

Artificial Intelligence

Tijmen Blankevoort



Scyfer

- Opgericht 2013
- Prof. dr. Max Welling
- Omzet+talent verdubbeld elk jaar

Scyfer maakt state-of-the-art Artificial Intelligence oplossingen voor o.a. Multinationals in Nederland. Onze expertise laat AI echt werken voor business.



UNIVERSITEIT VAN AMSTERDAM



Onze klanten



ING



NOS



LEIDEN UNIVERSITY
MEDICAL CENTER

TATA STEEL

Reap the benefits of AI

Een intelligente revolutie



Live vertalingen (Skype 2015)



Zelfrijdende auto's
(Veel bedrijven)



Professioneel Go
(Google 2016)

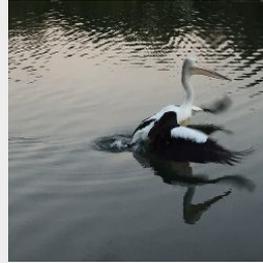
Jeopardy
(IBM 2013)



The AI jungle



Bovenmenselijke
beeldherkenning



AI Vertaalt



Stephen Hawking en Elon
Musk zijn bang voor AI

Moeilijke ethische vragen
over auto's



AI Neemt de wereld
over?



IBM Watson
verhelpt
kanker

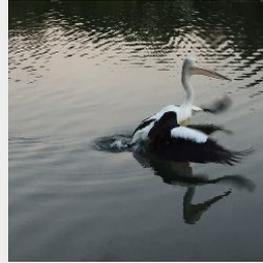


Google is
"AI-first"

The AI jungle



Bovenmenselijke
beeldherkenning



AI Vertaalt



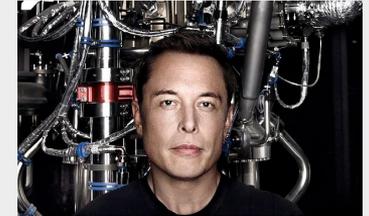
Stephen Hawking en Elon
Musk zijn bang voor AI

Moeilijke ethische vragen
over auto's



AI Neemt de wereld
over?

Wat is er nu eigenlijk
aan de hand?



IBM Watson
verhelpt
kanker



Google is
"AI-first"

Doel van deze presentatie

- Ik laat zien dat AI slim is
- Ik laat zien dat AI dom is
- Wat is er mogelijk in de business
- Hoe maken we algoritmes die zichzelf uitleggen
- Een kleine blik op de toekomst

Machine Learning

Machine learning

We willen dat een computer iets voor je doet

```
If Number = 1 Then
    Count1 = Count1 + 1
ElseIf Number = 2 Then
    Count2 = Count2 + 1
ElseIf Number = 3 then
    Count3 = Count3 + 1
Else
    CountX = CountX + 1
End If
```



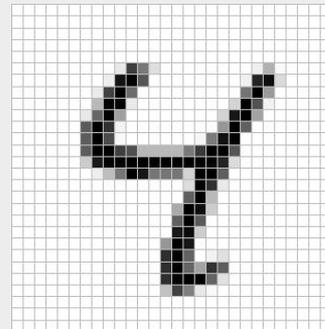
Waarom programmeren als de computer kan leren?

Voorbeeld



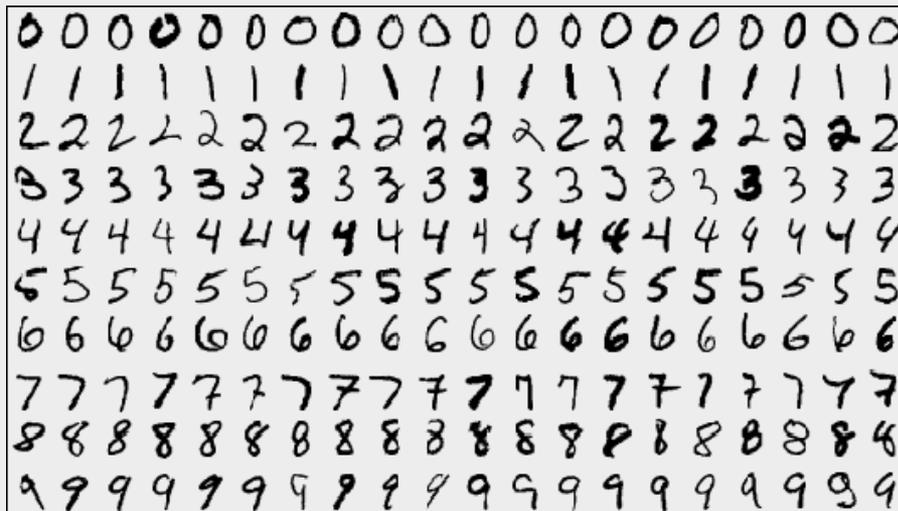
De computer leert om karakters op brieven te lezen

Dit is een 4



Herhaal veelvuldig

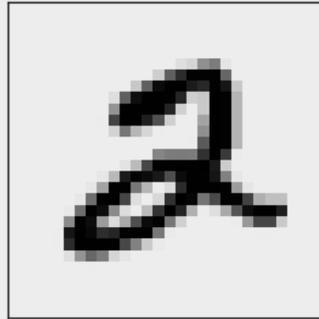
Dit zijn 0'en,
1'en, 2'en...



60.000 voorbeelden met een label (0, 1, 2, etc.)

