University of Toronto
Faculty of Arts and Science
Department of Computer Science

Reasoning with Neural Networks

Rodrigo Toro Icarte (rntoro@cs.toronto.edu)

March 08, 2016

Motivation

Could a crocodile run a steeplechase?[1]

[1] The example was borrowed from Levesque (2014)

Symbolic approach

KB:

$$\begin{aligned}
&\ldots\\
&\forall x. Crocodile(x) \supset WeakLegs(x)\\
&\ldots\\
&\forall x. WeakLegs(x) \supset \neg CanJump(x)\\
&\ldots\\
&\forall x. \neg CanJump(x) \supset \neg CanSteeplechase(x)\\
&\ldots
\end{aligned}$$

Query: $\neg \exists x. Crocodile(x) \wedge CanSteeplechase(x)$

Symbolic approach

Strategy: Contradiction proof.
Assume query is false: $Crocodile(a) \wedge CanSteeplechase(a)$

Symbolic approach

Strategy: Contradiction proof.
Assume query is false: $Crocodile(a) \wedge CanSteeplechase(a)$

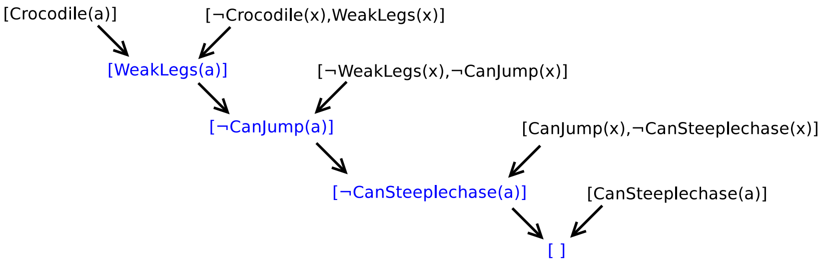

Symbolic approach

Strategy: Contradiction proof.
Assume query is false: $Crocodile(a) \wedge CanSteeplechase(a)$

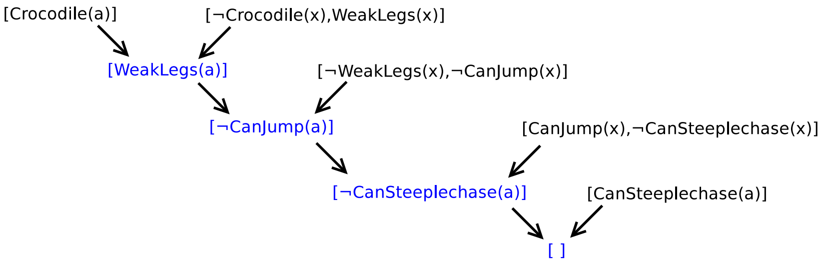

Observations:

- This is SLD resolution.

Symbolic approach

Strategy: Contradiction proof.
Assume query is false: $Crocodile(a) \wedge CanSteeplechase(a)$

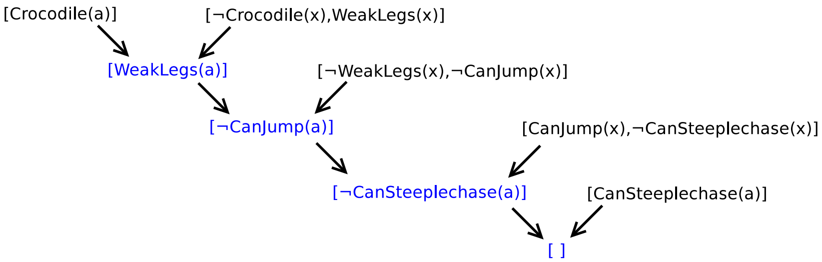

Observations:

- This is SLD resolution.
- Scalability problems.

A Neural Conversational Model Vinyals & Le (2015)

They trained a seq2seq:

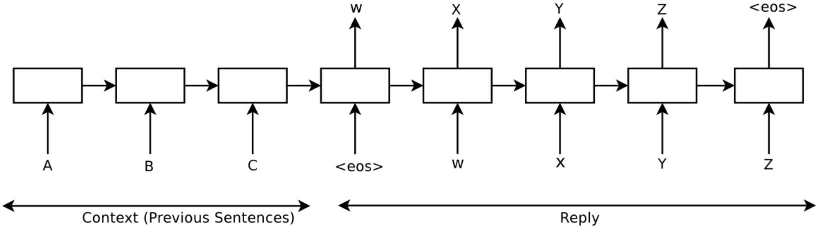

Credit: Vinyals & Le

... with movie subtitles.

