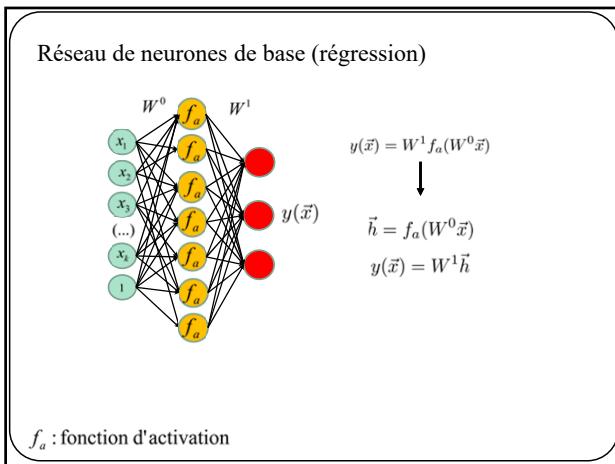
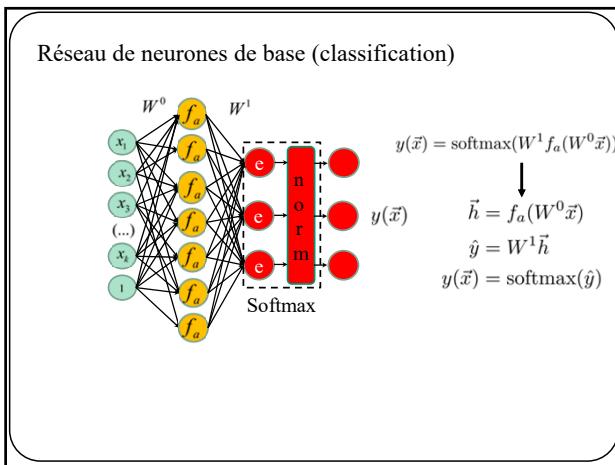


1

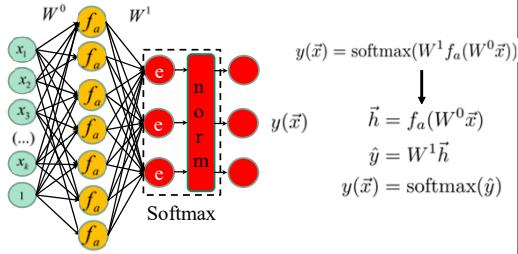


2



3

## Réseau de neurones de base (classification)

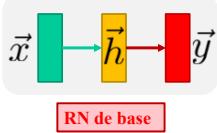


Ne permet que des tâches “1 pour 1”

- Classification (1 image = 1 étiquette)
  - Régression (1 donnée = 1 vecteur)
  - Localisation (1 boîte = 1 classification + 1 régression)

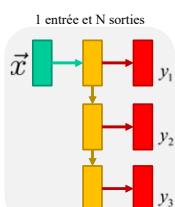
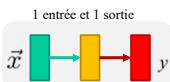
4

### Illustration simplifiée



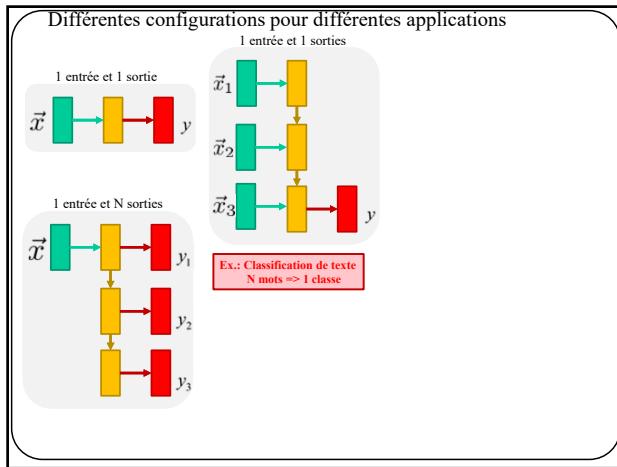
5

Différentes configurations pour différentes applications



**Ex.: description d'une image**  
1 image => N mots

6



7

---



---



---



---



---



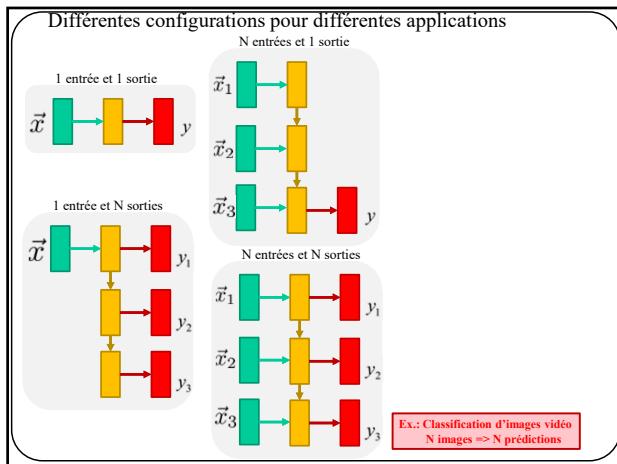
---



---



---



8

---



---



---



---



---



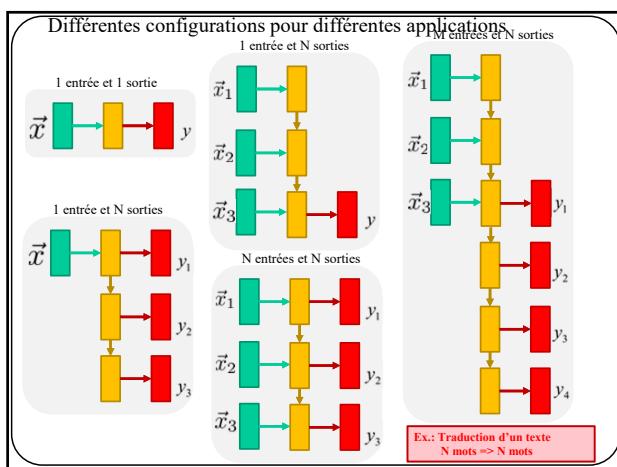
---



---



---



9

---



---



---



---



---



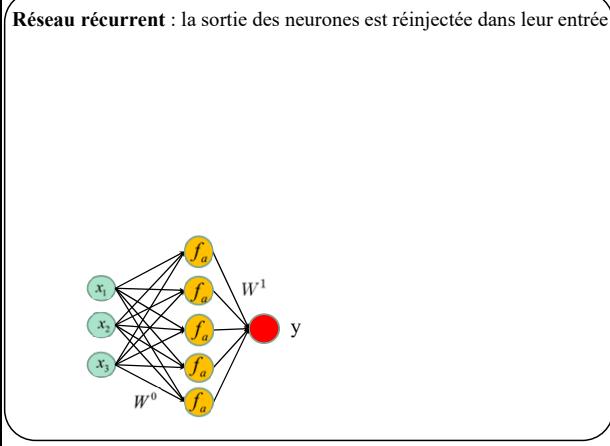
---



---



---



10

---

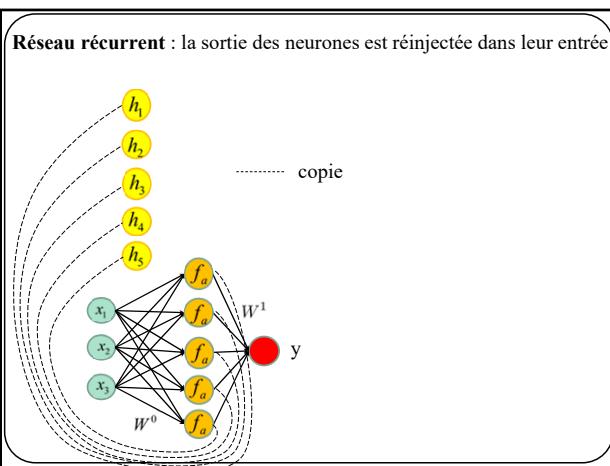
---

---

---

---

---



11

---

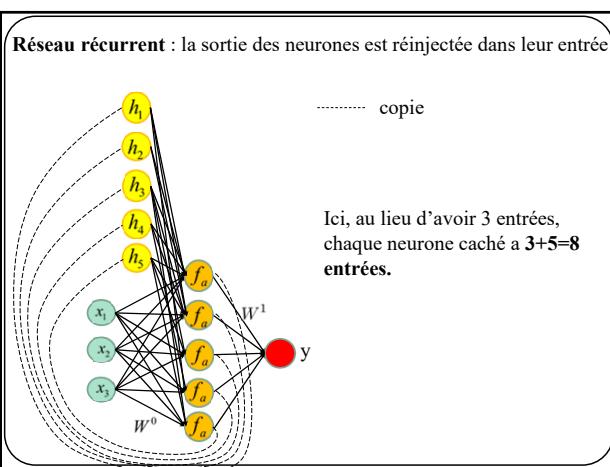
---

---

---

---

---



12

---

---

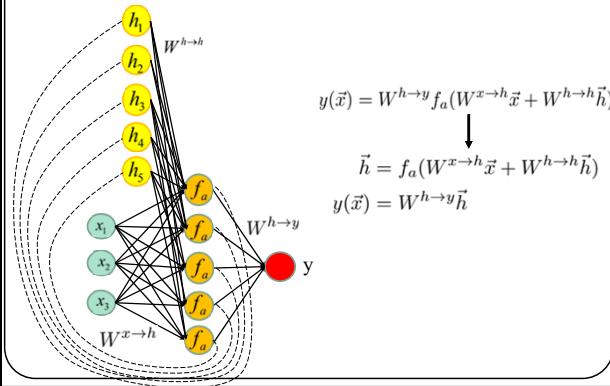
---

---

---

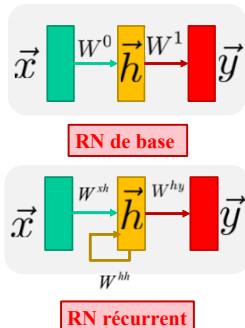
---

**Réseau récurrent** : la sortie des neurones est réinjectée dans leur entrée



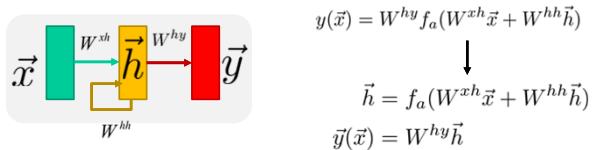
13

### Illustration simplifiée



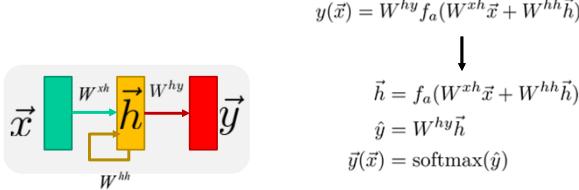
14

Dans le cas général avec K sorties (régression)



15

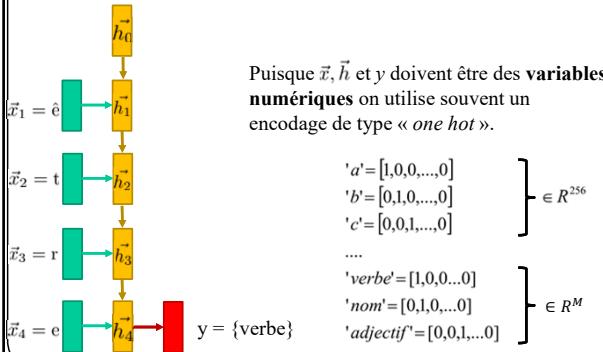
Dans le cas général avec K sorties (classification)



16

Exemple pour N entrées et 1 sortie:

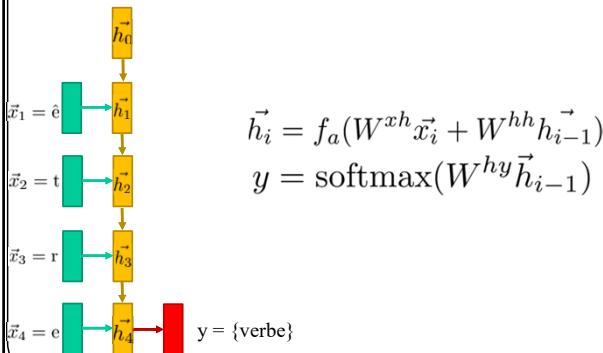
Analyse grammaticale (classification) : (é.t.r.e) => {verbe}



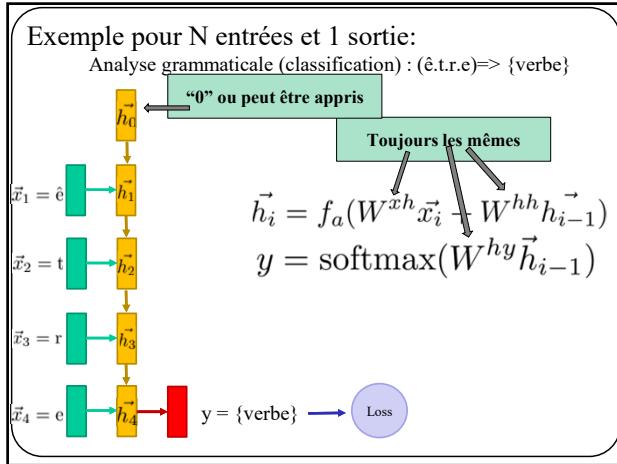
17

Exemple pour N entrées et 1 sortie:

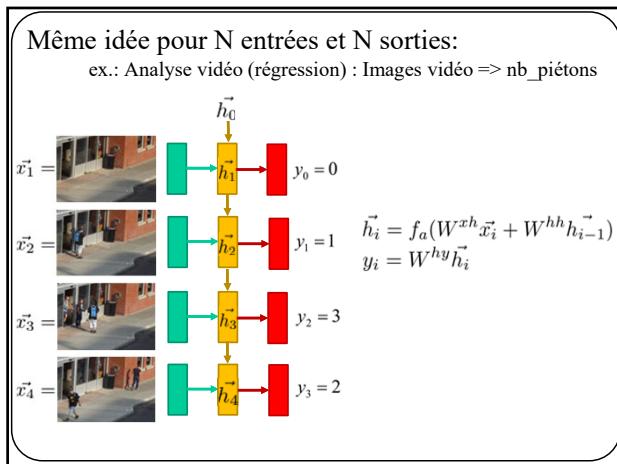
Analyse grammaticale (classification) : (é.t.r.e) => {verbe}



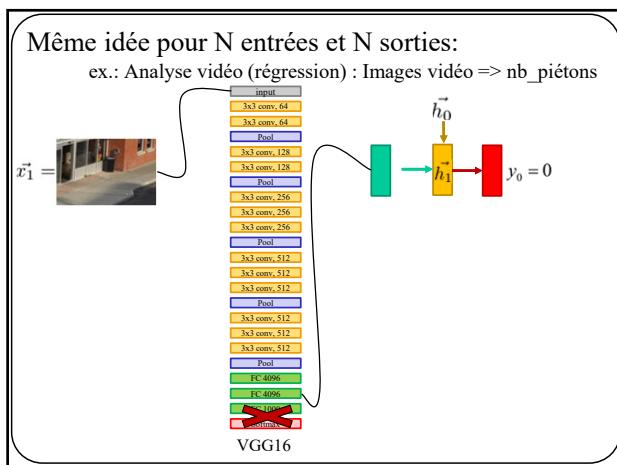
18



19



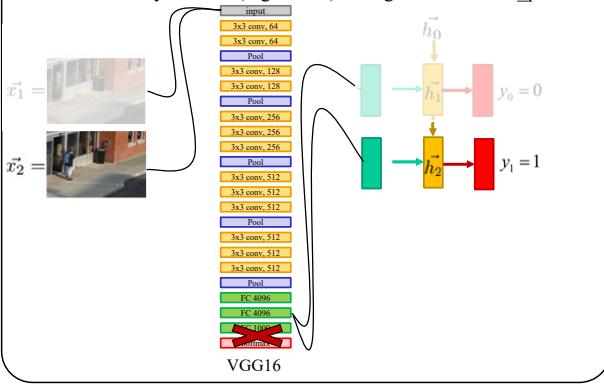
20



21

Même idée pour N entrées et N sorties:

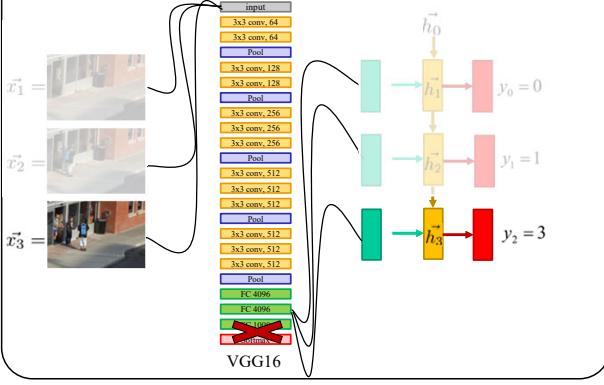
ex.: Analyse vidéo (régression) : Images vidéo => nb\_piétons



22

Même idée pour N entrées et N sorties:

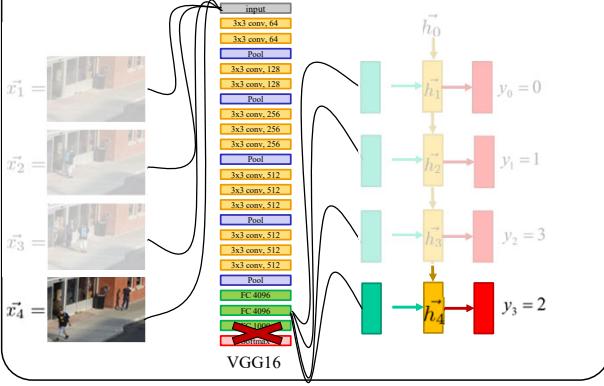
ex.: Analyse vidéo (régression) : Images vidéo => nb\_piétons



23

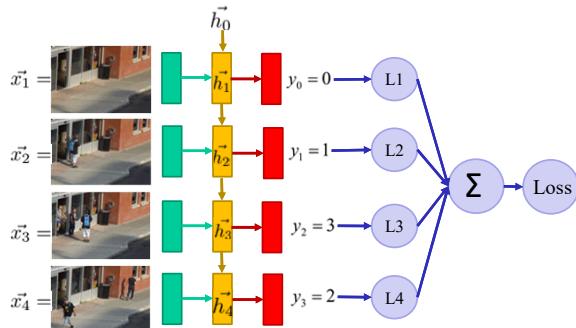
Même idée pour N entrées et N sorties:

ex.: Analyse vidéo (régression): Images vidéo => nb\_piétons



24

Même idée pour N entrées et N sorties:  
ex.: Analyse vidéo (régression) : Images vidéo => nb\_piétons



25

Autre exemple: **prédition de caractères** (modèle de langue)

## Alphabet jouet :[a,e,m,s]

### Représentation « one hot » jouet:

$$\begin{aligned} 'a' &= [1, 0, 0, 0] \\ 'e' &= [0, 1, 0, 0] \\ 'm' &= [0, 0, 1, 0] \\ 's' &= [0, 0, 0, 1] \end{aligned}$$

**But :** Entrainer un modèle à prédire les lettres du mot « **masse** ».

26

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]  
**Entraîner** un modèle

**QUESTION** What is the primary reason for the increase in mass media?

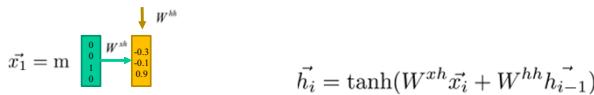
$$\vec{x_1} = \text{m}$$

27

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**Entrainer** un modèle à prédire les lettres du mot « **masse** ».

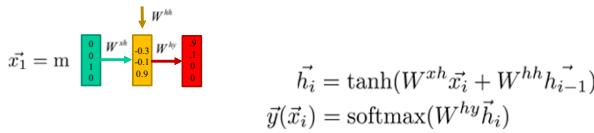


28

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[**a,e,m,s**]

**Entraîner** un modèle à prédire les lettres du mot « **masse** ».

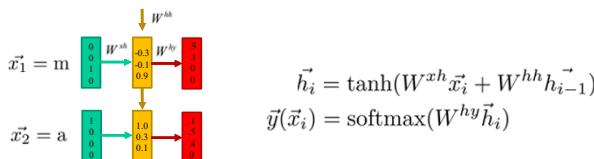


29

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**Entrainer** un modèle à prédire les lettres du mot « **masse** ».

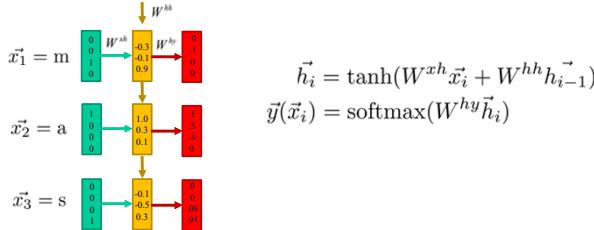


30

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**Entraîner** un modèle à prédire les lettres du mot « **masse** ».

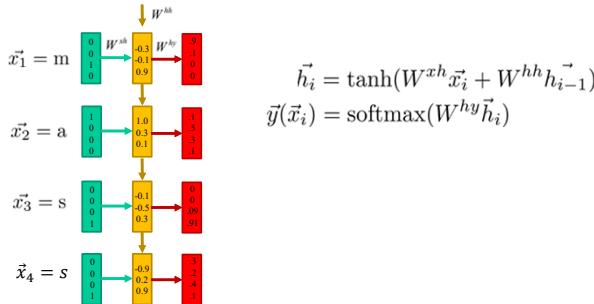


31

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**Entraîner** un modèle à prédire les lettres du mot « **masse** ».

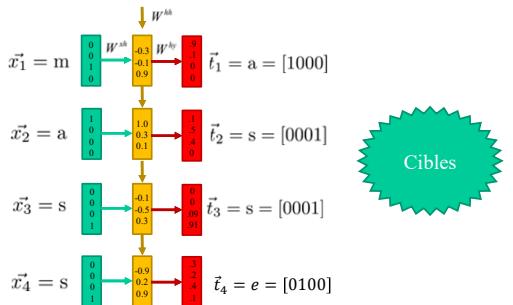


32

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**Entrainer** un modèle à prédire les lettres du mot « **masse** ».

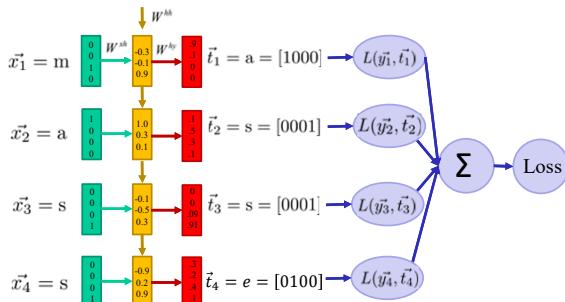


33

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[**a,e,m,s**]

**Entraîner** un modèle à prédire les lettres du mot « **masse** ».

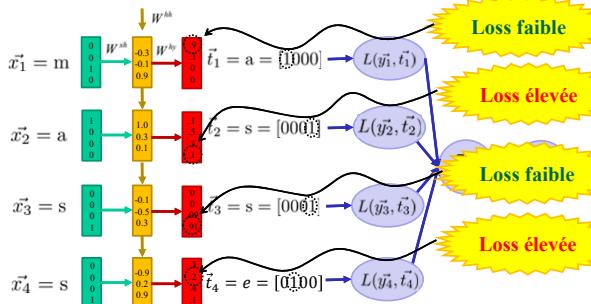


34

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**Entraîner** un modèle à prédire les lettres du mot « **masse** ».