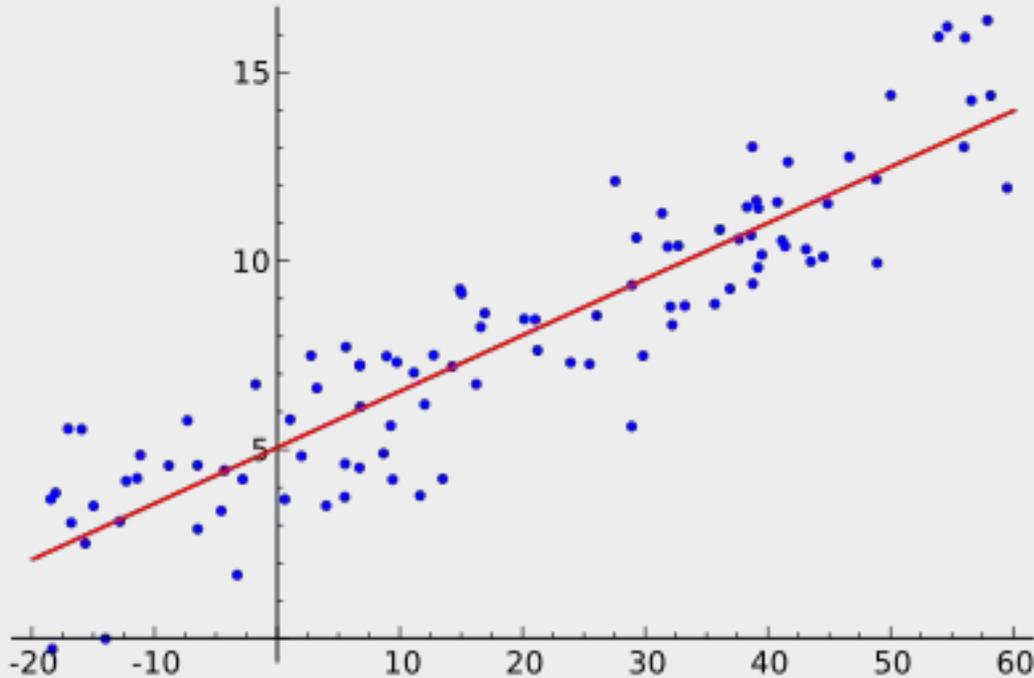


Na veel studeren



Dit is een 2

Lijkt erg op statistiek!



De punten worden na het maken van het model weggegooid. Alleen de rode lijn blijft over

Erg lastig om te doen



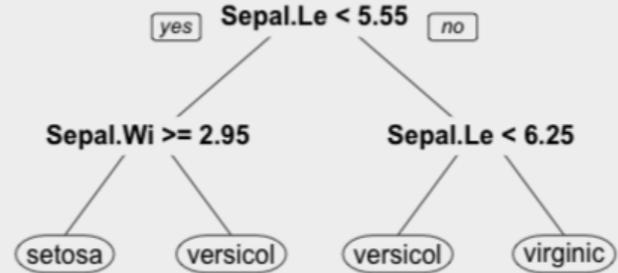
Virginica



Versicolor



Setosa



Een beslisboom kan niet goed een plaatje uitleggen

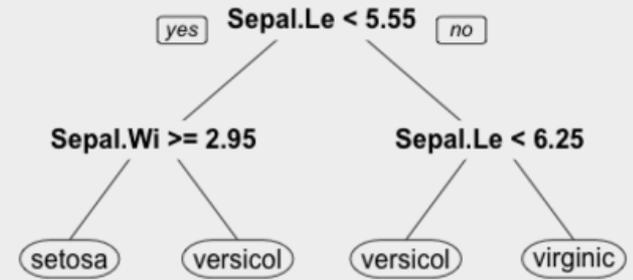
Oude methode



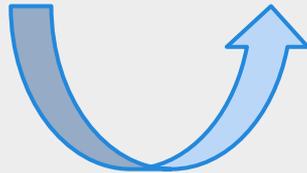
Ruwe data

Sepal.Length	Sepal.Width	Petal.Length	Petal.Width	Species
5.1	3.5	1.4	0.2	setosa
4.9	3.	1.4	0.2	setosa
4.7	3.2	1.3	0.2	setosa
4.6	3.1	1.5	0.2	setosa
5.	3.6	1.4	0.2	setosa
5.4	3.9	1.7	0.4	setosa
4.6	3.4	1.4	0.3	setosa
5.	3.4	1.5	0.2	setosa
4.4	2.9	1.4	0.2	setosa
4.9	3.1	1.5	0.1	setosa

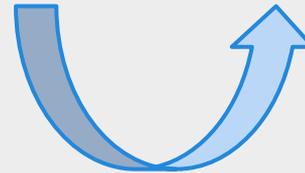
Features



Model

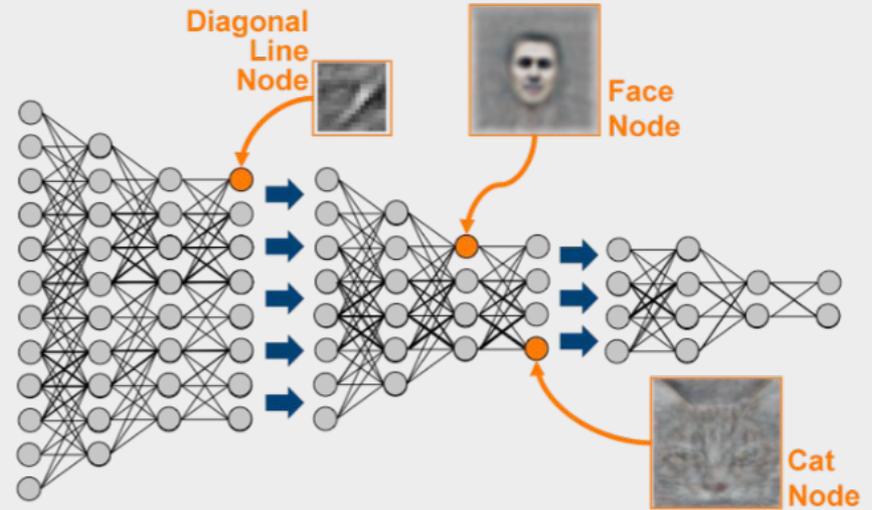
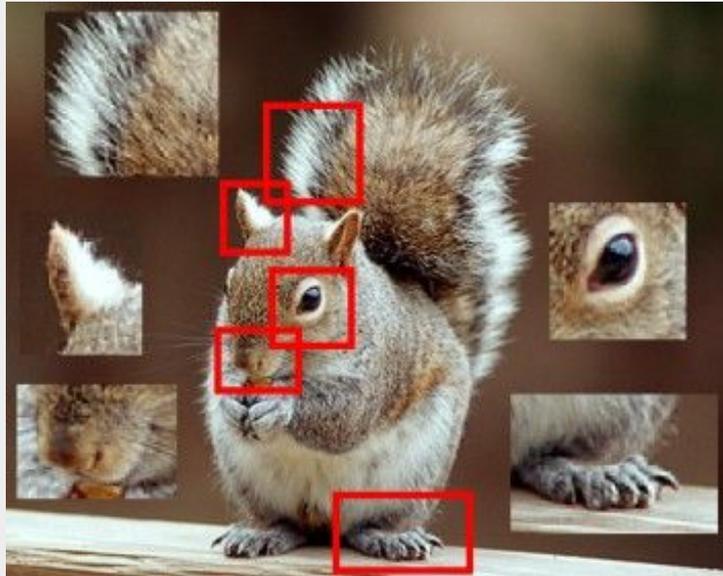