A Neural Conversational Model Vinyals & Le (2015)

Human: who is skywalker ?
Machine: he is a hero .

Human: what is the color of the sky ?
Machine: blue .

Human: how many legs does a cat have ?
Machine: four , i think .

Human: what is the purpose of living ?
Machine: to live forever

Research question

Can we use neural networks for common sense reasoning?

Long term goal

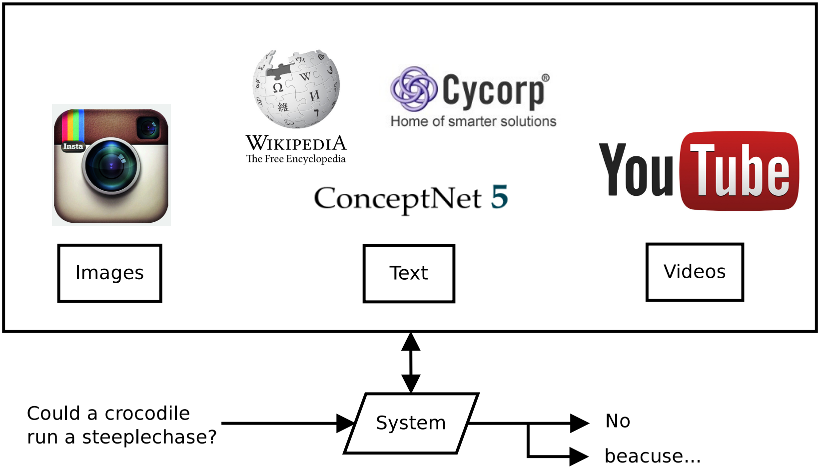

Reasoning with Neural Networks

Two main branches:

- Common Sense embeddings.
- Neural Reasoners.

Common Sense embeddings Hinton (1990)

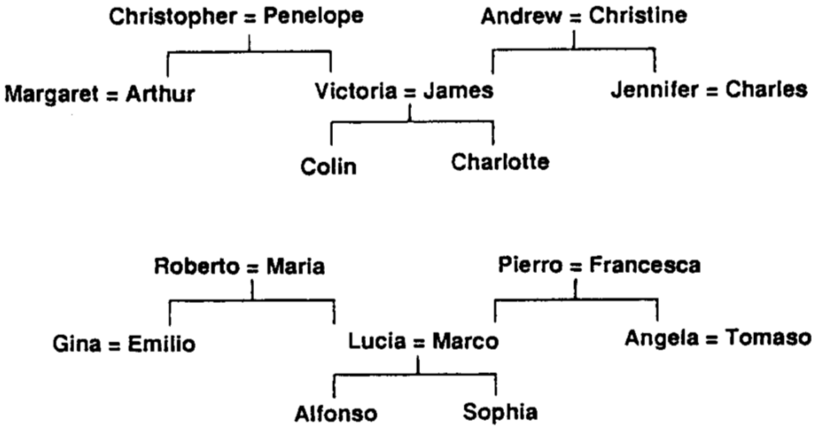

Common Sense embeddings Hinton (1990)

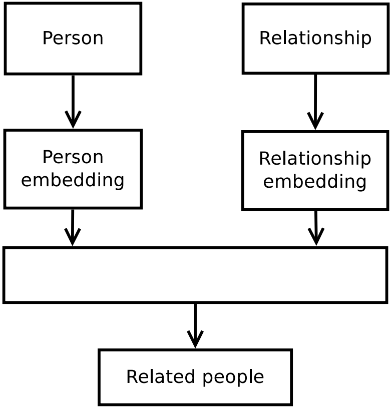

Common Sense embeddings McClelland & Rogers (2003)

