

Essentials of Deep Learning

Carlos Cotrini

2019

Agenda

- Many machine learning models cannot make “insights” on data.
- Neural networks.
- How neural networks are built.
- Convolutional neural networks for images.
- Recurrent neural networks for text.



POKER

HAND RANKINGS



#1 ROYAL FLUSH



#6 STRAIGHT



#2 STRAIGHT FLUSH



#7 THREE OF A KIND



#3 FOUR OF A KIND



#8 TWO PAIR



#4 FULL HOUSE



#9 ONE PAIR



#5 FLUSH



#10 HIGH CARD

Many ML algorithms cannot make insights



Full house







Royal flush

- Logistic model:

49%

Of course, you
don't need ML for
this

Rank 1	Suit 1	...	Rank 5	Suit 5	Type
2		...	8		FH
10		...	A		RF

Many ML algorithms cannot make insights



Full house



Royal flush

- Logistic model:

99.5%

Rank 1	Suit 1	...	Rank 5	Suit 5	# A	#2	...	# ♠	# ♥	Type
2	♦	...	8	♥	0	3		1	2	FH
10	♥	...	A	♥	1	0		0	5	RF

Many ML algorithms cannot make insights







Full house



Royal flush

- Neural network:


99%

Rank 1	Suit 1	...	Rank 5	Suit 5	Type
2		...	8		FH
10		...	A		RF


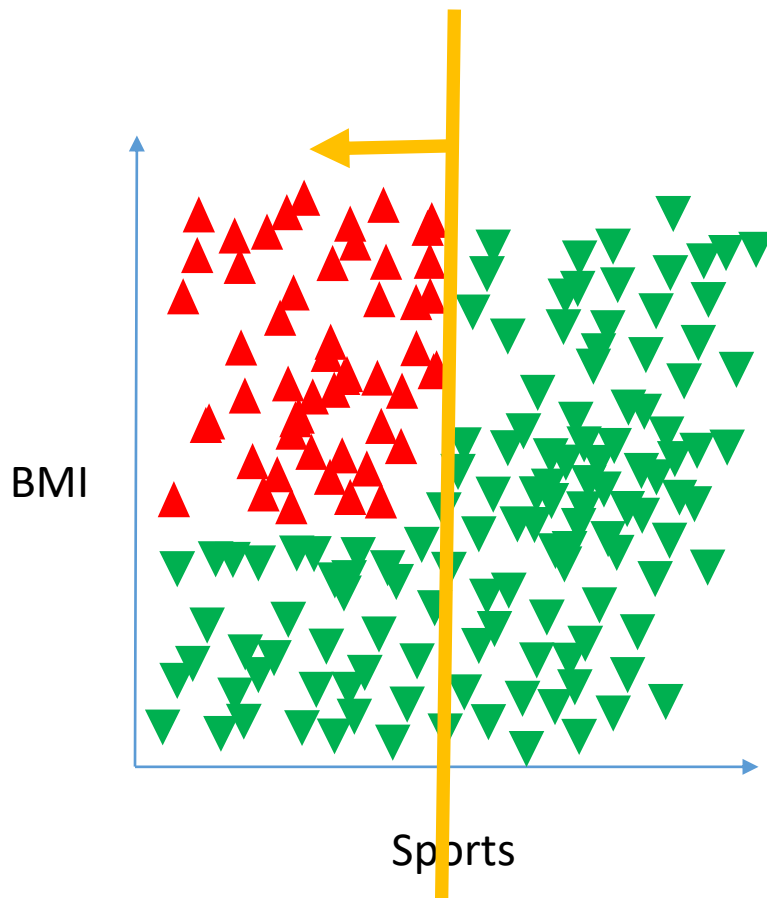
Agenda

- Many machine learning models cannot make “insights” on data.
- Neural networks.
- How neural networks are built.
- Convolutional neural networks for images.
- Recurrent neural networks for text.

