

# 9. Sequential Neural Models


CS 519 Deep Learning, Winter 2018

Fuxin Li

*With materials from Andrej Karpathy, Bo Xie, Zolt Kira*

# Sequential and Temporal Data

- Many applications exhibited by dynamically changing states
  - Language (e.g. sentences)
  - Temporal data
    - Speech
    - Stock Market



this deep learning course is great!  
this deep learning course is great!  
this deep learning course is great!



```
AAGTCAAGCTGCTCTGTGGGCTGTGATCTGCCTCAAACCCACAGCCTGGGTAGCAGG
AGGACCTTGATGCTCCTGGCACAGATGAGGAGAATCTCTTTTCTCCTGCTGAAG
GACAGACATGACTTTGGATTTCCCCAGGAGGAGTTTGGCAACCAGTCCAAAAGGCT
GAAACCATCCCTGTCTCCATGAGATGATCCAGCAGATCTCAATCTTTCAGACACA
AAGGACTCATCTGCTGCTTGGGATGAGACCCCTCCTAGACAAATTTACACTGAATC
TACCAGCAGCTGAATGACCTGGAAGCCTGTGTGATACAGGGGTGGGGGTGACAGAG
ACTCCCTGATGAAGGAGGACTCCATTCTGGCTGTGAGGAAATACTTCCAAAGAATC
ACTCTCTATCTGAAAGAGAAGAAATACAGCCCTTGTGCTGGGAGGTTGTGAGAGCA
GAAATCATGAGATCTTTTCTTTGTCAACAACTTGCAAGAAGTTTAAGAAGTAAG
GAATGA, TGTGATCTGCCTCAAACCCACAGCCTGGGTAGCAGGAGGACCTTGTGC
TCCTGGCACAGATGAGGAGAATCTCTCTTTTCTCCTGCTTGAAGGACAGACATGAT
TTGGATTTCCCCAGGAGGAGTTTGGCAACCAGTCCAAAAGGCTGAAACCATCCCTG
TCCTCCATGAGATGATCCAGCAGATCTTCAATCTCTTTCAGACAAAGGACTCATCTG
CTGCTTGGGATGAGACCCCTCCTAGACAAATTTACACTGAATCTTACCAGCAGCTGA
ATGACCTGGAAGCCTGTGTGATACAGGGGTGGGGGTGACAGAGACTCCCTGATGA
AGGAGGACTTCCATTCTGCTGTGAGGAAATACITCCAAAGAATCACTCTCTATCTGA
AAGAGAAGAAATACAGCCCTTGTGCTGGGAGGTTGTGAGAGCAGAAATCATGAGAT
CTTTTCTTTGTCAACAACTTGAAGAAGTTTAAGAAGTAAGGAATGA and
```

# Image Captioning



"man in black shirt is playing guitar."



"construction worker in orange safety vest is working on road."



"two young girls are playing with lego toy."



"boy is doing backflip on wakeboard."



"girl in pink dress is jumping in air."



"black and white dog jumps over bar."



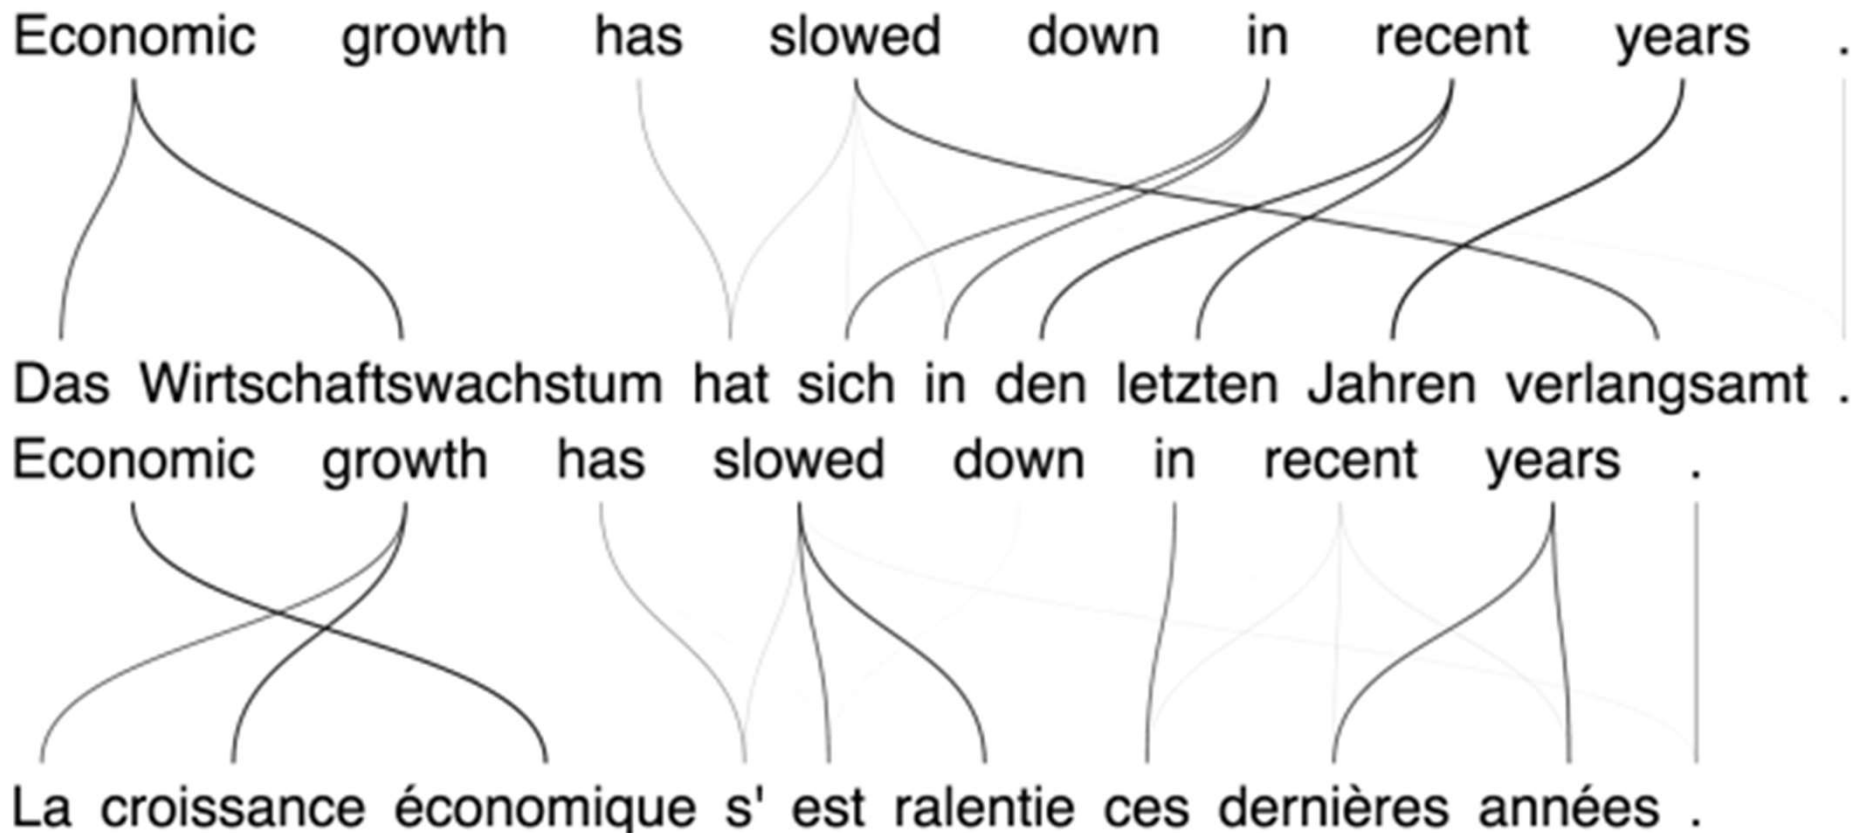
"young girl in pink shirt is swinging on swing."



"man in blue wetsuit is surfing on wave."

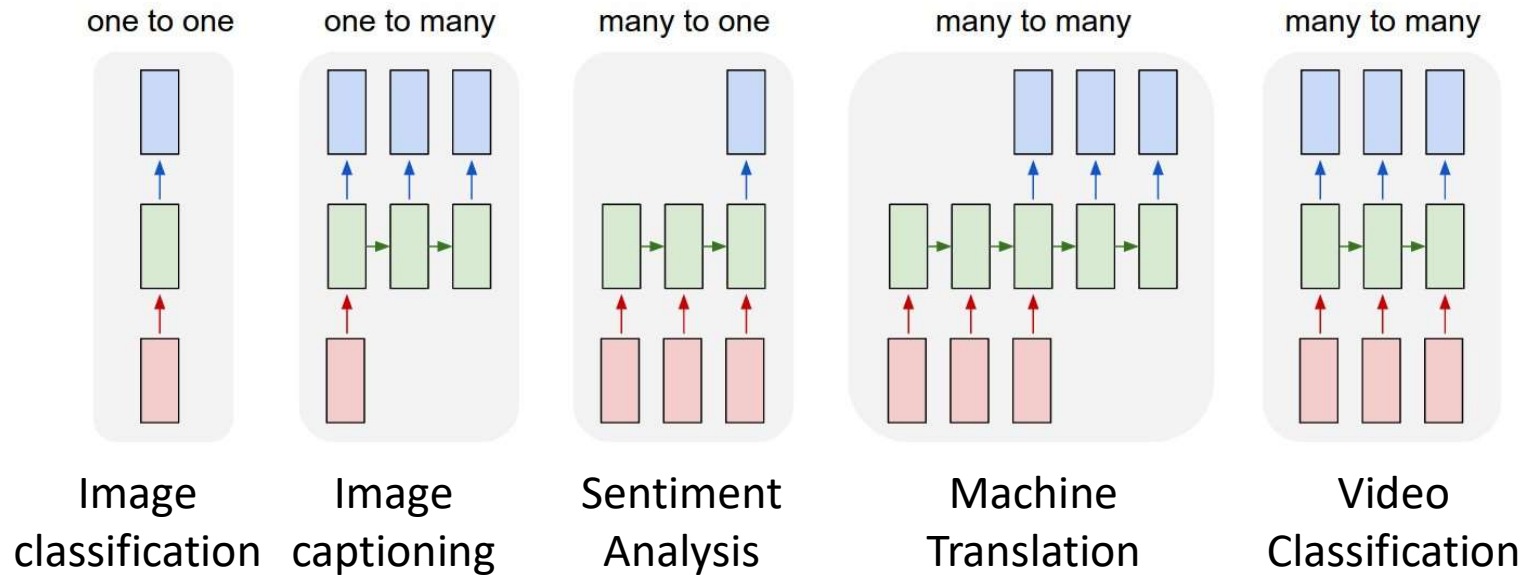
# Machine Translation

- Have to look at the entire sentence (or, many sentences)



# Sequence Data


- Many data are sequences and have different inputs/outputs



(cf. Andrej Karpathy blog)


# Previous: Autoregressive Models

- Autoregressive models
  - Predict the next term in a sequence from a fixed number of previous terms using “delay taps”.
- Neural Autoregressive models
  - Use neural net to do so



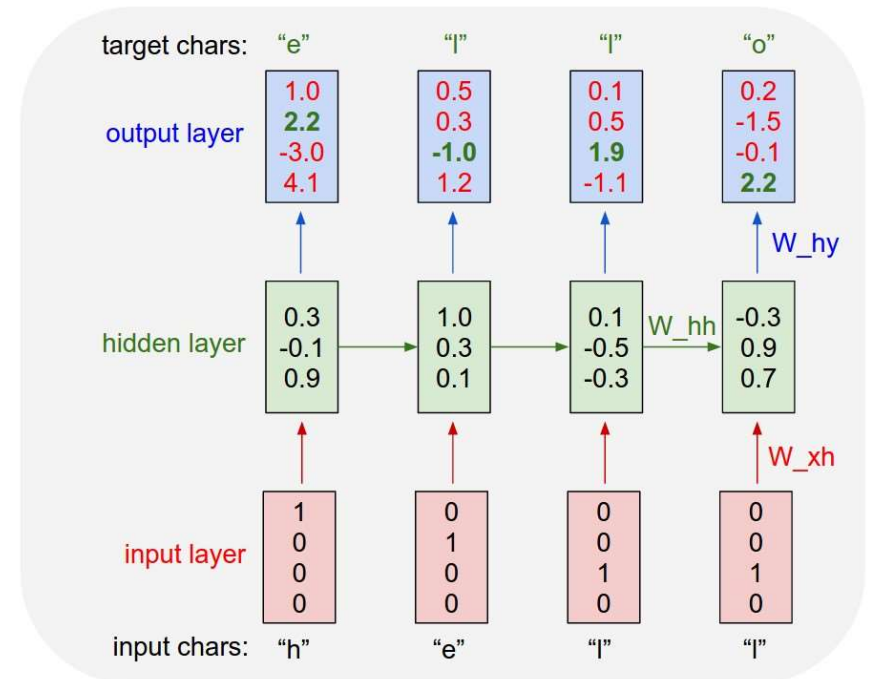
# Previous: Hidden Markov Models

- Hidden states
- Outputs are generated from hidden states
  - Does not accept additional inputs
  - Discrete state-space
    - Need to learn all discrete transition probabilities!



