Deep Learning for Natural Language Processing

Stephen Clark et al.
University of Cambridge and DeepMind

DeepMind

UNIVERSITY OF CAMBRIDGE

11. Machine Comprehension

Edward Grefenstette
DeepMind & UCL

What is machine comprehension?

- Too general a question to answer? `<end of lecture>`
- Being able to retrieve information from a text?
- Being able detect the type of relation between sentences?
- We'll explore a few of these in this lecture.

DeepMind

UNIVERSITY OF CAMBRIDGE

What do we care about in Natural Language Understanding?

- Machines that can **understand us** (👈 machine comp.?)
- Machines that **we can understand**

How do we test this?

DeepMind

UNIVERSITY OF CAMBRIDGE

Reading Comprehension (RC)

- Candidate **reads** a document
- Candidate sees a **question about the document**
- Candidate must **select/generate an answer**
  - Span selection
  - Multiple Choice
  - Free form answers

DeepMind

UNIVERSITY OF CAMBRIDGE

RC Example: MCTest (MSR)

James the Turtle was always getting in trouble. Sometimes he'd reach into the freezer and empty out all the food. Other times he'd sled on the deck and get a splinter. His aunt Jane tried as hard as she could to keep him out of trouble, but he was sneaky and got into lots of trouble behind her back. One day, James thought he would go into town and see what kind of trouble he could get into. He went to the grocery store and pulled all the pudding off the shelves and ate two jars. Then he walked to the fast food restaurant and ordered 15 bags of fries. He didn't pay, and instead headed home.

Where did James go after he went to the grocery store?

1. his deck
2. his freezer
3. a fast food restaurant
4. his room

DeepMind

UNIVERSITY OF CAMBRIDGE

RC Example: MCTest (MSR)

What's great about it?

- Real data
- Relatively hard questions (many distractors)
- Some need for paraphrase, coreference resolution

What could be better?

- Little data (<3k questions from <1k articles)
- Multiple choice vs. free form.

DeepMind

UNIVERSITY OF CAMBRIDGE

RC Example: bAbI QA (FAIR)

John picked up the apple.
John went to the office.
John went to the kitchen.
John dropped the apple.

Q: Where was the apple before the kitchen?

A: office

DeepMind

UNIVERSITY OF CAMBRIDGE

RC Example: bAbI QA (FAIR)

What's so great about it?

- Lots of data (can generate more)
- Many sub-tasks based on modes of reasoning

What could be better?

- Synthetic data (better as a unit test)
- Very predictable structure
- Lack of grammatical or lexical diversity

DeepMind

UNIVERSITY OF CAMBRIDGE

RC Example: CNN/DailyMail😱 (DeepMind)

The BBC producer allegedly struck by Jeremy Clarkson will not press charges against the “Top Gear” host, his lawyer said Friday. Clarkson, who hosted one of the most-watched television shows in the world, was dropped by the BBC Wednesday after an internal investigation by the British broadcaster found he had subjected producer Oisin Tymon “to an unprovoked physical and verbal attack.” …

Q: Producer *X* will not press charges against Jeremy Clarkson, his lawyer says.

A: Oisin Tymon

DeepMind

UNIVERSITY OF CAMBRIDGE

RC Example: CNN/DailyMail😱 (DeepMind)

What's so great about it?

- Lots of data (easy to generate more)
- Natural text (Cloze-questions from abstractive summaries)

What could be better?

- Semi-synthetic (lack of variety, unanswerable questions)
- Answers basically are just pointers to entities

DeepMind

UNIVERSITY OF CAMBRIDGE

RC Example: SQuAD (Stanford)

In 1271, Kublai Khan imposed the name Great Yuan (Chinese: 大元; pinyin: Dà Yuán; Wade–Giles: Ta-Yüan), establishing the Yuan dynasty. "Dà Yuán" (大元) is from the sentence "大哉乾元" (dà zai Qián Yuán / "Great is Qián, the Primal") in the Commentaries on the Classic of Changes (I Ching) section regarding Qián (乾). The counterpart in Mongolian language was Dai Ön Ulus, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus. In Mongolian, Dai Ön (Great Yuan) is often used in conjunction with the "Yeke Mongghul Ulus" (lit. "Great Mongol State"), resulting in Dai Ön Yeke Mongghul Ulus (Mongolian script:), meaning "Great Yuan Great Mongol State". The Yuan dynasty is also known as the "Mongol dynasty" or "Mongol Dynasty of China", similar to the names "Manchu dynasty" or "Manchu Dynasty of China" for the Qing dynasty. Furthermore, the Yuan is sometimes known as the "Empire of the Great Khan" or "Khanate of the Great Khan", which particularly appeared on some Yuan maps, since Yuan emperors held the nominal title of Great Khan. Nevertheless, both terms can also refer to the khanate within the Mongol Empire directly ruled by Great Khans before the actual establishment of the Yuan dynasty by Kublai Khan in 1271.