Feature Extraction



Machine Learning

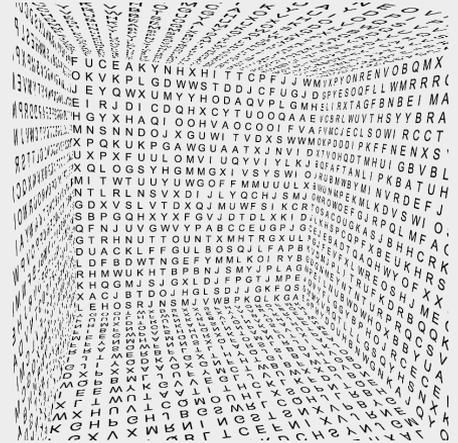
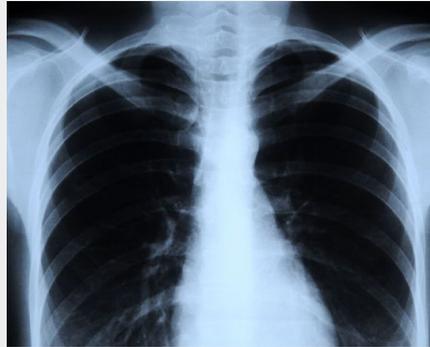
Deep Learning (2012 - nu)



Plots: Mogelijkheden

Nu kunnen we Machine Learning toepassen op complexe data!

- Tekst
- Plaatjes
- Geluid
- Sensor data
- ...



Geef mij genoeg (input, output), en Deep Learning kan dezelfde voorspellingen doen!

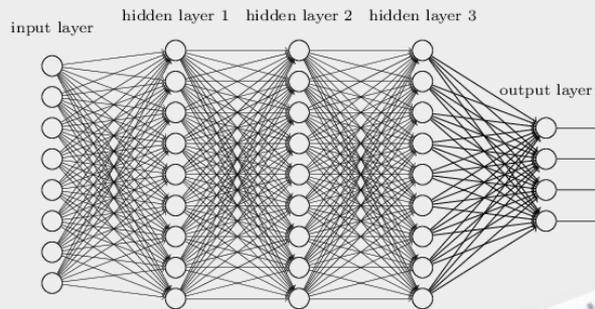
Deep Learning

We train very large neural networks

E.g. Each pixel is an input

Each neuron has parameters

Tweak and tune parameters for each example to teach it something new



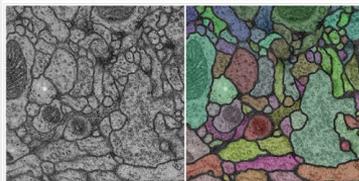
ResNet: 152 layers

Netwerk zoekt het zelf uit

Let the machine figure it out!

Best results on almost all competitions

- ICDAR 2011 Chinese handwriting recognition
- IJCNN 2011 Traffic signs
- Cell segmentation challenge 2012
- Kaggle



Deep learning voor plaatjes



IMAGENET



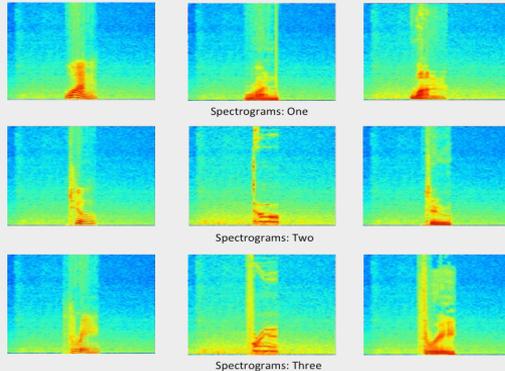
Speech recognition



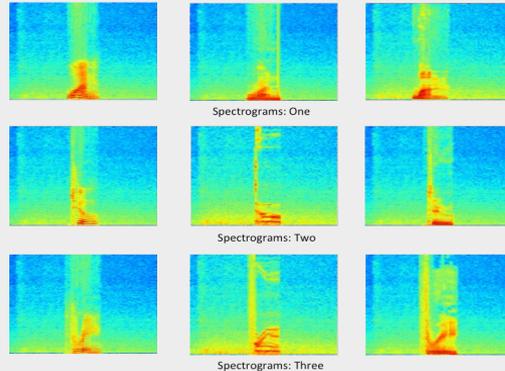
Neem 5.000 audiobooks
Data zijn losse zinnen

Input: Spraak
Output: Tekst

Geef dit aan Deep Learning:



Speech recognition



Neem 5.000 audiobooks
Data zijn losse zinnen

Input: Spraak
Output: Tekst

Geef dit aan Deep Learning:



Siri
Herkennt spraak!

E-mails

Verzamel miljoen e-mails

Input: e-mails

Output: Korte antwoorden

Geef aan Deep Learning:

Really love your startup idea



 **Oleg Campbell** <oleg@replyapp.io>
to Jason ▾

2:01 AM (5 minutes ago) ☆



Hey Jason,

My name is Oleg. I found your startup on AngelList and really liked the idea.
Quick question: what are your future plans as for attracting new users?

I though I may have some ideas for you. Let me know what to you think.

Oleg

Yes, I'd love to attend the meeting. See you there!



E-mails

Really love your startup idea

 **Oleg Campbell** <oleg@replyapp.io>

2:01 AM (5 minutes ago) ☆

to Jason ▾

Hey Jason,

My name is Oleg. I found your startup on AngelList and really liked the idea. Quick question: what are your future plans as for attracting new users?

I though I may have some ideas for you. Let me know what to you think.

Oleg

Yes, I'd love to attend the meeting. See you there!

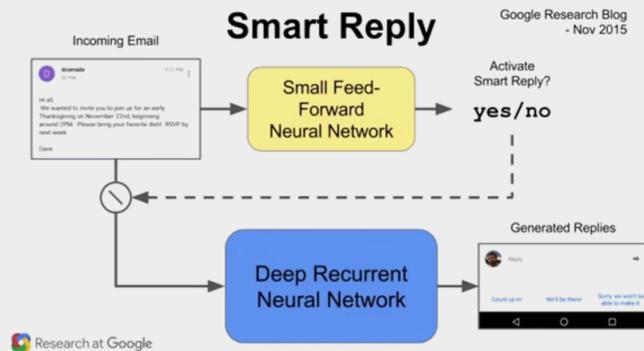


Gather millions of e-mails

Input: e-mails

Output: short responses

Give it to Deep Learning:



Inbox smart reply
Automatically generate simple possible responses to e-mails

Combine text and images

What happens if you give the algorithm context?