# Recurrent Neural Networks


- Similar to
  - Linear Dynamic Systems
    - E.g. Kalman filters
  - Hidden Markov Models
  - But not generative
- “Turing-complete”



(cf. Andrej Karpathy blog)



# Vanilla RNN Flow Graph



$$\mathbf{a}_t = \mathbf{b} + \mathbf{W}\mathbf{h}_{t-1} + \mathbf{U}\mathbf{x}_t$$

$$\mathbf{h}_t = \tanh(\mathbf{a}_t)$$

$$\mathbf{y}_t = \mathbf{c} + \mathbf{V}\mathbf{h}_t$$

U – input to hidden

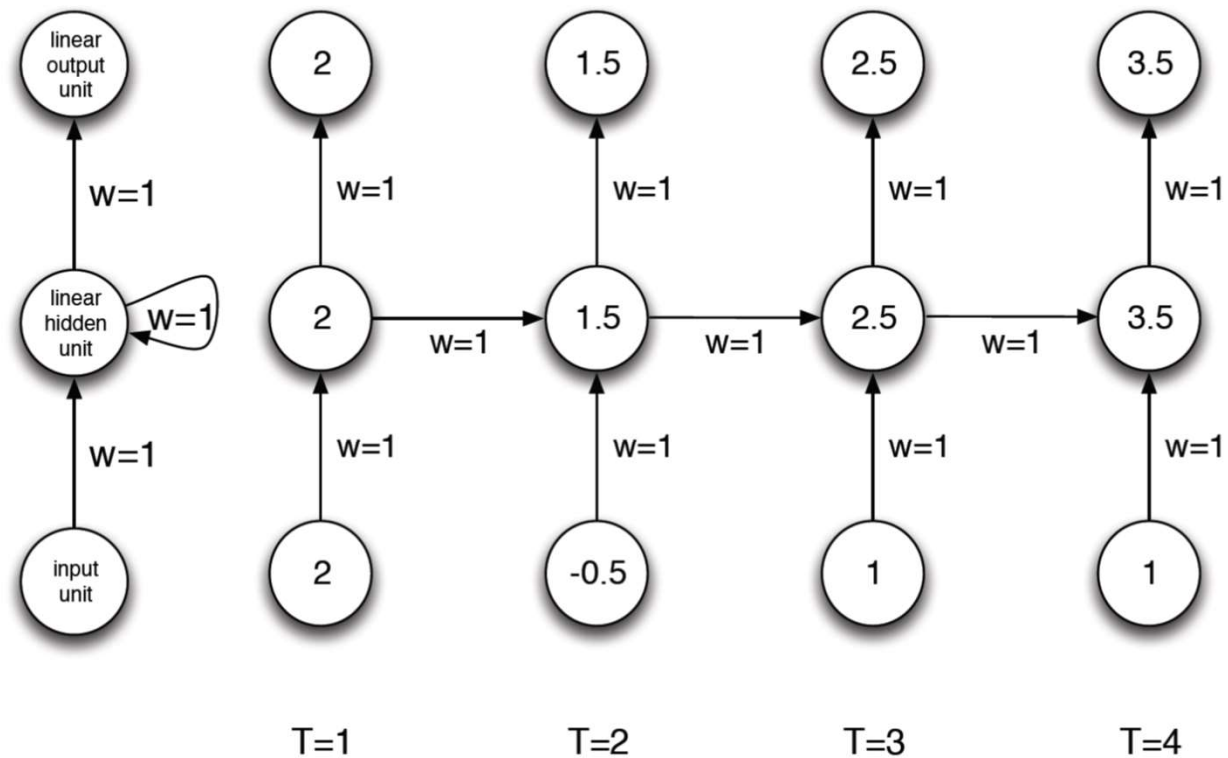
V – hidden to output

W – hidden to hidden

# Examples

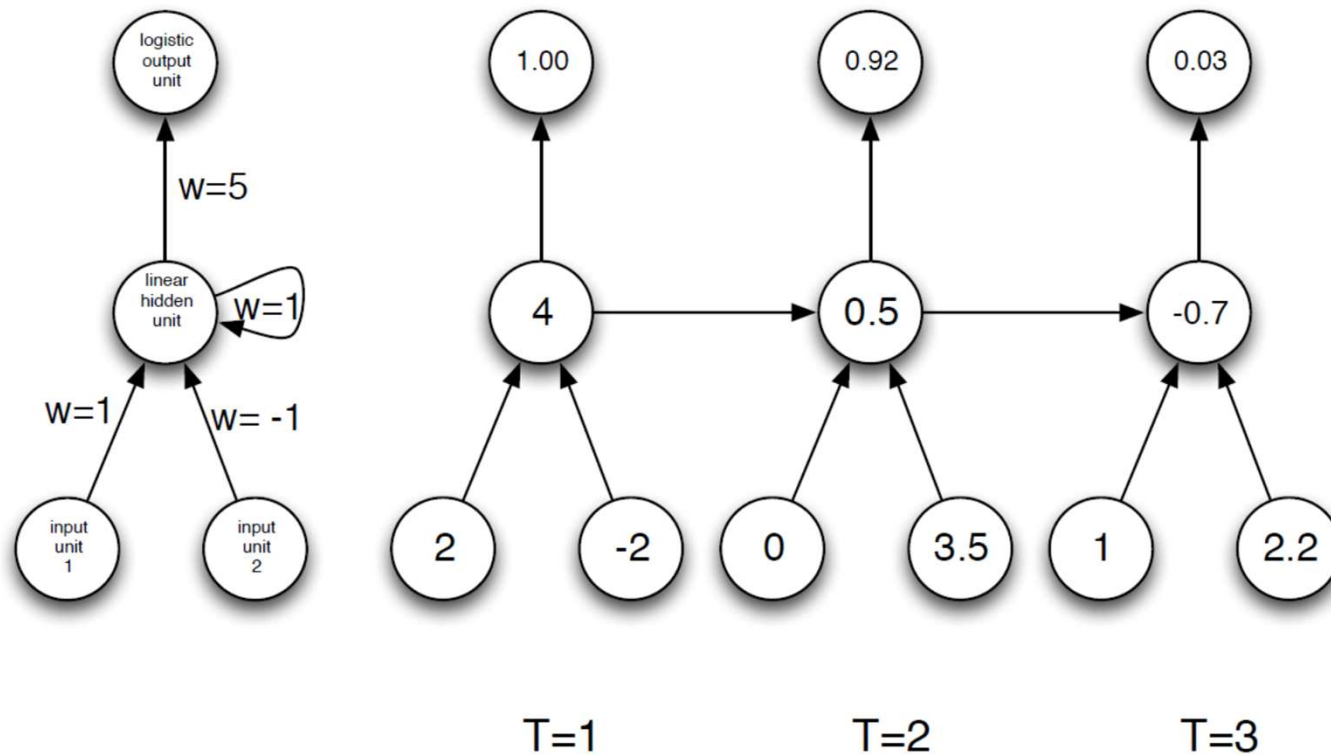
Now let's look at some simple examples of RNNs.

This one sums its inputs:




# Examples

This one determines if the total values of the first or second input are larger:



# Finite State Machines

- Each node denotes a state
- Reads input symbols one at a time
- After reading, transition to some other state
  - e.g. DFA, NFA
- States = hidden units




# The parity Example

Assume we have a sequence of binary inputs. We'll consider how to determine the **parity**, i.e. whether the number of 1's is even or odd.

We can compute parity incrementally by keeping track of the parity of the input so far:

Parity bits: 0 1 1 0 1 1  $\longrightarrow$   
Input: 0 1 0 1 1 0 1 0 1 1


Each parity bit is the XOR of the input and the previous parity bit.



# RNN Parity

- At each time step, compute parity between input vs. previous parity bit

$y^{(t-1)}$	$x^{(t)}$	$h_1^{(t)}$	$h_2^{(t)}$	$y^{(t)}$
0	0	0	0	0
0	1	0	1	1
1	0	0	1	1
1	1	1	1	0



# RNN Universality

- RNN can simulate any finite state machines
  - is Turing complete with infinite hidden nodes (Siegelmann and Sontag, 1995)
  - e.g., a computer (Zaremba and Sutskever 2014)

Training data:

```
Input:  
j=8584  
for x in range(8):  
    j+=920  
b=(1500+j)  
print((b+7567))  
Target: 25011.
```

```
Input:  
i=8827  
c=(i-5347)  
print((c+8704) if 2641<8500 else 5308)  
Target: 12184.
```

# RNN Universality

- Testing programs

**Input:**

```
d=8640;  
print((7135 if 6710>((d+7080)*14) else 7200)).
```

<b>Target:</b>	7200.
<b>"Baseline" prediction:</b>	7200.
<b>"Naive" prediction:</b>	7200.
<b>"Mix" prediction:</b>	7200.
<b>"Combined" prediction:</b>	7200.

**Input:**

```
print((((841 if 2076<7326 else 1869)*10) if 7827<317 else 7192)).
```

<b>Target:</b>	7192.
<b>"Baseline" prediction:</b>	7192.
<b>"Naive" prediction:</b>	7192.
<b>"Mix" prediction:</b>	7192.
<b>"Combined" prediction:</b>	7192.



# RNN Universality (if only you can train it!)

**Input:**

```
print((4*7054)).
```

<b>Target:</b>	28216.
<b>"Baseline" prediction:</b>	28216.
<b>"Naive" prediction:</b>	28116.
<b>"Mix" prediction:</b>	28216.
<b>"Combined" prediction:</b>	28216.

**Input:**

```
print((4635-5257)).
```

<b>Target:</b>	-622.
<b>"Baseline" prediction:</b>	-688.
<b>"Naive" prediction:</b>	-628.
<b>"Mix" prediction:</b>	-692.
<b>"Combined" prediction:</b>	-632.

**Input:**

```
e=1079  
for x in range(10):e+=4729  
print(e).
```

<b>Target:</b>	48369.
<b>"Baseline" prediction:</b>	48017.
<b>"Naive" prediction:</b>	48011.
<b>"Mix" prediction:</b>	48101.
<b>"Combined" prediction:</b>	48009.

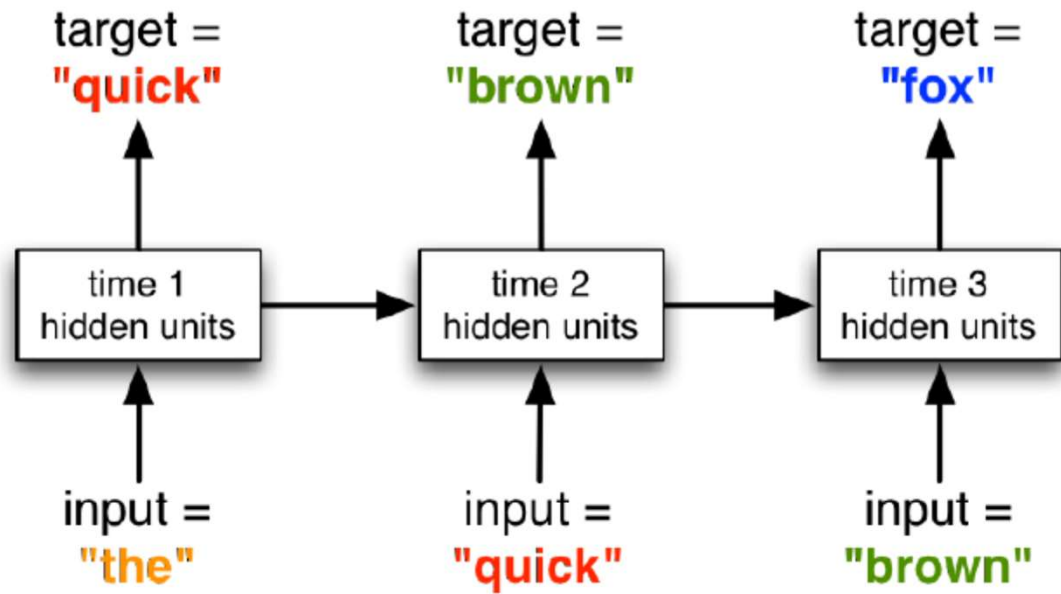
**Input:**

```
print((8*(5051-648))).
```

<b>Target:</b>	35224.
<b>"Baseline" prediction:</b>	34044.
<b>"Naive" prediction:</b>	32180.
<b>"Mix" prediction:</b>	33284.
<b>"Combined" prediction:</b>	33004.