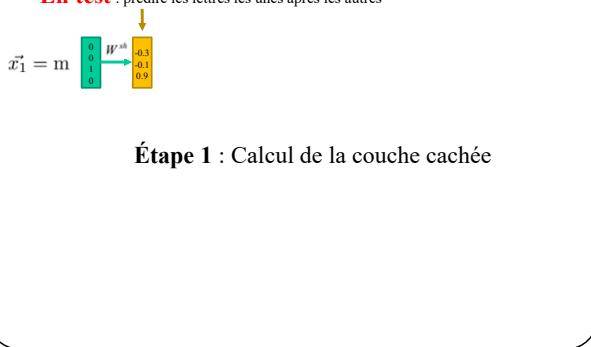


35

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[**a,e,m,s**]

**En test** : prédire les lettres les unes après les autres

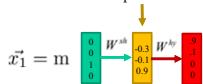


36

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**En test** : prédire les lettres les unes après les autres



**Étape 2** : Calcul de la sortie (softmax)

---



---



---



---



---



---



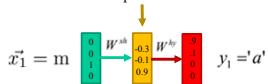
---

37

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**En test** : prédire les lettres les unes après les autres



**Étape 3** : Sélectionner le caractère le plus probable

---



---



---



---



---



---



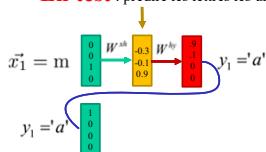
---

38

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**En test** : prédire les lettres les unes après les autres



**Étape 4** : Injecter le caractère prédit au début du réseau

---



---



---



---



---



---



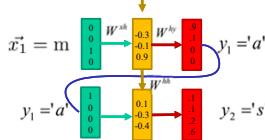
---

39

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**En test** : prédire les lettres les unes après les autres



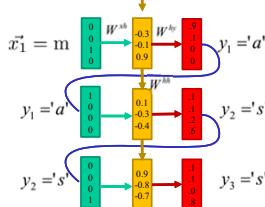
Et on recommence!

40

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**En test** : prédire les lettres les unes après les autres

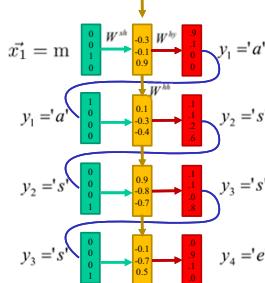


41

Autre exemple: **prédition de caractères** (modèle de langue)

Alphabet :[a,e,m,s]

**En test** : prédire les lettres les unes après les autres



42



Texte généré une fois le modèle entraîné

```
PANDARUS:
Alas, I think he shall be come apprehended and the day
When little strain would be attain'd into being never fed,
And who is but a chain and subjects of his death,
I should not sleep.

Second Senator:
They are away this miseries, produced upon my soul,
Breaking and strongly should be buried, when I perish
The earth and thoughts of many states.

DUKE VINCENTIO:
Well, your wit is in the care of side and that.

Second Lord:
They would be ruled after this chamber, and
my fair mess begun out of the fact, to be conveyed,
Whose noble souls I'll have the heart of the wars.

Clown:
Come, sir, I will make did behold your worship.

VIOLENT:
I'll drink it.
```

Crédit: A. Karpathy, CS231

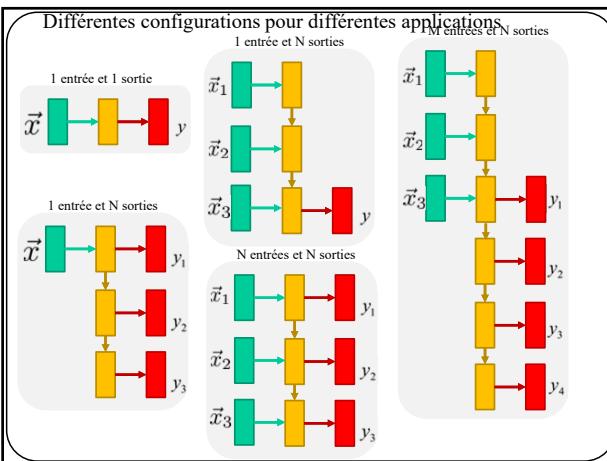
46

Entraînement sur le code source de Linux en C++

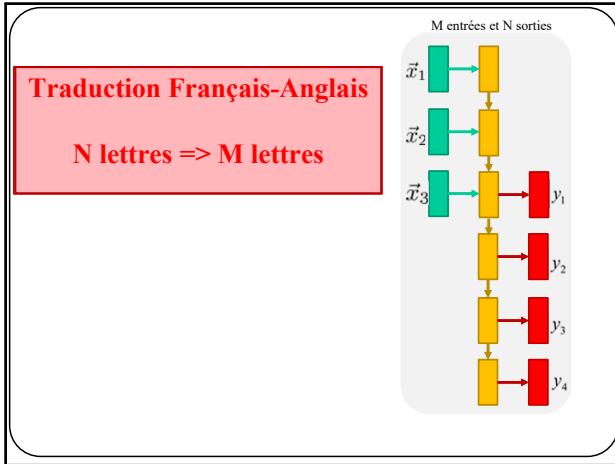
```
static void do_command(struct seq_file *m, void *v)
{
    int column = 32 << (cmd[2] & 0x8);
    if (column >= 32)
        cmd = (int)(int_state ^ (in_8(ech->oh_flags) & Cmd) ? 2 : 1);
    else
        seq = (char *)v;
    for (i = 0; i < 16 * 16*4;)
        if (i < 4 * 16)
            pipe = (in_use & UNXTHREAD_UNCC) +
                ((counts & 0xfffffffffffff8) & 0x0000000f) << 8;
        if (counts & 0x00000001)
            subpid, ppd_md.kexec.handle, 0x20000000;
        pipe_set_bit(pipe, 0);
    }
    /* From now on we can ignore pipeline to place column if all dash */
    subsystem_info |= Adf_Change.PageSize;
    rnk_controls(offset, idx, &offset);
    /* Now we want to deliberately put it to device */
    outstrncpy(&outstr, "policy%08x", val, 8);
    for (i = 0; i < COUNTER; i++)
        seq_puts(&seq, "policy ");
}
seq_puts(&seq, "\n");
```

Crédit: A. Karpathy, CS231

47



48



49

---

---

---

---

---

---

---

---

**Autre exemple: traduction**

Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.

Traduire 'assez' => 'enough'  
 Alphabet fr : [<BoS>, a, e, s, z, <EoS>]  
 Alphabet en: [<BoS>, e, g, h, n, o, u, <EoS>]

Pas le même nombre d'entrées que de sorties !  
 (BoS : Beginning of Sentence, EoS:End of Sentence).

---

---

---

---

---

---

---

---

50

**Autre exemple: traduction**

Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.

Traduire 'assez' => 'enough'  
 Alphabet fr : [<BoS>, a, e, s, z, <EoS>]  
 Alphabet en: [<BoS>, e, g, h, n, o, u, <EoS>]

$\vec{x}_1 = \text{<BoS>}$

---

---

---

---

---

---

---

---

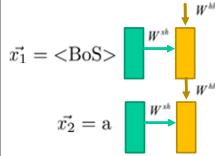
51

Autre exemple: traduction

Traduire ‘assez’ -> ‘enough’

Alphabet fr :[<BoS>,a,e,s,z,<EoS>]

Alphabet en: [<BoS>,e,g,h,n,o,u,<EndOfSequence>]



Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.

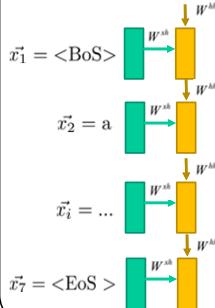
52

Autre exemple: traduction

Traduire ‘assez’ -> ‘enough’

Alphabet fr :[<BoS>,a,e,s,z,<EoS>]

Alphabet en: [<BoS>,e,g,h,n,o,u,<EoS>]



Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.

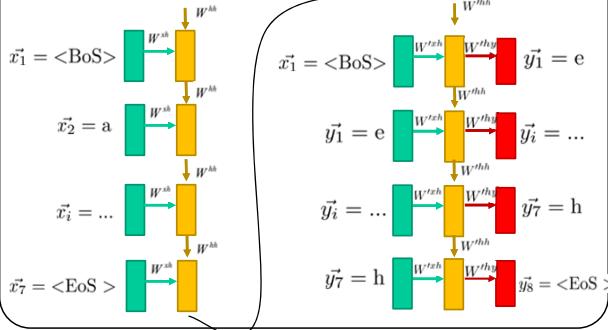
53

Autre exemple: traduction

Traduire ‘assez’ -> ‘enough’

Alphabet fr :[<BoS>,a,e,s,z,<EoS>]

Alphabet en: [<BoS>,e,g,h,n,o,u,<BoE>]

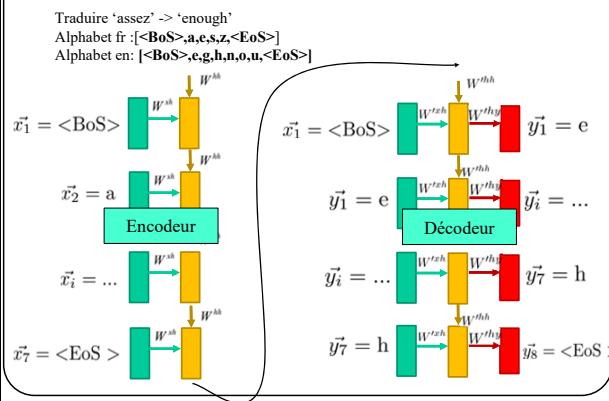


Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.

54

Autre exemple: traduction

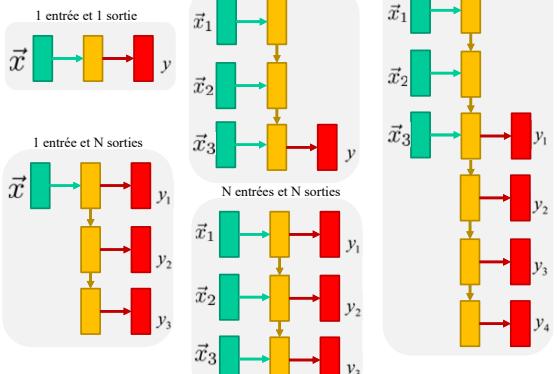
Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.



55

Différentes configurations pour différentes applications

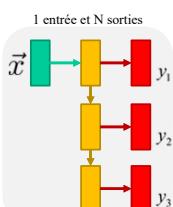
1 entrée et N sorties



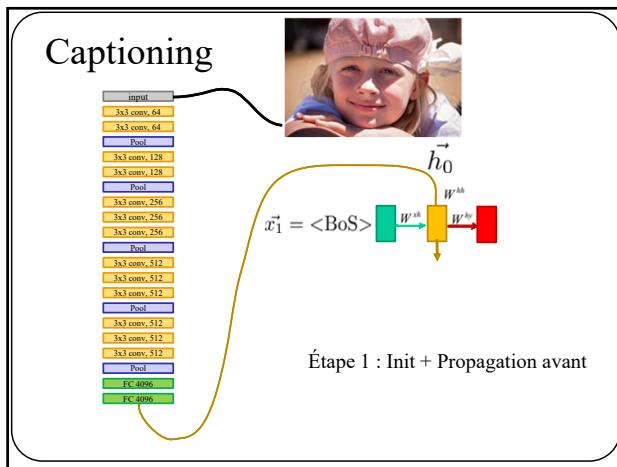
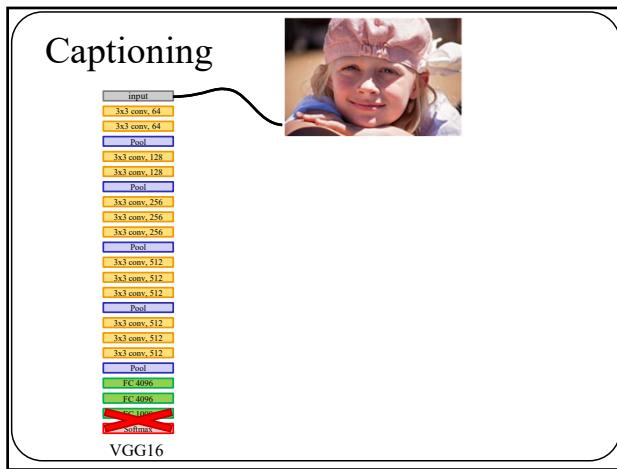
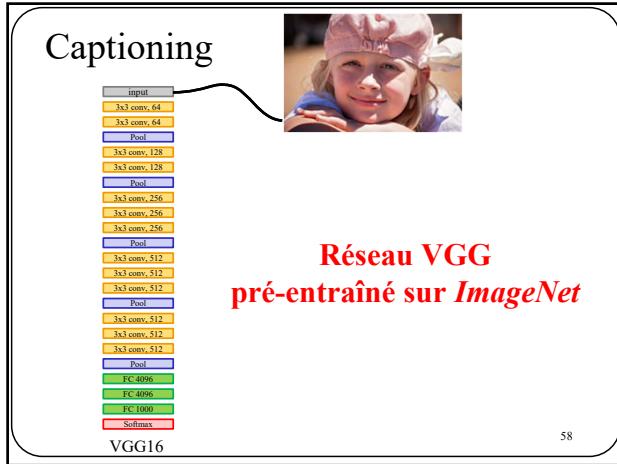
56

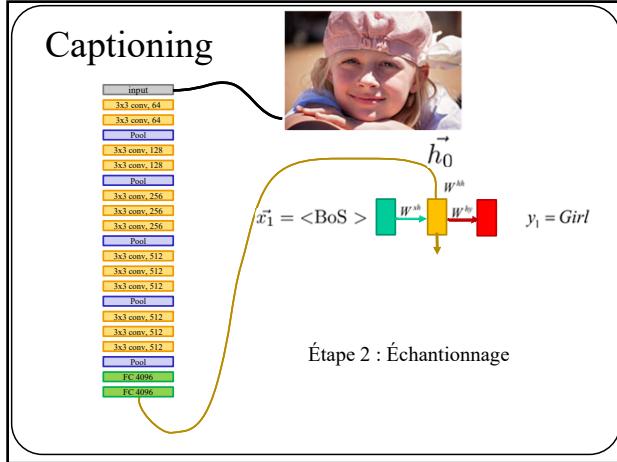
Différentes configurations pour différentes applications

## Description du contenu d'une image (``captioning``)

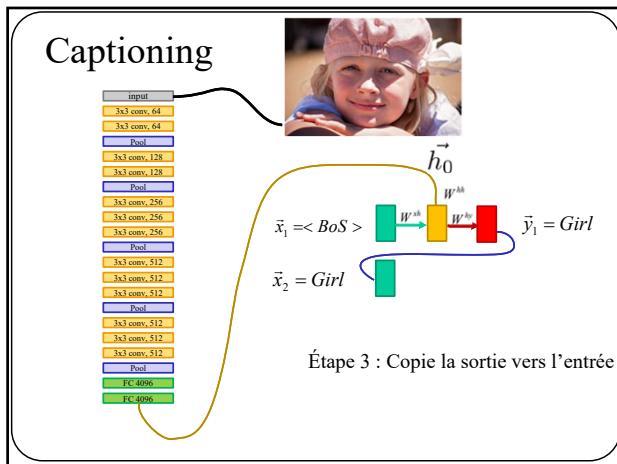


57

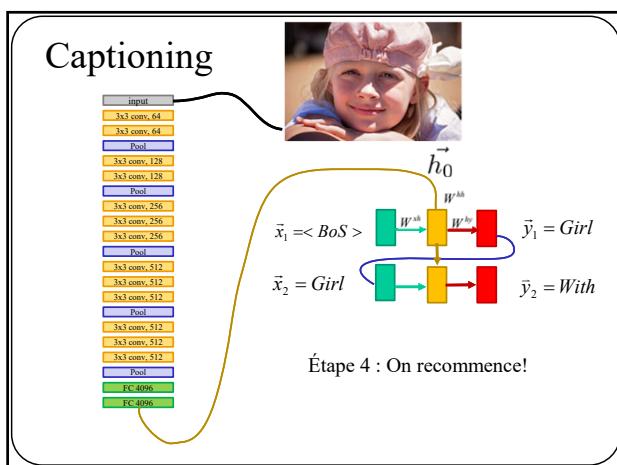




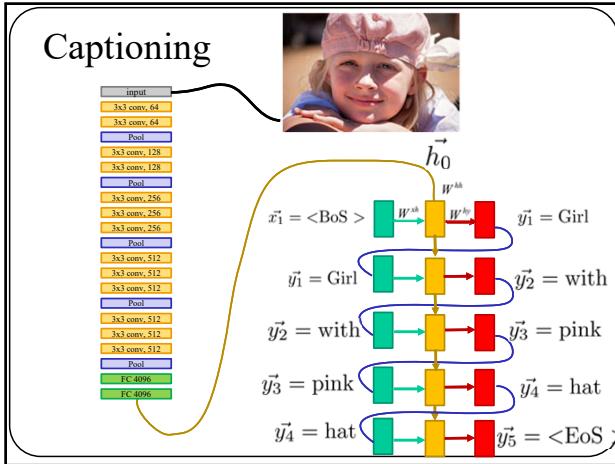
61



62



63



64



65



66

---

---

---

---

---

---

---

---

---

---



---

---

---

---

---

---

---

---

---

---



---

---

---

---

---

---

---

---

---

---

## NeuralTalk and Walk [kwsy2yp hr1frp 24797<5334](#)



67

67

## Analyse de texte

Souvent les modèles de langue utilisent l'encodage « one hot »

Pour des **caractères**...

$$\begin{aligned} 'a' &= [1,0,0,\dots,0] \\ 'b' &= [0,1,0,\dots,0] \\ 'c' &= [0,0,1,\dots,0] \\ &\dots \end{aligned} \quad \left. \right\} \in R^{256}$$

68

68

## Analyse de texte

Souvent les modèles de langue utilisent l'encodage « one hot »

Pour des **mots**...

$$\begin{aligned} \dots \\ 'grand' &= [...,1,0,0,\dots,0] \\ 'grandement' &= [...,0,1,0,\dots,0] \\ 'grandeur' &= [...,0,0,1,\dots,0] \\ &\dots \end{aligned} \quad \left. \right\} \in R^{10,000}$$

69

69

## Prédiction sur des lettres vs. mots

$$\left. \begin{array}{l} 'a' = [1, 0, 0, \dots, 0] \\ 'b' = [0, 1, 0, \dots, 0] \\ 'c' = [0, 0, 1, \dots, 0] \end{array} \right\} \in R^{256} \quad \text{Prédiction sur des lettres}$$

## Prédiction sur des lettres

$$\left. \begin{array}{l} \text{'grand'} = [ \dots, 1, 0, 0, \dots, 0 ] \\ \text{'grandement'} = [ \dots, 0, 1, 0, \dots, 0 ] \\ \text{'grandeur'} = [ \dots, 0, 0, 1, \dots, 0 ] \end{array} \right\} \in R^{10,000} \quad \text{Prédiction sur des mots}$$

## Prédiction sur des mots

70

## Prédiction sur des lettres vs. mots

'*a*'=[1,0,0,...,0]      ]  
'*b*'=[0,1,0,...,0]  
'*c*'=[0,0,1,...,0]  
...  
...  
'grand'  
'granden'  
'grandeur'  
[1,...,0,0,0,1,...,0]      ]  
  
En analyse des langues, un vecteur numérique associé à une séquence de caractères se nomme « JETON » (« token »)

En analyse des langues, un vecteur numérique associé à une séquence de caractères se nomme « JETON » (« token »)

71

## Limites des Jetons « one-hot »

Bien que simple, cet encodage a plusieurs **inconvénients**

- 1- Peu efficace en mémoire lorsque non compressés  
ex.: 10,000 bits pour encoder le mot « je » dans une langue à 10,000 mots!

## 2- Pas de distance sémantique entre les Jetons:

`distance[one-hot('bon'), one-hot('bien')] = distance[one-hot('bon'), one-hot('trottoir')]`

Or, on souhaiterait un **code** tel que

```

Or, on souhaiterait un code tel que
distance[code('bon'), code('bien')] << distance[code('bon'), code('trottoir')]
distance[code('Jean'), code('Chantal')] << distance[code('bon'), code('trottoir')]
distance[code('Inde'), code('Liban')] << distance[code('bon'), code('trottoir')]

```

74

Une solution est d'utiliser l'encodage **Word2Vec** de [Mikolov et al. '13]

Exemple jouet: on veut représenter ces 8 mots par des jetons à 4 éléments

Jeton « one-hot »	Dictionnaire de Jetons
'the'	2 3 4 5
'quick'	-1 -3 -2 2
'brown'	11 6 4 -3
'fox'	-4 8 -4 4
'jumps'	24 -6 42 17
'over'	91 13 14 -5
'lazy'	0 36 4 56
'dog'	-1 0 1 35

1 ligne = code pour 1 mot

75

**Word2Vec** s'appuie sur 2 idées fondamentales

**Idée 1:** Dictionnaire = matrice d'encodage

Comment sélectionner le jeton d'un mot? En multipliant son vecteur One-hot par la matrice d'encodage (le dictionnaire!)

