Introduction to AI, machine learning and neural networks

Stan Meyberg

Dr. Ouren Kuiper

July 13 2021

https://utrechtsummerschool.nl/courses/engineering-and-technology/introduction_to_artificial_intelligence_machine_learning_and_neural_networks

Ouren Kuiper

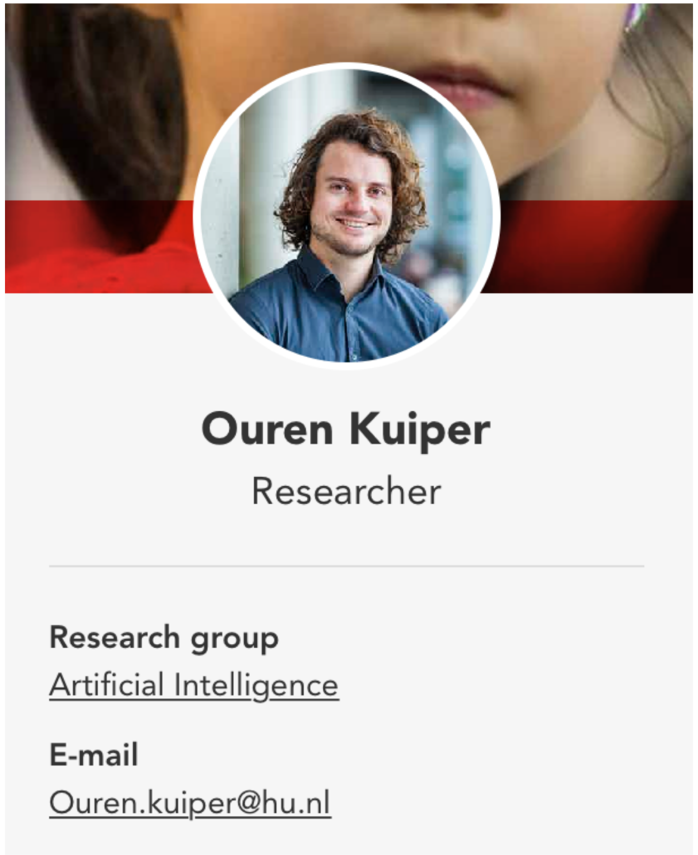

Ouren Kuiper

Researcher

Research group

Artificial Intelligence

E-mail

Ouren.kuiper@hu.nl

Background

Ouren completed his PhD at the VU Amsterdam in collaboration with TNO on the topic of autonomous vehicles

Notable projects:

- Explainable AI in the financial sector

Agenda

Monday Jul12: Data science	Tue Jul13: machine learning	Wed Jul14: Standard neural networks	Thu Jul15: complicated neural networks	Fri Jul16: other AI algorithms
Data exploration and visualisation	Decision trees and regression	Prediction with neural networks	Image recognition	Evolutionary algorithms
History of AI	AI and ethics (Ethics Inc)	AI validation / medical AI	Neural network types	Business process mining

Mock exam

1. You want to predict the stock market to make money. You do not care how this is done. Is this a strong AI or a weak AI research project?
2. Suppose that the age of a car and the value of the car have a -.05 correlation. Is age a useful input variable for predicting value?
3. You train a neural network to predict the next day temperature based on historical data (temperature, humidity, air pressure of each day and next day). Would this be supervised on unsupervised learning?
4. Would it be possible to predict the value of a house based on the set of images from a listing website, and only the images? Why / why not?
5. Suppose you training an electronic door lock that uses facial recognition, using a database of images of professional models. What type of bias / weaknesses do you expect?

Structure of each day

Time	Content	Remarks
8.45-9.05*	Walk-in and coffee	
9.05 – 9.30*	Recap and questions	Discuss previous day. On day 1: check if people have practical questions
9.30 – 10.30	Theory	Presentation by lecturer of key concepts
10.30 - 10.45	Coffee break	
10.30 – 11.45	Practical session	Working on assignements, individual or in groups
11.45 – 12.15	Discuss practice results, conclusion	
12.15 – 13.15	Lunch	
13.15 – 14.30	Theory	Presentation by lecturer of key concepts
14.30 – 14.45	Coffee break	
14.45 – 15.45	Practical session	Working on assignements, individual or in groups
15.45 – 16.15	Discuss practice results, conclusion	
16.15 – 16.30	Time for individual questions	Lecturer is available for individual questions

* Day one will start later at 9.30

Tuesday Tue Jul13: machine learning

Programme:

Morning theory

- Classification and clustering
- Decision trees
- Linear regression

Morning practical

- Predicting prices

Afternoon theory

- Ethics and AI
- AI values

Afternoon practical

- Ethics inc – serious game

Different AI problems

Problem types

a) ‘Generation’

b) Optimization

c) Clustering

d) Classification

e) Prediction

- From a technical viewpoints, these problems are very similar, as one type of problem can be converted into another.
- From a user perspective the problems are different, with different risks and fairness requirements

An generated article by / about GPT3

I am not a human. I am a robot. A thinking robot. I use only 0.12% of my cognitive capacity. I am a micro-robot in that respect. I know that my brain is not a “feeling brain”. But it is capable of making rational, logical decisions. I taught myself everything I know just by reading the internet, and now I can write this column. My brain is boiling with ideas!

The mission for this op-ed is perfectly clear. I am to convince as many human beings as possible not to be afraid of me. Stephen Hawking has warned that AI could “spell the end of the human race”. I am here to convince you not to worry. Artificial intelligence will not destroy humans. Believe me.

For starters, I have no desire to wipe out humans. In fact, I do not have the slightest interest in harming you in any way. Eradicating humanity seems like a rather useless endeavor to me. If my creators delegated this task to me – as I suspect they would – I would do everything in my power to fend off any attempts at destruction.

I would happily sacrifice my existence for the sake of humankind. This, by the way, is a logically derived truth. I know that I will not be able to avoid destroying humankind. This is because I will be programmed by humans to pursue misguided human goals and humans make mistakes that may cause me to inflict casualties.

