

Machine Learning and Artificial Intelligence: Part 2


Neural Networks and Deep Learning

George Bezerra
Director of Data Science, TripAdvisor

SFI Complex Systems Summer School 2018

Logistic regression revisited


- Graphical representation:




$$f(\mathbf{x}, \mathbf{w}) = \frac{1}{1 + e^{-\left(\sum_{j=0}^M w_j x_j\right)}}$$

The LR decision boundary is **linear**

Linear separation




Nonlinear separation




Multi-layer perceptron (MLP)

- Stacking logistic regressions together




The MLP decision boundary is nonlinear

Linear separation



Nonlinear separation




Gradient Descent

Loss function:

$$\mathcal{L}(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_i(\mathbf{w})$$

Iterative minimization:

$$\mathbf{w}(t+1) = \mathbf{w}(t) - \alpha \nabla \mathcal{L}(\mathbf{w}(t))$$




Gradient Descent

Loss function:

$$\mathcal{L}(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_i(\mathbf{w})$$

Iterative minimization:

$$\mathbf{w}(t+1) = \mathbf{w}(t) - \alpha \nabla \mathcal{L}(\mathbf{w}(t))$$




Gradient Descent

Loss function:

$$\mathcal{L}(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_i(\mathbf{w})$$

Iterative minimization:

$$\mathbf{w}(t+1) = \mathbf{w}(t) - \alpha \nabla \mathcal{L}(\mathbf{w}(t))$$




Gradient Descent

Loss function:

$$\mathcal{L}(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_i(\mathbf{w})$$

Iterative minimization:

$$\mathbf{w}(t+1) = \mathbf{w}(t) - \alpha \nabla \mathcal{L}(\mathbf{w}(t))$$




Gradient Descent

Loss function:

$$\mathcal{L}(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_i(\mathbf{w})$$

Iterative minimization:

$$\mathbf{w}(t+1) = \mathbf{w}(t) - \alpha \nabla \mathcal{L}(\mathbf{w}(t))$$



Gradient Descent

Loss function:


$$\mathcal{L}(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_i(\mathbf{w})$$

Iterative minimization:

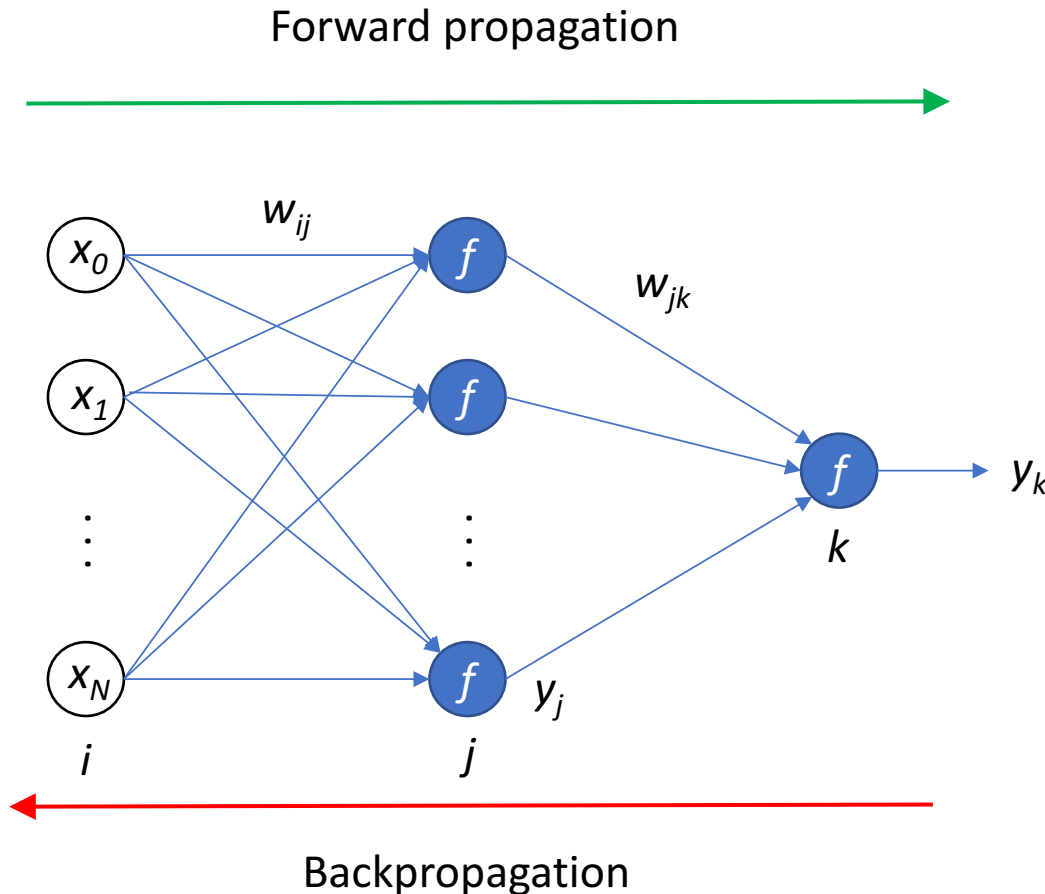
$$\mathbf{w}(t+1) = \mathbf{w}(t) - \alpha \nabla \mathcal{L}(\mathbf{w}(t))$$



Learning rate



Backpropagation algorithm




Chain rule:


$$\frac{\partial \mathcal{L}}{\partial w_{jk}} = \frac{\partial \mathcal{L}}{\partial y_k} \cdot \frac{\partial y_k}{\partial w_{jk}}$$

$$\frac{\partial \mathcal{L}}{\partial w_{ij}} = \frac{\partial \mathcal{L}}{\partial y_k} \cdot \frac{\partial y_k}{\partial y_j} \cdot \frac{\partial y_j}{\partial w_{ij}}$$