Neural networks are assemblies of logistic models





Neural networks are assemblies of logistic models





$$R = \sigma(0 \times BMI - 1.5 \times Sports + 3)$$

Neural networks are assemblies of logistic models





$$R = \sigma(0 \times 26 - 1.5 \times 4 + 3)$$

Neural networks are assemblies of logistic models



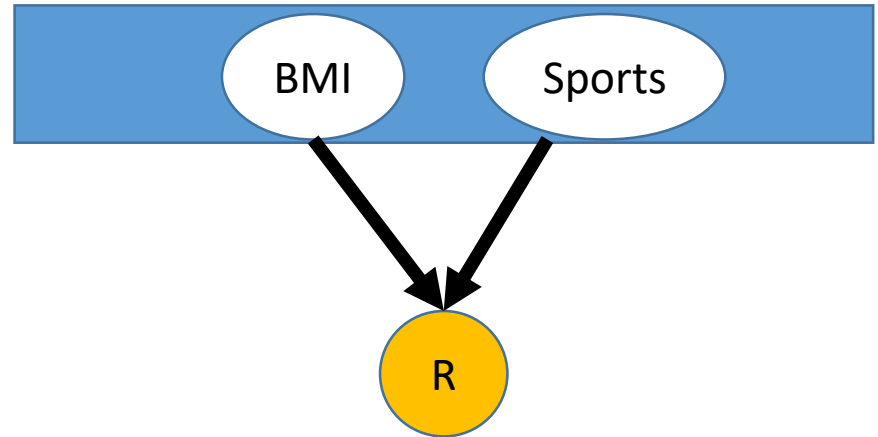
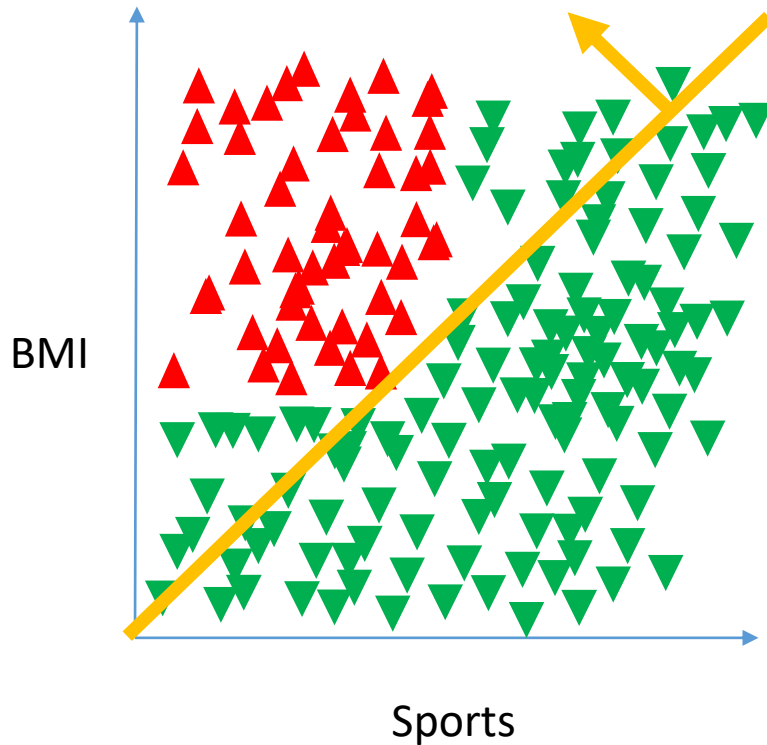
$$R = \sigma(-3)$$