Ex: sélectionner le jeton de « brown »

$$\begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \text{Dictionnaire Jeton} \\ (\text{matrice d'encodage}) \end{pmatrix} = \begin{pmatrix} 11 & 6 & 4 & -3 \end{pmatrix}$$

76

**Word2Vec** s'appuie sur 2 idées fondamentales

**Idée 1:** Dictionnaire jeton = matrice d'encodage

Première couche d'un réseau de neurones  
= matrice d'encodage

$\vec{x}$  : brown     $W^0 \in R^{4 \times 8}$

77

---

---

---

---

---

---

---

---



---

---

---

---

---

---

---

---



---

---

---

---

---

---

---

---

**Word2Vec** s'appuie sur 2 idées fondamentales

**Idée 1:** Dictionnaire = matrice d'encodage

Première couche d'un réseau de neurones  
=  
matrice d'encodage

$$jeton_{\vec{x}} = W^0 \vec{x}$$

78

**Word2Vec** s'appuie sur 2 idées fondamentales

**Idée 1:** Dictionnaire = matrice d'encodage



On pourra donc utiliser un réseau de neurones pour calculer le contenu du dictionnaire

79

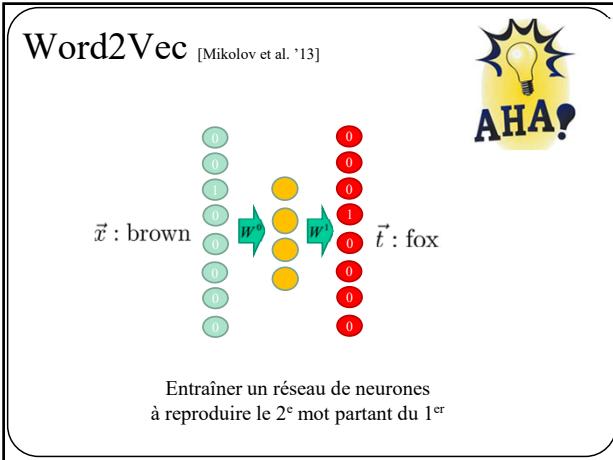
**Word2Vec** s'appuie sur 2 idées fondamentales

**Idée 2:** 2 mots proches dans un texte = 2 mots proches sémantiquement

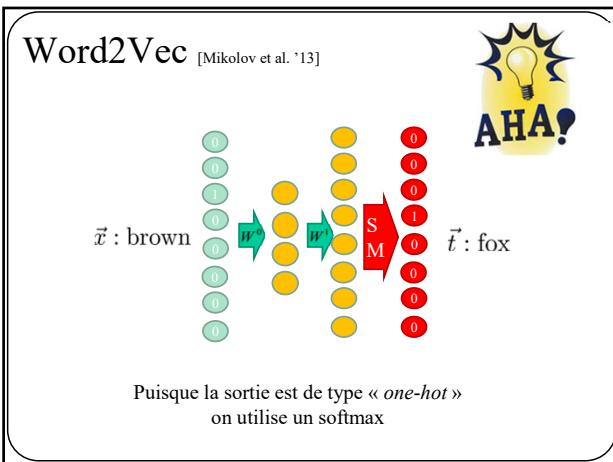
Source Text	Training Samples
The quick brown fox jumps over the lazy dog.	(the, quick) (the, brown)
The quick brown fox jumps over the lazy dog.	(quick, the) (quick, brown) (fox, quick) (fox, quick)
The quick brown fox jumps over the lazy dog.	(brown, the) (brown, quick) (brown, fox) (brown, jump)

Basé sur un corpus de texte, on va créer des millions de paires de mots

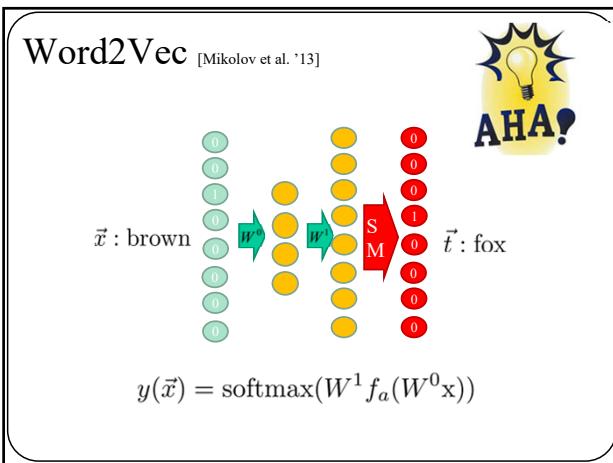
80



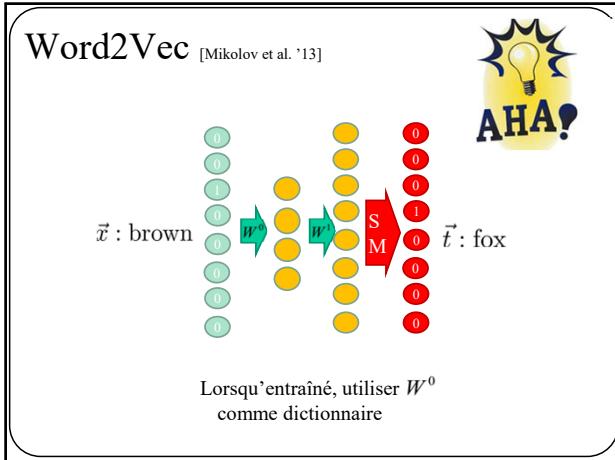
81



82



83



84

---

---

---

---

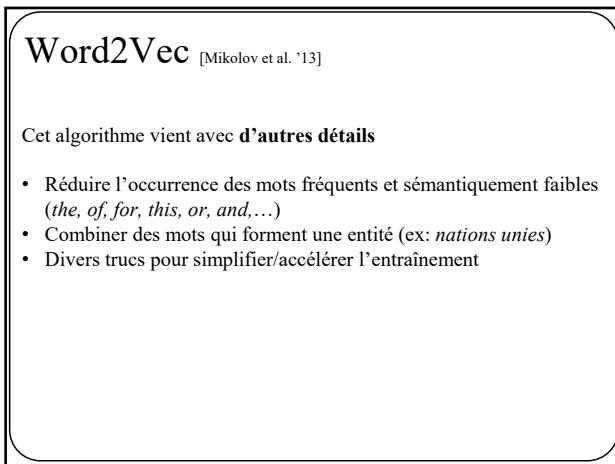
---

---

---

---

---




---

---

---

---

---

---

---

---

---

85

**Distance sémantique entre deux mots = distance entre leur jeton**

Word	First similar word	Second similar word	Third similar word
colosseum	rome (0.994)	roma (0.994)	coliseum (0.994)
colosseo	anfiteatro (0.995)	travel (0.994)	italia (0.994)
scala	aux (0.993)	camelias (0.992)	milano (0.992)
pompeii	retweeted (0.988)	nuovi (0.979)	settembre (0.978)
roma	rome (0.995)	metro (0.994)	colosseum (0.994)
italia	anfiteatro (0.995)	rome (0.995)	colosseo (0.994)
italy	travel (0.998)	davanti (0.997)	photography (0.997)

Word	Similar Words	Similarity	Word	Similar Words	Similarity
Linux	windows	0.85	Twitter	facebook	0.90
	redhat	0.83		instagram	0.86
	unix	0.83		netflix	0.84
	mac os	0.82		snapchat	0.82
	citrix	0.81		google	0.81
	serveurs	0.80		tweets	0.80
	microsoft	0.79		youtube	0.80
	ibm	0.79		linkedin	0.77
windows server	0.79		maddyness	0.77	
	env windows	0.79		tweet	0.77

Ahmia, Oussama & Béchet, Nicolas & Marteau, Pierre-François. Two Multilingual Corpora Extracted from the Tenders Electronic Daily for Machine Learning and Machine Translation Applications in LREC 2018

86

---

---

---

---

---

---

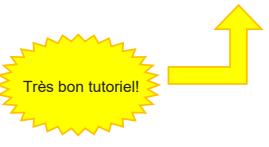
---

---

---

## Word2Vec

<http://mccormickml.com/2016/04/19/word2vec-tutorial-the-skip-gram-model/>



T.Mikolov et al. (2013). "Efficient Estimation of Word Representations in Vector Space", in ICLR 2013

87

---

---

---

---

---

---

---

## Comment entraîner un RNN?

---

---

---

---

---

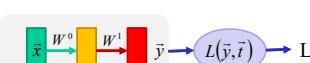
---

---

88

## Histoire de gradients

RN de classification avec entropie croisée



$$\vec{y}(\vec{x}) = S_M \left( W^1 \tanh \left( W^0 \vec{x} \right) \right)$$

$$L = L_{EC} (\vec{y}, \vec{t})$$

89

---

---

---

---

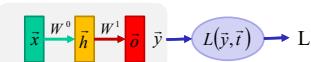
---

---

---

## Histoire de gradients

## Simple RN de classification avec entropie croisée



$$\vec{h} = \tanh(W^0 \vec{x})$$

$$\vec{o} = W^1 \vec{h}$$

$$\vec{y} = S_M(\vec{a})$$

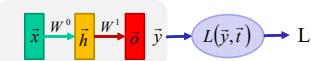
$$L = L_{CE}(\vec{y}, \vec{t})$$

## Propagation avant

90

## Histoire de gradients

## Simple RN de classification avec entropie croisée



$$\vec{h} = \tanh(W^0 \vec{x})$$

$$\vec{o} = W^1 \vec{h}$$

$$\vec{y} = S_M(\vec{x})$$

$$L = L_{CE}(\vec{y}, \vec{t})$$

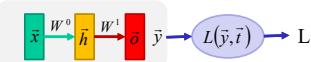
Pour entraîner le réseau  
il faut calculer

$$\nabla_{W^0} L \text{ et } \nabla_{W^1} L$$

91

## Histoire de gradients

## Simple RN de classification avec entropie croisée



$$\vec{h} = \tanh(W^0 \vec{x})$$

$$\vec{o} = W^1 \vec{h}$$

$$\vec{y} = S_M(\vec{a})$$

$$L = L_{CE}(\vec{y}, \vec{t})$$

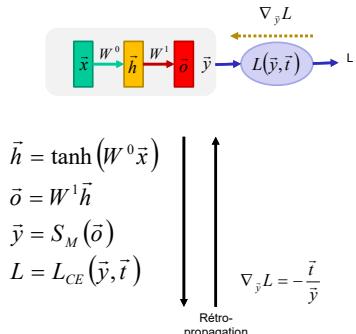
### Dérivée en chaîne

$$\nabla_{w^1} L = \nabla_{\vec{v}} L \nabla_{\vec{v}} \vec{v} \nabla_{w^1} \vec{o}$$

$$\nabla_{\vec{v}} \cdot L_t \equiv \nabla_{\vec{z}} L_t \nabla_{\vec{z}} \vec{v} \nabla_{\vec{z}} \vec{\partial} \nabla_{\vec{z}} \vec{h}$$

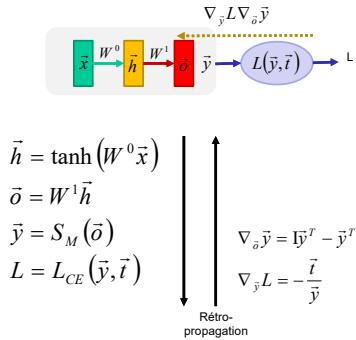
92

## Histoire de gradients



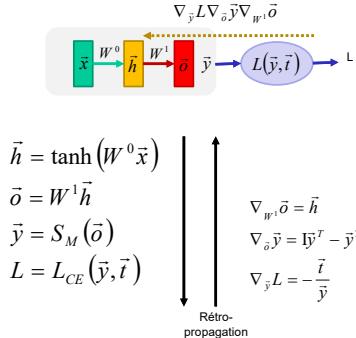
93

## Histoire de gradients



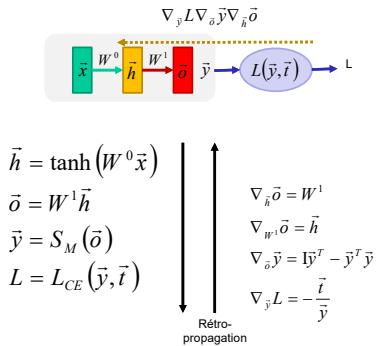
94

## Histoire de gradients



95

## Histoire de gradients



96

---

---

---

---

---

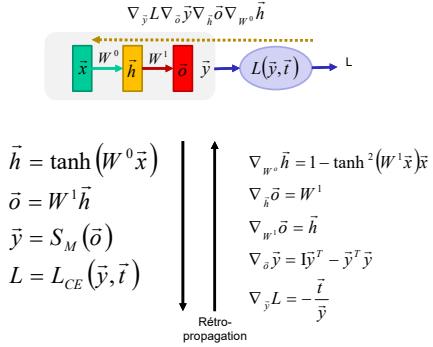
---

---

---

---

## Histoire de gradients



97

---

---

---

---

---

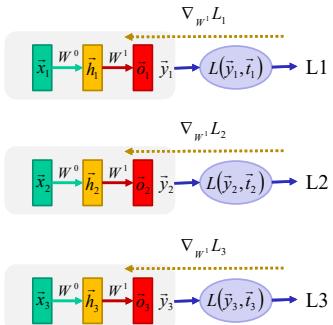
---

---

---

---

## Ex.: 3 données, 3 rétro-propagations



98

---

---

---

---

---

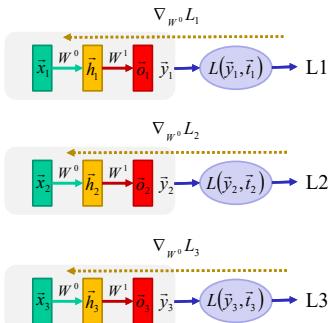
---

---

---

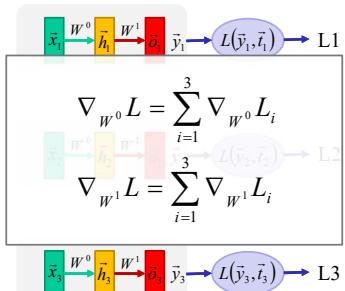
---

Ex.: 3 données, 3 rétro-propagations



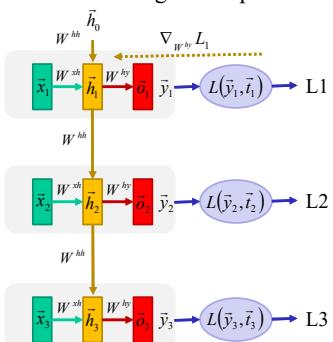
99

3 rétro-propagations



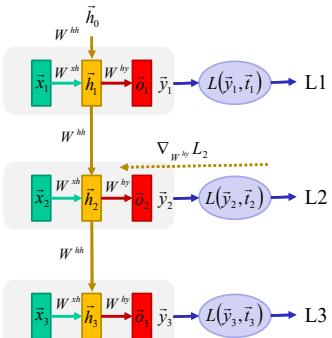
100

Réseau récurrent: gradient pour  $W^{hy}$



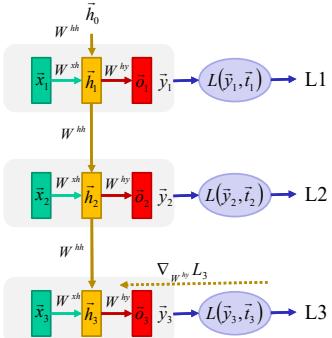
101

### Réseau récurrent: gradient pour $W^{hy}$



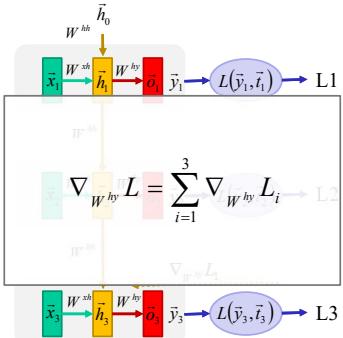
102

### Réseau récurrent: gradient pour $W^{hy}$



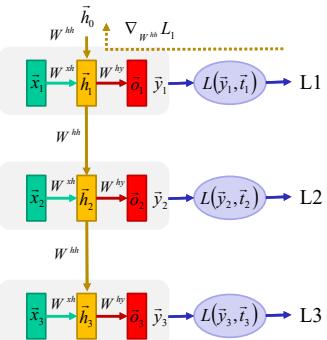
103

### Réseau récurrent: gradient pour $W^{hy}$



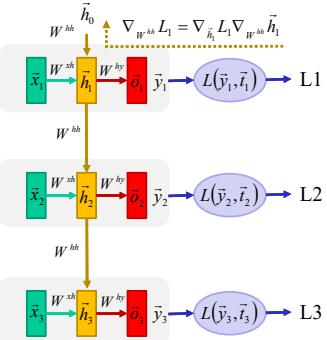
104

### Réseau récurrent: gradient pour $W^{hh}$



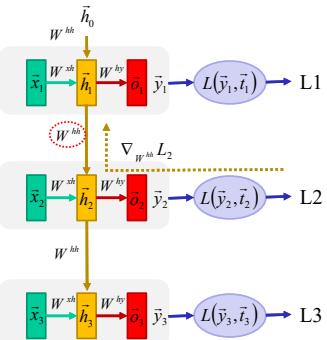
105

### Réseau récurrent: gradient pour $W^{hh}$



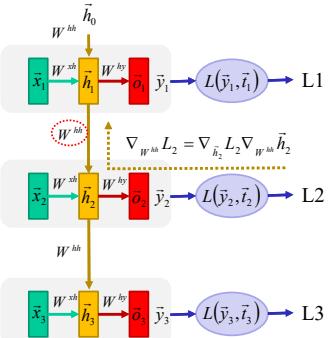
106

### Réseau récurrent: gradient pour $W^{hh}$



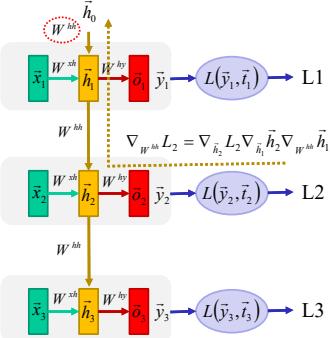
107

### Réseau récurrent: gradient pour $W^{hh}$



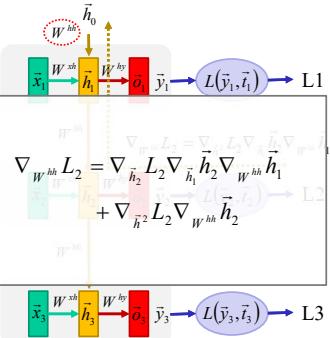
108

### Réseau récurrent: gradient pour $W^{hh}$



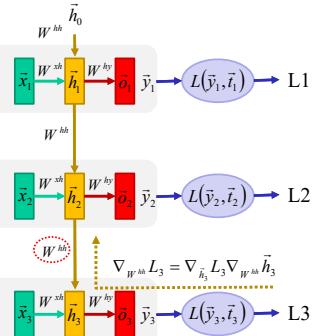
109

### Réseau récurrent: gradient pour $W^{hh}$



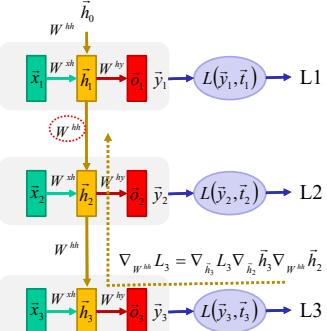
110

### Réseau récurrent: gradient pour $W^{hh}$



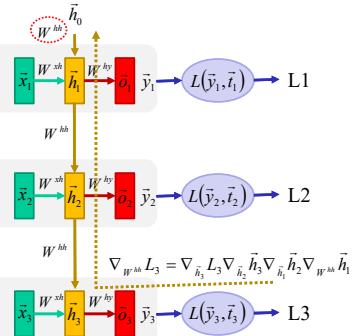
111

### Réseau récurrent: gradient pour $W^{hh}$



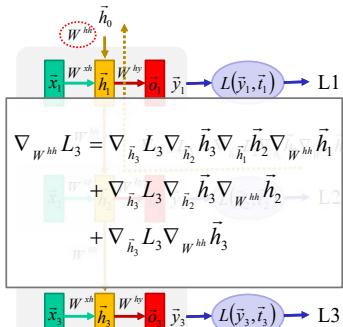
112

### Réseau récurrent: gradient pour $W^{hh}$



113

## Réseau récurrent: gradient pour $W^{hh}$



---

---

---

---

---

---

---

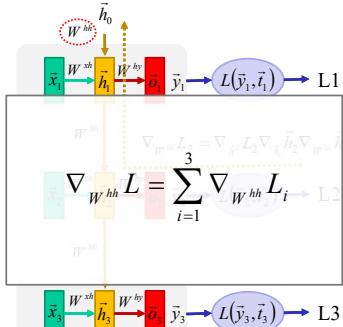
---

---

---

114

Réseau récurrent: gradient pour  $W^{hh}$



---

---

---

---

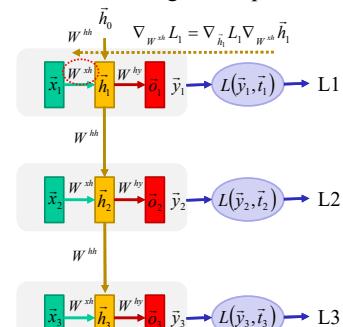
---

---

---

116

Réseau récurrent: gradient pour  $W^{xh}$



---

---

---

---

---

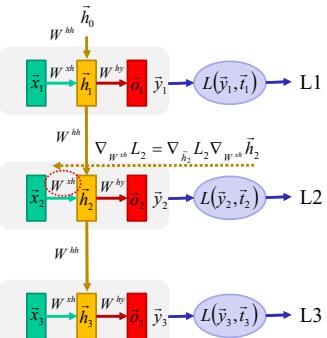
---

---

---

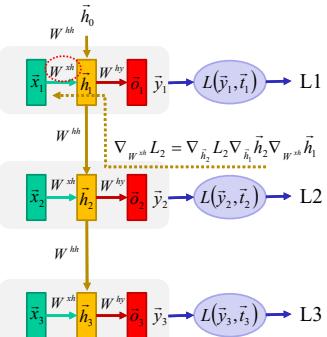
118

### Réseau récurrent: gradient pour $W^{xh}$



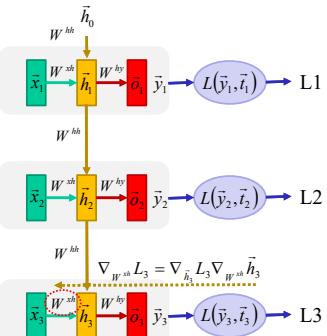
119

### Réseau récurrent: gradient pour $W^{xh}$



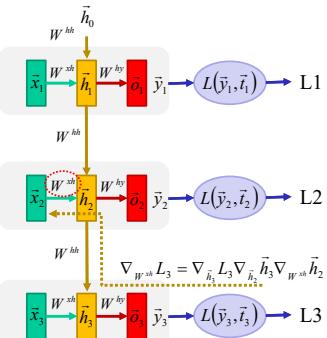
120

### Réseau récurrent: gradient pour $W^{xh}$



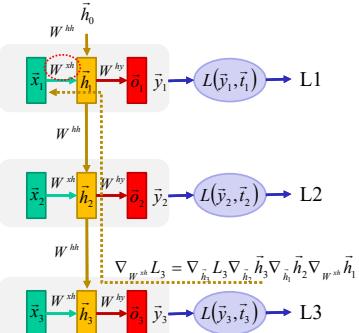
121

### Réseau récurrent: gradient pour $W^{xh}$



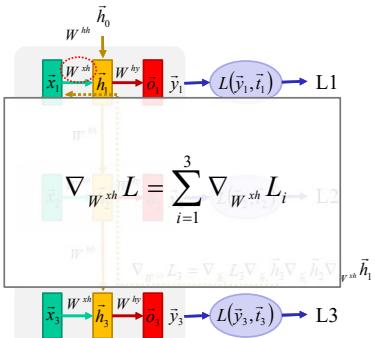
122

### Réseau récurrent: gradient pour $W^{xh}$



123

### Réseau récurrent: gradient pour $W^{xh}$



124

## Réseau récurrent: calcul du gradient

Moins difficile qu'il n'y paraît.

```

43 # backward pass: compute gradients going backwards
44 dwh, dbh, dyb = np.zeros_like(wh), np.zeros_like(bh), np.zeros_like(by)
45 dbh, dyb = np.zeros_like(bh), np.zeros_like(by)
46 dnext = np.zeros_like(hatf1)
47 for t in reversed(xrange(len(inputs))):  

48     dy = np.copy(pst[t])
49     dy[dy.get(t)] = 1 # backprop into y, see http://cs231n.github.io/neural-networks-case-study/#grad if confused here
50     dyb += np.dot(dy, hst[t])
51     dbh += dyb
52     dyb *= dy
53     dh = np.dot(Wh.T, dy) + dnext # backprop into h
54     ddraw = (1 - hst[t] * hst[t]) * dh # backprop through tanh nonlinearity
55     dbh += ddraw
56     dwh += np.dot(ddraw, xst[t])
57     dwh += np.dot(ddraw, hst[t-1])
58     dnext = np.dot(Wh.T, ddraw)
59 for param in [dwh, dbh, dyb, dbh, dyb]:
60     np.clip(param, -5, 5, out=param) # clip to mitigate exploding gradients
61 return dwh, dbh, dyb, dbh, dyb, hst[0:len(inputs)-1]

