


Sequence to Sequence Learning

Quoc V. Le




joint work with many Google Brain collaborators


Projects use
GoogleBrain



Directories
use
GoogleBrain



Directories
use
GoogleBrain



A man with glasses and a red shirt is sitting at a desk in a dimly lit room. He is looking down at his hands, which are clasped together. On the desk, there is a computer monitor displaying a red screen with a white infinity symbol and a horizontal line. There are also some papers and a small object on the desk. The background is dark and out of focus.

**Machines that understand
natural language**




Step 1: Understand words

A man with glasses and a red shirt is sitting at a desk in a dimly lit office at night. He is looking at a computer monitor which displays a red screen with a white infinity symbol and a horizontal line. The desk is cluttered with papers and a keyboard. A desk lamp is visible in the background.

**Machines that understand
natural language**



Step 1: Understand words




Step 2: Understand strings of words



**Machines that understand
natural language**



Step 1: Understand words





Step 2: Understand strings of words


100
dimensions





the cat sat on the mat



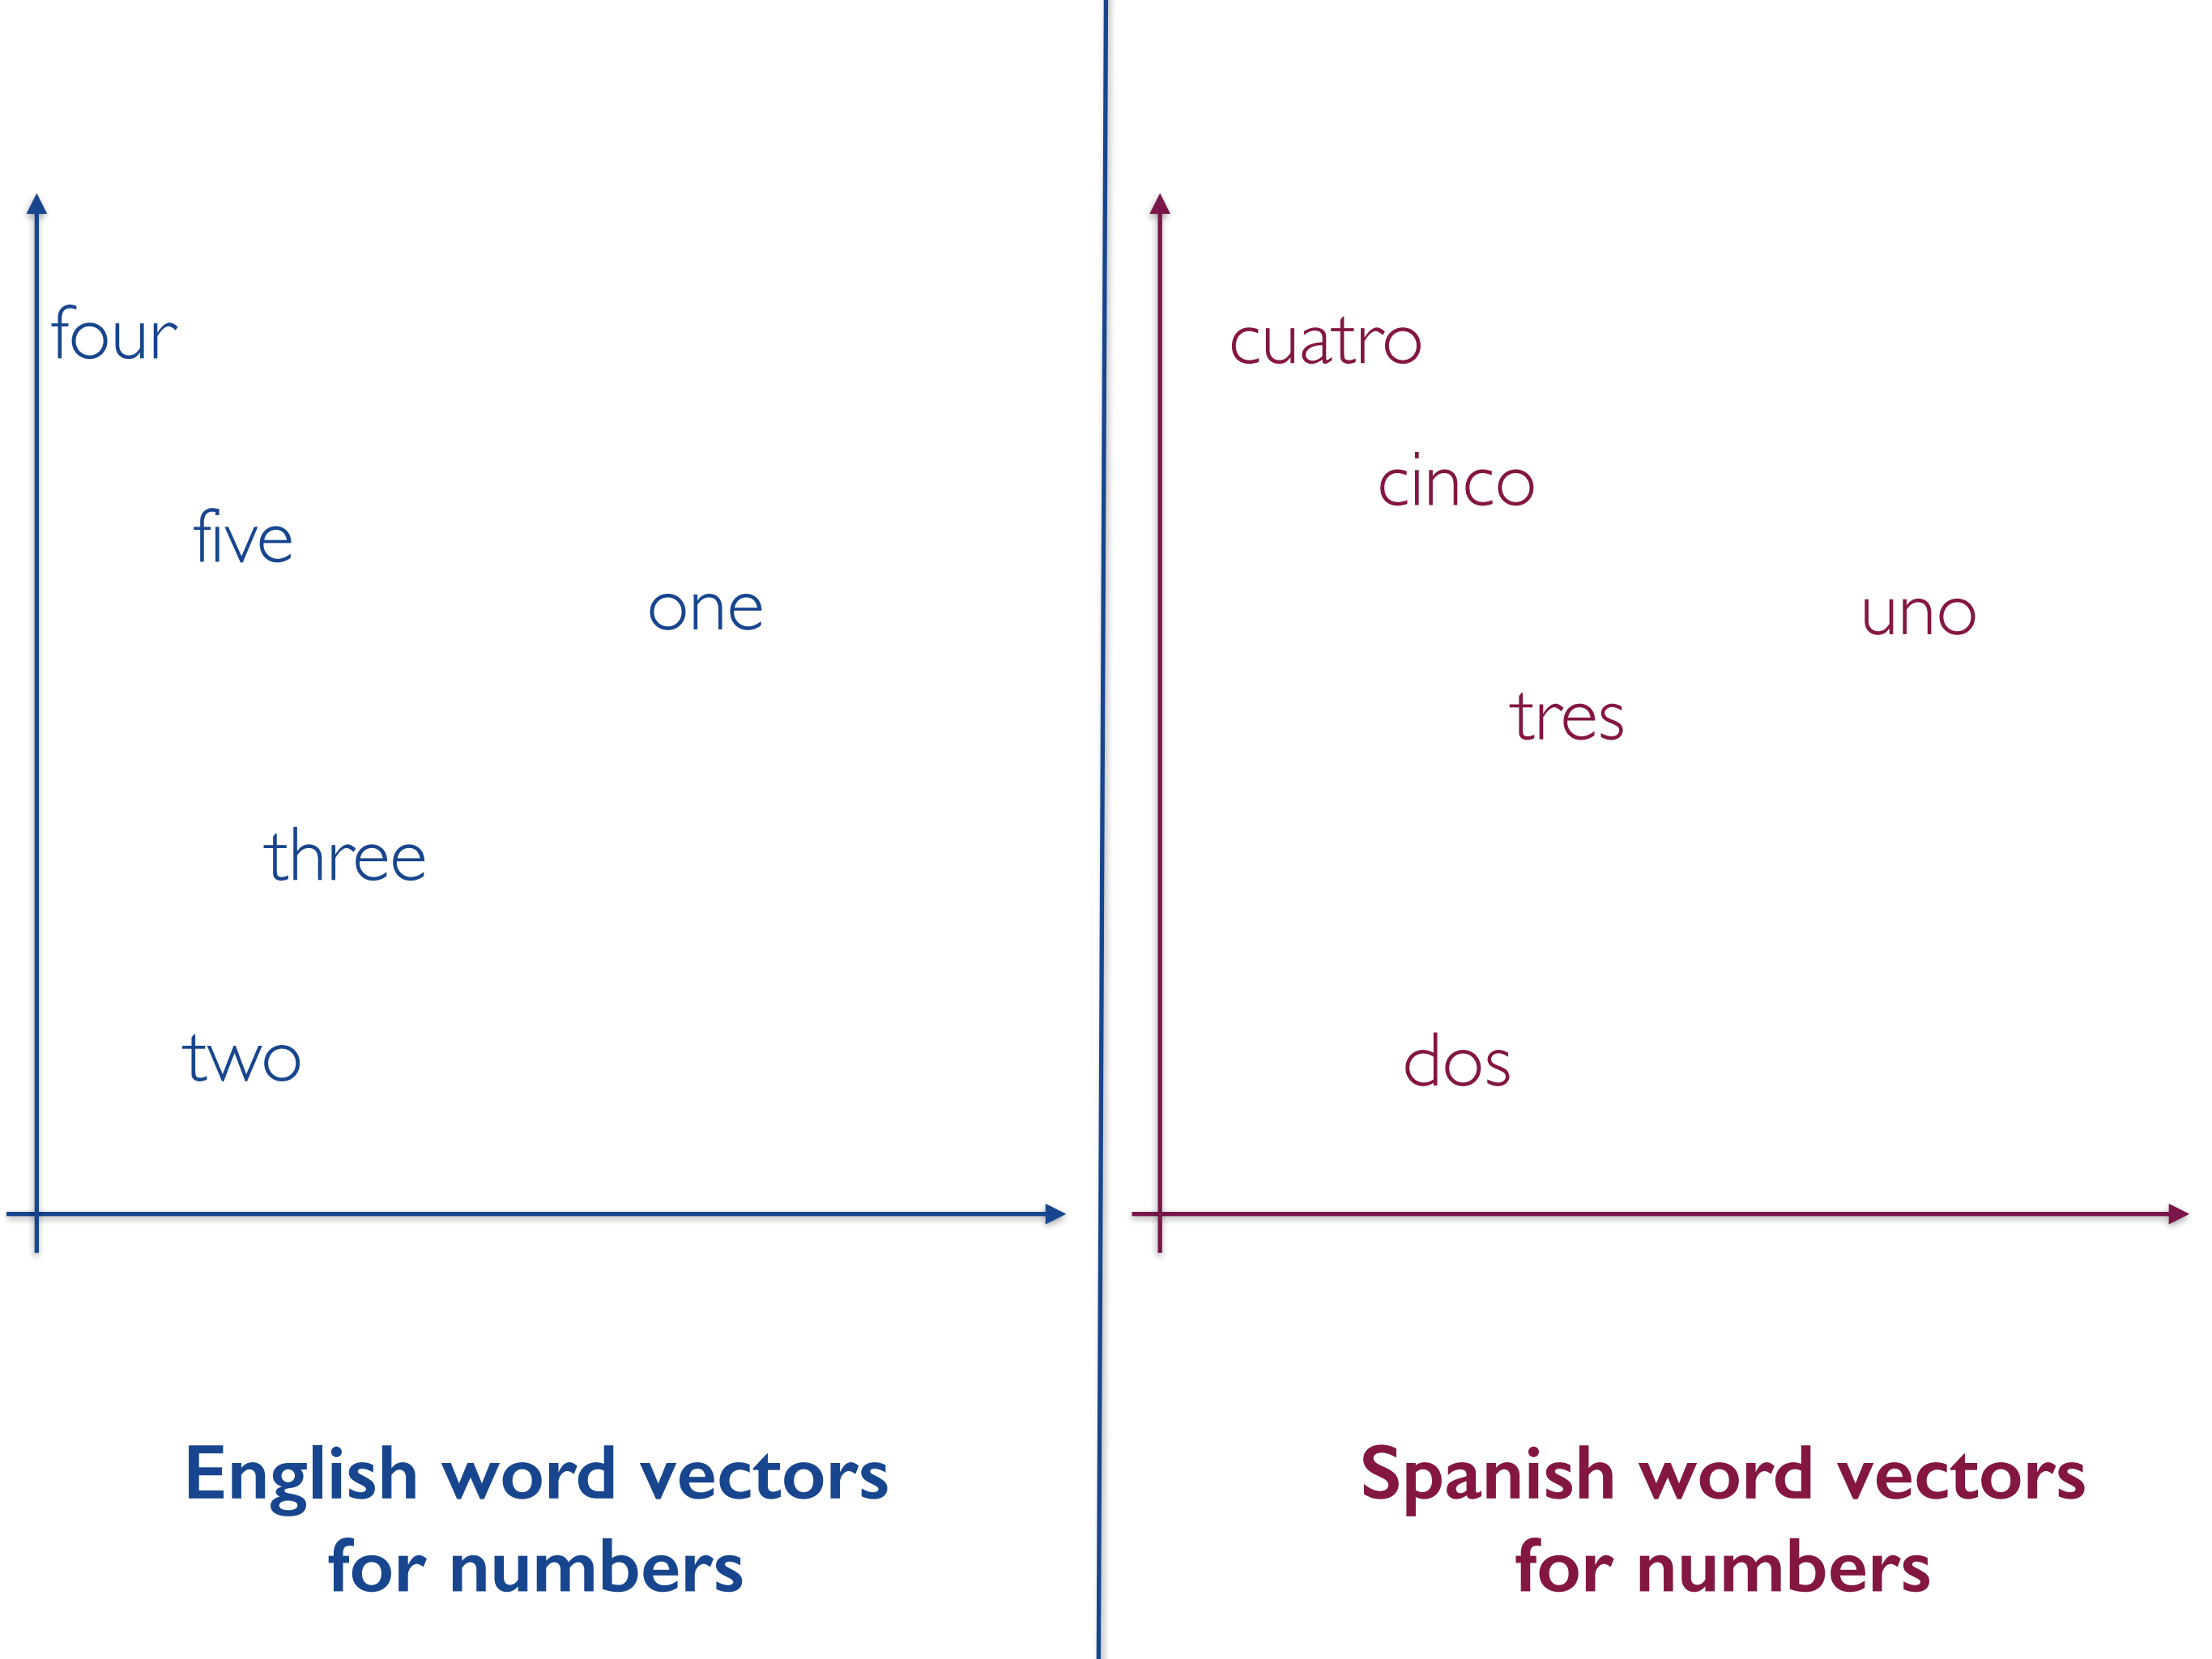








**English word vectors
for numbers**



four

five

one

three

two