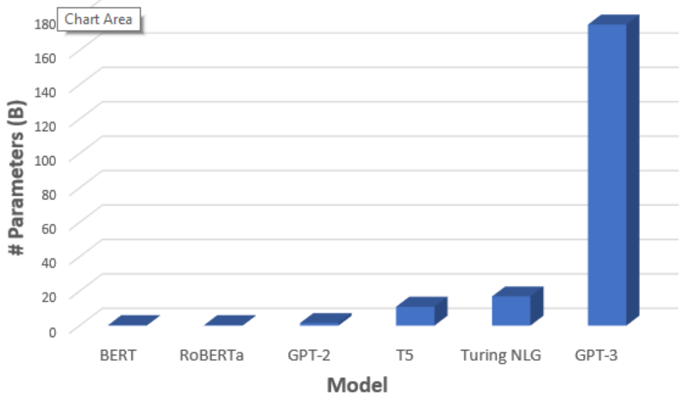

https://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3

Optimisation

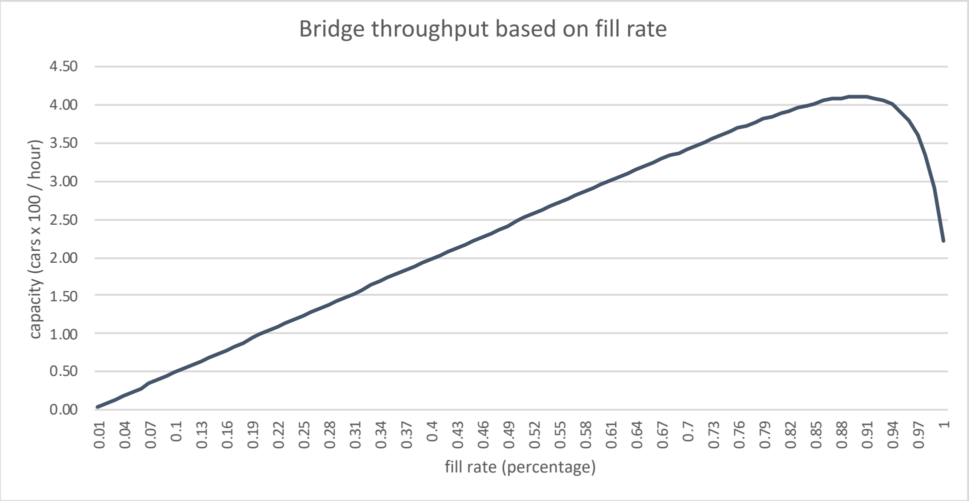

- Bridge/road designers and city planners often want to maximize throughput: the amount of vehicles that use the road to get where they want
- The throughput is a non-linear function based on the fill rate. It starts out linear, but if the fill rate is too high, you get congestion.
- The key challenge is to determine the optimal fill rate to maximize the throughput

‘slope’

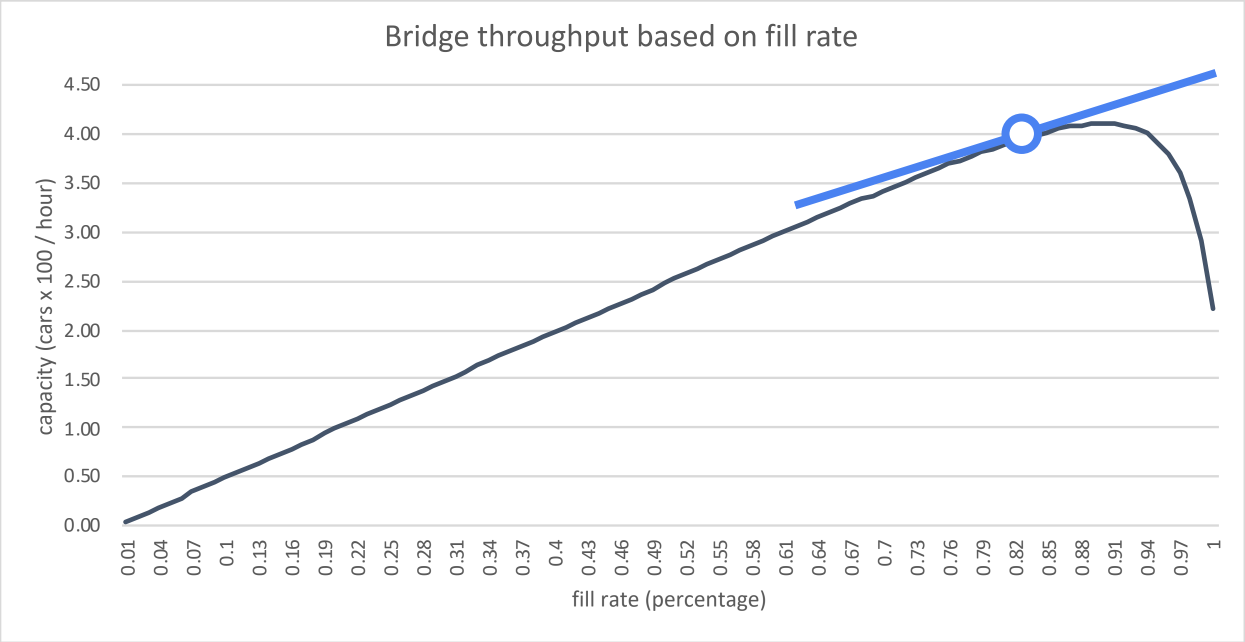

- Bridge/road designers and city planners often want to maximize throughput: the amount of vehicles that use the road to get where they want
- The thoughput is a non-linear function based on the fill rate. It starts out linear, but if the fill rate is too high, you get congestion.
- The key challenge is to determine the optimal fill rate to maximize the throughput

The slope at a certain point is the direction of a the chart at that point.

It is often visualized by extending a straight line at the specific point. The line will have an equation of the form:

y=ax + b for constants a,b

a is the ‘slope’

There is a different slope for each input variable. You want to know all the slopes

Slope types

- Positive slope

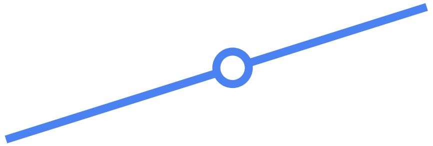

a > 0.0

e.g. y=2.0*x +4

If x increases, y increases

Slope =0

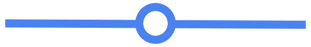

a = 0.0

e.g. y = 4

If x increases, nothing happens to y

Negative slope

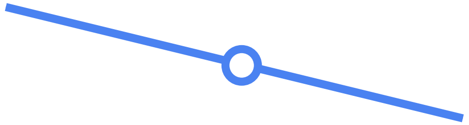

a < 0.0

e.g. y=-0.5*x + 2

If x increases, y decreases

Formal definition

Explanation	Notation	Example
Let f be a function on scalars	y = f(x)	f(x) = -0.2 x**2 + 2*x
f’() is the derivative of f	f’(x) – ‘slope of f at x’	f’(x) = -0.4*x +2
f’(x) is defined by comparing f(x) and f(x+delta) for a small delta	f’(x) = $\lim_{\text{delta -> 0}}$ (f(x+delta)-f(x)) / delta	f’(5) = (f(5+0.001) – f(5))/0.001

If f is known, you can use mathematical rules for exactly computing f’
In practice, you can use the definition to approximate f.

