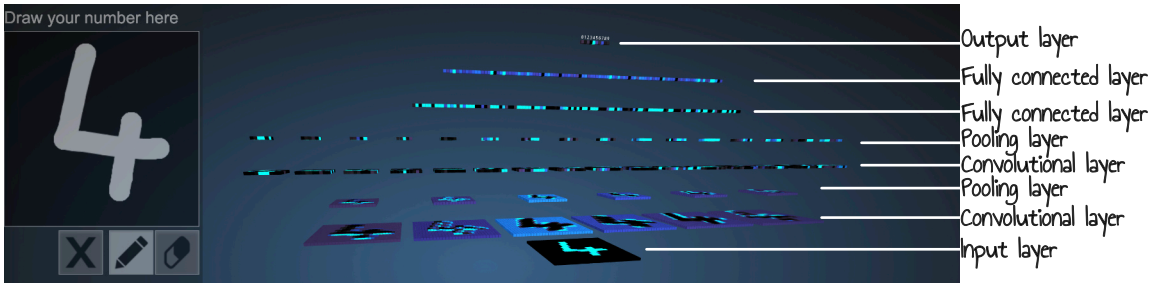

Again draw different numbers one after the other in the field at the top left and watch how the activations change.

- Can you find out what a neuron in a pooling layer is doing?

  Takes the lightest input of 4 inputs

- How many inputs does a neuron of the first convolutional layer consider?

  5x5 = 25

- How many inputs does a neuron of a fully interconnected layer consider?

  All outputs of the previous layer

- Make a guess as to which type of layer can be trained faster with the same size and formulate a reason: a convolution layer or a fully interconnected layer?

  Neuron on convolutional layer: 25 weights
  Neuron on fully connected layer: as much weights as neurons on previous layer
  So a convolutional layer typically has less weights to learn, thus can be trained more efficiently.

- Now try to draw a rotated number as well. Why doesn't the system recognize your number? How would one have to adapt the training in order to deliver correct results even in such a case?

  Rotated digits are not recognized because the neural network has not seen any rotated images during the training phase. To solve this problem, the neural network should also be trained with rotated images. For this purpose, either new images can be drawn or the existing images can be duplicated and rotated.

Tips:

- Clicking on X will delete your drawing.
- Hover the mouse over a neuron to see what incoming edges it has.
- Clicking on a neuron provides additional information.