Dataset: Images -> Text

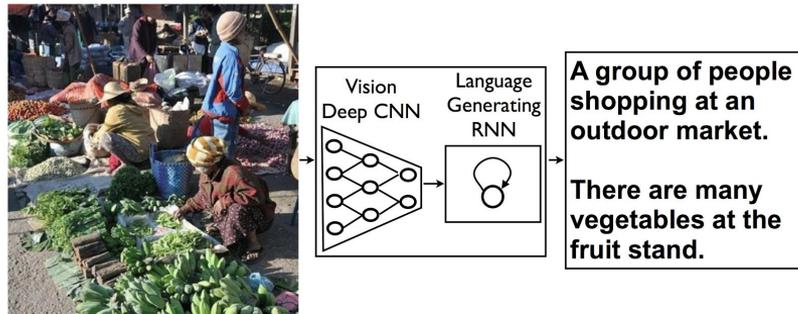
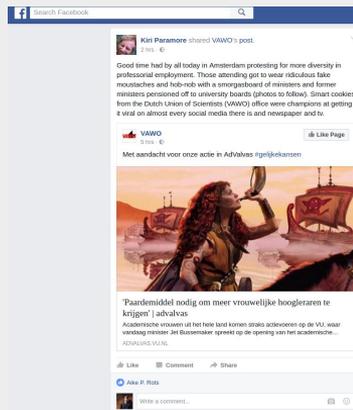
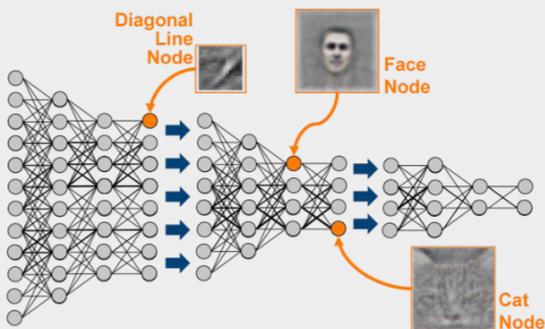


Image to text (Google 2014)

Deep learning is overal!



Google Translate



Facebook Timeline

De meest gebruikte technologie waar weinig van hebben gehoord!

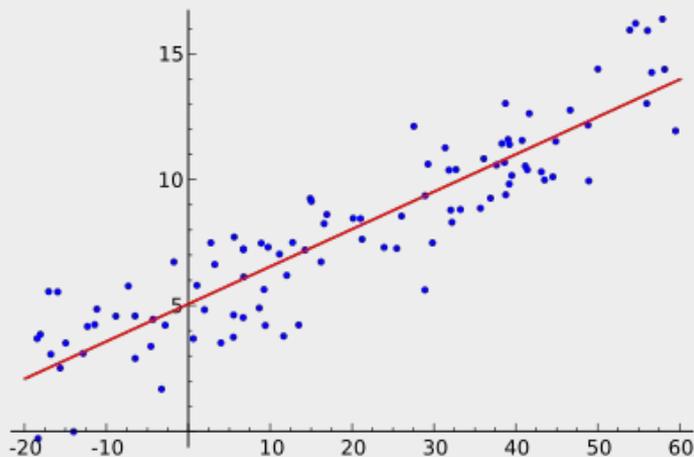


Apple Siri



Terug naar AI

AI is geen magie, maar functie approximatie!



Bruikbare AI oplossingen zijn momenteel bijna altijd een combinatie van:

- Machine Learning
- Een dataset met veel voorbeelden (input en output)
- Een hele hoop handwerk in programmeren en implementeren.

Leer computer games



Input: Pixels

Output: Score, Actie

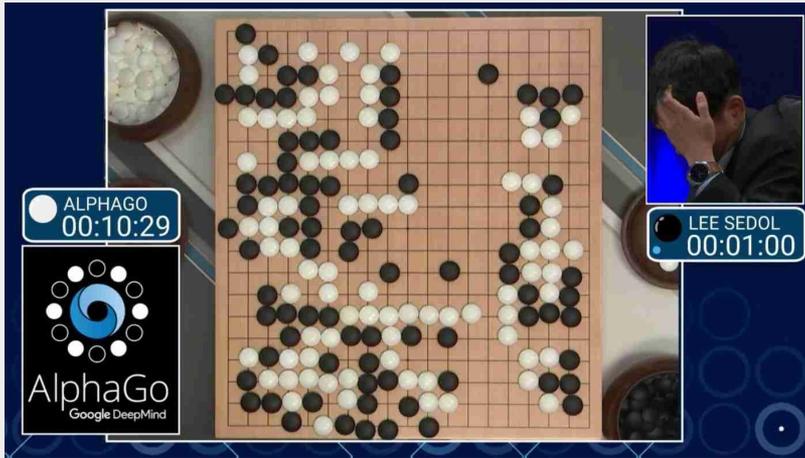
Neem elke tijdsstap de actie die de score optimaliseert

Heb je de voorspelling verkeerd, leer van die fout

Ook AlphaGo is zo gemaakt!



AI zijn punttoplossingen



AlphaGo verslaat Lee Sedol! Veel beter in Go dan de mens!

Maar... kan geen Stratego spelen!

Huidige AI zijn allemaal punttoplossingen die 1 ding erg goed kunnen

AI is eigenlijk niet zo slim

- AI maakt niet menselijke fouten
- AI heeft erg veel voorbeelden nodig om iets te leren
- AI heeft geen common-sense
- AI generaliseert totaal niet
- AI kan niet goed conceptualiseren en redeneren
- Geen creativiteit
- Geen geheugen
- Kan niet uitleggen waarom beslissingen worden gemaakt
-

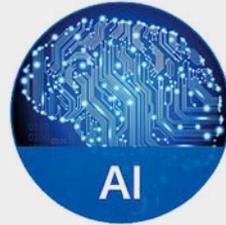
AI als 'teachable software tool'



Maar de potentie is geweldig!



1.25 miljoen mensen
verongelukken in het verkeer.
Zelfrijdende auto's gaan dit
sterk reduceren!