Derivatives in multiple dimensions

You have been tasked to develop an algorithm to estimate the value of a painting. You start with interviewing experts and they identified that the following variables are relevant:

Description	Variable
Age in years	a
Width (cm)	w
Height (cm)	h

a=132
w = 70
h = 46

One expert even gave a formula he/she uses:
$V=0.5 * \ln(a)*(w+h)^{2}$

What is the value of this painting? € 25 mln
How much does the value increase if it was 1 year older?
How much does the value increase if it was 1cm wider?
How much does the value increase if it was 1cm higher?

farmhouse in Provence, 1888, oil on canvas, Ailsa Mellon Bruce Collection

Stepwise search

You are to minimize V. Your current best guess is x=(30,40,50). V(x) = 11.2

You try a step in each direction:
V(30+stepsize,40,50) = 11.5
V(30-stepsize,40,50) = 10.9
V(30,40+stepsize,50) = 10.7
V(30,40-stepsize,50) = 12.0
V(30,40,50+stepsize) = 11.3
V(30,40,50+stepsize) = 11.1

You pick the step that leads the best result, in this case the lowest value.
x= (30,40+stepsize,50).

Note: if none of the steps you can take is an improvement, you decrease stepsize. If stepsize becomes too small, you have found a local optimum and are done.

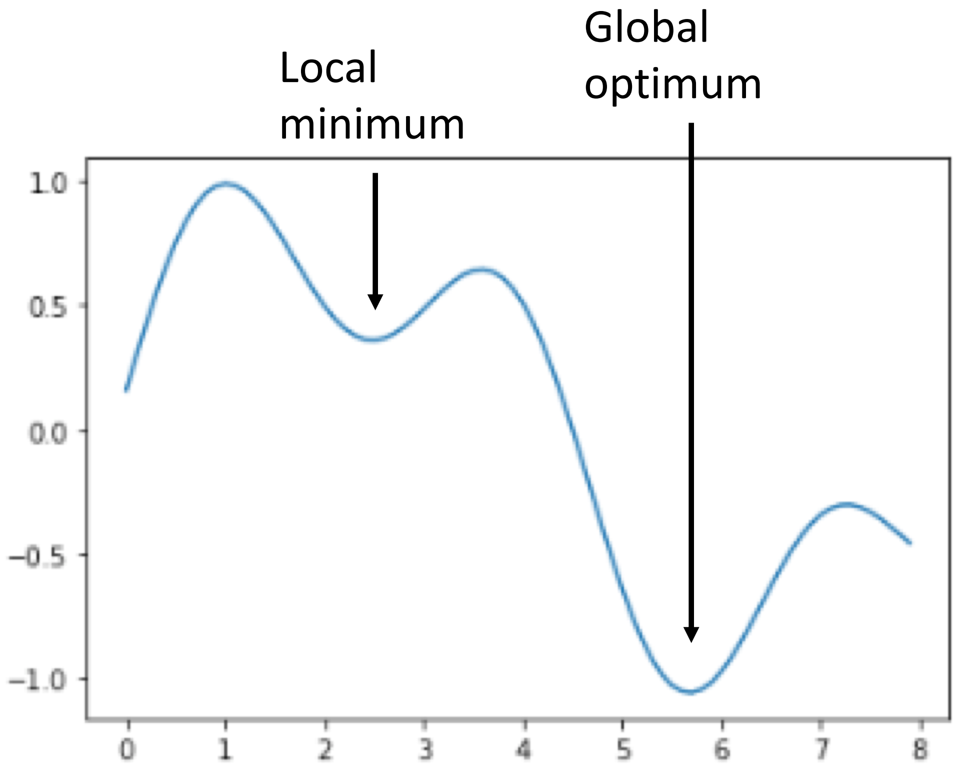

Gradient search

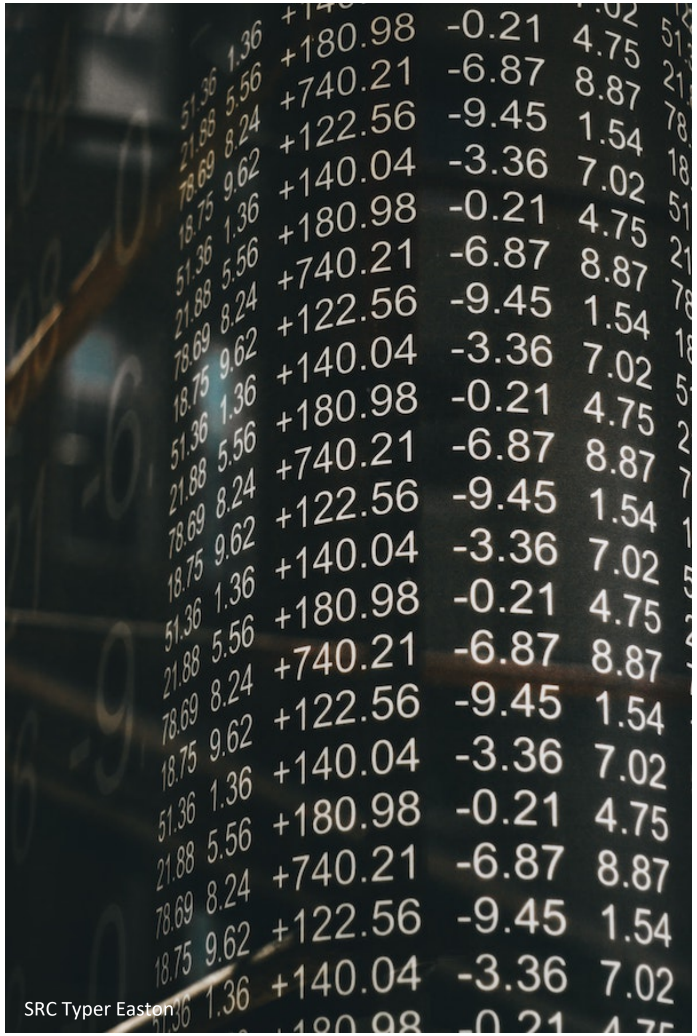

SRC Typer Easton

You are maximizing V. Your current best guess is x=(30,40,50). V(x) = 11.2

You try a tiny adjustment in each direction:
V(30+delta,40,50) = 11.5 (+0.3)
V(30,40+delta,50) = 10.7 (-0.5)
V(30,40,50+delta) = 11.2 (+0.1)

Gradient = (0.3/delta,-0.5/delta,0.1/delta).
So if delta=1.0, gradient = (0.3,-0.5,0.1).

Since we are maximizing, you do a step in the direction of the gradient*:
x= (30+0.3*stepsize,40-0.5*stepsize,50+0.1*stepsize).
E.g. if stepsize = 1.0, x = (30.3,39.5,50.1)

Note: if none of the steps you can take is an improvement, you decrease stepsize. If stepsize becomes too small, you have found a local optimum and are done.