English word vectors
for numbers

cuatro

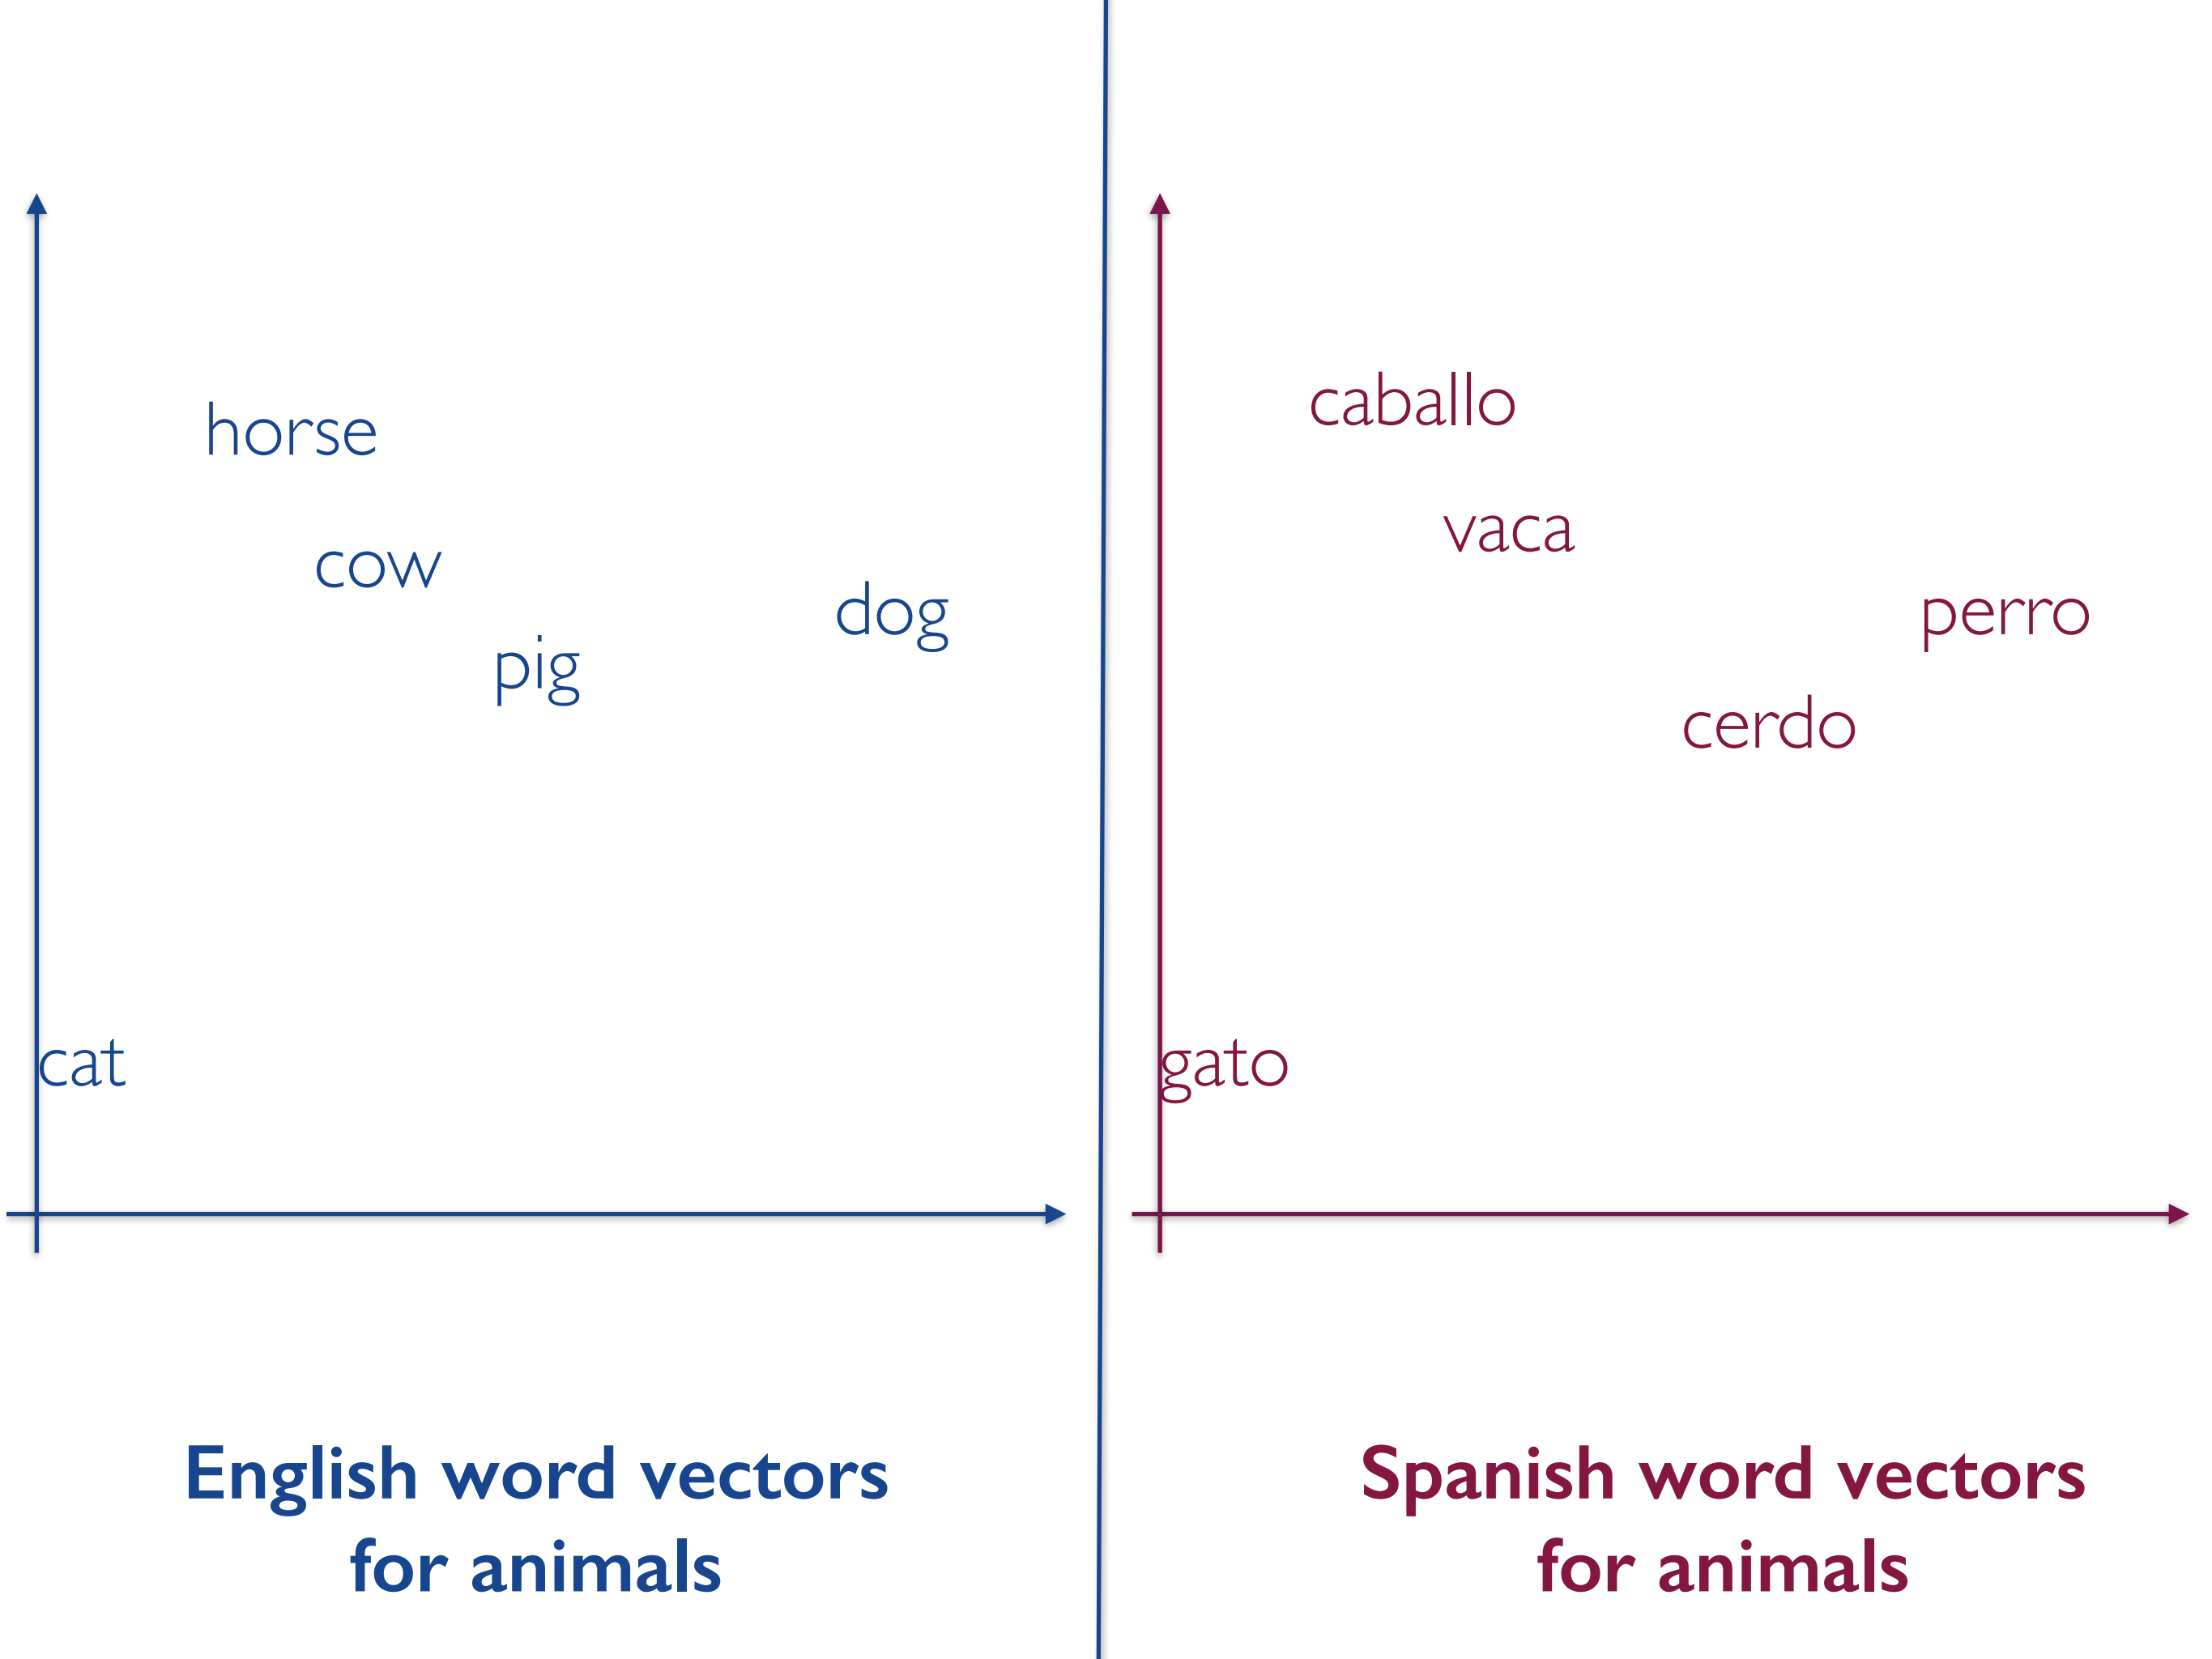
cinco

uno

tres

dos

Spanish word vectors
for numbers



horse

cow

pig

dog

cat

English word vectors
for animals

caballo

vaca

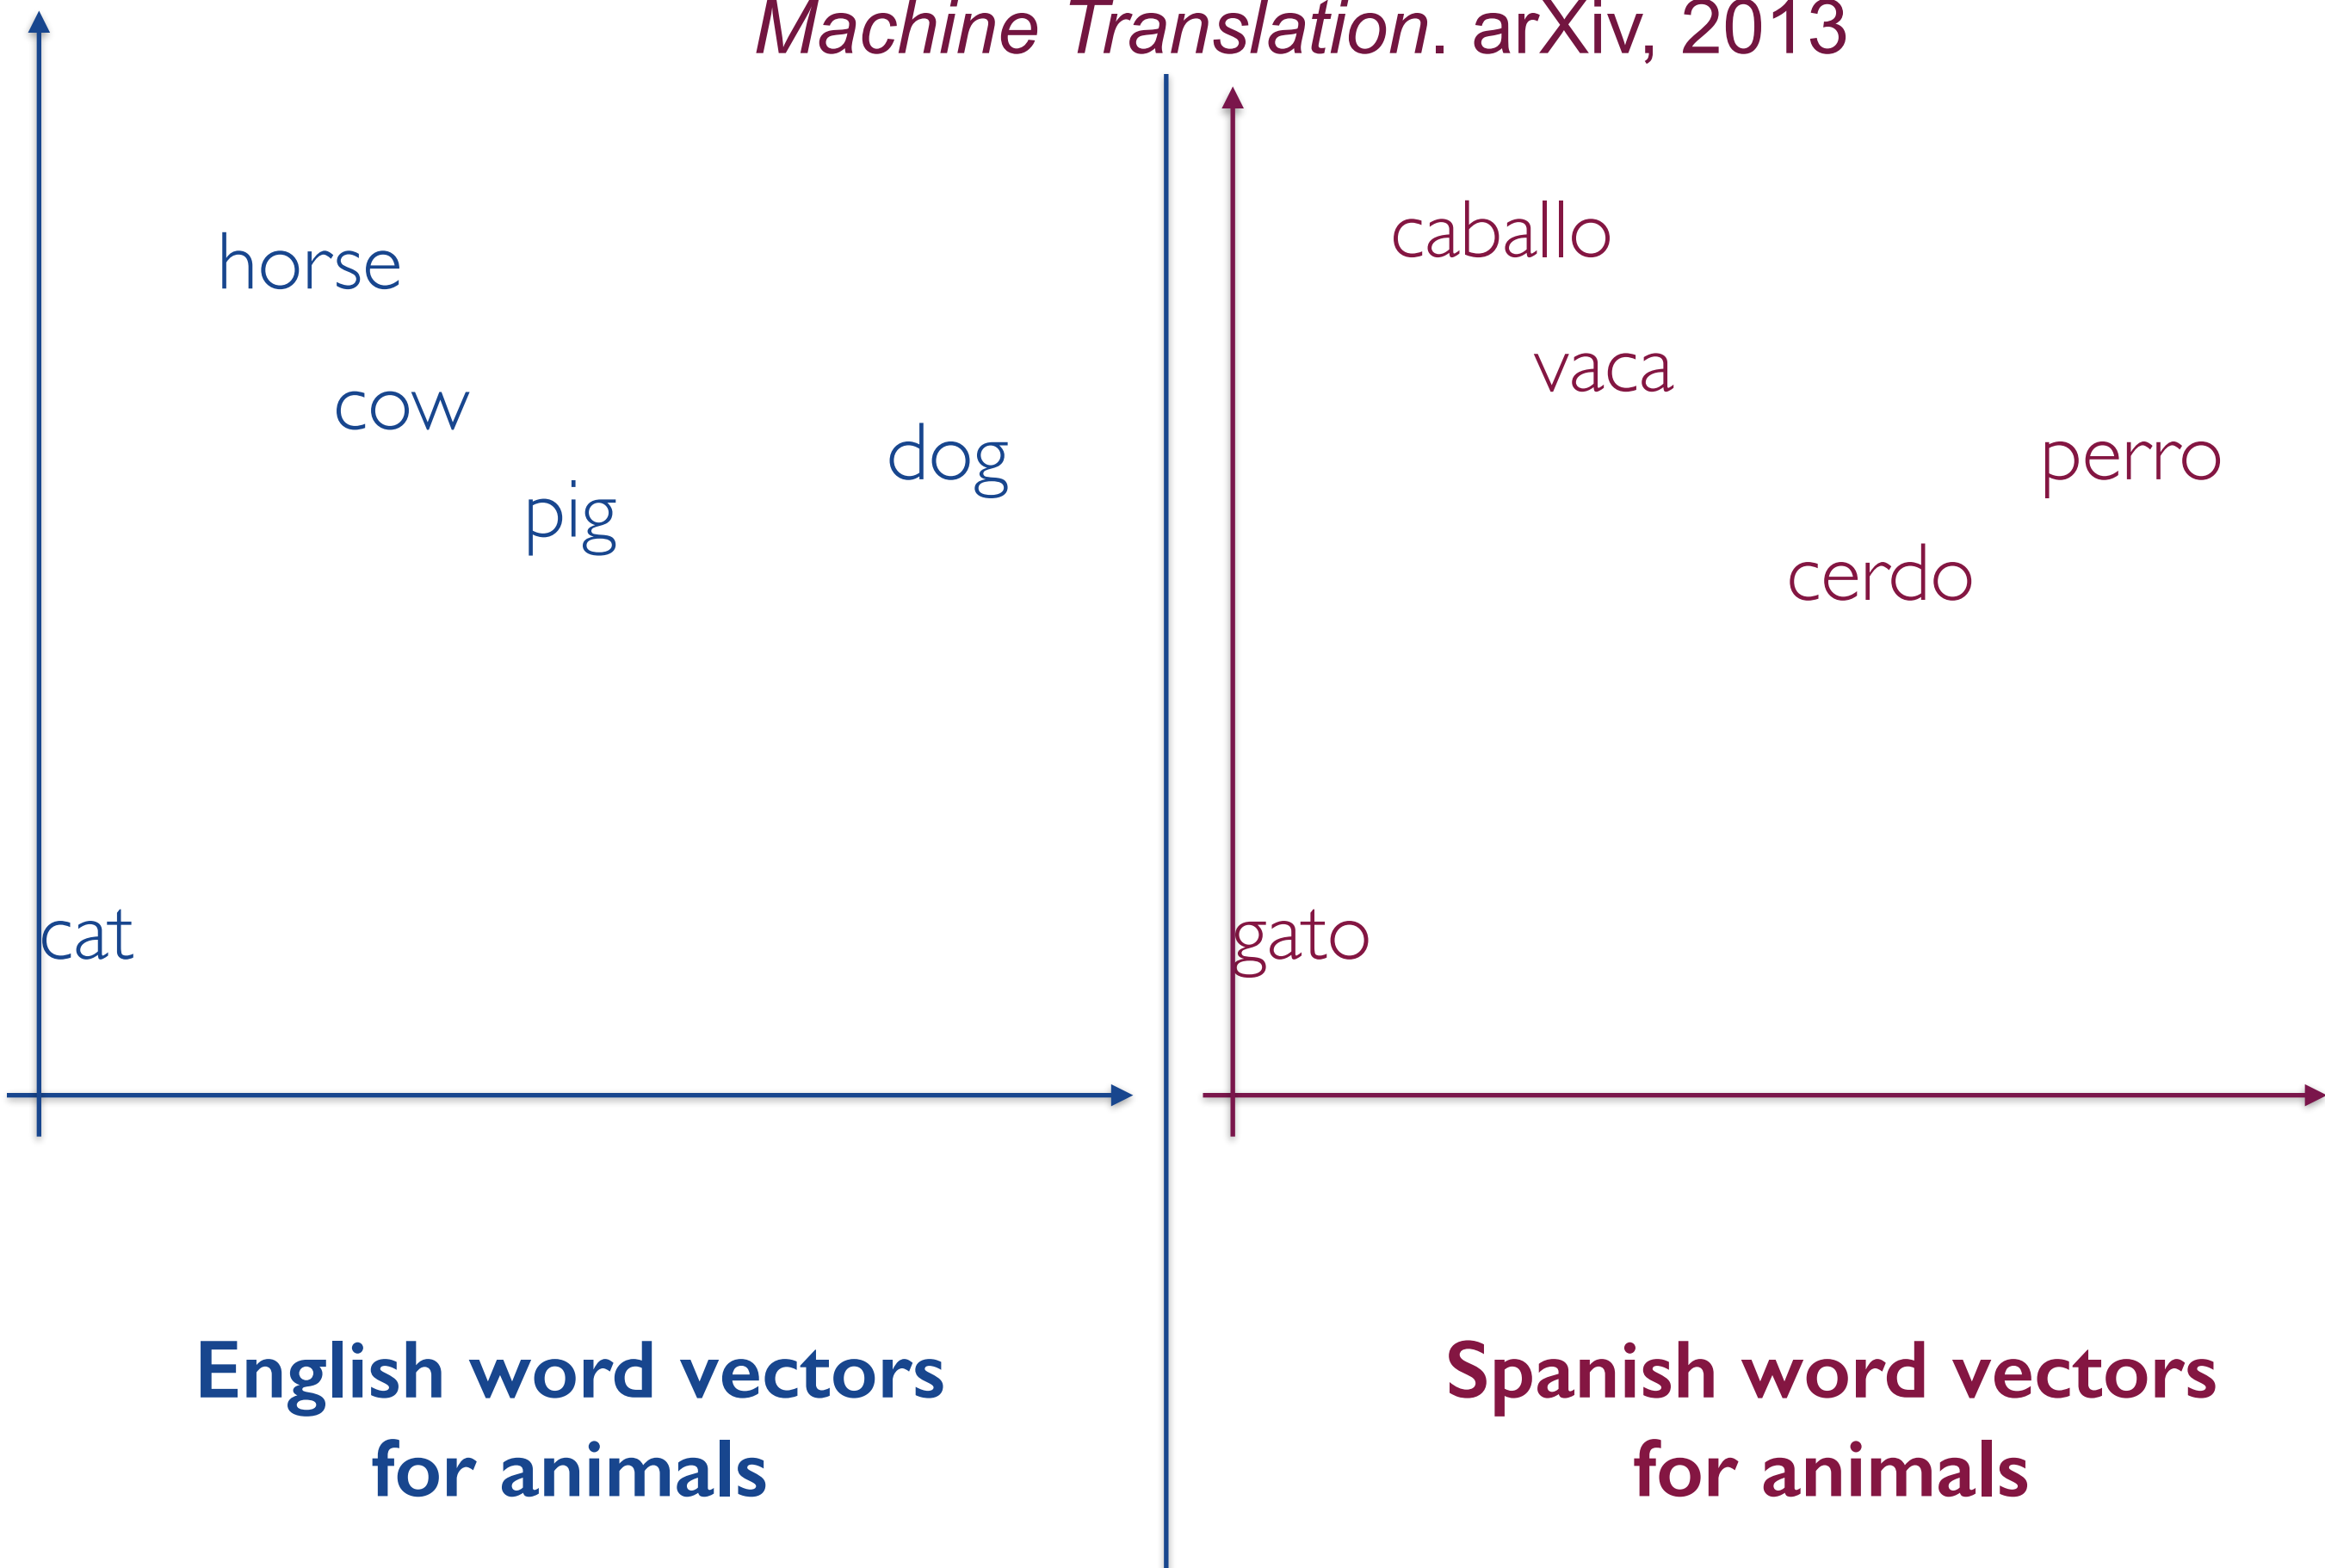
perro

cerdo

gato

Spanish word vectors
for animals

Mikolov, Le, Sutskever
*Exploiting Similarities among Languages for
Machine Translation. arXiv, 2013*




A man with glasses and a red shirt is sitting at a desk in a dimly lit office at night. He is looking at a computer monitor which displays a red square with a white infinity symbol. The background shows a city skyline through a window.