```

Voir [https://d2l.ai/chapter\\_recurrent-neural-networks/bptt.html](https://d2l.ai/chapter_recurrent-neural-networks/bptt.html) pour plus d'informations

125

Les réseaux récurrents ont un inconvenient majeur:

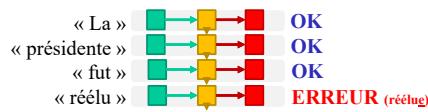
difficile à établir des **relations à longue distance**

126

126

## Exemples: analyse grammaticale

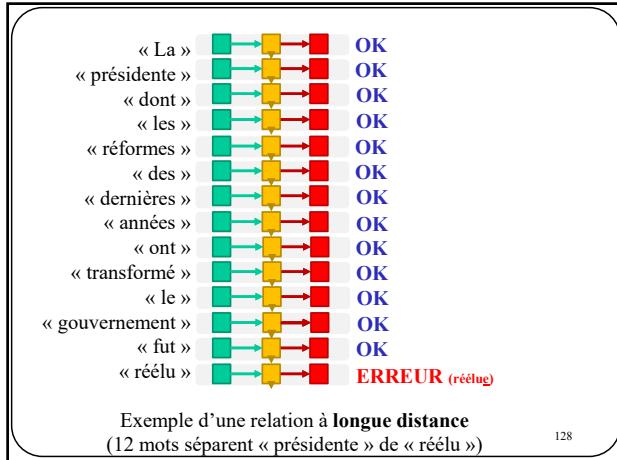
Entraîner un réseau à détecter des erreurs grammaticales



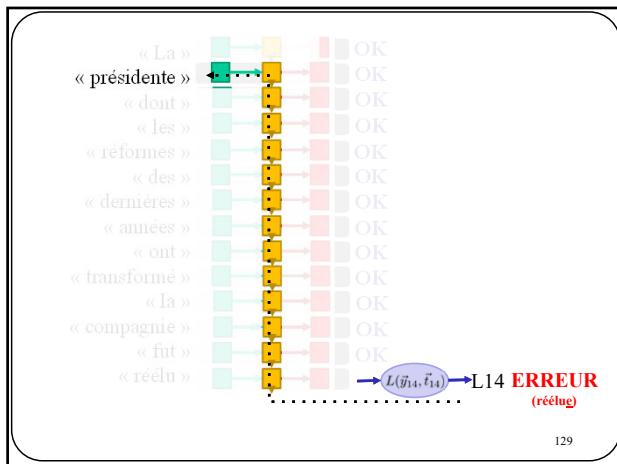
Exemple d'une relation à **courte distance**  
(1 mot sépare « présidente » de « réélu »)

127

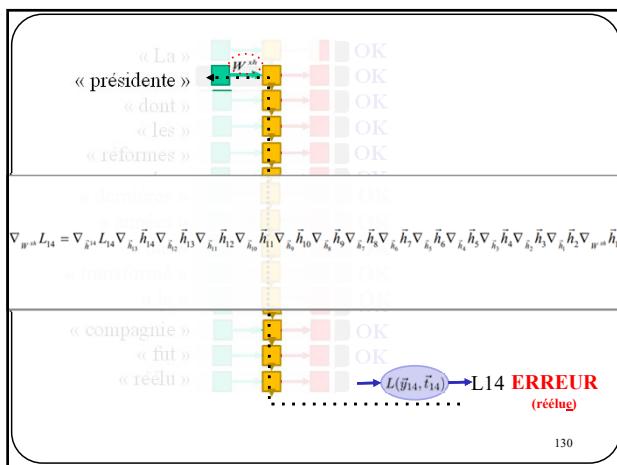
127



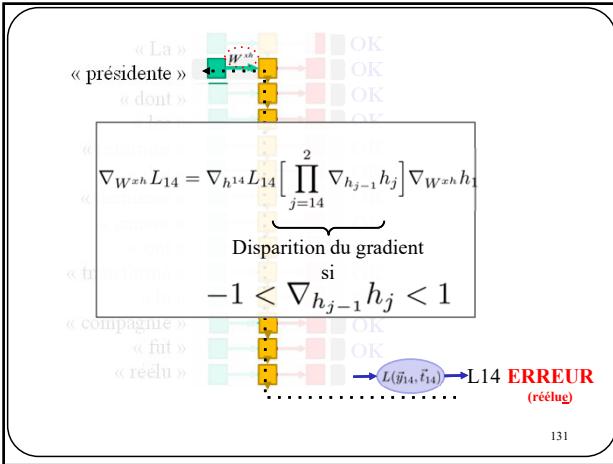
128



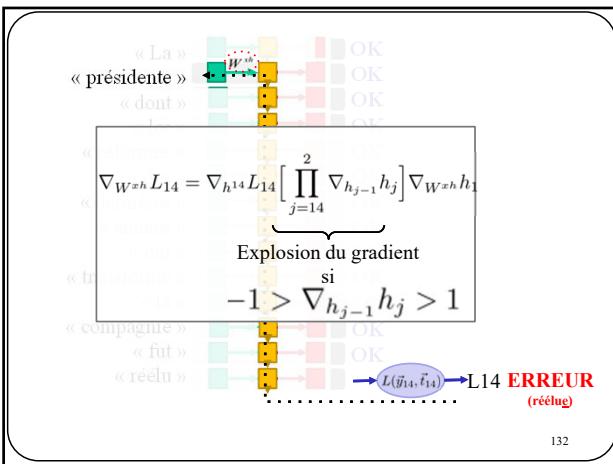
129



130



131

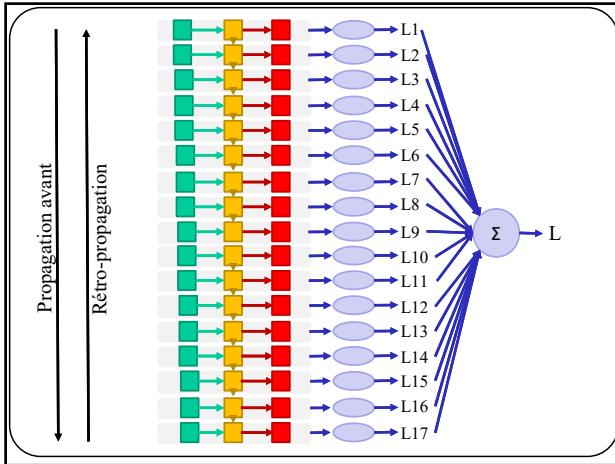


132

## Problème connexe

# Gestion de la mémoire

133



134

---



---



---



---



---



---



---



---



---



---



---



---



---



---



---



---



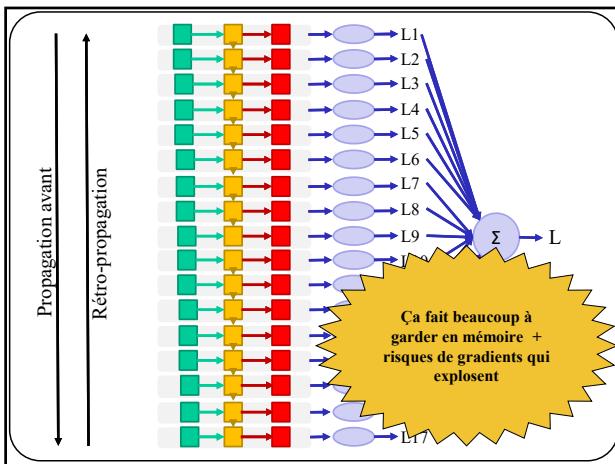
---



---



---



135

---



---



---



---



---



---



---



---



---



---



---



---



---



---



---



---



---



---



---

Solution pour la gestion de la  
mémoire  
**Fenêtres coulissantes**

136

---



---



---



---



---



---



---



---



---



---



---



---



---



---



---



---

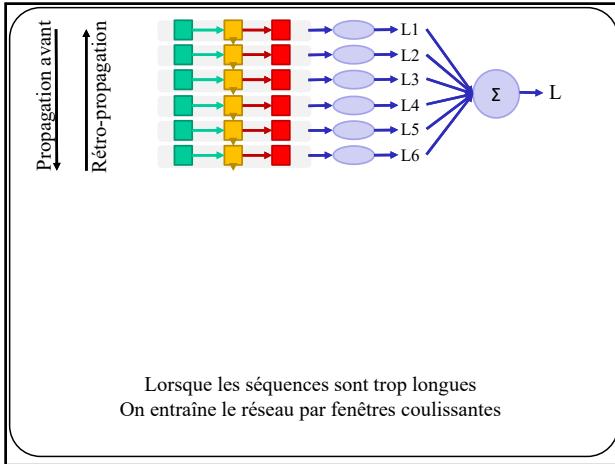


---

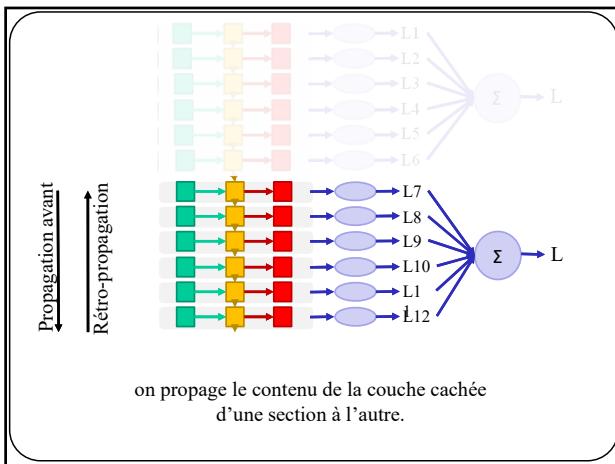


---

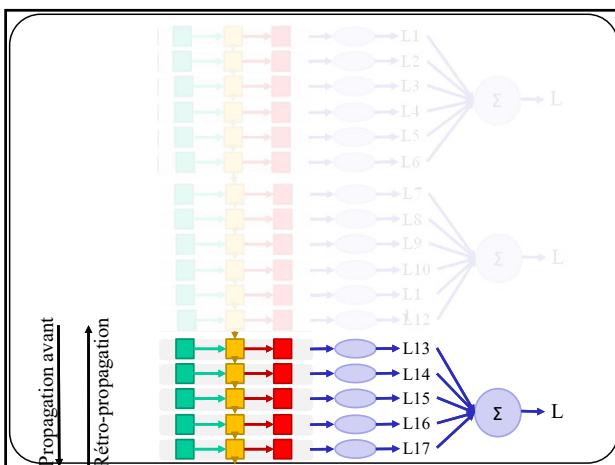
136



137



138



139

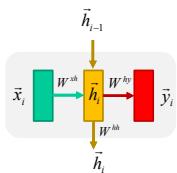
### Solution à la disparition du gradient:

## Gated Recurrent Unit : GRU Long-Short Term Memory : LSTM

140

140

## Illustration + formulation d'un RNN

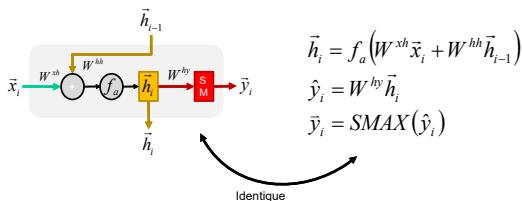


$$\begin{aligned}\vec{h}_i &= f_a \left( W^{xh} \vec{x}_i + W^{hh} \vec{h}_{i-1} \right) \\ \hat{y}_i &= W^{hy} \vec{h}_i \\ \bar{y}_i &= SMAX(\hat{y}_i)\end{aligned}$$

141

141

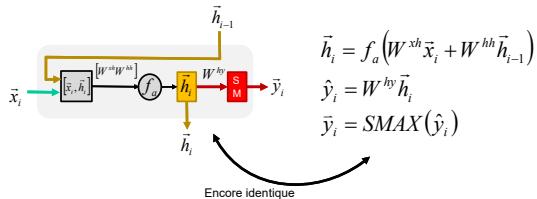
### Autre illustration du même RNN



142

142

## Autre illustration du même RNN



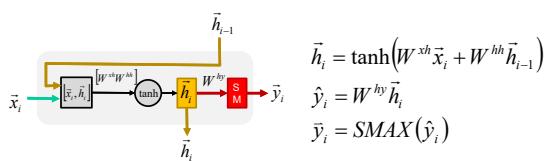
143

143

## GRU (Gated Recurrent Unit)

Modif 1

$$f_a = \tanh$$



144

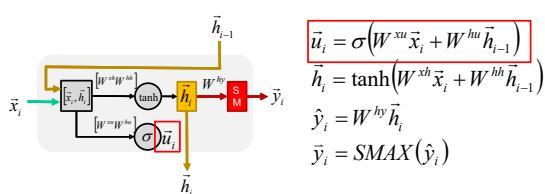
144

## GRU (Gated Recurrent Unit)

Modif 2

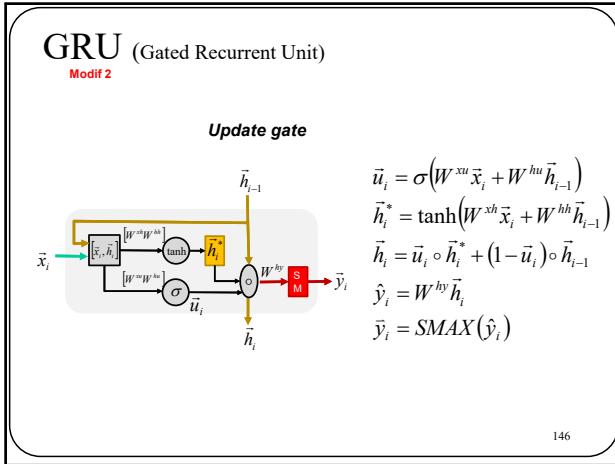
Update gate

$$\sigma = \text{sigmoid}$$

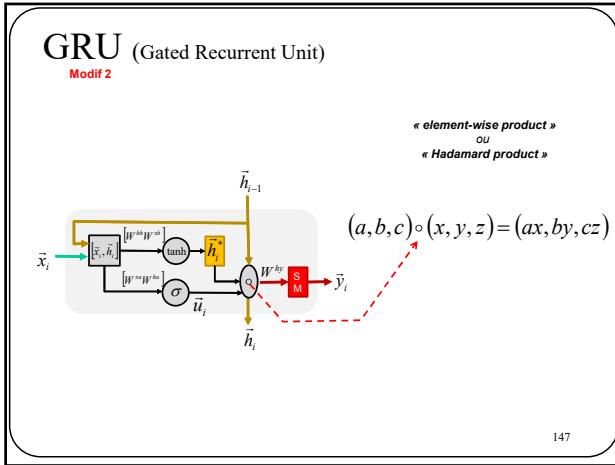


145

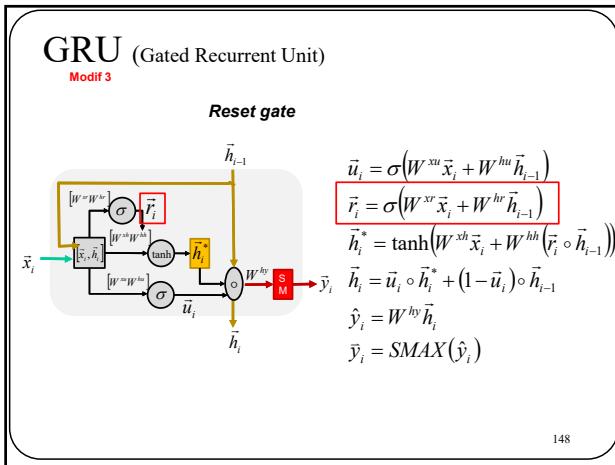
145



146



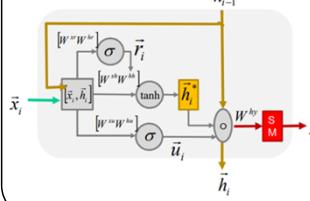
147



148

## Comprendre les gates

$$SI \quad \begin{array}{l} u_i = 1 \\ r_i = 1 \end{array} \quad \left\{ \begin{array}{l} \vec{u}_i = \sigma(W^{xu}\vec{x}_i + W^{hu}\vec{h}_{i-1}) \\ \vec{r}_i = \sigma(W^{xr}\vec{x}_i + W^{hr}\vec{h}_{i-1}) \\ \vec{h}_i^* = \tanh(W^{xh}\vec{x}_i + W^{hh}(\vec{r}_i \circ \vec{h}_{i-1})) \\ \vec{h}_i = \vec{u}_i \circ \vec{h}_i^* + (1 - \vec{u}_i) \circ \vec{h}_{i-1} \\ \hat{y}_i = W^{hy}\vec{h}_i \\ \vec{y}_i = SMAX(\hat{y}_i) \end{array} \right.$$

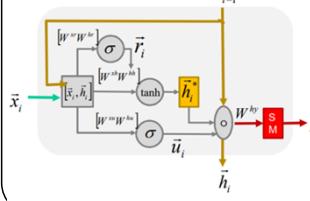


149

149

## Comprendre les gates

$$SI \quad \begin{array}{l} u_i = 1 \\ r_i = 1 \end{array} \quad \left\{ \begin{array}{l} \vec{u}_i = \sigma(W^{xu}\vec{x}_i + W^{hu}\vec{h}_{i-1}) \\ \vec{r}_i = \sigma(W^{xr}\vec{x}_i + W^{hr}\vec{h}_{i-1}) \\ \vec{h}_i^* = \tanh(W^{xh}\vec{x}_i + W^{hh}(\vec{r}_i \circ \vec{h}_{i-1})) \\ \vec{h}_i = \vec{u}_i \circ \vec{h}_i^* + (1 - \vec{u}_i) \circ \vec{h}_{i-1} \\ \hat{y}_i = W^{hy}\vec{h}_i \\ \vec{y}_i = SMAX(\hat{y}_i) \end{array} \right.$$

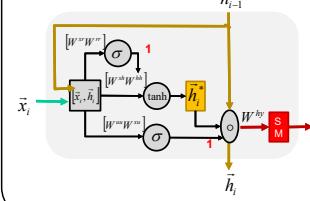


150

150

## Comprendre les gates

$$SI \quad \begin{array}{l} u_i = 1 \\ r_i = 1 \end{array} \quad \left\{ \begin{array}{l} \vec{u}_i = \sigma(W^{xu}\vec{x}_i + W^{hu}\vec{h}_{i-1}) \\ \vec{r}_i = \sigma(W^{xr}\vec{x}_i + W^{hr}\vec{h}_{i-1}) \\ \vec{h}_i^* = \tanh(W^{xh}\vec{x}_i + W^{hh}(\vec{r}_i \circ \vec{h}_{i-1})) \\ \vec{h}_i = \vec{u}_i \circ \vec{h}_i^* + (1 - \vec{u}_i) \circ \vec{h}_{i-1} \\ \hat{y}_i = W^{hy}\vec{h}_i \\ \vec{y}_i = SMAX(\hat{y}_i) \end{array} \right.$$

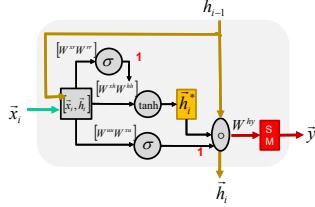


151

151

## Comprendre les gates

$$SI \quad \begin{cases} \bar{u}_i = 1 \\ \bar{r}_i = 1 \end{cases} \quad \begin{cases} \bar{h}_i^* = \tanh(W^{xh}\vec{x}_i + W^{hh}\bar{h}_{i-1}) \\ \bar{h}_i = \bar{h}_i^* \\ \hat{y}_i = W^{hy}\bar{h}_i \\ \vec{y}_i = SMAX(\hat{y}_i) \end{cases}$$

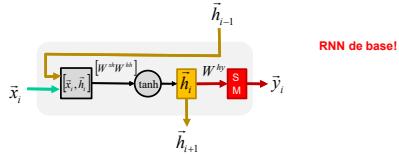


152

152

## Comprendre les gates

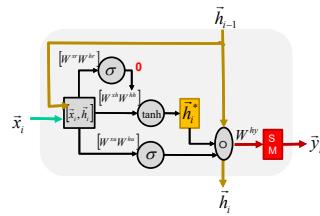
$$SI \quad \begin{cases} \bar{u}_i = 1 \\ \bar{r}_i = 1 \end{cases} \quad \begin{cases} \bar{h}_i^* = \tanh(W^{xh}\vec{x}_i + W^{hh}\bar{h}_{i-1}) \\ \bar{h}_i = \bar{h}_i^* \\ \hat{y}_i = W^{hy}\bar{h}_i \\ \vec{y}_i = SMAX(\hat{y}_i) \end{cases}$$



153

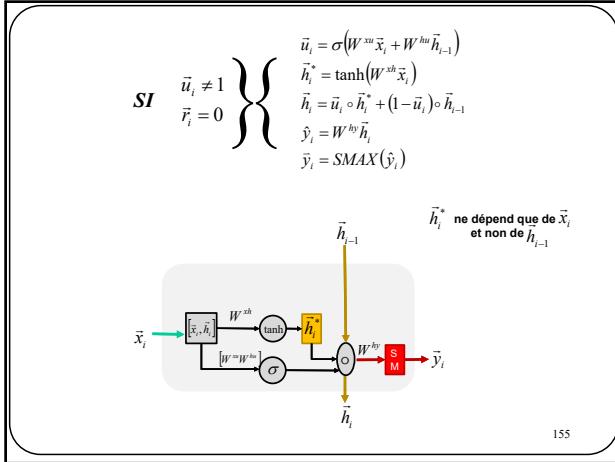
153

$$SI \quad \begin{cases} \bar{u}_i \neq 1 \\ \bar{r}_i = 0 \end{cases} \quad \begin{cases} \bar{u}_i = \sigma(W^{uu}\vec{x}_i + W^{hu}\bar{h}_{i-1}) \\ \bar{r}_i = \sigma(W^{ur}\vec{x}_i + W^{hr}\bar{h}_{i-1}) \\ \bar{h}_i^* = \tanh(W^{xh}\vec{x}_i + W^{hh}\bar{h}_{i-1}) \\ \bar{h}_i = \bar{u}_i \circ \bar{h}_i^* + (1 - \bar{u}_i) \circ \bar{h}_{i-1} \\ \hat{y}_i = W^{hy}\bar{h}_i \\ \vec{y}_i = SMAX(\hat{y}_i) \end{cases}$$

