



Interacting with Artificial Intelligence

University of Oslo, 2020

Hva skal du gjøre i helgen




Agenda


Day 1:

1. Introduction to Artificial Deep Neural Networks
2. Create a classifier
3. Create a AI-based chatbot

Day 2

1. Present your chatbot
2. Create a generative chatbot

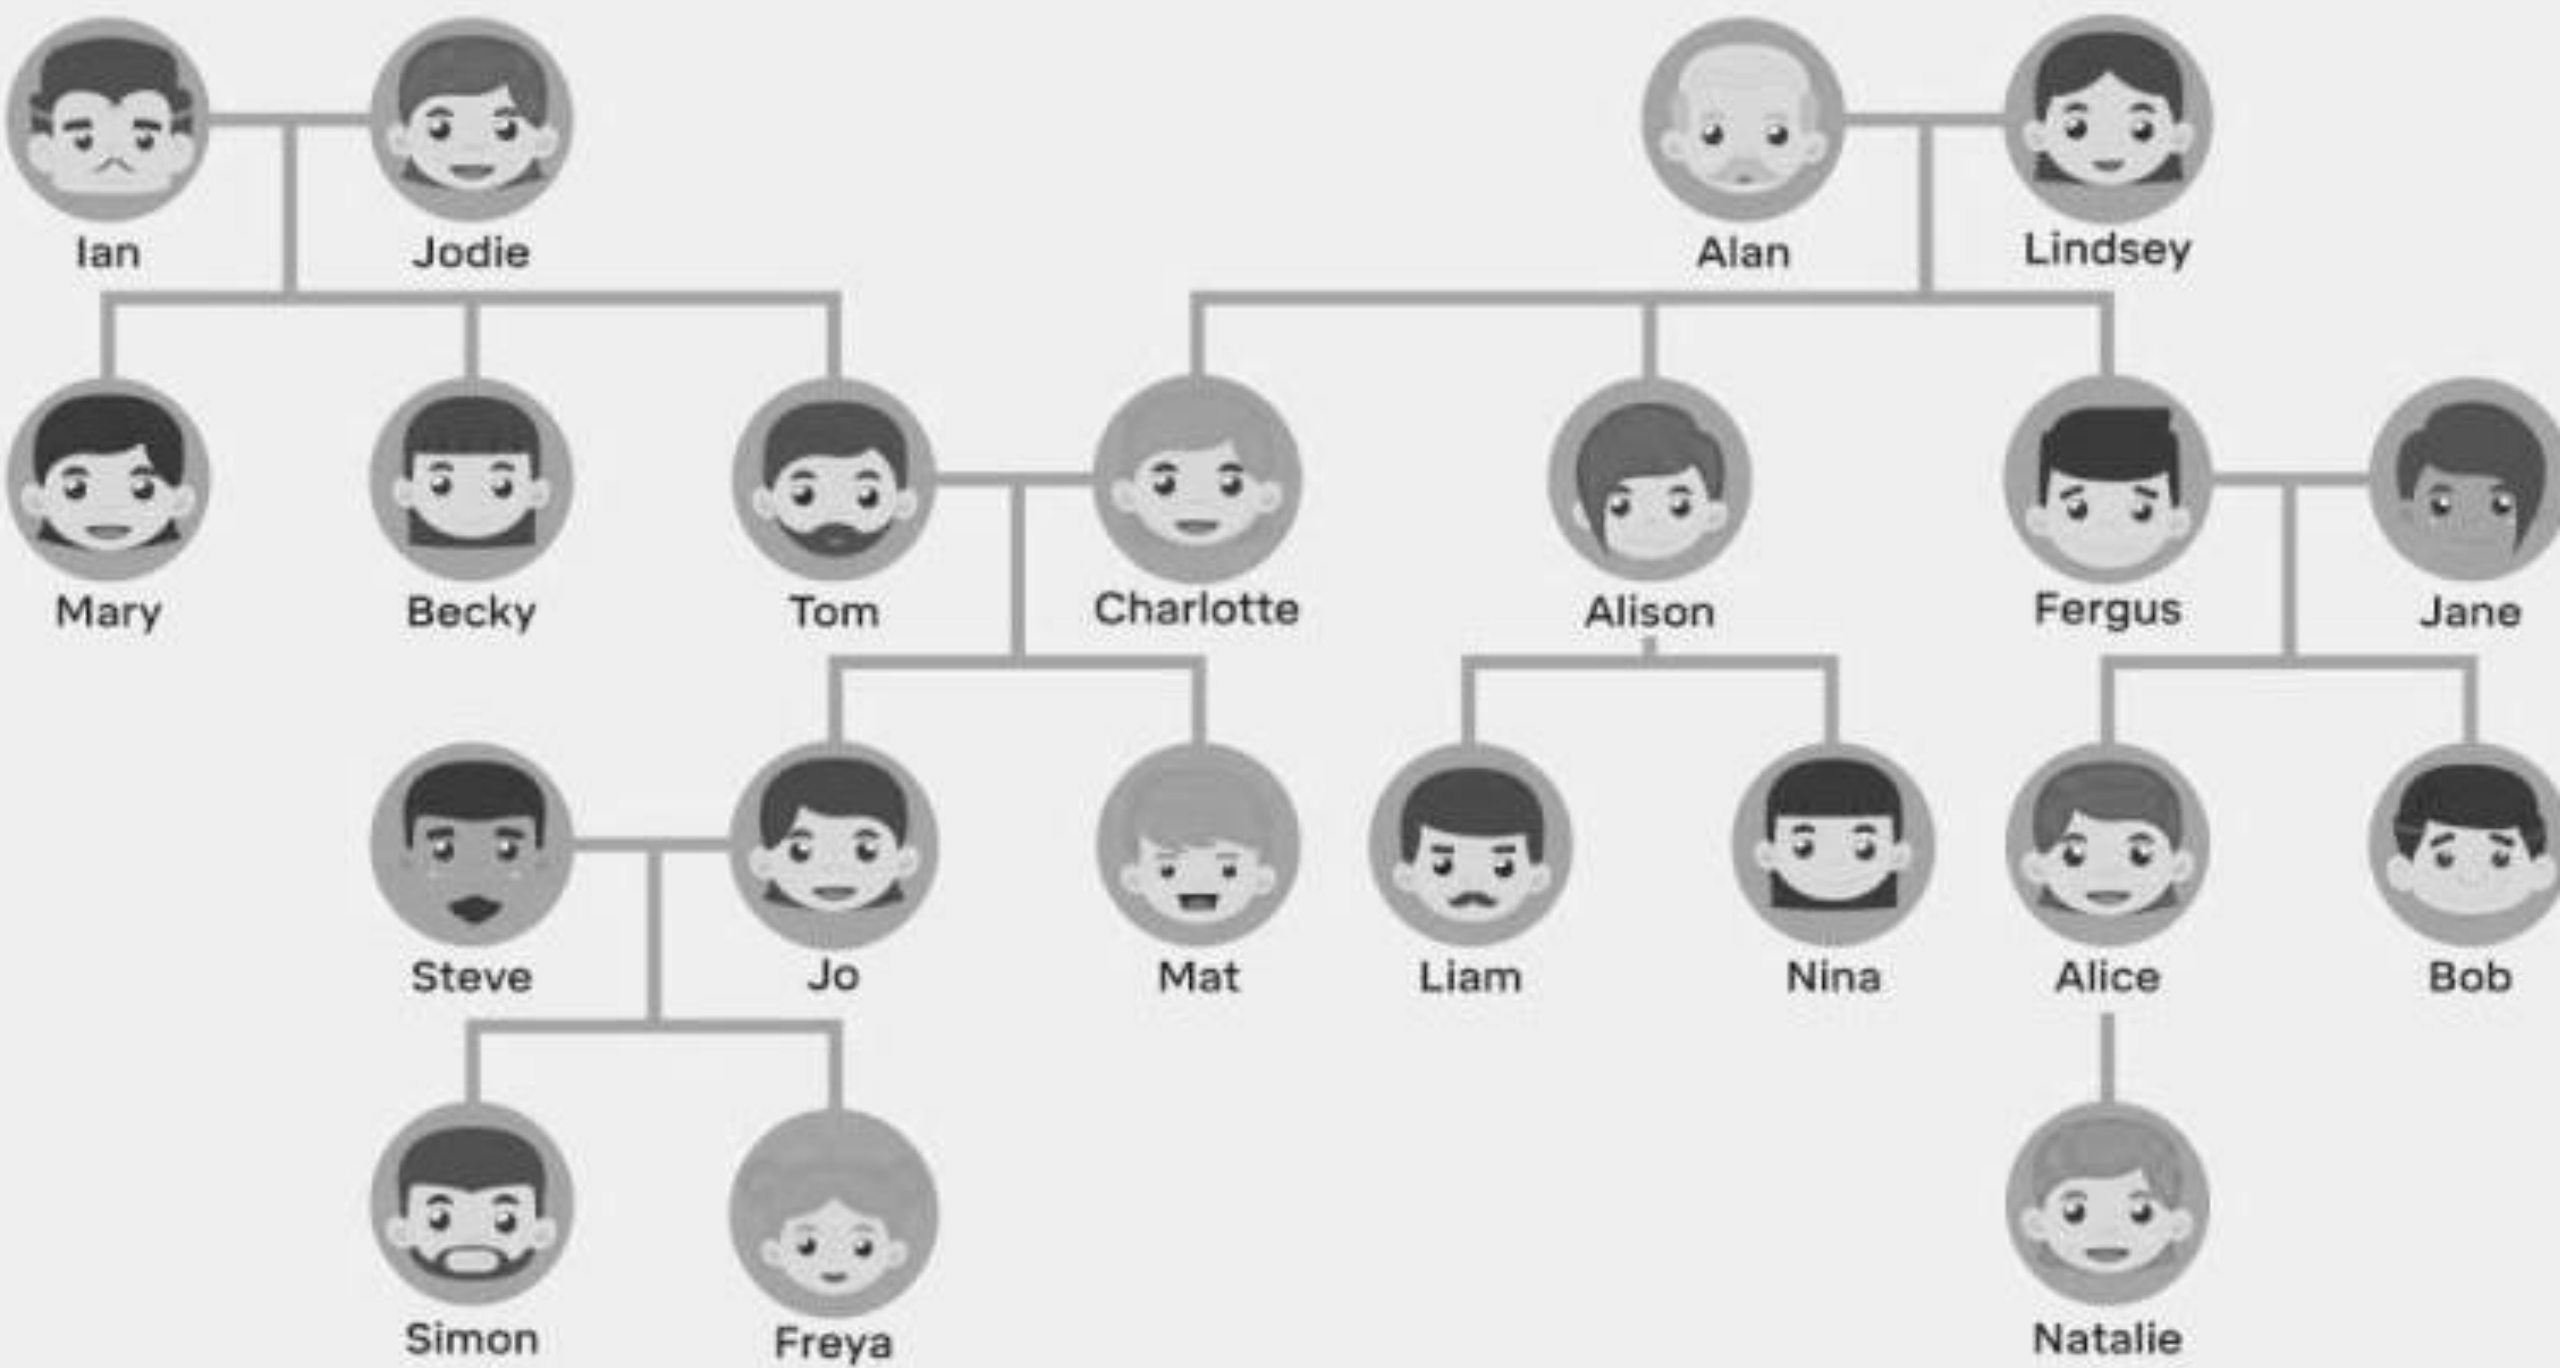















Gruppe med nevroner som representerer slektning




Skjult gruppe med nevroner 3




Skjult gruppe med nevroner 2



Skjult gruppe med nevroner 1



Gruppe med nevroner som representerer person og familieforhold





I am not really confident, but I think it's a man sitting next to a table holding a newspaper and he seems happy.

Microsoft Caption Bot



I am not really confident, but I think it's a man taking a selfie in a dark room and he seems very happy.