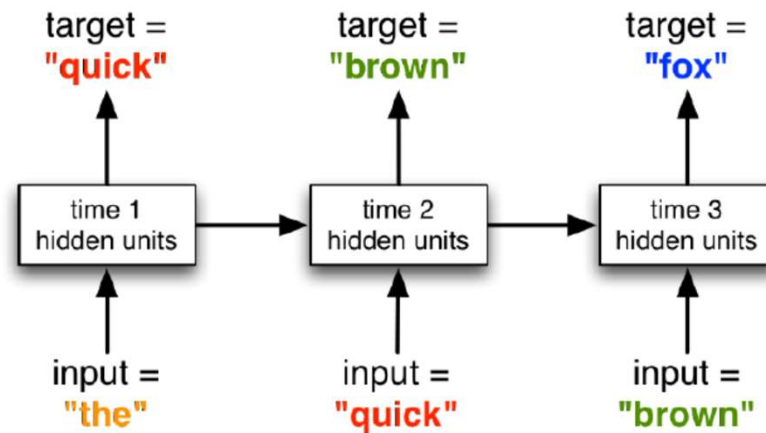
# RNN Text Model

One way to use RNNs to model text:



# Generate Text from RNN

One way to use RNNs to model text:



# RNN Sentence Model

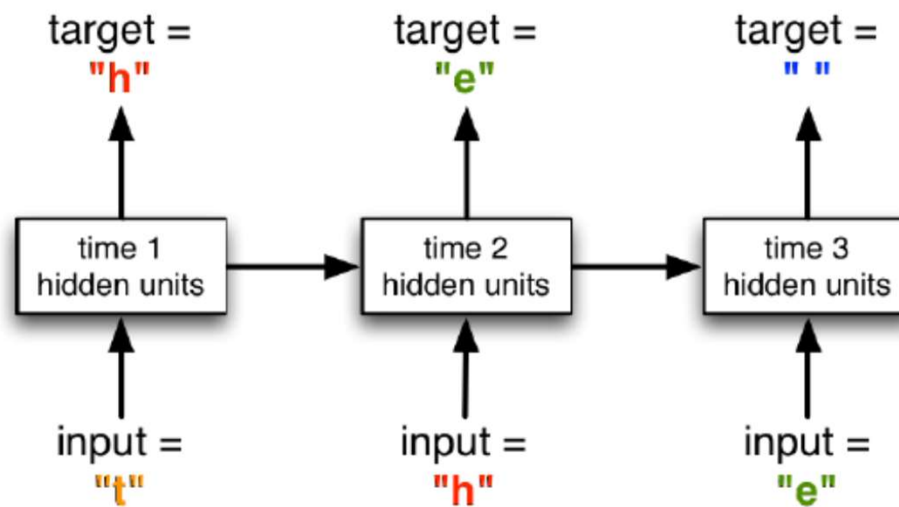
- Hypothetical: Different hidden units for:
  - Subject
  - Verb
  - Object (different type)

# Realistic Ones

Query	cell 25	cell 26	cell 27	cell 30
al yo yo sauce				
atkins diet lasagna				diet
blender recipes		recipes		
cake bakery edinburgh		bakery		
canning corn beef hash		corn, beef		
torre de pizza		pizza		
famous desserts				
fried chicken		chicken		
smoked turkey recipes		recipes		
italian sausage hoagies		sausage		
do you get allergy				
much pain will after total knee replacement				
how to make whiter teeth		whiter		
illini community hospital				hospital
implant infection	infection			
introductory psychology				
narcotics during pregnancy side effects				
fight sinus infections	infections			
health insurance high blood pressure				insurance,high
all antidepressant medications	antidepressant			medications

# RNN Character Model

Another approach is to model text *one character at a time!*



# Realistic Wiki Hidden Unit

t	f	p	:	/	/	w	w	.	y	n	e	t	n	e	w	s	.	c	o	m	/	]	E	n	g	l	i	s	h	-	l	a	n	g	u	a	g	e	w	e	b	s	i	t	e	o	f	I	s	r	a	e	l	'	s	l	a	r			
t	p	:	/	/	w	w	.	b	a	c	a	h	e	t	s	.	c	o	m	/			-	x	g	l	i	s	h	l	i	n	g	u	a	g	e	s	a	i	r	s	i	t	e	o	f	t	s	l	a	e	l	i	s	s	i	n	g		
d	:	x	n	e	.	w	a	e	a	.	a	w	a	t	o	a	.	s	&	n	t	i	a	c	a	-	s	a	r	d	e	e	l	h	o	a	n	t	b	i	s	a	n	f	a	n	r	e	i	f	'	a	a	t	d						
m	w	-	2		◆	p	i	i	s	o	e	s	s	i	s	.	/	e	r	n	.	c	]	(	d	c	e	e	n	e	p	e	s	a	a	i	k	i	i	e	e	l	e	d	h	,	i	r	t	h	r	a	o	n	s	e	,	c	o	s	e
d	r	.	<	:	a	h	b	-	n	p	t	w	t	.	x	i	g	h	/	m	a	)	T	v	d	r	y	z	i	c	o	u	e	d	I	s	u	:	t	h	a	-	o	o	t	u	,	s	t	u	i	f	l	v	e	p	e	r	y		
s	t	p	,	t	c	o	a	2	d	r	u	l	w	o	c	l	e	n	s	r	]	p	.	l	l	v	a	o	d	,	,	e	y	t	c	-	n	d	m	-	o	i	b	u	v	s	]	b	b	i	m	s	u	l	t	a	t	l	y	b	n

g	e	s	t	n	e	w	s	p	a	p	e	r		'	'	[	[	Y	e	d	i	o	t	h	A	h	r	o	n	o	t	h	]	]	'	'	'	H	e	b	r	e	w	-	l	a	n	g	u	a	g	e	p	e	r	i	o	d				
e	t	a	a	w	s	p	a	p	e	r	s	o	'	[	[	T	e	l	t	i	(	f	e	a	n	e	m	t	i	]	'	'	*	'	'	[	e	r	r	e	w	s	l	e	n	g	u	a	g	e	:	a	r	o	s	o	d	i				
i	r	s	c	o	e		e	n	a		i	T	T	h	A	o	a	i	n	n	h	S	r	m	u	w	]	e	y		s			[	'	i	n	e	i	a	'	s	i	w	d	d	e	'	h	s	o	l	r	i	f	r	:					
u	s	.	s	e	t	l	g	o	r	s	.	a	s	a	t	C	a	r	e	e	g	'	a	C	l	r	i	s	z	]	i	e	'	:	:	,	#	:	T	A	a	a	a	a	t	B	a	s	e	e	i	l	o	'	i	a	n	f	v	l		
-		t	u	a	e	v	r	t	i	d	,	t	B	A	m	S	u	s	y	u	t	]	]	A	s	a	o	i	g	s	]	]	,	.	.	.	:	s	M	B	o	l	o	u	s	:	T	o	u	a	-	n	:	d	w	o	a	p	n	u		
a	,	d	,	i	i	u	i	t	i	c	p	.	]	(	l	S	v	H	v	t	u	s	u	i	e	D	n	o	e	g	a	n	o	.	,	]	:	{	C	C	u	i	b	o	h	e	C	y	b	k	s	l	s	:	r	-	e	p	c	n	t	s

i	c	a	s	:	'	'	'	*	'	'	[	[	G	l	o	b	e	s	]	]	'	'	[	[	h	t	t	p	:	/	/	w	w	.	g	l	o	b	e	s	.	c	o	.	i	l	/	]	b	u	s	i	n	e	s	s	d	a						
c	a	l	:	'	'	'	*	'	'	'	[	[	T	a	a	b	a	]	]	'	'	(	[	[	t	t	p	:	/	/	w	w	.	b	u	o	b	a	l	.	c	o	m	u	n	/	s	A	-	y	t	i	n	e	s	s	a	e	t					
s	t	l	'								[	h	A	e	o	v	e	l	t		s	a	h	a	d	:	x	g	e	.	w	a	o	i	r	.	r	t	o	a	.	e	l	.	i	T	&	a	i	e	g		e	o	o	y								
t	t	'									&	[	&	&	m	C	o	e	r	o	n	e	'	:	:	,	i	'	o	d	w	.	:	n	i	i	i	s	a	a	u	e	.	e	n	i	/	o	m	l	c	C	.	(	e	f	t	g	i	r		i	i	u
a	'	n	:	,	C	:	&	:	#	*	:	a	f	D	r	u	s	u	]	l	,	.	o	m	e	l	p	<	,	d	h	a	:	d	e	u	o	o	t	/	i	h	n	c	s	i	f	S	,	u	r	h	o	s	t	,	t	u	n					
n	k	i	<	]	:	&	1	1	s	T	G	u	i	t	r	s	i	,			:	b	a	c	m	r	-	x	t	p	o	b	-	g	r	e	s	i	s	l	e	r	l	n	a	f	a	D	]	l	o	s	p	t	a	d	,	i	f	r	m			

i	l	y	*	'	'	[	[	H	a	a	r	e	t	z	]	H	a	'	A	r	e	t	z	]	]	'	'	[	[	h	t	t	p	:	/	/	w	w	.	h	a	a	r	e	t	z	.	c	o	.	i	l	/	]	R	e	l	a	t	i	v		
l	y	*	'	'	[	[	T	e	r	r	d	n	F	e	r	a	n	t	a	h	]	]	'	'	(	[	[	t	t	p	:	/	/	w	w	.	b	o	n	m	d	s	t	.	c	o	m	u	n	/	s	-	e	s	a	t	e	o	i				
r	e	'									h	A	i	l	n	n	t	t	e	H	a	l	s	r	c	n	o	l	'	s	a	h	a	d	:	x	n	e	.	w	a	a	m	r	t	d	h	e	o	h	.	o	l	.	c	&	o	p	i	n	i	v	e
k	i	:	*	s	C	O	S	a	n	t	h	i	T	i	m	'	l	i	e		:	,	i	m	c	d	w	-	2		◆	p	h	i	s	e	r	d	i	t	.	i	n	a	/	c	m	f	i	.	(	a	f	l	c	a	n	a					
d	s	-	!	[	t	B	T	C	o	m	m	g	d	]	]	W	o	n		a	a	e	,	:	.	b	a	e	r	.	<	t	a	i	b	-	d	u	l	c	n	n	c	/	a	r	n	e	s	i	]	]	l	i	c	e	y	s	t	o			
n	d	s	#	:	&	:	G	l	D	u	v	c	c	s	a	o	S	u	c	l	t	e	l	]	z	]	,	:	o	'	o	m	t	]	,	:	e	o	a	2	n	i	v	f	s	r	o	o	e	i	u	n	a	l	a	)	u	v	v	r	o		

First row: Green for excited, blue for not excited  
 Next 5 rows: top-5 guesses for the next character

# Realistic Wiki Hidden Unit

\*' '[ [Jerusalem Report] ]' '[ http://www.jrep.com/ ] Left-of-center En  
 \*' '[ hToausal m aogurt ] ]' ( http://www.bsinioom/ -iat aftenenter (ng  
 [ ' [Cassmene] Beaonds s sa [ad :xne.waaaaoca. s &ato -nfhhlsu m-ouc  
 ' s mFurnls iaet allsa ' : : , i' cdw-2tpi i isoeg.er/. a] (oseswr -ci ddrs [mt  
 : \* : AqDenebi utn | C ipre e | , . b1emr . 9 : ahb - npumugh nmp ) Teiretu : eoseodsald  
 # T&Tf Si wrp e ] ] aluvelru , s : - mprt s < moa2deys h ilr ] c . Augl , 1p , larc : fae

glish [ [weekly newspaper] ] \*' '[ [YNet News] ] \*' '[ http://www.ynetnews.c  
 lish c [Caakly] cawspaper ] ] \*' '[ hTaA at ] ]' ( http://www.bacahets.co  
 i aci - lhSoip ] i sec ] enp ] s . ' ' [ Co \*wess ] s a [ad :xne.waea . awat oa  
 eena , pCci etnedlox ] gic i | | s' [ sAmFeSahon ] t' : : , i momw-2 pi i isoessis . / er  
 syz . sfpenn a | ruel | rra . ' # \* : oDuFrei uep , : b1edr . < : ahb - nptwt . xi gh  
 a dpeamAr bdeor pitee ] dts - | T { [ BaAvTp oSwao , . . oacstp , tcoa2drulwo clens

om/ ] English-language website of Israel's largest newspaper ' [ [Yed  
 m/ - xgl ish languages a irs site of t slaelis singe t aawspapers o [ [Tel  
 . s &nt iaca - sardeelh o an t bisanfanreif ' aatdir scoe ena i TThAoai  
 n.c ] ( dceen epesaai ki i eele dh , irthraonse , coseus . setlgors . asatCare  
 / ma ) Tvdryzi couedlsu : tha - oo tu , stui f lvepery - tuaevrtid , tBAmSusy  
 r ] p . llvaod , , eytc - n dm - oi buvs ] bb imsul ta t lybna , d , i i u i t i c p . ] ( l S v H v t u


ioth Ahronoth ] ' ' Hebrew-language periodicals : ' ' \*' '[ [Globes] ]'  
 t i ( feanemt i ] ' ' \*' '[ errewslanguage : arosodical : ' ' \*' '[ Taaba ] ]'  
 nnh Srmuw ] ey ] s [ ' ineia ' si wdde ' h sol r i fr : s t l ' [ hAeovelt ] s  
 eg ' aClrisz ] i e' : : , # : T Aaaaat Baseeil o ' i anfv l t t ' ' & [ & & mCoerone ' : :  
 ut ] ] Asaoi gs ] ] , . : s MBolous : Toua - n : d woa pnua ' n : , C : & : # \* : afDrusu ] l ,  
 sui e Dnoegano . , ] : { CCuiboh e Cybksls : r - epcnt snk i < ] : & 11s TGuitr si ,