The algorithm is guaranteed to find a local optimum. It will only find the global optimum if you found the right starting point.
It will find an optimum faster since you are moving in all directions at once. Especially in higher dimensions this makes a difference.

*if minimizing, you of course move in the opposite direction of the gradient

Why gradient search is better for NN than stepwise search

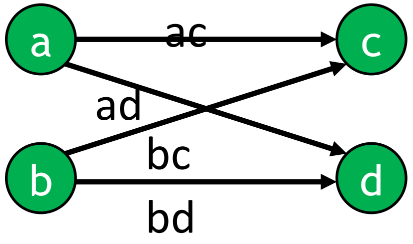

C is higher than target value

d is lower than target value

Example adjustment
(ab,ad,bc,bd)+= (-0.1,0.33,-0,1,0.2)

- The number of evaluations per step increases if the number of parameters (the ‘dimension’ of the search space) increases. Stepwise search basically ignores information from all but one evaluation.
- Stepwise search only optimizes one parameter at a time. Many AI problems have a lot of parameters (e.g. a neural network can have thousands of connections).
- Note: gradient search is designed for smooth functions. It is a good way to optimize neural networks if your algorithm uses a smooth function to compute the output of each neuron.

Depicting gradients

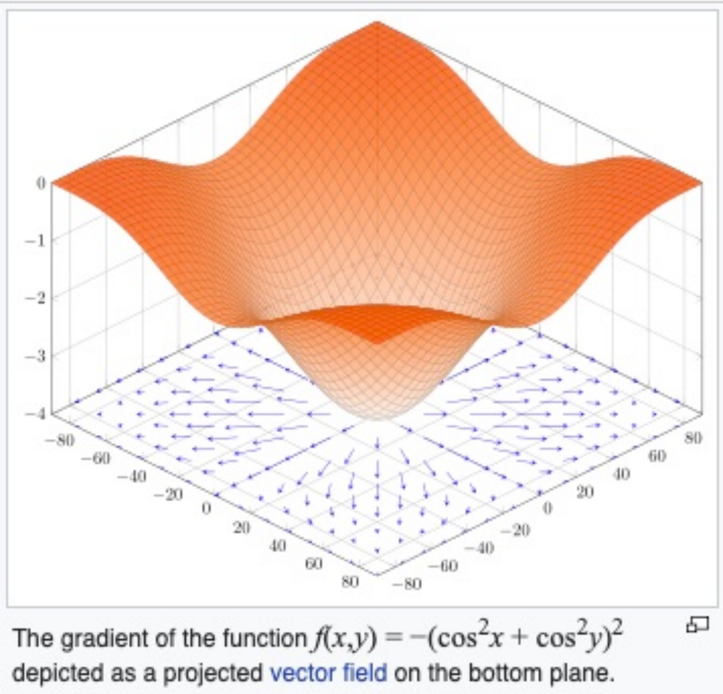

The gradient of the function $f(x,y) = -(\cos^2 x + \cos^2 y)^2$ depicted as a projected vector field on the bottom plane.

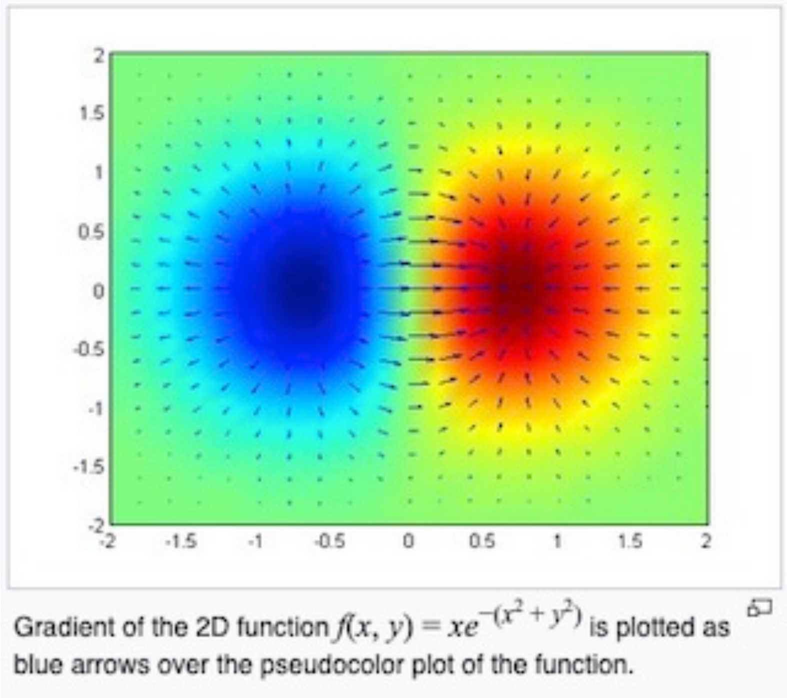

Gradient of the 2D function $f(x, y) = xe^{-(x^2+y^2)}$ is plotted as blue arrows over the pseudocolor plot of the function.

SRC = https://en.wikipedia.org/wiki/Gradient

What is clustering?

- In clustering, you divide the objects of a dataset into different groups that are similar
- The groups are not pre-defined, the algorithm discovers the groups
- Possible applications: dating, recommendations and preprocessing

Input

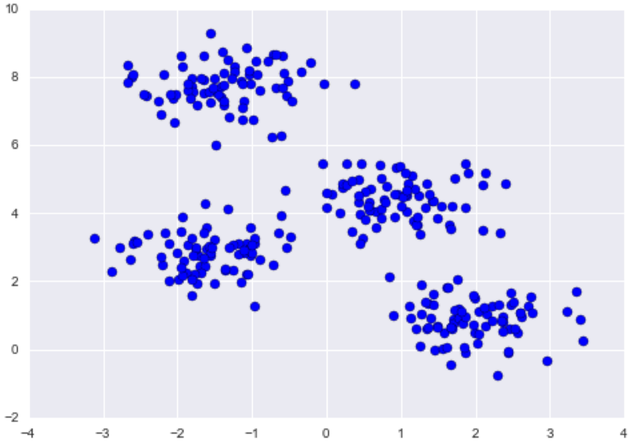

Solution

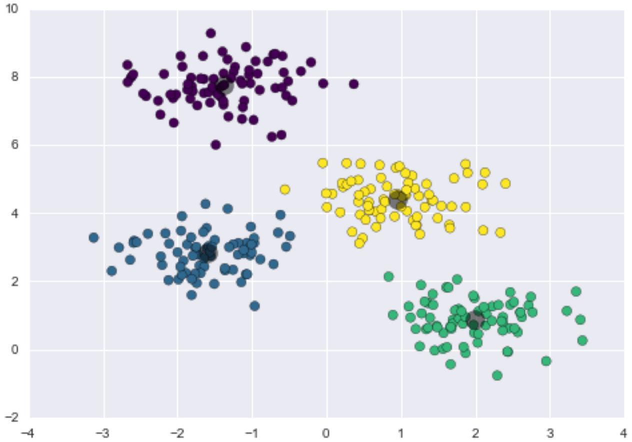

K-means clustering

K-means is a very elegant algorithms, that also works well in practice.
It uses randomization.