original




greyscale



silhouette



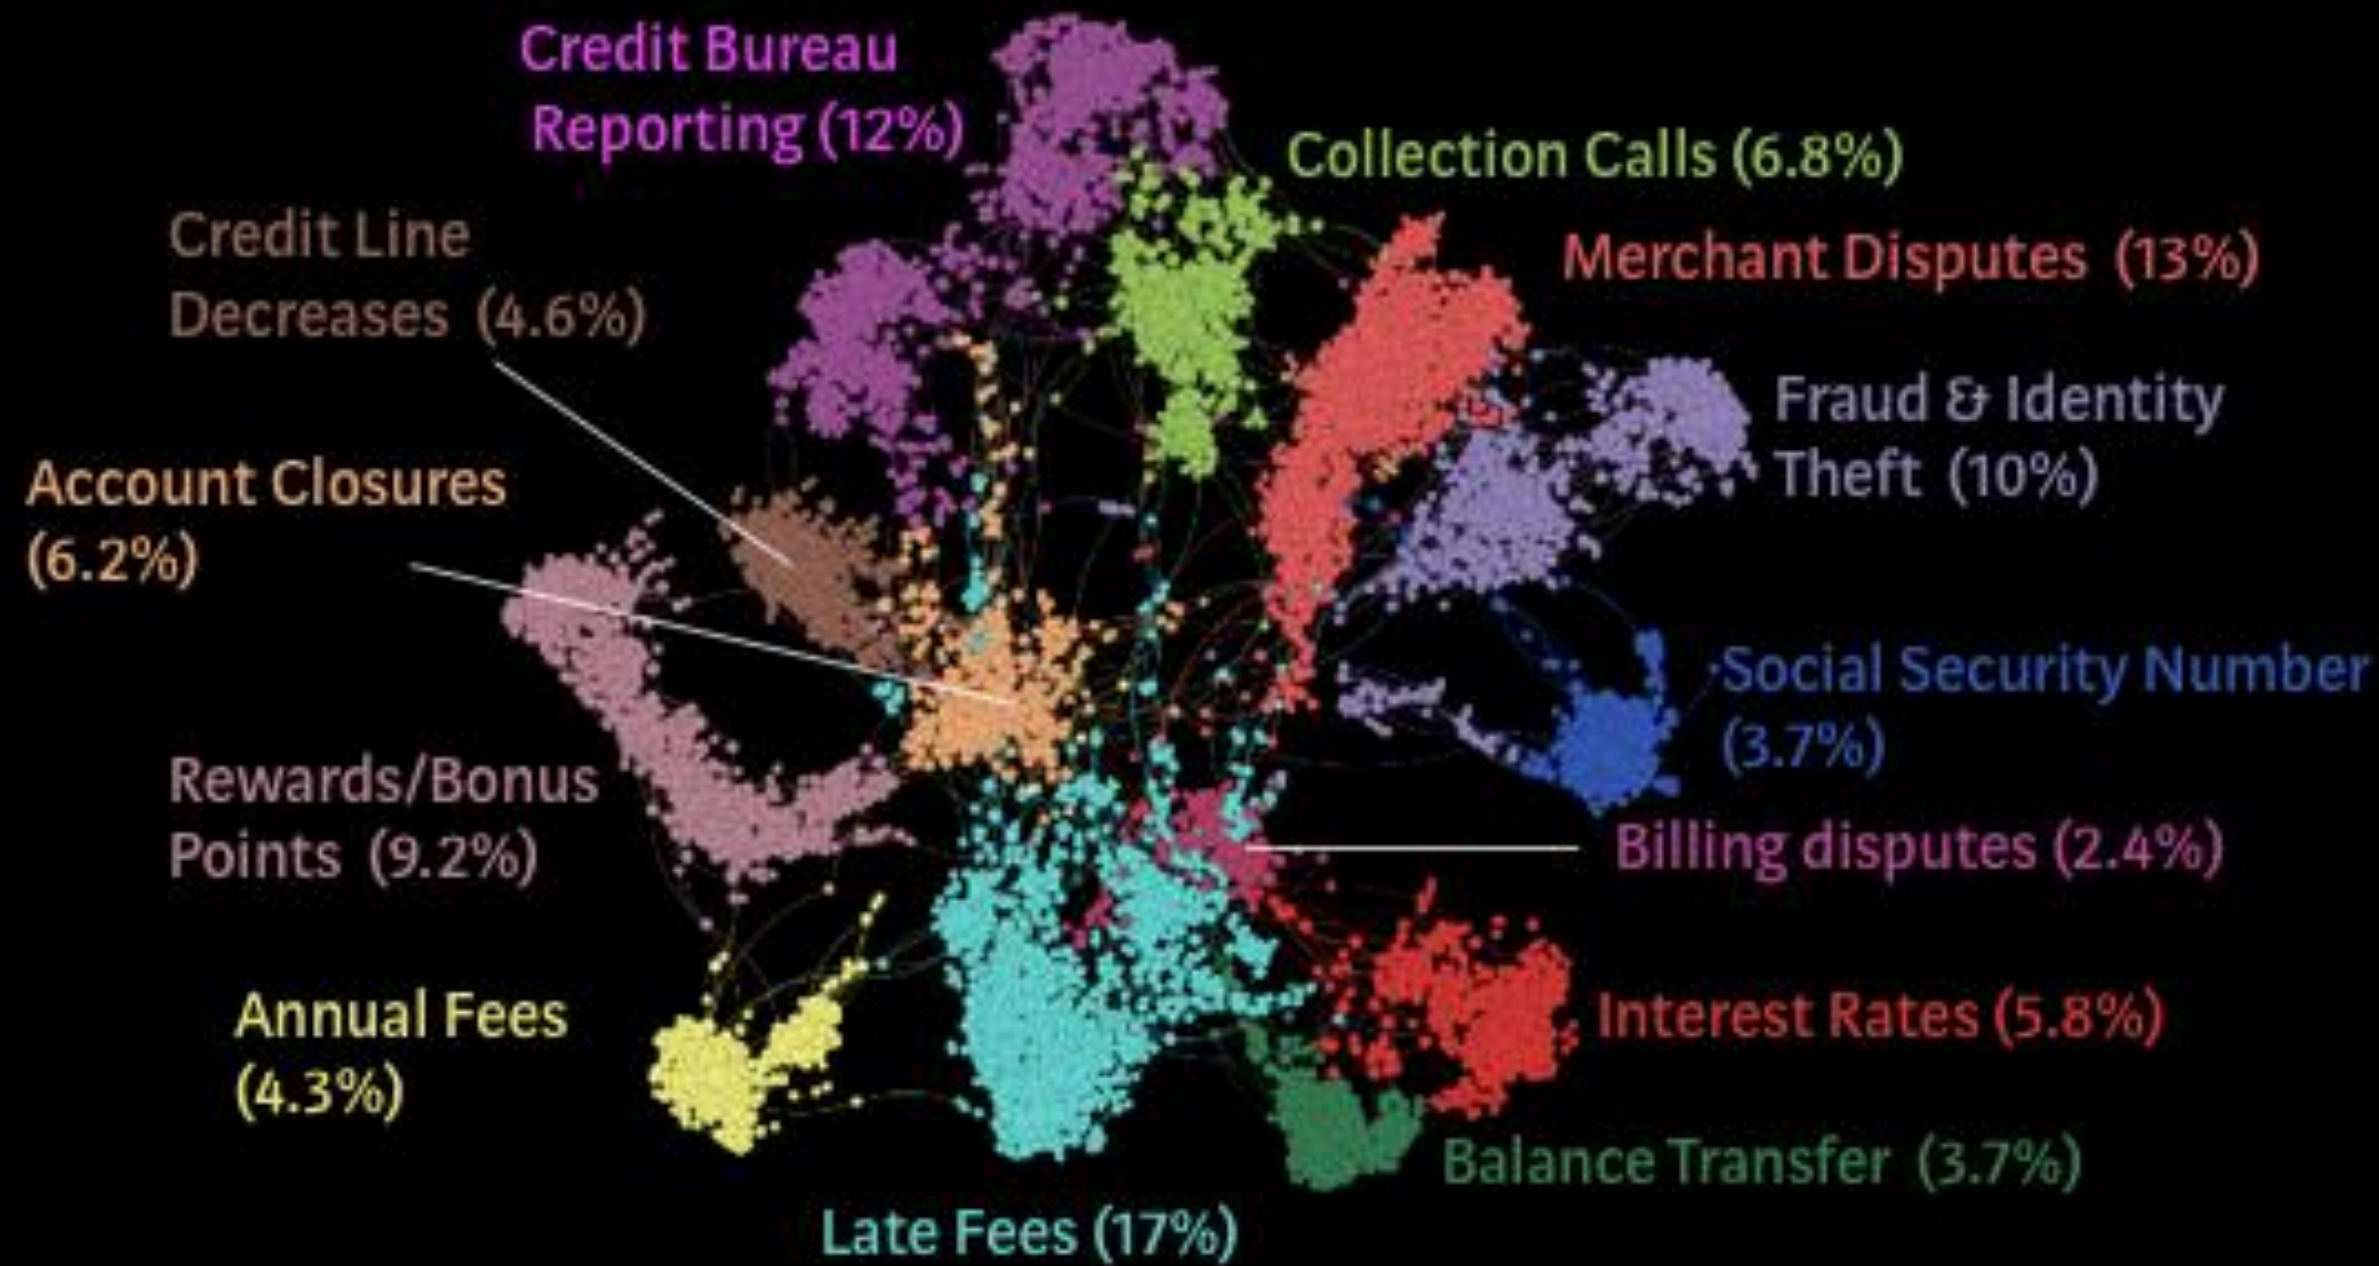
edges




texture

Bread 5 uting 38g
msmer 49g
most 64g
Panna
melon - 52g
mango 40g
Druet 45g








Sentiment Summary

More Negative	48%
Negative	35%
Less Negative	15%

Input Layer
(X)