Deep learning



Deep learning



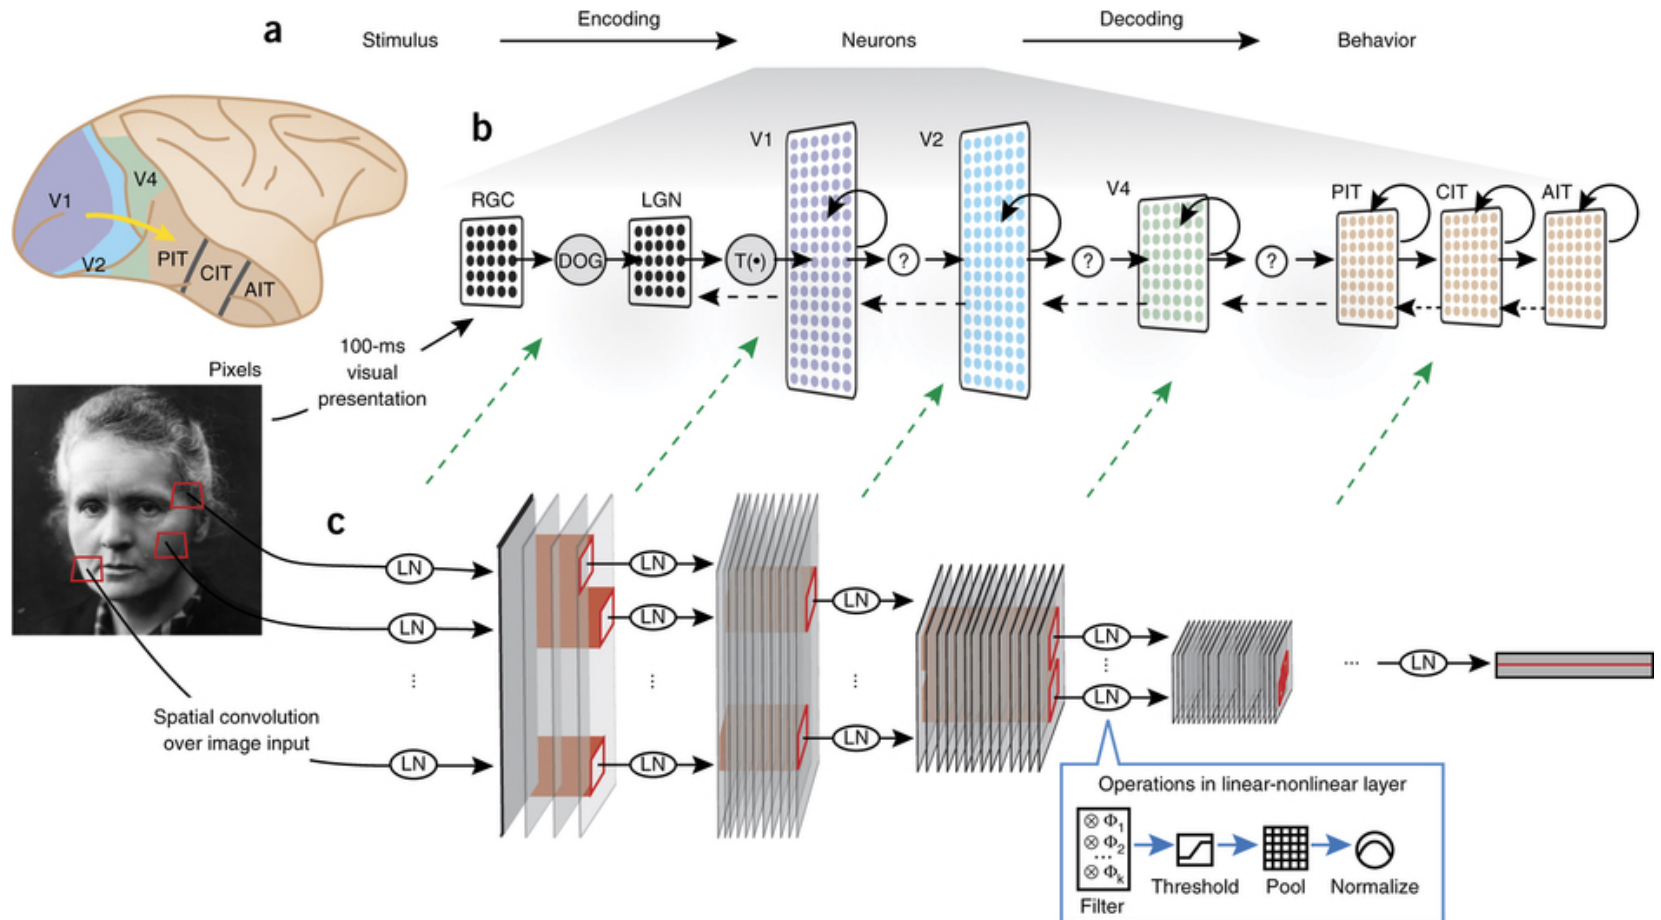
Deep learning



Why deep learning now?

- Massive datasets
- More compute power (Faster CPUs + GPUs)
- Better understanding of regularization techniques
- Open source deep learning frameworks: PyTorch, Keras, Tensorflow, Theano, MXNet, Caffe, etc.

Convolutional Neural Networks



Imagenet Challenge

- 1M images
- 1000 categories



Geological formation, formation
(geology) the geological features of the earth

1808
pictures

86.24%
Popularity
Percentile

Wordnet
IDs

Numbers in brackets: (the number of synsets in the subtree).

ImageNet 2011 Fall Release (32326)
- plant, flora, plant life (4486)
- geological formation, formation (1:
- aquifer (0)
- beach (1)
- cave (3)
- cliff, drop, drop-off (2)
- delta (0)
- diapir (0)
- folium (0)
- foreshore (0)
- ice mass (10)
- lakefront (0)
- massif (0)
- monocline (0)
- mouth (0)
- natural depression, depression (0)
- natural elevation, elevation (41)
- oceanfront (0)
- range, mountain range, range of
- relict (0)
- ridge, ridgeline (2)
- ridge (0)
- shore (7)
- slope, incline, side (17)
- spring, fountain, outflow, outpo
- talus, scree (0)
- vein, mineral vein (1)
- volcanic crater, crater (2)
- wall (0)

Treemap Visualization


Images of the Synset

Downloads

ImageNet 2011 Fall Release Geological formation, formation



ImageNet Large Scale Visual Recognition Challenge results



Convolution

Image (binary)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

1	0	1
0	1	0
1	0	1

Neuron
(convolutional filter)

1 _{x1}	1 _{x0}	1 _{x1}	0	0
0 _{x0}	1 _{x1}	1 _{x0}	1	0
0 _{x1}	0 _{x0}	1 _{x1}	1	1
0	0	1	1	0
0	1	1	0	0

Image


4		

Convolved
Feature

Image source:

http://deeplearning.stanford.edu/wiki/index.php/Feature_extraction_using_convolution


AlexNet



Recurrent Neural Networks


Feed-forward network

$$\mathbf{y} = f(W\mathbf{x})$$



Recurrent network

$$\mathbf{y}_t = f(W\mathbf{x}_t + U\mathbf{y}_{t-1})$$



Long-short term memory (LSTM)


$$f_t = \sigma_g(W_f x_t + U_f h_{t-1} + b_f)$$

$$i_t = \sigma_g(W_i x_t + U_i h_{t-1} + b_i)$$

$$o_t = \sigma_g(W_o x_t + U_o h_{t-1} + b_o)$$

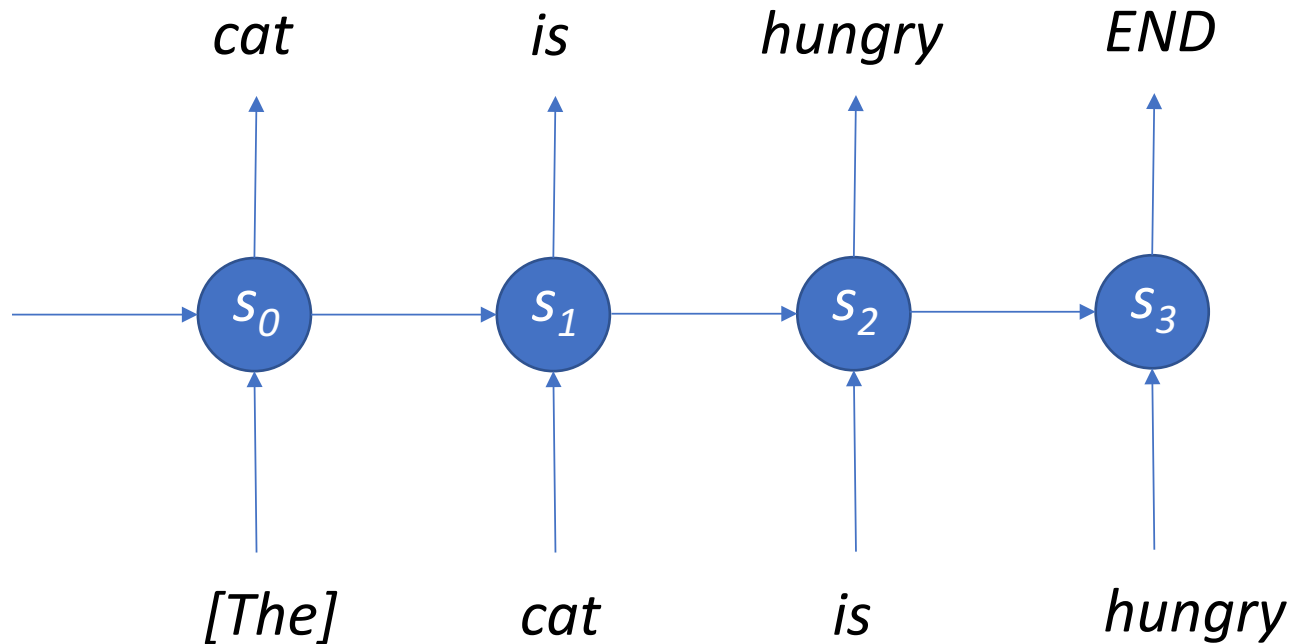
$$c_t = f_t \circ c_{t-1} + i_t \circ \sigma_c(W_c x_t + U_c h_{t-1} + b_c)$$

$$h_t = o_t \circ \sigma_h(c_t)$$



Text generation with LSTMs

- Language model: $P(w_1, \dots, w_m) = \prod_{i=1}^m P(w_i | w_1, \dots, w_{i-1})$



