CS885 Reinforcement Learning
Lecture 19c: July 4, 2018

Memory Augmented Networks

[GBC] Chap. 10

Memory Augmented Networks

- LSTM (Long Short Term Memory)
  - Hochreiter, Schmidhuber (1997)
- Attention Mechanism
- Addressable memory
  - End-to-end Memory Networks (Sukhbaatar, Szlam, Weston, Fergus; NIPS 2015)
- General read/write memory
  - Differentiable Neural Computers (Graves et al.; Nature 2016)

Long Short Term Memory (LSTM)

- Special gated structure to control memorization and forgetting in RNNs
- Mitigate gradient vanishing
- Facilitate long term memory

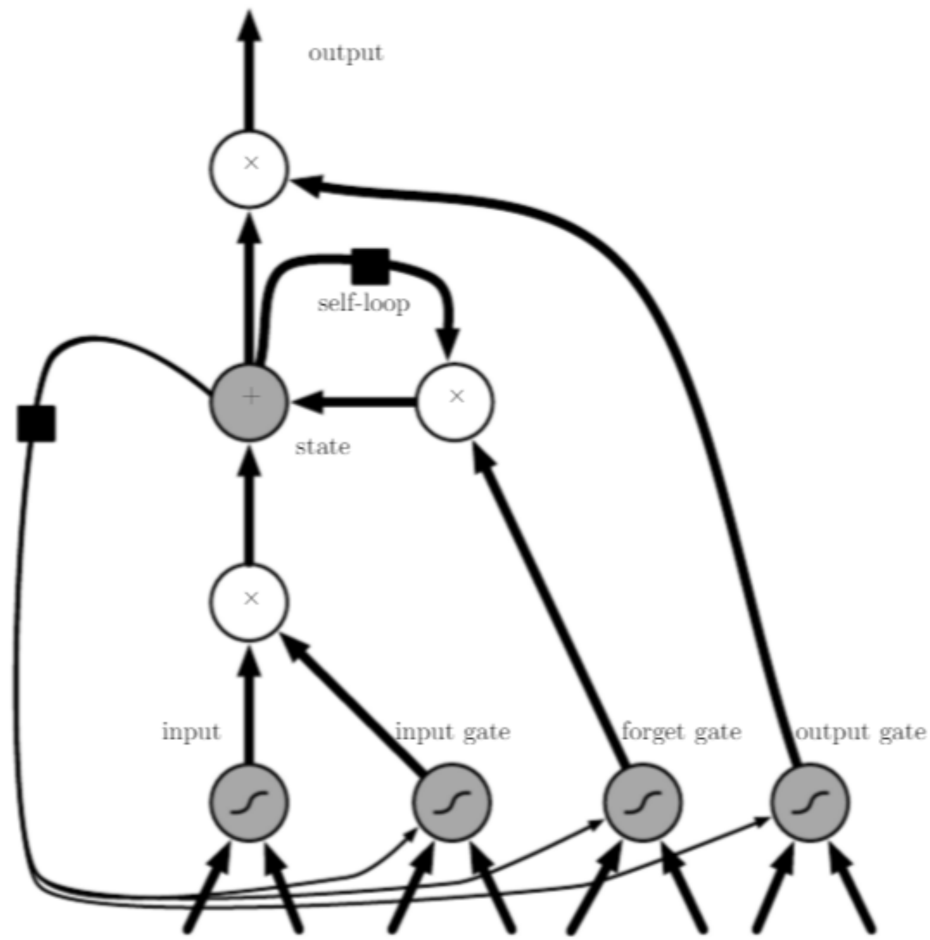

Unrolled long short term memory

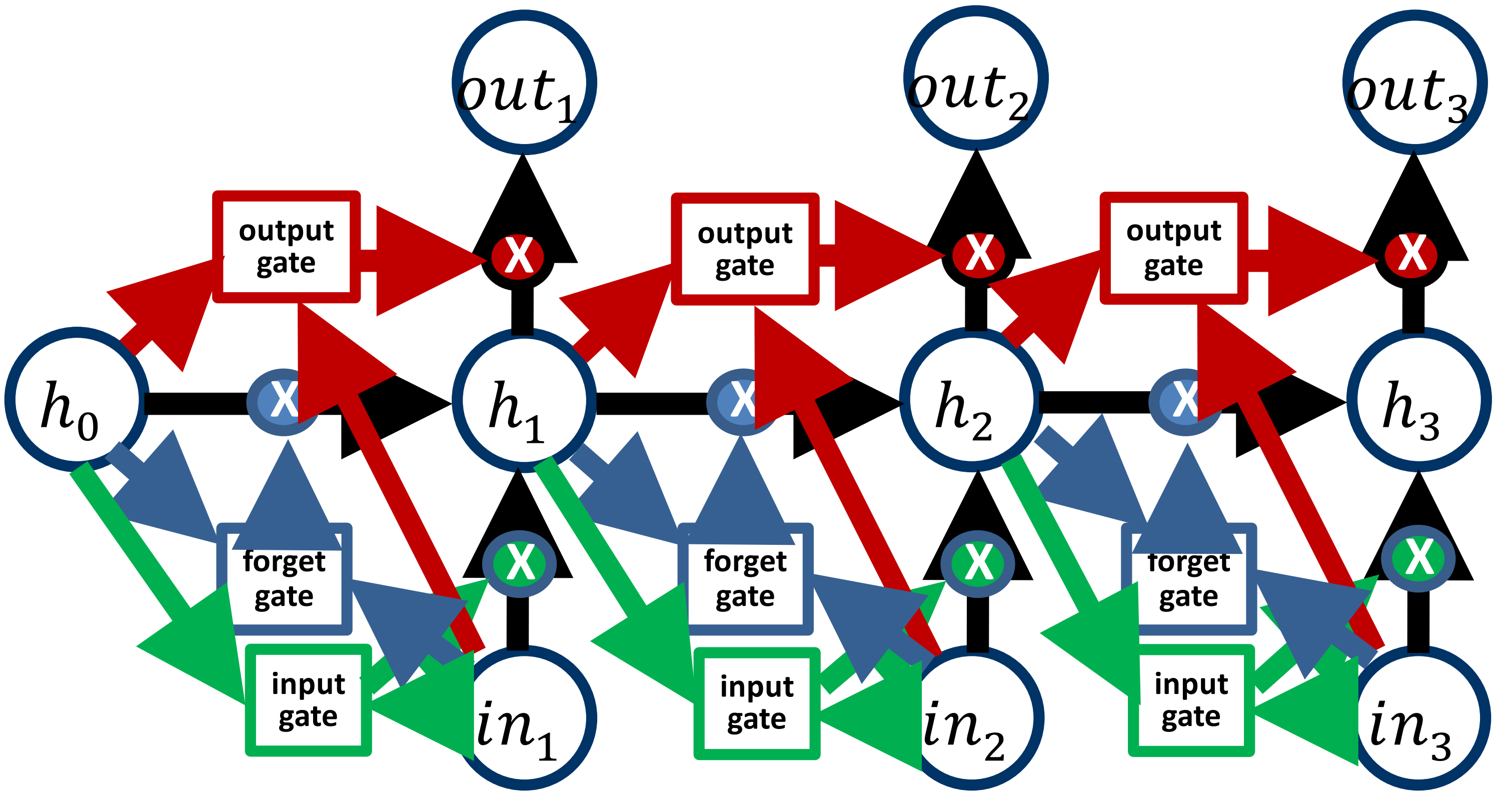

Encoder-Decoder Model

- Also known as sequence2sequence
  - $x^{(i)}$: i^{th} input
  - $y^{(i)}$: i^{th} output
  - c: context (embedding)