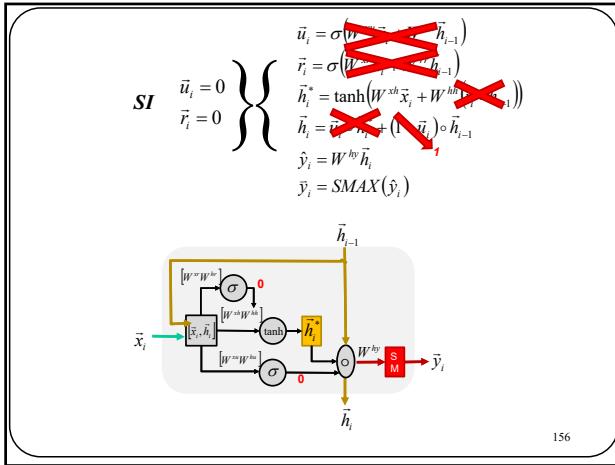


154

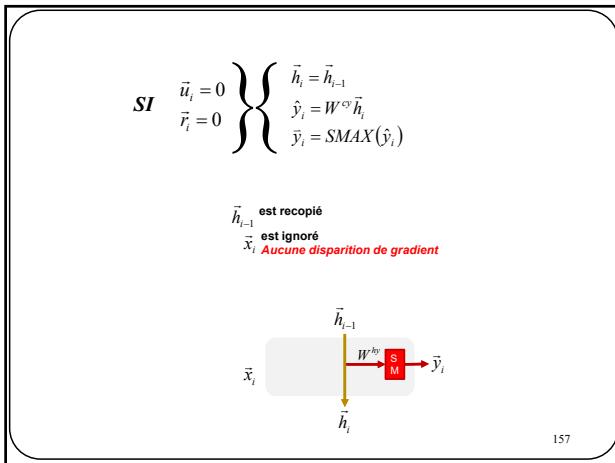
154



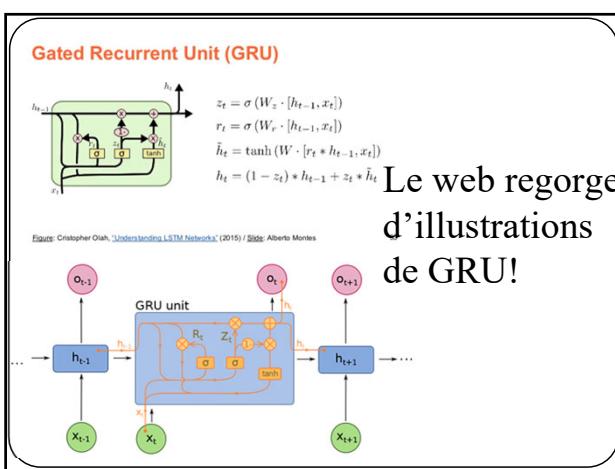
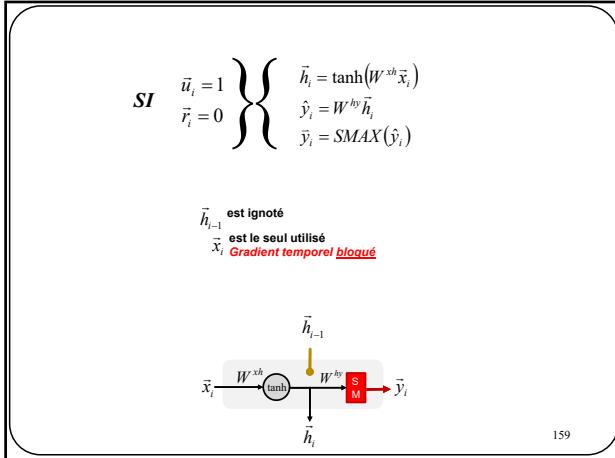
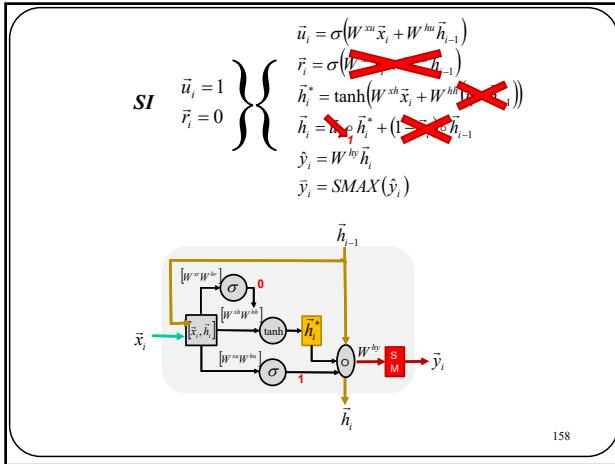
155



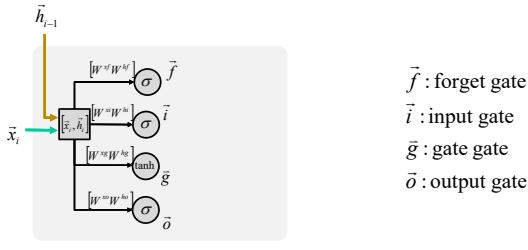
156



157



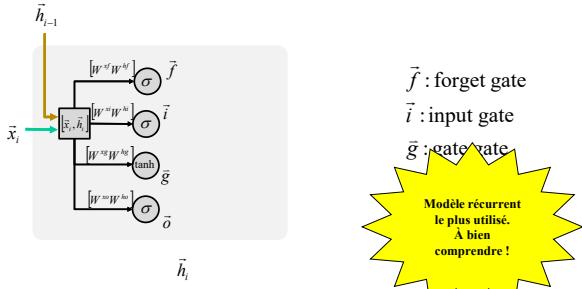
## LSTM (Long Short Term Memory)



Hochreiter and Schmidhuber, "Long Short Term Memory", Neural Computation, 1997

161

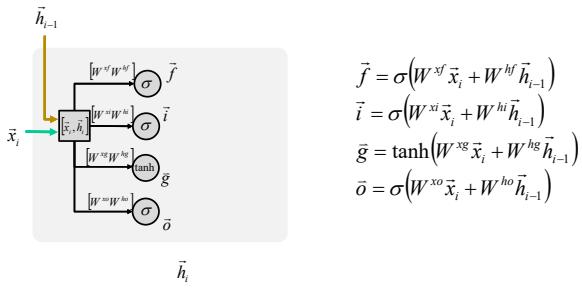
## LSTM (Long Short Term Memory)



Hochreiter and Schmidhuber, "Long Short Term Memory", Neural Computation, 1997

162

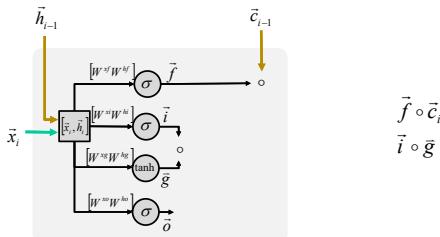
## LSTM (Long Short Term Memory)



163

163

## LSTM (Long Short Term Memory)



164

---



---



---



---



---



---



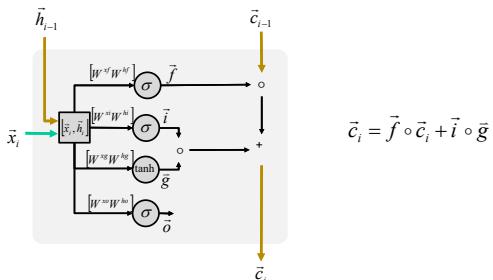
---



---

164

## LSTM (Long Short Term Memory)



165

---



---



---



---



---



---



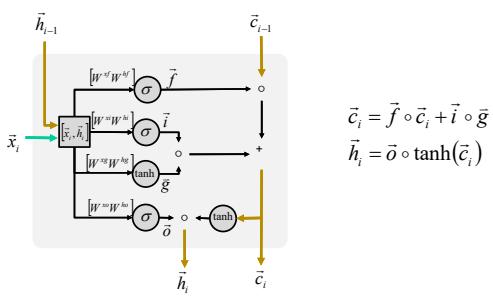
---



---

165

## LSTM (Long Short Term Memory)



166

---



---



---



---



---



---



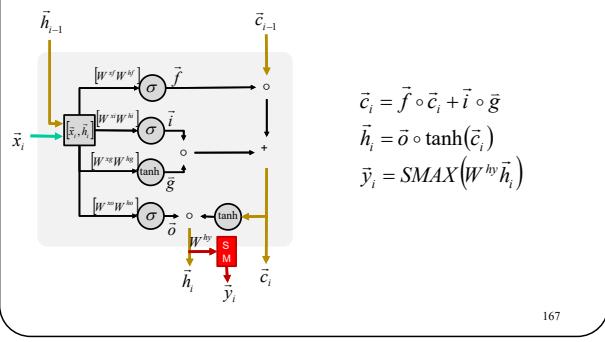
---



---

166

## LSTM (Long Short Term Memory)

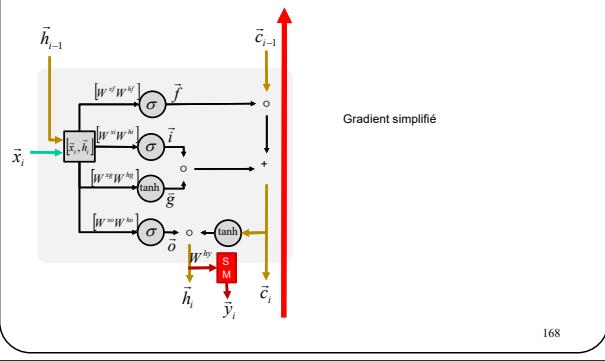


$$\begin{aligned}\vec{c}_i &= \vec{f} \circ \vec{c}_i + \vec{i} \circ \vec{g} \\ \vec{h}_i &= \vec{o} \circ \tanh(\vec{c}_i) \\ \vec{y}_i &= SMAX\left(W^{hy}\vec{h}_i\right)\end{aligned}$$

167

167

## LSTM (Long Short Term Memory)

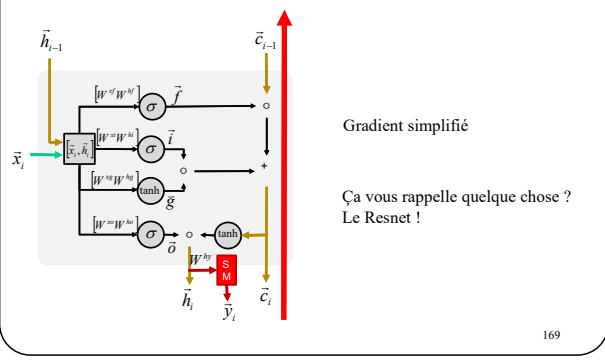


## Gradient simplifié

168

168

## LSTM (Long Short Term Memory)

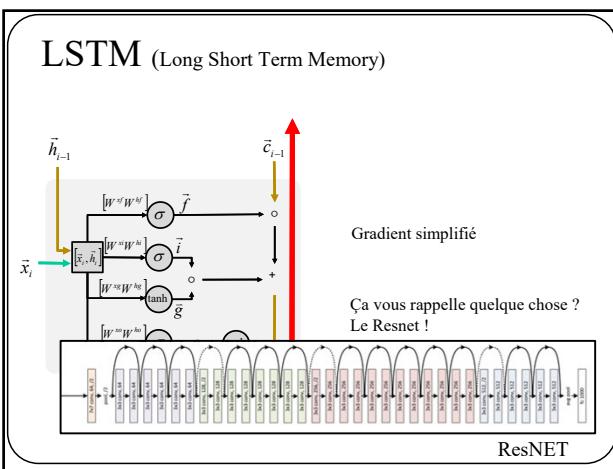


## Gradient simplifié

Ça vous rappelle quelque chose ?  
Le Resnet !

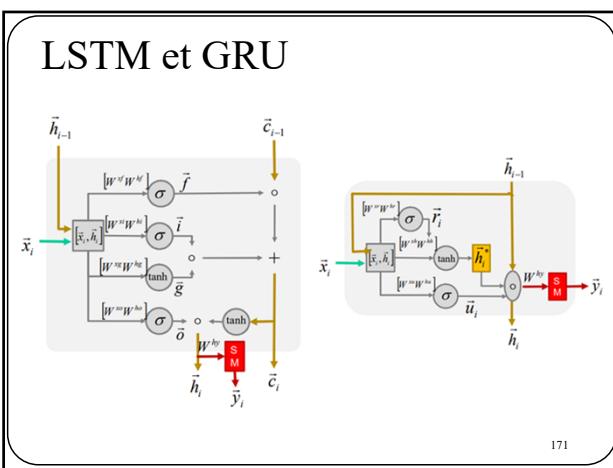
169

## LSTM (Long Short Term Memory)



170

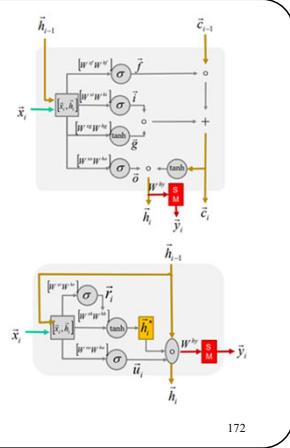
LSTM et GRU



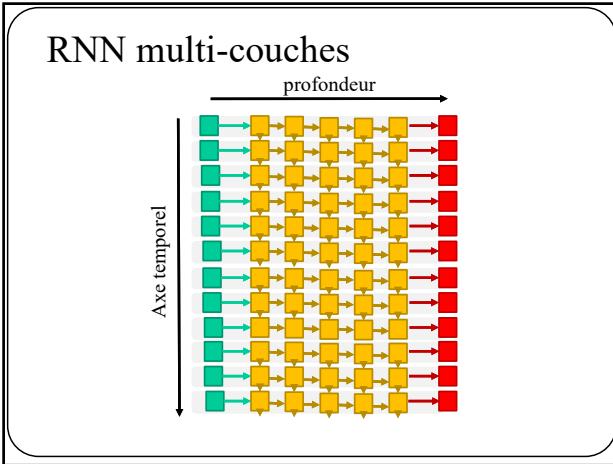
171

## LSTM et GRU

- Servent à protéger le gradient
  - Conçus empiriquement
  - GRU légèrement plus simple
  - Les "gates" ne servent qu'à bloquer ou permettre à l'information (données ou temporelle) de passer