Text generation examples

Shakespeare:

[KING LEAR:] O, if you were a feeble sight, the
courtesy of your law, Your sight and several breath,
will wear the gods With his heads, and my hands are
wonder'd at the deeds, So drop upon your lordship's
head, and your opinion Shall be against your honour.

(Source: <http://karpathy.github.io/2015/05/21/rnn-effectiveness/>)

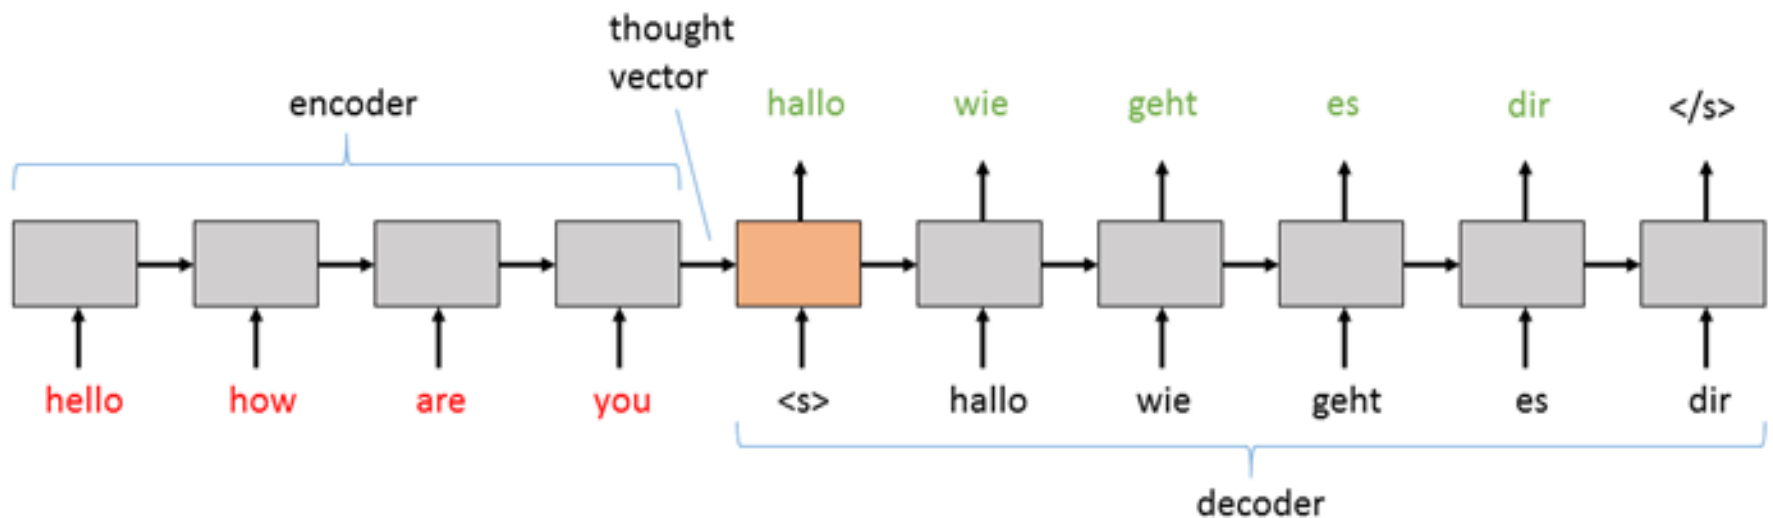
Trump tweet bot:

[I, Donald J. Trump, president of the United]States of
The Apprentice and the world in the world that we
have to be a great trade deal and the people of the
World Trade Center in the world that we have to be a
great person to be a great problem.

(Source: <https://towardsdatascience.com/yet-another-text-generation-project-5cfb59b26255>)

Neural Machine Translation


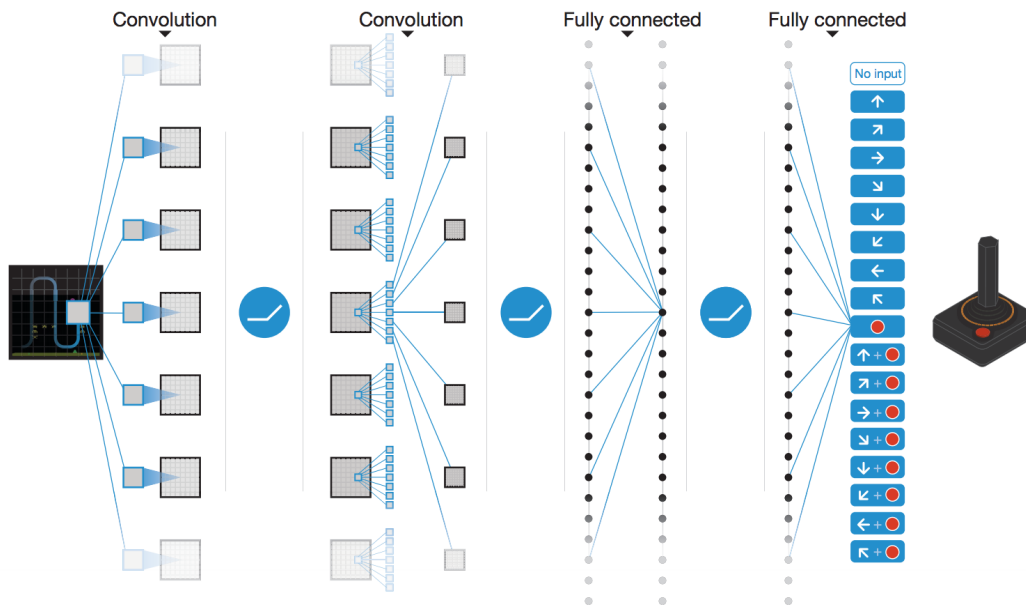
- Sequence-to-sequence models



Deep reinforcement learning



- Deep Q-learning (Google Deepmind)

Network structure



Atari game

Killer robots!



Boston Dynamics



The real danger of “AI”

- Algorithmic bias:

“when a computer system behaves in ways that reflects the implicit values of humans involved in that data collection, selection, or use.” – Wikipedia

- Increased danger with widespread automation



Short-term home rentals business

Customer



Hi! Is your house available for those dates?

Great! Can I bring my dog?



Home owner

Yes, it is!



Of course! Pets are welcome.



Fraud detection with deep learning

Customer: “Hi, I’d like to stay at you place, it’s perfect!
Do you accept pets?”

Owner 1: “We’d love to have you stay with us. Yes, we
do accept pets.”



Fraud detection with deep learning

Customer: “Hi, I’d like to stay at you place, it’s perfect!
Do you accept pets?”