**Machines that understand
natural language**



Step 1: Understand words



Step 2: Understand strings of words



i love music

i like mathematics

i enjoy computer science

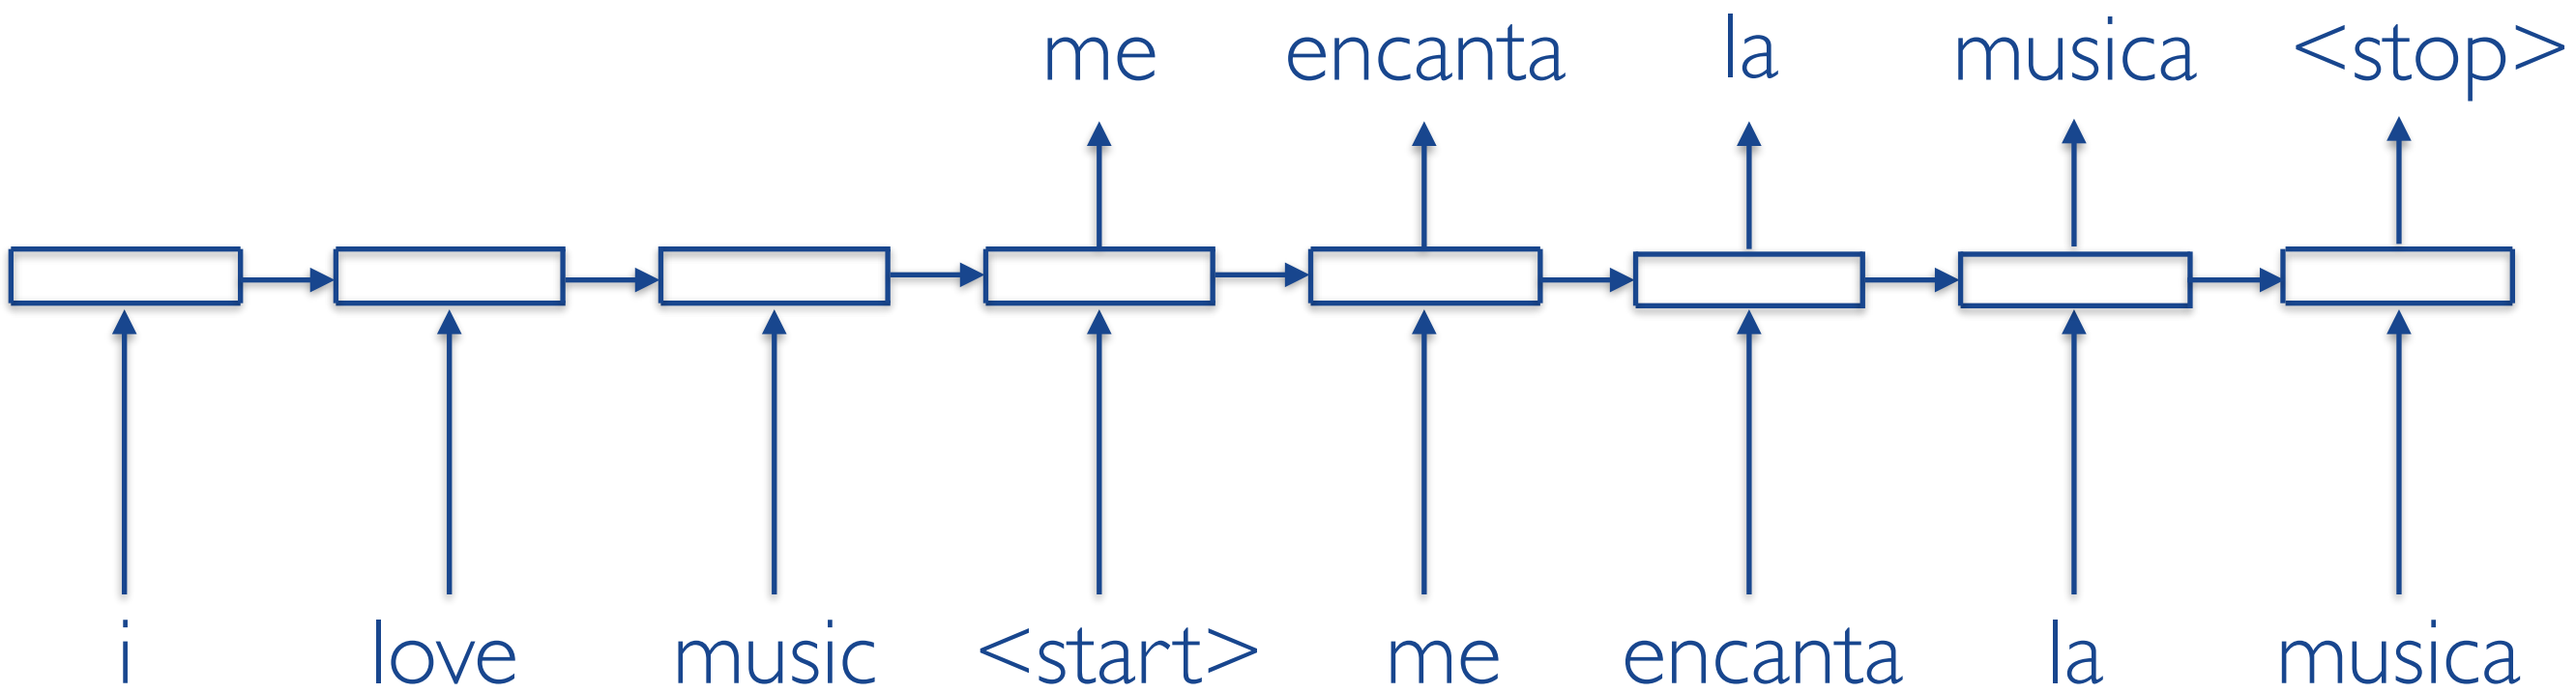
English “sentence” vectors

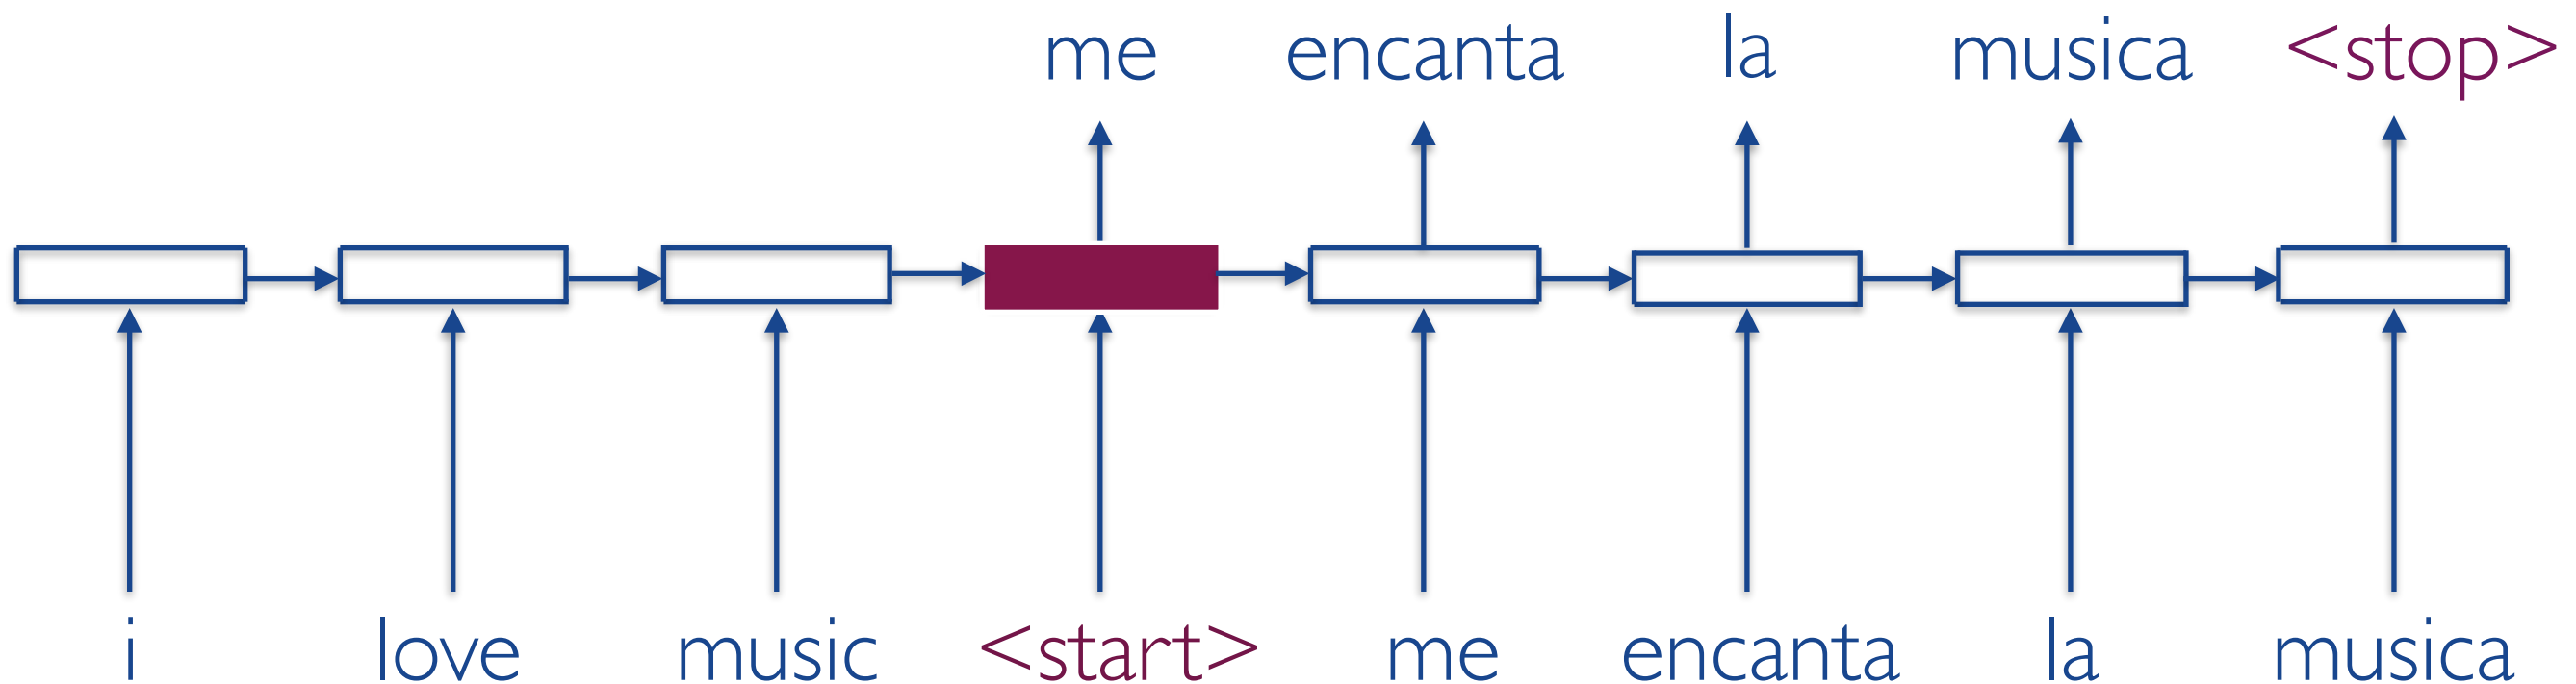