- Usage:
  - Machine translation
  - Question answering
  - Dialog

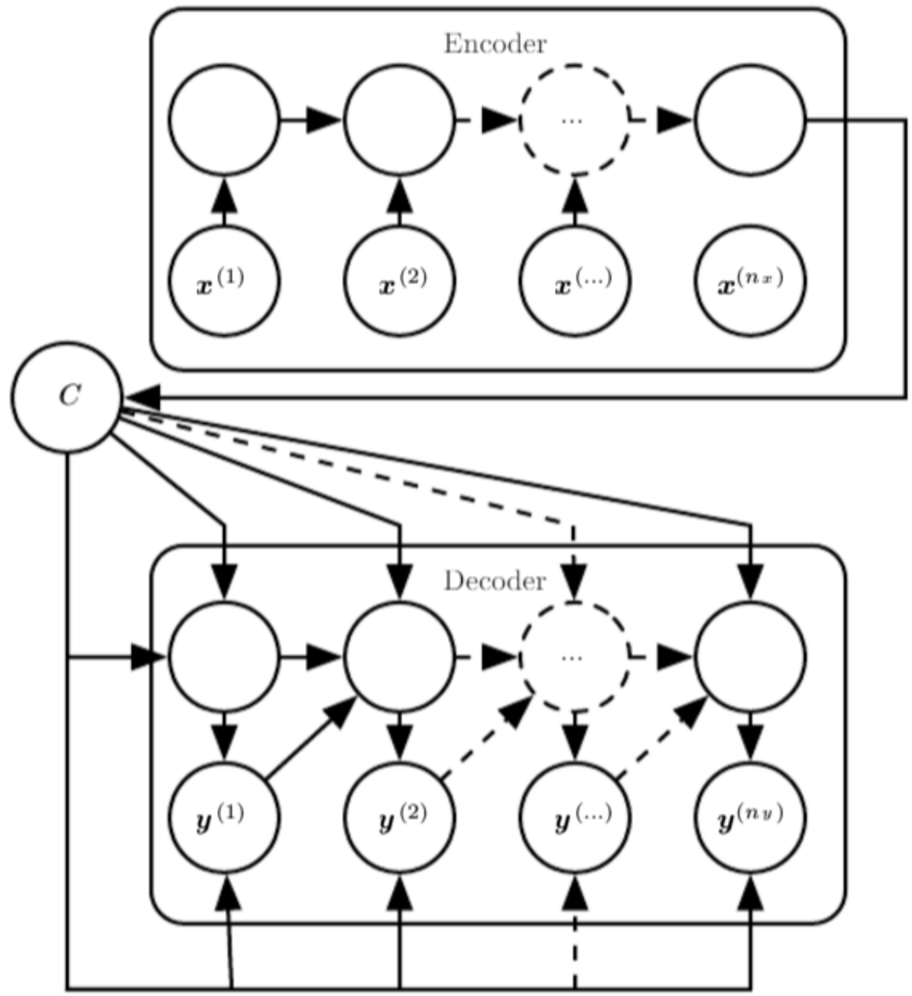

Machine Translation

- Cho, van Merrienboer, Gulcehre, Bahdanau, Bougares, Schwenk, Bengio (2014) Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation

Source	Translation Model	RNN Encoder–Decoder
at the end of the	[a la fin de la] [ŕ la fin des années] [être supprimés à la fin de la]	[à la fin du] [à la fin des] [à la fin de la]
for the first time	[r © pour la premirëre fois] [été donnés pour la première fois] [été commémorée pour la première fois]	[pour la première fois] [pour la première fois ,] [pour la première fois que]
in the United States and	[? aux ?tats-Unis et] [été ouvertes aux États-Unis et] [été constatées aux États-Unis et]	[aux Etats-Unis et] [des Etats-Unis et] [des États-Unis et]
, as well as	[?s , qu'] [?s , ainsi que] [?re aussi bien que]	[, ainsi qu'] [, ainsi que] [, ainsi que les]
one of the most	[?t ?l' un des plus] [?l' un des plus] [être retenue comme un de ses plus]	[l' un des] [le] [un des]