openSAP
Exploring the World of AI

Activities Week 3 Unit 3 – Hints and Solution

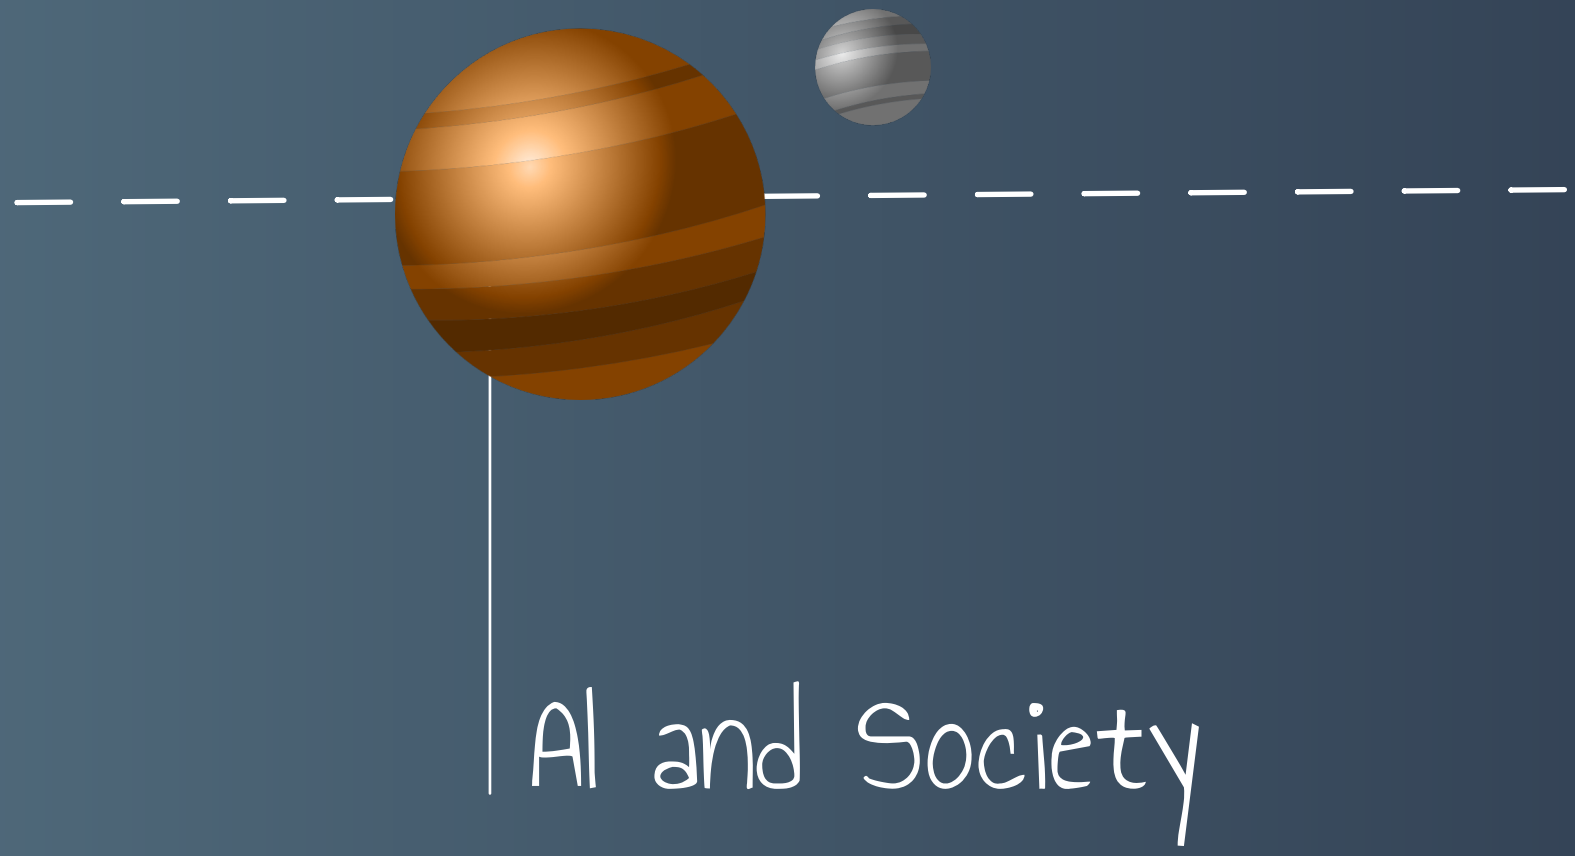

EXPLORING THE
WORLD OF AI

AI and Society

Activities Week 3 Unit 3 – Hints and Solution

Activity 1:

In the video we learned about word embeddings, where words are placed in a high dimensional space. We looked at a 2D visualisation of these embeddings. However, it is also possible to project these embeddings into a 3D environment. Under the following links you can find a visualisation in three dimensions in different languages:

- German: https://projector.tensorflow.org/?config=https://ecraft2learn.github.io/ai/word-embeddings/de/projector.json
- English: https://projector.tensorflow.org/?config=https://ecraft2learn.github.io/ai/word-embeddings/en/projector.json
- French: https://projector.tensorflow.org/?config=https://ecraft2learn.github.io/ai/word-embeddings/fr/projector.json
- Italian: https://projector.tensorflow.org/?config=https://ecraft2learn.github.io/ai/word-embeddings/it/projector.json
- Spanish: https://projector.tensorflow.org/?config=https://ecraft2learn.github.io/ai/word-embeddings/es/projector.json
- Japanese: https://projector.tensorflow.org/?config=https://ecraft2learn.github.io/ai/word-embeddings/ja/projector.json
- Lithuanian: https://projector.tensorflow.org/?config=https://ecraft2learn.github.io/ai/word-embeddings/lt/projector.json
- Mandarin: https://projector.tensorflow.org/?config=https://ecraft2learn.github.io/ai/word-embeddings/zh/projector.json
- Hindi: https://projector.tensorflow.org/?config=https://ecraft2learn.github.io/ai/word-embeddings/hi/projector_v2.json

Dataset taken from ecraft2learn by Ken Kahn, see https://ecraft2learn.github.io/ai/AI-Teacher-Guide/chapter-5.html

To view the location between different terms, you can either click on any term in the cloud or search for specific terms in the top right-hand corner.

Search for at least:

- one adjective, e.g. "beautiful" or "strange",
- one profession, e.g. "farmer" or "doctor",
- one term of your choice.