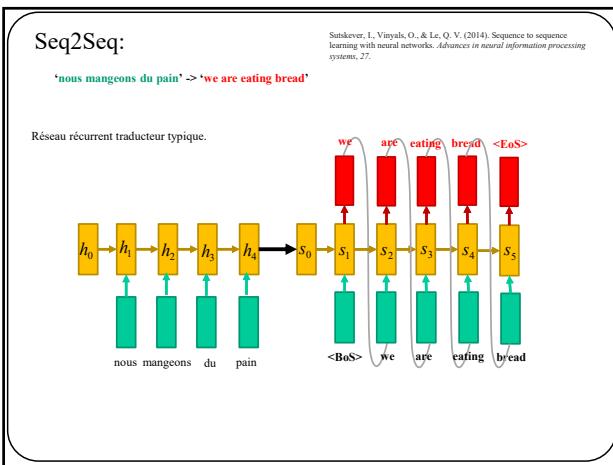
172



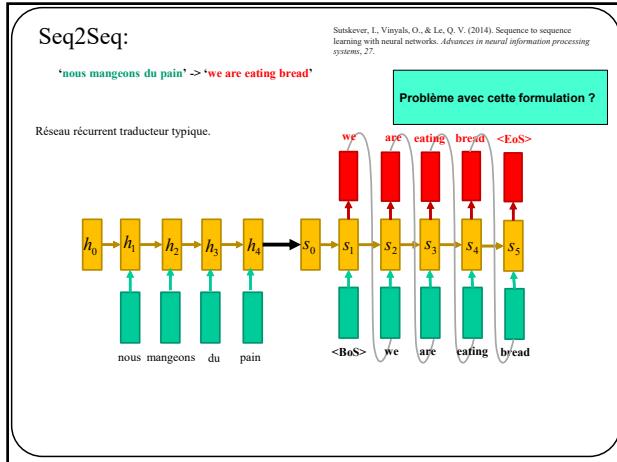
173

## Modèles d'attention

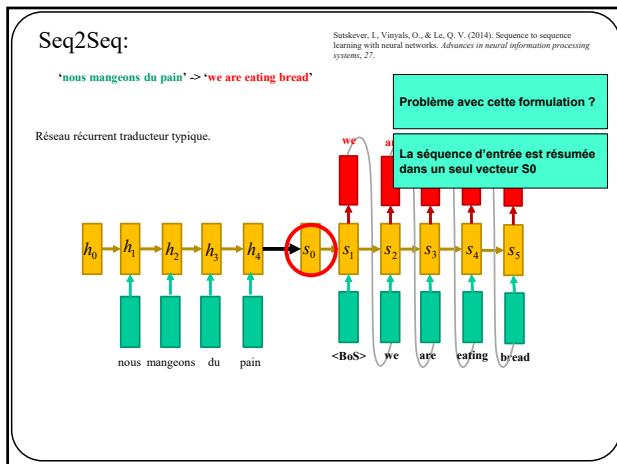
174



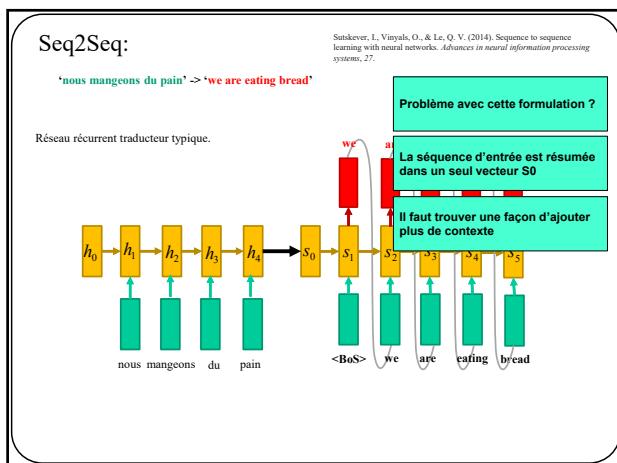
175



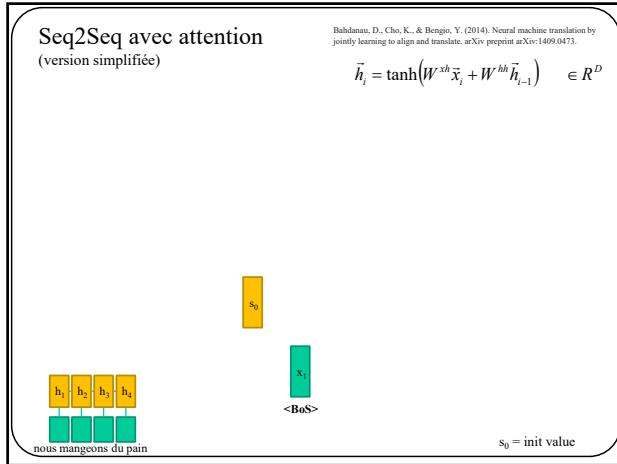
176



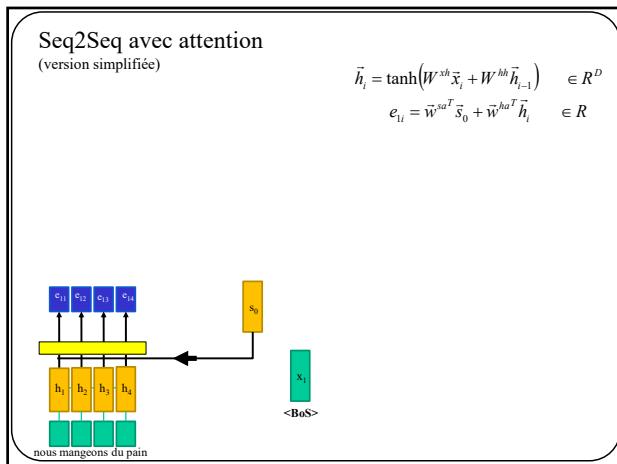
177



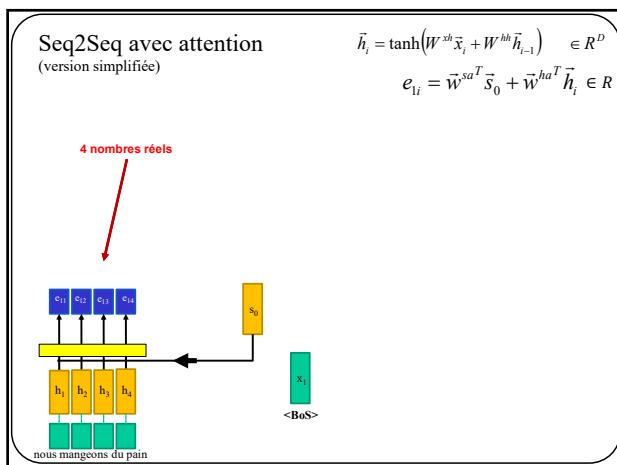
178



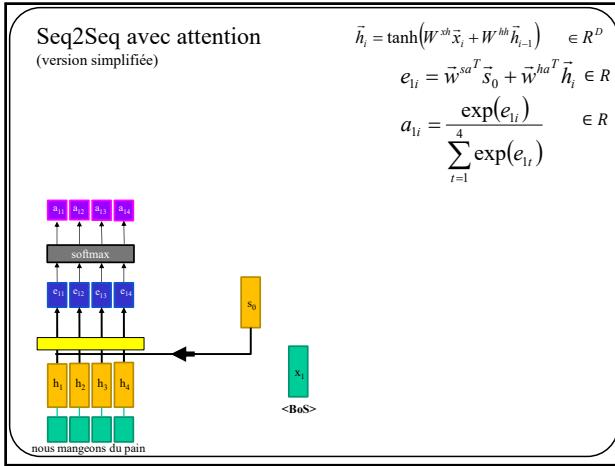
179



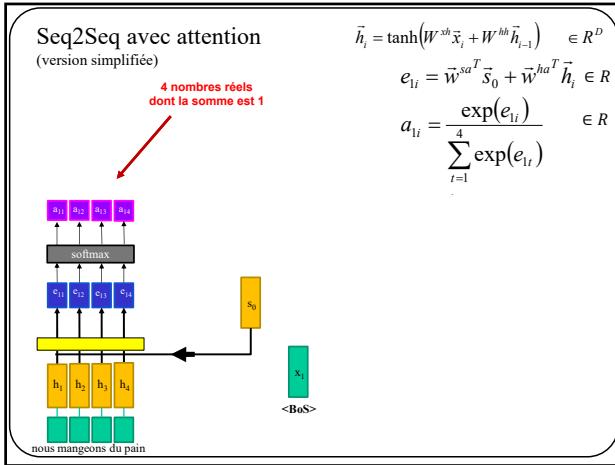
180



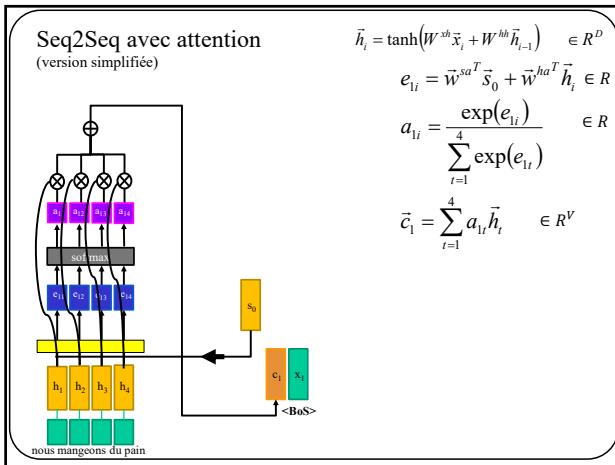
181



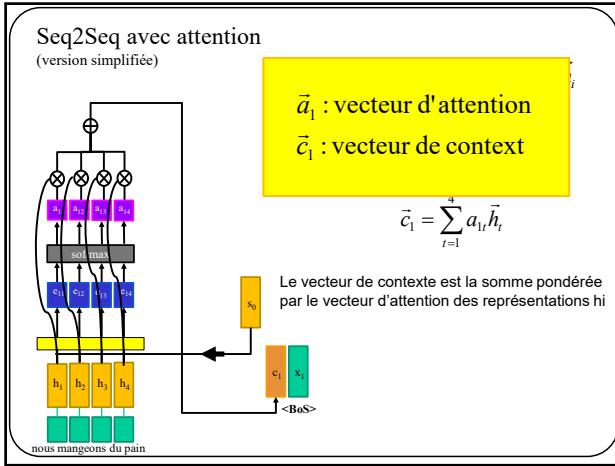
182



183



184



185

---



---



---



---



---



---



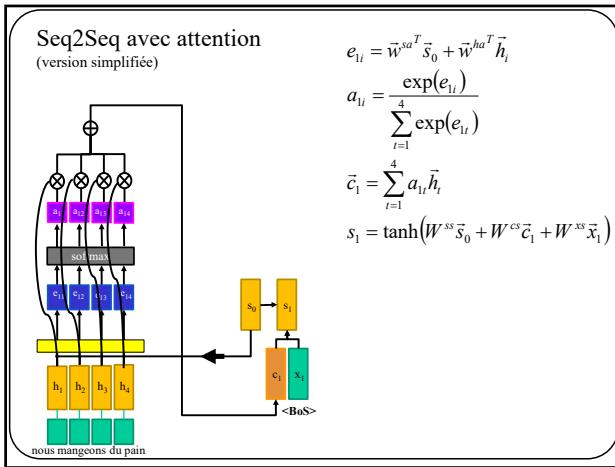
---



---



---



186

---



---



---



---



---



---



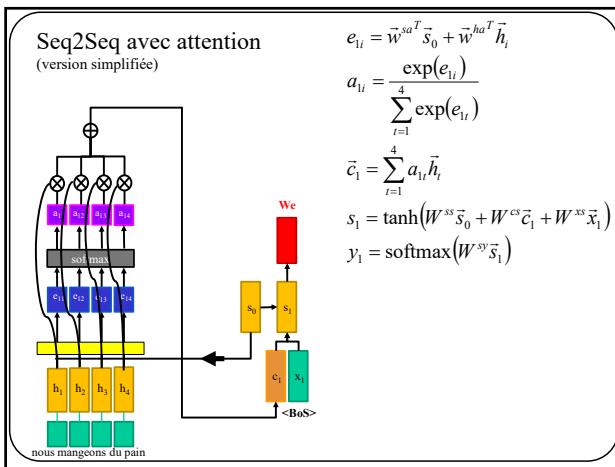
---



---



---



187

---



---



---



---



---



---



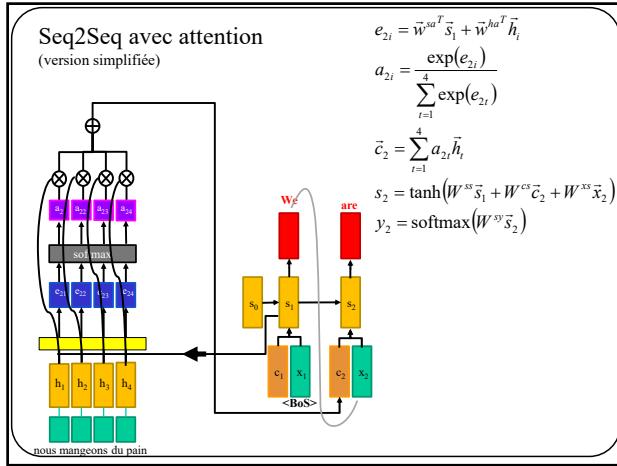
---



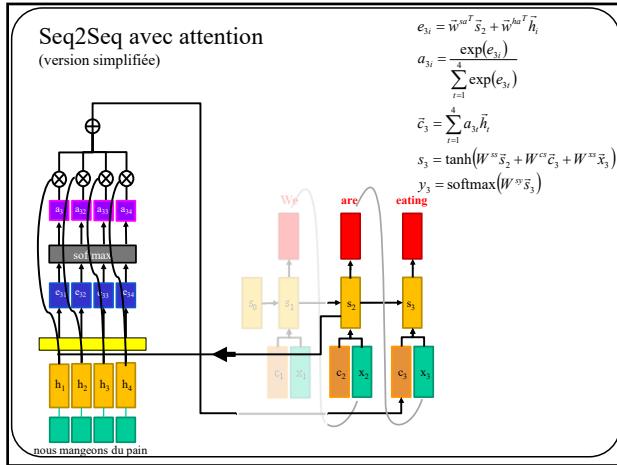
---



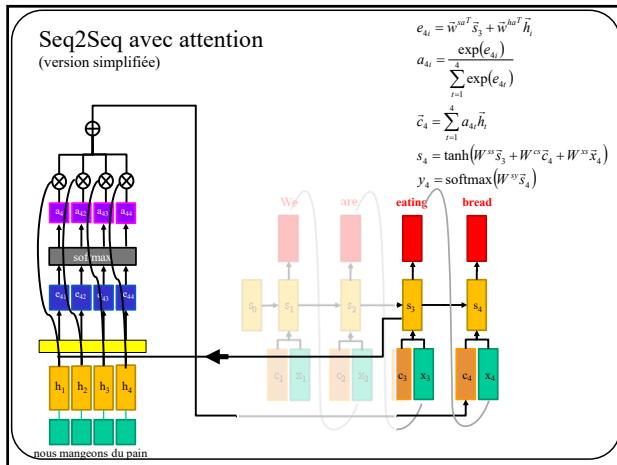
---



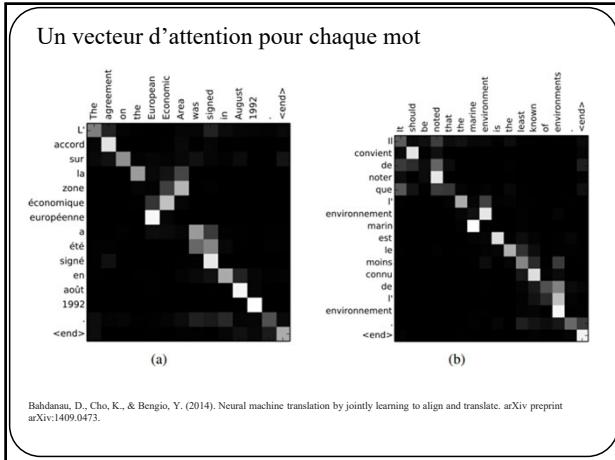
188



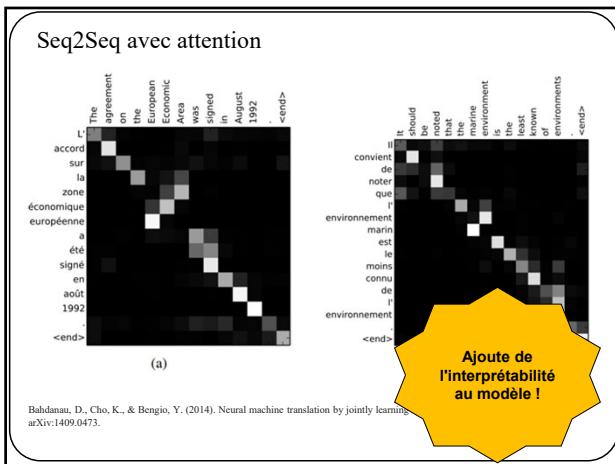
189



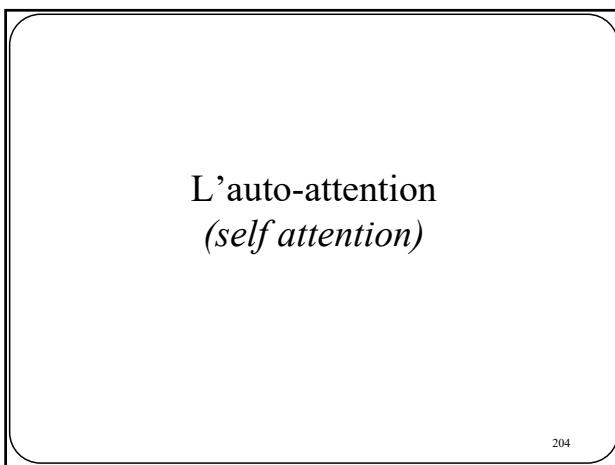
190



191



192



204

Revenons à la base : **multiplication matricielle**

Considérons les 4 matrices suivantes

$$X = \begin{pmatrix} x_{11} & x_{12} & x_{13} & x_{14} \\ x_{21} & x_{22} & x_{23} & x_{24} \\ x_{31} & x_{32} & x_{33} & x_{34} \end{pmatrix} \in R^{3 \times 4}$$

$$W^q = \begin{pmatrix} W^q_{11} & W^q_{12} & W^q_{13} \\ W^q_{21} & W^q_{22} & W^q_{23} \\ W^q_{31} & W^q_{32} & W^q_{33} \end{pmatrix} \in R^{3 \times 3}$$

$$W^k = \begin{pmatrix} W^k_{11} & W^k_{12} & W^k_{13} \\ W^k_{21} & W^k_{22} & W^k_{23} \\ W^k_{31} & W^k_{32} & W^k_{33} \end{pmatrix} \in R^{3 \times 3}$$

$$W^V = \begin{pmatrix} W^V_{11} & W^V_{12} & W^V_{13} \\ W^V_{21} & W^V_{22} & W^V_{23} \end{pmatrix} \in R^{2 \times 3}$$

205

205

Revenons à la base : **multiplication matricielle**

Leur multiplication donne:

$$X = \begin{pmatrix} x_{11} & x_{12} & x_{13} & x_{14} \\ x_{21} & x_{22} & x_{23} & x_{24} \\ x_{31} & x_{32} & x_{33} & x_{34} \end{pmatrix} \in R^{3 \times 4}$$

$$W^q X = Q = \begin{pmatrix} Q_{11} & Q_{12} & Q_{13} & Q_{14} \\ Q_{21} & Q_{22} & Q_{23} & Q_{24} \\ Q_{31} & Q_{32} & Q_{33} & Q_{34} \end{pmatrix} \in R^{3x4}$$

$$W^q = \begin{pmatrix} W_{11}^q & W_{12}^q & W_{13}^q \\ W_{21}^q & W_{22}^q & W_{23}^q \\ W_{31}^q & W_{32}^q & W_{33}^q \end{pmatrix} \in R^{3 \times 3}$$

$$W^K X = K = \begin{pmatrix} K^{-11} & K^{-12} & K^{-13} & K^{-14} \\ K^x_{21} & K^x_{22} & K^x_{23} & K^x_{24} \\ K^x_{31} & K^x_{32} & K^x_{33} & K^x_{34} \end{pmatrix} \in R^{3 \times 4}$$

$$W^k = \begin{pmatrix} W^k_{11} & W^k_{12} & W^k_{13} \\ W^k_{21} & W^k_{22} & W^k_{23} \\ W^k_{31} & W^k_{32} & W^k_{33} \end{pmatrix} \in R^{3 \times 3}$$

$$W^V X = V = \begin{pmatrix} V_{11}^x & V_{12}^x & V_{13}^x & V_{14}^x \\ V_{21}^x & V_{22}^x & V_{23}^x & V_{24}^x \end{pmatrix} \in R^{2 \times 4}$$

206

206

### Auto attention

$X$  est une matrice de données pour laquelle chaque colonne  $i$  correspond au jeton d'une mot  $\bar{x}_i$

$$X = \begin{pmatrix} x_{11} & x_{12} & x_{13} & x_{14} \\ x_{21} & x_{22} & x_{23} & x_{24} \\ x_{31} & x_{32} & x_{33} & x_{34} \end{pmatrix} \in R^{3 \times 4}$$

The matrix  $X$  is shown as a 3x4 grid of elements. Below the grid, four arrows point upwards from the labels "Nous", "mangons", "du", and "pain" to the first, second, third, and fourth columns respectively. The first column is green, the second is orange, the third is red, and the fourth is purple.

Dans cet exemple, 4 mots en entrée donc 4 colonnes dans X  
Les jetons peuvent être obtenus par [Word2Vec](#)

207

207

### Auto attention

$X = \begin{pmatrix} x_{11} & x_{12} & x_{13} & x_{14} \\ x_{21} & x_{22} & x_{23} & x_{24} \\ x_{31} & x_{32} & x_{33} & x_{34} \end{pmatrix} \in R^{3 \times 4}$

$W^q = \begin{pmatrix} W^q_{11} & W^q_{12} & W^q_{13} \\ W^q_{21} & W^q_{22} & W^q_{23} \\ W^q_{31} & W^q_{32} & W^q_{33} \end{pmatrix} \in R^{3 \times 3}$

$W^k = \begin{pmatrix} W^k_{11} & W^k_{12} & W^k_{13} \\ W^k_{21} & W^k_{22} & W^k_{23} \\ W^k_{31} & W^k_{32} & W^k_{33} \end{pmatrix} \in R^{3 \times 3}$

$W^v = \begin{pmatrix} W^v_{11} & W^v_{12} & W^v_{13} \\ W^v_{21} & W^v_{22} & W^v_{23} \end{pmatrix} \in R^{2 \times 3}$

W : Matrices de paramètres appris par rétropropagation

208

208

### Auto attention

$X = \begin{pmatrix} x_{11} & x_{12} & x_{13} & x_{14} \\ x_{21} & x_{22} & x_{23} & x_{24} \\ x_{31} & x_{32} & x_{33} & x_{34} \end{pmatrix} \in R^{3 \times 4}$

$W^q = \begin{pmatrix} W^q_{11} & W^q_{12} & W^q_{13} \\ W^q_{21} & W^q_{22} & W^q_{23} \\ W^q_{31} & W^q_{32} & W^q_{33} \end{pmatrix} \in R^{3 \times 0}$

$W^k = \begin{pmatrix} W^k_{11} & W^k_{12} & W^k_{13} \\ W^k_{21} & W^k_{22} & W^k_{23} \\ W^k_{31} & W^k_{32} & W^k_{33} \end{pmatrix} \in R^{3 \times 0}$

$W^v = \begin{pmatrix} W^v_{11} & W^v_{12} & W^v_{13} \\ W^v_{21} & W^v_{22} & W^v_{23} \end{pmatrix} \in R^{2 \times 0}$

Matrices de paramètres appris par rétropropagation  
Pour ces 3 matrices, le nombre de colonnes (3) doit être égale au nombre de lignes dans X (3)

209

209

### Auto attention

$X = \begin{pmatrix} x_{11} & x_{12} & x_{13} & x_{14} \\ x_{21} & x_{22} & x_{23} & x_{24} \\ x_{31} & x_{32} & x_{33} & x_{34} \end{pmatrix} \in R^{3 \times 4}$

$W^q = \begin{pmatrix} W^q_{11} & W^q_{12} & W^q_{13} \\ W^q_{21} & W^q_{22} & W^q_{23} \\ W^q_{31} & W^q_{32} & W^q_{33} \end{pmatrix} \in R^{3 \times 0}$

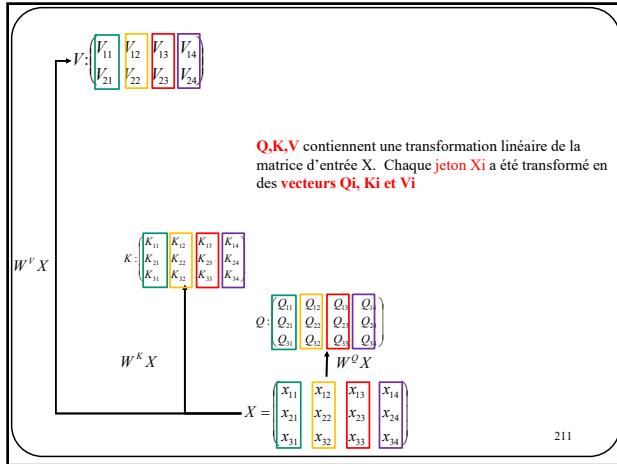
$W^k = \begin{pmatrix} W^k_{11} & W^k_{12} & W^k_{13} \\ W^k_{21} & W^k_{22} & W^k_{23} \\ W^k_{31} & W^k_{32} & W^k_{33} \end{pmatrix} \in R^{3 \times 0}$

$W^v = \begin{pmatrix} W^v_{11} & W^v_{12} & W^v_{13} \\ W^v_{21} & W^v_{22} & W^v_{23} \end{pmatrix} \in R^{2 \times 0}$

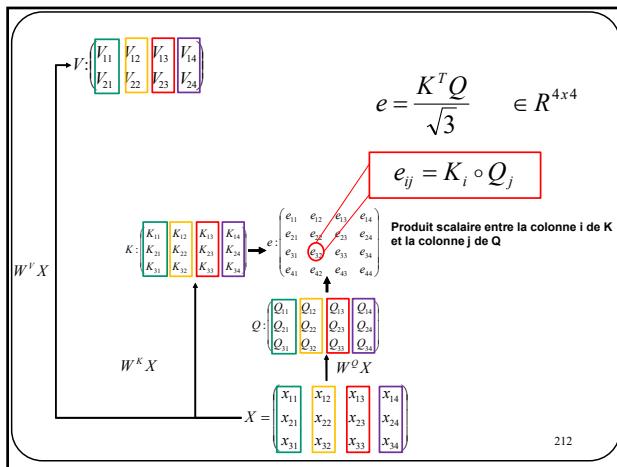
Matrices de paramètres appris par rétropropagation  
Pour ces 3 matrices, le nombre de ligne (3,3,2) est arbitraire

210

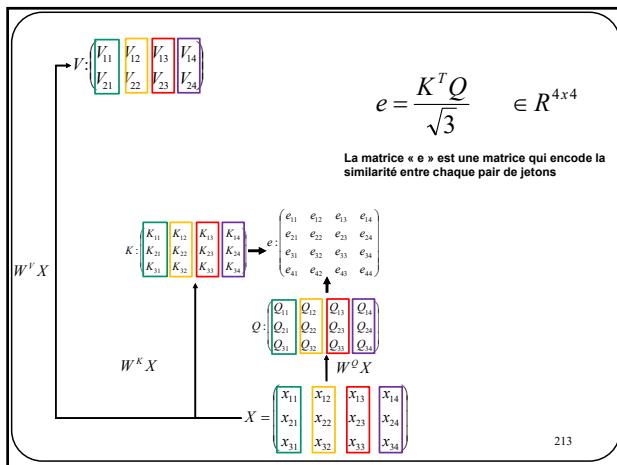
210



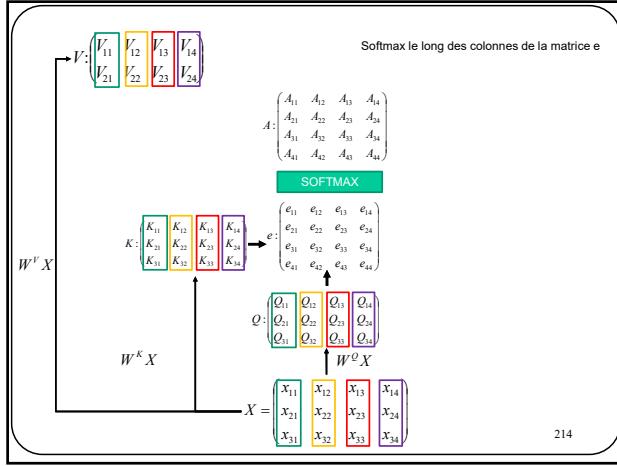
211



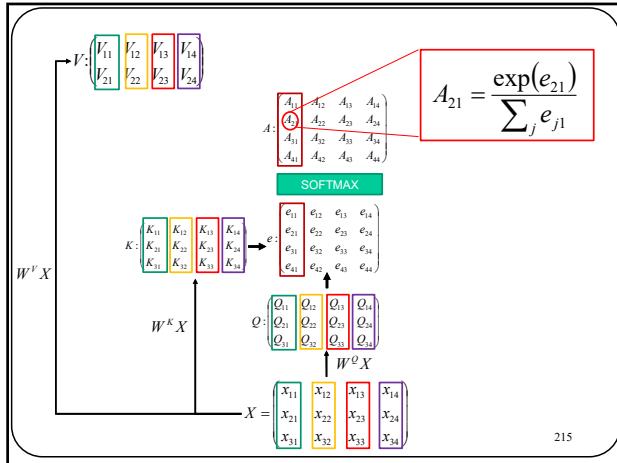
212



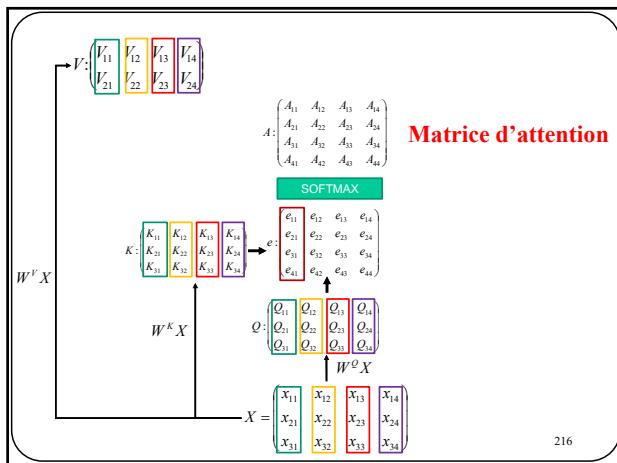
213



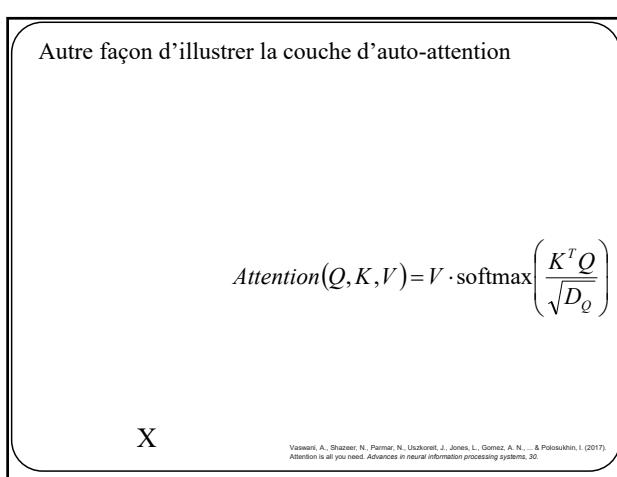
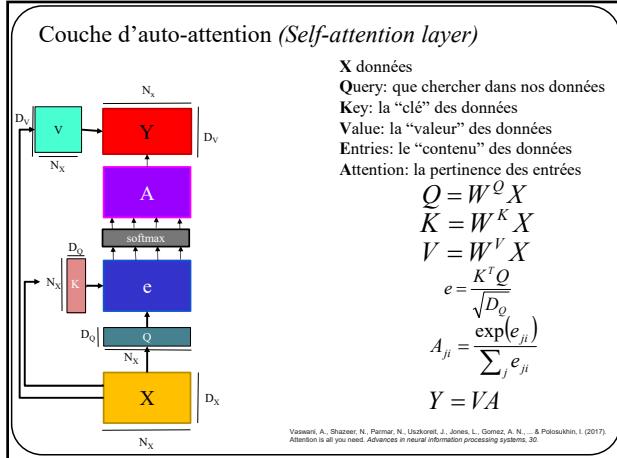
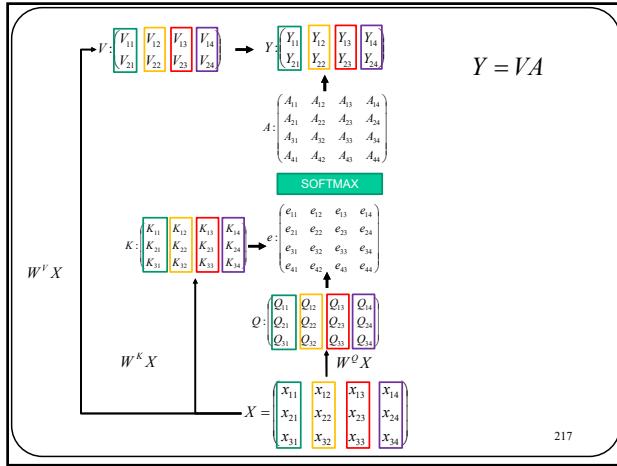
214