Attention

- Mechanism for alignment in machine translation, image captioning, memory addressing, etc.
- Attention in machine translation: align each output word with relevant input words by computing a softmax of the inputs
  - Context vector c_i: weighted sum of input encodings h_j

    $$c_i = \sum_j a_{ij} h_j$$

  - Where a_{ij} is an alignment weight between input encoding h_j and output encoding s_i

$$a_{ij} = \frac{\exp(alignment(s_i, h_j))}{\sum_{j'} \exp(alignment(s_i, h_{j'}))} \quad \text{(softmax)}$$

  - Alignment example: $alignment(s_i, h_j) = s_i^T h_j$

Attention

- Picture

Machine Translation with LSTM units and attention

- Bahdanau, Cho, Bengio (ICLR-2015)

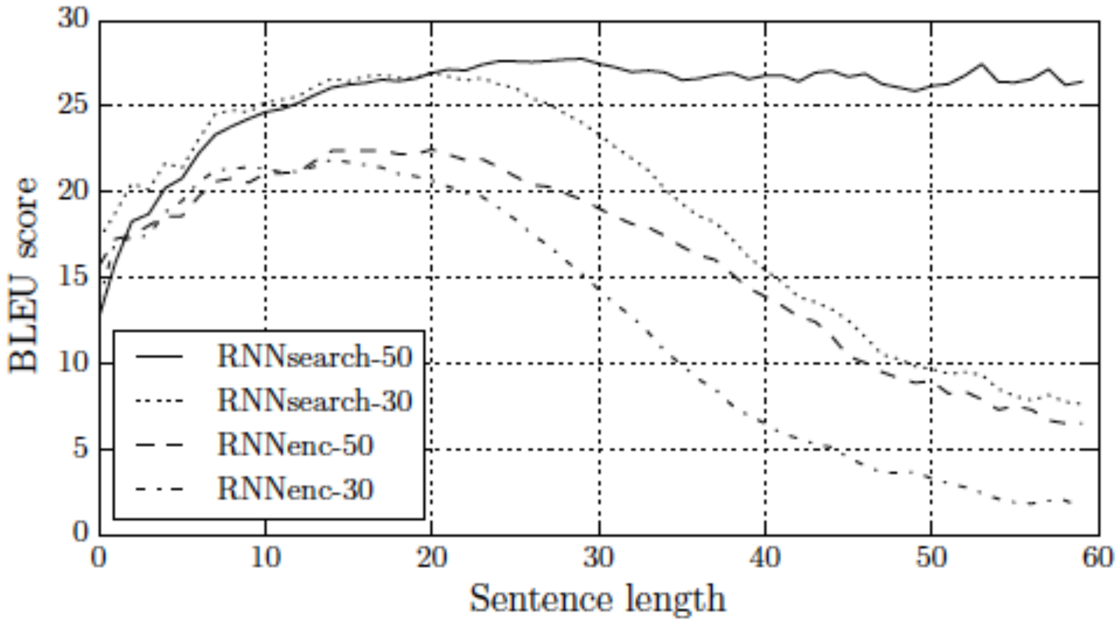

RNNsearch: with attention
RNNenc: no attention

- Bleu: BiLingual Evaluation Understudy
  – Percentage of words in common between translation and ground truth

Alignment example

- Bahdanau, Cho, Bengio (ICLR-2015)