Above: Green for excited, blue for not excited

Below: top-5 guesses for the next character



# Vanilla RNN Flow Graph



$$\mathbf{a}_t = \mathbf{b} + \mathbf{W}\mathbf{h}_{t-1} + \mathbf{U}\mathbf{x}_t$$

$$\mathbf{h}_t = \tanh(\mathbf{a}_t)$$

$$\mathbf{y}_t = \mathbf{c} + \mathbf{V}\mathbf{h}_t$$

U – input to hidden

V – hidden to output

W – hidden to hidden


# Training RNN

- “Backpropagation through time”  
= Backpropagation

- What to do with this if

$$\mathbf{w} = \mathbf{w}_1 = \mathbf{w}_2 = \mathbf{w}_3?$$

$$\frac{\partial E}{\partial \mathbf{w}} = \frac{\partial E}{\partial \mathbf{y}_1} \frac{\partial \mathbf{y}_1}{\partial \mathbf{w}} + \frac{\partial E}{\partial \mathbf{y}_2} \frac{\partial \mathbf{y}_2}{\partial \mathbf{w}} + \frac{\partial E}{\partial \mathbf{y}_3} \frac{\partial \mathbf{y}_3}{\partial \mathbf{w}}$$




# Training RNN

- Again, assume

$$w = w_1 = w_2$$


$$\frac{\partial E}{\partial w} = \frac{\partial E}{\partial h_2} \frac{\partial h_2}{\partial w} + \frac{\partial E}{\partial h_3} \frac{\partial h_3}{\partial w}$$

$$\frac{\partial E}{\partial h_2} = \frac{\partial E}{\partial y_2} \frac{\partial y_2}{\partial h_2} + \frac{\partial E}{\partial h_3} \frac{\partial h_3}{\partial h_2}$$



# k timesteps?

$$\frac{\partial E}{\partial \mathbf{h}_2} = \frac{\partial E}{\partial \mathbf{y}_k} \frac{\partial \mathbf{y}_k}{\partial \mathbf{h}_k} \frac{\partial \mathbf{h}_k}{\partial \mathbf{h}_{k-1}} \cdots \frac{\partial \mathbf{h}_3}{\partial \mathbf{h}_2} + \frac{\partial E}{\partial \mathbf{y}_{k-1}} \frac{\partial \mathbf{y}_{k-1}}{\partial \mathbf{h}_{k-1}} \cdots \frac{\partial \mathbf{h}_3}{\partial \mathbf{h}_2} + \cdots$$



- What's the problem?

$$\frac{\partial \mathbf{h}_k}{\partial \mathbf{h}_{k-1}} = \frac{\partial \mathbf{h}_{k-1}}{\partial \mathbf{h}_{k-2}} = \cdots = \tanh(\mathbf{a}_t)' \mathbf{W}$$

- There are terms like  $\mathbf{W}^k$  in the gradient

$$\begin{aligned} \mathbf{a}_t &= \mathbf{b} + \mathbf{W}\mathbf{h}_{t-1} + \mathbf{U}\mathbf{x}_t \\ \mathbf{h}_t &= \tanh(\mathbf{a}_t) \\ \mathbf{y}_t &= \mathbf{c} + \mathbf{V}\mathbf{h}_t \end{aligned}$$

# What's wrong with $\mathbf{W}^k$ ?

- Suppose  $\mathbf{W}$  is diagonalizable for simplicity

$$\mathbf{W} = \mathbf{U}\mathbf{D}\mathbf{U}^T$$
$$\mathbf{W}^k = \mathbf{U}\mathbf{D}^k\mathbf{U}^T$$


- What if,
  - $\mathbf{W}$  has an eigenvalue of 4?
  - $\mathbf{W}$  has an eigenvalue of 0.25?
  - Both?

# Cannot train it with backprop

## Vanishing gradients

---

- Vanishing long term gradient ( $g^t$  is very small if  $g < 1$ )
- Strong short term gradient



# Do we need long-term gradients?

- Long-term dependency is one main reason we want temporal models

– Example:

Rob Ford told the flabbergasted reporters assembled at the press conference that \_\_\_\_\_.


German for “travel”

Die Koffer waren gepackt, und er *reiste*, nachdem er seine Mutter und seine Schwestern geküsst und noch ein letztes Mal sein angebetetes Gretchen an sich gedrückt hatte, das, in schlichten weißen Musselin gekleidet und mit einer einzelnen Nachthyazinthe im üppigen braunen Haar, kraftlos die Treppe herabgetaumelt war, immer noch blass von dem Entsetzen und der Aufregung des vorangegangenen Abends, aber voller Sehnsucht, ihren armen schmerzenden Kopf noch einmal an die Brust des Mannes zu legen, den sie mehr als ihr eigenes Leben liebte, *ab.*

Only now we are sure the travel started, not ended (reiste an)

# LSTM: Long short-term Memory

- Need memory!
  - Vanilla RNN has volatile memory (automatically transformed every time-step)
  - More “fixed” memory stores info longer so errors don’t need to be propagated very far
- Complex architecture with memory





# LSTM Starting point

- Instead of using volatile state transition

$$\mathbf{h}_t = \tanh(\mathbf{W}\mathbf{h}_{t-1} + \mathbf{U}\mathbf{x}_t + b)$$

- Use fixed transition and learn the difference

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

- Now we can truncate the BPTT safely after several timesteps
- However, this has the drawback of  $\mathbf{c}_t$  being stored for too long
  - Add a weight? (subject to vanishing as well)
  - Add an “adaptive weight”

# Forget Gate

- Decide how much of  $c_{t-1}$  should we forget

$$\begin{aligned}c_t &= f_t \odot c_{t-1} + z_t \\z_t &= \tanh(\mathbf{W}\mathbf{y}_{t-1} + \mathbf{U}\mathbf{x}_t + b)\end{aligned}$$

- Forget neurons also trained

$$f_t = \sigma(\mathbf{W}_f\mathbf{x}_t + \mathbf{R}_f\mathbf{y}_{t-1} + \mathbf{p}_f \odot c_{t-1} + b_f)$$

- How much we forget is dependent on:
  - Previous output
  - Current input
  - Previous memory

# Input Modulation

- Memory is supposed to be “persistent”
- Some input might be corrupt and should not affect our memory
- We may want to decide which input affects our memory

- Input Gate:

$$\mathbf{i}^t = \sigma(\mathbf{W}_i \mathbf{x}^t + \mathbf{R}_i \mathbf{y}^{t-1} + \mathbf{p}_i \odot \mathbf{c}^{t-1} + \mathbf{b}_i) \quad \textit{input gate}$$

- Final memory update:

$$\mathbf{c}^t = \mathbf{i}^t \odot \mathbf{z}^t + \mathbf{f}^t \odot \mathbf{c}^{t-1}$$

# Output Modulation


- Do not always “tell” what we remembered

$$\begin{aligned} \mathbf{o}^t &= \sigma(\mathbf{W}_o \mathbf{x}^t + \mathbf{R}_o \mathbf{y}^{t-1} + \mathbf{p}_o \odot \mathbf{c}^t + \mathbf{b}_o) && \textit{output gate} \\ \mathbf{y}^t &= \mathbf{o}^t \odot h(\mathbf{c}^t) && \textit{block output} \end{aligned}$$

- Only output if we “feel like it”
- The output part can vary a lot depending on applications

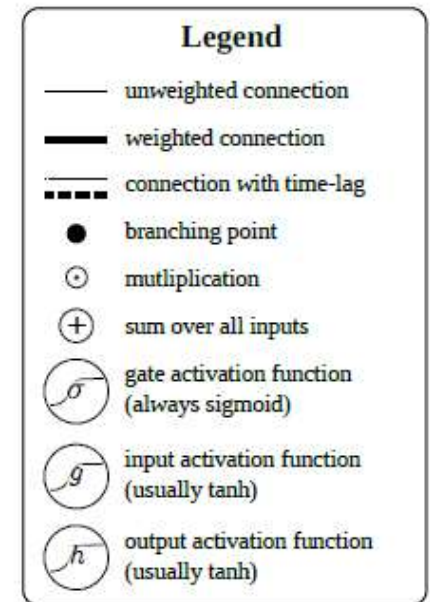
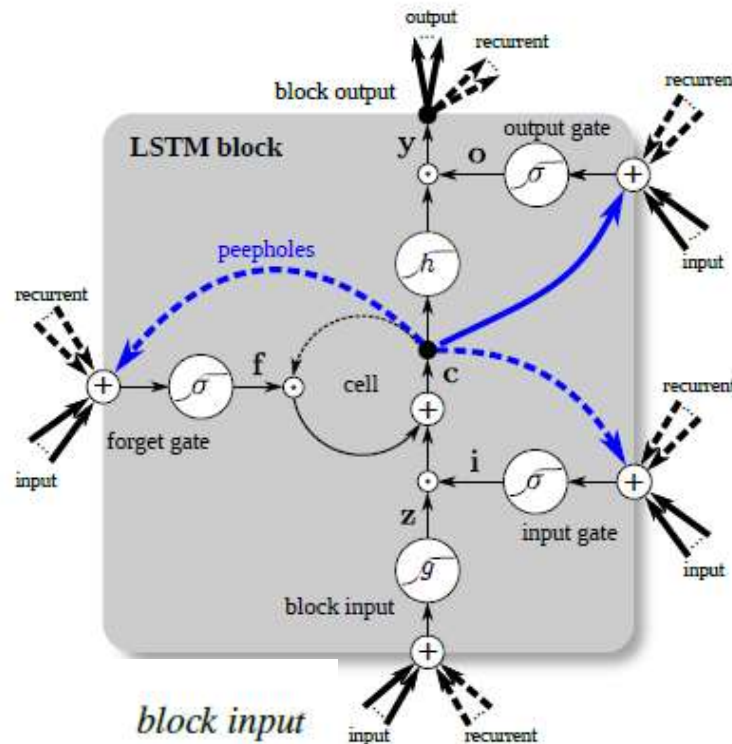
# LSTM

- Hochreiter & Schmidhuber (1997)
- Use gates to remember things for a long period of time
- Use gates to modulate input and output



# LSTM Architecture

- “Official version” with a lot of peepholes



$$z^t = g(W_z x^t + R_z y^{t-1} + b_z)$$

$$i^t = \sigma(W_i x^t + R_i y^{t-1} + p_i \odot c^{t-1} + b_i) \quad \text{input gate}$$

$$f^t = \sigma(W_f x^t + R_f y^{t-1} + p_f \odot c^{t-1} + b_f) \quad \text{forget gate}$$

$$c^t = i^t \odot z^t + f^t \odot c^{t-1} \quad \text{cell state}$$

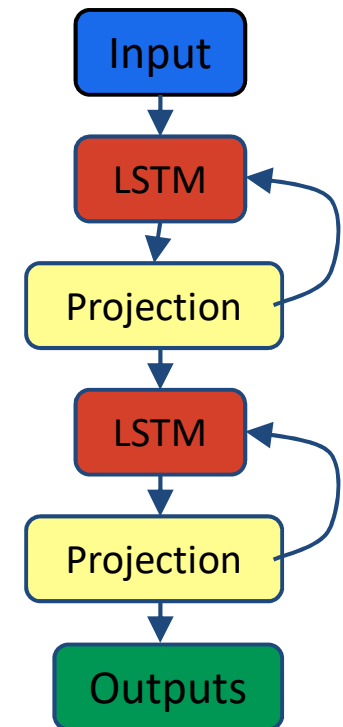
$$o^t = \sigma(W_o x^t + R_o y^{t-1} + p_o \odot c^t + b_o) \quad \text{output gate}$$

$$y^t = o^t \odot h(c^t) \quad \text{block output}$$

Cf. LSTM: a search space odyssey

# Google Speech recognition

- Task:
  - Google Now/Voice search / mobile dictation
  - Streaming, real-time recognition in 50 languages
- Model:
  - Deep Projection Long-Short Term Memory Recurrent Neural networks
  - Distributed training with asynchronous gradient descent across hundreds of machines.
  - Cross-entropy objective (truncated backpropagation through time) followed by sequence discriminative training (sMBR).
  - 40-dimensional filterbank energy inputs
  - Predict 14,000 acoustic state posteriors



# LSTM Large vocabulary speech recognition



Models	Parameters	Cross-Entropy	sMBR sequence training
ReLU DNN	85M	11.3	10.4
Deep Projection LSTM RNN (2 layer)	13M	10.7	9.7