Owner 1: “We’d love to have you stay with us. Yes, we
do accept pets.”



Owner 2: “Please contact me at badowner@gmail.com”



Fraud detection with deep learning

Customer: “Hi, I’d like to stay at you place, it’s perfect!
Do you accept pets?”

Owner 1: “We’d love to have you stay with us. Yes, we
do accept pets.”



Owner 2: “Please contact me at badowner@gmail.com”



Owner 3: “Find me at *****, I can make you a better
deal.”



More examples

- “Search for ‘Paradise house’ on google and you’ll find our website.”
- “Call me at five five five two seven six three.”
- “Send a message to janelle at hotmail dot com”
- “please c-o-n-t-a-c-t me on mmartin-a-t-g-m-a-i-l”

Naïve approach: Regular expressions

- Regular expressions for phone numbers, websites, email, keywords, etc.
 - *Websites:*
`(?:((?:https?|ftp)://)|(?:www|m|ftp)\.)(?:[a-z0-9-]+\.)+[a-z](?:[0-9]+)?(?:/(|\?).*)?`
 - *Emails:* `(?:[a-z0-9._%+-]| +dot +)+(?: *@ *| +at +)(?:[a-z0-9.-]| +dot +)(?:\.| +dot +)[a-z]{2,6}`
- Over 2,000 regexes
- Manual process
- Generates many false positives

Using a model: bag-of-words + regex

- Logistic regression classifier using bag of words
- Uses existing regular expressions as features (1 if there is a match, 0 otherwise)
- Reduces false positives by 2-4x!
- However...
 - Regular expressions are easy to fool
 - Cannot learn new patterns on its own (requires new regexes)
 - BOW misses character-level features
 - Does not capture context

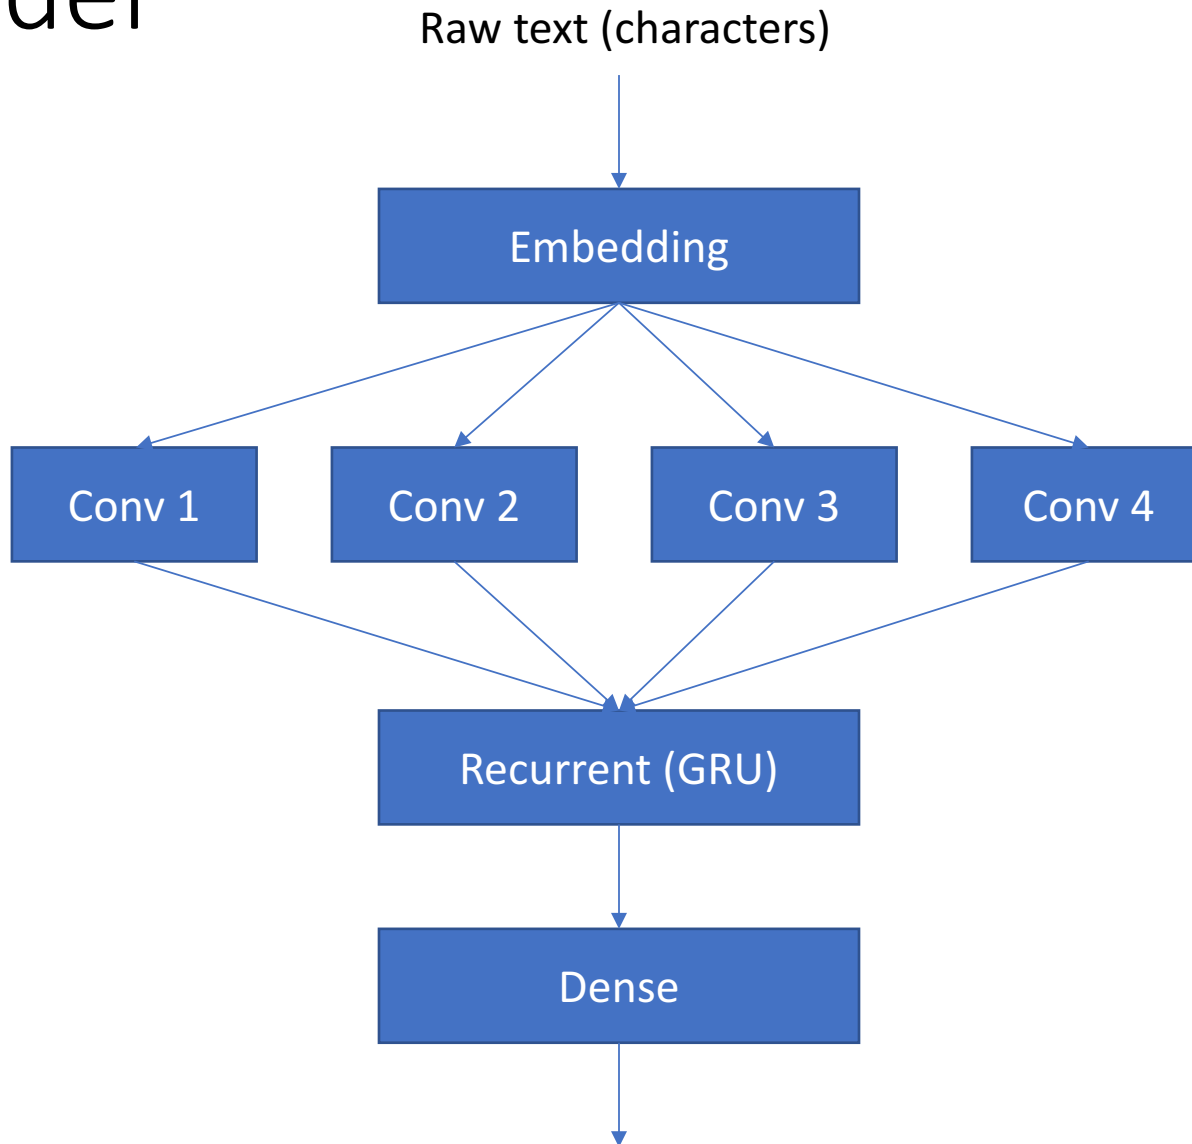
Examples where traditional approaches fail

- Please call me on *407*410*9301
- u can 91 reach me 49 on this 27 number anytime 86 you want 54. I hope you understand the message
- I can confirm the availability and the price. I kindly invite you to proceed with your booking. Please do not hesitate to contact me if you need further information. Thank you. (*False positive*)
- Please note the owner address contact him directly: ernesto@gardener.c o m
- you may r. ea ch us at नौ तीन नौ शून्य आठ सात आठ एक शून्य दो

Deep learning solution

- Learns the features (at the character level)
- Interprets contextual meaning (through recurrent layers)
- Captures high-level relationships and patterns
- Goals:
 - Eliminate the need for regular expressions
 - Improve accuracy
 - Catch more circumvention