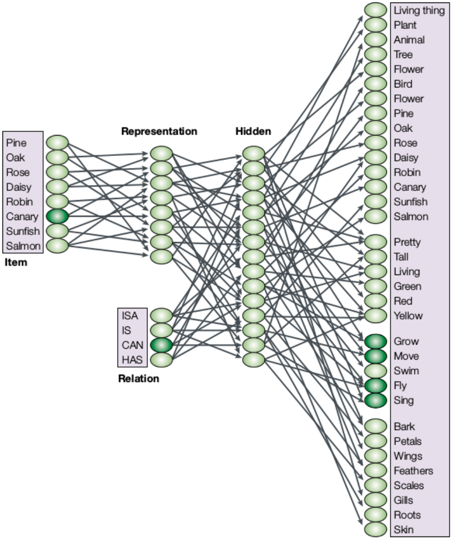

Common Sense embeddings McClelland & Rogers (2003)

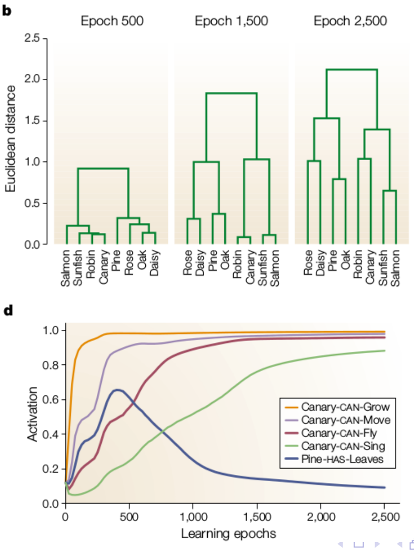

Common Sense embeddings McClelland & Rogers (2003)

DC's delayed copy of a swan

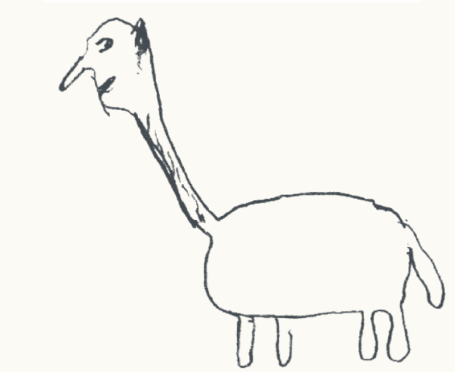

Common Sense embeddings Socher et al. (2013)

Reasoning with neural tensor networks for knowledge base completion.

Common Sense embeddings Socher et al. (2013)

Reasoning with neural tensor networks for knowledge base completion.

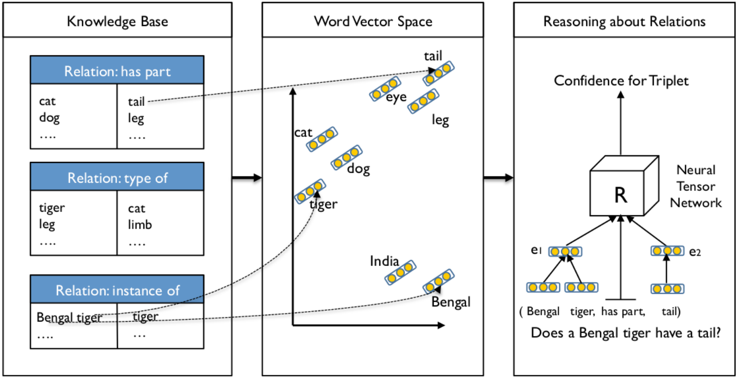

Common Sense embeddings
Socher et al. (2013)

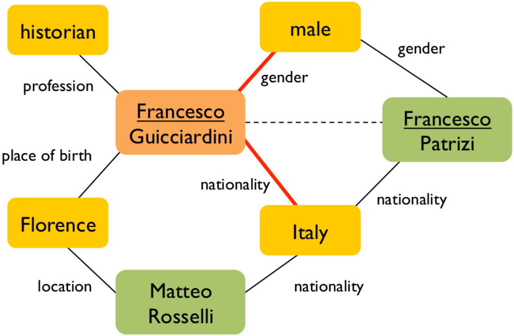

Common Sense embeddings Bowman et al. (2014)

Recursive neural networks can learn logical semantics.