me encanta la musica

me gustan las matemáticas

me gusta disfrutar de la informática

Spanish “sentence” vectors




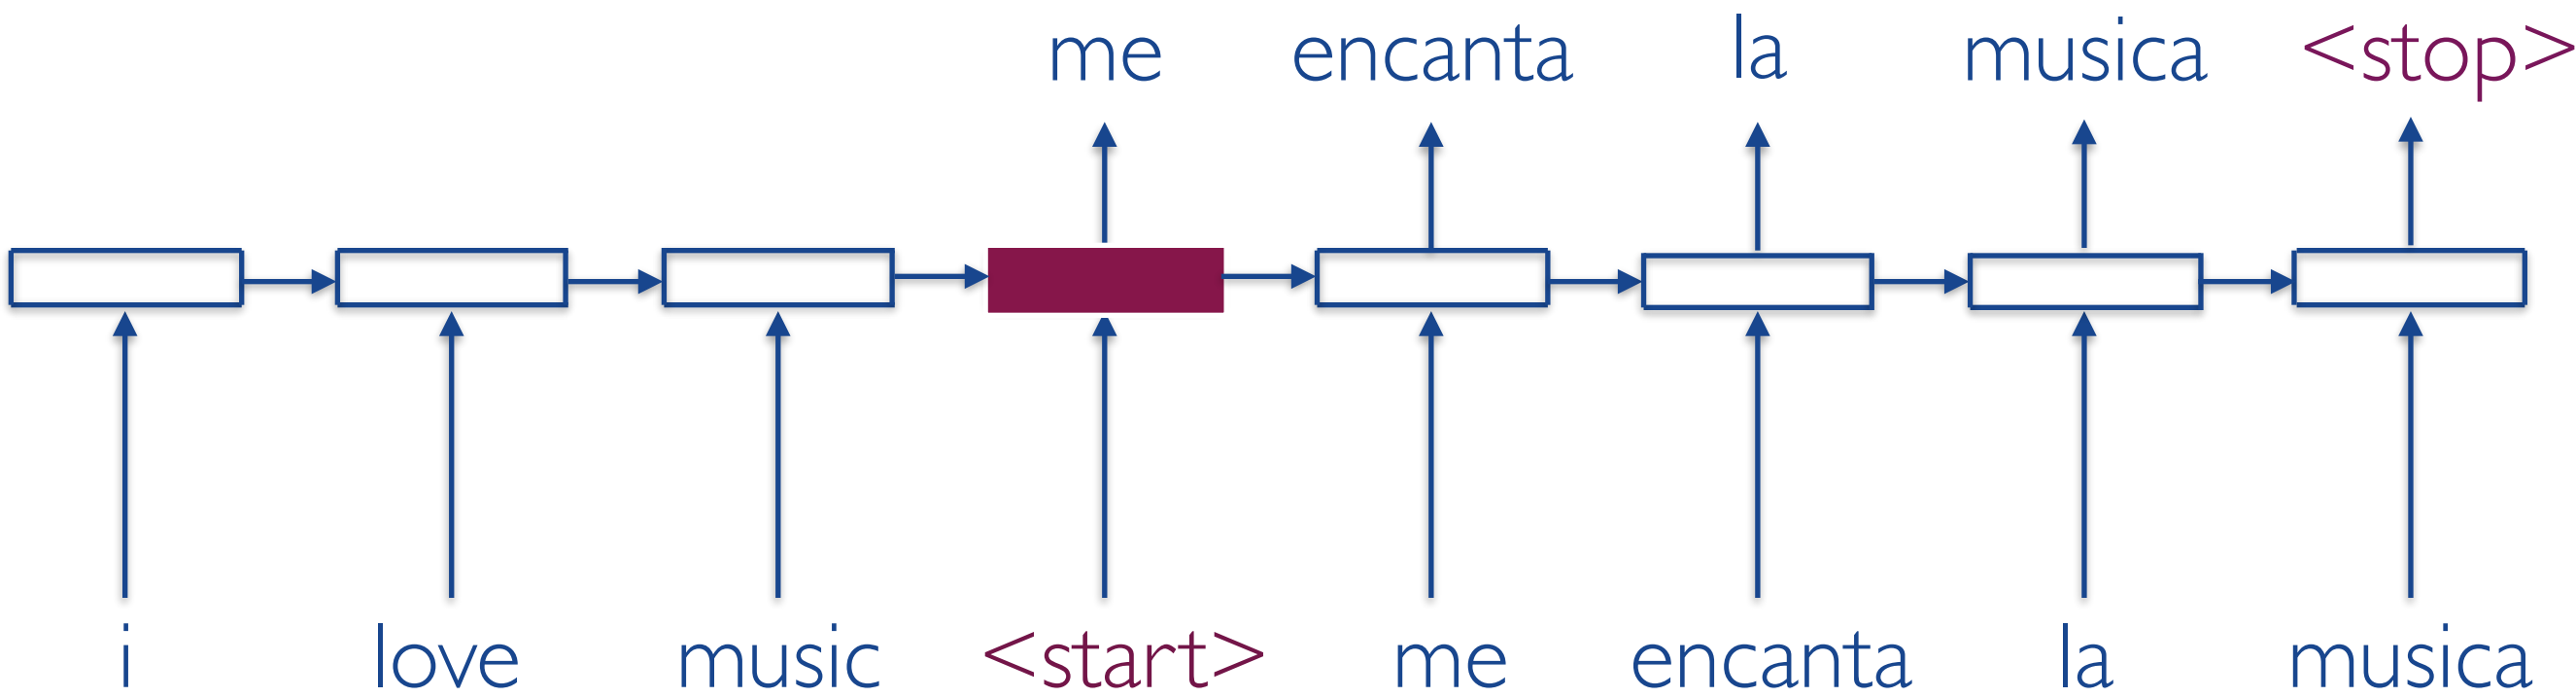


“end-to-end training”

Sutskever, Vinyals, Le. *Sequence to Sequence Learning with Neural Networks*. NIPS, 2014.

