COS 484/584

L17: Neural Machine Translation - 2

Spring 2021

Last time: Sequence to Sequence learning (Seq2seq)

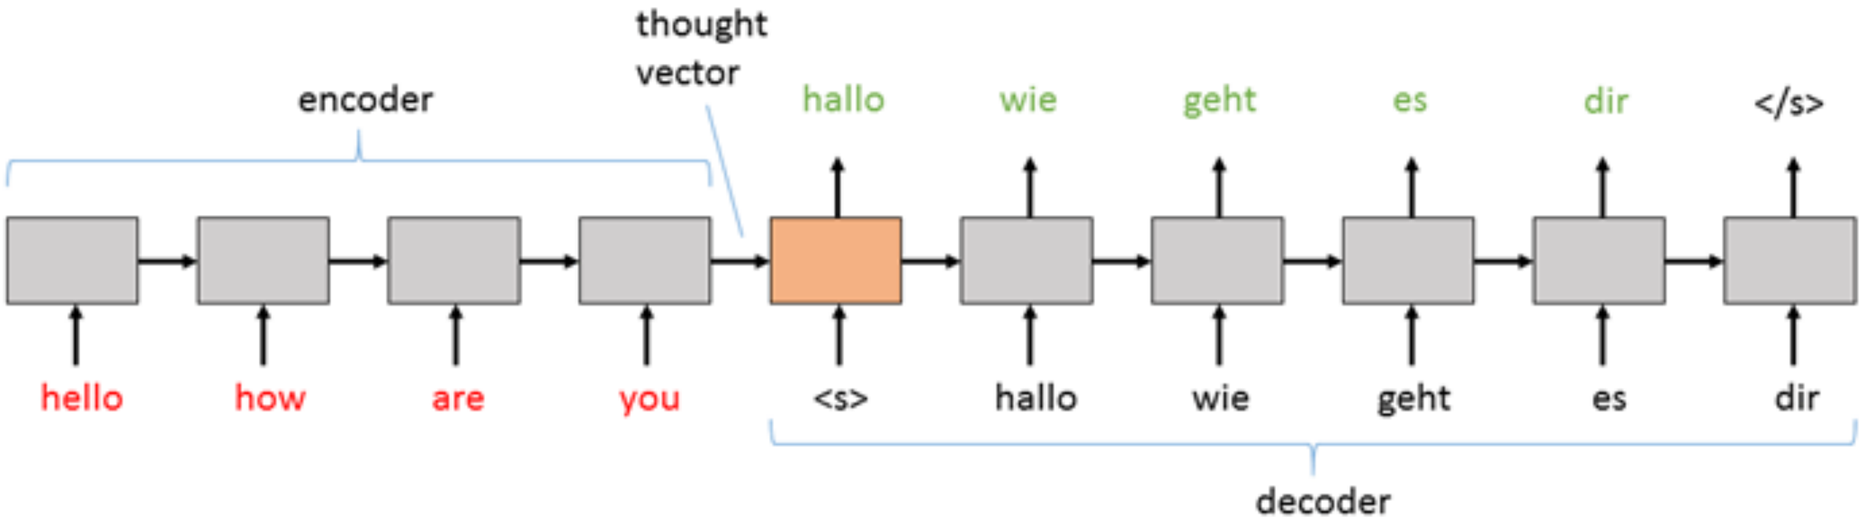

- Encode entire input sequence into a single vector **(using an RNN)**
- Decode one word at a time **(again, using an RNN!)**

(Sutskever et al., 2014)

How seq2seq changed the MT landscape

seq2seq
Search term

+ Compare

United States | 2004 - present | All categories | Web Search

Interest over time

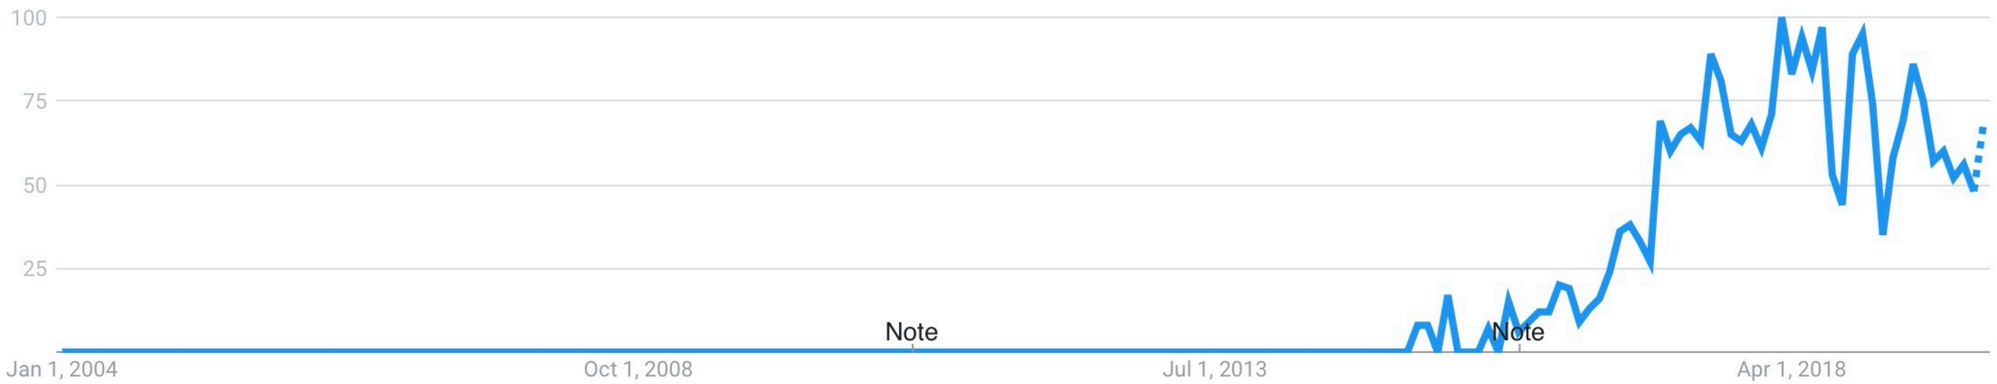

MT Progress

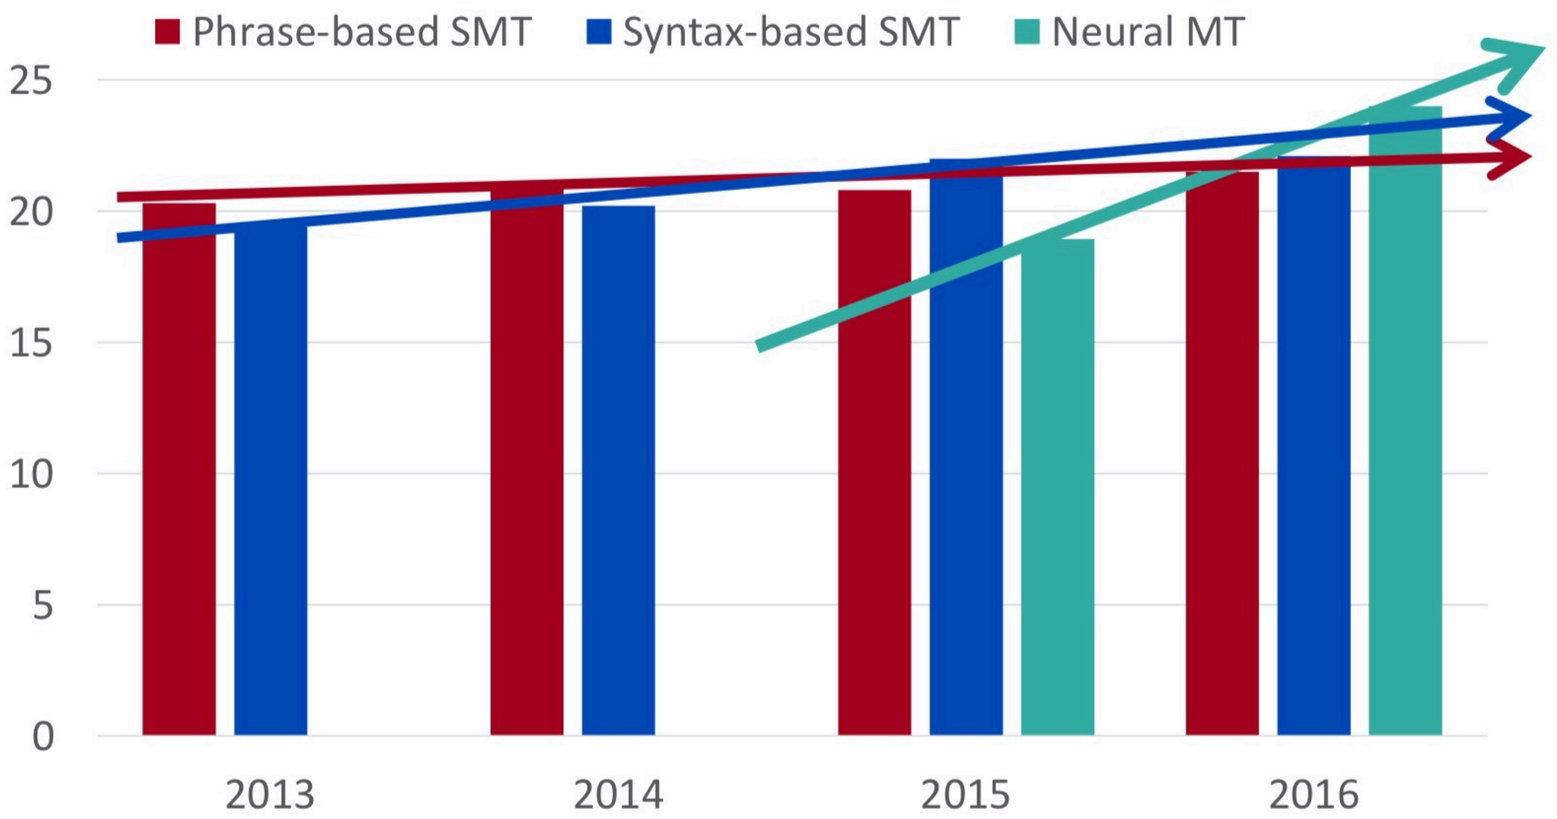

(source: Rico Sennrich)

RESEARCH › PUBLICATIONS ›

Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation

Table 10: Mean of side-by-side scores on production data

	PBMT	GNMT	Human	Relative Improvement
English → Spanish	4.885	5.428	5.504	87%
English → French	4.932	5.295	5.496	64%
English → Chinese	4.035	4.594	4.987	58%
Spanish → English	4.872	5.187	5.372	63%
French → English	5.046	5.343	5.404	83%
Chinese → English	3.694	4.263	4.636	60%

(Wu et al., 2016)

Versatile seq2seq

- Seq2seq finds applications in many other tasks!
- Any task where inputs and outputs are sequences of words/ characters
  - Summarization (input text → summary)
  - Dialogue (previous utterance → reply)
  - Parsing (sentence → parse tree in sequence form)
  - Question answering (context+question → answer)

Issues with vanilla seq2seq

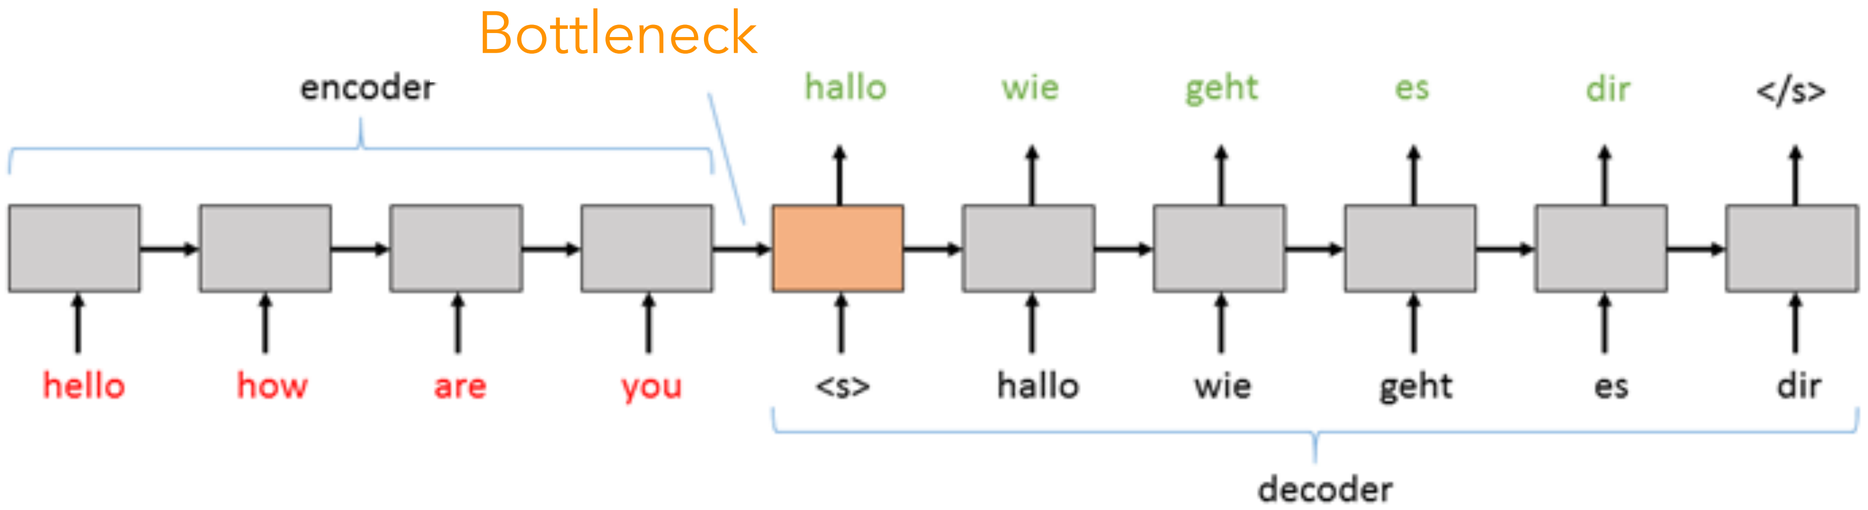

- A single encoding vector, h^{enc}, needs to capture all the information about source sentence
- Longer sequences can lead to vanishing gradients
- Model may “overfit” to training sequences