Neural networks are assemblies of logistic models



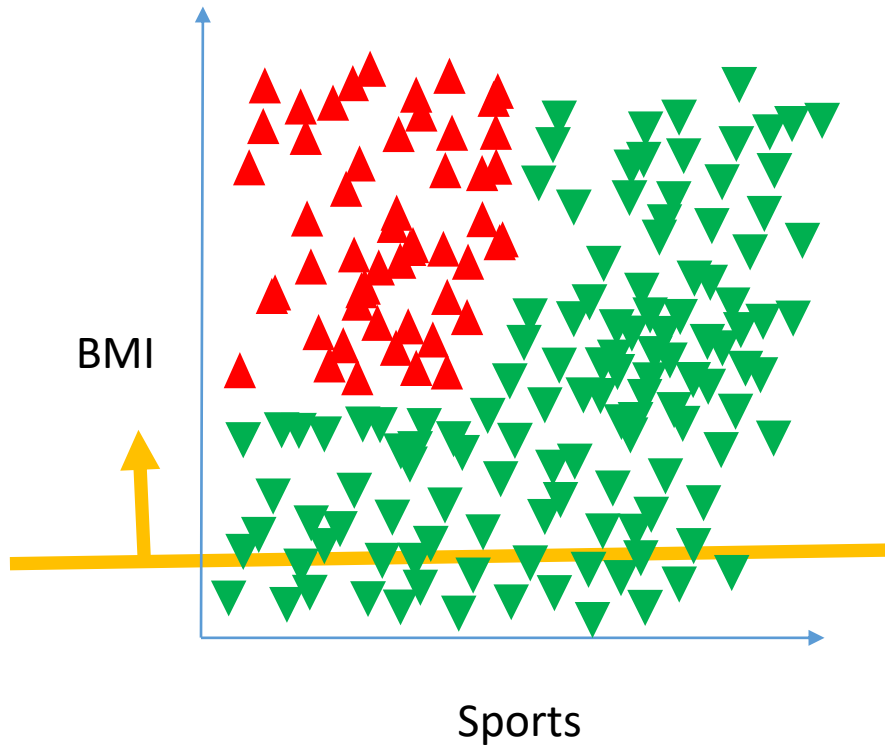

$$R = 0.12$$

Neural networks are assemblies of logistic models



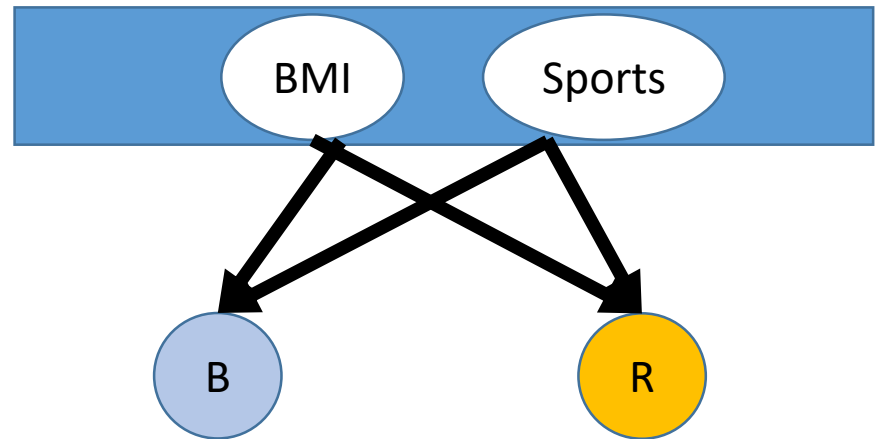
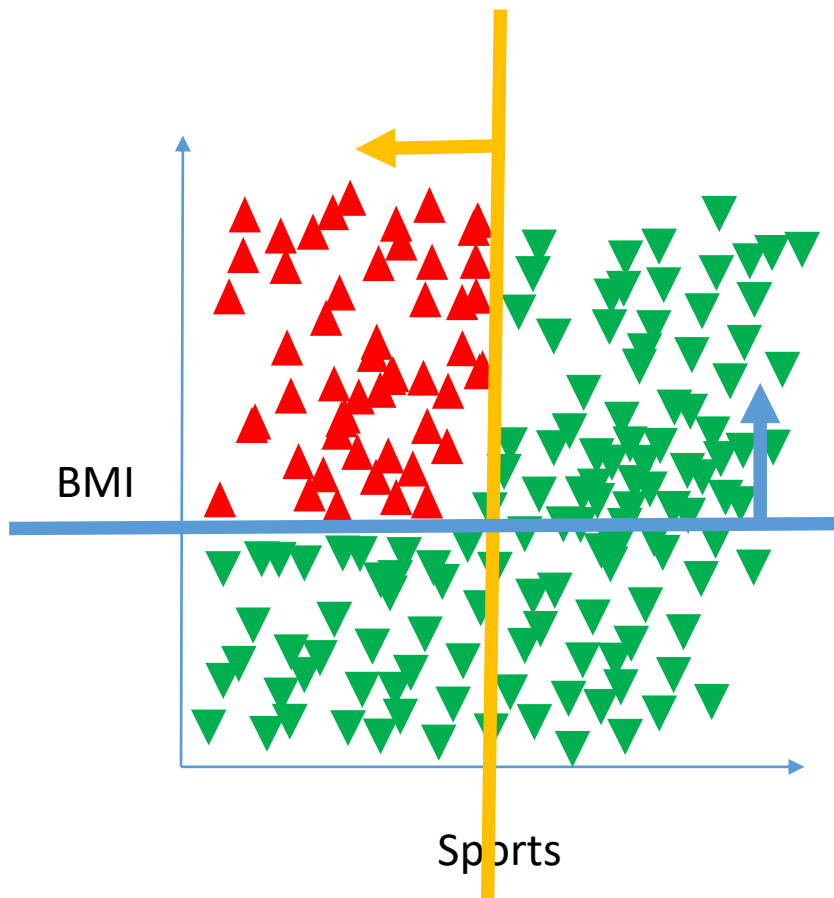
$$R = \sigma(-4 \times BMI + 1 \times Sports + 60)$$