IoT en AI, devices
worden 'slimmer'

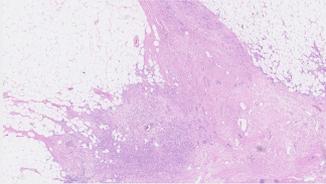


AI kan veel processen
in bijv. De healthcare
verbeteren



AI assistenten helpen
met dagelijkse taken,
en maken informatie
veel makkelijker
beschikbaar

Wat kan onze business met AI?



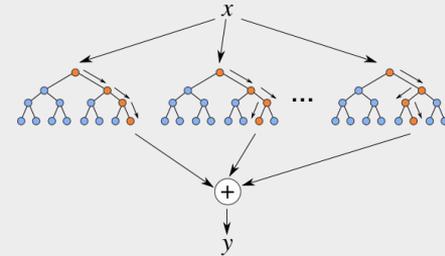
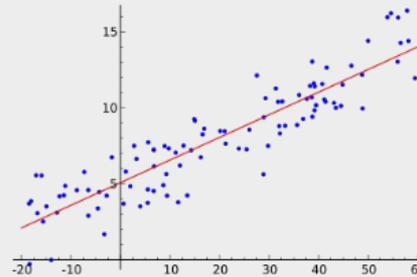
Complexe repetitieve taken
assisteren/overnemen

‘Gebruik maken’ van meer data
Huidige modellen verbeteren
Voorspellen

The black box

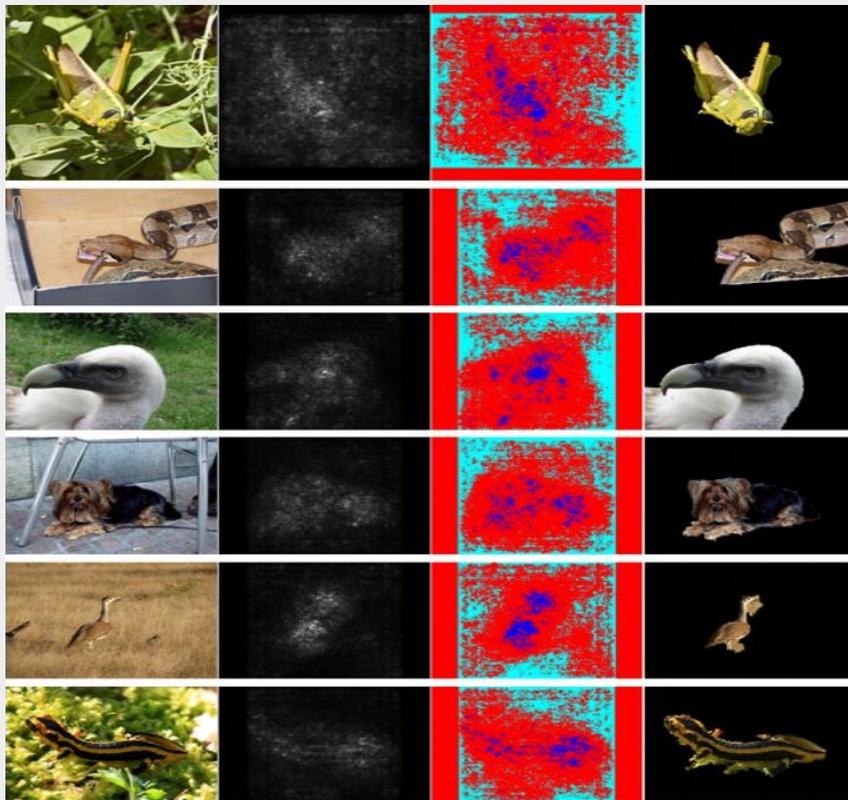


*Complex correlations
and decisions are
difficult to explain in an
easy way.*

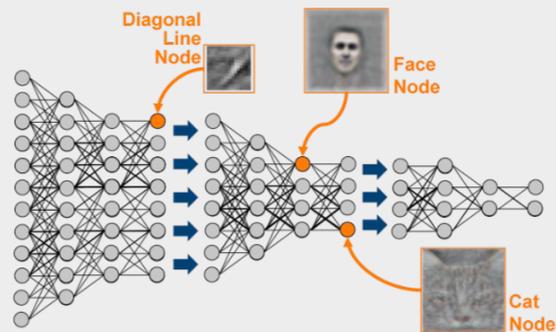


Deep learning feels like a black box, or does it?

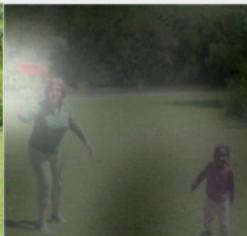
Input features



The 'concepts' make complex algorithms easier to understand!



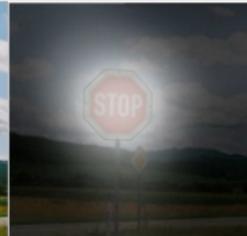
Where in the image do we look



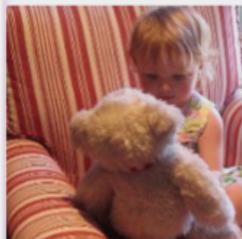
A woman is throwing a frisbee in a park.



A dog is standing on a hardwood floor.



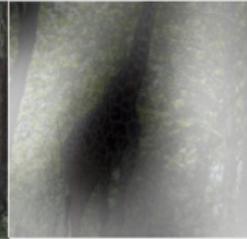
A stop sign is on a road with a mountain in the background.



A little girl sitting on a bed with a teddy bear.



A group of people sitting on a boat in the water.