Issues with vanilla seq2seq

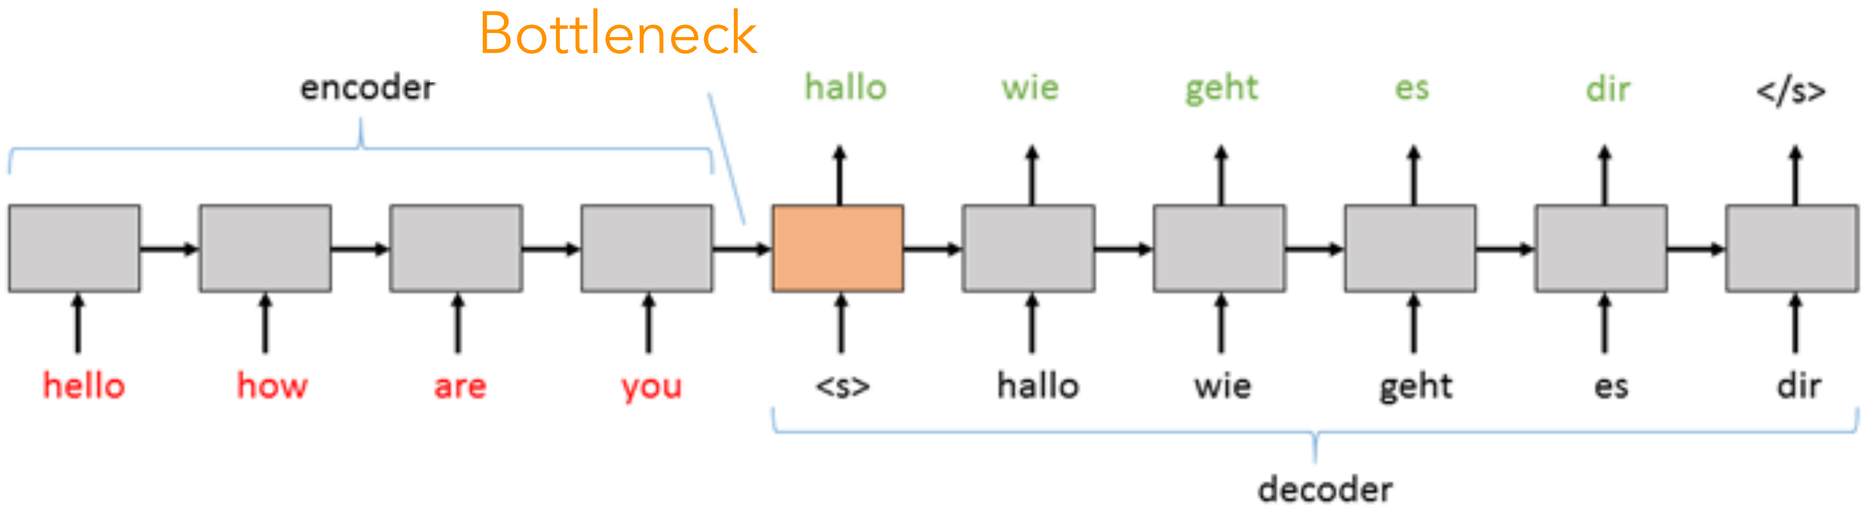

- A single encoding vector, h^{enc}, needs to capture all the information about source sentence
- **Longer sequences can lead to vanishing gradients**
- Model may "overfit" to training sequences

Remember alignments?

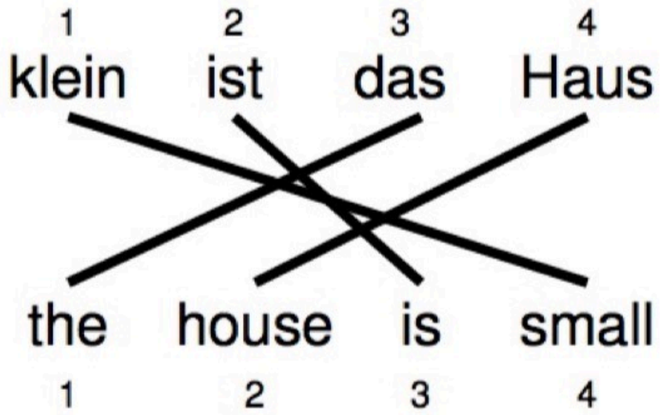

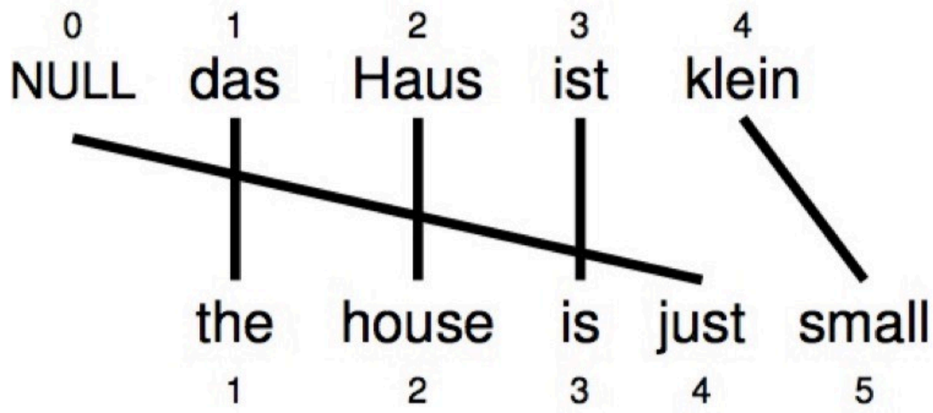

$$\mathbf{a} = (3, 4, 2, 1)^\top$$

$$\mathbf{a} = (1, 2, 3, 0, 4)^\top$$

Attention

- The neural MT equivalent of alignment models
- Key idea: At each time step during decoding, **focus on a particular part** of source sentence
  - This depends on the decoder's current hidden state h^{dec} (i.e. an idea of what you are trying to decode)
  - Usually implemented as a probability distribution over the hidden states of the encoder (h_i^{enc})

Seq2seq with attention

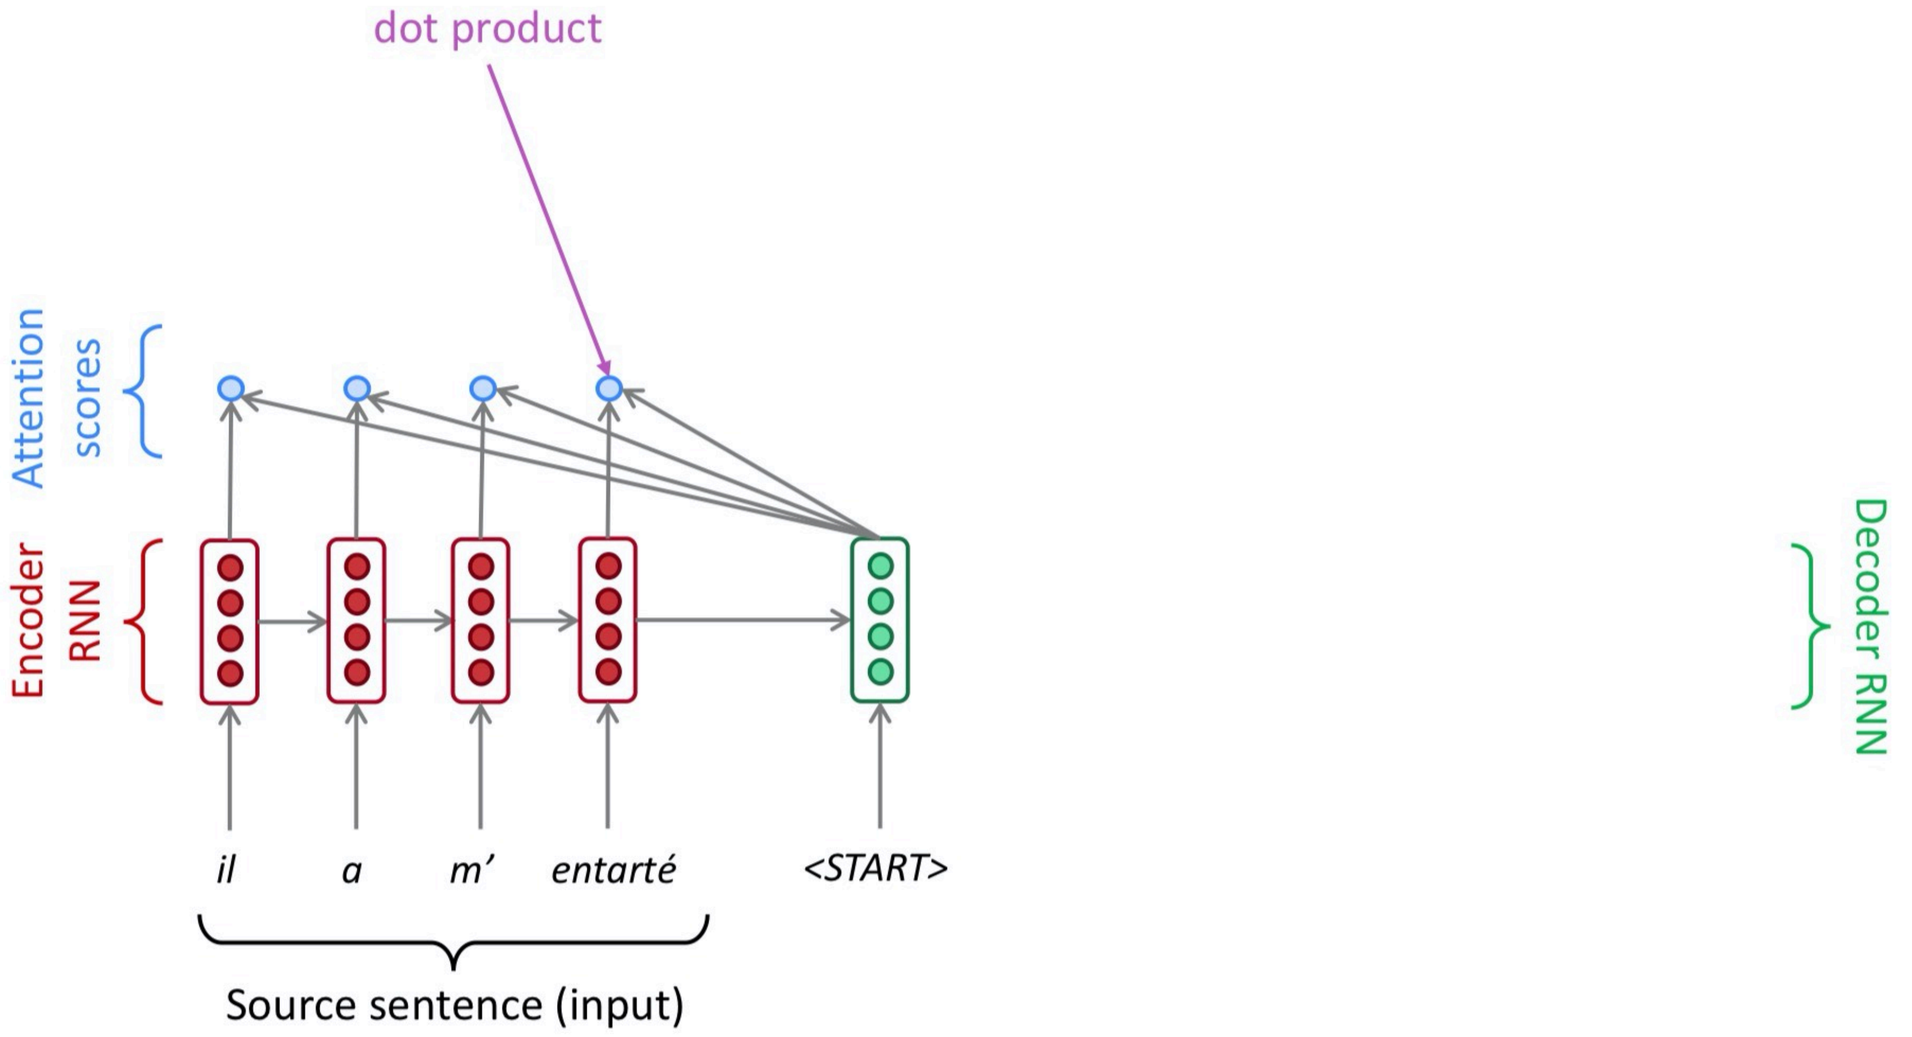

(slide credit: Abigail See)