Q: What writing inspired the name Great Yuan?
A (spans):
the Commentaries on the Classic of Changes (I Ching)
the Commentaries on the Classic of Changes

Q: What was the Yuan dynasty called in Mongolian?
A (spans):
Dai Ön Ulus, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus
Dai Ön Ulus

RC Example: SQuAD (Stanford)

In 1271, Kublai Khan imposed the name Great Yuan (Chinese: 大元; pinyin: Dà Yuán; Wade–Giles: Ta-Yüan), establishing the Yuan dynasty. "Dà Yuán" (大元) is from the sentence "大哉乾元" (dà zai Qián Yuán / "Great is Qián, the Primal") in **the Commentaries on the Classic of Changes (I Ching)** section regarding Qián (乾). The counterpart in Mongolian language was Dai Ön Ulus, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus. In Mongolian, Dai Ön (Great Yuan) is often used in conjunction with the "Yeke Mongghul Ulus" (lit. "Great Mongol State"), resulting in Dai Ön Yeke Mongghul Ulus (Mongolian script:), meaning "Great Yuan Great Mongol State". The Yuan dynasty is also known as the "Mongol dynasty" or "Mongol Dynasty of China", similar to the names "Manchu dynasty" or "Manchu Dynasty of China" for the Qing dynasty. Furthermore, the Yuan is sometimes known as the "Empire of the Great Khan" or "Khanate of the Great Khan", which particularly appeared on some Yuan maps, since Yuan emperors held the nominal title of Great Khan. Nevertheless, both terms can also refer to the khanate within the Mongol Empire directly ruled by Great Khans before the actual establishment of the Yuan dynasty by Kublai Khan in 1271.

Q: What writing inspired the name Great Yuan?
A (spans):
the Commentaries on the Classic of Changes (I Ching)
the Commentaries on the Classic of Changes

Q: What was the Yuan dynasty called in Mongolian?
A (spans):
Dai Ön Ulus, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus
Dai Ön Ulus

DeepMind

UNIVERSITY OF CAMBRIDGE

RC Example: SQuAD (Stanford)

In 1271, Kublai Khan imposed the name Great Yuan (Chinese: 大元; pinyin: Dà Yuán; Wade–Giles: Ta-Yüan), establishing the Yuan dynasty. "Dà Yuán" (大元) is from the sentence "大哉乾元" (dà zai Qián Yuán / "Great is Qián, the Primal") in **the Commentaries on the Classic of Changes** (I Ching) section regarding Qián (乾). The counterpart in Mongolian language was Dai Ön Ulus, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus. In Mongolian, Dai Ön (Great Yuan) is often used in conjunction with the "Yeke Mongghul Ulus" (lit. "Great Mongol State"), resulting in Dai Ön Yeke Mongghul Ulus (Mongolian script:), meaning "Great Yuan Great Mongol State". The Yuan dynasty is also known as the "Mongol dynasty" or "Mongol Dynasty of China", similar to the names "Manchu dynasty" or "Manchu Dynasty of China" for the Qing dynasty. Furthermore, the Yuan is sometimes known as the "Empire of the Great Khan" or "Khanate of the Great Khan", which particularly appeared on some Yuan maps, since Yuan emperors held the nominal title of Great Khan. Nevertheless, both terms can also refer to the khanate within the Mongol Empire directly ruled by Great Khans before the actual establishment of the Yuan dynasty by Kublai Khan in 1271.

Q: What writing inspired the name Great Yuan?
A (spans):
the Commentaries on the Classic of Changes (I Ching)
the Commentaries on the Classic of Changes

Q: What was the Yuan dynasty called in Mongolian?
A (spans):
Dai Ön Ulus, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus
Dai Ön Ulus

RC Example: SQuAD (Stanford)