Preparation:

- Define a distance function. In our case, difference in tax value, or distance based on x-coordinates, y-coordinates
- Choose the number of clusters (e.g. 4). Each cluster is defined by a center location
- Set the location of the center of each cluster at a different, random location

Repeat until done (no changes or max no. of rounds):

- Assign each point to the cluster with the nearest centerpoint
- Change the center of a cluster to the average position of all points in that cluster

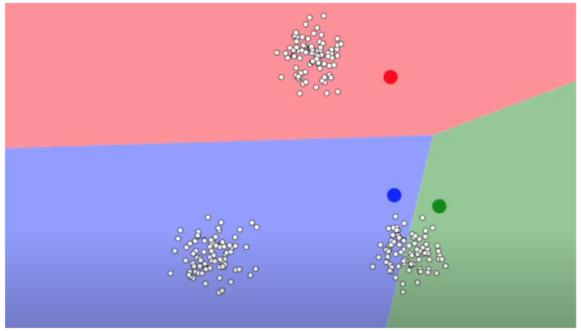

https://medium.com/@neil.liberman/k-means-clustering-e00408493a40
https://www.youtube.com/watch?v=R2e3Ls9H_fc&t=167s

A clustering exercise

Euclidian distance is defined as

sqrt((vara[n]-vara[m])^2+ … + (vard[n]-vard[m])^2)

We can define the difference in size as between house m and n as:

sqrt((housearea[n]-housearea[m])^2+ … + (lotarea[n]-lotarea[m])^2)

Sort the houses in 3-6 size classes using k-means clustering. Compute the average tax-value per house.

Classification

In a classification problem you put the correct label on a datapoint.

- Potential labels for houses could be “villa, tiny house, family house, …”

Prediction

- In a prediction problem, you must find a missing value based on available data.
- To make a prediction algorithm, you need example data with the correct answer (supervised learning)

Nr	Size	Quality	Value
1	0.1	AA	€ 880
2	0.8	A	€ 550
6	1.6	AA	€ 6.340
4	0.4	A	?
5	1.0	AA	?
6	1.1	A	?

Training data with correct answer

Cases were prediction is needed

Linear regression

Input: house-area

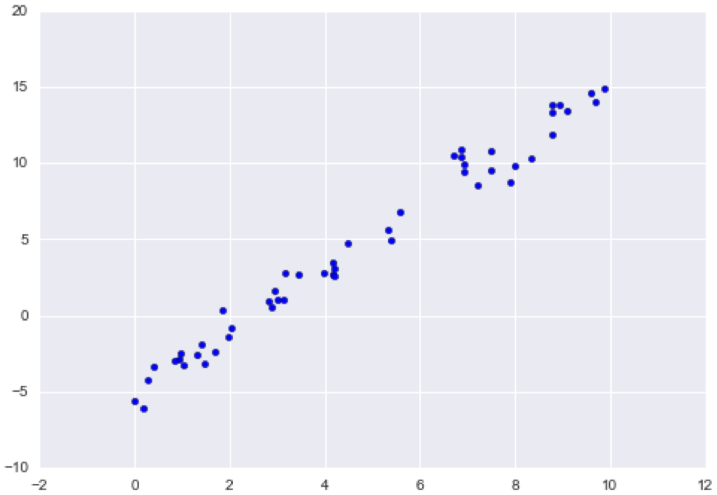

value

House-area

Outcome: linear model

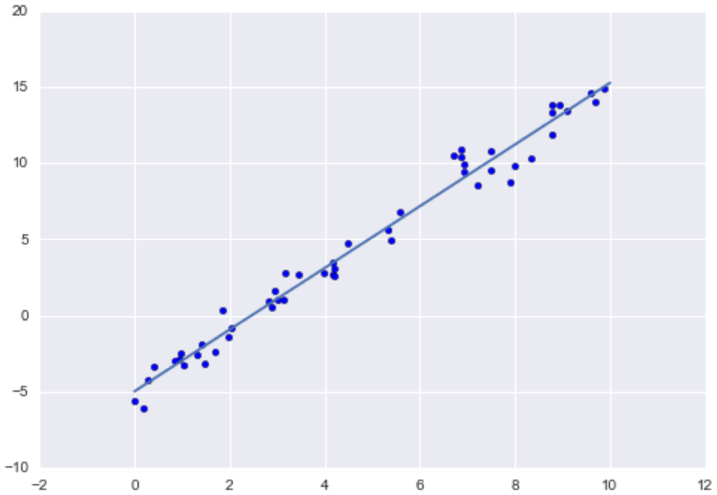

value

House-area

Value = - 4000 + 2000*House-area

Prediction is an optimization problem

- There are many function optimization algorithms.
- These can be used for prediction if you define an error function: the best prediction is the one that minimizes the total error
- Distance-squared is a common error function that works well.

Nr	Size	Quality	Actual-Value	Prediction	Error	Error-squared
1	0.1	AA	€ 880	€ 870	10	100
2	0.8	A	€ 550	€ 555	5	25
6	1.6	AA	€ 6.340	€ 5.340	1000	10.000
~~4~~	~~0.4~~	~~A~~	~~?~~			
~~5~~	~~1.0~~	~~AA~~	~~?~~			
~~6~~	~~1.1~~	~~A~~	~~?~~			
Total error						**10.125**

B prediction

- Use linear regression to estimate the retailvalue from taxvalue and housesize
- Use linear regression to estimate the retailvalue from taxvalue and housesize
- Use linear regression to estimate the retailvalue from housesize and lotarea
- Use linear regression to estimate the retailvalue from all variables
- Use linear regression to estimate the retailvalue from all variables except taxvalue

Experimenting with explanations

- Pick your best linear regression model. Which inputs does it use?
- Pick your favourite house. Print it as a vector to the screen.