- [\*Long Short-Term Memory Recurrent Neural Network Architectures for Large Scale Acoustic Modeling\*](#) H. Sak, A. Senior, F. Beaufays to appear in Interspeech 2014
- [\*Sequence Discriminative Distributed Training of Long Short-Term Memory Recurrent Neural Networks\*](#) H. Sak, O. Vinyals, G. Heigold A. Senior, E. McDermott, R. Monga, M. Mao to appear in Interspeech 2014

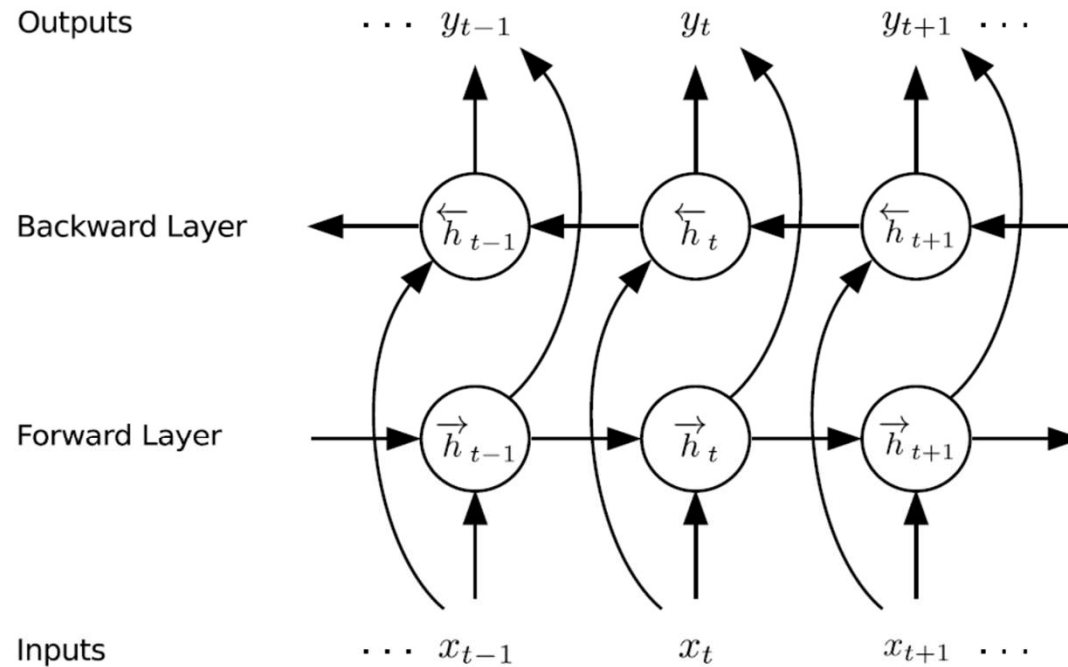
Voice search task; Training data: 3M utterances (1900 hrs); models trained on CPU clusters

Slide provided by Andrew Senior, Vincent Vanhoucke, Hasim Sak (June 2014)



# Bidirectional LSTM

Both forward and backward paths  
Still DAG!



**Fig. 2.** Bidirectional Recurrent Neural Network

# Pen trajectories

- Task: generate pen trajectories by predicting one  $(x,y)$  point at a time
- Data: IAM **online** handwriting, 10K training sequences, **many writers**, unconstrained style, captured from a whiteboard

So you say to your neighbour,  
would find the bus safe and sound  
would be the vineyards

- First problem: what to use for the **density model**?

# Recurrent Mixture Density Networks

- Network outputs parameterise a **mixture distribution** (usually Gaussian)
- Every prediction conditioned on all inputs so far

$$\Pr(x_{t+1}|x_{1:t}) = \sum_k w_k(x_{1:t}) \mathcal{N}(x_{t+1}|\mu_k(x_{1:t}), \Sigma_k(x_{1:t}))$$

- Number of components is number of **choices** for what comes next
- M. Schuster, “Better Generative Models for Sequential Data Problems: Bidirectional Recurrent Mixture Density Networks”, NIPS 1999

# Network details

$$x_t \in \mathbb{R} \times \mathbb{R} \times \{0, 1\}$$

$$y_t = \left( e_t, \{ \pi_t^j, \mu_t^j, \sigma_t^j, \rho_t^j \}_{j=1}^M \right)$$

$$\hat{y}_t = \left( \hat{e}_t, \{ \hat{w}_t^j, \hat{\mu}_t^j, \hat{\sigma}_t^j, \hat{\rho}_t^j \}_{j=1}^M \right) = b_y + \sum_{n=1}^N W_{h^n y} h_t^n$$

$$e_t = \frac{1}{1 + \exp(\hat{e}_t)} \quad \implies e_t \in (0, 1)$$

$$\pi_t^j = \frac{\exp(\hat{\pi}_t^j)}{\sum_{j'=1}^M \exp(\hat{\pi}_t^{j'})} \quad \implies \pi_t^j \in (0, 1), \quad \sum_j \pi_t^j = 1$$



$$\mu_t^j = \hat{\mu}_t^j \quad \implies \mu_t^j \in \mathbb{R}$$

$$\sigma_t^j = \exp(\hat{\sigma}_t^j) \quad \implies \sigma_t^j > 0$$

$$\rho_t^j = \tanh(\hat{\rho}_t^j) \quad \implies \rho_t^j \in (-1, 1)$$

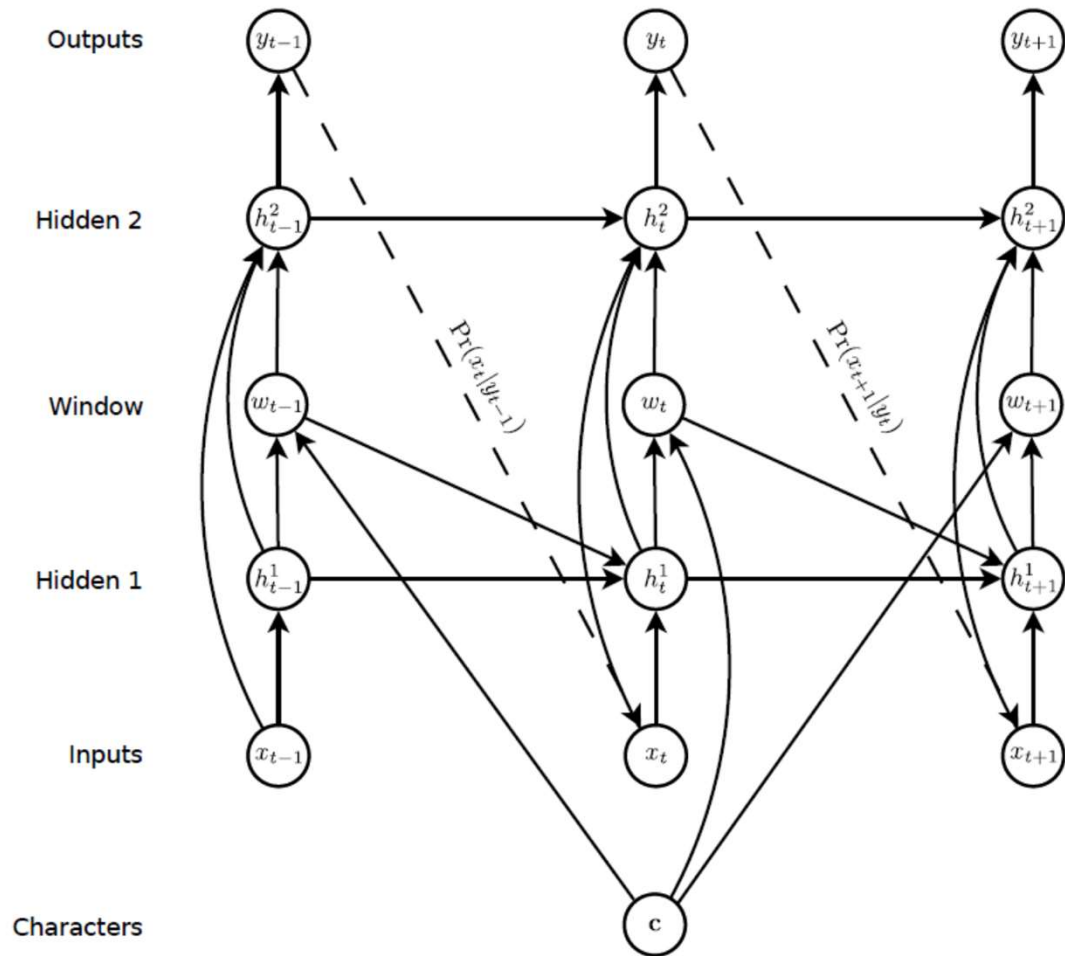
A. Graves, "Generating Sequences with Recurrent Neural Networks,  
arXiv:1308.0850v5

# Illustration of mixture density




# Synthesis

- Adding text input



# Learning text windows

$$\phi(t, u) = \sum_{k=1}^K \alpha_t^k \exp \left( -\beta_t^k (\kappa_t^k - u)^2 \right)$$
$$w_t = \sum_{u=1}^U \phi(t, u) c_u$$



Thought that the muster from

A demonstration of online handwriting recognition by an RNN  
with Long Short Term Memory (from Alex Graves)

this deep learning course is great!

this deep learning course is great!


this deep learning course is great!

<http://www.cs.toronto.edu/~graves/handwriting.html>



# LSTM Architecture Explorations

- “Official version” with a lot of peepholes



$$z^t = g(W_z x^t + R_z y^{t-1} + b_z)$$

$$i^t = \sigma(W_i x^t + R_i y^{t-1} + p_i \odot c^{t-1} + b_i)$$

$$f^t = \sigma(W_f x^t + R_f y^{t-1} + p_f \odot c^{t-1} + b_f)$$

$$c^t = i^t \odot z^t + f^t \odot c^{t-1}$$

$$o^t = \sigma(W_o x^t + R_o y^{t-1} + p_o \odot c^t + b_o)$$

$$y^t = o^t \odot h(c^t)$$

*block input*

*input gate*

*forget gate*

*cell state*

*output gate*

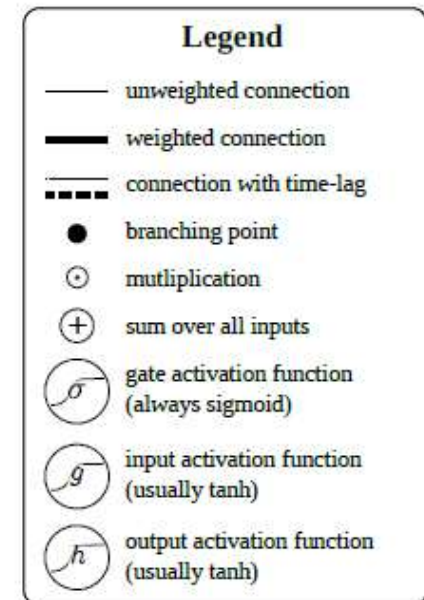

*block output*

Cf. LSTM: a search space odyssey

# A search space odyssey

- What if we remove some parts of this?

1. No Input Gate (NIG)
2. No Forget Gate (NFG)
3. No Output Gate (NOG)
4. No Input Activation Function (NIAF)
5. No Output Activation Function (NOAF)
6. No Peepholes (NP)
7. Coupled Input and Forget Gate (CIFG)
8. Full Gate Recurrence (FGR)




Cf. LSTM: a search space odyssey

# Datasets

- TIMIT
  - Speech data
  - Framewise classification
  - 3696 sequences, 304 frames per sequence
- IAM
  - Handwriting stroke data
  - Map handwriting strokes to characters
  - 5535 sequences, 334 frames per sequence
- JSB
  - Music Modeling
  - Predict next note
  - 229 sequences, 61 frames per sequence

# Results



1. No Input Gate (NIG)
2. No Forget Gate (NFG)
3. No Output Gate (NOG)
4. No Input Activation Function (NIAF)


5. No Output Activation Function (NOAF)
6. No Peepholes (NP)
7. Coupled Input and Forget Gate (CIFG)
8. Full Gate Recurrence (FGR)

Cf. LSTM: a search space odyssey

# Impact of Parameters

- Analysis method: fANOVA (Hutter et al. 2011, 2014)
- (Random) Decision forests trained on the parameter space to partition the parameter space and find the best parameter
- Given trained (random) decision forest, can go to each leaf node and count the impact of missing one predictor


# Impact of Parameters



*Figure 3.* Pie charts showing which fraction of variance of the test set performance can be attributed to each of the hyperparameters. The percentage of variance that is due to interactions between multiple parameters is indicated as “higher order.”