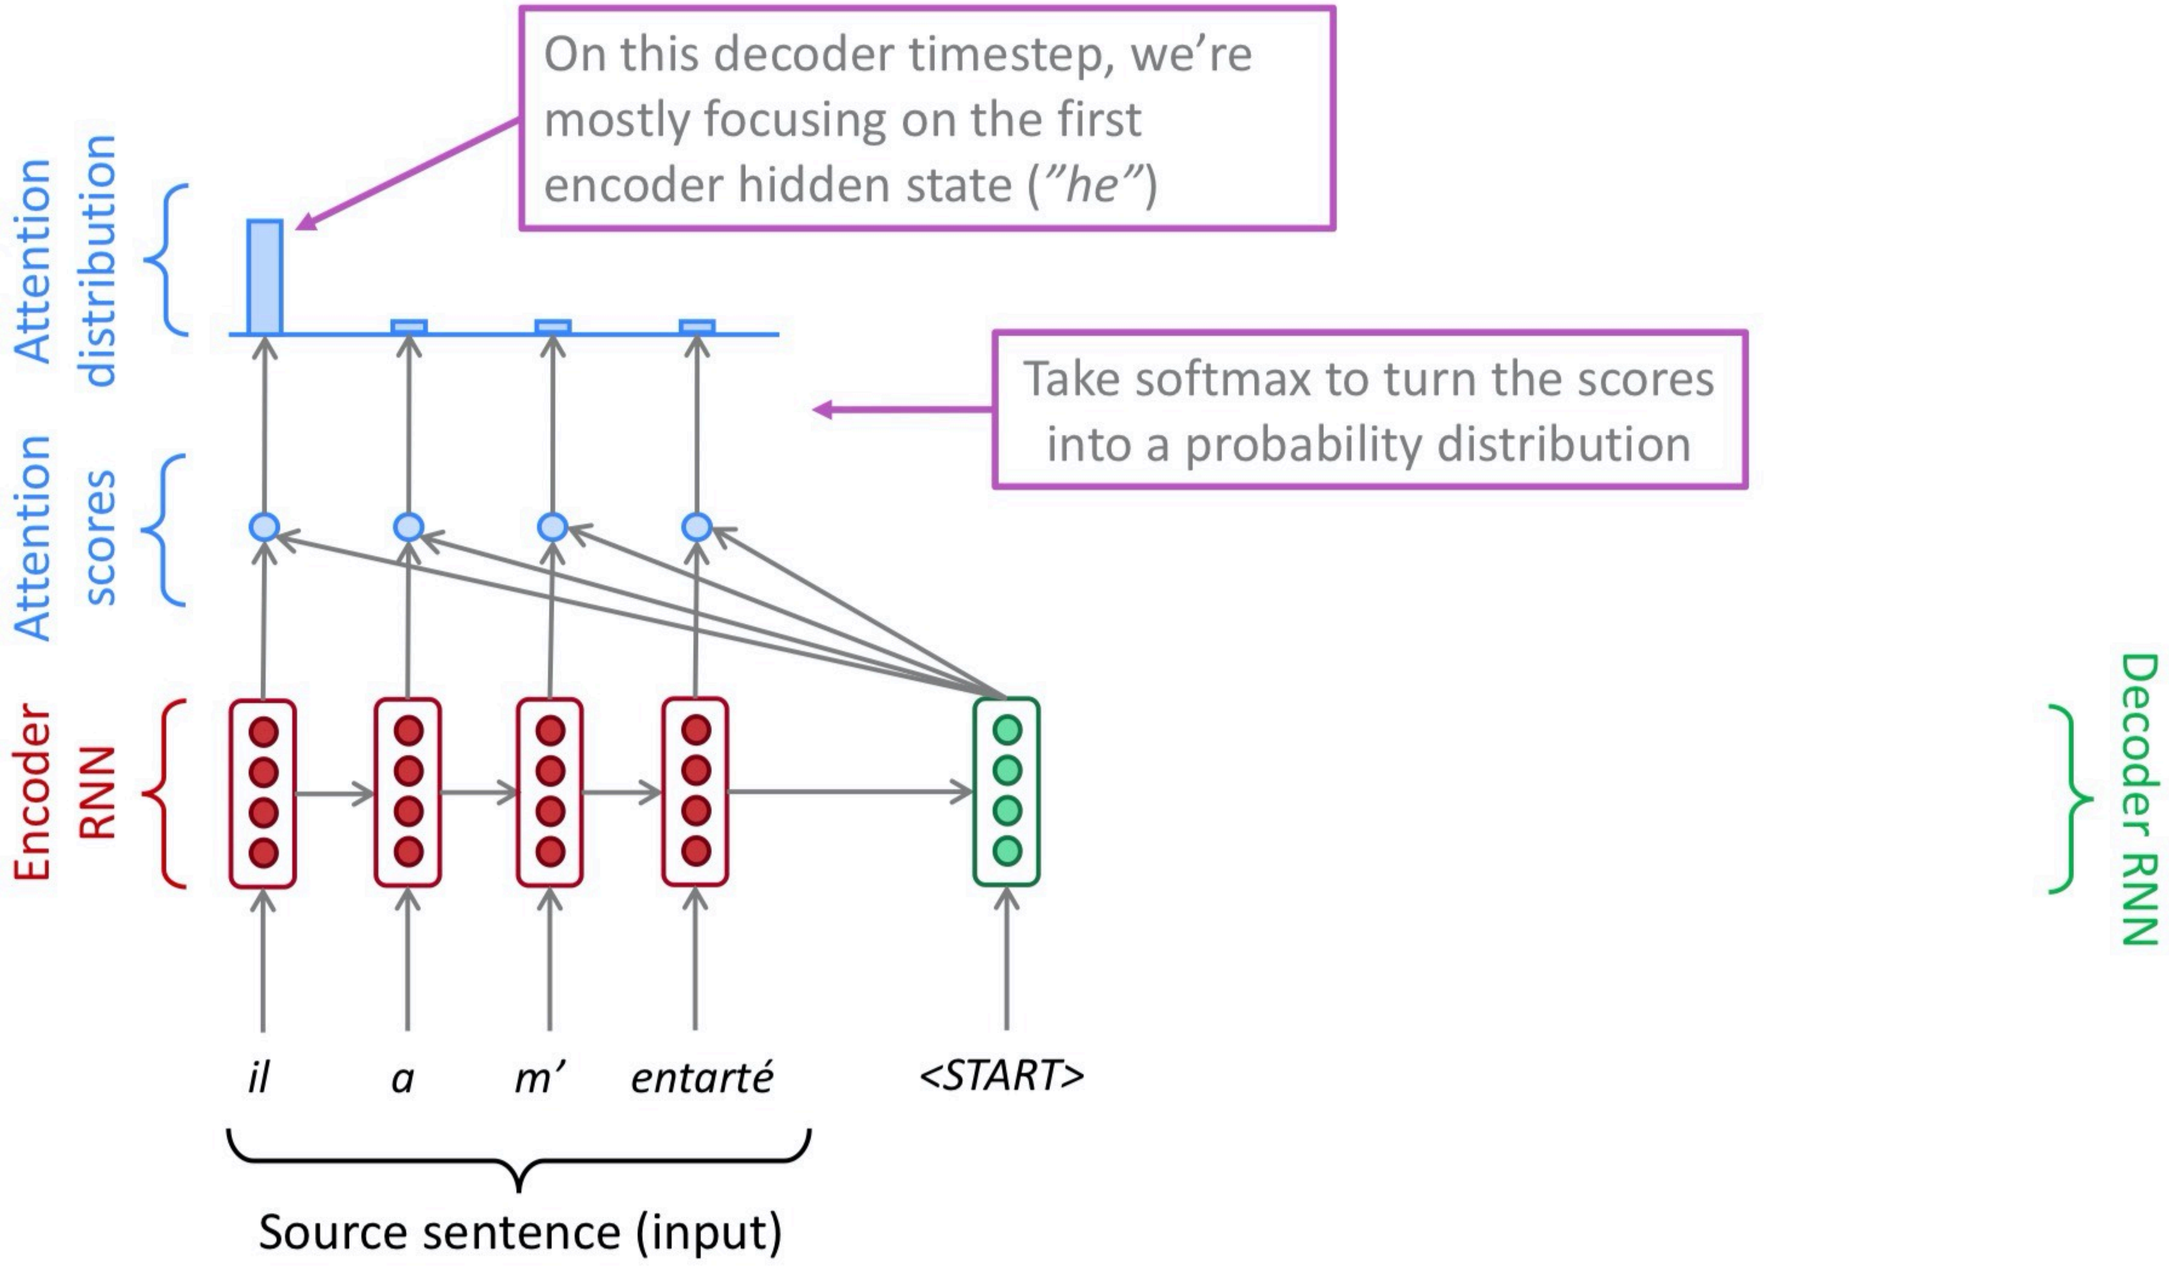
On this decoder timestep, we're mostly focusing on the first encoder hidden state ("*he*")
Take softmax to turn the scores into a probability distribution
Attention distribution
Attention scores
Encoder RNN
Decoder RNN
il
a
m'
entarté
<START>
Source sentence (input)

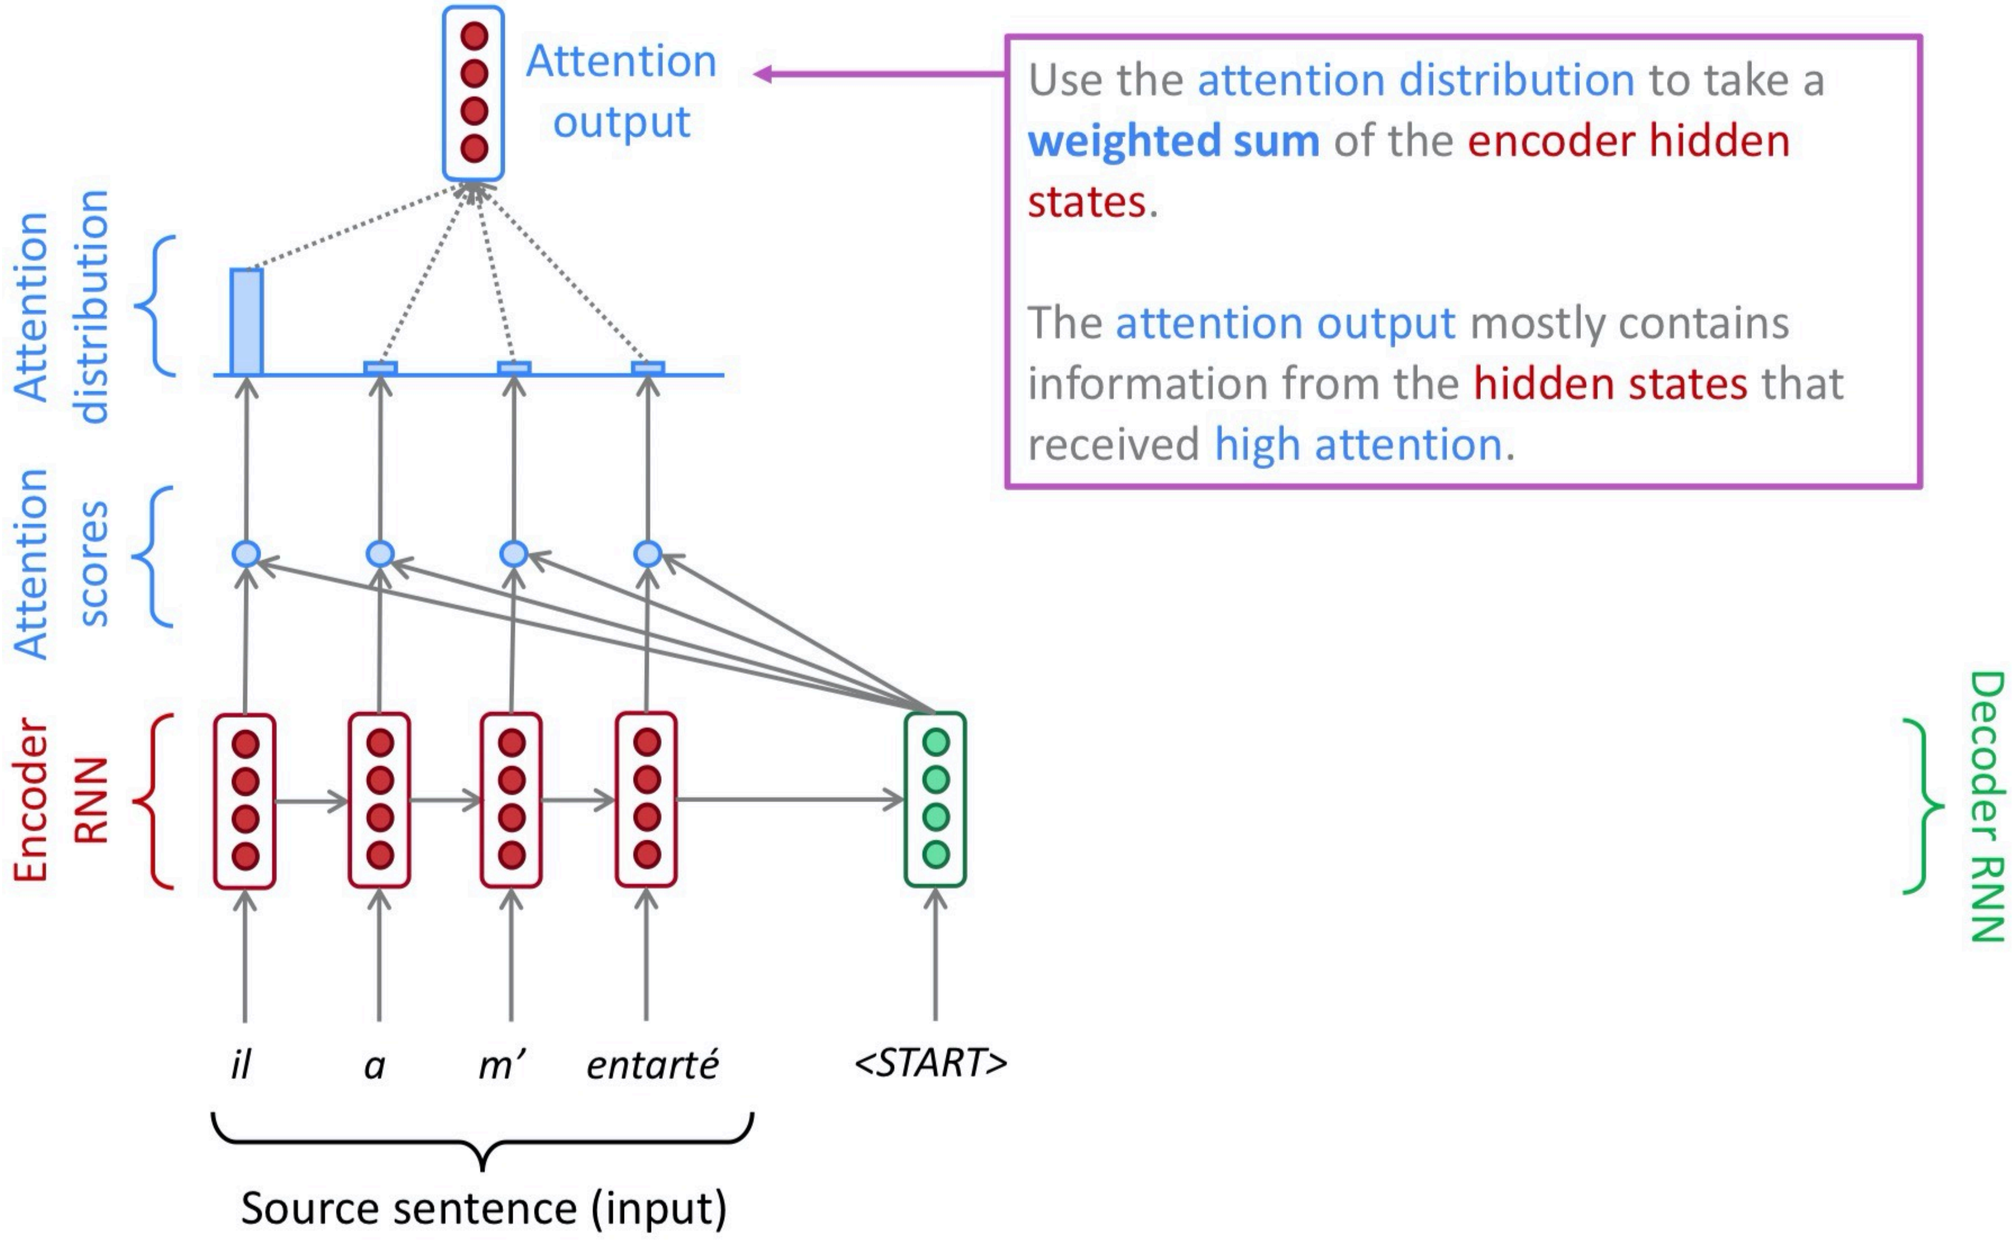

Use the attention distribution to take a **weighted sum** of the encoder hidden states.

The attention output mostly contains information from the hidden states that received high attention.