Common Sense embeddings Bowman et al. (2014)

Recursive neural networks can learn logical semantics.

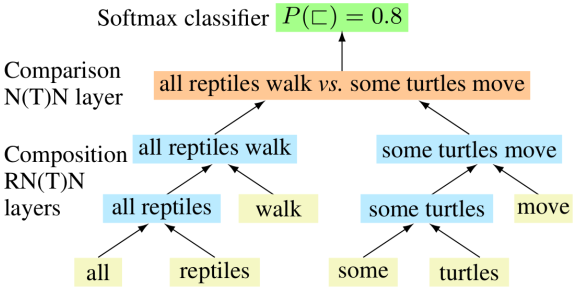

Common Sense embeddings Bowman et al. (2014)

$$\vec{y}_{\text{TreeRNN}} = f\left(\mathbf{M}\begin{bmatrix}\vec{x}^{(l)}\\ \vec{x}^{(r)}\end{bmatrix} + \vec{b}\right)$$

$$\vec{y}_{\text{TreeRNTN}} = \vec{y}_{\text{TreeRNN}} + f(\vec{x}^{(l)T}\mathbf{T}^{[1\ldots n]}\vec{x}^{(r)})$$

Common Sense embeddings
Bowman et al. (2014)

Name	Symbol	Set-theoretic definition	Example
(strict) entailment	$x \sqsubset y$	$x \subset y$	*turtle, reptile*
(strict) reverse entailment	$x \sqsupset y$	$x \supset y$	*reptile, turtle*
equivalence	$x \equiv y$	$x = y$	*couch, sofa*
alternation	$x \mid y$	$x \cap y = \emptyset \wedge x \cup y \neq \mathcal{D}$	*turtle, warthog*
negation	$x \wedge y$	$x \cap y = \emptyset \wedge x \cup y = \mathcal{D}$	*able, unable*
cover	$x \smile y$	$x \cap y \neq \emptyset \wedge x \cup y = \mathcal{D}$	*animal, non-turtle*
independence	$x \# y$	(else)	*turtle, pet*

Common Sense embeddings
Bowman et al. (2014)

	$\equiv$	$\sqsubset$	$\sqsupset$	$\wedge$	$\mid$	$\smile$	$\#$
$\equiv$	$\equiv$	$\sqsubset$	$\sqsupset$	$\wedge$	$\mid$	$\smile$	$\#$
$\sqsubset$	$\sqsubset$	$\sqsubset$	·	$\mid$	$\mid$	·	·
$\sqsupset$	$\sqsupset$	·	$\sqsupset$	$\smile$	·	$\smile$	·
$\wedge$	$\wedge$	$\smile$	$\mid$	$\equiv$	$\sqsupset$	$\sqsubset$	$\#$
$\mid$	$\mid$	·	$\mid$	$\sqsubset$	·	$\sqsubset$	·
$\smile$	$\smile$	$\smile$	·	$\sqsupset$	$\sqsupset$	·	·
$\#$	$\#$	·	·	$\#$	·	·	·

Common Sense embeddings Bowman et al. (2014)

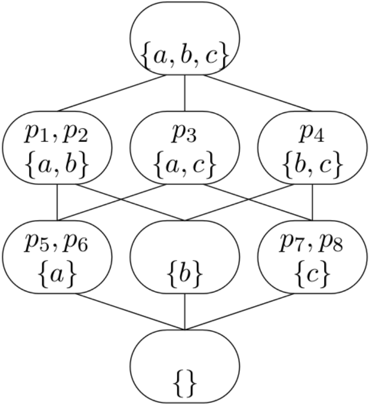

Common Sense embeddings Bowman et al. (2014)

Train	Test
$p_1 \equiv p_2$	$p_2 \wedge p_7$
$p_1 \sqsupset p_5$	$p_2 \sqsupset p_5$
$p_4 \sqsupset p_8$	~~$p_5 \equiv p_6$~~
$p_5 \mid p_7$	~~$p_7 \sqsubset p_4$~~
$p_7 \wedge p_1$	$p_8 \sqsubset p_4$

Common Sense embeddings
Bowman et al. (2014)