Cf. LSTM: a search space odyssey

# Impact of Parameters



*Figure 5.* Left: The predicted marginal error for combinations of learning rate and hidden size. Right: The component that is solely due to the interaction of the two and cannot be attributed to changes in one of them alone. In other words the difference to the case of them being perfectly independent. (Blue is better than red.)

learning rate  $\times$  hidden size = 6.7%

learning rate  $\times$  input noise = 4.4%

hidden size  $\times$  input noise = 2.0%

learning rate  $\times$  momentum = 1.5%

momentum  $\times$  hidden size = 0.6%

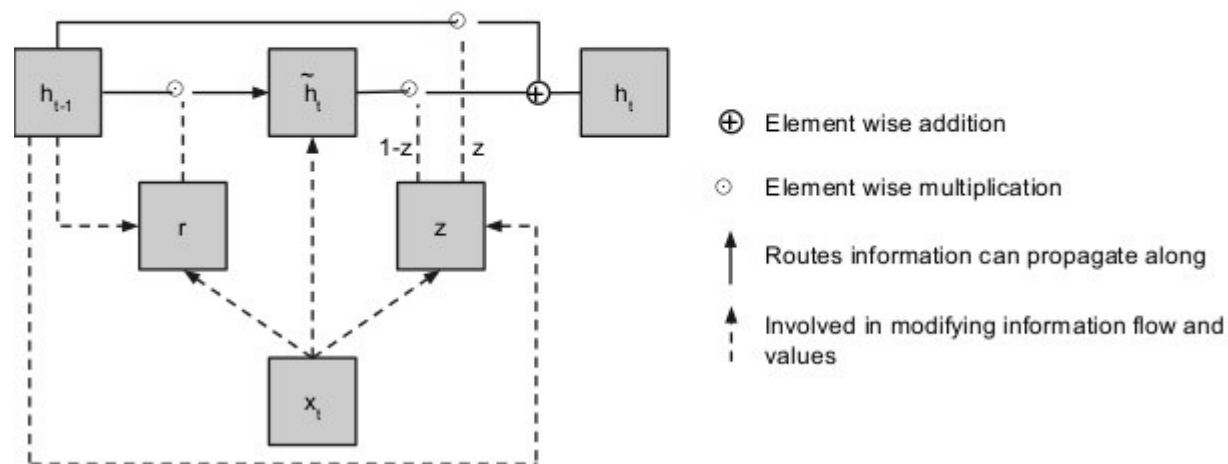
momentum  $\times$  input noise = 0.4%

Cf. LSTM: a search space odyssey

# GRU: Gated Recurrence Unit

- Much simpler than LSTM
  - No output gate
  - Coupled input and forget gate

## Gated Recurrent Unit - GRU





# Data


- Music Datasets:
  - Nottingham, 1200 sequences
  - MuseData, 881 sequences
  - JSB, 382 sequences
- Ubisoft Data A
  - Speech, 7230 sequences, length 500
- Ubisoft Data B
  - Speech, 800 sequences, length 8000

# Results

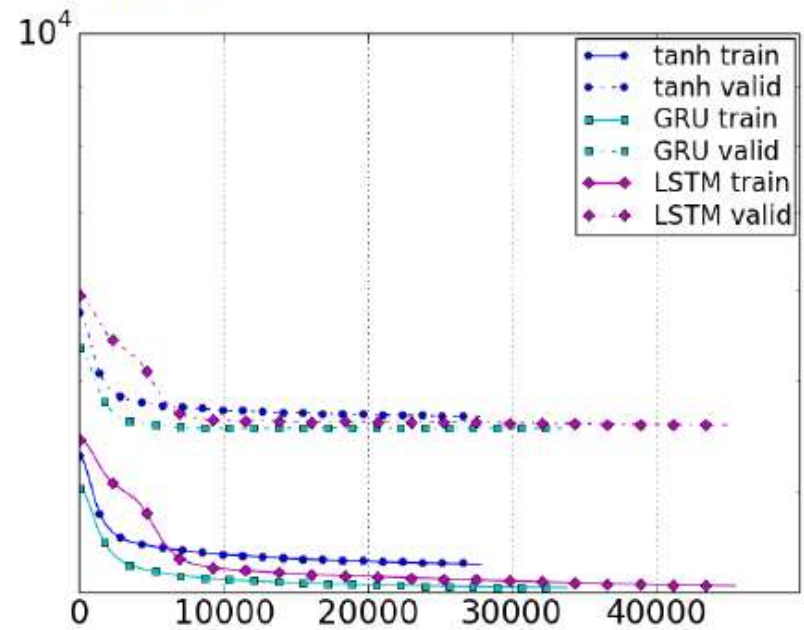
Nottingham  
Music, 1200 sequences

MuseData  
Music, 881 sequences

Wall Clock Time (seconds)



(a) Nottingham Dataset



(b) MuseData Dataset

Cf. Empirical Evaluation of Gated Recurrent Neural Network Modeling

# Results


Ubisoft Data A

Speech, 7230 sequences, length 500


Ubisoft Data B

Speech, 800 sequences, length 8000

Wall Clock Time (seconds)



(a) Ubisoft Dataset A

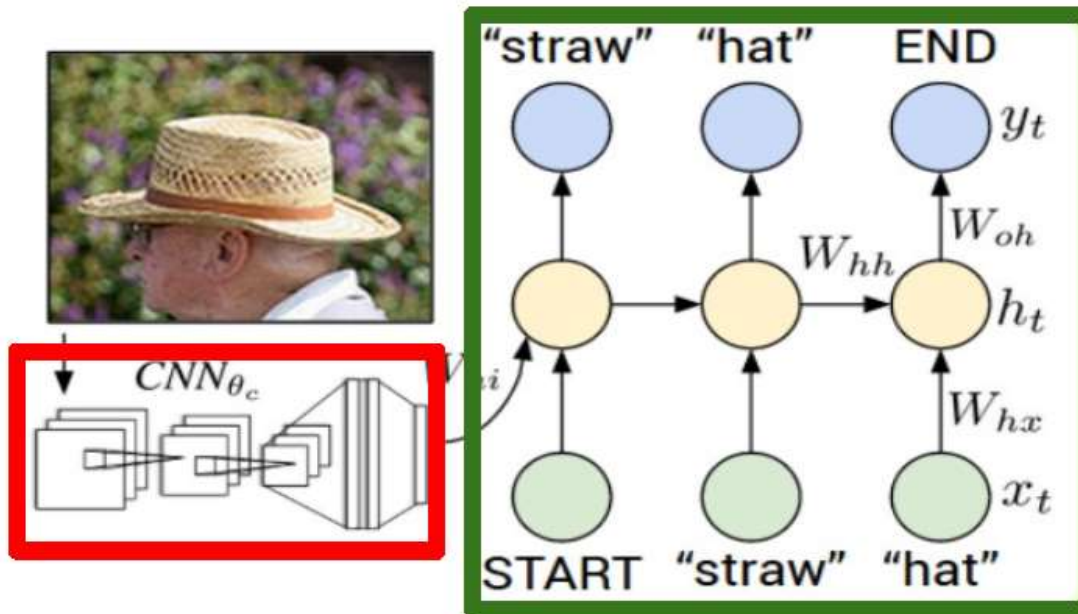


(b) Ubisoft Dataset B

Cf. Empirical Ev:

# CNN+RNN Example

## Recurrent Neural Network



## Convolutional Neural Network



“straw hat”

training example

image

conv-64

conv-64

maxpool

conv-128

conv-128

maxpool

conv-256

conv-256

maxpool

conv-512

conv-512

maxpool

conv-512

conv-512

maxpool

FC-4096

FC-4096

FC-1000

softmax



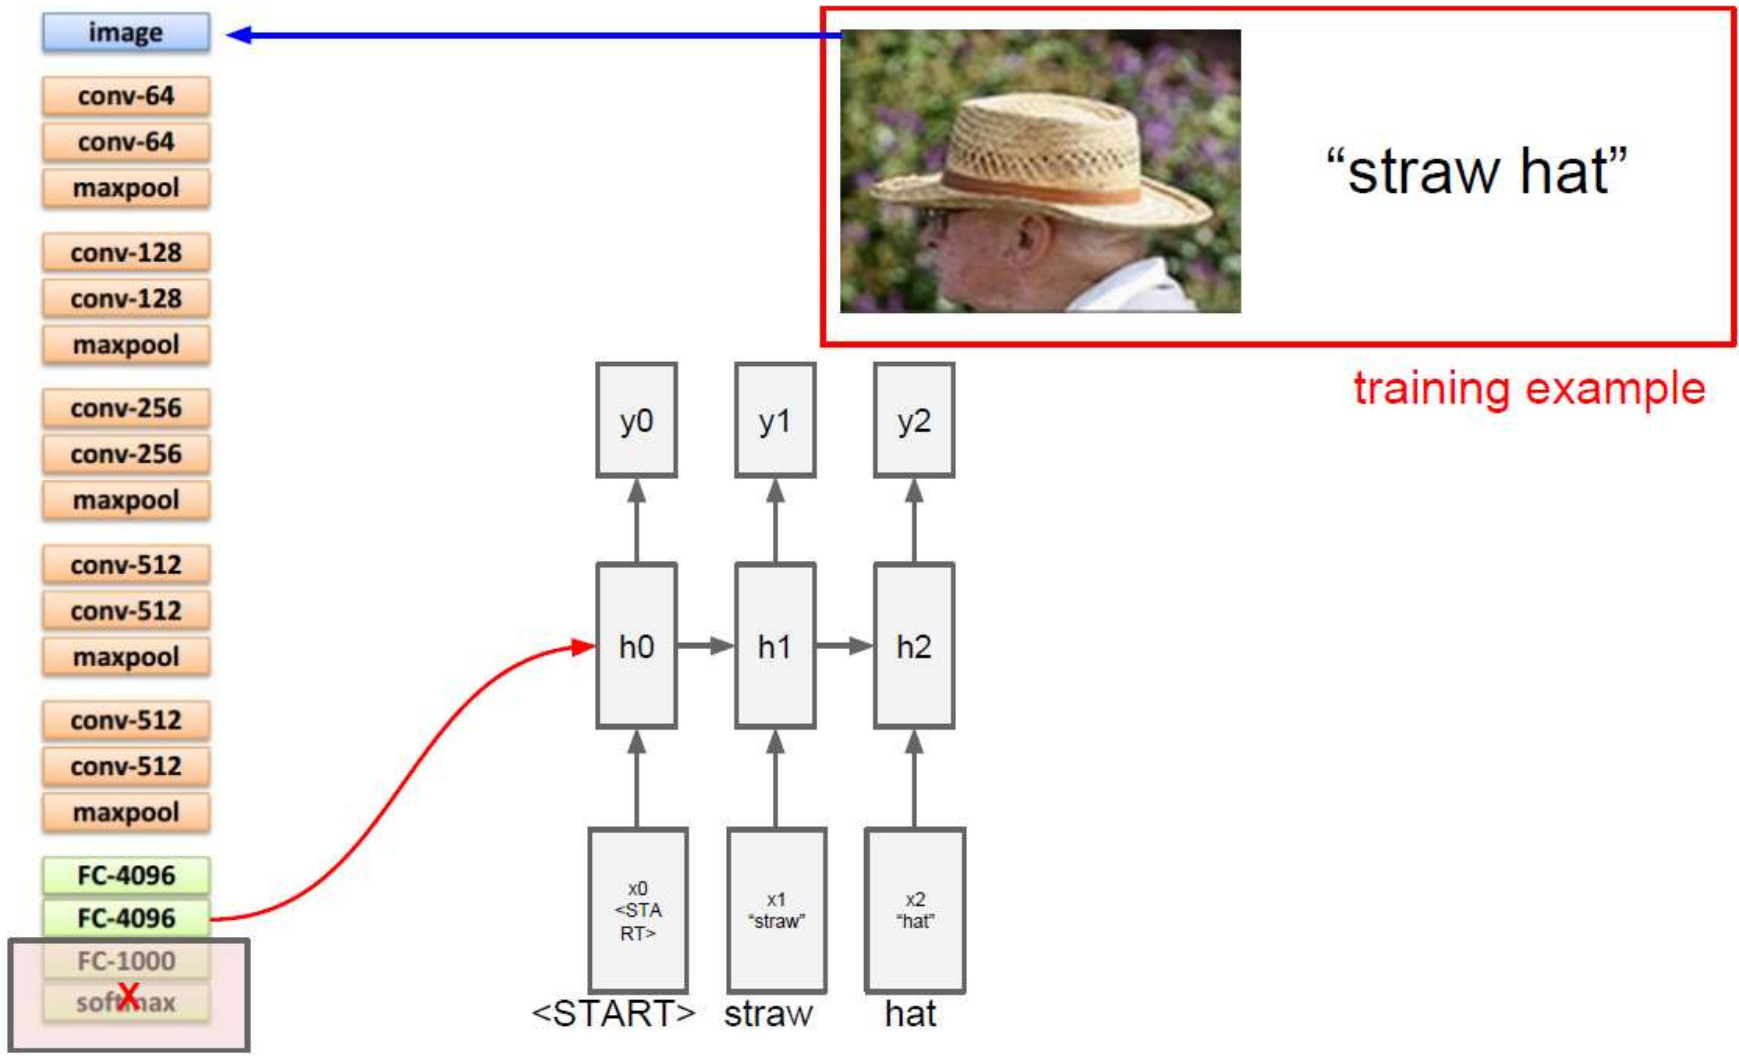
“straw hat”