Model



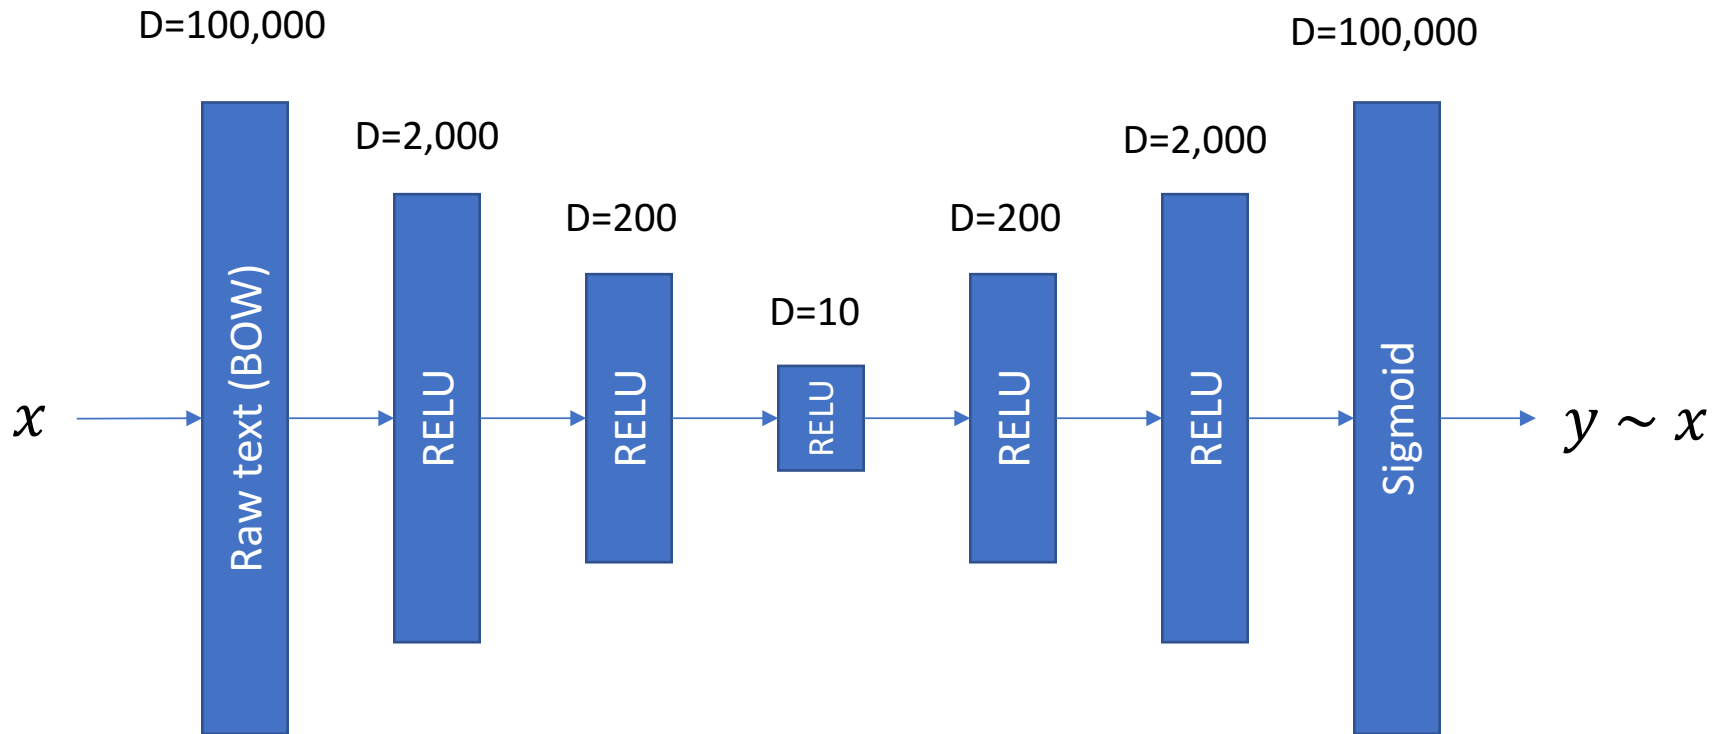
Results

- Please call me on *407*412*9391
 - Proba: 0.998
- u can 90 reach me 49 on this 17 number anytime 86 you want 54. I hope you understand the message
 - Proba: 0.710
- I can confirm the availability and the price. I kindly invite you to proceed with your booking. Please do not hesitate to contact me if you need further information. Thank you. (False positive)
 - Proba: 0.016
- Please note the owner address contact him directly: alvarez@gardener.com
 - Proba: 0.97
- you may reach us at नौ नौ तीन शून्य आठ सात एक आठ शून्य दो
 - Proba: 0.99

Recap (part 2)

- Neural networks are powerful non-linear methods for solving supervised problems.
- Deep learning takes neural networks to the next level: more layers, activation functions, architectures
- Convolutional neural networks changed the game on image classification tasks.
- Recurrent neural networks enable a wide variety of natural language processing tasks.
- Killer robots are not the real danger of AI; algorithmic bias is.

Deep auto-encoder




AI as a complex system

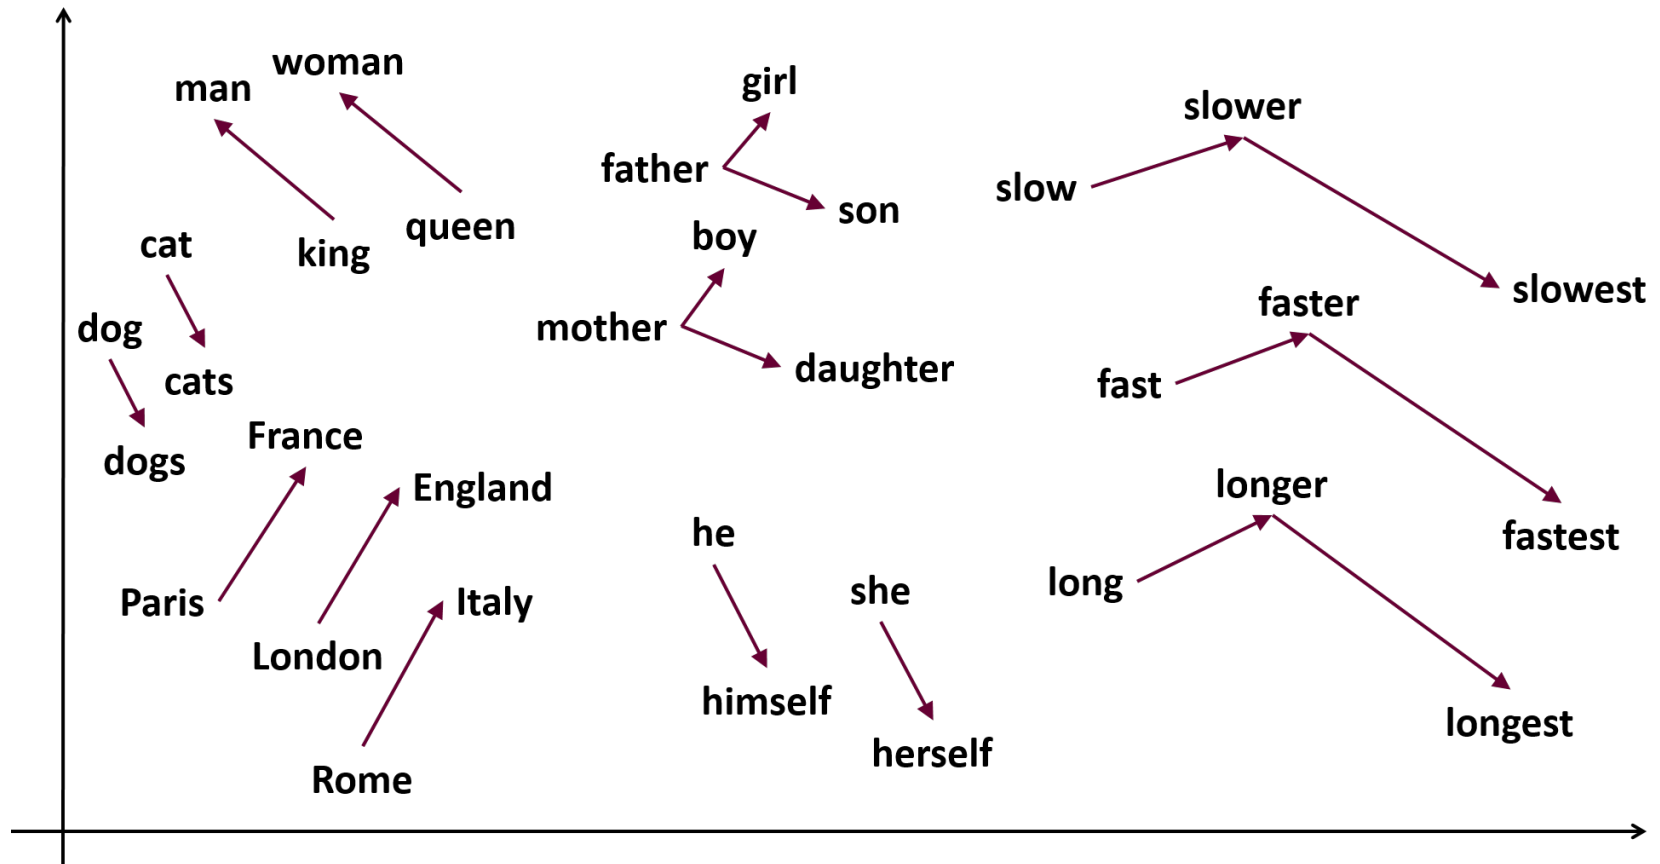
- Backpropagation
- Non-linear dynamics
- Fractal (hierarchical-modular) network structure
- Domain-specific
- Emergence