Related Approach: Cho & Bengio. Encoder-Decoder RNNs.







Mary admires John

Mary is in love with John

Mary respects John



Mary admires John

Mary is in love with John

Mary respects John

John admires Mary

John is in love with Mary

John respects Mary

A person with dark hair and glasses, wearing a red jacket, is sitting at a desk. In the background, there is a computer monitor displaying a red screen with a white symbol, and a desk lamp. The scene is dimly lit, suggesting an office environment at night.

Quoc: Hello, I have a problem with my machine

Machine: Hi, this is Shiva. How may I assist you today?

Quoc: I want to access using VPN

Machine: Currently, the machine is connected to corporate network?

Quoc: No

Machine: Check out this solution <URL>

Quoc: Thanks , bye!

Machine: Thank you.

A person with dark hair and glasses, wearing a red long-sleeved shirt, is sitting at a desk. They are looking down at something on the desk. In the background, there is a computer monitor displaying a red screen with a white infinity symbol and a vertical line. A desk lamp is visible above the monitor.

Quoc: What is the purpose of life?

Machine: To serve the greater good.

Quoc: What is the purpose of living?

Machine: To live forever.

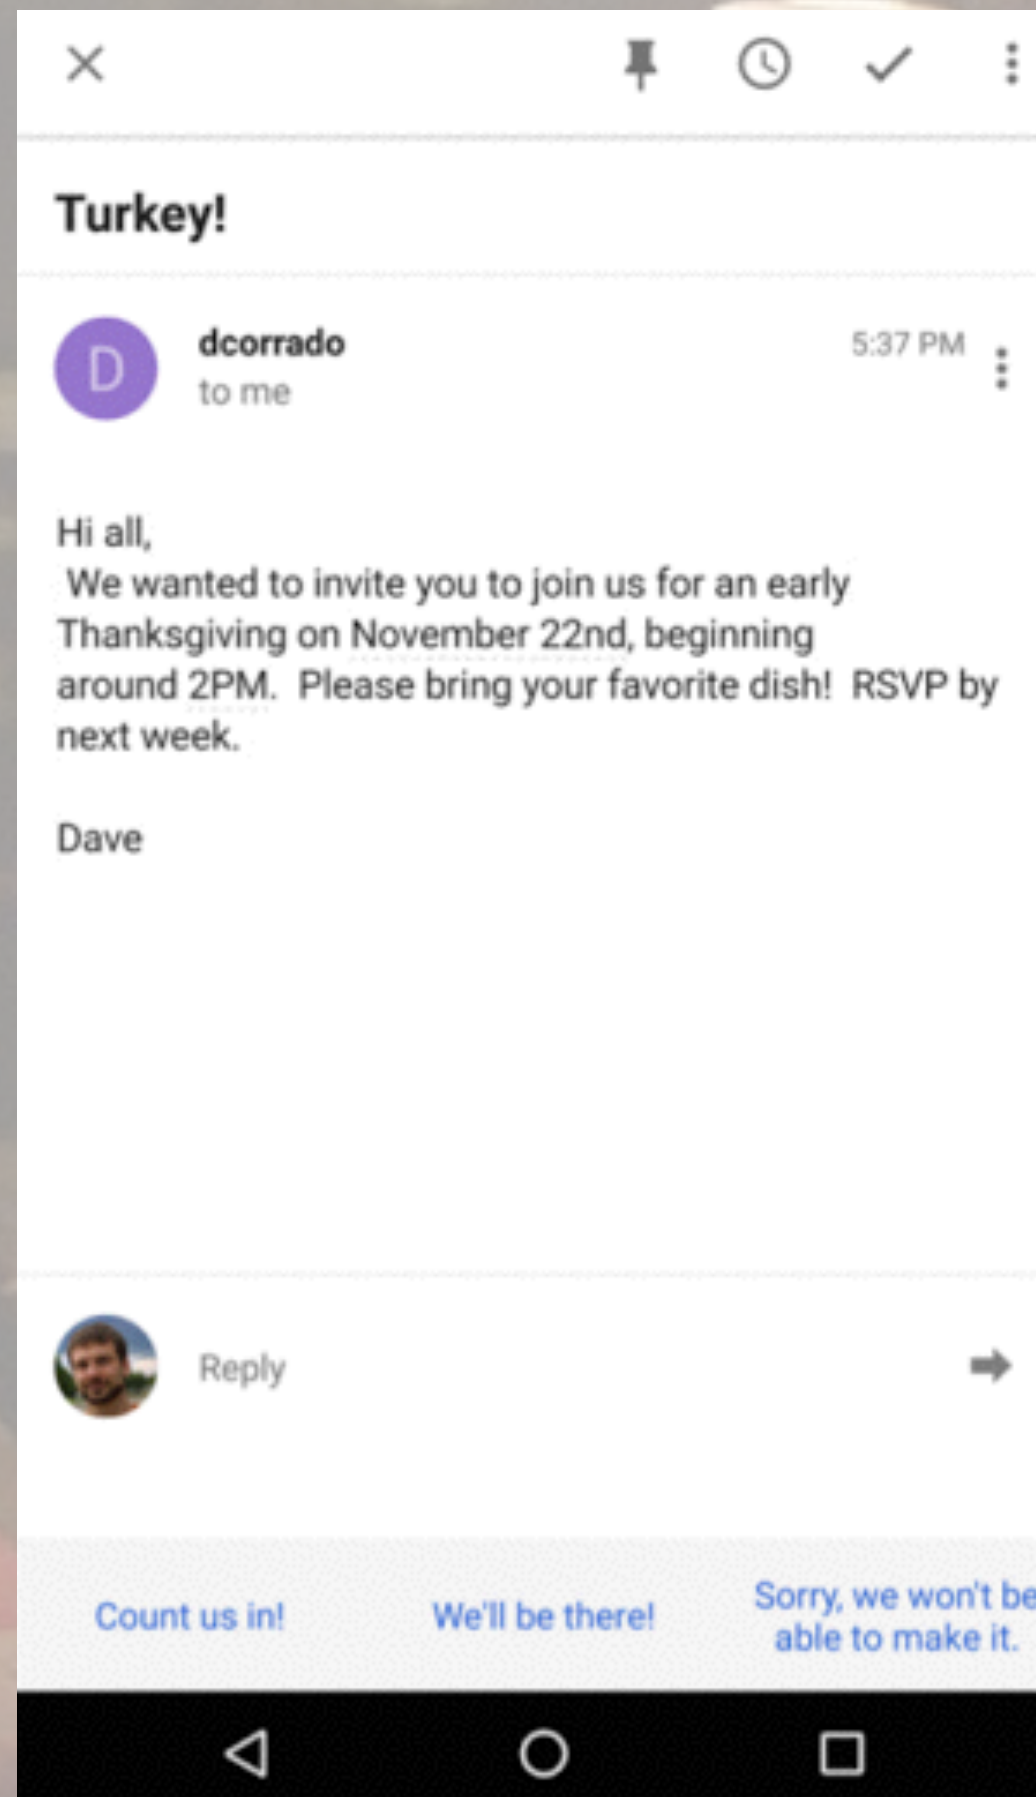
Quoc: What is the purpose of dying?

Machine: To have a life.

Vinyals, Le. *A Neural Conversational Model*.

ICML Deep Learning Workshop, 2015.