Train	Test
$p_1 \equiv p_2$	$p_2 \wedge p_7$
$p_1 \sqsupset p_5$	$p_2 \sqsupset p_5$
$p_4 \sqsupset p_8$	~~$p_5 \equiv p_6$~~
$p_5 \mid p_7$	~~$p_7 \sqsubset p_4$~~
$p_7 \wedge p_1$	$p_8 \sqsubset p_4$

	Train	Test
# only	53.8 (10.5)	53.8 (10.5)
15d NN	99.8 (99.0)	94.0 (87.0)
15d NTN	**100** **(100)**	**99.6** **(95.5)**

Common Sense embeddings Bowman et al. (2014)

$not\ p_3$	$\wedge$	p_3
$not\ not\ p_6$	$\equiv$	p_6
p_3	$\sqsubset$	$(p_3\ or\ p_2)$
$(p_1\ or\ (p_2\ or\ p_4))$	$\sqsupset$	$(p_2\ and\ not\ p_4)$
$not\ (not\ p_1\ and\ not\ p_2)$	$\equiv$	$(p_1\ or\ p_2)$

Common Sense embeddings Bowman et al. (2014)

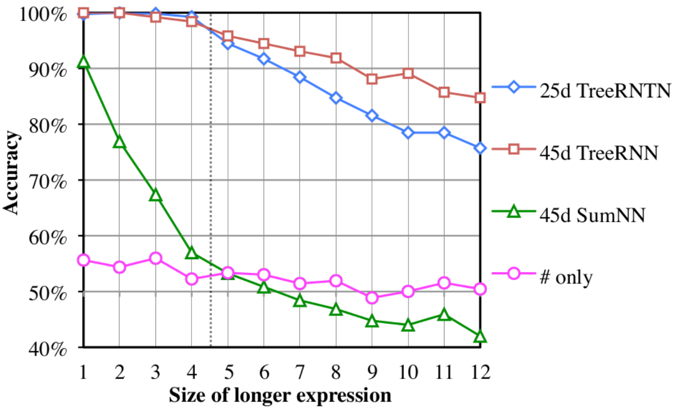

Common Sense embeddings Bowman et al. (2014)

(most turtle) swim | (no turtle) move

(all lizard) reptile $\sqsubset$ (some lizard) animal

(most turtle) reptile | (all turtle) (not animal)

Common Sense embeddings Bowman et al. (2014)

(most turtle) swim | (no turtle) move

(all lizard) reptile ⊏ (some lizard) animal

(most turtle) reptile | (all turtle) (not animal)

	Train		Test	
# only	35.4	(7.5)	35.4	(7.5)
25d SumNN	96.9	(97.7)	93.9	(95.0)
25d TreeRNN	99.6	(99.6)	99.2	(99.3)
25d TreeRNTN	**100**	**(100)**	**99.7**	**(99.5)**

Common Sense embeddings Bowman et al. (2014)

SICK textual entailment challenge

The patient is being helped by the doctor	*entailment*	The doctor is helping the patient (PASSIVE)
A little girl is playing the violin on a beach	*contradiction*	There is no girl playing the violin on a beach (NEG)
The yellow dog is drinking water from a bottle	*contradiction*	The yellow dog is drinking water from a pot (SUBST)
A woman is breaking two eggs in a bowl	*neutral*	A man is mixing a few ingredients in a bowl (MULTIED)
Dough is being spread by a man	*neutral*	A woman is slicing meat with a knife (DIFF)

Common Sense embeddings
Bowman et al. (2014)

	neutral only	30d SumNN	30d TrRNN	50d TrRNTN
DG Train	50.0	68.0	67.0	**74.0**
SICK Train	56.7	96.6	95.4	**97.8**
SICK Test	56.7	73.4	74.9	**76.9**
PASSIVE (4%)	0	76	68	**88**
NEG (7%)	0	96	**100**	**100**
SUBST (24%)	28	**72**	64	**72**
MULTIED (39%)	**68**	61	66	64
DIFF (26%)	**96**	68	79	**96**
SHORT (47%)	50.0	73.9	73.5	**77.3**

Reasoning about facts

Reasoning about facts

The bAbI project (Weston et al. (2015)).