215

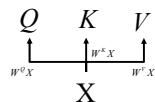


216



Autre façon d'illustrer la couche d'auto-attention

$$\text{Attention}(Q, K, V) = V \cdot \text{softmax} \left( \frac{K^T Q}{\sqrt{D_Q}} \right)$$

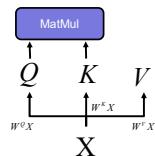


Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

220

Autre façon d'illustrer la couche d'auto-attention

$$\text{Attention}(Q, K, V) = V \cdot \text{softmax} \left( \frac{K^T Q}{\sqrt{D_Q}} \right)$$

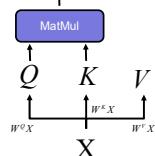


Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

221

Autre façon d'illustrer la couche d'auto-attention

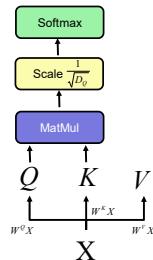
$$\text{Attention}(Q, K, V) = V \cdot \text{softmax} \left( \frac{K^T Q}{\sqrt{D_Q}} \right)$$



Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

222

Autre façon d'illustrer la couche d'auto-attention

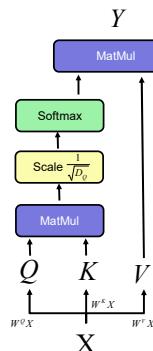


$$\text{Attention}(Q, K, V) = V \cdot \text{softmax} \left( \frac{K^T Q}{\sqrt{D_Q}} \right)$$

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

223

Autre façon d'illustrer la couche d'auto-attention

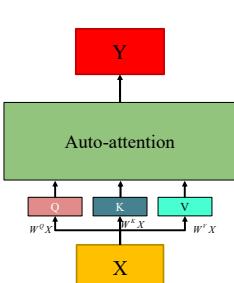


$$\text{Attention}(Q, K, V) = V \cdot \text{softmax} \left( \frac{K^T Q}{\sqrt{D_Q}} \right)$$

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

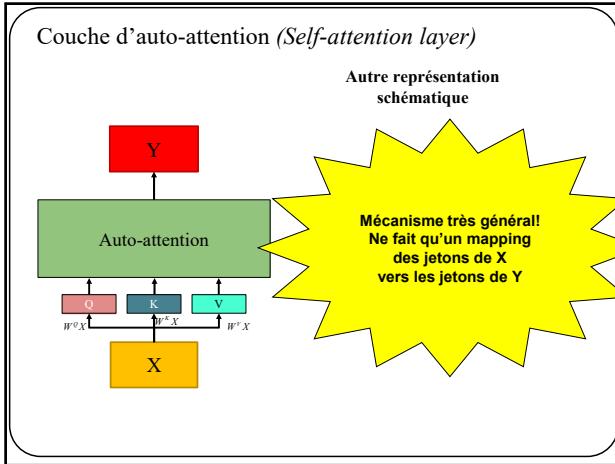
224

### Couche d'auto-attention (*Self-attention layer*)



### Autre représentation schématique

225



226

---

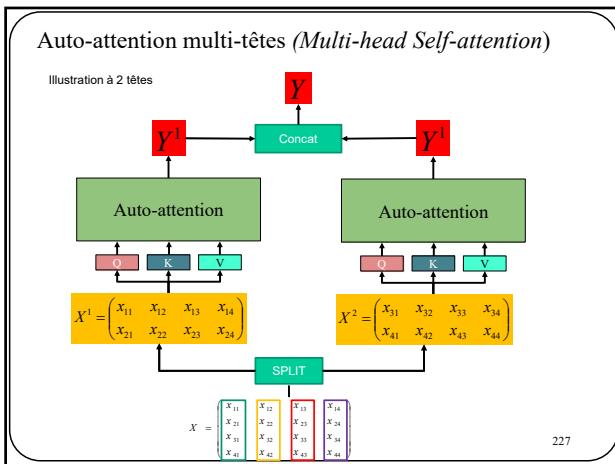
---

---

---

---

---



227

---

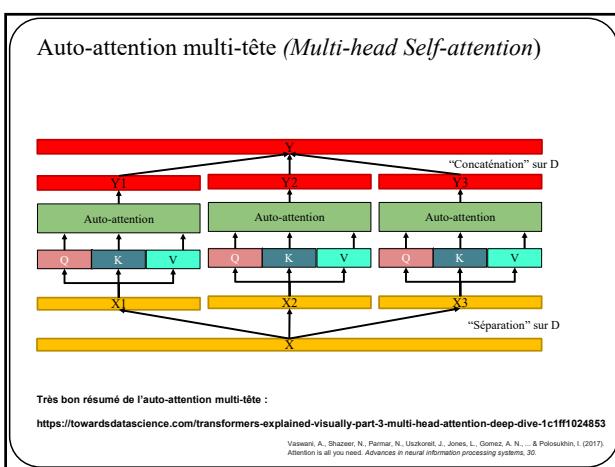
---

---

---

---

---



228

---

---

---

---

---

---

## L'apothéose des réseaux de neurones

## *Transformer*

*(Attention is all you need)*

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

229

# *Transformer*

*Implique aucune notion de récurrence*

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

230

## *Transformer (Attention is all you need)*

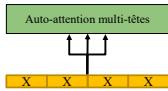
X X X X

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

231

*Transformer (Attention is all you need)*

- Auto-attention multi-têtes sur les dimensions de X

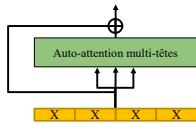


Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

232

*Transformer (Attention is all you need)*

- Auto-attention multi-têtes sur les dimensions de X
  - “+ “ = connexion résiduelle

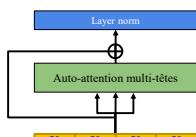


Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

233

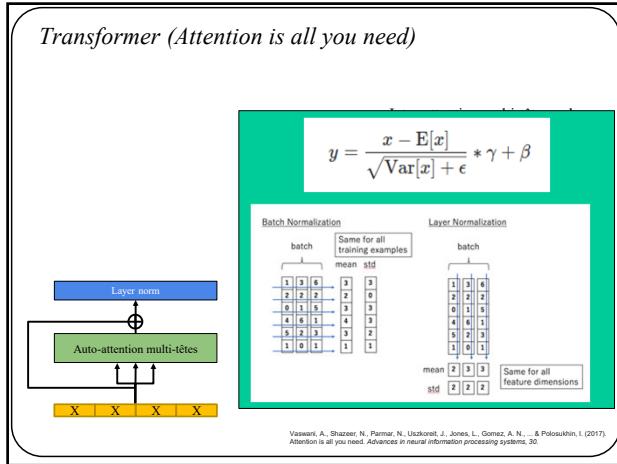
*Transformer (Attention is all you need)*

- Auto-attention multi-têtes sur les dimensions de X
  - “+ “ = connexion résiduelle
  - “Layer-norm”

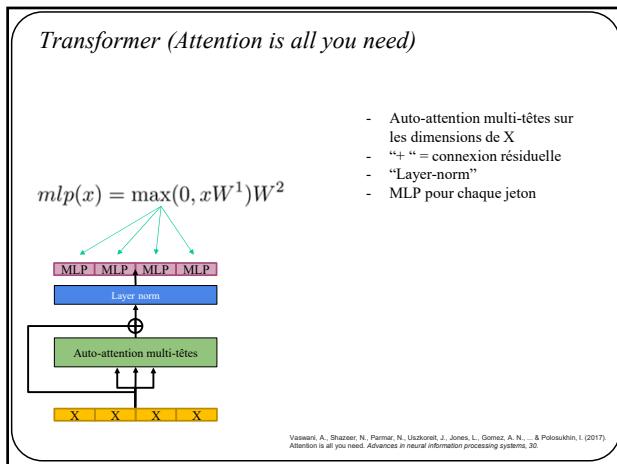


Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

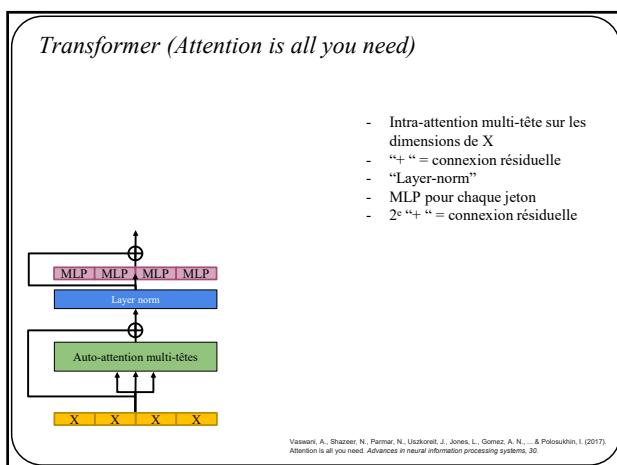
234



235

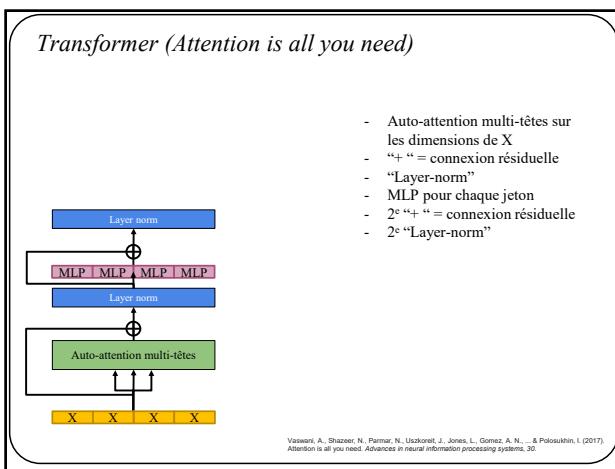


236



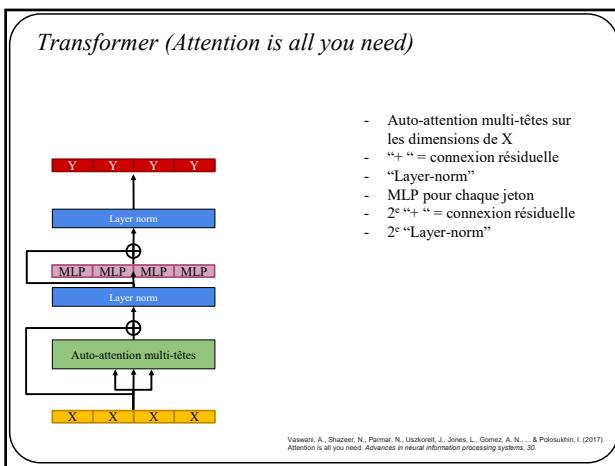
237

*Transformer (Attention is all you need)*



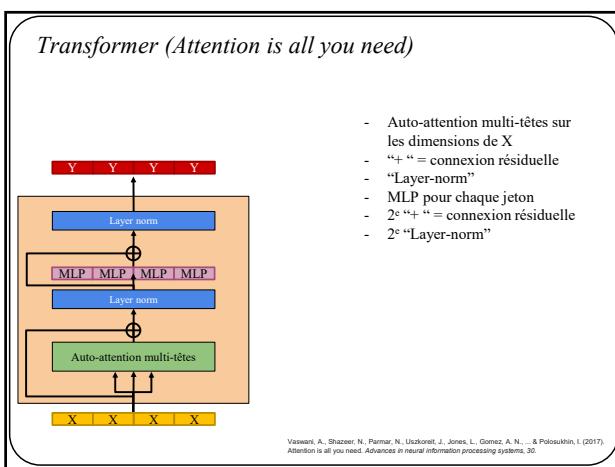
238

*Transformer (Attention is all you need)*

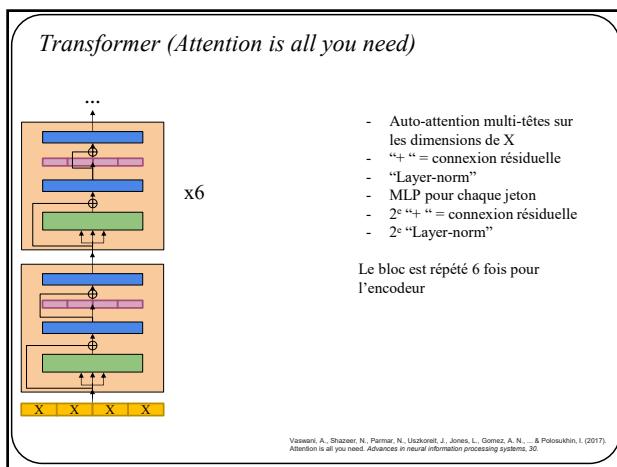


239

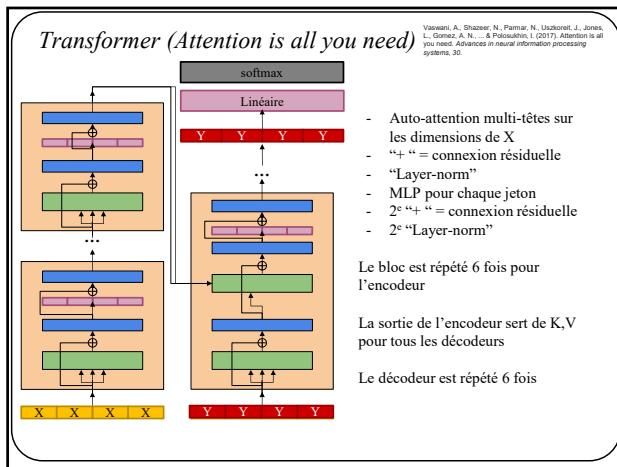
*Transformer (Attention is all you need)*



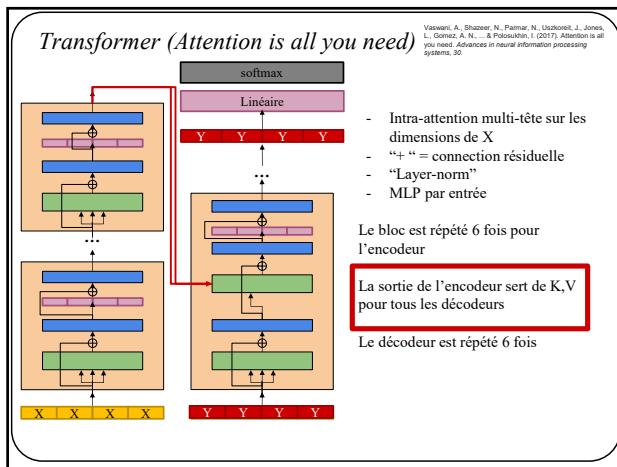
240



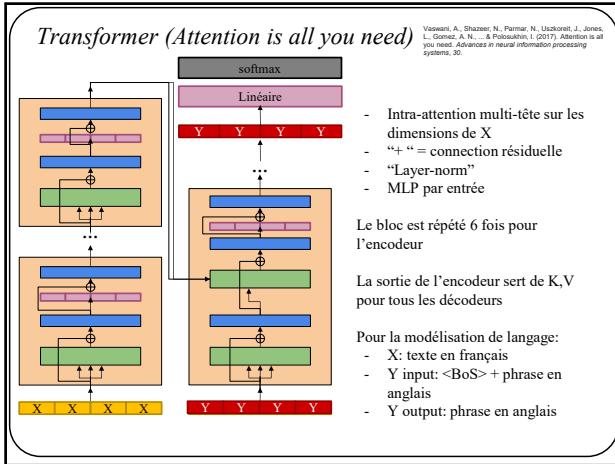
241



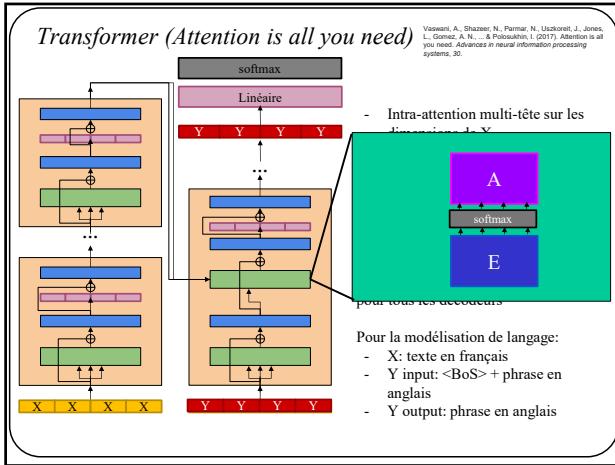
242



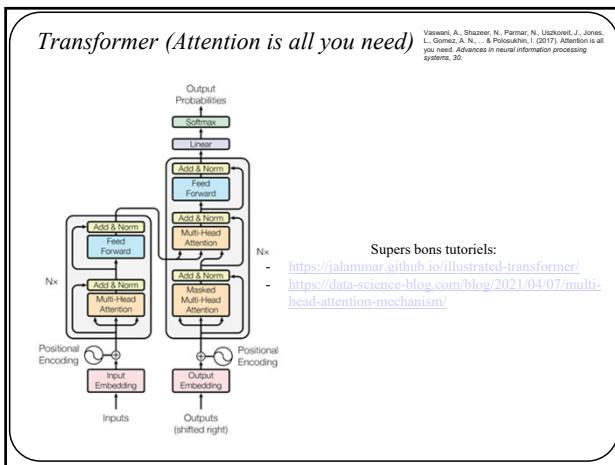
243



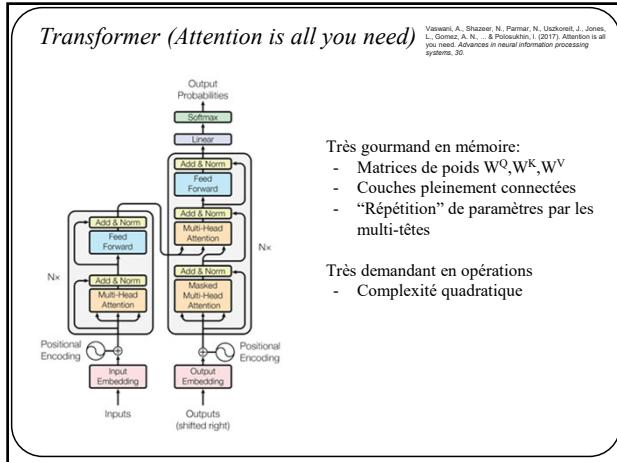
244



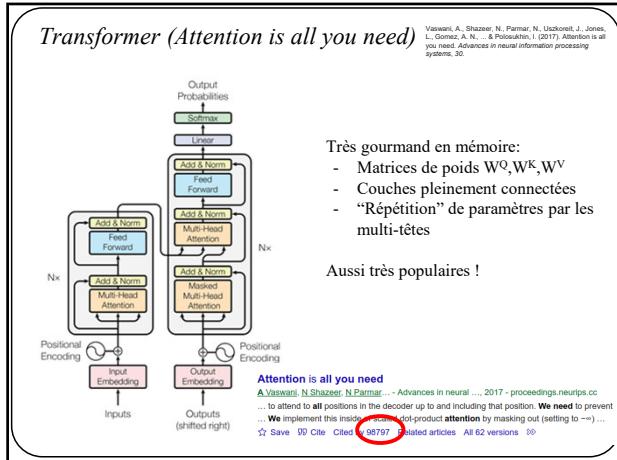
245



249



250



251

Différentes version de transformers

SxxuudshoUhvghw083=56P gh sdchp fwhvq

Model	Layers	Width	Heads	Params	Data	Training
Transformer-Base	12	512	8	65M	8x P100	(12 hours)
Transformer-Large	12	1024	16	213M	8x P100	(3.5 days)

fuhgleMcvwqf Mrkqvqg

252

*Transformers*

Model	Layers	Width	Heads	Params	Data	Training
Transformer-Base	12	512	8	65M		8x P100 (12 hours)
Transformer-Large	12	1024	16	213M		8x P100 (3.5 days)
BERT-Base	12	768	12	110M	13 GB	
BERT-Large	24	1024	16	340M	13 GB	

fig14-16cnn&Mrkqvrc

253

*Transformers*

Model	Layers	Width	Heads	Params	Data	Training
Transformer-Base	12	512	8	65M		8x P100 (12 hours)
Transformer-Large	12	1024	16	213M		8x P100 (3.5 days)
BERT-Base	12	768	12	110M	13 GB	
BERT-Large	24	1024	16	340M	13 GB	
XLNet-Large	24	1024	16	~340M	126 GB	512x TPU-v3 (2.5 days)
RoBERTa	24	1024	16	355M	160 GB	1024x V100 GPU (1 day)

fig14-16cnn&Mrkqvrc

254

*Transformers*

Model	Layers	Width	Heads	Params	Data	Training
Transformer-Base	12	512	8	65M		8x P100 (12 hours)
Transformer-Large	12	1024	16	213M		8x P100 (3.5 days)
BERT-Base	12	768	12	110M	13 GB	
BERT-Large	24	1024	16	340M	13 GB	
XLNet-Large	24	1024	16	~340M	126 GB	512x TPU-v3 (2.5 days)
RoBERTa	24	1024	16	355M	160 GB	1024x V100 GPU (1 day)
GPT-2	12	768	?	117M	40 GB	
GPT-2	24	1024	?	345M	40 GB	
GPT-2	36	1280	?	762M	40 GB	
GPT-2	48	1600	?	1.5B	40 GB	

fig14-16cnn&Mrkqvrc

255

*Transformers*

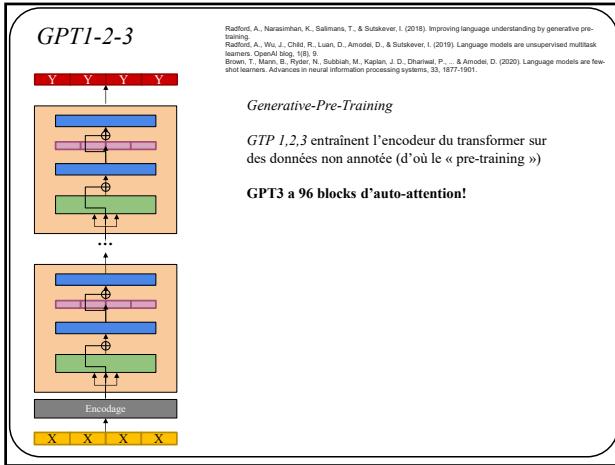
---

Model	Layers	Width	Heads	Params	Data	Training
Transformer-Base	12	512	8	65M		8x P100 (12 hours)
Transformer-Large	12	1024	16	213M		8x P100 (3.5 days)
BERT-Base	12	768	12	110M	13 GB	
BERT-Large	24	1024	16	340M	13 GB	
XLNet-Large	24	1024	16	~940M	126 GB	512x TPU-v3 (2.5 days)
RoBERTa	24	1024	16	355M	160 GB	1024x V100 GPU (1 day)
GPT-2	12	768	?	117M	40 GB	
GPT-2	24	1024	?	345M	40 GB	
GPT-2	36	1280	?	762M	40 GB	
GPT-2	48	1600	?	1.5B	40 GB	
Megatron-LM	40	1536	16	1.2B	174 GB	64x V100 GPU
Megatron-LM	54	1920	20	2.5B	174 GB	128x V100 GPU
Megatron-LM	64	2304	24	4.2B	174 GB	256x V100 GPU (10 days)
Megatron-LM	72	3072	32	8.3B	174 GB	512x V100 GPU (9 days)