training example



“straw hat”

training example





image

conv-64

conv-64

maxpool

conv-128

conv-128

maxpool

conv-256

conv-256

maxpool

conv-512

conv-512

maxpool

conv-512

conv-512

maxpool

FC-4096

FC-4096

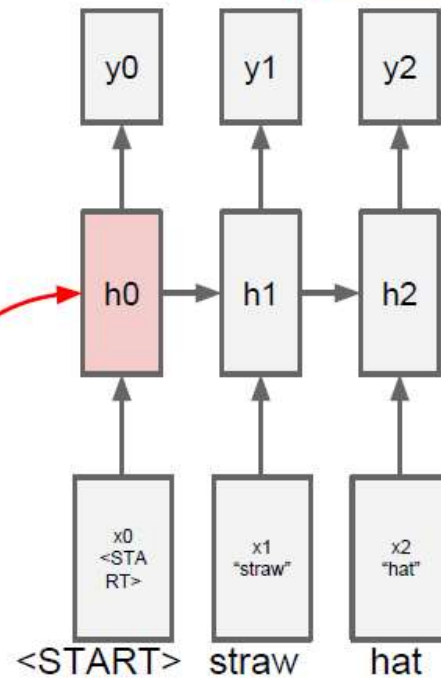
FC-1000

softmax



“straw hat”

training example



before:

$$h_0 = \max(0, W_{xh} * x_0)$$

now:

$$h_0 = \max(0, W_{xh} * x_0 + W_{ih} * v)$$

image

conv-64

conv-64

maxpool

conv-128

conv-128

maxpool

conv-256

conv-256

maxpool

conv-512

conv-512

maxpool

conv-512

conv-512

maxpool

FC-4096

FC-4096

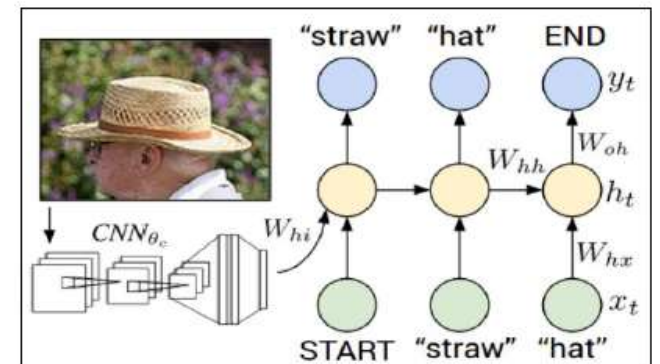
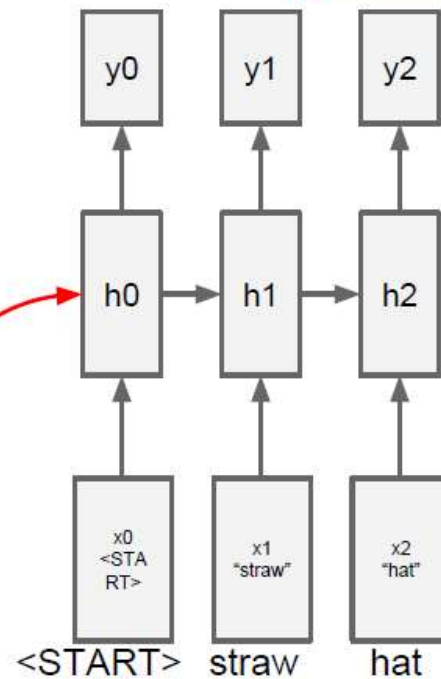
FC-1000

softmax



“straw hat”

training example





test image

image

conv-64

conv-64

maxpool

conv-128

conv-128

maxpool

conv-256

conv-256

maxpool

conv-512

conv-512

maxpool

conv-512

conv-512


maxpool

FC-4096

FC-4096



test image



image

conv-64

conv-64

maxpool

conv-128

conv-128

maxpool

conv-256

conv-256

maxpool

conv-512

conv-512

maxpool

conv-512

conv-512


maxpool

FC-4096

FC-4096



test image



image

conv-64

conv-64

maxpool

conv-128

conv-128

maxpool

conv-256

conv-256

maxpool

conv-512

conv-512

maxpool

conv-512

conv-512

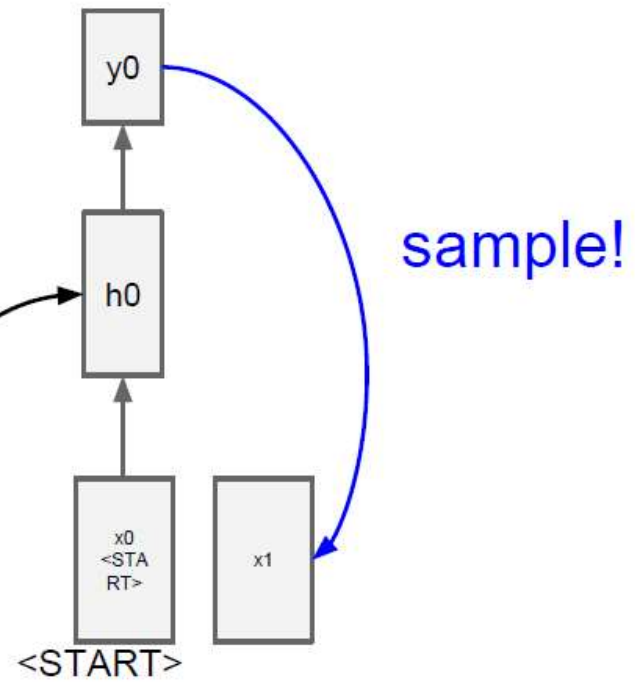
maxpool

FC-4096

FC-4096




test image






test image






test image



sample!





image

conv-64

conv-64

maxpool

conv-128

conv-128

maxpool

conv-256

conv-256

maxpool

conv-512

conv-512

maxpool

conv-512

conv-512


maxpool

FC-4096

FC-4096

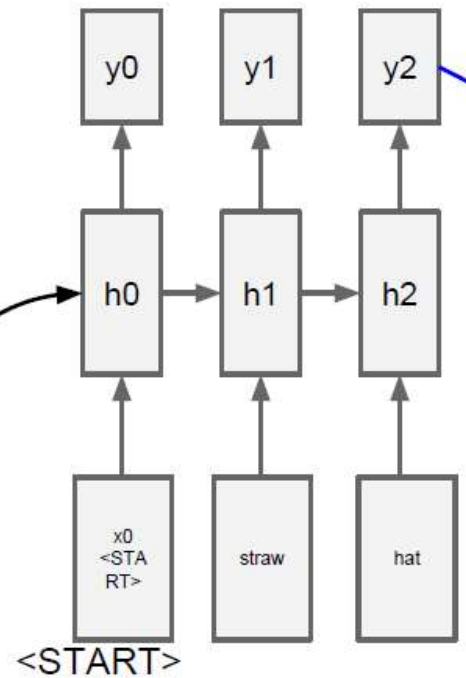


test image







test image

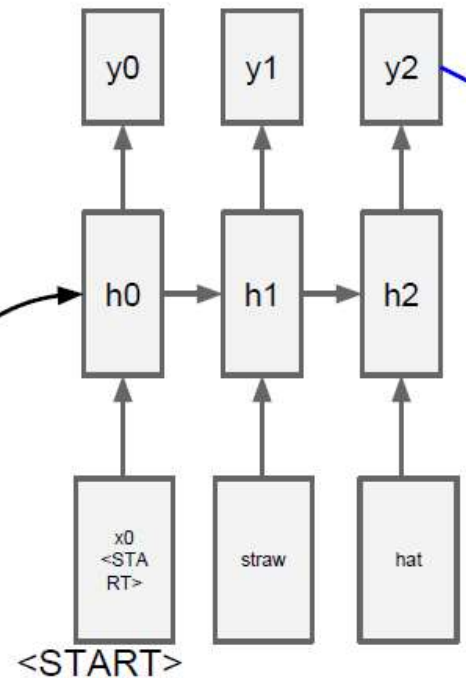


sample!  
 <END> token  
 => finish.






test image



sample!  
 <END> token  
 => finish.


- Don't have to do greedy word-by-word sampling, can also search over longer phrases with **beam search**

# RNN vs. LSTM



“hidden” representation  
(e.g. 200 numbers)  
 $h_1 = \max(0, W_{xh} * x_1 + W_{hh} * h_0)$

# RNN vs. LSTM

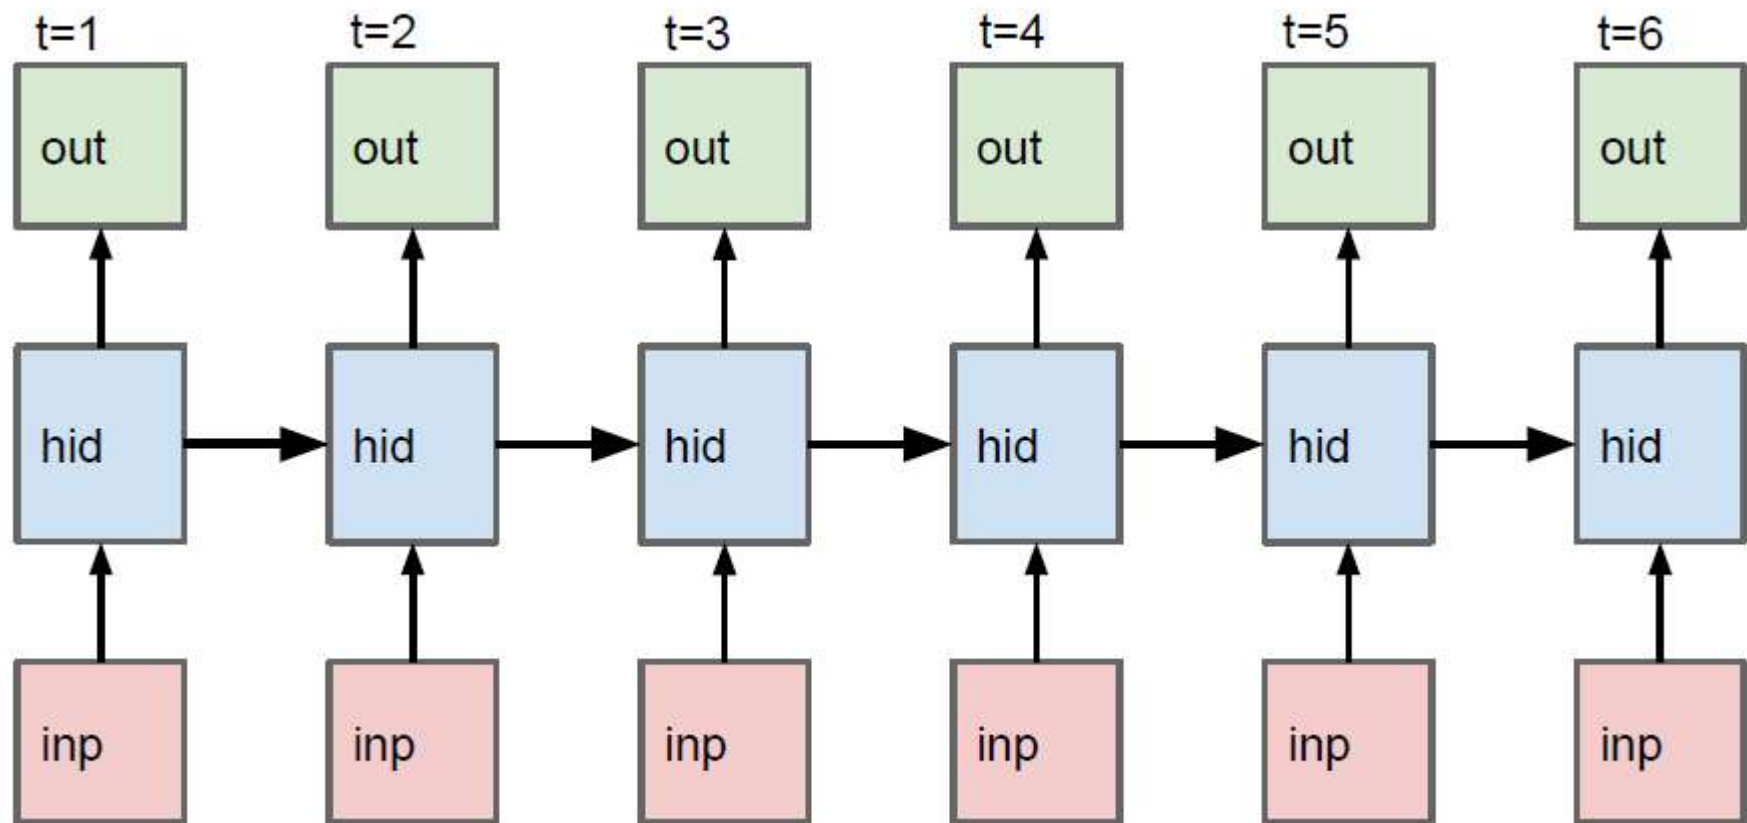


“hidden” representation  
(e.g. 200 numbers)  
 $h1 = \max(0, W_{xh} * x1 + W_{hh} * h0)$


LSTM changes the form of the equation for **h1** such that:

1. more expressive multiplicative interactions
2. gradients flow nicer
3. network can explicitly decide to reset the hidden state

# RNN



# I STM



# Image Sentence Datasets

a man riding a bike on a dirt path through a forest.  
bicyclist raises his fist as he rides on desert dirt trail.  
this dirt bike rider is smiling and raising his fist in triumph.  
a man riding a bicycle while pumping his fist in the air.  
a mountain biker pumps his fist in celebration.