Task 15: Basic Deduction
Sheep are afraid of wolves.
Cats are afraid of dogs.
Mice are afraid of cats.
Gertrude is a sheep.
What is Gertrude afraid of? A:wolves

Task 16: Basic Induction
Lily is a swan.
Lily is white.
Bernhard is green.
Greg is a swan.
What color is Greg? A:white

Task 17: Positional Reasoning
The triangle is to the right of the blue square.
The red square is on top of the blue square.
The red sphere is to the right of the blue square.
Is the red sphere to the right of the blue square? A:yes
Is the red square to the left of the triangle? A:yes

Task 18: Size Reasoning
The football fits in the suitcase.
The suitcase fits in the cupboard.
The box is smaller than the football.
Will the box fit in the suitcase? A:yes
Will the cupboard fit in the box? A:no

Reasoning about facts

Three models have been proposed:

- Dynamic Networks (Kumar et al. (2015))
- Memory Networks (Sukhbaatar et al. (2015))
- Neural Reasoner (Peng et al. (2015))

Reasoning about facts

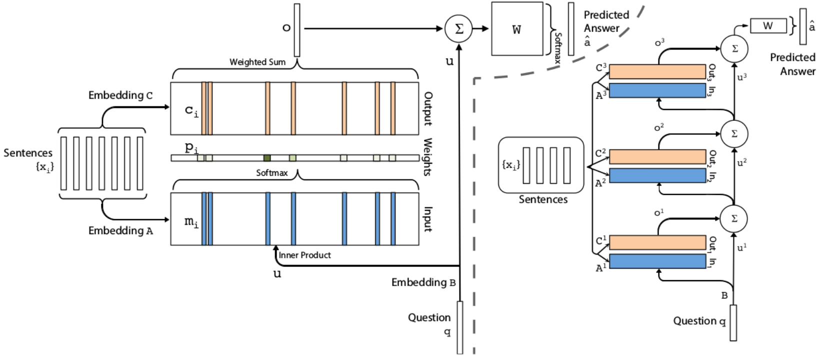

Credit: Sukhbaatar et al. (2015)

Reasoning about facts

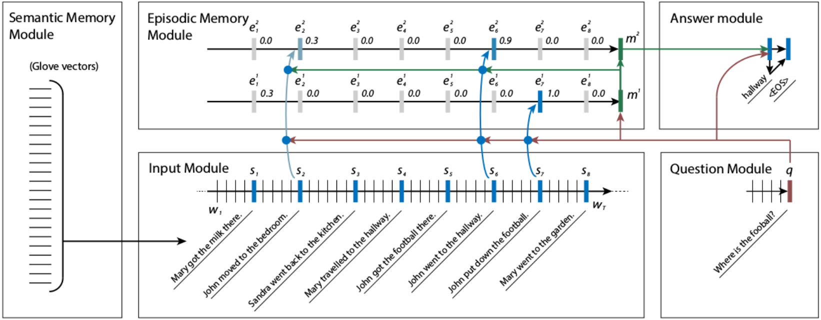

Credit: Kumar et al. (2015)

Reasoning about facts

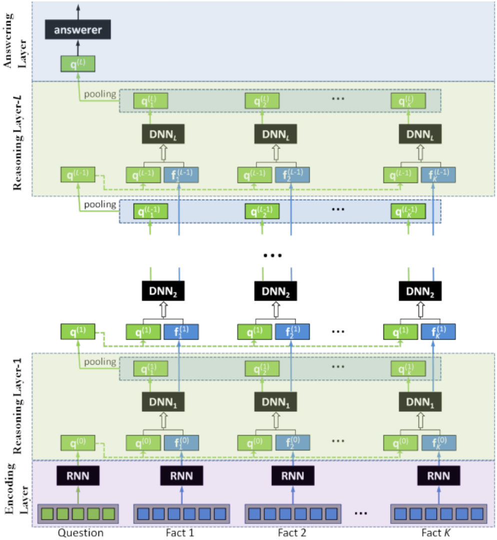

Credit: Peng et al. (2015)