Neural networks are assemblies of logistic models





$$R = \sigma(1.2 \times BMI + 0 \times Sports - 21)$$

Neural networks are assemblies of logistic models



$$R = \sigma(0 \times BMI + 1.5 \times Sports - 3)$$
$$B = \sigma(-2.5 \times BMI + 0 \times Sports - 27)$$

Neural networks are assemblies of logistic models





$$R = \sigma(0 \times BMI + 1.5 \times Sports - 3)$$

$$B = \sigma(-2.5 \times BMI + 0 \times Sports - 27)$$

$$y = \sigma(1 \times R + 1 \times B)$$

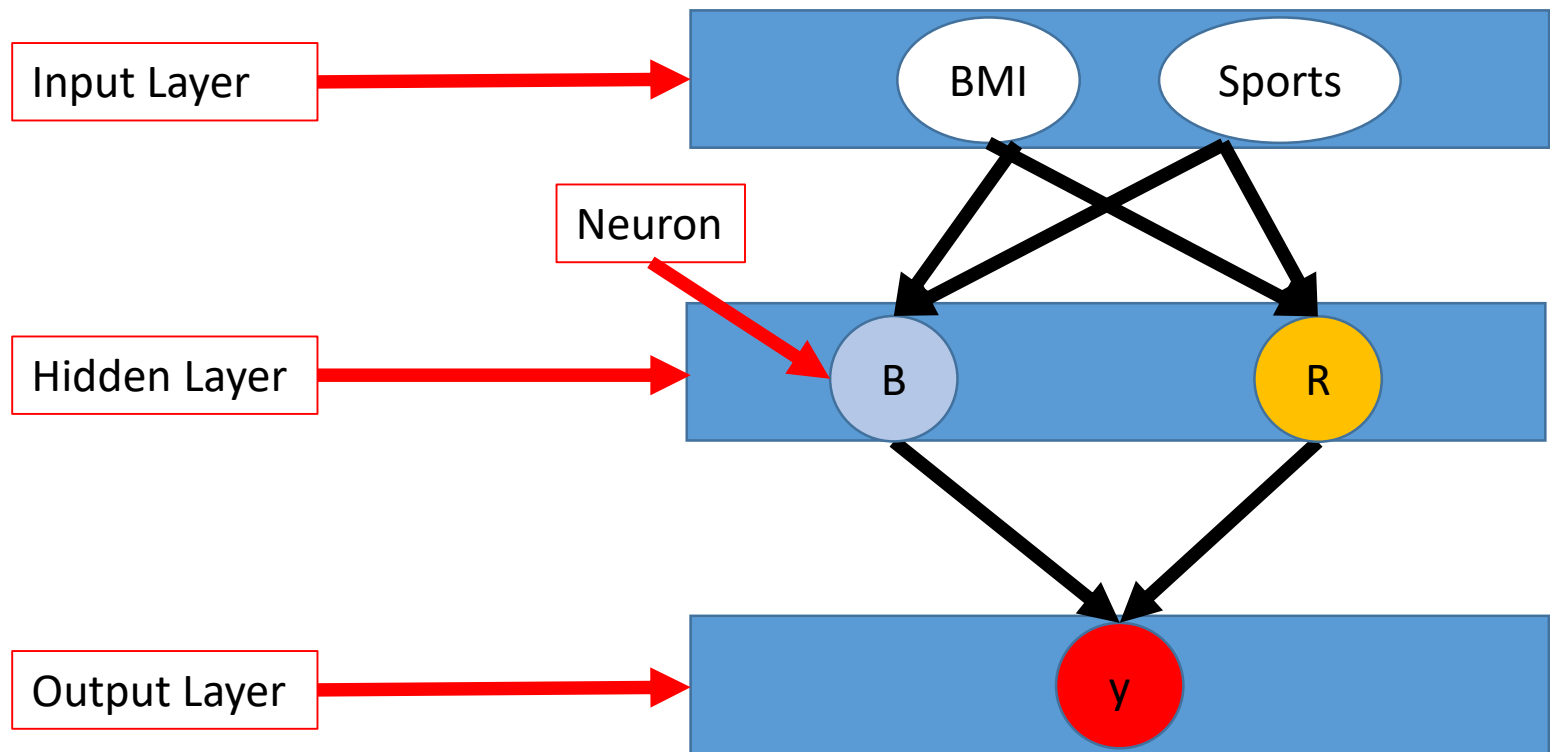
Neural networks are assemblies of logistic models