In 1271, Kublai Khan imposed the name Great Yuan (Chinese: 大元; pinyin: Dà Yuán; Wade–Giles: Ta-Yüan), establishing the Yuan dynasty. "Dà Yuán" (大元) is from the sentence "大哉乾元" (dà zai Qián Yuán / "Great is Qián, the Primal") in the Commentaries on the Classic of Changes (I Ching) section regarding Qián (乾). The counterpart in Mongolian language was **Dai Ön Ulus, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus**. In Mongolian, Dai Ön (Great Yuan) is often used in conjunction with the "Yeke Mongghul Ulus" (lit. "Great Mongol State"), resulting in Dai Ön Yeke Mongghul Ulus (Mongolian script:), meaning "Great Yuan Great Mongol State". The Yuan dynasty is also known as the "Mongol dynasty" or "Mongol Dynasty of China", similar to the names "Manchu dynasty" or "Manchu Dynasty of China" for the Qing dynasty. Furthermore, the Yuan is sometimes known as the "Empire of the Great Khan" or "Khanate of the Great Khan", which particularly appeared on some Yuan maps, since Yuan emperors held the nominal title of Great Khan. Nevertheless, both terms can also refer to the khanate within the Mongol Empire directly ruled by Great Khans before the actual establishment of the Yuan dynasty by Kublai Khan in 1271.

Q: What writing inspired the name Great Yuan?
A (spans):
the Commentaries on the Classic of Changes (I Ching)
the Commentaries on the Classic of Changes

Q: What was the Yuan dynasty called in Mongolian?
A (spans):
Dai Ön Ulus, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus
Dai Ön Ulus

RC Example: SQuAD (Stanford)

In 1271, Kublai Khan imposed the name Great Yuan (Chinese: 大元; pinyin: Dà Yuán; Wade–Giles: Ta-Yüan), establishing the Yuan dynasty. "Dà Yuán" (大元) is from the sentence "大哉乾元" (dà zai Qián Yuán / "Great is Qián, the Primal") in the Commentaries on the Classic of Changes (I Ching) section regarding Qián (乾). The counterpart in Mongolian language was **Dai Ön Ulus**, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus. In Mongolian, Dai Ön (Great Yuan) is often used in conjunction with the "Yeke Mongghul Ulus" (lit. "Great Mongol State"), resulting in Dai Ön Yeke Mongghul Ulus (Mongolian script:), meaning "Great Yuan Great Mongol State". The Yuan dynasty is also known as the "Mongol dynasty" or "Mongol Dynasty of China", similar to the names "Manchu dynasty" or "Manchu Dynasty of China" for the Qing dynasty. Furthermore, the Yuan is sometimes known as the "Empire of the Great Khan" or "Khanate of the Great Khan", which particularly appeared on some Yuan maps, since Yuan emperors held the nominal title of Great Khan. Nevertheless, both terms can also refer to the khanate within the Mongol Empire directly ruled by Great Khans before the actual establishment of the Yuan dynasty by Kublai Khan in 1271.

Q: What writing inspired the name Great Yuan?
A (spans):
the Commentaries on the Classic of Changes (I Ching)
the Commentaries on the Classic of Changes

Q: What was the Yuan dynasty called in Mongolian?
A (spans):
Dai Ön Ulus, also rendered as Ikh Yuan Üls or Yekhe Yuan Ulus
Dai Ön Ulus

RC Example: SQuAD (Stanford)

What's so great about it?

- Lots of data
- Well-recognised (or well-marketed) benchmark
- Human-produced questions

What could be better?

- Answers are contiguous spans, limiting *sort* of question
- Overfitting is a massive issue

DeepMind

UNIVERSITY OF CAMBRIDGE

What's missing in datasets?

- Lots of natural, human-written, question/answer pairs.
- Questions that require linking information across different parts of the text (i.e. beyond simple anaphora resolution)
- More diverse questions. Not just who/what/where, but:
  - "how?", "in what manner …?".
  - Temporal questions.
  - Questions about abstract relations, narrative structure.

DeepMind

UNIVERSITY OF CAMBRIDGE

Some recent work

- Some recent datasets, e.g. SearchQA (Dunn *et al.* 2017) or NarrativeQA (Kočiský *et al.* 2017) try to incorporate the need to search/aggregate information, answer high-level questions.

- Ultimately, tension between making task feasible (often over-simplifying) and realistic (often too hard).

DeepMind

UNIVERSITY OF CAMBRIDGE

Deep Learning Methods

Multiple-Choice:

1. Embed question given document (or vice versa).
2. Embed answers.
3. Compute similarity. Rank. Select.

DeepMind

UNIVERSITY OF CAMBRIDGE

Deep Learning Methods

Entity answers:

1. Embed document (or document words).
2. Embed question (possibly conditioned on doc).
3. Softmax over entity markers / attention / pointer net

DeepMind

UNIVERSITY OF CAMBRIDGE

Deep Learning Methods

Spans:

1. Enumerate spans. There are O(Length^2). Typically only consider spans within sentences. Optionally further restrict e.g. using constituency parse.
2. Embed spans.
3. Embed question.
4. Compute similarity. Rank. Select.