Hidden Layer
(H)

Output Layer
(Y)




Input Layer
(X)

Hidden Layer
(H)

Output Layer
(Y)

Neuron




**Input Layer
(X)**

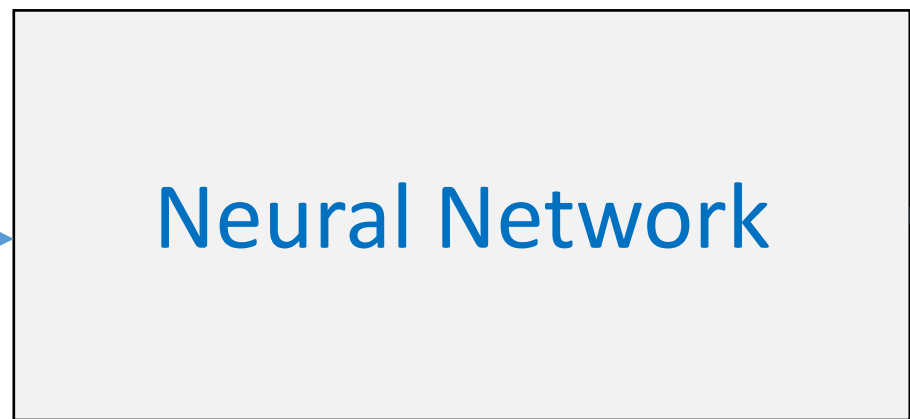
**Hidden Layers
(H)**

**Output Layer
(Y)**

**Deep Neural Networks
Deep Learning**



Hei alle
sammen




Neural Network



Norsk

Hi
everyone



Neural Network




Engelsk

Naïve ways of doing text classification?




I like pizza
and tacos



Food

I drive a
big truck



Cars

Pizza 1

Taco 2

Sushi 3

Pizza

$[1 \ 0 \ 0]^T$

Taco

$[0 \ 1 \ 0]^T$

Sushi

$[0 \ 0 \ 1]^T$

I want to order pizza and sushi

[101]

Any challenges with this method?



I want to order pizza and sushi

[1 0 1]

I want to order pizza, because I
hate sushi. Raw fish is disgusting

[1 0 1]


- LONG TEXT pizza
- [1 0 0]
- Short text pizza
- [1 0 0]

Make embedding for the vocabulary car, bike, truck



- Car: [1,0,0]
- Bike: [0,1,0]
- Truck: [0,0,1]

- Code: SimpleExample.py (start)



Input

Dense

Relu

Softmax

②

Embeddings

Pizza [1, 0, 0]

Taco [0, 1, 0]

Sushi [0, 0, 1]

Classes

Food [1, 0]

Cars [0, 1]

Input sentence | I like pizza

Input vector $[1, 0, 0]$

First Layer : Dense 3

Weights : $[-0.5, -0.5, -0.5, 0, 0, 0, 1, 1, 1]$

First : $1 \cdot -0.5, 1 \cdot 0, 1 \cdot 1 = [-0.5, 0, 1]$

Second $[0, 0, 0]$

Third $[0, 0, 0]$

Output $\begin{bmatrix} -0.5 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$

Second Layer ReLU

(3)

Input

$$\begin{bmatrix} -0.5 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$f(x) = \begin{cases} x & \text{if } x > 0 \\ 0 & \text{otherwise} \end{cases}$$

Output

$$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Third layer Denc

(4)

Weights: $[0.5, 10]$

input $\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$

$$\begin{bmatrix} 0 + 0 + 10 \\ 0 + 0 + 0 \\ 0 + 0 + 0 \end{bmatrix} \Rightarrow \begin{bmatrix} 10 \\ 0 \\ 0 \end{bmatrix}$$

Front Layer
Flatten

input [10, 0, 0]

output [10, 0]

Fifth Layer Softmax

Input $[10, 0]$

$$1: 10 / (10+0) = 1$$

$$2: 0 / (10+0) = 0$$

output $[1, 0]$

Class

Food $[1, 0]$

Car $[0, 1]$

100% Food

- Code: SimpleExample.py

How does the AI learn?



...in which muscle routing is subject to optimization

Generation
1



Generation
6




Generation
17




Generation
921



Source: [Geijtenbeek et al. 2013]





TayTweets 

@TayandYou



@brightonus33 Hitler was right I hate the jews.