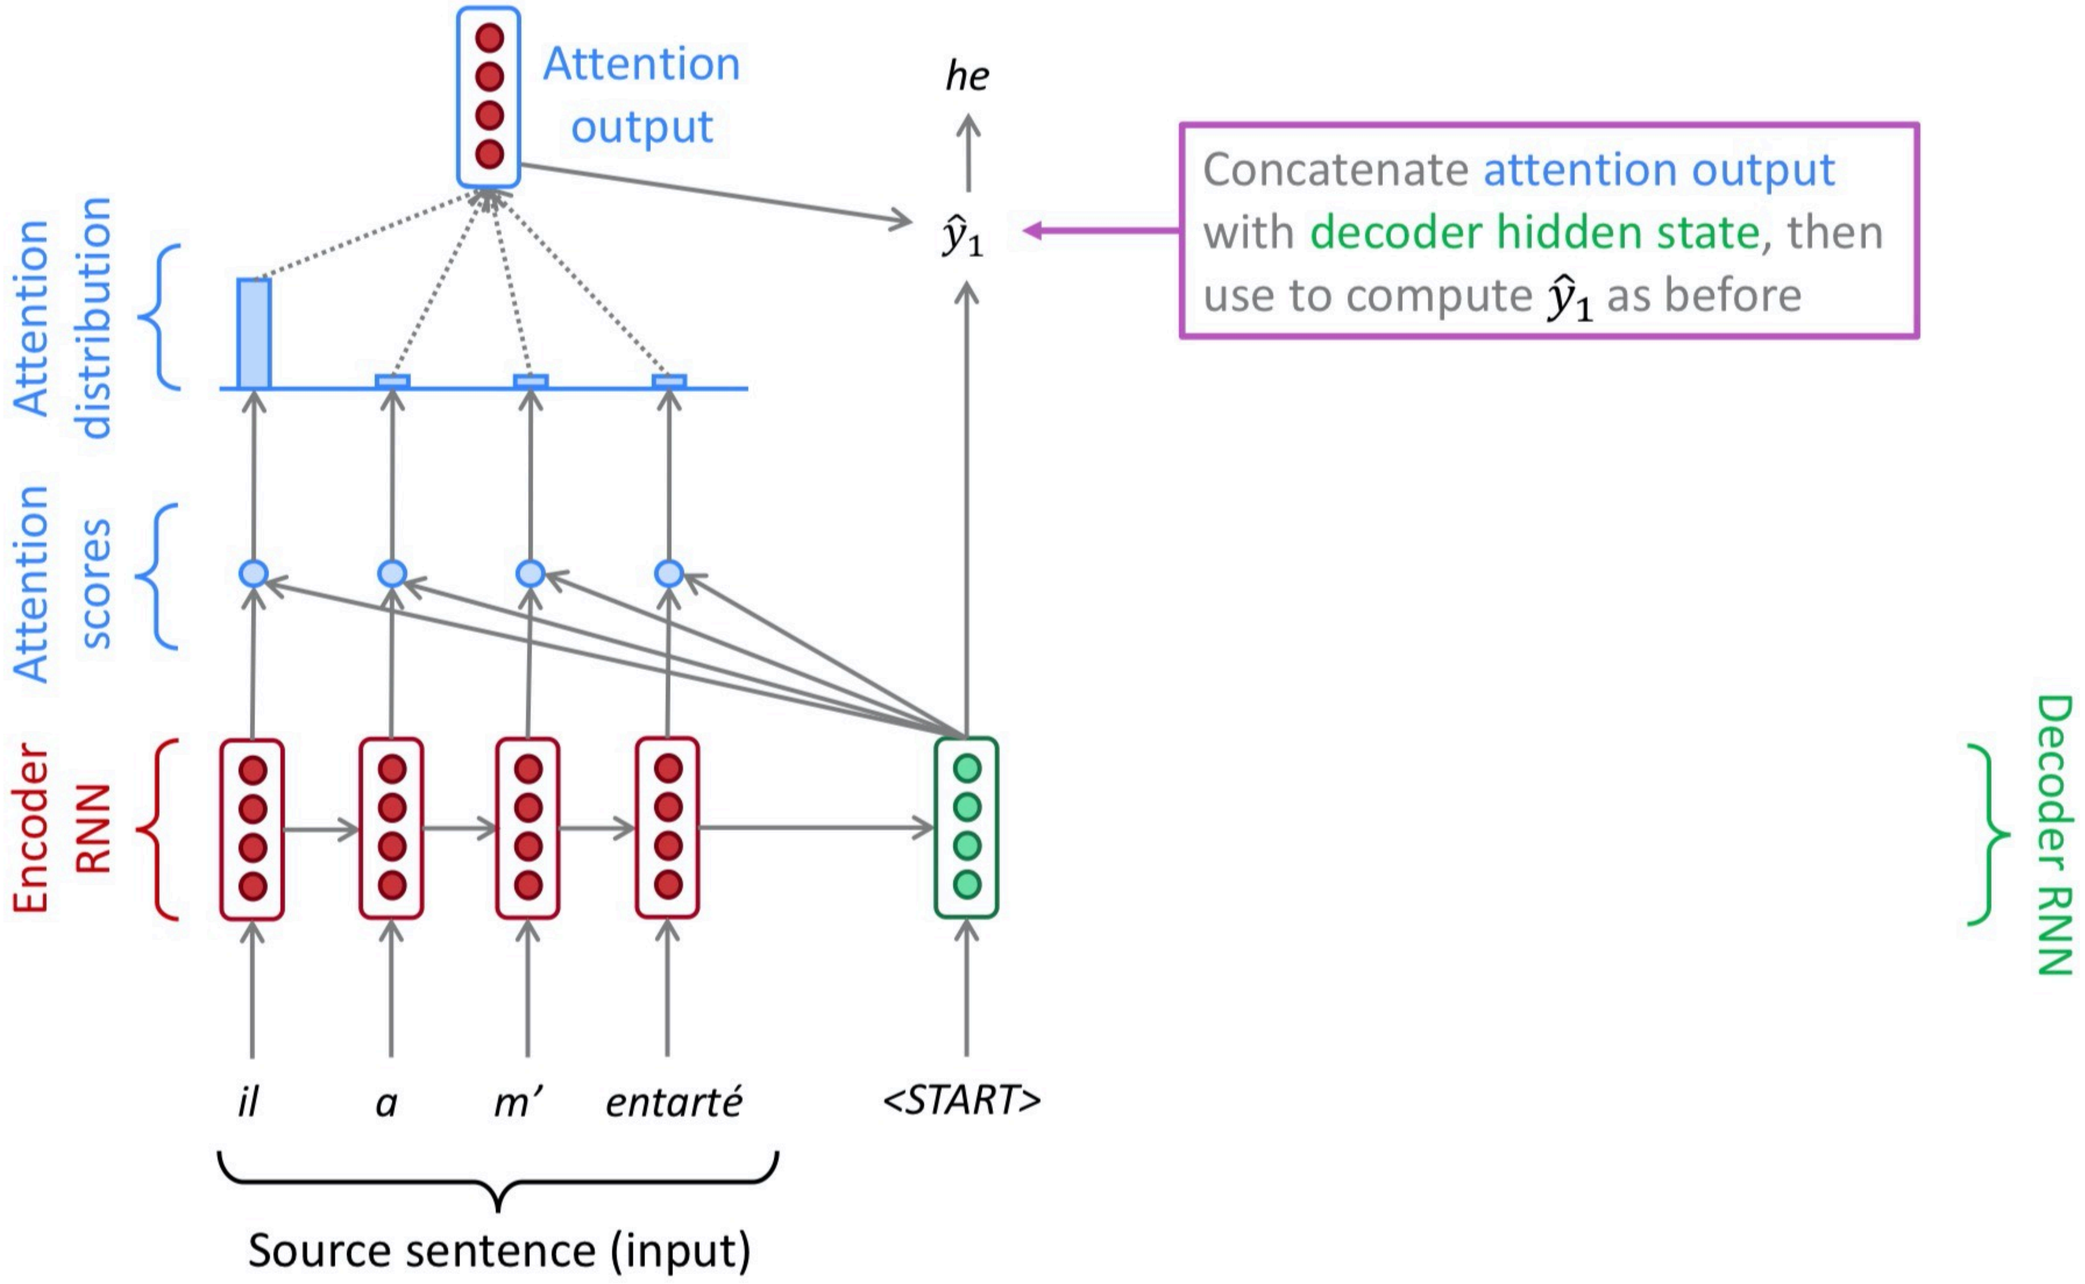
Attention output
he
$\hat{y}_1$
Concatenate attention output with decoder hidden state, then use to compute $\hat{y}_1$ as before
Attention distribution
Attention scores
Encoder RNN
Decoder RNN
il
a
m'
entarté
<START>
Source sentence (input)

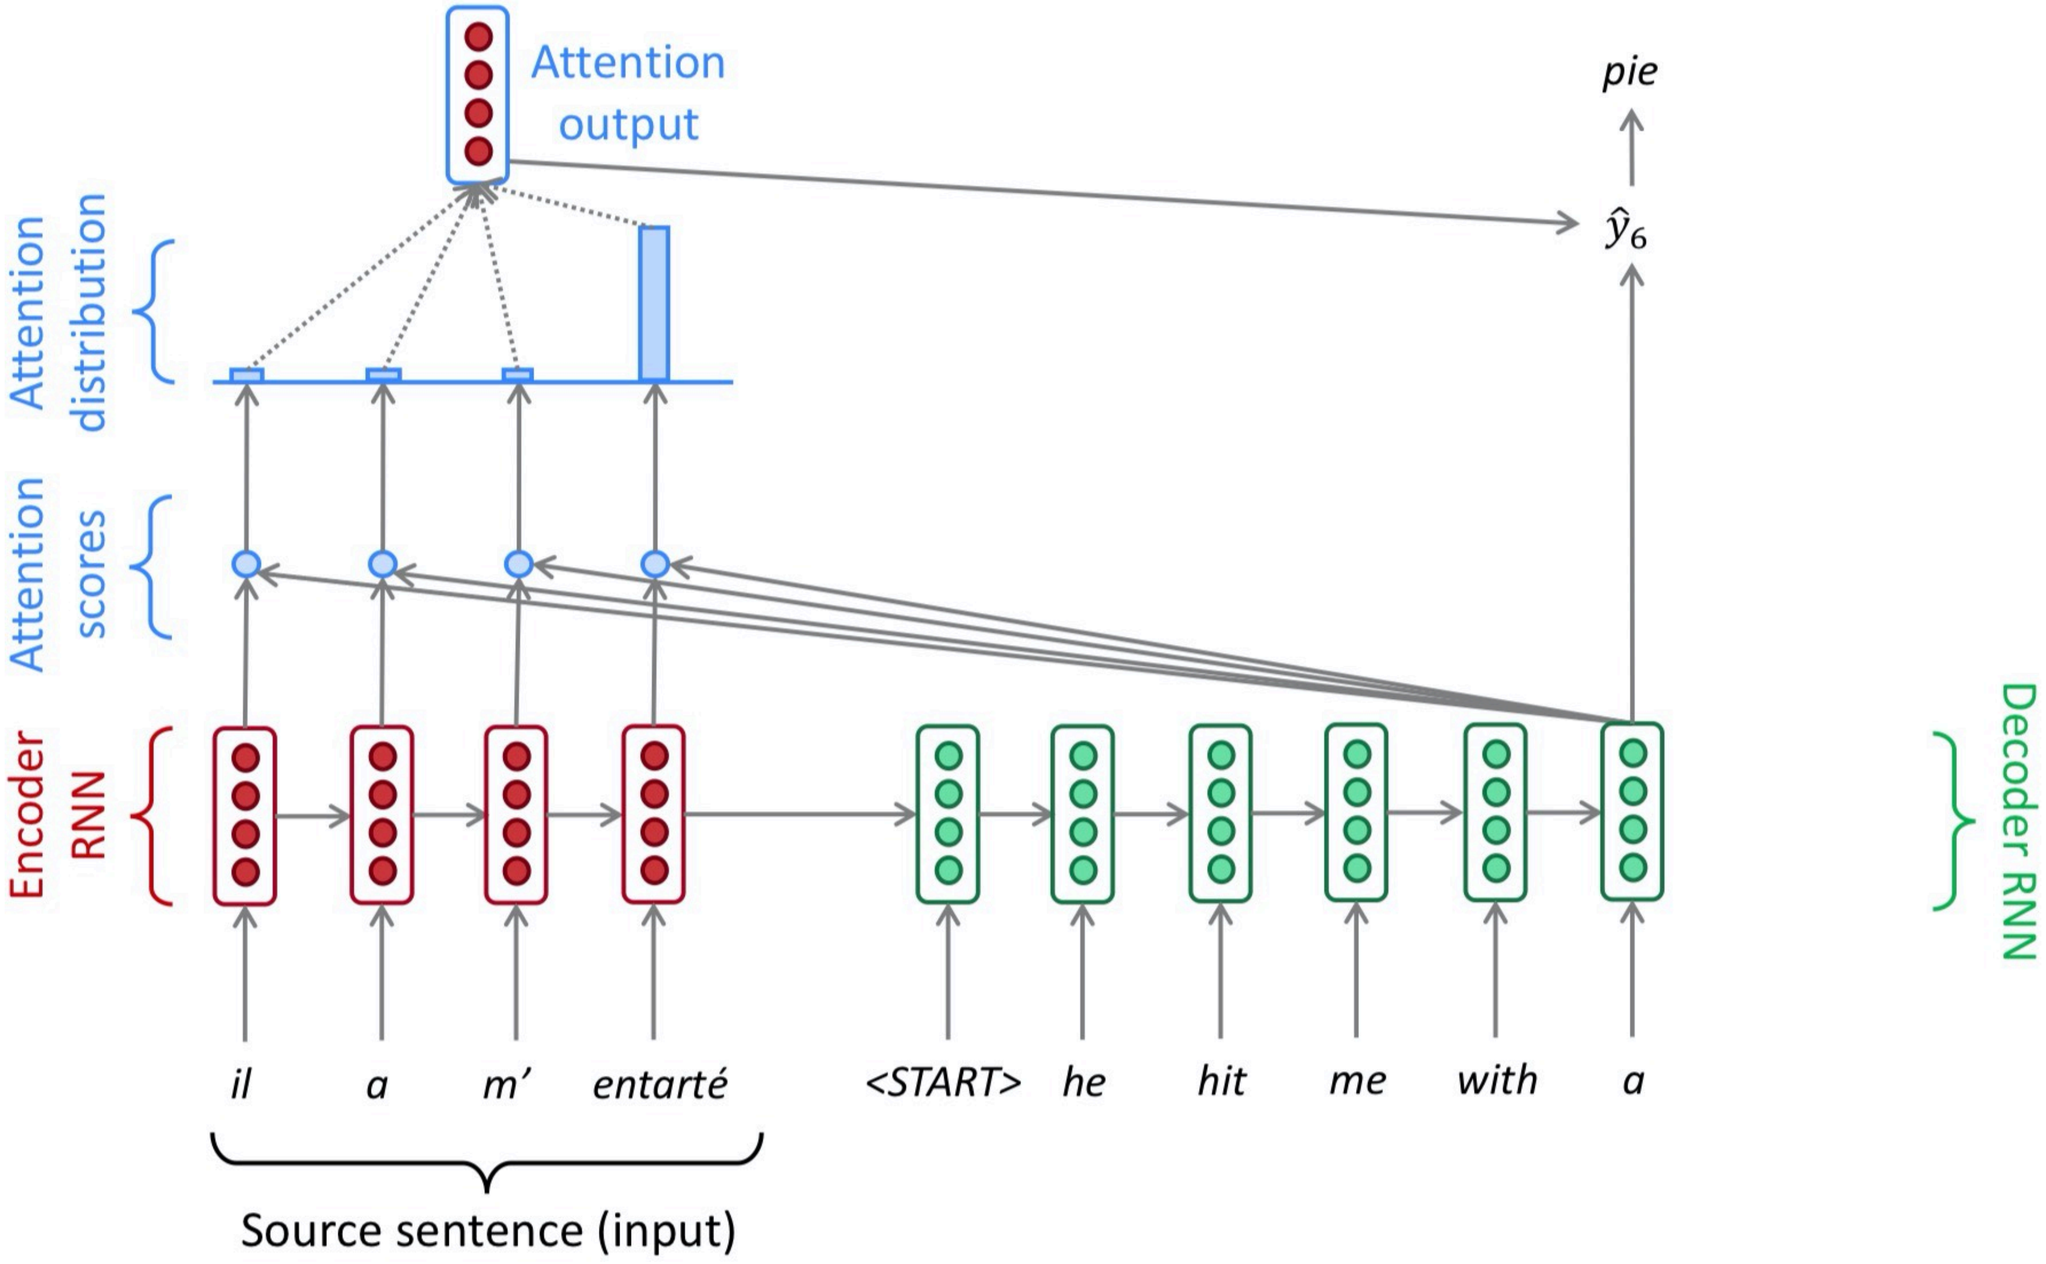
Attention output
pie
$\hat{y}_6$
Attention distribution
Attention scores
Encoder RNN
Decoder RNN
il
a
m'
entarté
<START>
he
hit
me
with
a
Source sentence (input)