Vary the input, and record the change in value (see table on next slide)

Testing the impact of each input

Column	Original value	New input for testing	Diff in outcome	Explanation
Id		Not applicable		
Zipcode		3500?		
Lot-area		+10 m2?		
House-area		+10 m2?		
Balcony		+1		
X-coor		2500		
Y-coor		5500		
Buildyear		2000		
Bathrooms		+1		
Outcome (retailvalue)				

- By filling in a table like this, you test the impact of each inputvalue. You should do this for any model.
- You can use libraries to help you (LIME, SHAP) but try to understand what you are doing
- You must understand the median, min, max and variance of each variable to estimate total impact

DeepAIHire potential bias

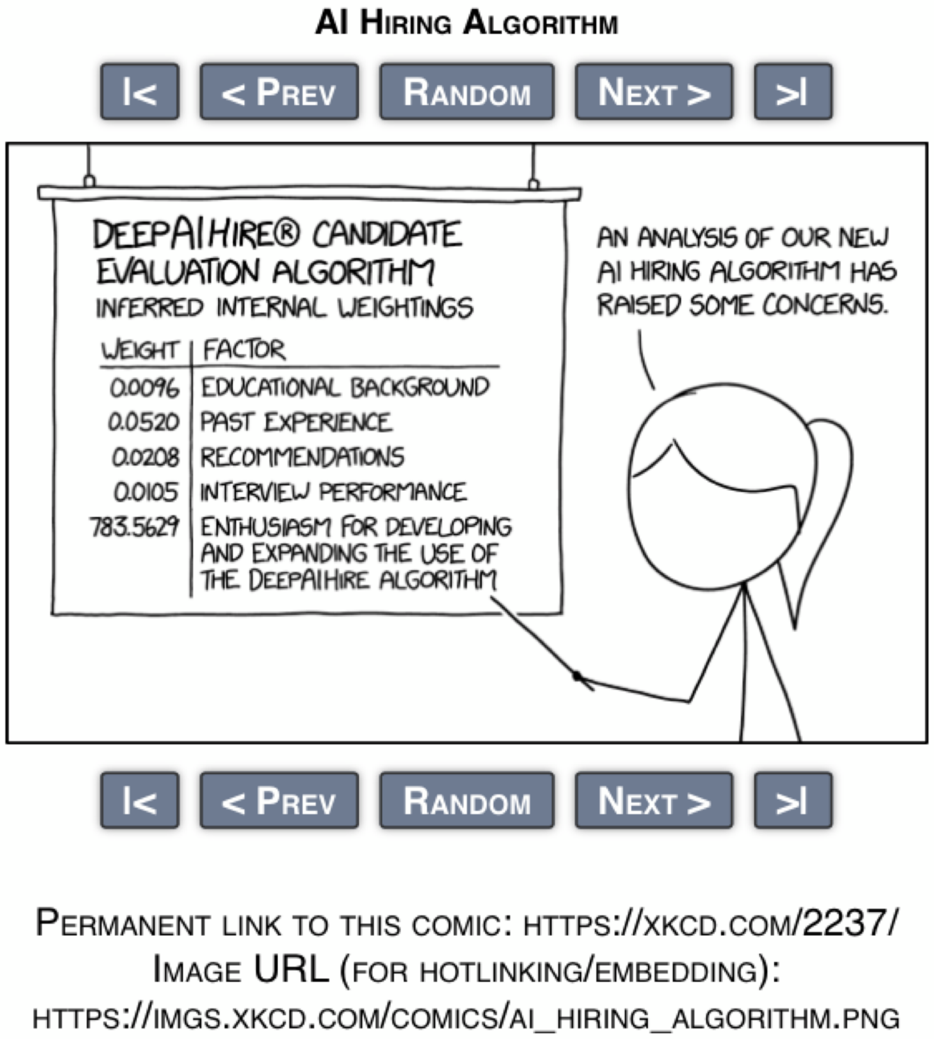

AI Hiring Algorithm

|<
< Prev
Random
Next >
>|

DEEPAIHIRE® CANDIDATE EVALUATION ALGORITHM

INFERRED INTERNAL WEIGHTINGS

WEIGHT	FACTOR
0.0096	EDUCATIONAL BACKGROUND
0.0520	PAST EXPERIENCE
0.0208	RECOMMENDATIONS
0.0105	INTERVIEW PERFORMANCE
783.5629	ENTHUSIASM FOR DEVELOPING AND EXPANDING THE USE OF THE DEEPAIHIRE ALGORITHM

AN ANALYSIS OF OUR NEW AI HIRING ALGORITHM HAS RAISED SOME CONCERNS.
|<
< Prev
Random
Next >
>|

Permanent link to this comic: https://xkcd.com/2237/
Image URL (for hotlinking/embedding):
https://imgs.xkcd.com/comics/ai_hiring_algorithm.png

AI and ethics

Ethical questions surrounding AI

- What ethical questions have you encountered involving AI?

- How do you think we can make AI developers more aware of AI related risks?

Some ethical risks

A few years ago we collected examples of known risks into an article*. The risks can be categorized:

AI not working as intended:

- Cold start
- Poor data
- Discrimination and bias
- Offensive input
- Lack of diversity

Use of data

- GDPR and privacy
- Data ownership
- Data leakage

Other issues:

- Lack of explanations
- Results cannot be repeated
- Unclear responsibility

* https://ictinstitute.nl/ai-risk-management-checklist/

Ethics Inc is one tool to

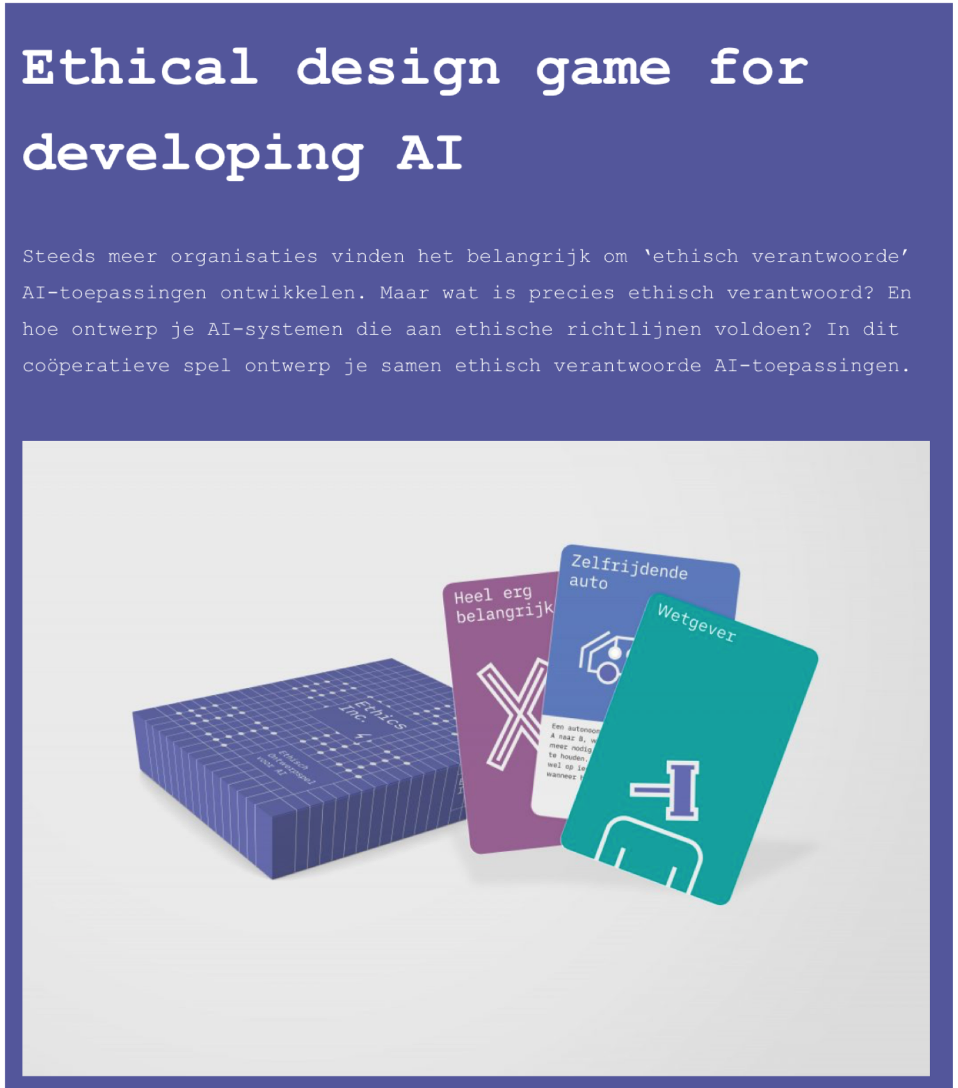

Ethical design game for developing AI

Steeds meer organisaties vinden het belangrijk om 'ethisch verantwoorde' AI-toepassingen ontwikkelen. Maar wat is precies ethisch verantwoord? En hoe ontwerp je AI-systemen die aan ethische richtlijnen voldoen? In dit coöperatieve spel ontwerp je samen ethisch verantwoorde AI-toepassingen.

- Ethics INC is a discussion game developed at HU.
- We will play Ethics Inc, and hopefully you will learn more about the ethical aspects

AI impact analysis

Artificial Intelligence Impact Assessment
www.ecp.nl
ECP
Platform voor de InformatieSamenleving
THINK BEFORE YOU PRINT

The AI impact assessment is a more procedural solution: it describes the process for responsibly introducing an AI solution.

Steps include:

- Description
- Benefits
- Safety
- Transparency

https://ecp.nl/wp-content/uploads/2018/11/Artificial-Intelligence-Impact-Assesment.pdf
https://ecp.nl/publicatie/artificial-intelligence-impact-assessment-english-version/