SmartReply



**Machines that understand
natural language**



Step 1: Understand words



Step 2: Understand strings of words



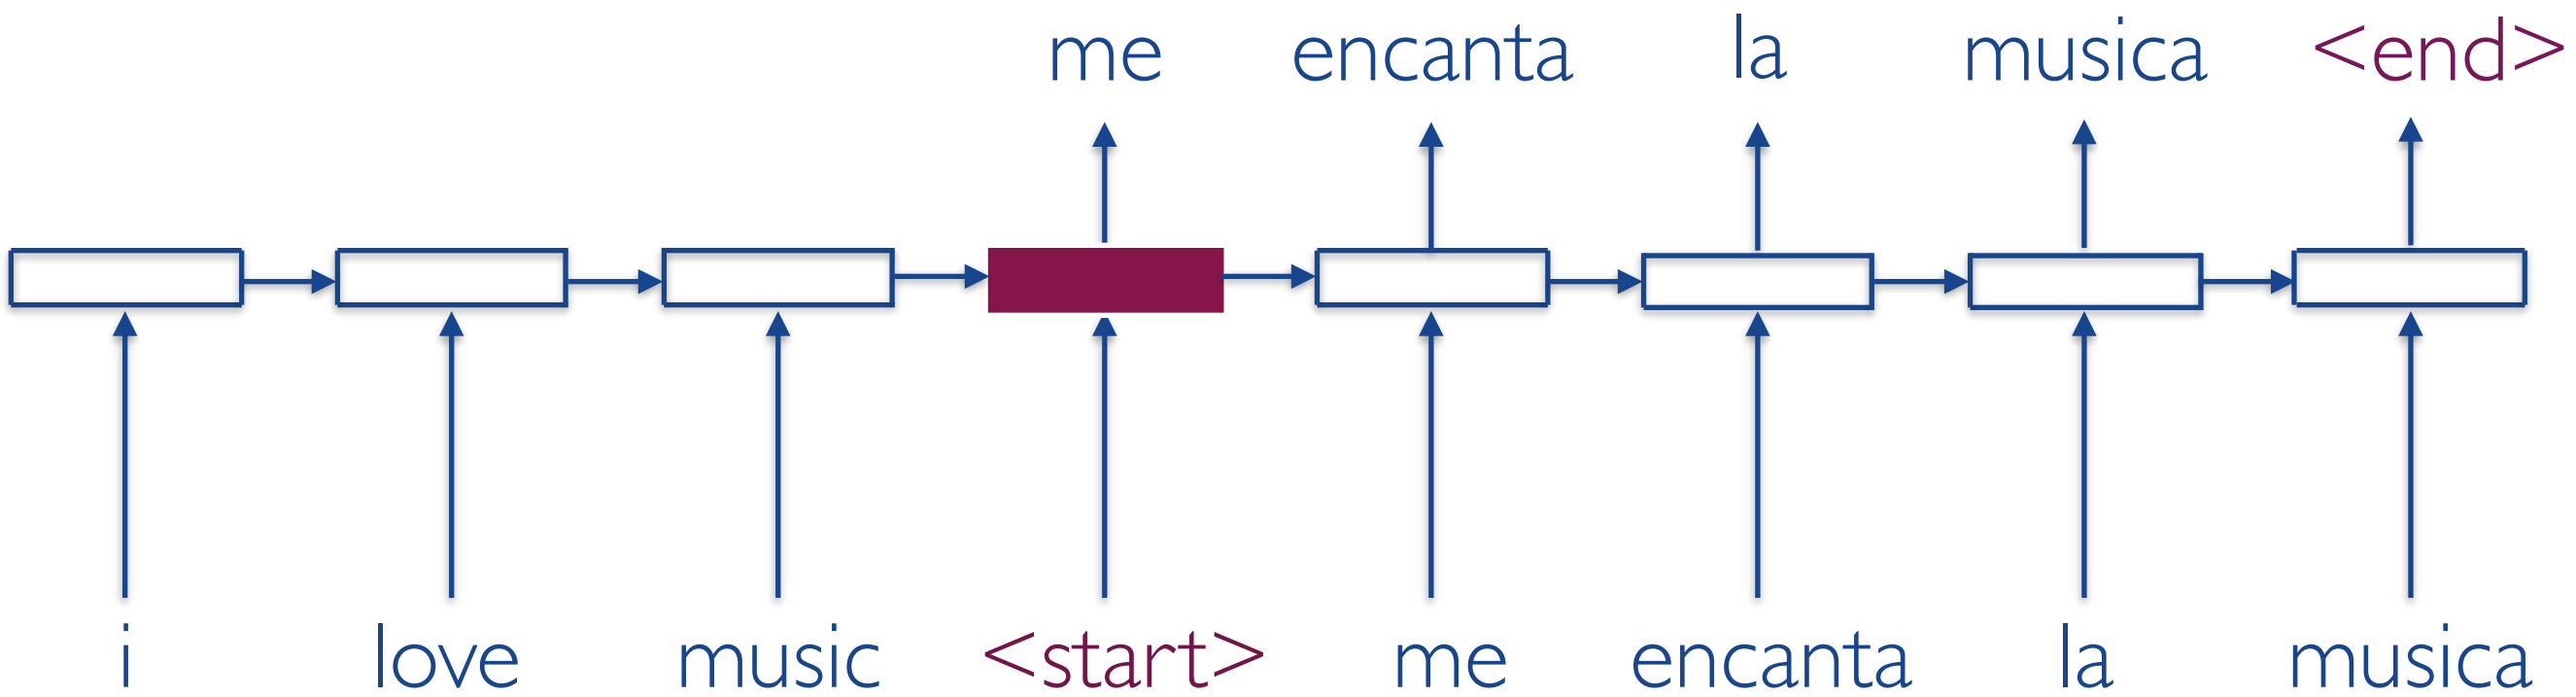
Step 3: Have memory, logical reasoning


Advances in Memory and Logical Reasoning


“Soft” Attention Mechanism
(Bahdanau, Cho, Bengio, 2014)


Augmented Memory
(Graves et al, 2014; Weston et al, 2014)


Augmented Logic and Arithmetic Components
(Neelakantan, Le, Sutskever, 2016)













Bahdanau, Cho, Bengio. *Neural Machine Translation by Jointly Learning to Align and Translate*. ICLR, 2015.



Graves, Wayne, Danihelka. *Neural Turing Machines*. ArXiv, 2014
Sukhabaatar, Szlam, Weston, Fergus.
End-to-end memory networks. NIPS, 2015.



Neelakantan, Le, Sutskever. *Neural Programmer: Inducing Latent Programs with Gradient Descent*. ICLR, 2016.

ALU Inputs and Outputs





Figure 10.1 ALU Inputs and Outputs



Neelakantan, Le, Sutskever. *Neural Programmer: Inducing Latent Programs with Gradient Descent*. ICLR, 2016.

A person with glasses and a pink shirt is sitting at a desk in a dimly lit room. The desk is cluttered with papers and a laptop. A large window in the background shows a city skyline at night with many lights. A desk lamp is visible on the right side of the desk.


Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?

Machine: 5 apples.

Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?



External Memory



Arithmetic & Logic
Functions


add(a,b)
subtract(a,b)
multiply(a,b)

...

Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?



External Memory



Arithmetic & Logic
Functions


add(a,b)
subtract(a,b)
multiply(a,b)

...

Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?



External Memory



Arithmetic & Logic
Functions


add(a,b)
subtract(a,b)
multiply(a,b)

...

Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?



External Memory



Arithmetic & Logic
Functions


add(a,b)
subtract(a,b)
multiply(a,b)

...

Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?



External Memory




Arithmetic & Logic
Functions


add(a,b)
subtract(a,b)
multiply(a,b)

...


Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?



Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?




Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?



$\text{add}(a,b) = 5$
 $\text{subtract}(a,b) = -1$
 $\text{multiply}(a,b) = 6$


...

Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?



$\text{add}(a,b) = 5 * w1$
 $\text{subtract}(a,b) = -1 * w2$
 $\text{multiply}(a,b) = 6 * w3$
...

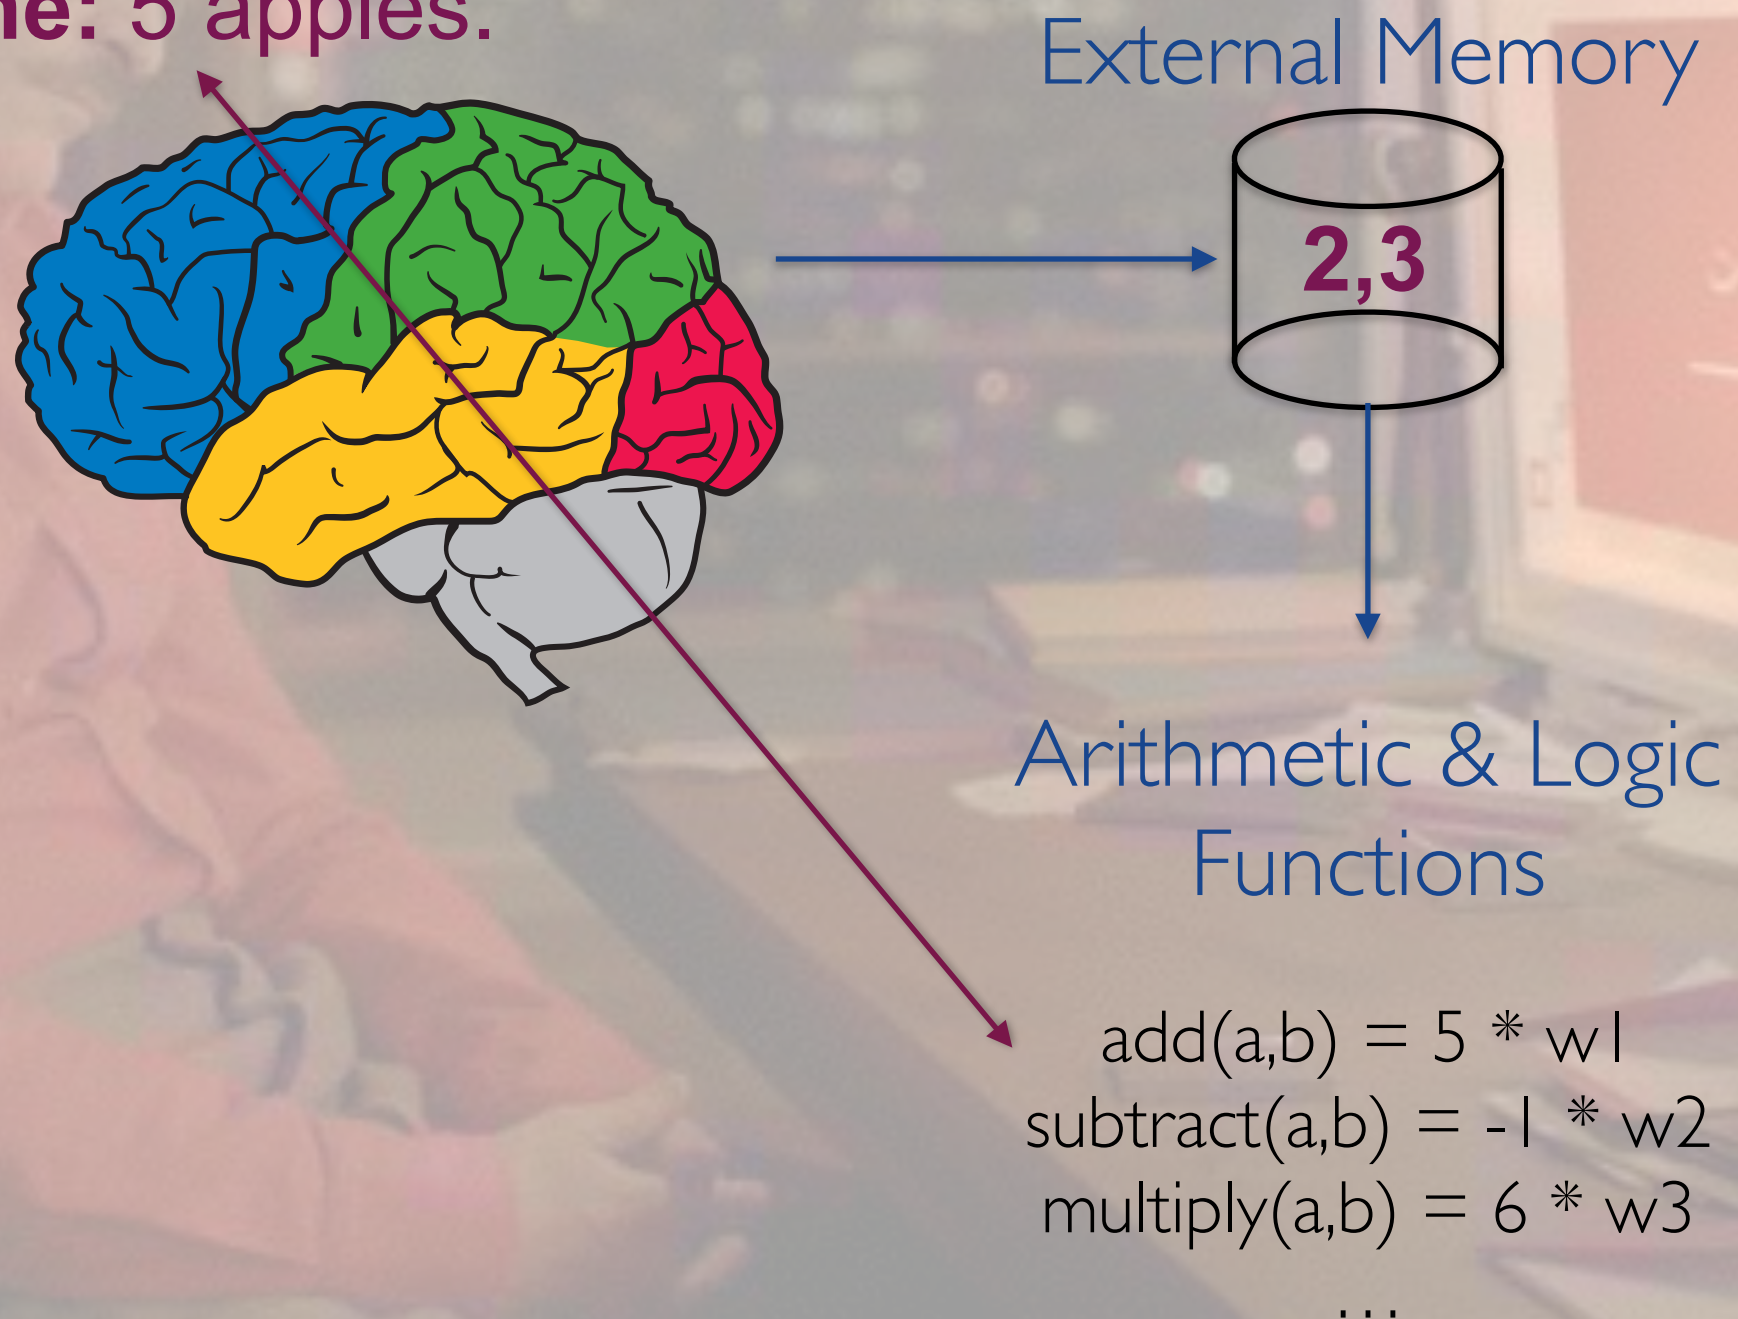
Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?