Computing attention

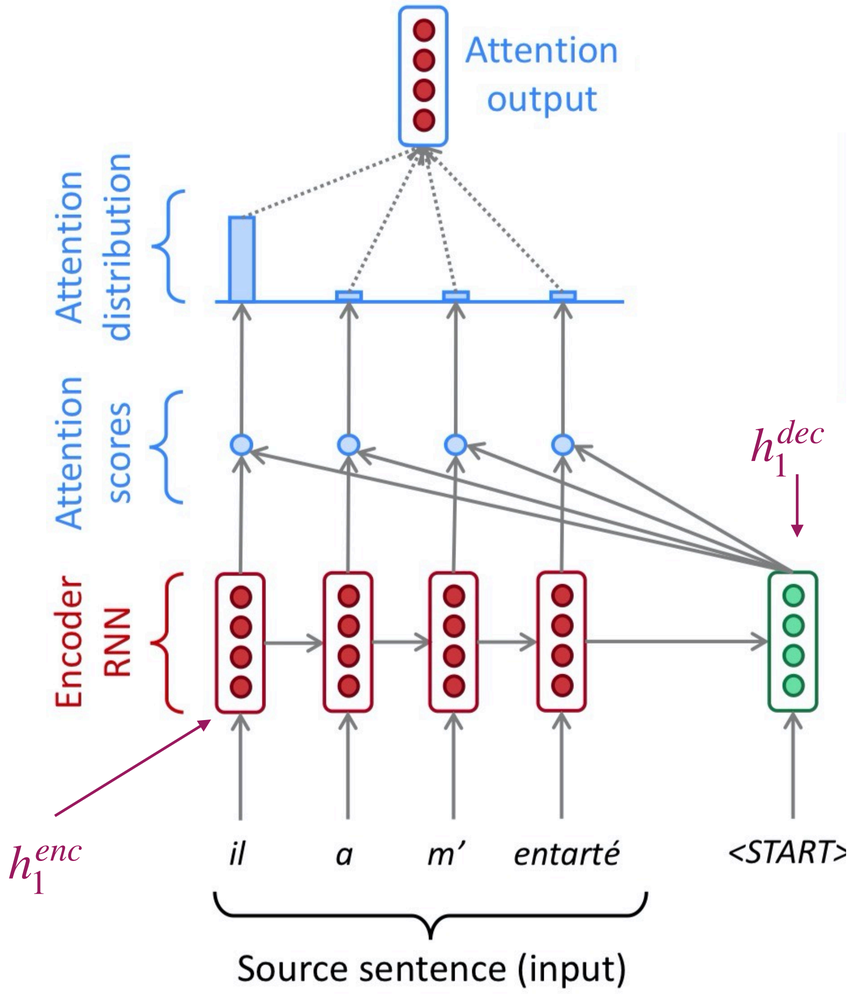

- Encoder hidden states: $h_1^{enc}, \ldots, h_n^{enc}$
- Decoder hidden state at time t: h_t^{dec}
- First, get attention scores for this time step of decoder (we'll define g soon):

$$e^t = [g(h_1^{enc}, h_t^{dec}), \ldots, g(h_n^{enc}, h_t^{dec})]$$

- Obtain the attention distribution using softmax:

$$\alpha^t = \text{softmax}\,(e^t) \in \mathbb{R}^n$$

- Compute weighted sum of encoder hidden states:

$$a_t = \sum_{i=1}^{n} \alpha_i^t h_i^{enc} \in \mathbb{R}^h$$

- Finally, concatenate with decoder state and pass on to output layer:

$$[a_t; h_t^{dec}] \in \mathbb{R}^{2h}$$

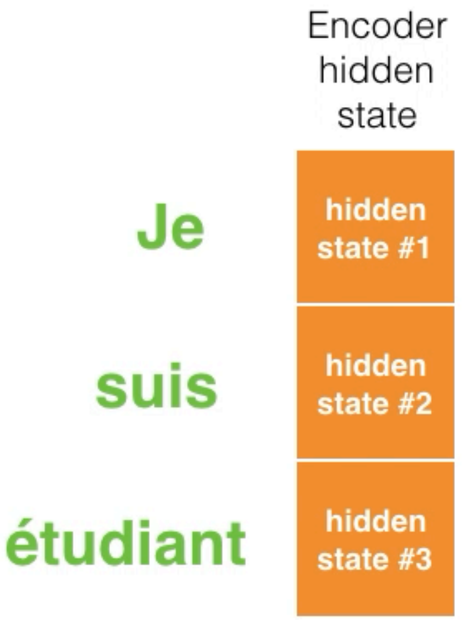

(credits: Jay Alammar)

Types of attention

- Assume encoder hidden states $h_1^{enc}, h_2^{enc}, \ldots, h_n^{enc}$ and a decoder hidden state h^{dec}

1. **Dot-product attention** (assumes equal dimensions for h^{enc} and h^{dec}):

$$e_i = g(h_i^{enc}, h^{dec}) = (h^{dec})^T \; h_i^{enc} \in \mathbb{R}$$

2. **Multiplicative attention:**

$$g(h_i^{enc}, h^{dec}) = (h^{dec})^T \; W \; h_i^{enc} \in \mathbb{R}, \text{ where } W \text{ is a weight matrix (learned)}$$

3. **Additive attention:**

$$g(h_i^{enc}, h^{dec}) = v^T \tanh\left(W_1 h_i^{enc} + W_2 h^{dec}\right) \in \mathbb{R}$$

where W_1, W_2 are weight matrices (learned) and v is a weight vector (learned)

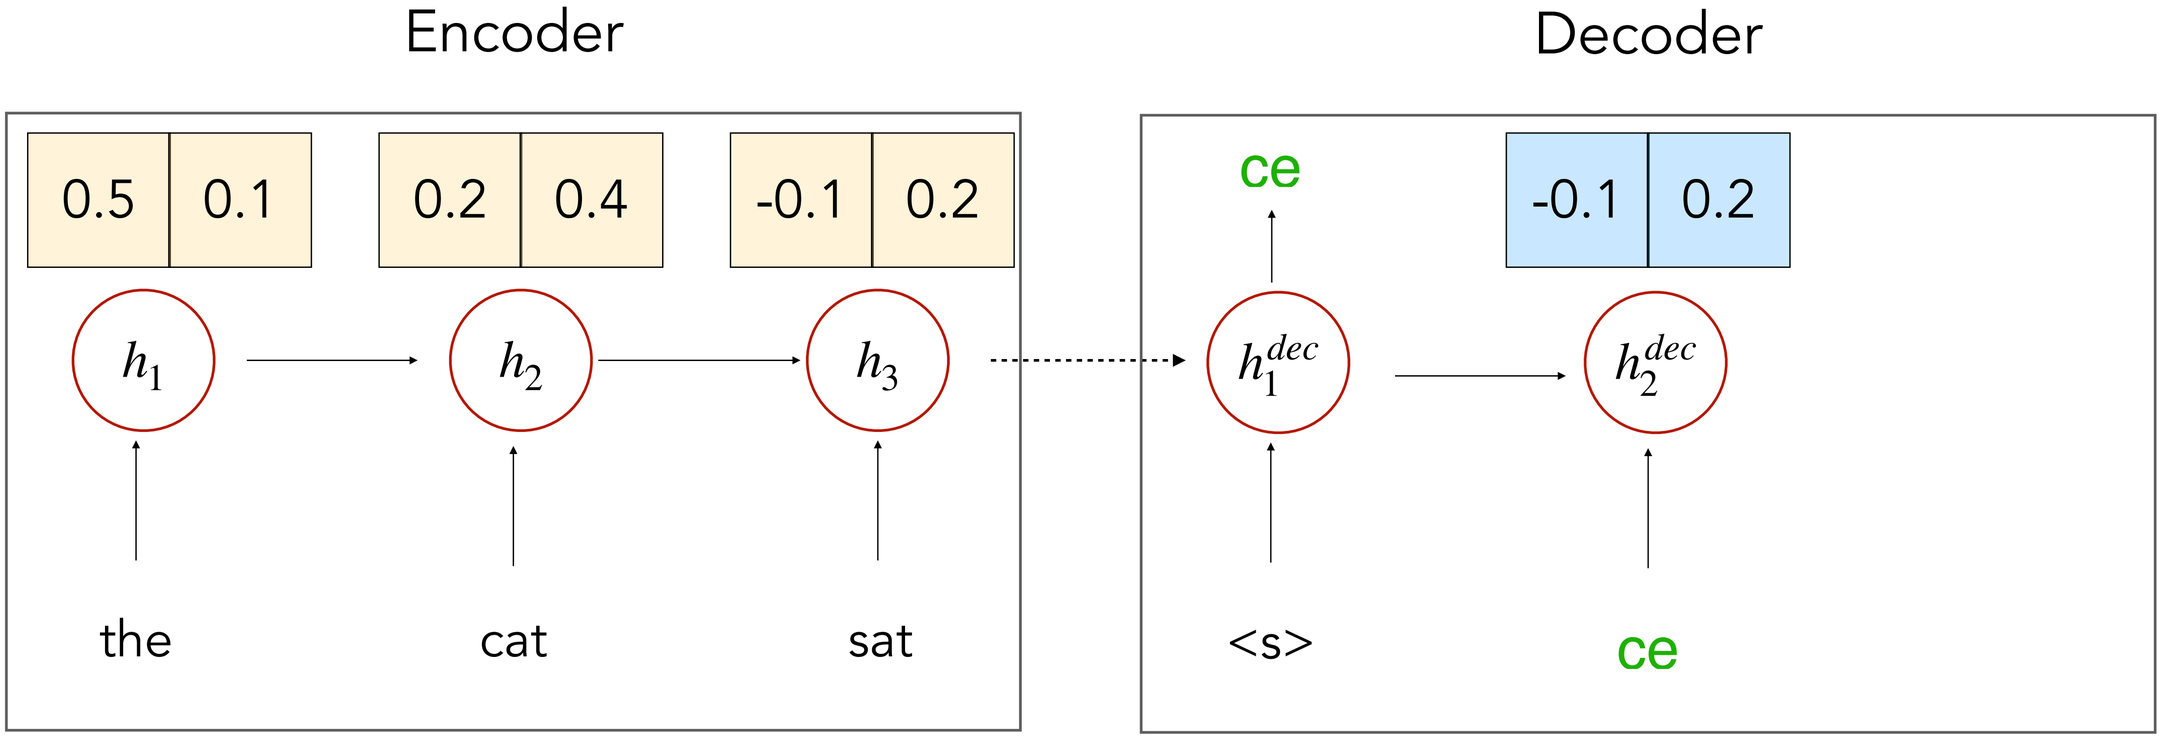

Dot-product attention:

$$g(h_i^{enc}, h^{dec}) = h^{dec} \cdot h^{enc}$$

Assuming we use dot product attention, which input word will have the highest attention value at current time step?