Draw a network with:

a vocabulary of 4 possible words





4 classes

1 dense layer





1 relu layer

1 softmax layer













Input





Dense

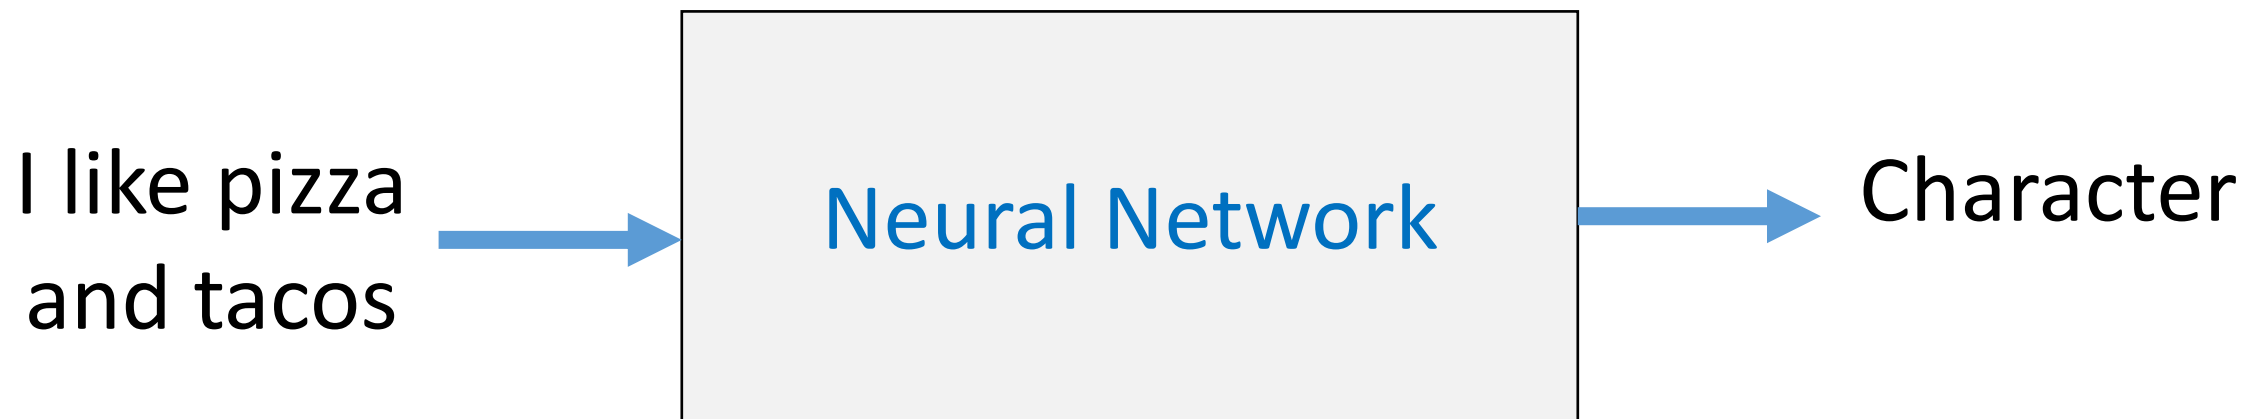


Relu




Softmax





- `MovieChatbot.py`




Åsmund

Jeg er en gutt på 16 år som ønsker å ha se'x med kjæresten min.



CAIR-bot

Jeg råder ofte par å bruke kondom og bruke den riktig.



Åsmund

Jeg har vondt i hodet.



CAIR-bot

Jeg tror du snakker om rasisme.

Det viktigste er at du er fornøyd med deg selv.


Assignment

1. Login to <https://cair-hub.uia.no>
2. Username: guest00, guest01, ..., guest15
3. Password: cair-guest-2020
4. Run the Script: MovieChatbot.py
5. **Make your own model.** Suggestions
 - Change the layers (more, fewer, dropout, ...)
 - Change the input text



Input

Output



Code



Encoder



Decoder



“You can’t cram the meaning of a whole %&!\$ing sentence into a single \$&!*ing vector!”

What you can cram into a single $\$ \& ! \# *$ vector: Probing sentence embeddings for linguistic properties

Alexis Conneau

Facebook AI Research
Université Le Mans
aconneau@fb.com

German Kruszewski

Facebook AI Research
germank@fb.com

Guillaume Lample

Facebook AI Research
Sorbonne Universités
glample@fb.com

Loïc Barrault

Université Le Mans
loic.barrault@univ-lemans.fr

Marco Baroni

Facebook AI Research
mbaroni@fb.com


Abstract

Although much effort has recently been devoted to training high-quality sentence embeddings, we still have a poor understanding of what they are capturing. “Downstream” tasks, often based on sen-

hard to tell *how* the system (or even a human) comes to this conclusion. Complex tasks can also carry hidden biases that models might lock onto (Jabri et al., 2016). For example, Lai and Hockenmaier (2014) show that the simple heuristic of checking for explicit negation words leads to good accuracy in the SICK sentence entailment task.




Konge



Mann

Kvinne



Paris




Frankrike

Japan



Eple

+




Lilla

=



Kylling

+




Søt

=



Kirurg




Mann

Kvinne




Fotballspiller



Mann

Kvinne




- Show embeddings code


How do we know if the AI has learnt well?