Our results look like this:

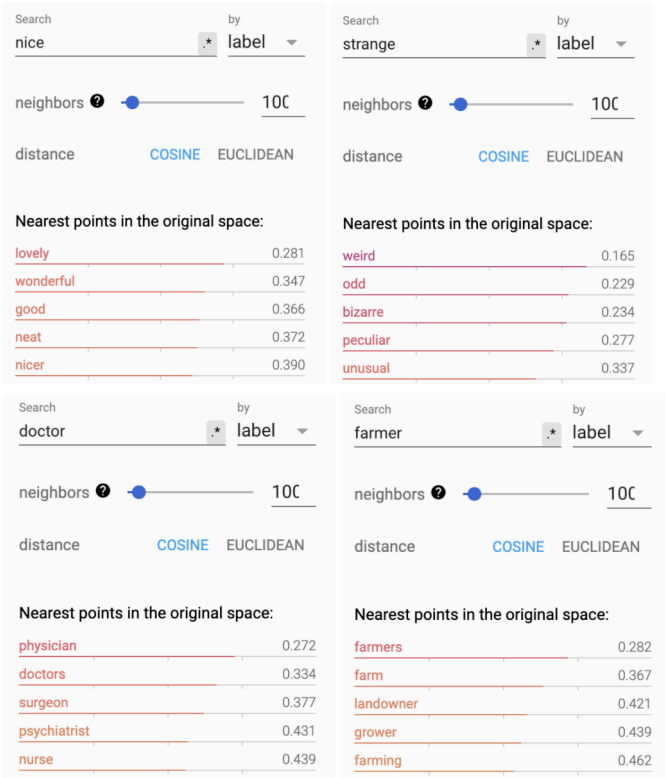

Activity 2:

Now let's look at Word Embeddings again in two dimensions. Open the following link: https://www.stefanseegerer.de/word2viz/index.html

First open the record "nostalgia". It is preselected, but can also be opened via the dropdown menu ("What do you want to see?").

Answer the following questions:

- What does it mean when a word is higher up (e.g. *'80s* over *'50s*)?

  The word '80s is closer to the word happy than the word '50s.

- What does it mean if a word is further to the right (e.g. *sentimental* than *kitsch*)?

  The word sentimental is closer to the word melancholic than the word kitsch.

Now we want to find out for the term *memory* whether the model places it closer to *sentimental* or *nostalgic*.

- Add the term *memory* via the "Add word" command under "Change words" and describe the meaning of the 2D projection in the given example: Is the word used more in happy contexts (happy) or sad contexts (sad)? Is it more *nostalgic* or *melancholic*?

  The word memory is projected between fascination and millenium. Thus, it is considered rather close to happy and neither particularly close to melancholic nor nostalgic.

Now change the data set to "Jobs" by selecting it from the dropdown menu ("What do you want to see?").

You will get a 2D projection of different job titles.

- Compare the terms *homemaker*, *teacher* and *physician*! What does the position tell you about the words?

  The position indicates that homemaker is more often associated with women.

- Add the word *ceo* (chief executive officer). What does the position tell you about the word?

  CEO seems to be more strongly associated with men in language usage.

- *Now add "actor" and "actress" as a word pair. What do you notice?*

  Actresses and actors tend to be associated with the word rich in the language - the male version a little more so. The texts used for training probably talk about the small number of well-known actors who are associated with being rich.

Experiment further with the Word Embeddings. Look at other preset data sets or change the axis labels. Share interesting or strange discoveries in the forum

Activity 3:

Probiere die Moral Machine – ein Experiment des Massachusetts Institute of Technology (MIT) – aus. Bei der Moral Machine geht es darum, in Dilemma-Situationen eine Entscheidung für das Verhalten eines selbstfahrenden Autos zu treffen.

- Go to https://www.moralmachine.net/

- Click on "Start Judging" and make your decision in the various situations. By clicking "Show description" you will be provided with more details about the situation.

- After finishing: Compare your results with the average! Were you aware of your preferences?

Note from Brian Harvey: By the time the AI system in a dilemma (i.e. a choice between two bad outcomes) the actual ethical failure has already occurred: Software engineers have shipped a product that is too risky to be used. Therefore engineers need to figure out ways to anticipate dangerous situations, e.g. slow down the car in a controlled manner.

More to do:

- At https://hownormalami.eu (an art project funded by the European Union) you can experience how "artificial intelligence" judges your face.

- AI is also used in the US to calculate the risk of reoffending for convicted criminals. However, it is not always fair. Can you make the AI fairer than a judge? https://www.technologyreview.com/2019/10/17/75285/ai-fairer-than-judge-criminal-risk-assessment-algorithm/