A) the
B) cat
C) sat

the -> -0.05 + 0.02
cat -> -0.02 + 0.08
sat -> 0.01 + 0.04

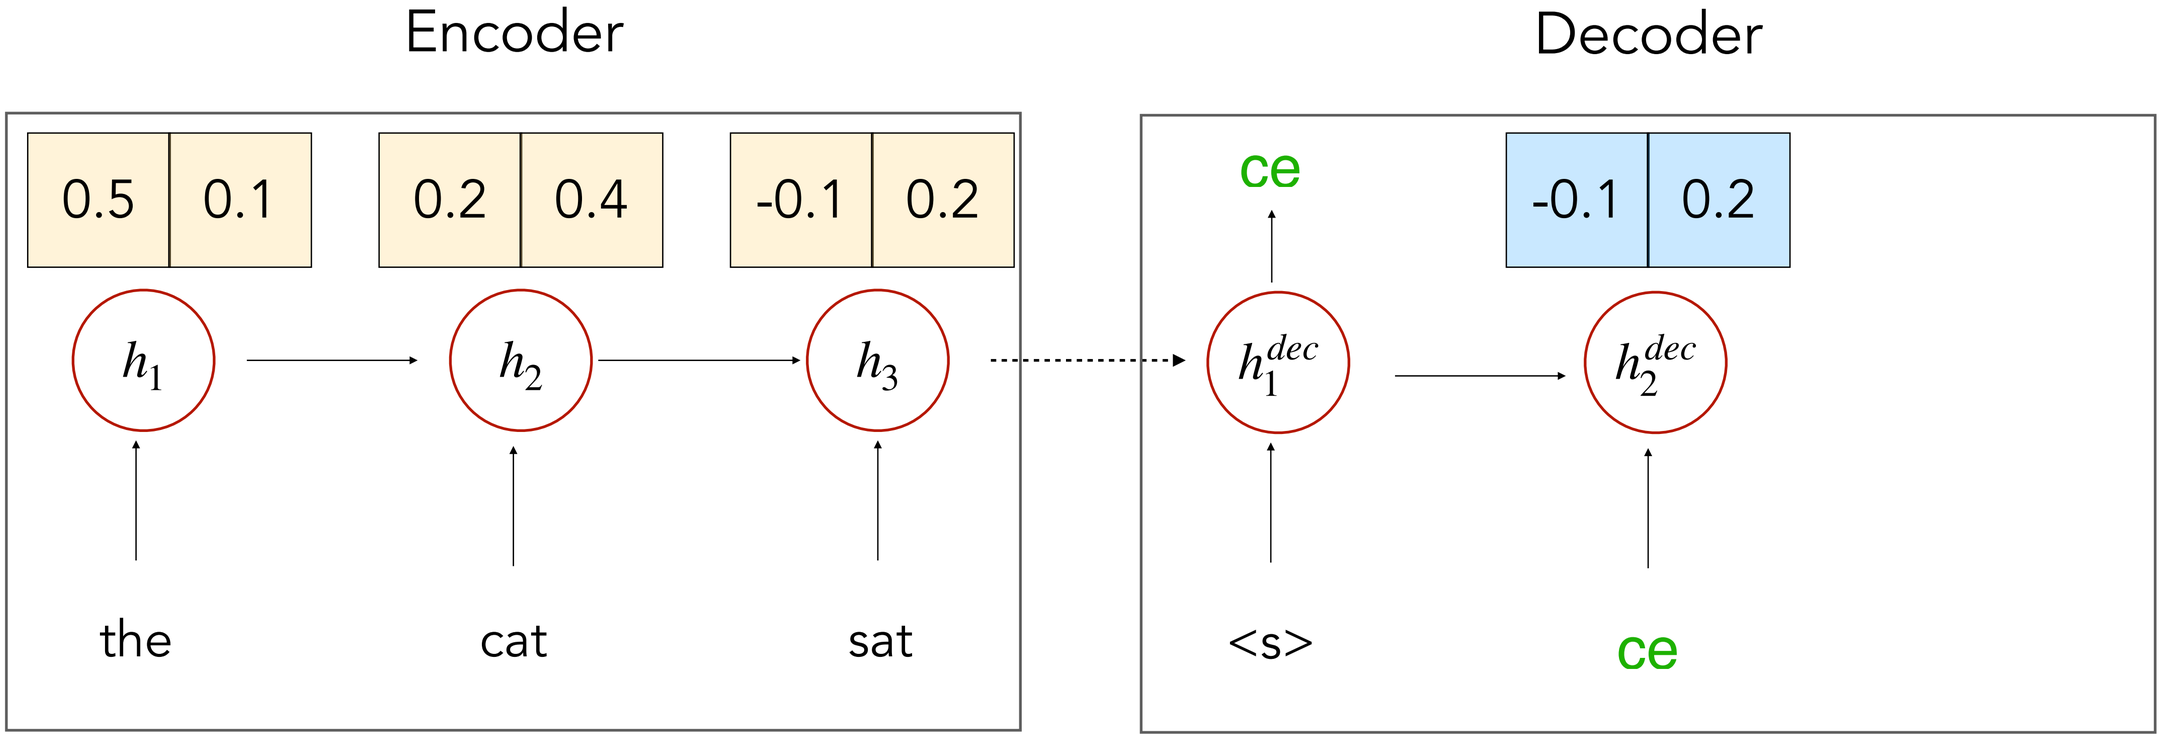

What if we use multiplicative attention with $W = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$?

Which input word will have the highest attention value at current time step?

A) the
B) cat
C) sat

the -> -0.05
cat -> -0.02
sat -> 0.01

Multiplicative attention:

$g(h_i^{enc}, h^{dec}) = (h^{dec})^T \; W \; h_i^{enc}$

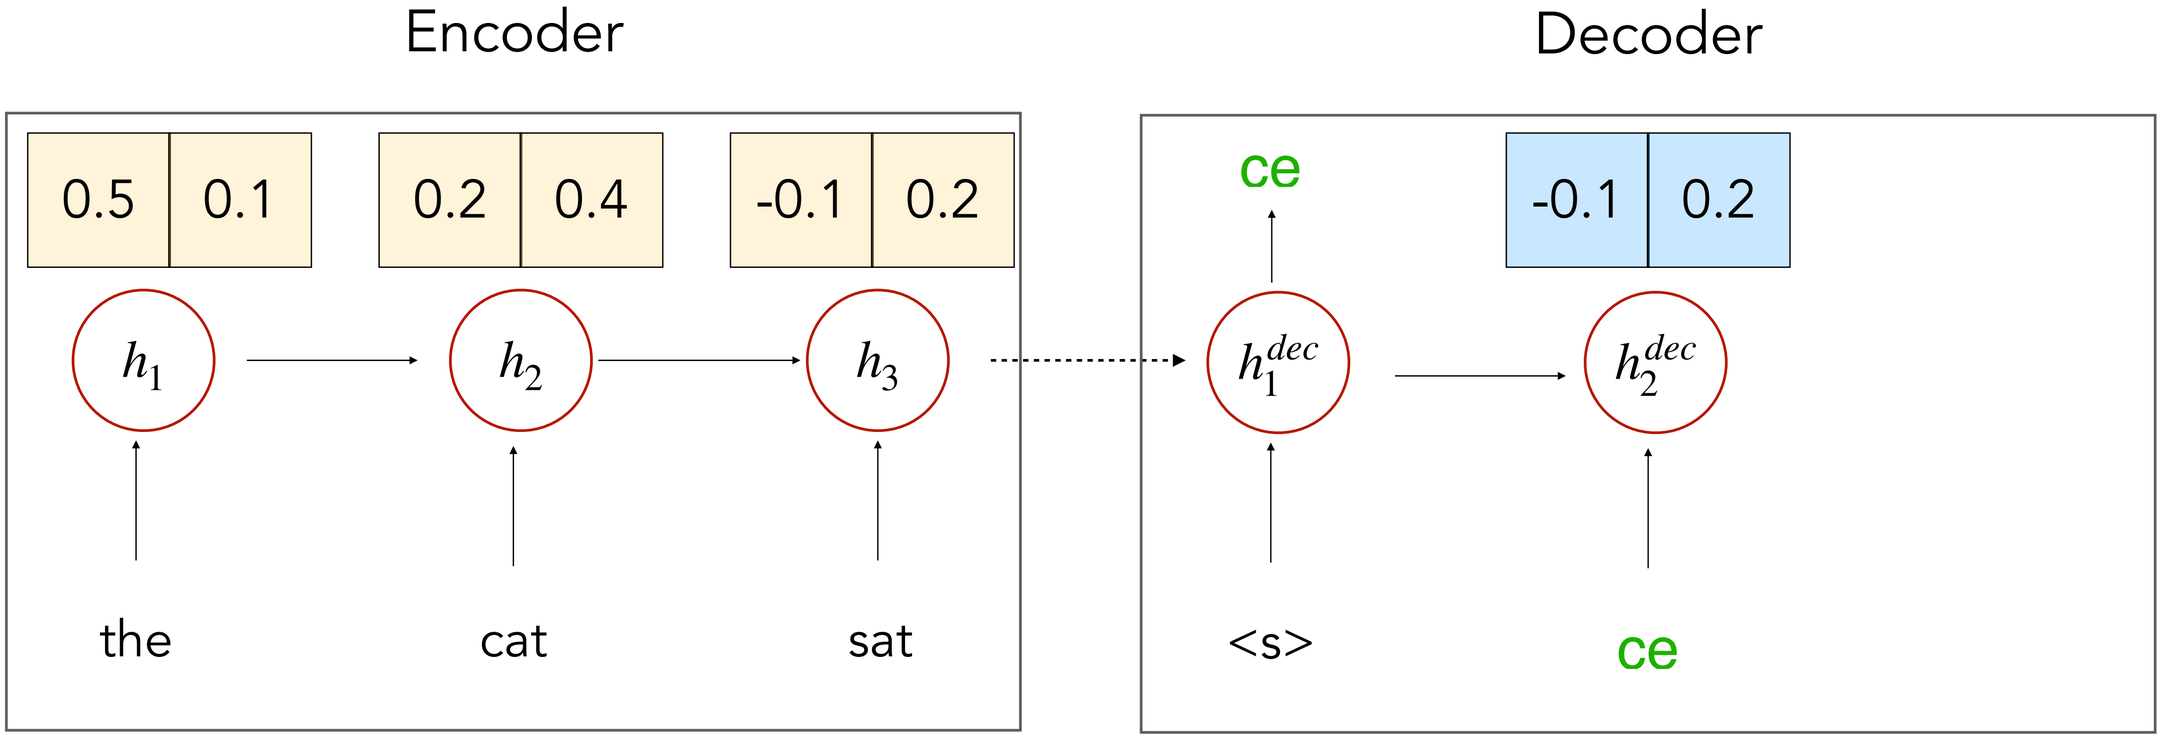

Which value of W in multiplicative attention will provide the same word with highest attention value as dot-product attention?

Multiplicative attention:

$g(h_i^{enc}, h^{dec}) = (h^{dec})^T \; W \; h_i^{enc}$

A) $W = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ B) $W = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.5 \end{bmatrix}$ C) both

Attention improves translation

System	Ppl	BLEU
Winning WMT'14 system – *phrase-based* + *large LM* (Buck et al., 2014)		20.7
Existing NMT systems		
RNNsearch (Jean et al., 2015)		16.5
RNNsearch + unk replace (Jean et al., 2015)		19.0
RNNsearch + unk replace + large vocab + *ensemble* 8 models (Jean et al., 2015)		**21.6**
Our NMT systems		
Base	10.6	11.3
Base + reverse	9.9	12.6 (+*1.3*)
Base + reverse + dropout	8.1	14.0 (+*1.4*)
Base + reverse + dropout + global attention (*location*)	7.3	16.8 (+*2.8*)
Base + reverse + dropout + global attention (*location*) + feed input	6.4	18.1 (+*1.3*)
Base + reverse + dropout + local-p attention (*general*) + feed input	5.9	19.0 (+*0.9*)
Base + reverse + dropout + local-p attention (*general*) + feed input + unk replace		20.9 (+*1.9*)
Ensemble 8 models + unk replace		**23.0** (+*2.1*)

(Luong et al., 2015)

Visualizing attention

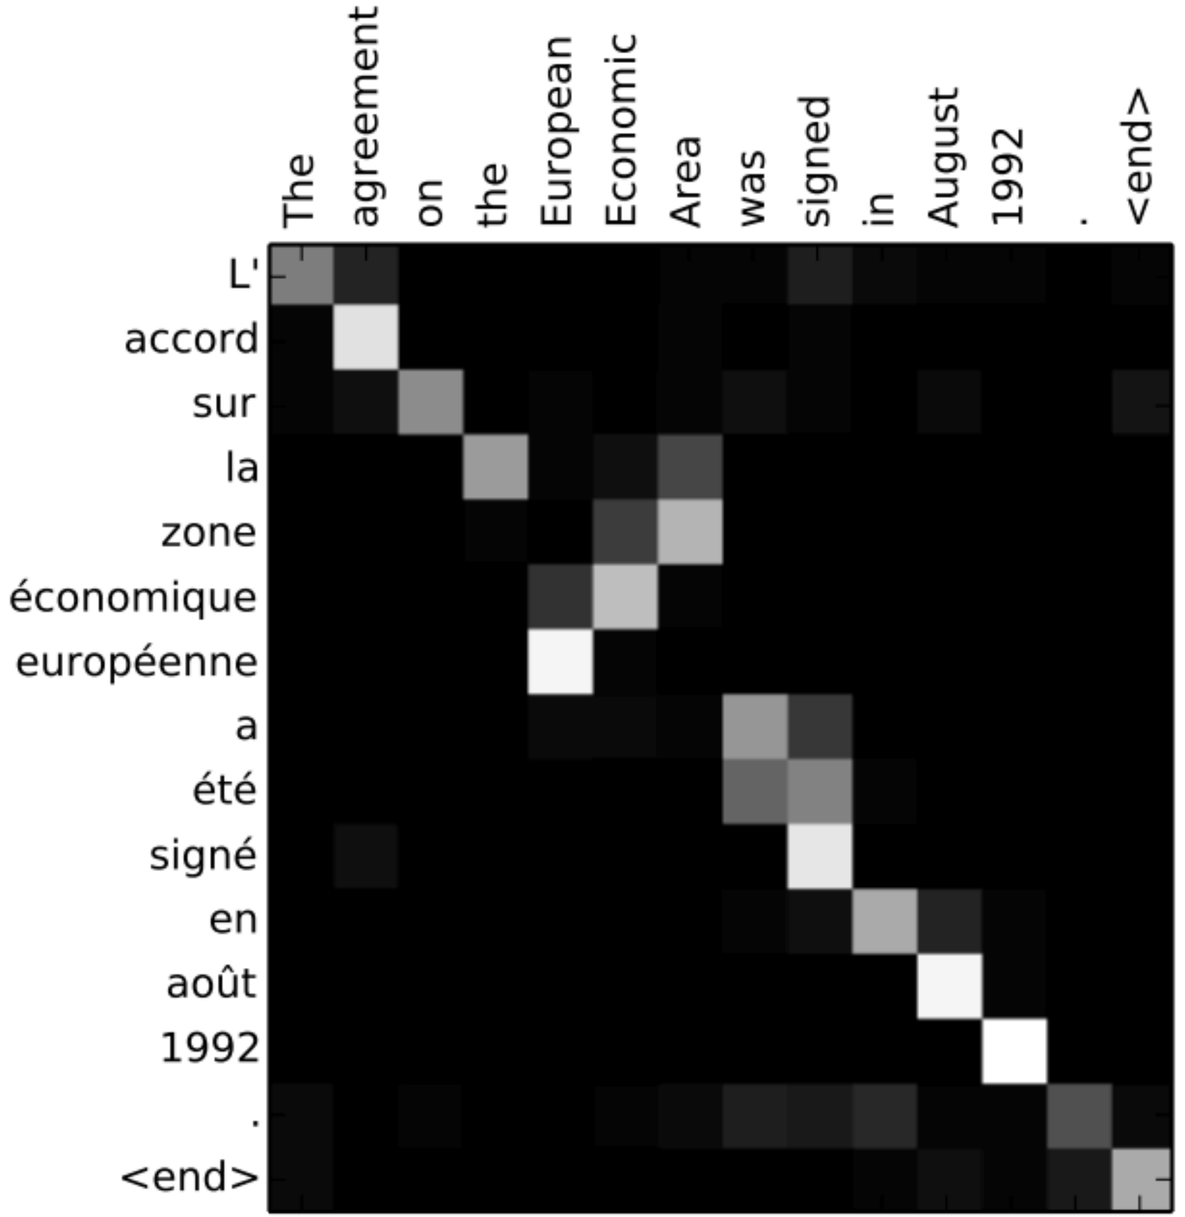

(credits: Jay Alammar)