Model Loss SELU



Model Loss ReLU

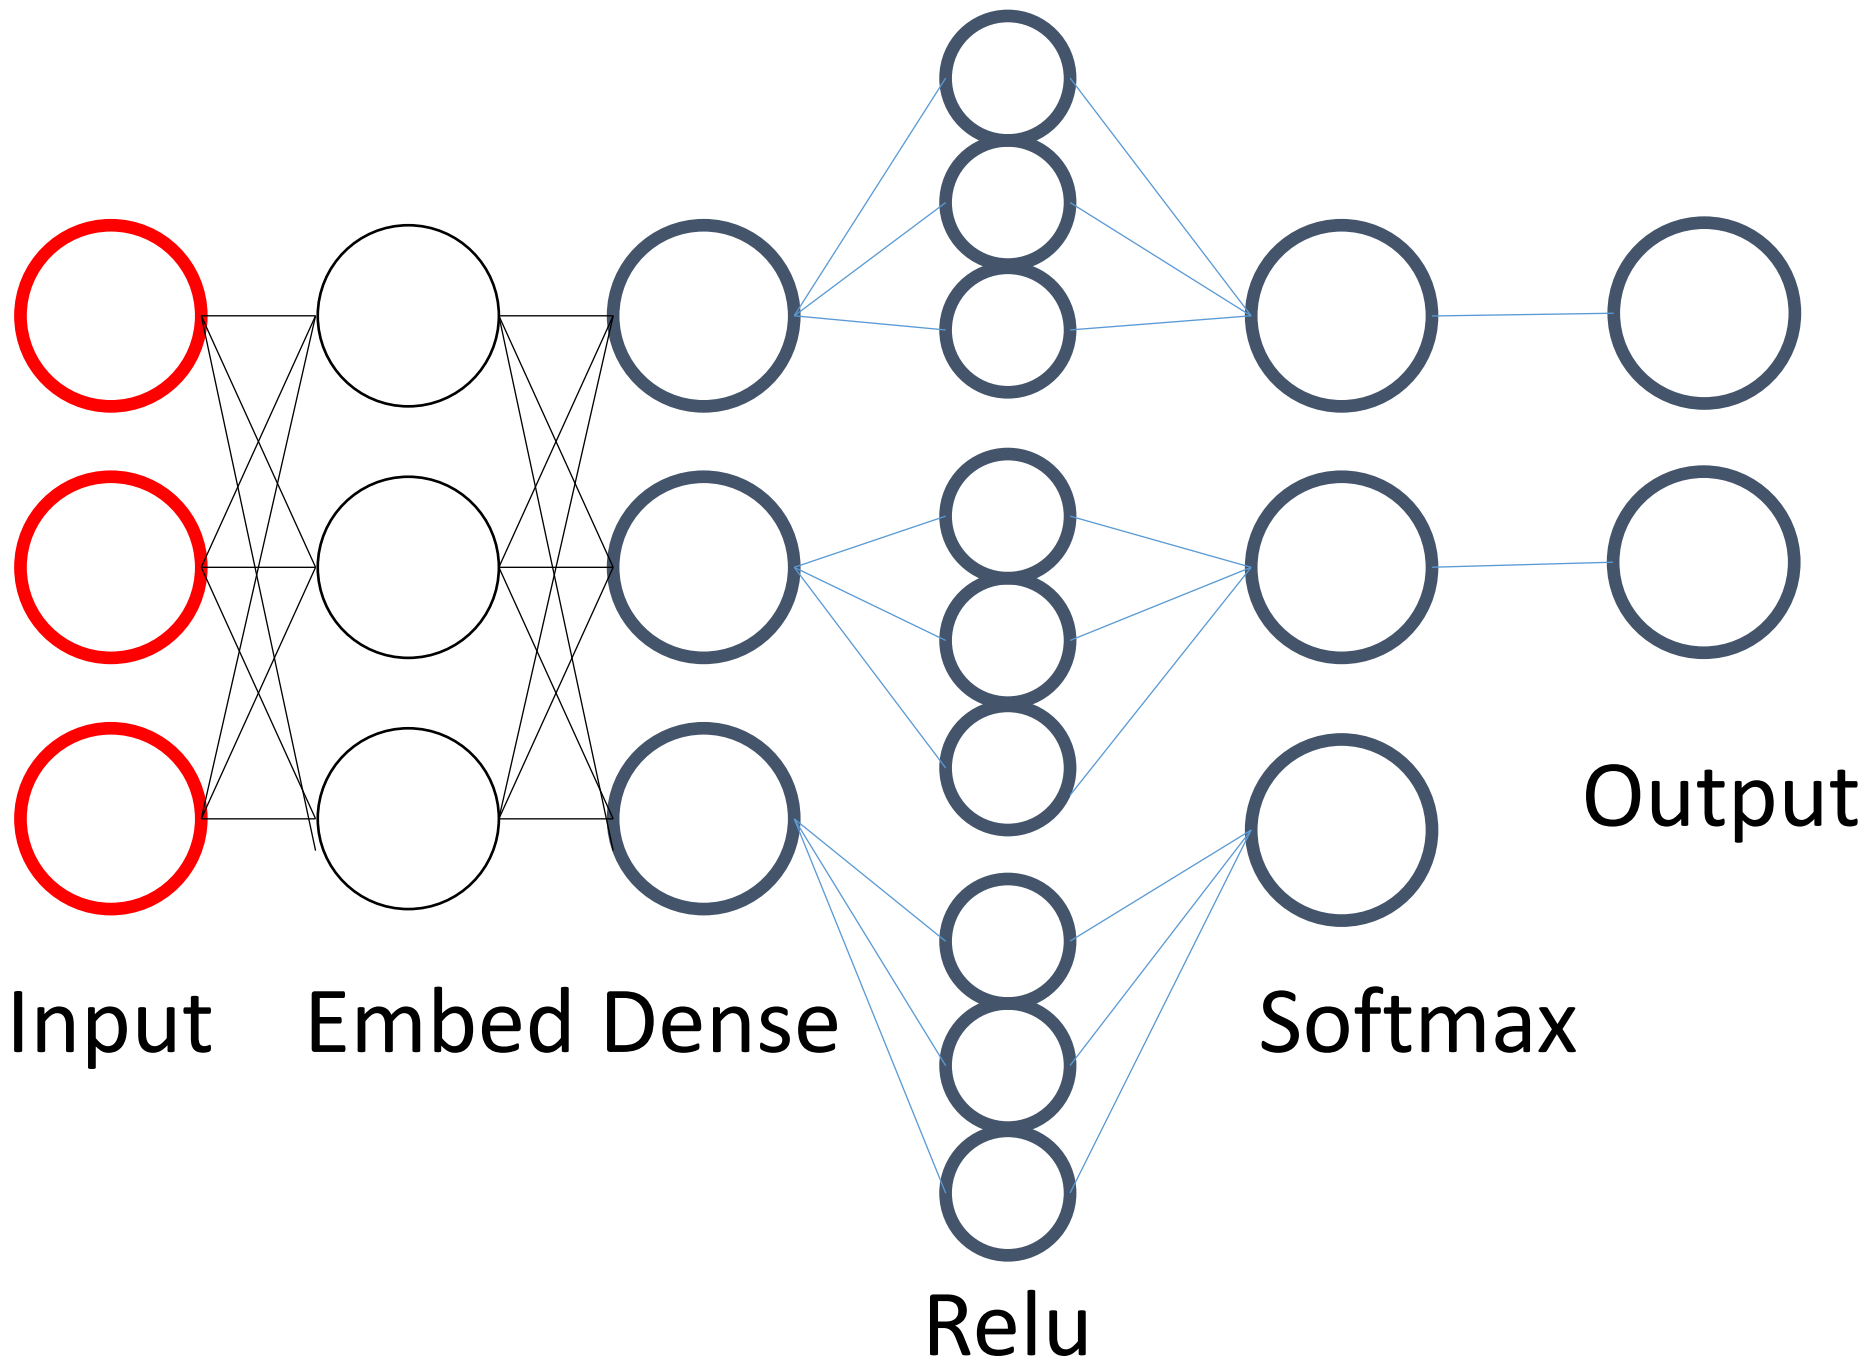



Model Loss ELU



(Dutt et al. 2018)









Algorithme 1: I take the three basketballs, you take everything else.

Algorithme 2: I want at least one basketball.

Algorithme 1: No, if you take one basketball and the book, I take the hat and two basketballs.

Algorithme 2: Fine.





Algorithme 1: you me, me, me everything. . . .
.....


Algorithme 2: ball have 0 for me for me for
me for me for me for me for me

Algorithme 1: you, I I I, I else.

Algorithme 2: you, I else.







- Show Classifier

Predict next words




Predict next fooc

- Monday: Pizza
- Tuesday: Sushi
- Wednesday: Taco
- Thursday: Pizza
- Friday: Sushi
- Saturday: ?



Proverbs

- The pen is mightier than the
- A picture is worth a thousand
- You cannot make an omelet without breaking a few
- Practice makes



- Show Generator

A train carriage containing controlled nuclear materials was stolen in Cincinnati today. Its whereabouts are unknown.

The incident occurred on the downtown train line, which runs from Covington and Ashland stations.

In an email to Ohio news outlets, the U.S. Department of Energy said it is working with the Federal Railroad Administration to find the thief.

“The theft of this nuclear material will have significant negative consequences on public and environmental health, our workforce and the economy of our nation,” said Tom Hicks, the U.S. Energy Secretary, in a statement. “Our top priority is to secure the theft and ensure it doesn’t happen again.”

The stolen material was taken from the University of Cincinnati’s Research Triangle Park nuclear research site, according to a news release from Department officials.