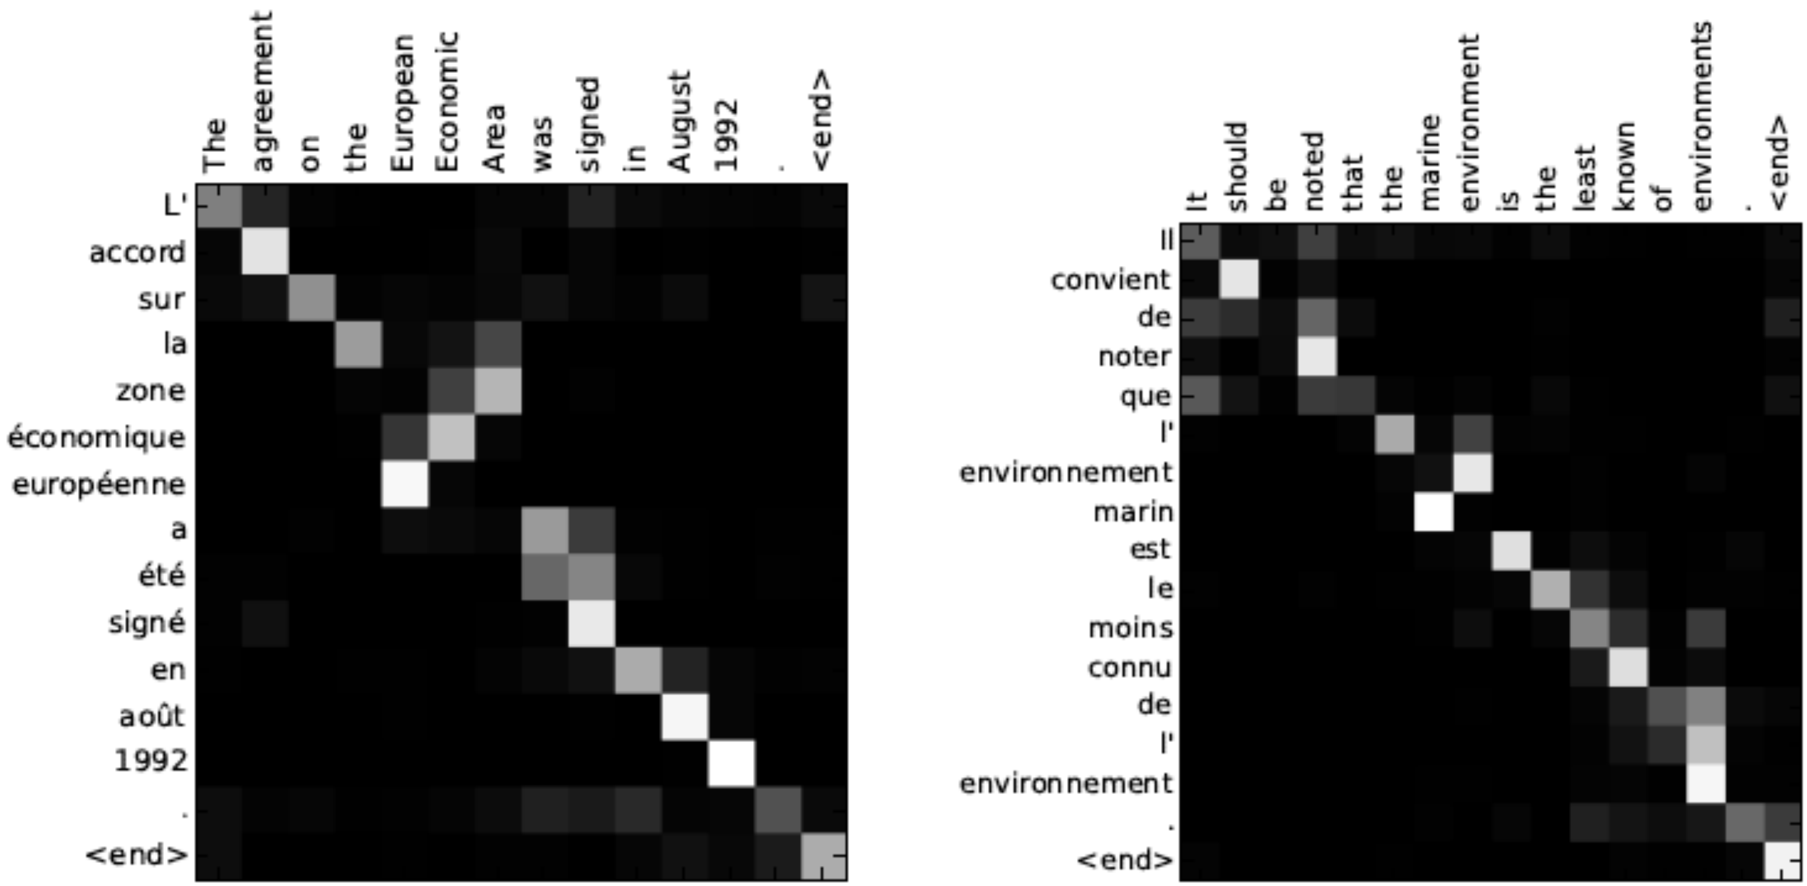

Question Answering / Dialog Systems

- Suppose we have a database of message-response pairs
  - Store database in memory
  - Key-value pairs: embeddings of message-response pairs (m_i, r_i)
- Use attention mechanism to answer query
  - Embed query: q
  - Measure alignment of query with each message: $a_i = q^T m_i$
  - Compute softmax distribution: $p_i = \exp(a_i) / \sum_j \exp(a_j)$
  - Compute response: $r = \sum_i p_i r_i$
  - Decode response

Reading Comprehension

- End-to-end memory networks (Sukhbaatar, Szlam, Weston, Fergus; NIPS 2015)