$$R = \sigma(? \times BMI + ? \times Sports + ?)$$

$$B = \sigma(? \times BMI + ? \times Sports + ?)$$

$$y = \sigma(? \times R + ? \times B)$$



$$R = \sigma(? \times BMI + ? \times Sports - ?)$$

$$B = \sigma(? \times BMI + ? \times Sports - ?)$$

$$y = \sigma(? \times R + ? \times B)$$

Predicting poker hands




One-hot
encoding

Are there more
than 2 aces in
the hand?

Are all cards of
the same suit?

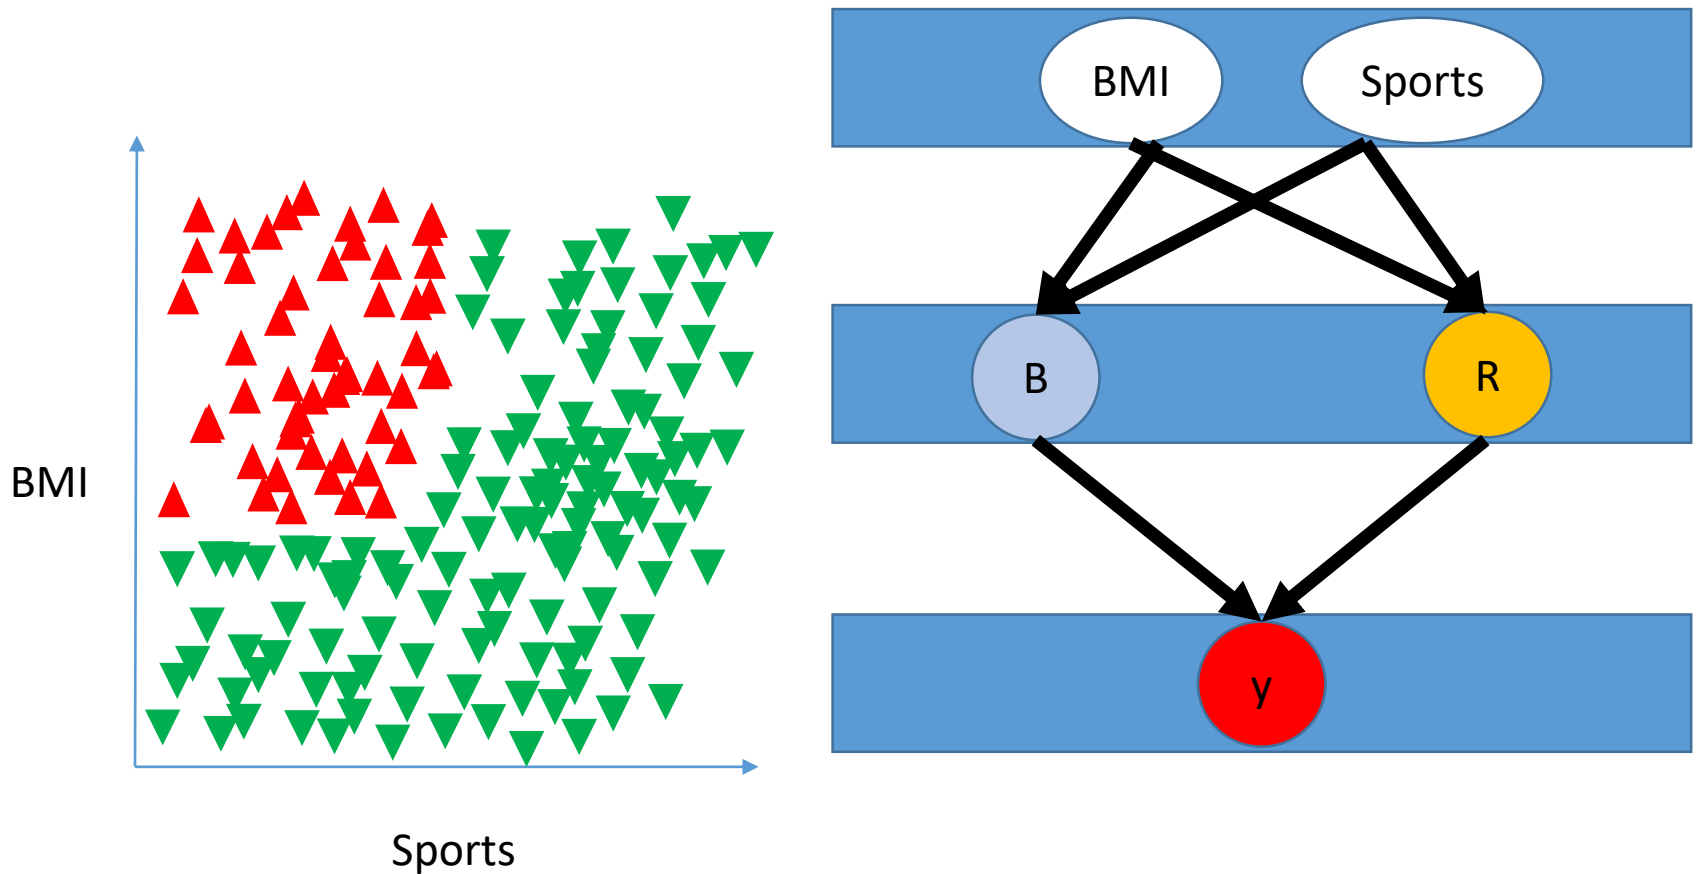
Is it a straight?