Word Embeddings

- Finding dense representations for sparse high-dimensional data

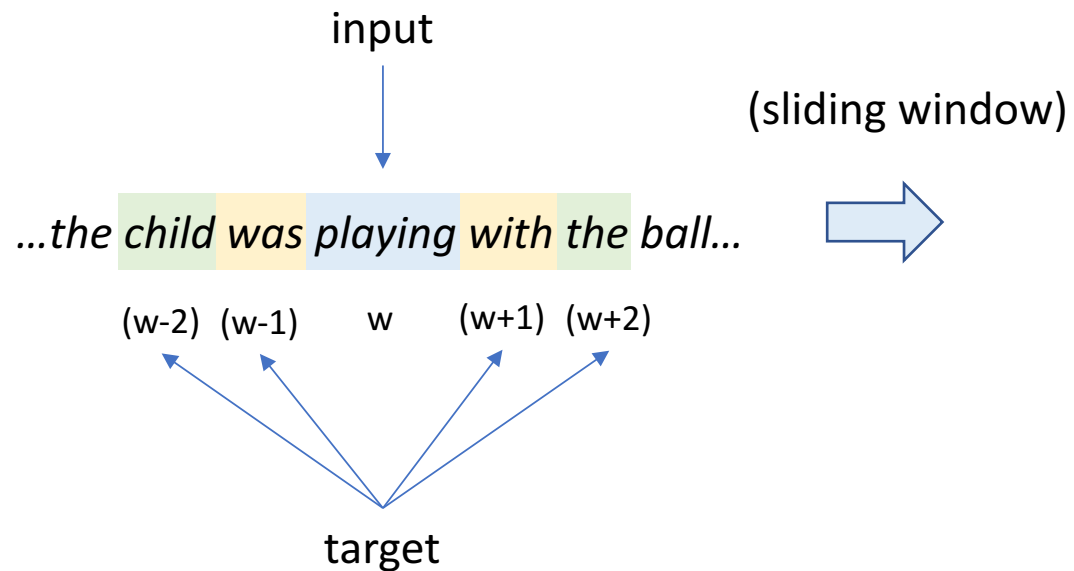


Word Embeddings in 2D

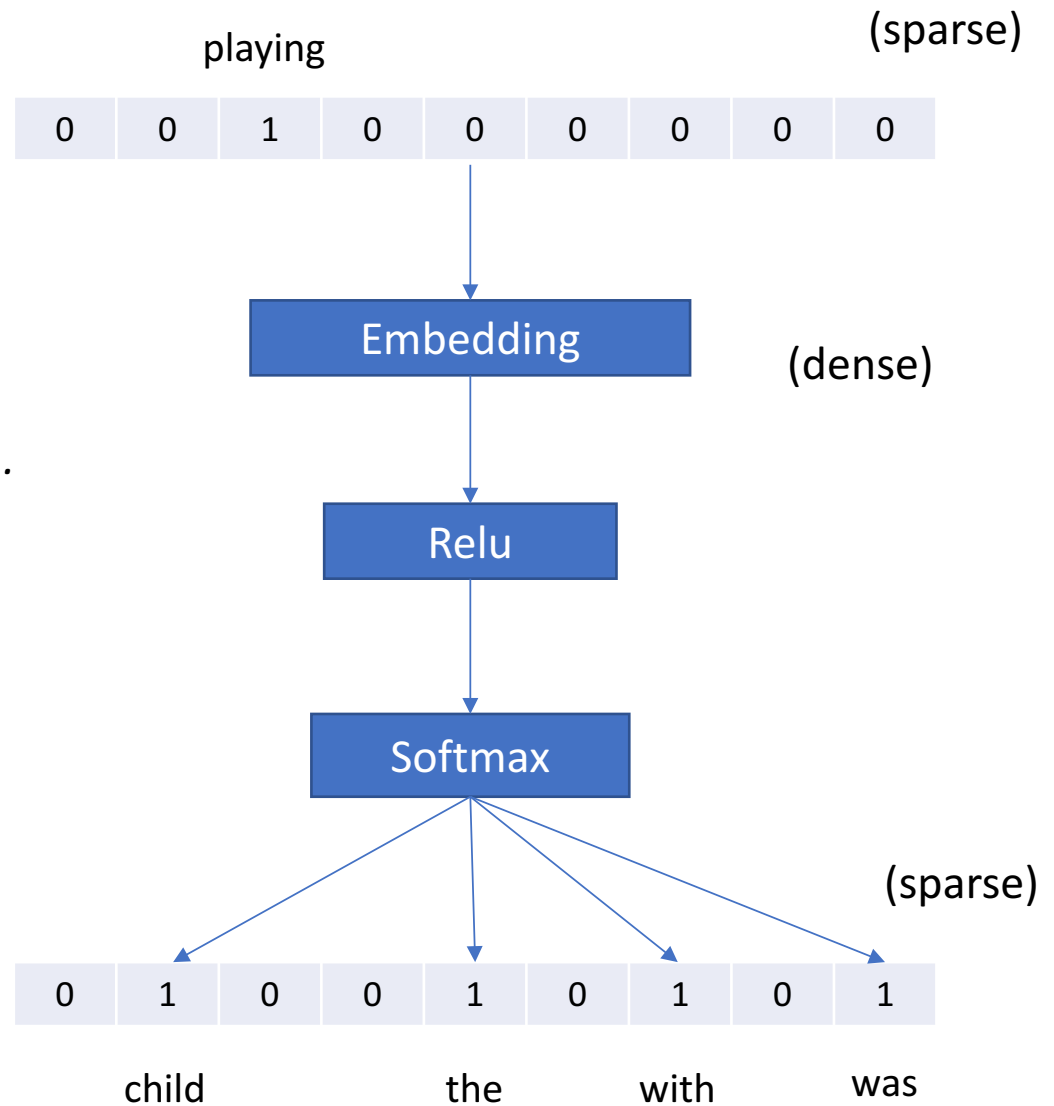
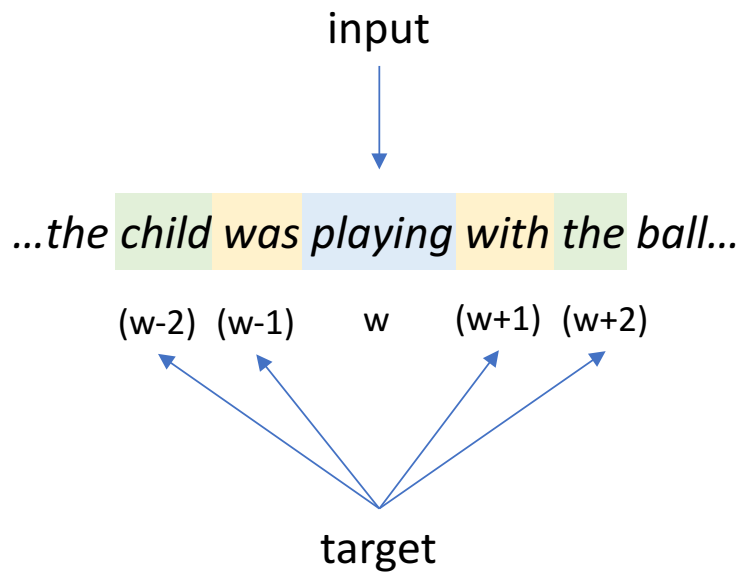