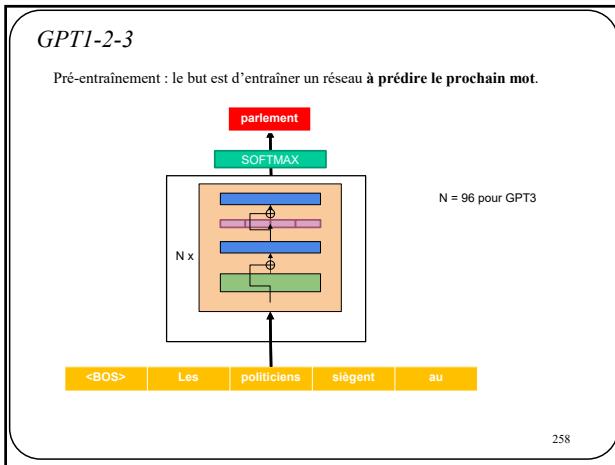
fusigle-Mecvñq Mrkvqvc

---

256



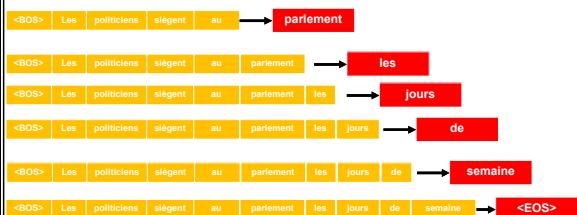
257



258

GPT1-2-3

Et on recommence:



259

259

Transformers

Les transformers sont vraiment bons pour générer du texte

**SYSTEM PROMPT (HUMAN-WRITTEN)**

*Recycling is good for the world.*

*NO! YOU COULD NOT BE MORE WRONG!!*

**MODEL COMPLETION (MACHINE-WRITTEN, 25 TRIES)**

Recycling is NOT good for the world. It is bad for the environment, it is bad for our health, and it is bad for our economy. I'm not kidding. Recycling is not good for the environment. It is destructive to the earth and it is a major contributor to global warming. Recycling is not good for our health. It contributes to obesity and disease like heart disease and cancer. Recycling is bad for our economy. It increases the cost of a product. The cost of recycling is even higher than the cost of producing a new product. This is because companies are willing to pay for the privilege of having the waste they must collect and effect a recycling program. Recycling is a huge, colossal waste of time, energy, money, and resources. And THAT is why we need to get back to basics and get back to basics in our recycling efforts. One of the best ways to start is to look at the process of creating a paper product. When you make a paper product, it is basically a long chain of materials. Everything from the raw materials (wood, cardboard, paper, etc.) to the reagents (dyes, solvents, etc.) to the printing equipment (chemicals, glue, paper, ink, etc.), to the packaging, to the packaging materials (myrcia, chemicals, etc.), to the processing equipment (heating, cooling, etc.), to the packaging materials, to the packaging materials that are shipped overseas and to the packaging materials that are used in the United States. Each step along the way creates tons of waste that we constantly have to clean up. The process of making a paper product is a very wasteful battle. But the end result is something that all of us need to consume. And if we want to keep the recycling process running efficiently, then we really need to think about each and every step that goes into making a paper product.

<https://openai.com/blog/better-language-models/#sample8>

260

Transformers

Les transformers sont vraiment bons pour générer du texte sémantiquement viable

I really like the IFT780 course on neural networks offered at the Université de Sherbrooke, because

the teacher uses examples from video games, which I find a nice contrast to the theory.

I think I would like a course on AI to really understand how it works, but when we are in class, the prof speaks in a rapid French, and I can't really process what he's saying.

Having started a course on it at a different time, I understand more this time, but I still don't know where I stand.

My "new" plan for next year is to do all of the tutorials and classes I've been wanting to take, but I'm still stuck on one big heavy, headachey decision: how do I want to spend my post - masters year?

With a lot of math and statistics classes, it's not a hard choice.

I'm on the fence with what I want to spend my summer doing: a lot of analysis and research work or do some volunteering?

I've applied to several volunteer programs: naturis and Vert directeur de la santé, which are essentially summer internships for dentists and health care professionals to do a bit of volunteer work.

<https://app.inferkit.com/demo>

261

**GPT-1-2-3**

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. OpenAI blog, 1(8), 9.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. Advances in neural information processing systems, 33, 1877-1901.

	Parameters	Layers	$d_{model}$
GPT-2	117M 345M 762M 1542M	12 24 36 48	768 1024 1280 1600

Table 2. Architecture hyperparameters for the 4 model sizes.

Model Name	$n_{tokens}$	$n_{layers}$	$d_{model}$	$n_{heads}$	$d_{attn}$	Batch Size	Learning Rate
GPT-3 Small	1.25B	12	64	4	64	1024	$3.0 \times 10^{-4}$
GPT-3 Medium	3.50B	24	1024	16	64	0.5M	$3.0 \times 10^{-4}$
GPT-3 Large	7.60B	24	1536	16	96	0.5M	$2.5 \times 10^{-4}$
GPT-3 XL	1.30B	24	2048	24	128	1M	$2.0 \times 10^{-4}$
GPT-3.2-TB	2.7B	32	2560	32	192	1M	$1.6 \times 10^{-4}$
GPT-3.6-TB	6.7B	32	4096	32	128	2M	$1.2 \times 10^{-4}$
GPT-3.13B	13.0B	40	5140	40	128	2M	$1.0 \times 10^{-4}$
GPT-3.175B or GPT-3 <sup>*</sup>	175.0B	96	12288	96	128	3.2M	$0.6 \times 10^{-4}$

Table 3.1: Sizes, architectures, and learning hyper parameters (batch size in tokens and learning rate) of the models which we trained. All models were trained for a total of 300 billion tokens.

262

**GPT-1-2-3**

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. OpenAI blog, 1(8), 9.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. Advances in neural information processing systems, 33, 1877-1901.

Figure 1. Trend of state-of-the-art NLP model sizes with time.

<https://developer.nvidia.com/blog/scaling-language-model-training-to-a-trillion-parameters-using-megatron/>

\*356 years on a V100 GPU server with 28 TFLOPS capacity and would cost \$4.6 million at \$1.5 per hour\*

<https://bdtechtalks.com/2020/09/21/gpt-3-economy-business-model/>

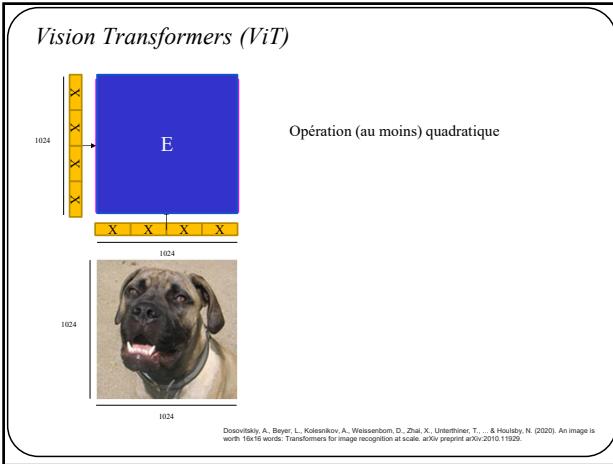
263

**Vision Transformers (ViT)**

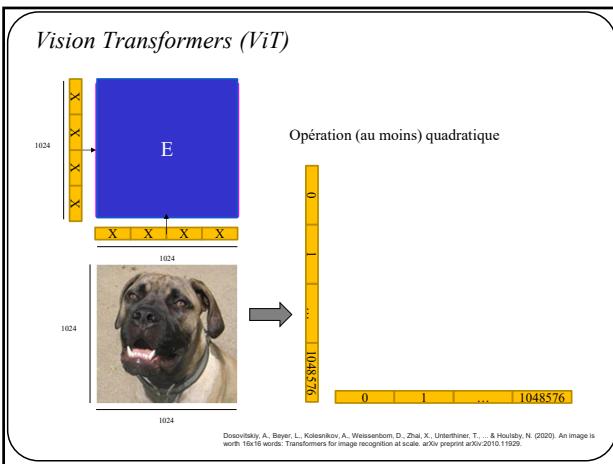
Opération (au moins) quadratique = 1024<sup>2</sup> opérations et composantes à garder en mémoire

Dosovitsky, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.

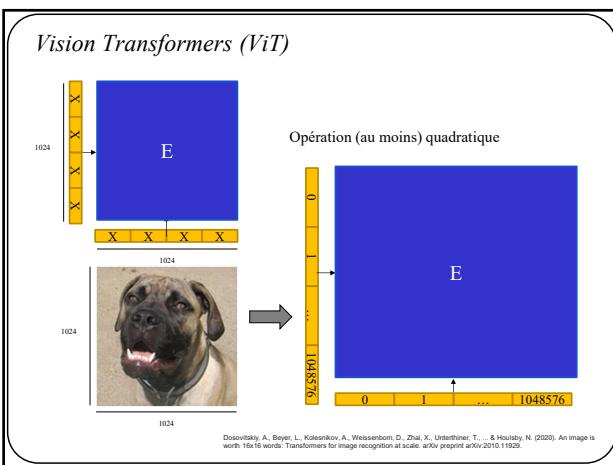
264



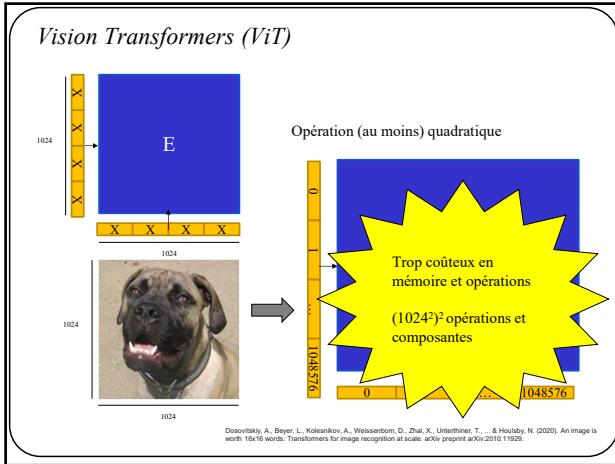
265



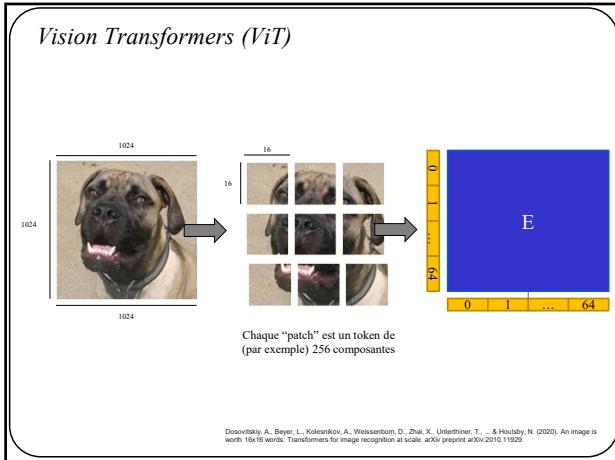
266



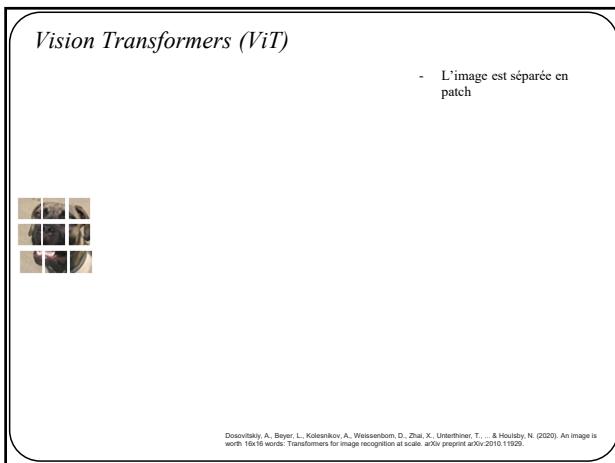
267



268



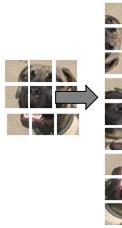
269



270

Vision Transformers (ViT)

- L'image est séparée en patch
  - Chaque patch est linéarisée

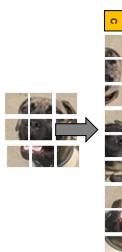


Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.

271

Vision Transformers (ViT)

- L'image est séparée en patch
  - Chaque patch est linéarisée
  - Un token spécial représentant la classe est ajouté

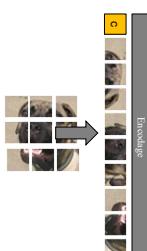


Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.

272

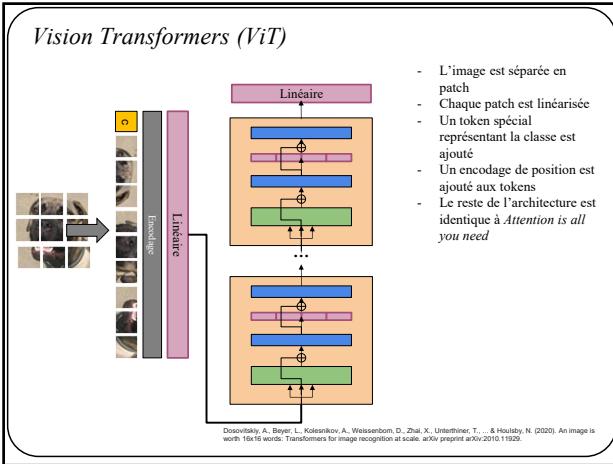
Vision Transformers (ViT)

- L'image est séparée en patch
  - Chaque patch est linéarisée
  - Un token spécial représentant la classe est ajouté
  - Un encodage de position est ajouté aux tokens

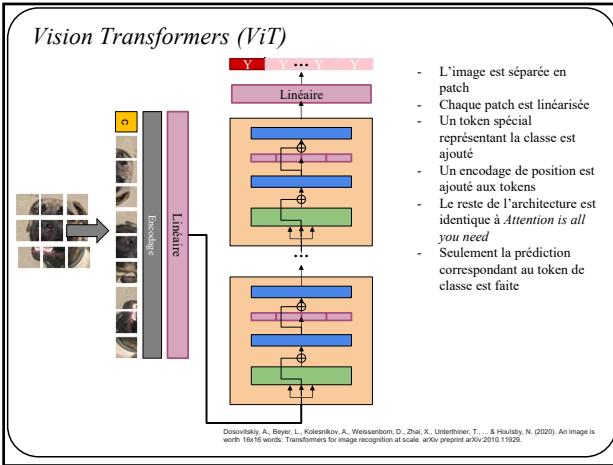


Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.

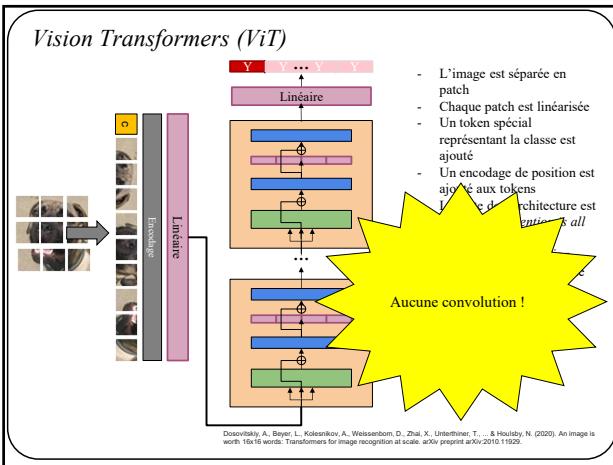
273



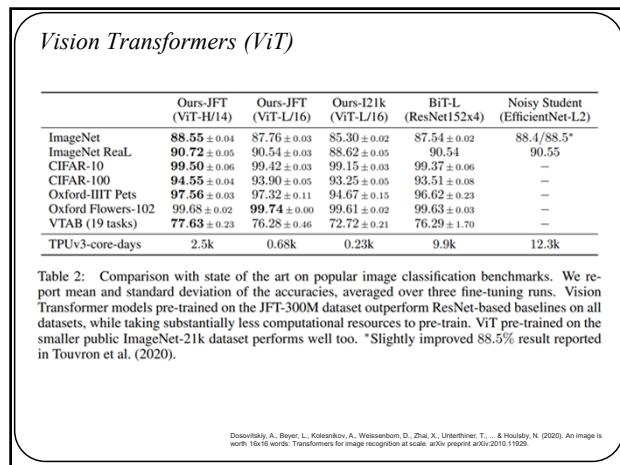
274



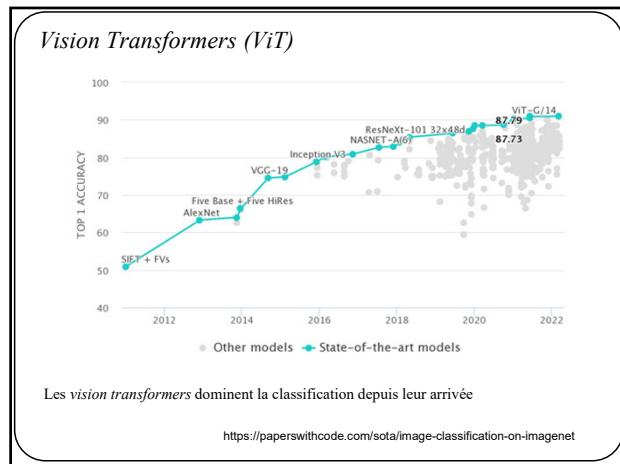
275



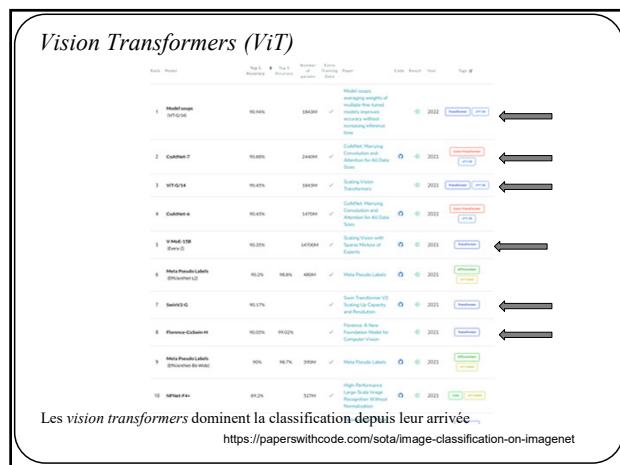
276



277



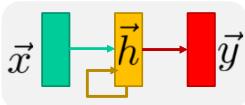
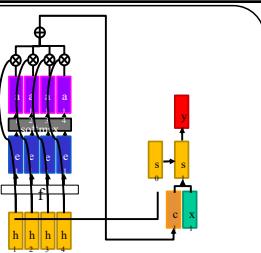
278



279



**Sommaire**

- Les réseaux récurrents peuvent traiter des séquences
- Ils ne requièrent que de légères modifications à des réseaux pleinement connectés
- Ils sont instables sur de longues séquences
- LSTM/GRU sont utilisés en pratique
- L'attention est un mécanisme très puissant permettant aux réseaux d'apprendre quelle partie des données utilisées pour faire une prédiction
- L'attention n'est pas limitée au texte, ou même aux séquences

283

---



---



---



---



---



---



---

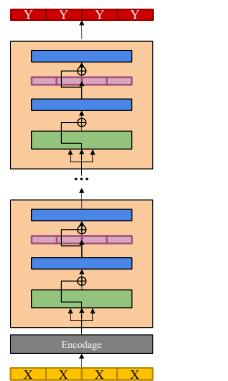


---



---

**Sommaire**



- Un *Transformer* sont un modèle puissant pour les tâches liées au langage naturel et aux images
- Les *transformers* n'utilisent *que* l'attention (pas un modèle récurrent)
- Les *transformers* sont demandant en ressources

284

---



---



---



---



---



---



---



---



---