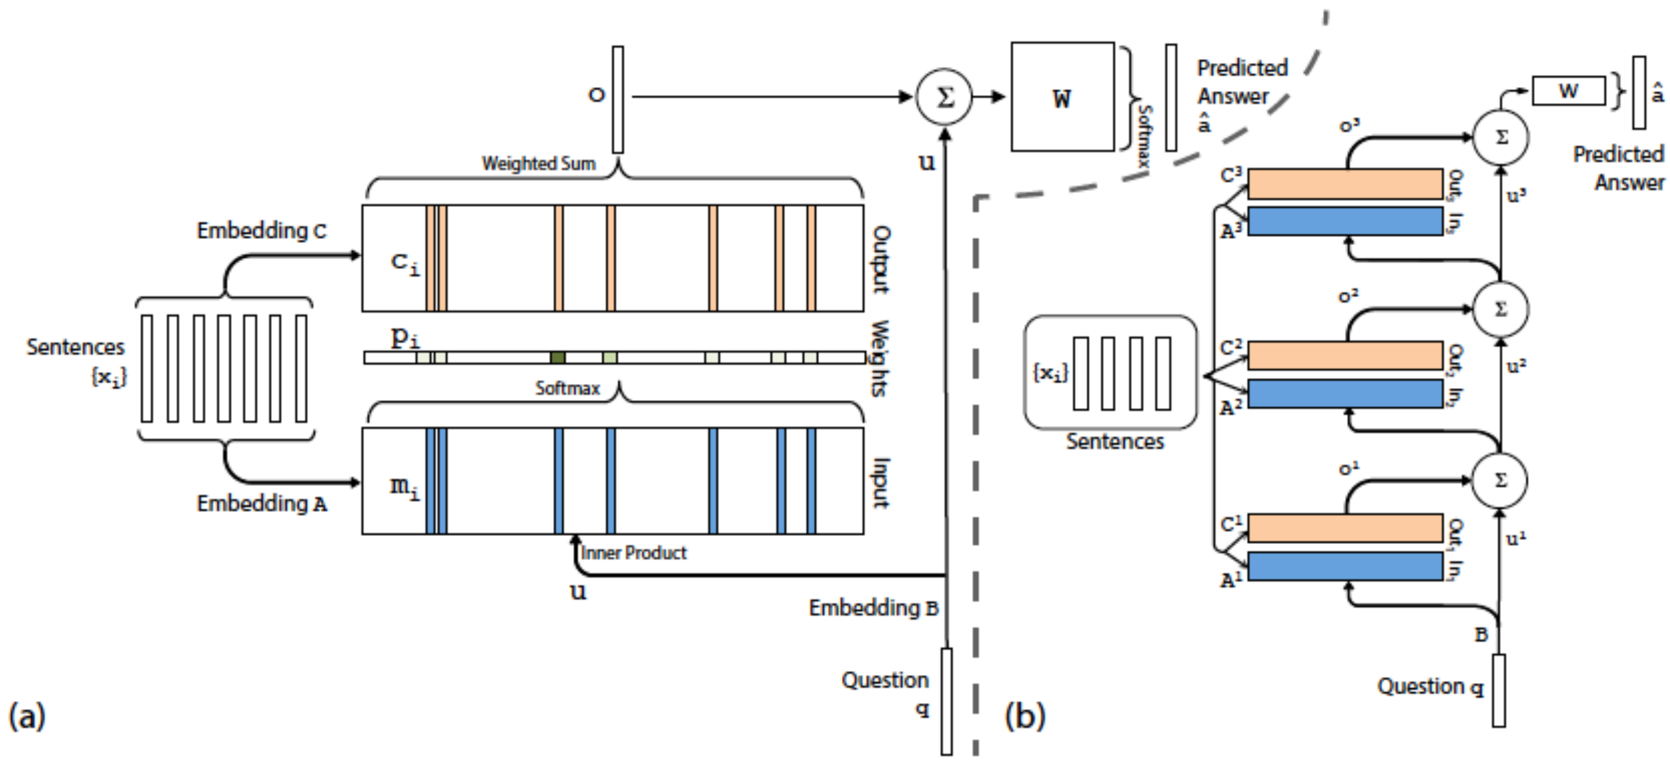

Reading Comprehension

- End-to-end memory networks (Sukhbaatar, Szlam, Weston, Fergus; NIPS 2015)

Story (1: 1 supporting fact)	Support	Hop 1	Hop 2	Hop 3
Daniel went to the bathroom.		0.00	0.00	0.03
Mary travelled to the hallway.		0.00	0.00	0.00
John went to the bedroom.		0.37	0.02	0.00
John travelled to the bathroom.	yes	0.60	0.98	0.96
Mary went to the office.		0.01	0.00	0.00
Where is John? Answer: bathroom Prediction: bathroom				

Story (2: 2 supporting facts)	Support	Hop 1	Hop 2	Hop 3
John dropped the milk.		0.06	0.00	0.00
John took the milk there.	yes	0.88	1.00	0.00
Sandra went back to the bathroom.		0.00	0.00	0.00
John moved to the hallway.	yes	0.00	0.00	1.00
Mary went back to the bedroom.		0.00	0.00	0.00
Where is the milk? Answer: hallway Prediction: hallway				

Story (16: basic induction)	Support	Hop 1	Hop 2	Hop 3
Brian is a frog.	yes	0.00	0.98	0.00
Lily is gray.		0.07	0.00	0.00
Brian is yellow.	yes	0.07	0.00	1.00
Julius is green.		0.06	0.00	0.00
Greg is a frog.	yes	0.76	0.02	0.00
What color is Greg? Answer: yellow Prediction: yellow				

Story (18: size reasoning)	Support	Hop 1	Hop 2	Hop 3
The suitcase is bigger than the chest.	yes	0.00	0.88	0.00
The box is bigger than the chocolate.		0.04	0.05	0.10
The chest is bigger than the chocolate.	yes	0.17	0.07	0.90
The chest fits inside the container.		0.00	0.00	0.00
The chest fits inside the box.		0.00	0.00	0.00
Does the suitcase fit in the chocolate? Answer: no Prediction: no				

General read/write memory

- Replace hidden units by addressable memory
- Replace output gate by attention mechanism
- Generalize input and forget gates to vectors that perform specific operations on each record in the memory