$\text{add}(a,b) = 5 * w1$
 $\text{subtract}(a,b) = -1 * w2$
 $\text{multiply}(a,b) = 6 * w3$
...

Quoc: I have 2 apples and Tom gives me 3 apples. How many apples do I have?

Machine: 5 apples.



**Machines that understand
natural language**



Step 1: Understand words



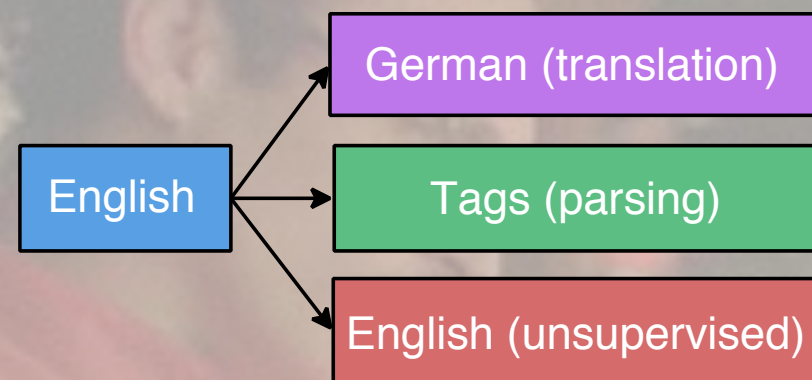
Step 2: Understand strings of words

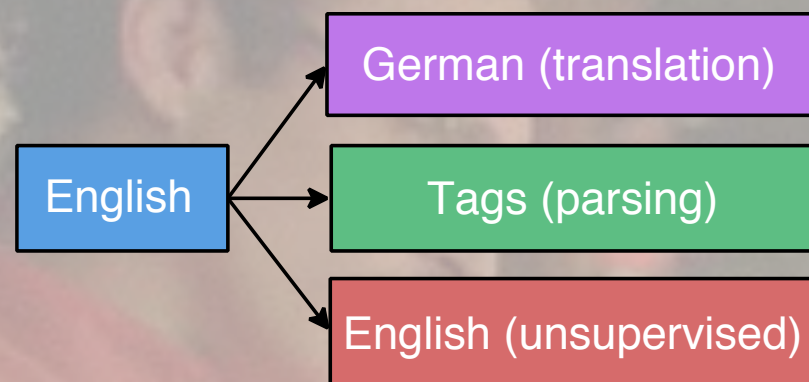


Step 3: Have memory, logical reasoning

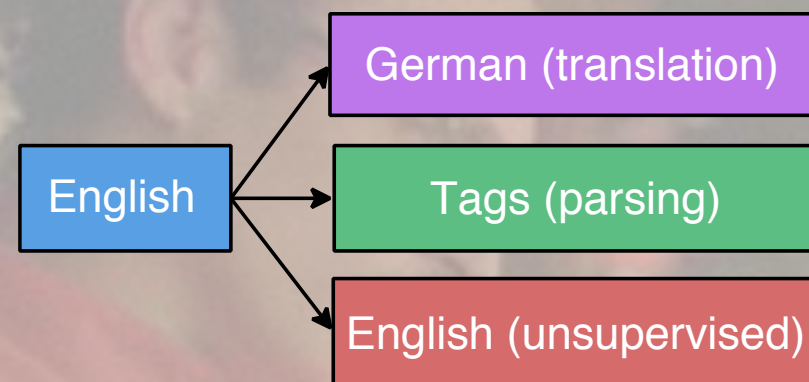


Step 4: Learn from many tasks

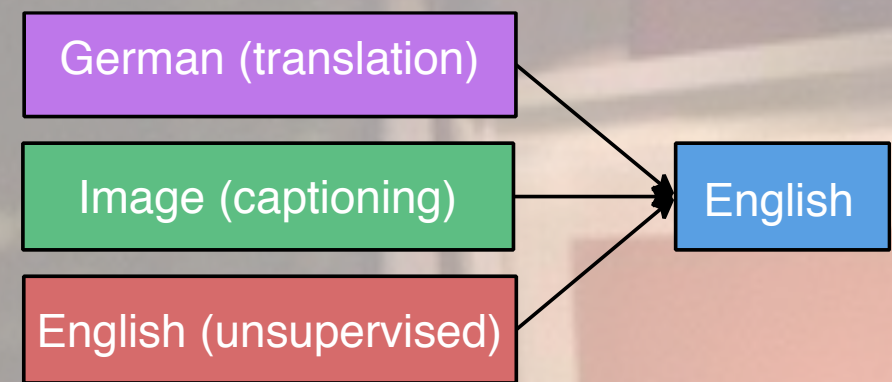





One-to-many



One-to-many

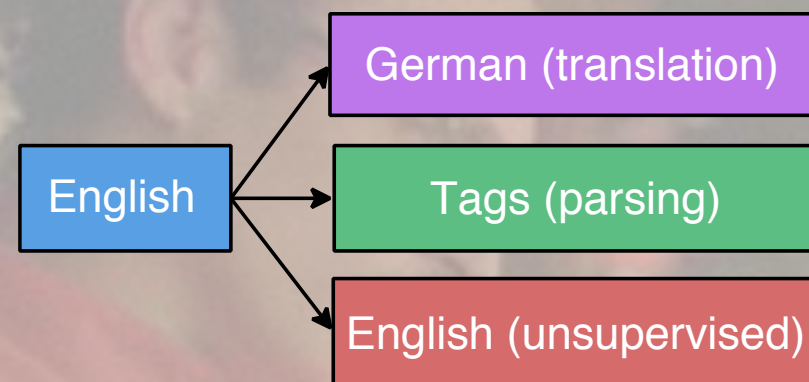


Many-to-one

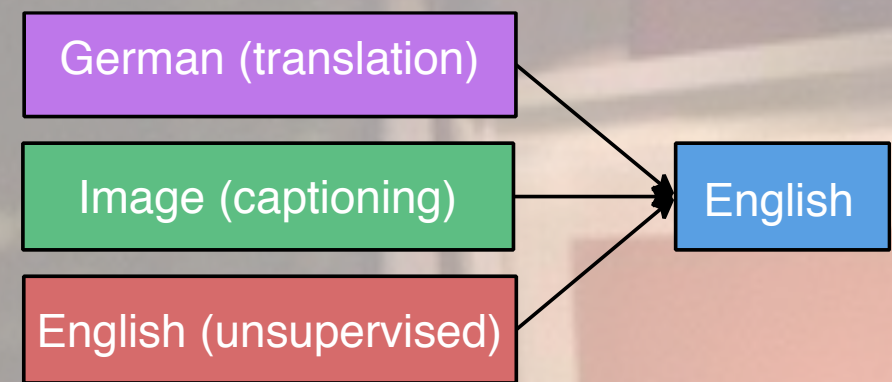


Many-to-many


Able to improve accuracies of all related tasks



One-to-many



Many-to-one



Many-to-many

Able to improve accuracies of all related tasks

Dai, Le. *Semi-supervised Sequence Learning*. NIPS, 2015.
Luong et al. *Multitask Sequence Learning*. ICLR, 2016

**Machines that understand
natural language**



Step 1: Understand words



Step 2: Understand strings of words



Step 3: Have memory, logical reasoning



Step 4: Learn from many tasks