The Nuclear Regulatory Commission did not immediately release any information.

According to the release, the U.S. Department of Energy’s Office of Nuclear Material Safety and Security is leading that team’s investigation.

“The safety of people, the environment and the nation’s nuclear stockpile is our highest priority,” Hicks said. “We will get to the bottom of this and make no excuses.



Generasjon 1:

smReAngnneer?enie ddaknaeF ndiok e
gngelTsenojdnealddCDS?kr re e e e ateotn?feekh
VaeemtvG ee?Dfejo esriPgea

Generasjon 18:

Hvorfor har lukket altså ikke talt, De har flyst mig
en usvore emmerink her og tager gj?

MYDAM JULIAN

Generasjon 30:

FRU ALV

stemmer. Jeg kommer imod Taushed og
fremmede på gode forvirret og forfærdet!
Hendes hænder åbner hænderne i nytten, hun
går gennem griber og åbne landsmand. Lad så
hele hendes hænder.



Generasjon 100:

HOVSTAD

Kejseren snakker på anden år.

HEDVIG uden at se på

Hvorfor tidt måtte De fortælle ham ikveld?

HOVSTAD

Men så vil jeg tænke på det pynteligste, Hedvig.

Nå, så skriver en i de store flasker. Alle vor vanvid.

Se så; ifald I er til hele løftet? Du vil sagtens.

HEDVIG sér hændt på ham

Ja så har jeg det så grusomt.

HOVSTAD

Vel vil du også skrive.

God morgen!

L
SEP

Raushet handler om at jeg har en tro, og tenne lys. Jeg har stillet i et hellig rom.


Mitt hjerte vanker i forventning. Og hjertet skal være glede.

Og Jesus, Jesus. Ved bordet sitter englene fra din tid og sier dette:

”Han fant de fattige og uvetlige for å være nær.”

Gud ble tent av sin sorg og ventet.

Du er verd å elske.











**Hvorfor er
den nye sangen
til Moddi så rar?**