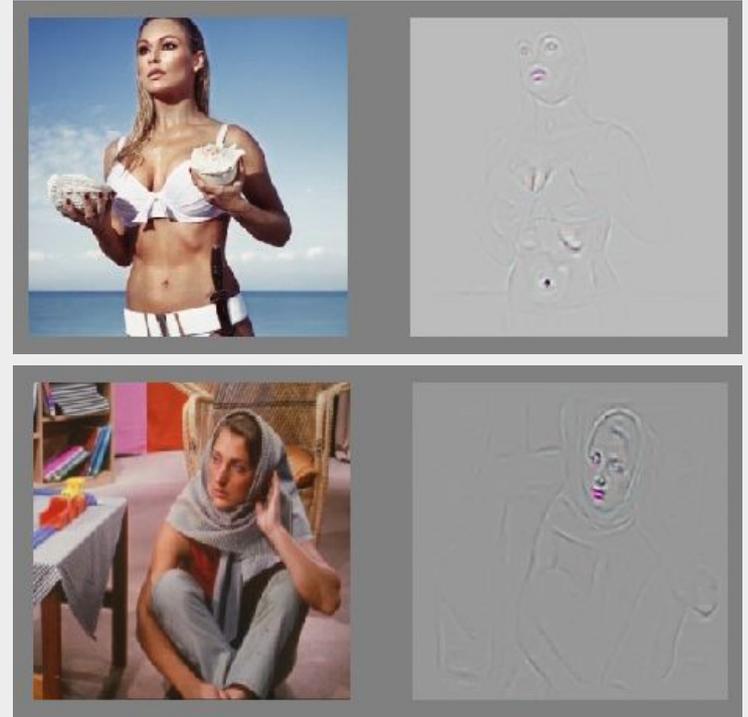


A giraffe standing in a forest with trees in the background.

Nudity filter



Filtered response maps (Zeiler 2014)



Text explanation of decisions

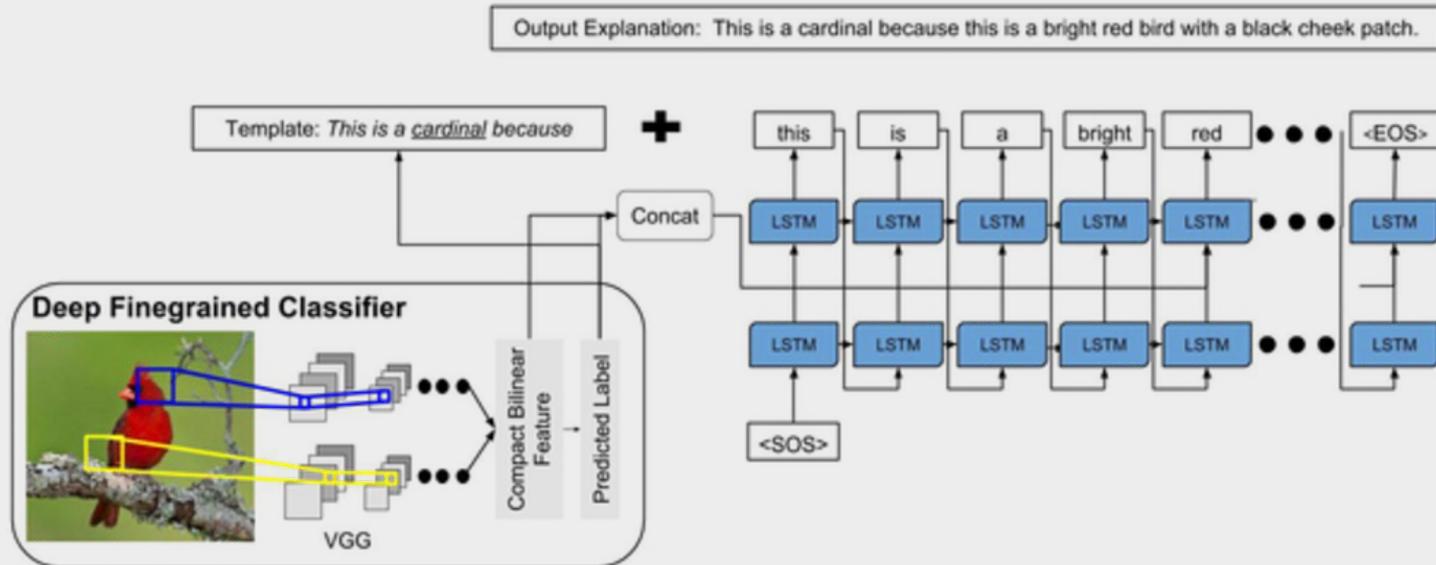


Image generation



redshank

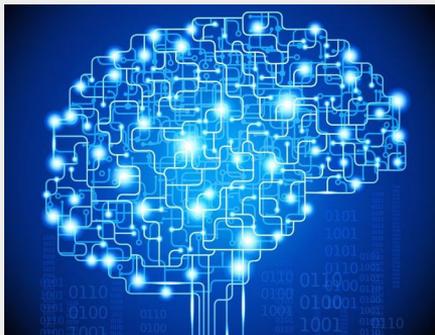
ant

monastery

volcano

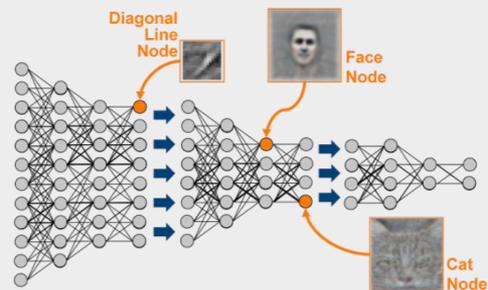
Evolutie van AI

```
If Number = 1 Then  
    Count1 = Count1 + 1  
ElseIf Number = 2 Then  
    Count2 = Count2 + 1  
ElseIf Number = 3 then  
    Count3 = Count3 + 1  
Else  
    CountX = CountX + 1  
End If
```



Programmeren
Niet te doen voor moeilijke
taken

Machine Learning
Maakt het ons mogelijk
makkelijker algoritmes te
maken voor complexe taken



Deep Learning
Nu kunnen we gebruik maken
van zeer complexe data voor
machine learning