## Microsoft COCO

*[Tsung-Yi Lin et al. 2014]*

[mscoco.org](http://mscoco.org)

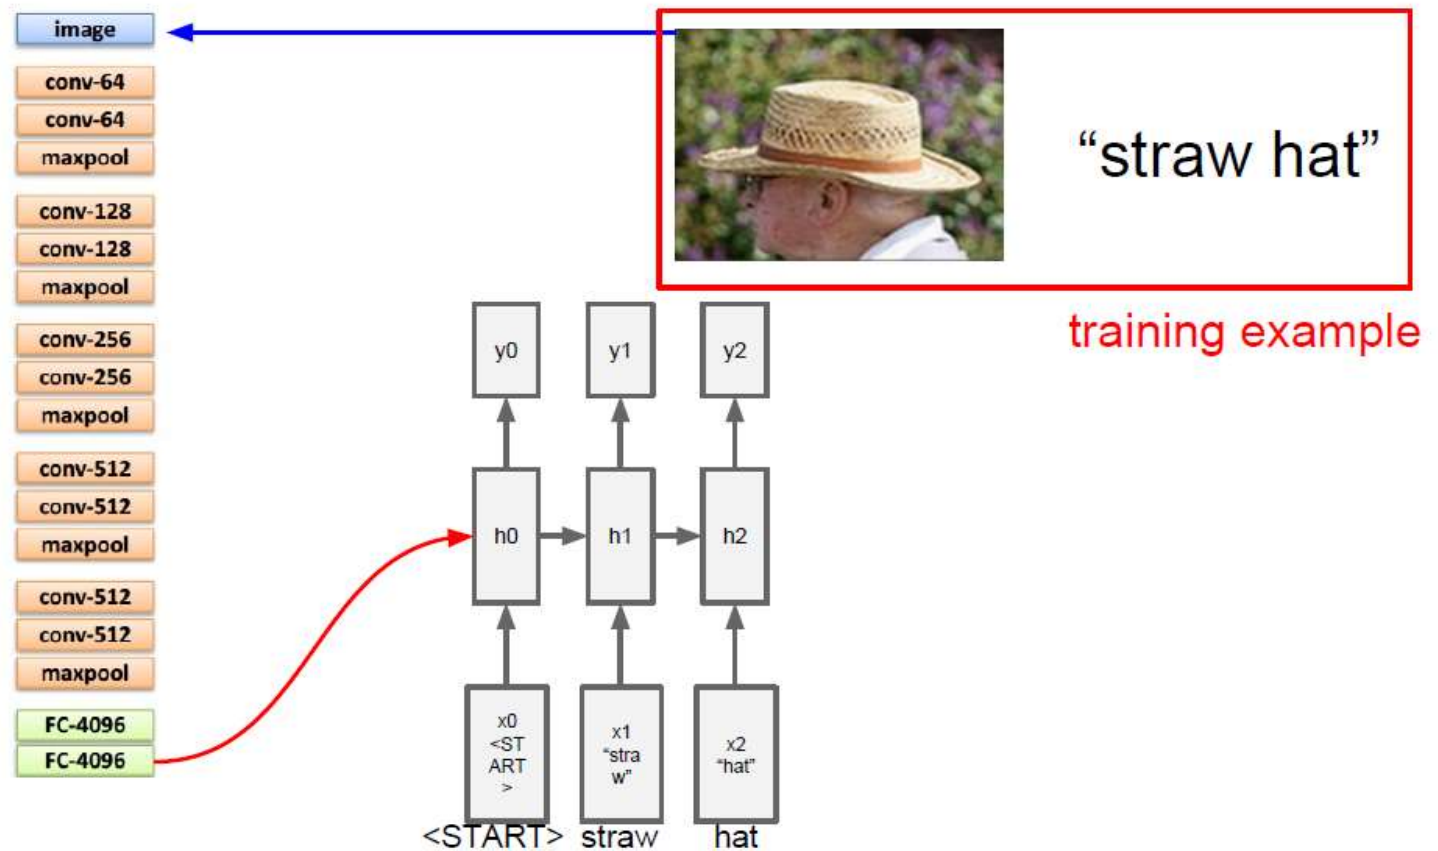
currently:

~120K images

~5 sentences each

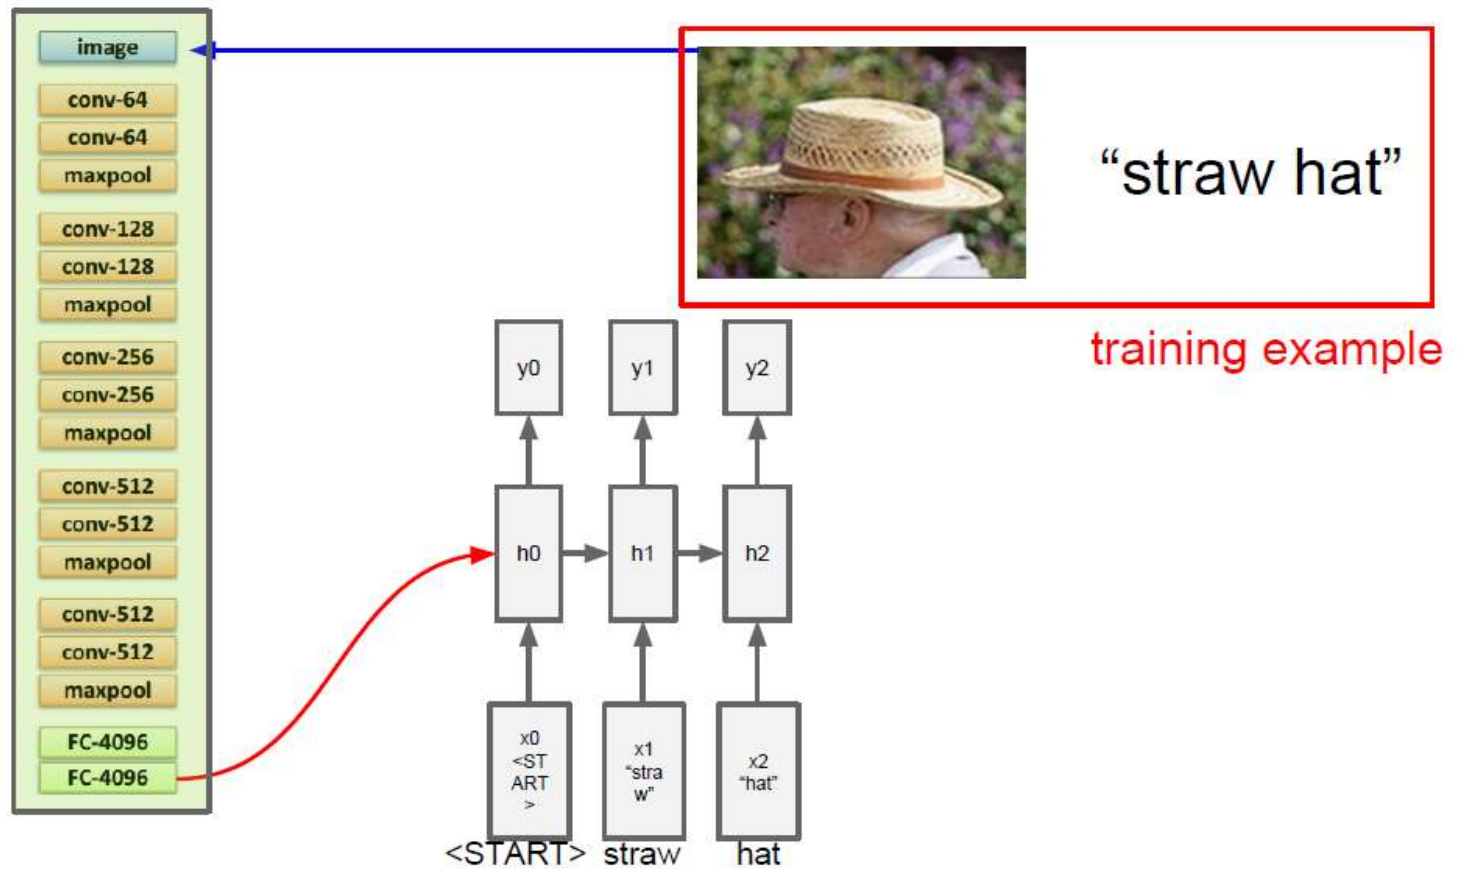


# + Transfer Learning



# Pre-training

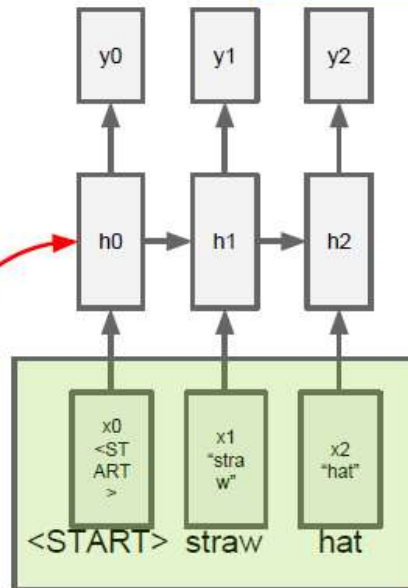
use weights  
pretrained from  
ImageNet



use weights pretrained from ImageNet



training example



use word vectors pretrained with word2vec [1]

[1] Mikolov et al., 2013

# Summary of the approach

We wanted to describe images with sentences.

1. Define a single function from input -> output
2. Initialize parts of net from elsewhere if possible
3. Get some data
4. Train with SGD

# Wow I can't believe that worked



a group of people standing  
around a room with  
remotes  
logprob: -9.17



a young boy is holding a  
baseball bat  
logprob: -7.61



a cow is standing in the middle of a street  
logprob: -8.84

Wow I can't believe that worked



a cat is sitting on a toilet seat  
logprob: -7.79



a display case filled with lots of different types of donuts  
logprob: -7.78



a group of people sitting at a table with wine glasses  
logprob: -6.71

Well, I can kind of see it



a man standing next to a clock on a wall  
logprob: -10.08



a young boy is holding a  
baseball bat  
logprob: -7.65



a cat is sitting on a couch with a remote control  
logprob: -12.45

Well, I can kind of see it



a baby laying on a bed with a stuffed bear  
logprob: -8.66



a table with a plate of food and a cup of coffee  
logprob: -9.93



a young boy is playing frisbee in the park  
logprob: -9.52



# Not sure what happened there...



a toilet with a seat up in a bathroom  
logprob: -13.44



a woman holding a teddy bear in front of a mirror  
logprob: -9.65



















a horse is standing in the middle of a road  
logprob: -10.34

See predictions on  
 1000 COCO images:  
<http://bit.ly/neuraltalkdemo>

**NeuraTalk Sentence Generation Results**  
 (Showing results for coco on 1000 images)

Dev params: name: /neuraltalk/params/result\_coco.json?team\_name=1,dataset\_path=1,src\_model\_checkpoint\_coco\_senteval\_03\_milford\_eda\_jun\_11\_14.pl?dump\_option=all; max\_image=1000  
 Real average complexity of ground truth words: 11.95

 a group of people walking down a street Suggested: -11.71	 a man is standing on a beach with a surfboard Suggested: -12.24	 a woman is holding a black umbrella in the rain Suggested: -11.54
 a plate of food with a sandwich and a salad Suggested: -8.96	 a woman sitting on a bench with a dog Suggested: -8.33	 a traffic light with a red light on top Suggested: -10.33
 a car lying on a street with a sign Suggested: -11.61	 a group of people sitting at a table with wine glasses Suggested: -11.1	 a bus is parked in a parking lot Suggested: -11.19
 a pizza with toppings on a white plate Suggested: -7.17	 a herd of sheep standing on top of a hill Suggested: -11.14	 a giraffe is standing in a field of grass Suggested: -11.12
 a pizza with toppings on a white plate Suggested: -7.08	 a large clock tower with a clock on top Suggested: -11.24	 a woman is holding a pink umbrella in the rain Suggested: -12.07
 a plate of food with a fork and a glass of wine Suggested: -11.16		

## **What this approach Doesn't do:**

- There is no *reasoning*
- A single glance is taken at the image, no objects are detected, etc.
- We can't just describe any image