Reasoning about facts

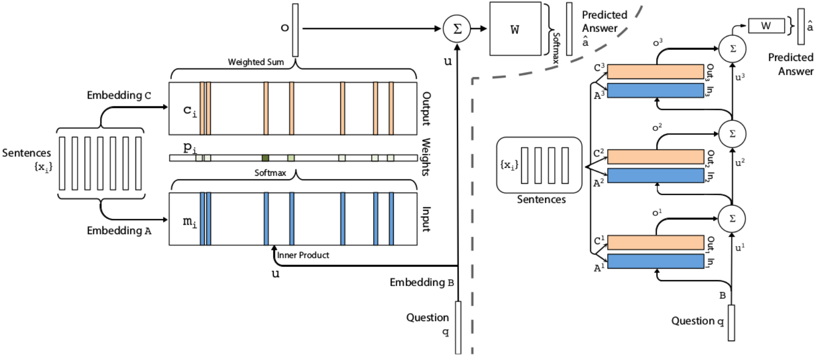

Credit: Sukhbaatar et al. (2015)

Reasoning about facts

SLD resolution.

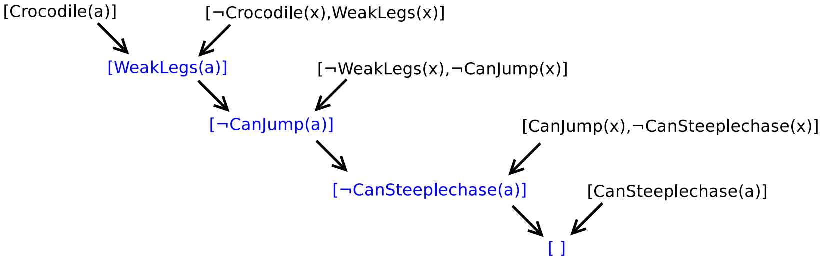

Reasoning about facts
Testing Memory Networks

Facts

Facts
mice are afraid of sheep
wolves are afraid of cats
jessica is a wolf
sheep are afraid of cats
winona is a mouse
cats are afraid of mice
gertrude is a cat
emily is a wolf

Questions
what is jessica afraid of?

Reasoning about facts
Testing Memory Networks

Facts
mice are afraid of sheep
wolves are afraid of cats
jessica is a wolf
sheep are afraid of cats
winona is a mouse
cats are afraid of mice
gertrude is a cat
emily is a wolf

Questions	
what is jessica afraid of?	A: cat (99.74%)

Reasoning about facts
Testing Memory Networks

Facts

mice are afraid of sheep
wolves are afraid of cats
jessica is a wolf
sheep are afraid of cats
winona is a mouse
cats are afraid of mice
gertrude is a cat
emily is a wolf

Questions

what is jessica afraid of? A: cat (99.74%)
is emily afraid of gertrude?

Reasoning about facts
Testing Memory Networks

Facts
mice are afraid of sheep
wolves are afraid of cats
jessica is a wolf
sheep are afraid of cats
winona is a mouse
cats are afraid of mice
gertrude is a cat
emily is a wolf

Questions	
what is jessica afraid of?	A: cat (99.74%)
is emily afraid of gertrude?	A: cat (71.79%)

Reasoning about facts
Testing Memory Networks

Facts

the triangle is to the left of the red square
the pink rectangle is below the triangle

Questions

is the red square to the right of the pink rectangle?

Reasoning about facts
Testing Memory Networks

Facts

the triangle is to the left of the red square
the pink rectangle is below the triangle

Questions

is the red square to the right of the pink rectangle? A: yes (87%)

Reasoning about facts
Testing Memory Networks

Facts

the triangle is to the left of the red square
the pink rectangle is below the triangle

Questions

is the red square to the right of the pink rectangle? A: yes (87%)
is the red square to the left of the pink rectangle?

Reasoning about facts
Testing Memory Networks

Facts
the triangle is to the left of the red square
the pink rectangle is below the triangle

Questions	
is the red square to the right of the pink rectangle?	A: yes (87%)
is the red square to the left of the pink rectangle?	A: yes (92%)

Reasoning about facts
Testing Memory Networks

Facts
sandra and daniel journeyed to the bedroom
john and sandra travelled to the garden
sandra and john travelled to the bedroom
mary and sandra went back to the kitchen
sandra and mary travelled to the bedroom
john and mary moved to the office

Questions
where is daniel?