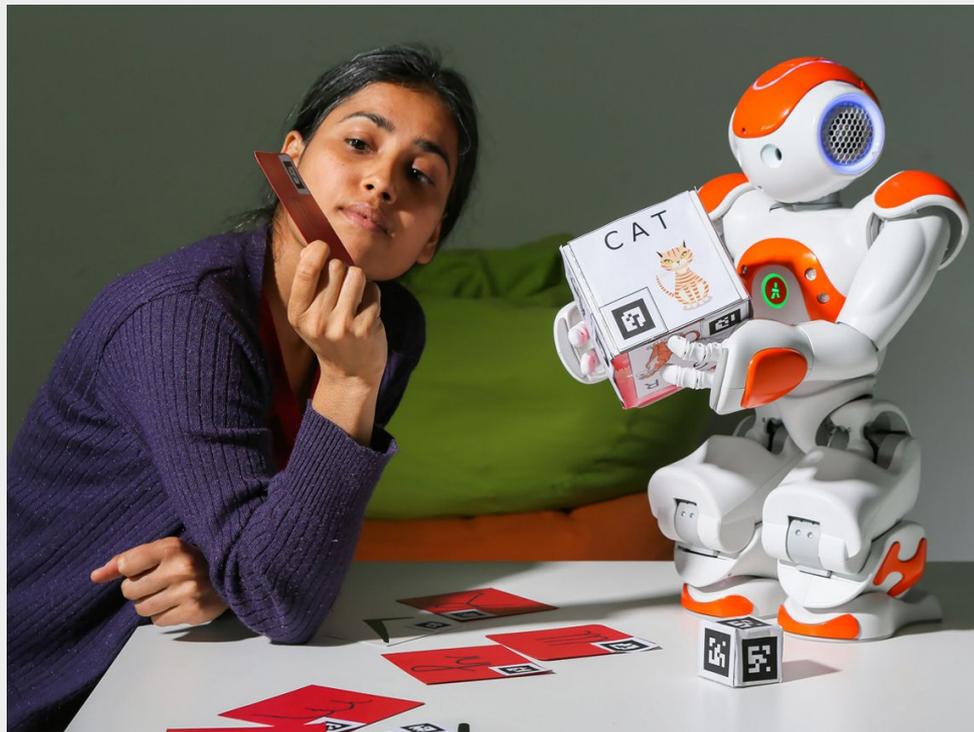


Active Learning
Maakt het mogelijk om zelf een
AI te leren een taak uit te
voeren

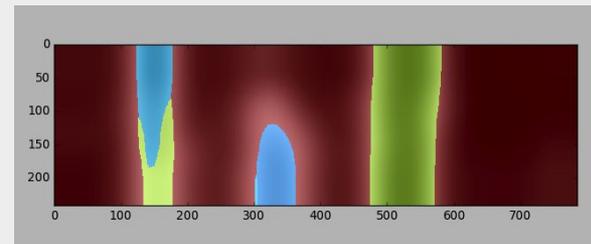
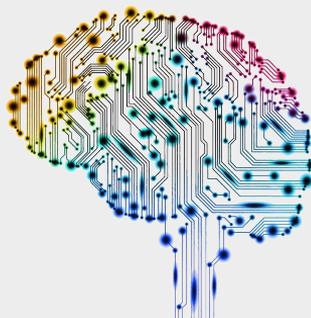
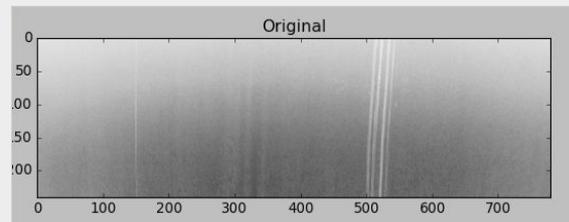
Active learning

De komende paar jaar gaat het mogelijk worden om zelf AI taken te leren

- AI doet met je mee
- Maakt de AI een fout dan kan je het verbeteren
- AI zal vragen stellen om slimmer te worden
- AI legt uit waarom beslissingen worden gemaakt
- De AI wordt steeds slimmer door interactie



Zeer accurate inspectieresultaten



Feedback geven



Vragen stellen



SIS expert



Vera

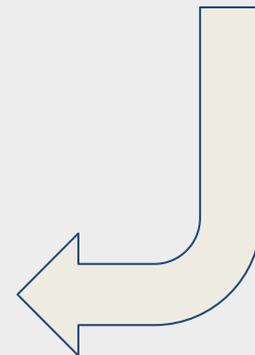
Deep Learning

Fouten aangeven



Gebruikers

Resultaten



Geweldige mogelijkheden



Kunde als toegankelijk product



Tijmen's DL/AL guide

Deep Learning/AI will help you when:

- 1) You want the computer to perform a (repetitive) task
- 2) There are a lot of examples
- 3) Each example has a label (action/prediction/class)
- 4) The samples have complex data as input

Tijmen's Letterbox

In which we will answer questions asked by real customers that wanted to do something with AI

Let's apply the guide together!

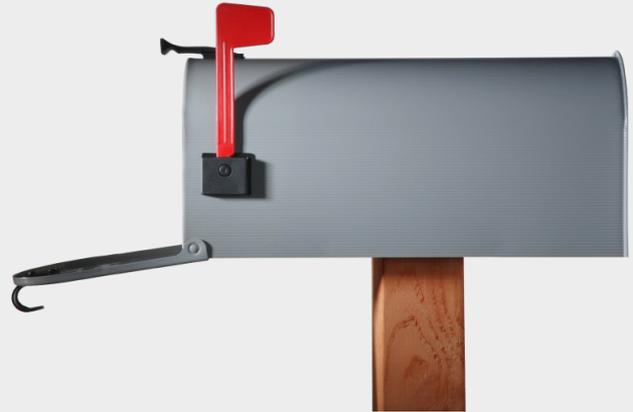


Image -> Caption

Dear Tijmen,

I want to make an algorithm that takes images as input, and generates explanatory captions underneath. I could use this to help blind people understand images online.

Love,

Samus

Is this possible?

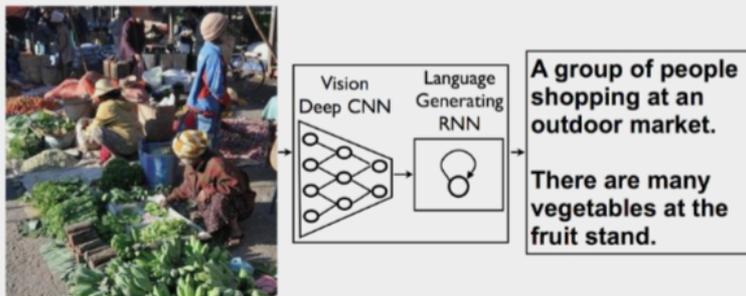


The man at bat readies to swing at the pitch while the umpire looks on.

Tijmen's guide:

- Repetitive task
- Complex input data
- A lot of examples
- Each example has a label

Image -> Caption



YES

Describes without errors	Describes with minor errors	Somewhat related to the image
 <p>A person riding a motorcycle on a dirt road.</p>	 <p>Two dogs play in the grass.</p>	 <p>A skateboarder does a trick on a ramp.</p>

- Repetitive task
- Complex input data
- A lot of examples
- Each example has a label

Image to text (Google 2014)

Boat engine data

Dear Tijmen,

We put 200 sensors in large boat engines we produce.