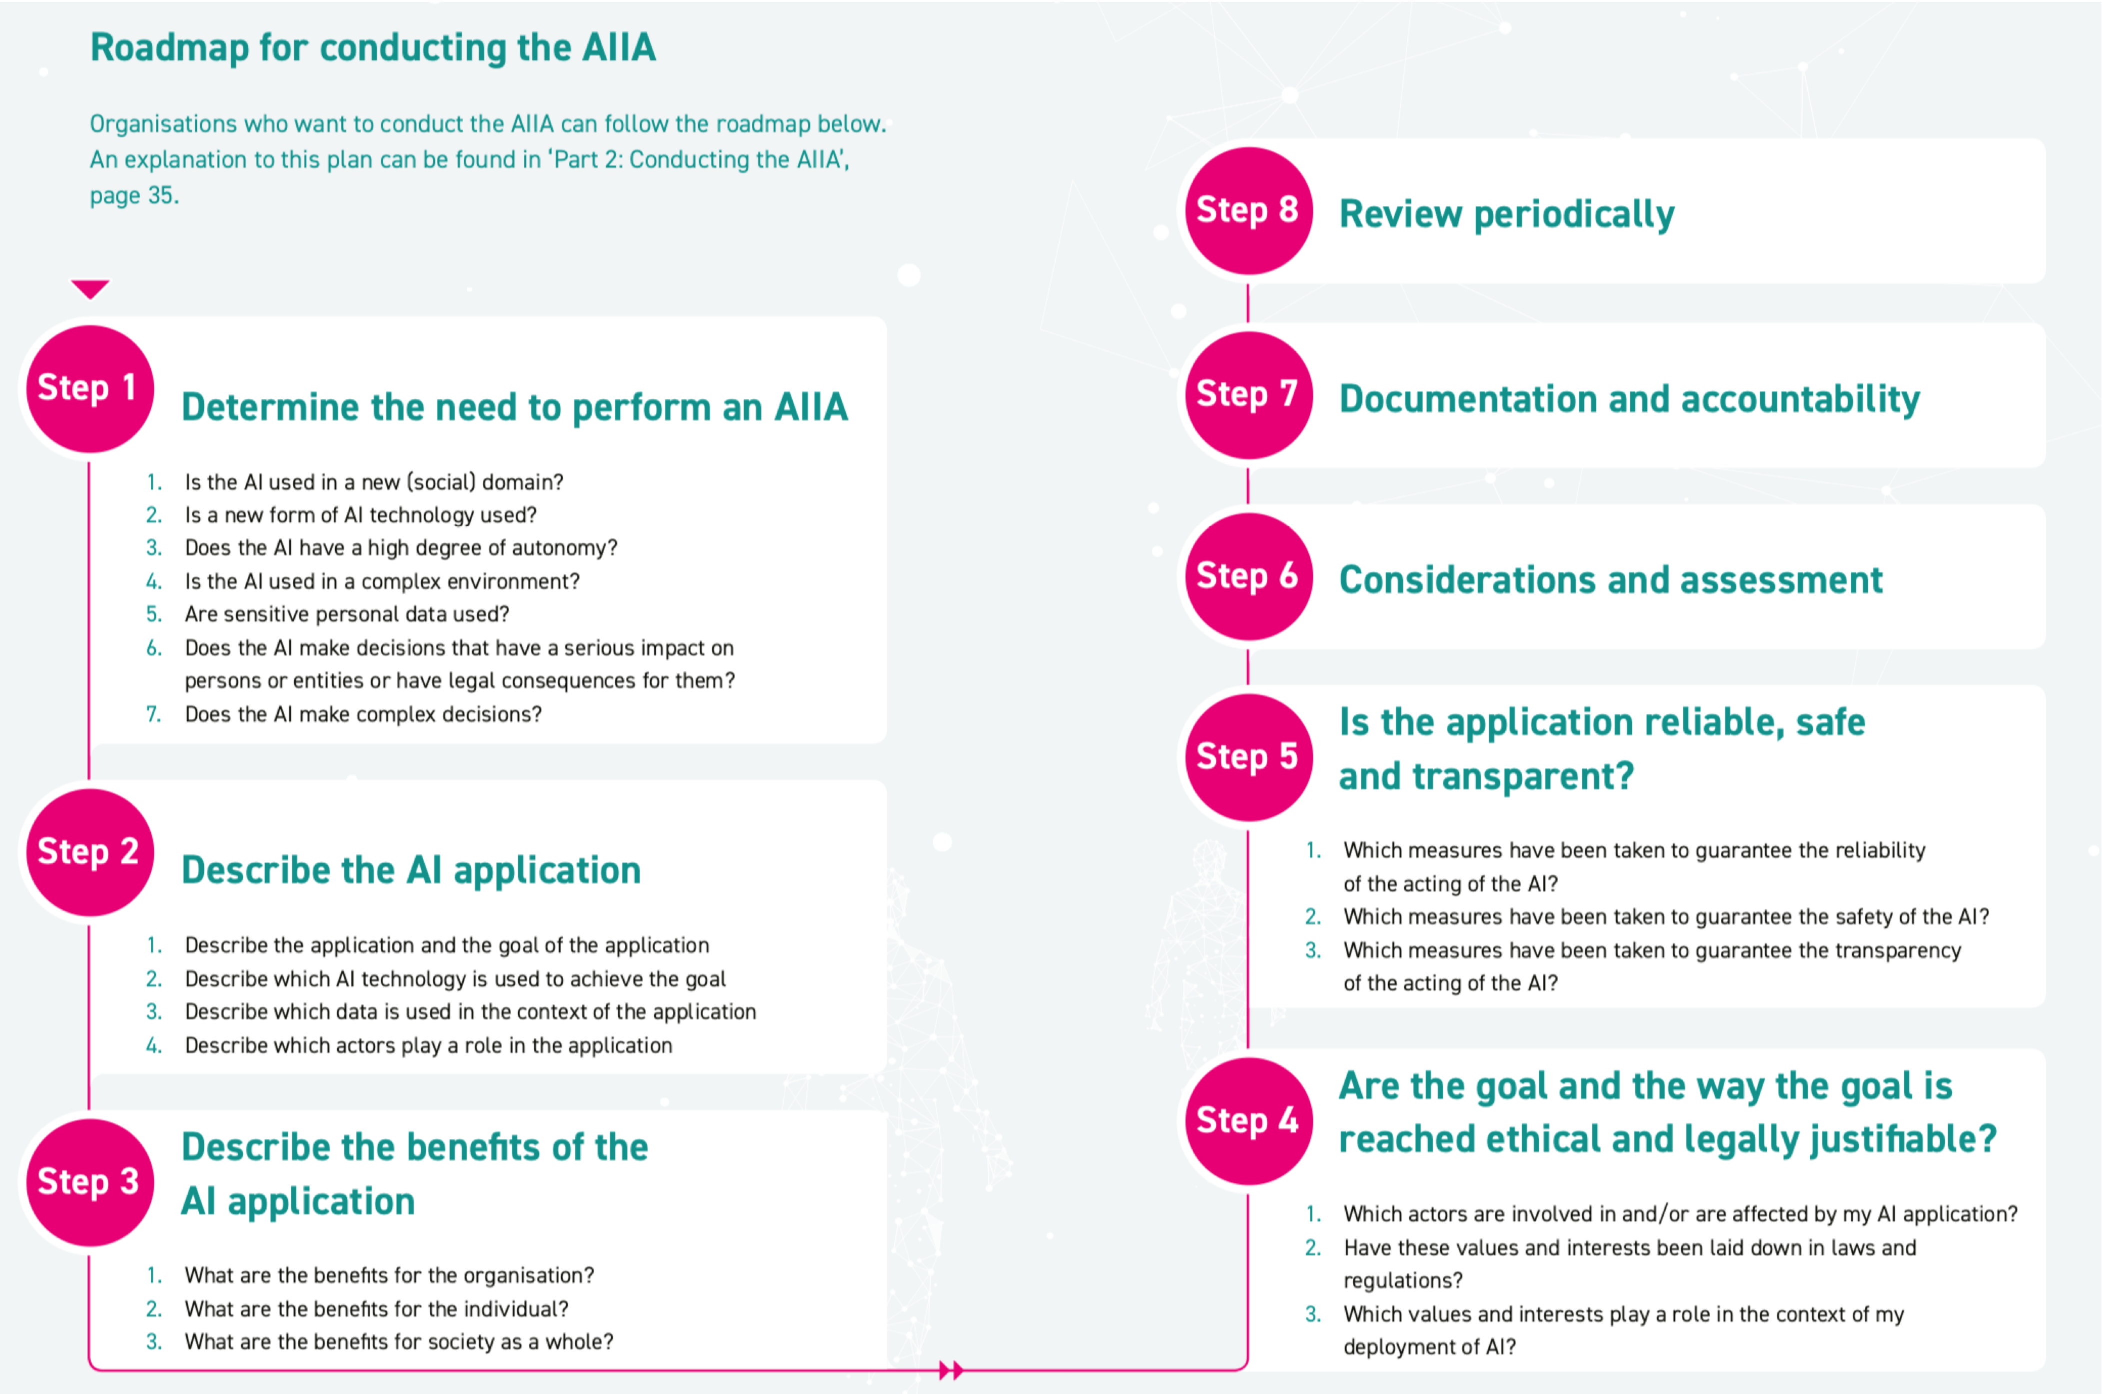

Roadmap for conducting the AIIA

Organisations who want to conduct the AIIA can follow the roadmap below. An explanation to this plan can be found in 'Part 2: Conducting the AIIA', page 35.

Step 1 Determine the need to perform an AIIA

1. Is the AI used in a new (social) domain?
2. Is a new form of AI technology used?
3. Does the AI have a high degree of autonomy?
4. Is the AI used in a complex environment?
5. Are sensitive personal data used?
6. Does the AI make decisions that have a serious impact on persons or entities or have legal consequences for them?
7. Does the AI make complex decisions?

Step 2 Describe the AI application

1. Describe the application and the goal of the application
2. Describe which AI technology is used to achieve the goal
3. Describe which data is used in the context of the application
4. Describe which actors play a role in the application

Step 3 Describe the benefits of the AI application

1. What are the benefits for the organisation?
2. What are the benefits for the individual?
3. What are the benefits for society as a whole?

Step 4 Are the goal and the way the goal is reached ethical and legally justifiable?

1. Which actors are involved in and/or are affected by my AI application?
2. Have these values and interests been laid down in laws and regulations?
3. Which values and interests play a role in the context of my deployment of AI?

Step 5 Is the application reliable, safe and transparent?

1. Which measures have been taken to guarantee the reliability of the acting of the AI?
2. Which measures have been taken to guarantee the safety of the AI?
3. Which measures have been taken to guarantee the transparency of the acting of the AI?

Step 6 Considerations and assessment

Step 7 Documentation and accountability

Step 8 Review periodically