Reasoning about facts
Testing Memory Networks

Facts
sandra and daniel journeyed to the bedroom
john and sandra travelled to the garden
sandra and john travelled to the bedroom
mary and sandra went back to the kitchen
sandra and mary travelled to the bedroom
john and mary moved to the office

Questions	
where is daniel?	A: bedroom (99.60%)

Reasoning about facts
Testing Memory Networks

Facts
sandra and daniel journeyed to the bedroom
john and sandra travelled to the garden
sandra and john travelled to the bedroom
mary and sandra went back to the kitchen
sandra and mary travelled to the bedroom
john and mary moved to the office

Questions	
where is daniel?	A: bedroom (99.60%)
is daniel in the bedroom?	

Reasoning about facts
Testing Memory Networks

Facts
sandra and daniel journeyed to the bedroom
john and sandra travelled to the garden
sandra and john travelled to the bedroom
mary and sandra went back to the kitchen
sandra and mary travelled to the bedroom
john and mary moved to the office

Questions	
where is daniel?	A: bedroom (99.60%)
is daniel in the bedroom?	A: no (91.38%)

Reasoning about facts

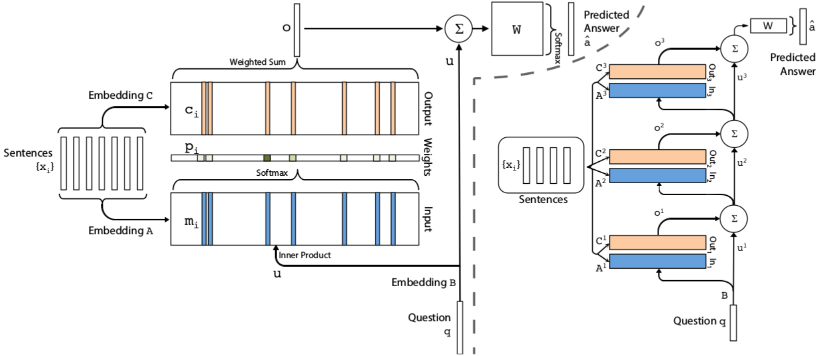

Credit: Sukhbaatar et al. (2015)

Proposals: Explanations

Example 1:

- julius is white.
- What is julius color? White.

Proposals: Explanations

Example 1:

- julius is white.
- What is julius color? White.

Example 2:

- julius is a lion.
- julius is white.
- greg is a lion.
- What is greg color? White.

Questions

References I

Bowman, S. R., Potts, C., & Manning, C. D. (2014). Recursive neural networks can learn logical semantics. *arXiv preprint arXiv:1406.1827*.

Hinton, G. E. (1990). Mapping part-whole hierarchies into connectionist networks. *Artificial Intelligence*, *46*(1), 47–75.

Kiros, R., Zhu, Y., Salakhutdinov, R. R., Zemel, R., Urtasun, R., Torralba, A., & Fidler, S. (2015). Skip-thought vectors. In *Advances in neural information processing systems* (pp. 3276–3284).

Kumar, A., Irsoy, O., Su, J., Bradbury, J., English, R., Pierce, B., … Socher, R. (2015). Ask me anything: Dynamic memory networks for natural language processing. *arXiv preprint arXiv:1506.07285*.

References II

Levesque, H. J. (2014). On our best behaviour. *Artificial Intelligence*, *212*, 27–35.

McClelland, J. L., & Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience*, *4*(4), 310–322.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Peng, B., Lu, Z., Li, H., & Wong, K.-F. (2015). Towards neural network-based reasoning. *arXiv preprint arXiv:1508.05508*.

Socher, R., Chen, D., Manning, C. D., & Ng, A. (2013). Reasoning with neural tensor networks for knowledge base completion. In *Advances in neural information processing systems* (pp. 926–934).

References III

Sukhbaatar, S., Weston, J., Fergus, R., et al. (2015). End-to-end memory networks. In *Advances in neural information processing systems* (pp. 2431–2439).

Vinyals, O., & Le, Q. (2015). A neural conversational model. *arXiv preprint arXiv:1506.05869*.

Weston, J., Bordes, A., Chopra, S., & Mikolov, T. (2015). Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.