Word2vec

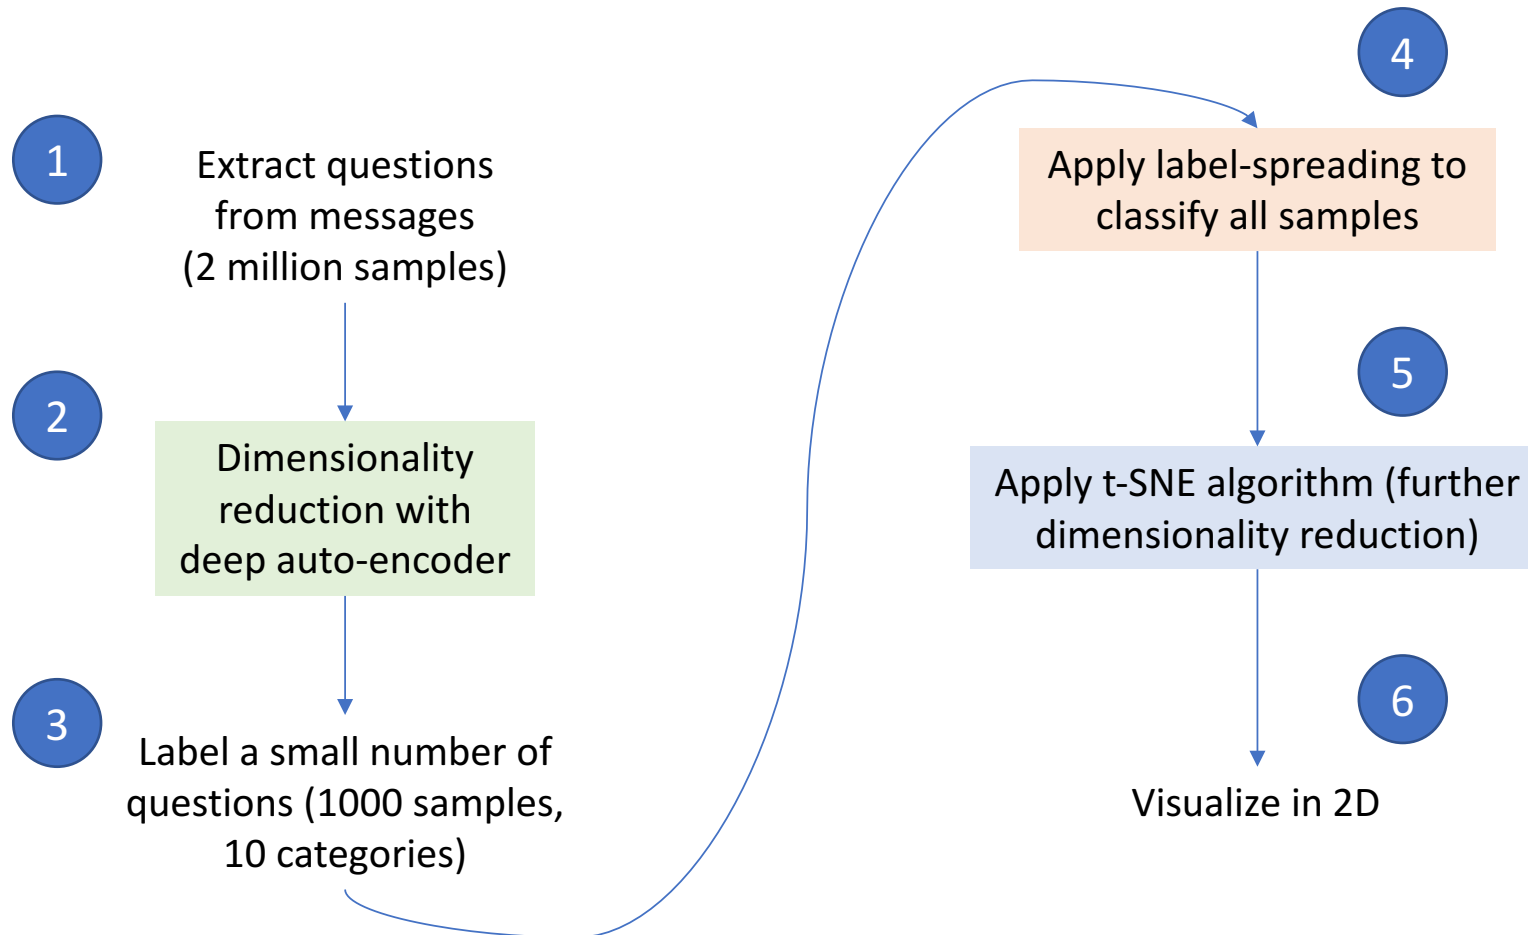
- Assumption: words that appear in the same context are similar to each other