They produce a lot of data. Parts of the engines break sometimes. Can you make an algorithm that predicts engine failure based on the sensor data?

Love,

Zelda

Is this possible?



Tijmen's guide:

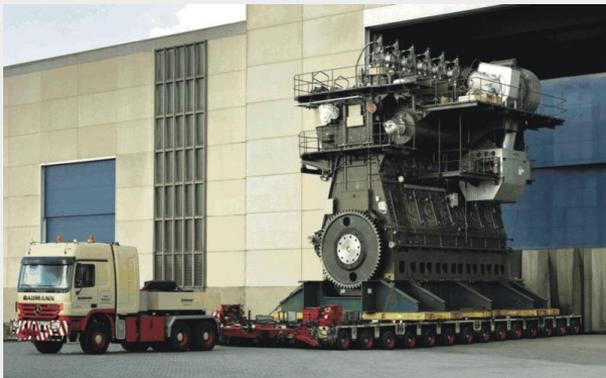
- Repetitive task
- Complex input data
- A lot of examples
- Each example has a label

Boat engine data

Reason:

There are not enough engine breakdowns for the terabytes of sensor data that are being produced!

Big difference between a lot of data and a lot of examples!



Tijmen's guide:

- Repetitive task
- Complex input data
- **A lot of examples**
- Each example has a label

Big Data

Dear Tijmen,

We are an insurance company and we were advised to store a lot of Big Data to prepare us for the future. We stored all sorts of information, like customer interactions, marketing campaigns and risk model predictions. Can you apply Deep Learning to help us find use for this data?

Love,

Jace

Is this possible?



Tijmen's guide:

- Repetitive task
- Complex input data
- A lot of examples
- Each example has a label

Big Data

There is no task in this question!
First think of what you want an algorithm to do!

Deep Learning does not magically find patterns in your
Terabytes of data

Just saving all data without a reason is a great recipe for storing
nonsense and missing important data



Tijmen's guide:

- Repetitive task
- Complex input data
- A lot of examples
- Each example has a label

Eye Scans

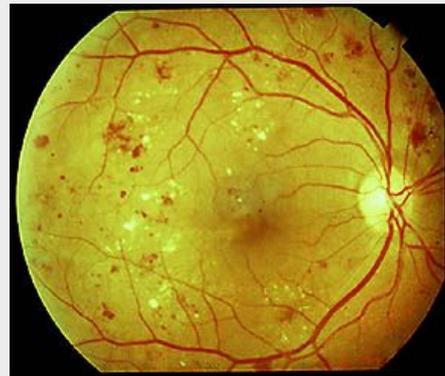
Dear Tijmen,

Our company has around 50.000 eye scans of Diabetes patients. We want an algorithm that automatically determines the severity of eye damage due to diabetes. Each scan has a severity indication made by an eye doctor. The problem is that the input is very complex: damage can be hard to spot as it is very small in the image

Love,

Thrall

Is this possible?



Tijmen's guide:

- Repetitive task
- Complex input data
- A lot of examples
- Each example has a label



Eye Scans

Complexity of the input does not matter. If a human can see it, Deep learning can see it.

The 50.000 images are on the edge of being enough, but if we gather more from practice we can make the algorithm!



Tijmen's guide:

- Computer performs a task
- Complex input data
- A lot of examples
- Each example has a label

Cancer death prediction

Dear Tijmen,

Our hospital chain has a lot of patient data about cancer therapy outcomes. For each patient we have 20 indicative variables, like blood pressure, age and medicine usage. There are about 1.000.000 patients that we have data of. Can you make an algorithm that predicts therapy outcome for patients?

Love,

Kvothe

Is this possible?



Tijmen's guide:

- Repetitive task
- Complex input data
- A lot of examples
- Each example has a label



Cancer death prediction



The input data is not complex.

Deep Learning adds nothing here.

A few thousand patients and simple machine learning algorithm will create the best results you can get

Tijmen's guide:

- Repetitive task
- **Complex input data**
- A lot of examples
- Each example has a label



Reading assistant

Dear Tijmen,

My company processes around 300 very thick digital text reports from other companies, reporting on several business aspects. We only have the manpower to read so much. There are always specific paragraphs and word usages we look for in the reports to make an assessment. A tool that could help us find these paragraphs/words quickly would be great. We do not mind tagging the paragraphs/words that we look for.

Love,

Charlotte

Is this possible?

...that it makes us...
...must feel that their ideas and suggestions...
...n. When they feel that they have l...
alled knowledge management, the 1...
ing, not at the end when you have p...
on their desktops one...
group of in...



Tijmen's guide:

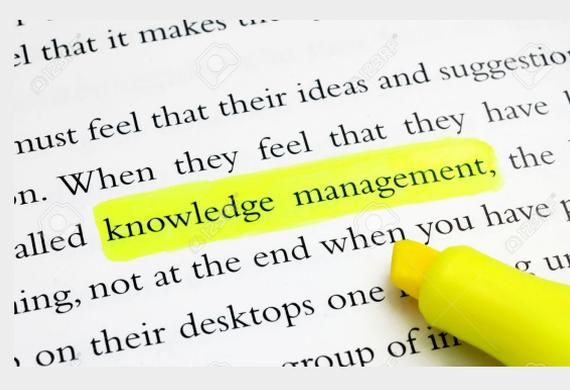
- Repetitive task
- Complex input data
- A lot of examples
- Each example has a label

Reading assistant



If enough examples get annotated this is doable

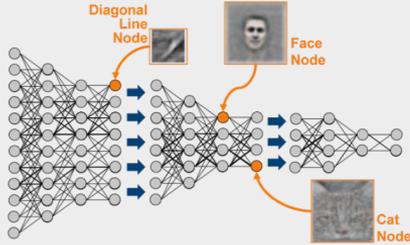
We might be able to use pretrained text algorithms for this...



Tijmen's guide:

- Repetitive task
- Complex input data
- A lot of examples
- Each example has a label

Conclusie



1) Veel van de huidige AI hype komt door Deep Learning/Machine learning

2) AI is niet zo slim, beter uit te leggen als 'teachable software tool'



4) De nabije toekomst van active learning, stelt mensen en AI in staat om steeds beter samen te werken

3) Added-value, zoek naar:

- Repetetieve taken
- Beter gebruik van meer data
 - Huidige modellen verbeteren
 - Beter voorspellingen doen