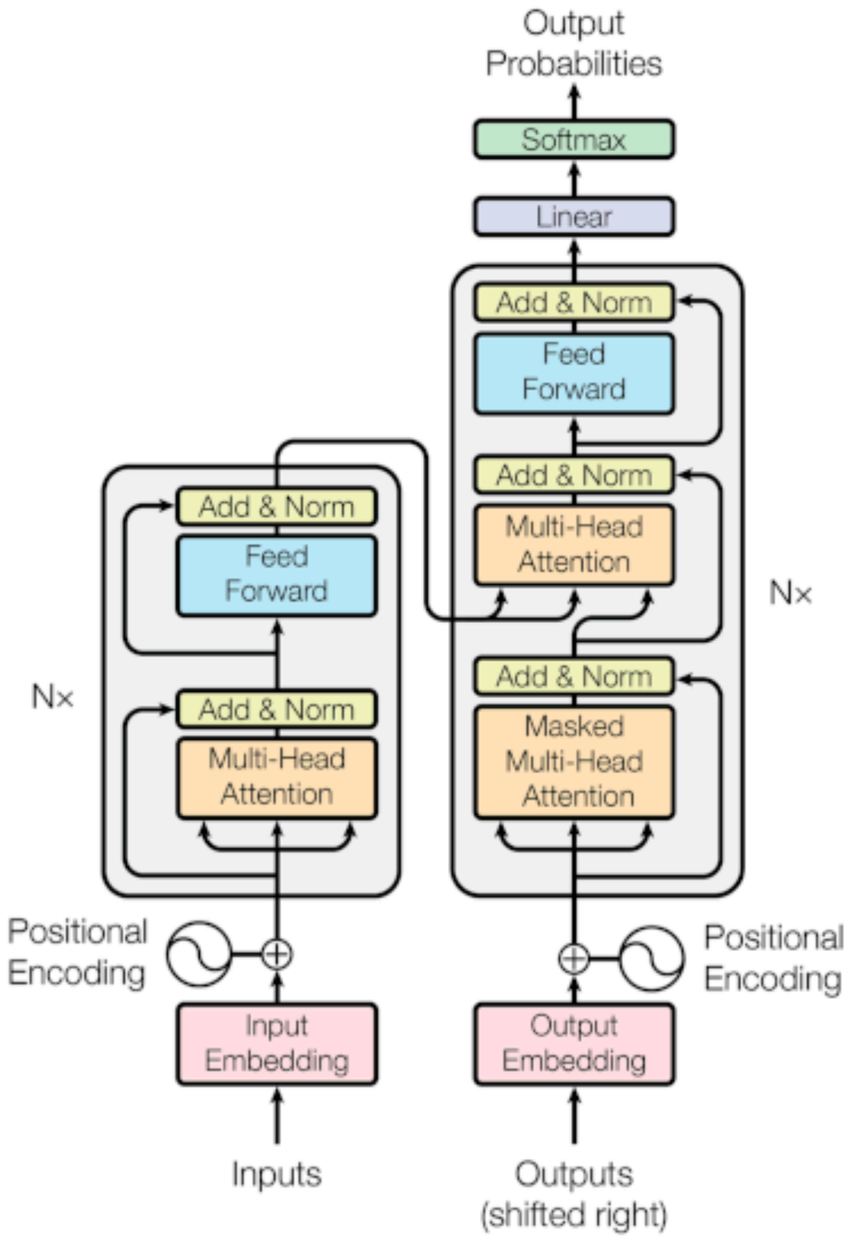

Going all in on attention

- More recent models (e.g. Transformer, Vaswani et al., 2017) have replaced RNNs entirely with attention mechanisms
- Theoretically limiting (since recurrence can help handle arbitrarily long sequences)
- Huge gains in practical performance

WMT 2014, English-German

Issues with vanilla seq2seq

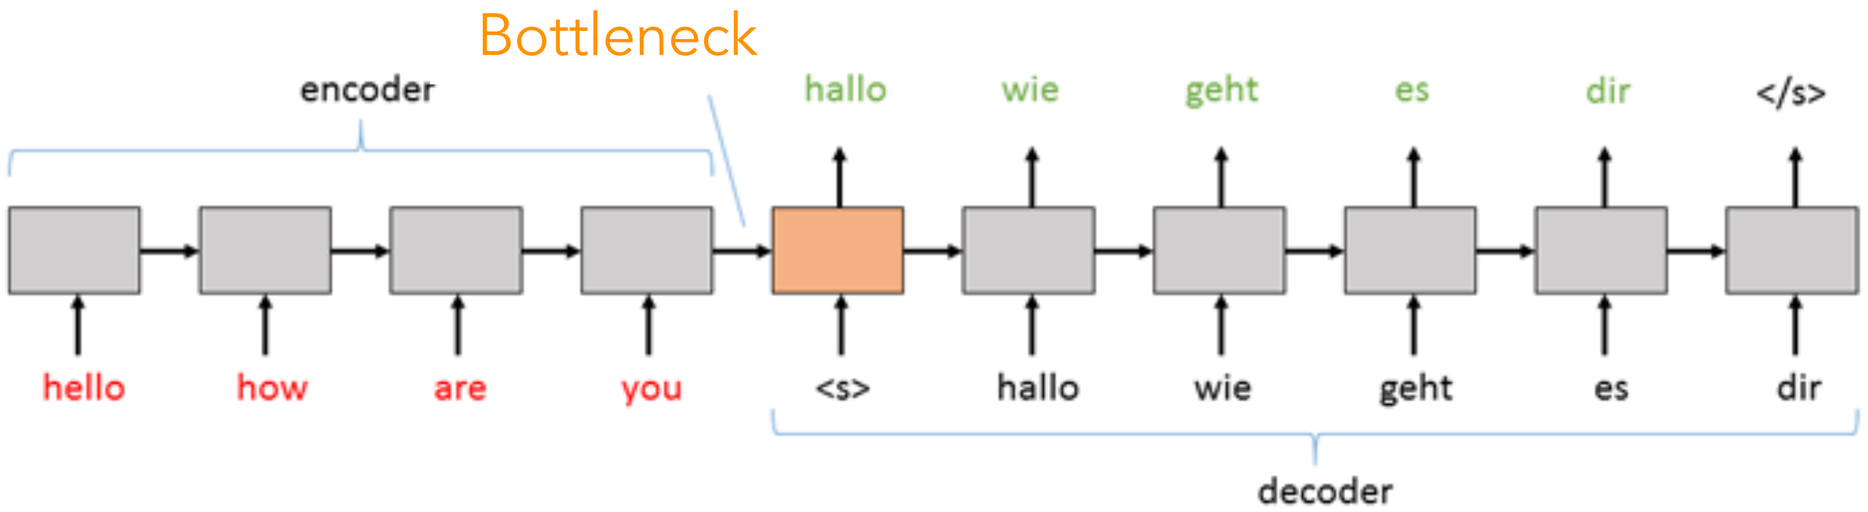

- A single encoding vector, h^{enc}, needs to capture all the information about source sentence
- Longer sequences can lead to vanishing gradients
- **Model may "overfit" to training sequences**

Dropout

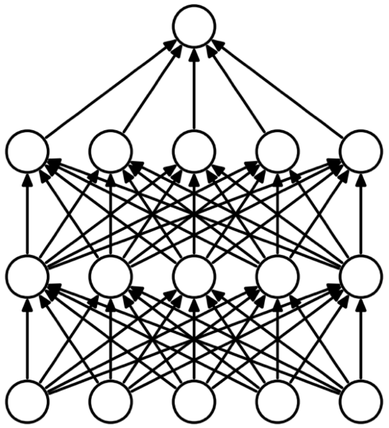

(a) Standard Neural Net

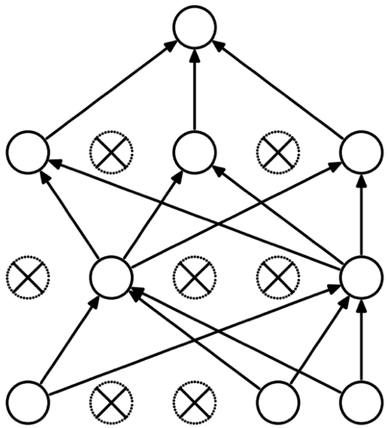

(b) After applying dropout.

- Form of regularization for RNNs (and any NN in general)
- Idea: "Handicap" NN by removing hidden units **stochastically**
  - set each hidden unit in a layer to 0 with probability p during training ($p = 0.5$ usually works well)
  - scale outputs by $1/(1-p)$
  - hidden units forced to learn more general patterns and improve redundancy
- Test time: Simply compute identity

(Srivastava et al., 2014)

ID	System	BLEU					
		5k	10k	20k	40k	80k	165k
1	Transformer-big	3.3	3.4	4.3	4.7	5.1	5.5
2	Transformer-base	8.3	11.9	16.8	23.2	28.0	32.1
3	2 + feed-forward dimension (2048 → 512)	8.8	12.0	16.7	22.3	27.7	31.7
4	3 + attention heads (8→ 2)	9.2	12.7	19.0	23.6	28.7	32.3
5	4 + dropout (0.1→ 0.3)	10.6	17.0	21.9	26.7	**31.0**	**33.4**
6	5 + layers (6 → 5)	10.9	16.9	21.9	26.0	30.2	33.0
7	6 + label smoothing (0.1→ 0.6)	11.3	16.5	22.0	26.9	30.4	33.3
8	7 + decoder layerDrop (0 → 0.3)	12.9	17.3	22.5	26.9	30.3	33.1
9	8 + target word dropout (0 → 0.1)	13.7	18.1	23.1	27.0	30.7	33.0
10	9 + activation dropout (0 → 0.3)	**14.3**	**18.3**	**23.6**	**27.4**	30.4	32.6

Table 2: Results of Transformer optimized on the 5k dataset for different subsets and full corpus of IWSLT14 German → English. Averages over three runs from three different samples are reported.

(Araabi and Monz, 2020)

Other challenges with NMT

- Out-of-vocabulary words
- Low-resource languages
- Long-term context
- Common sense knowledge (e.g. *hot dog*, *paper jam*)
- Fairness and bias
- Uninterpretable

Massively multilingual MT

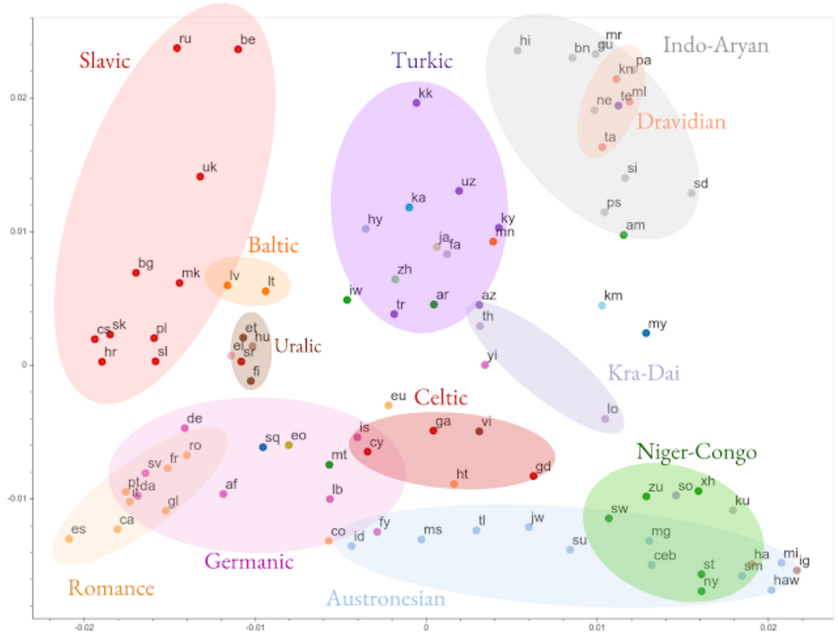

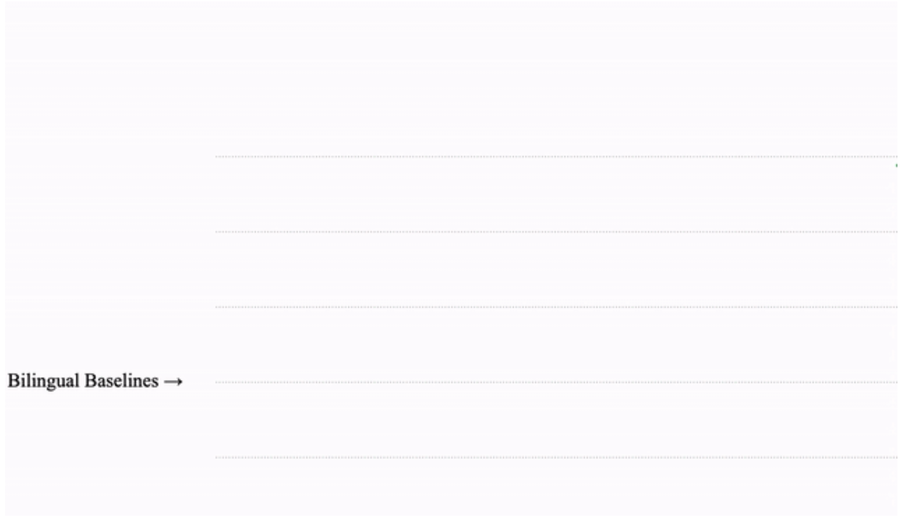

- Train a *single* neural network on 103 languages paired with English (remember Interlingua?)
- Massive improvements on low-resource languages

(Arivazhagan et al., 2019)