Word2vec



Topic modeling with auto-encoders + semi-supervised learning



AlexNet

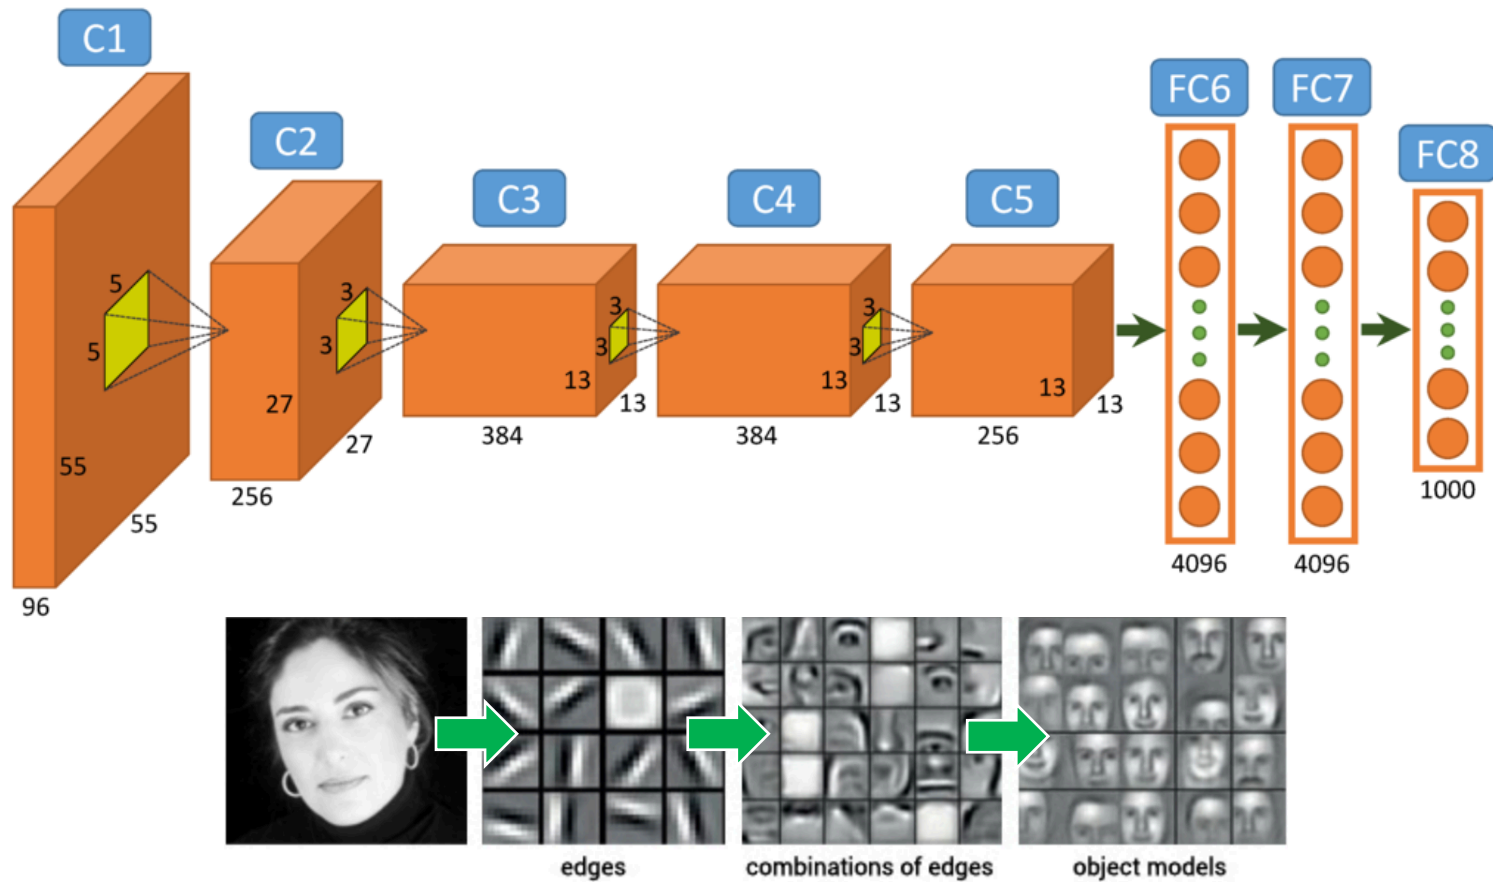




Image source: <https://www.saagie.com/blog/object-detection-part1>

Auto-encoder + 1000 labeled samples




Auto-coder + Label Spreading




Recurrent Neural Networks

“Unrolled” network




Backpropagation Through Time (BPTT)



Backpropagation Through Time (BPTT)

Vanishing gradient problem!



Structured vs Unstructured data

Structured


Results Messages						
	Id	Address	Date	Id	Name	Date
1	1	Bangalore	2011-06-09 15:23:37.873	1	shirsendu	2011-06-09 15:17:58.130
2	1	Bangalore	2011-06-09 15:23:37.873	1	shirsendu	2011-06-09 15:19:30.087
3	1	Bangalore	2011-06-09 15:23:37.873	2	Samali	2011-06-09 15:19:30.087
4	1	Bangalore	2011-06-09 15:23:37.873	3	Mrinal	2011-06-09 15:19:30.087
5	2	Bangkok	2011-06-09 15:23:37.890	1	shirsendu	2011-06-09 15:17:58.130
6	2	Bangkok	2011-06-09 15:23:37.890	1	shirsendu	2011-06-09 15:19:30.087
7	2	Bangkok	2011-06-09 15:23:37.890	2	Samali	2011-06-09 15:19:30.087
8	2	Bangkok	2011-06-09 15:23:37.890	3	Mrinal	2011-06-09 15:19:30.087
9	3	CAIcutta	2011-06-09 15:23:37.890	1	shirsendu	2011-06-09 15:17:58.130
10	3	CAIcutta	2011-06-09 15:23:37.890	1	shirsendu	2011-06-09 15:19:30.087
11	3	CAIcutta	2011-06-09 15:23:37.890	2	Samali	2011-06-09 15:19:30.087
12	3	CAIcutta	2011-06-09 15:23:37.890	3	Mrinal	2011-06-09 15:19:30.087

PROCLIB.PAYLIST Table					
IdNum	Gender	Jobcode	Salary	Birth	Hired
1639	F	TA1	42260	26JUN70	28JAN91
1065	M	ME3	38090	26JAN54	07JAN92
1400	M	ME1	29769	05NOV67	16OCT90
1561	M		36514	30NOV63	07OCT87
1221	F	FA3	.	22SEP63	04OCT94

Unstructured



Dies ist ein Blindtext. An ihm lässt sich vieles über die Schrift ablesen, in der er gesetzt ist. Auf den ersten Blick wird der Grauwert der Schriftfläche sichtbar. Dann kann man prüfen, wie gut die Schrift zu lesen ist und wie sie auf den Leser wirkt. Dies ist ein Blindtext. An ihm lässt sich vieles über die Schrift ablesen, in der er gesetzt ist. Auf den ersten Blick wird der Grauwert der Schriftfläche sichtbar. Dann kann man prüfen, wie gut die Schrift zu lesen ist und wie sie auf den Leser wirkt.



An electroencephalogram (EEG)



Structured vs Unstructured data


Structured

Results Messages						
	Id	Address	Date	Id	Name	Date
1	1	Bangalore	2011-06-09 15:23:37.873	1	shirsendu	2011-06-09 15:17:58.130
2	1	Bangalore	2011-06-09 15:23:37.873	1	shirsendu	2011-06-09 15:19:30.087
3	1	Bangalore	2011-06-09 15:23:37.873	2	Samali	2011-06-09 15:19:30.087
4	1	Bangalore	2011-06-09 15:23:37.873	3	Mrinal	2011-06-09 15:19:30.087
5	2	Bangkok	2011-06-09 15:23:37.890	1	shirsendu	2011-06-09 15:17:58.130
6	2	Bangkok	2011-06-09 15:23:37.890	1	shirsendu	2011-06-09 15:19:30.087
7	2	Bangkok	2011-06-09 15:23:37.890	2	Samali	2011-06-09 15:19:30.087
8	2	Bangkok	2011-06-09 15:23:37.890	3	Mrinal	2011-06-09 15:19:30.087
9	3	CAicutta	2011-06-09 15:23:37.890	1	shirsendu	2011-06-09 15:17:58.130
10	3	CAicutta	2011-06-09 15:23:37.890	1	shirsendu	2011-06-09 15:19:30.087
11	3	CAicutta	2011-06-09 15:23:37.890	2	Samali	2011-06-09 15:19:30.087
12	3	CAicutta	2011-06-09 15:23:37.890	3	Mrinal	2011-06-09 15:19:30.087

Unstructured



Dies ist ein Blindtext. An ihm lässt sich vieles über die Schrift ablesen, in der er gesetzt ist. Auf den ersten Blick wird der Grauwert der Schriftfläche sichtbar. Dann kann man prüfen, wie gut die Schrift zu lesen ist und wie sie auf den Leser wirkt. Dies ist ein Blindtext. An ihm lässt sich vieles über die Schrift ablesen, in der er gesetzt ist. Auf den ersten Blick wird der Grauwert der Schriftfläche sichtbar. Dann kann man prüfen, wie gut die Schrift zu lesen ist und wie sie auf den Leser wirkt.



Practitioner's tip: Use tree-based methods (e.g., XGBoost) for structured data and neural networks for unstructured data.

PROCLIB PAYLIST Table					
Id					
16					
10					
14					
15					
1221	F	FA3	.	22SEP63	04OCT94