Agenda

- Many machine learning models cannot make “insights” on data.
- Neural networks.
- How neural networks are built.
- Convolutional neural networks for images.
- Recurrent neural networks for text.



Gradient descent: How neural networks are built




1) Make a guess for the lines





2) Compute gradients for all lines




3) Slightly move lines as indicated by the gradients





4) Go to step 2



2) Compute gradients for all lines



3) Slightly move lines as indicated by the gradients



Algorithm for training a neural network

1. Make a guess for the lines.
2. Compute gradients.
3. Slightly move lines.
4. Go to (2).

Algorithm for training a neural network

1. Make a guess for the lines.
2. For $i = 1 \dots \text{NumEpochs}$
 - a) Compute gradients.
 - b) Slightly move lines.

Each iteration in step (2) is called an **epoch**.

You must indicate (ahem! guess) the number of epochs before calling this algorithm.

Flower classification



Iris setosa



Iris tectorum




Iris latifolia

Data representation

Sepal length	Sepal width	Petal length	Petal width	Is setosa?
5.1	3.5	1.4	0.2	1
2.1	1.2	3.3	3.2	0
3.1	1.6	2.2	4.1	1
2.2	4.1	1.3	1.4	1


Data representation

- $X[i, j]$: Value of column j for flower i .
- $y[i]$: 1 if flower i is an iris setosa and 0 otherwise.



Script organization

1. **Input parsing:** Read X and y.
2. **Network architecture:** Define neural network layers.




3. **Compilation and training:** Compile and train neural network.


Agenda

- Many machine learning models cannot make “insights” on data.
- Neural networks.
- How neural networks are built.
- Convolutional neural networks for images.
- Recurrent neural networks for text.


Image recognition




Happy




Happy




Happy



Not happy




Not happy




Not happy

Image recognition






0	1	0	0	1	0	0	1	1	0	1	0	1	0	0	1	1	...	1
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	-----	---


Image recognition





0	1	0	0	1	0	0	1	1	0	1	0	1	0	0	1	1	...	1
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	-----	---




Convolutional neural networks





Convolutional neural networks




filter





Convolutional neural networks




filter





Convolutional neural networks




filter




Convolutional neural networks




filter




Convolutional neural networks




Convolutional neural networks




Convolutional neural networks




Convolutional neural networks




Convolutional neural networks






Convolutional neural networks




Convolutional neural networks




Convolutional neural networks




Convolutional neural networks





Convolutional networks




Convolutional neural networks





Convolutional neural networks



Author guessing

Text

- O Romeo, Romeo, wherefore art thou Romeo? Deny thy father and refuse thy name; Or if thou wilt not, be but sworn my love.
- It is better to be feared, than to be loved, if you cannot have both.
- When I was young I thought that money was the most important thing in life; now that I am old I know that it is.