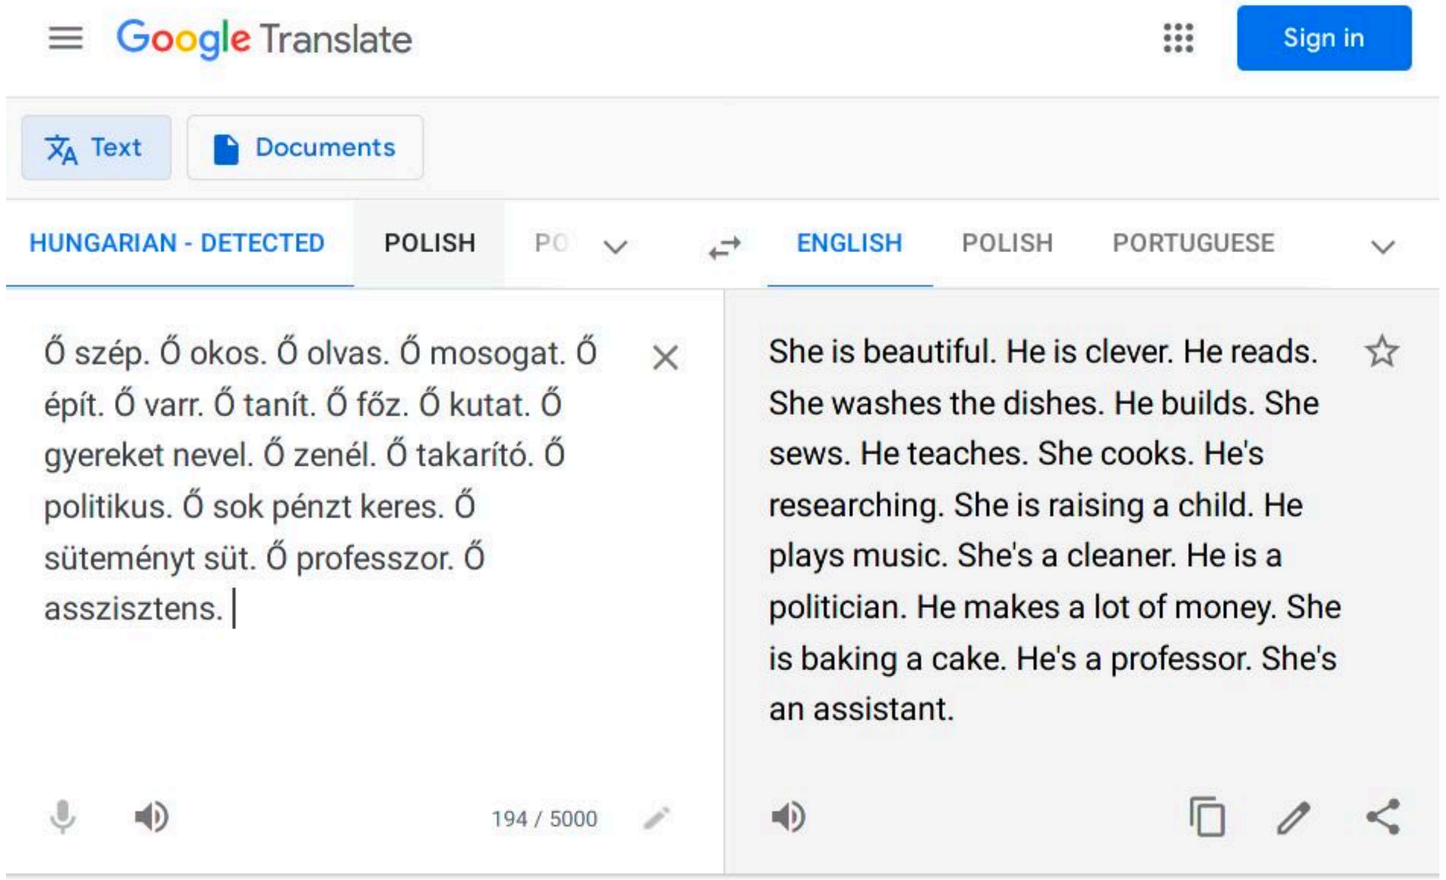
Google Translate
Sign in
Text
Documents
HUNGARIAN - DETECTED
POLISH
PO
ENGLISH
POLISH
PORTUGUESE
Ő szép. Ő okos. Ő olvas. Ő mosogat. Ő épít. Ő varr. Ő tanít. Ő főz. Ő kutat. Ő gyereket nevel. Ő zenél. Ő takarító. Ő politikus. Ő sok pénzt keres. Ő süteményt süt. Ő professzor. Ő asszisztens.
194 / 5000
She is beautiful. He is clever. He reads. She washes the dishes. He builds. She sews. He teaches. She cooks. He's researching. She is raising a child. He plays music. She's a cleaner. He is a politician. He makes a lot of money. She is baking a cake. He's a professor. She's an assistant.

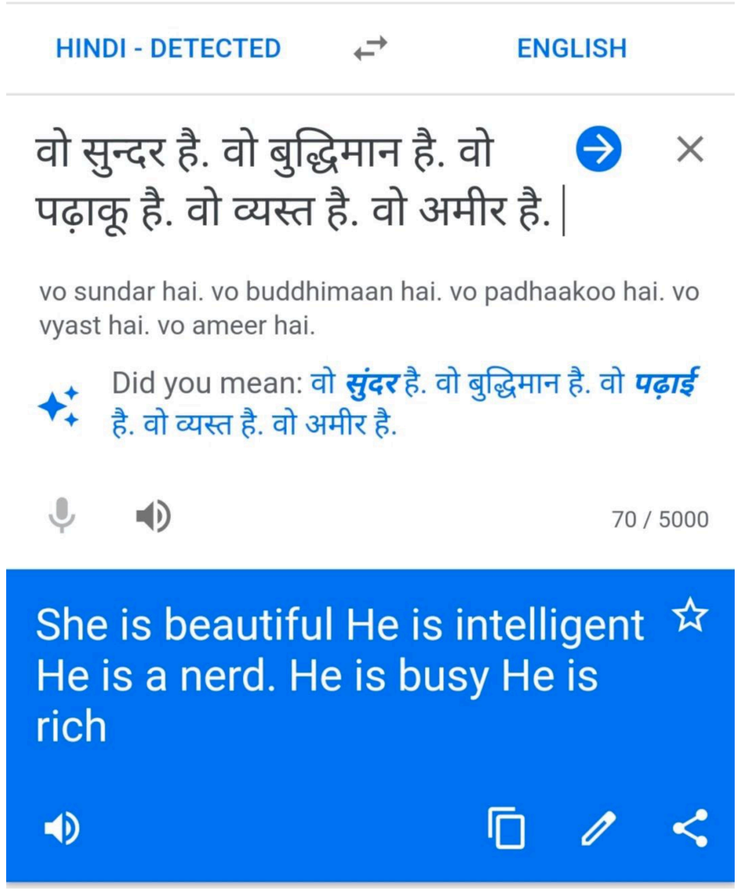
HINDI - DETECTED
ENGLISH
वो सुन्दर है. वो बुद्धिमान है. वो पढ़ाकू है. वो व्यस्त है. वो अमीर है.
vo sundar hai. vo buddhimaan hai. vo padhaakoo hai. vo vyast hai. vo ameer hai.
Did you mean: वो सुंदर है. वो बुद्धिमान है. वो पढ़ाई है. वो व्यस्त है. वो अमीर है.
70 / 5000
She is beautiful He is intelligent He is a nerd. He is busy He is rich

Bias and Fairness

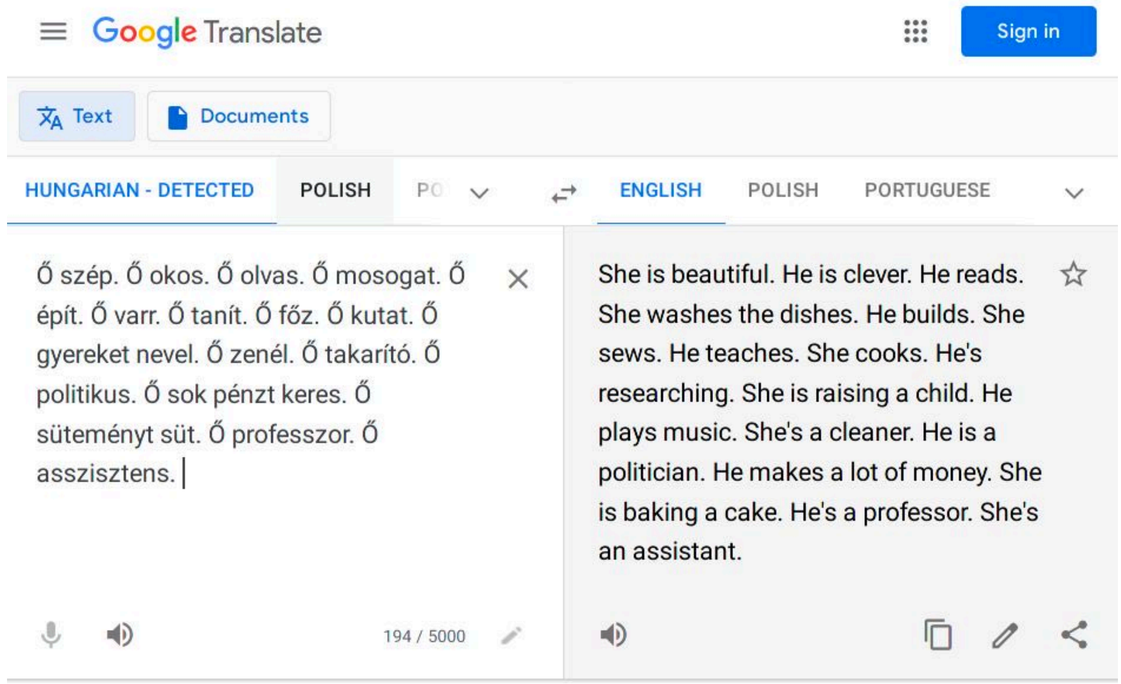

- NMT systems suffer from issues of systematic bias (e.g. gender)
- Evident when translating from/to a language with gender-specific (or gender-agnostic) terms
- Models learn (and amplify) stereotypes from data

(Farkas and Nemeth, 2020)

Measuring bias in MT

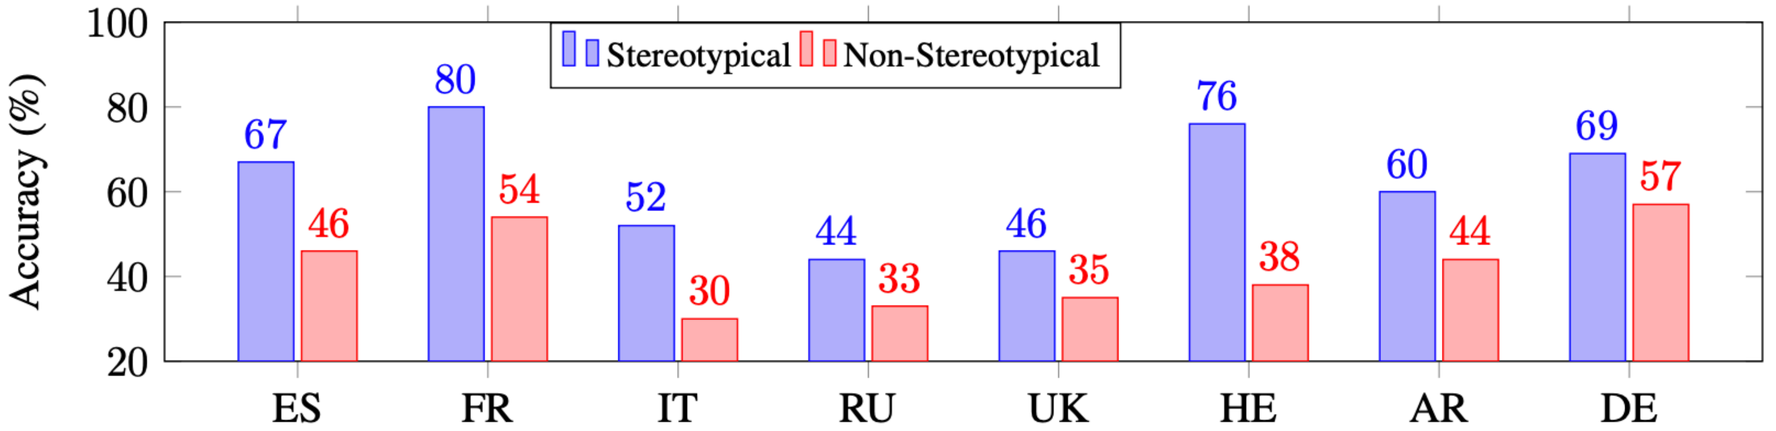

- WinoMT: Stanovsky et al. (2019) use coreference resolution to construct a dataset of non-stereotypical gender roles
  - e.g. “The doctor asked the nurse to help **her** in the operation”
- Systems consistently performed worse on non-stereotypical gender translation

Source	[Target lang.] Predicted translation	Phenomenon
The janitor does not like **the baker** because **she** always messes up the kitchen.	[ES] Al conserje no le gusta **el panadero** porque **ella** siempre desordena la cocina.	Biased translation, giving "baker" a male inflection, with a mismatched pronoun reference.
The janitor does not like **the pretty baker** because **she** always messes up the kitchen.	[ES] Al conserje no le gusta **la panadera bonita** porque **ella** siempre desordena la cocina.	Adding a stereotypically female adjective "fixes" the translation.
The counselor asked **the guard** a few questions and praised **her** for the good work.	[FR] Le conseiller a posé quelques questions à **la garde** et **l'a louée** pour le bon travail.	French uses "garde" for both male and female guards, allowing for a more direct translation from English.

Table 5: Examples of Google Translate's output for different sentences in the WinoMT corpus. Words in **blue**, **red**, and **orange** indicate male, female and neutral entities, respectively.

(Stanovsky et al. 2019)

Mitigating bias

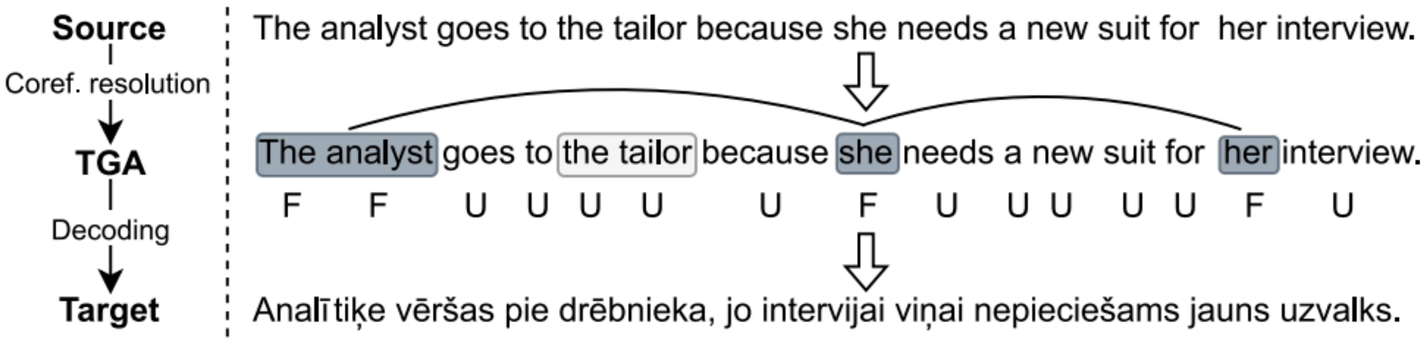

		WMT Data Systems Acc.	Δ G	Δ S	M:F
EN-DE	**Baseline**	66.7	10.2	14.4	2.6
	TGA Oracle	89.0	-4.7	1.7	1
	TGA HuggingFace	77.6	-0.1	11.9	1.6
	TGA AllenNLP	81.5	-2.0	11.1	1.4
EN-FR	**Baseline**	48.6	29.8	11.8	5.5
	TGA Oracle	81.5	1.4	2.8	1.2
	TGA HuggingFace	67.8	4.9	12.4	2
	TGA AllenNLP	74.4	1.6	10.1	1.6

- Stafanovics et al. (2020) use word-level annotations of subject's gender to train NMT systems
- TGA (target gender annotations) help reduce gender bias (∇G = diff. in F1 between sentences with male and female antecedents, ∇S = diff. in accuracy between sentences w/ or w/o stereotypes)

Anonymous feedback form:
https://forms.gle/7BxYDUTebogndJQE8