Author

- Shakespeare
- Machiavelli
- Oscar Wilde

Why not...

- Standard neural networks? Text has variable size and with very long texts, we would need very complex neural networks.
- Convolutional neural networks? A filter may miss important information! See Oscar Wilde's quote.

Recurrent neural networks

Review	Positive review?
"Nice film"	1
"OK film"	1
"Bad movie"	0
"Terrible!"	0

Bag-of-words vectorization


Review	Positive review?
"Nice film"	1
"OK film"	1
"Bad movie"	0
"Terrible!"	0

bad	film	movie	nice	ok	terrible	Positive?
0	1	0	1	0	0	1
0	1	0	0	1	0	1
1	0	1	0	0	0	0
0	0	0	0	0	1	0

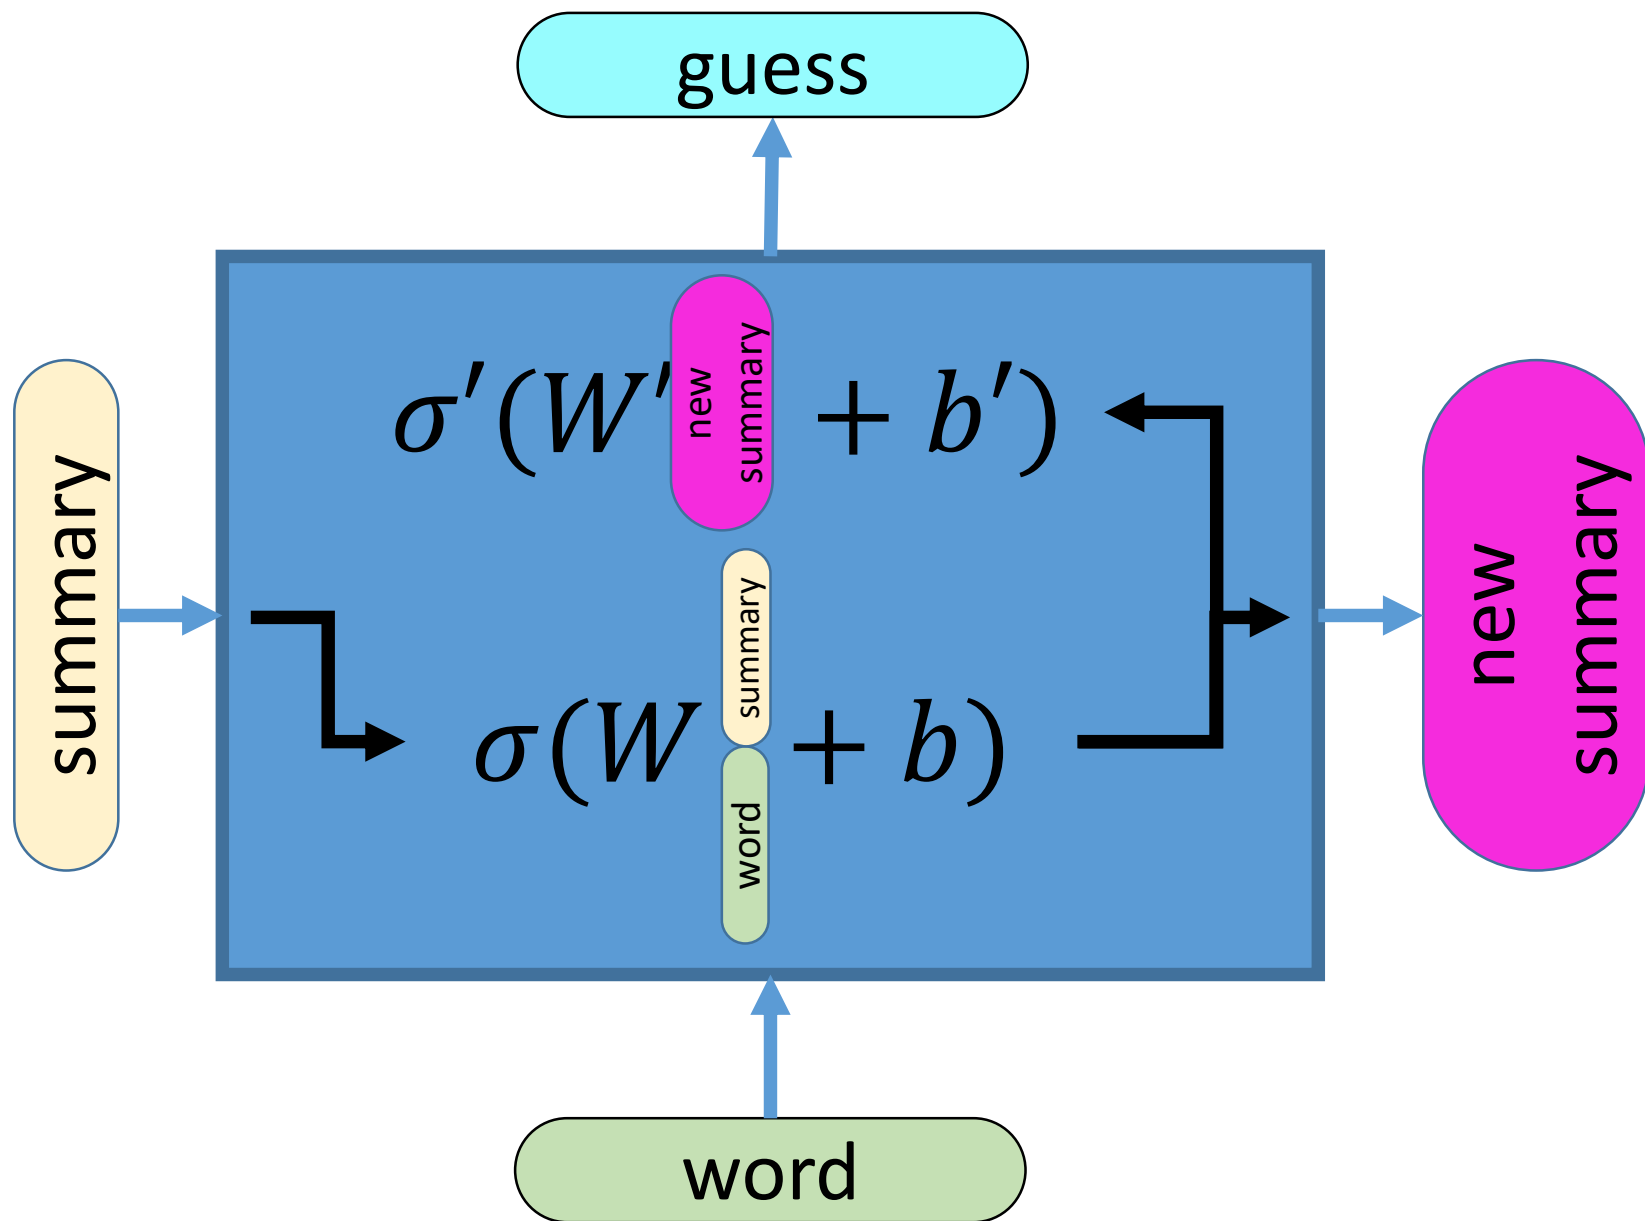
Analogously, words are vectors

	bad	film	movie	nice	ok	terrible
bad	1	0	0	0	0	0
film	0	1	0	0	0	0
movie	0	0	1	0	0	0
nice	0	0	0	1	0	0


Recurrent neural networks



* Recall that words can be represented as vectors



Recurrent neural networks



* Recall that words can be represented as vectors

Author guessing

Text

- O Romeo, Romeo, wherefore art thou Romeo? Deny thy father and refuse thy name; Or if thou wilt not, be but sworn my love.
- Mr. and Mrs. Dursley, of number four Privet Drive, were proud to say that they were perfectly normal, thank you very much.
- It is better to be feared, than to be loved, if you cannot have both.
- When I was young I thought that money was the most important thing in life; now that I am old I know that it is.

Author

- Shakespeare
- J.K. Rowling (Harry Potter)
- Machiavelli
- Oscar Wilde

What we learned

- Many machine learning models cannot make “insights” on data.
- What are neural networks?
- What are convolutional neural networks?
- What are recurrent neural networks?