Cf. models at https://rajpurkar.github.io/SQuAD-explorer/

DeepMind

UNIVERSITY OF CAMBRIDGE

Deep Learning Methods

Incorporating search

Latent boolean decision to search or answer.

No gradient info on this decision, or on search terms.

Must use gradient estimator, e.g. REINFORCE (Williams 1992) to compute gradients of loss w.r.t. discrete decisions.

Free form answers from large documents

Output is just a conditional language model.

Difficulty usually lies elsewhere, but objective is a good question. Minimise NLL, or use other metrics e.g. BLEU/ROUGE/etc?

Textual Entailment (SNLI)

A man is crowd surfing at a concert

- The man is at a football game **Contradiction**
- The man is drunk **Neutral**
- The man is at a concert **Entailment**

A wedding party is taking pictures

- There is a funeral **Contradiction**
- They are outside **Neutral**
- Someone got married **Entailment**

DeepMind

UNIVERSITY OF CAMBRIDGE

Learning to detect entailment

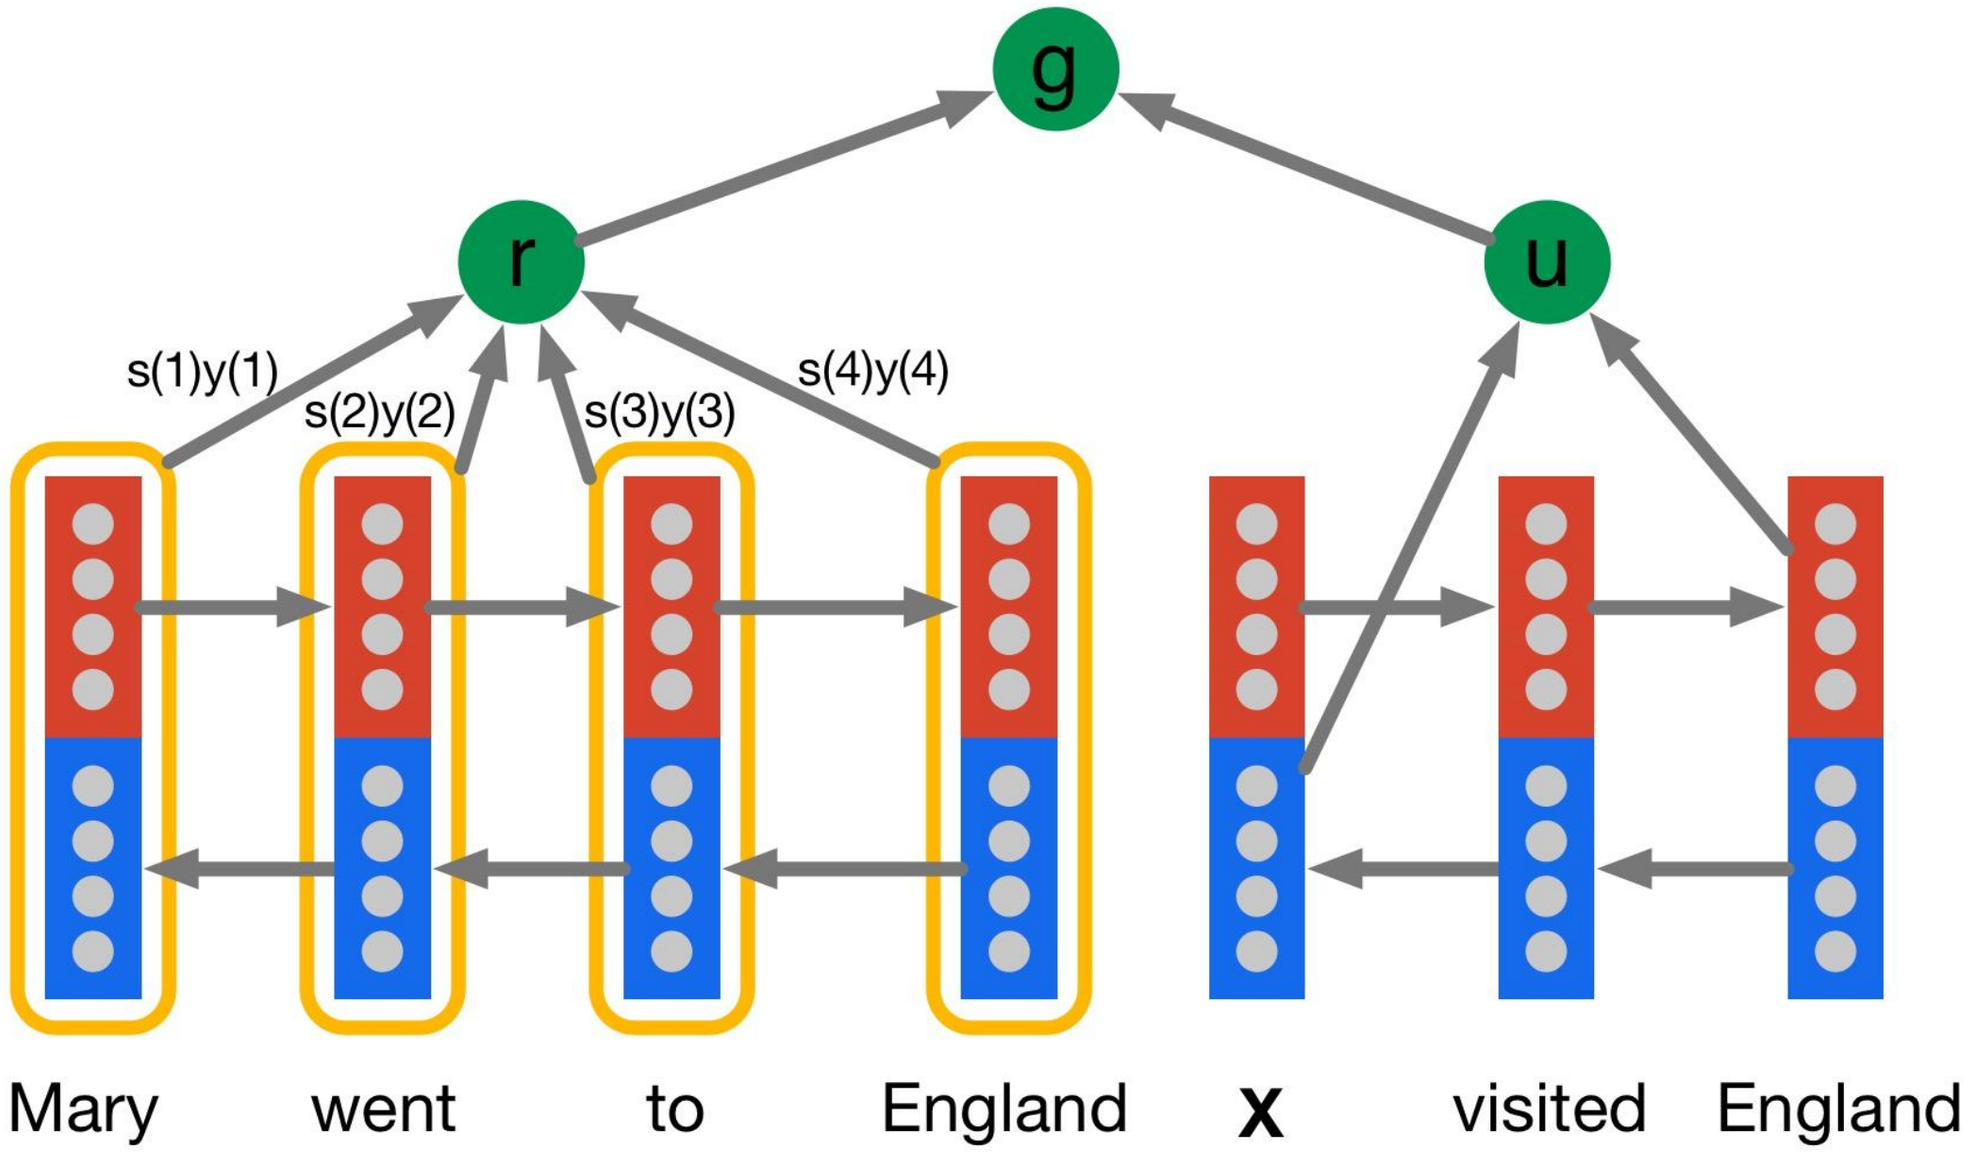

DeepMind

UNIVERSITY OF CAMBRIDGE

Learning to detect entailment

Hypothesis: A girl is wearing a blue jacket.

Premise

A young girl wearing a pink coat plays with a yellow toy golf club.

DeepMind

UNIVERSITY OF CAMBRIDGE

Learning to detect entailment

Hypothesis: Two mimes sit in complete silence.

A man wearing a yellow striped shirt laughs while seated next to another man who is wearing a light blue shirt and clasping his hands together.

Premise

DeepMind

UNIVERSITY OF CAMBRIDGE

Learning to detect entailment

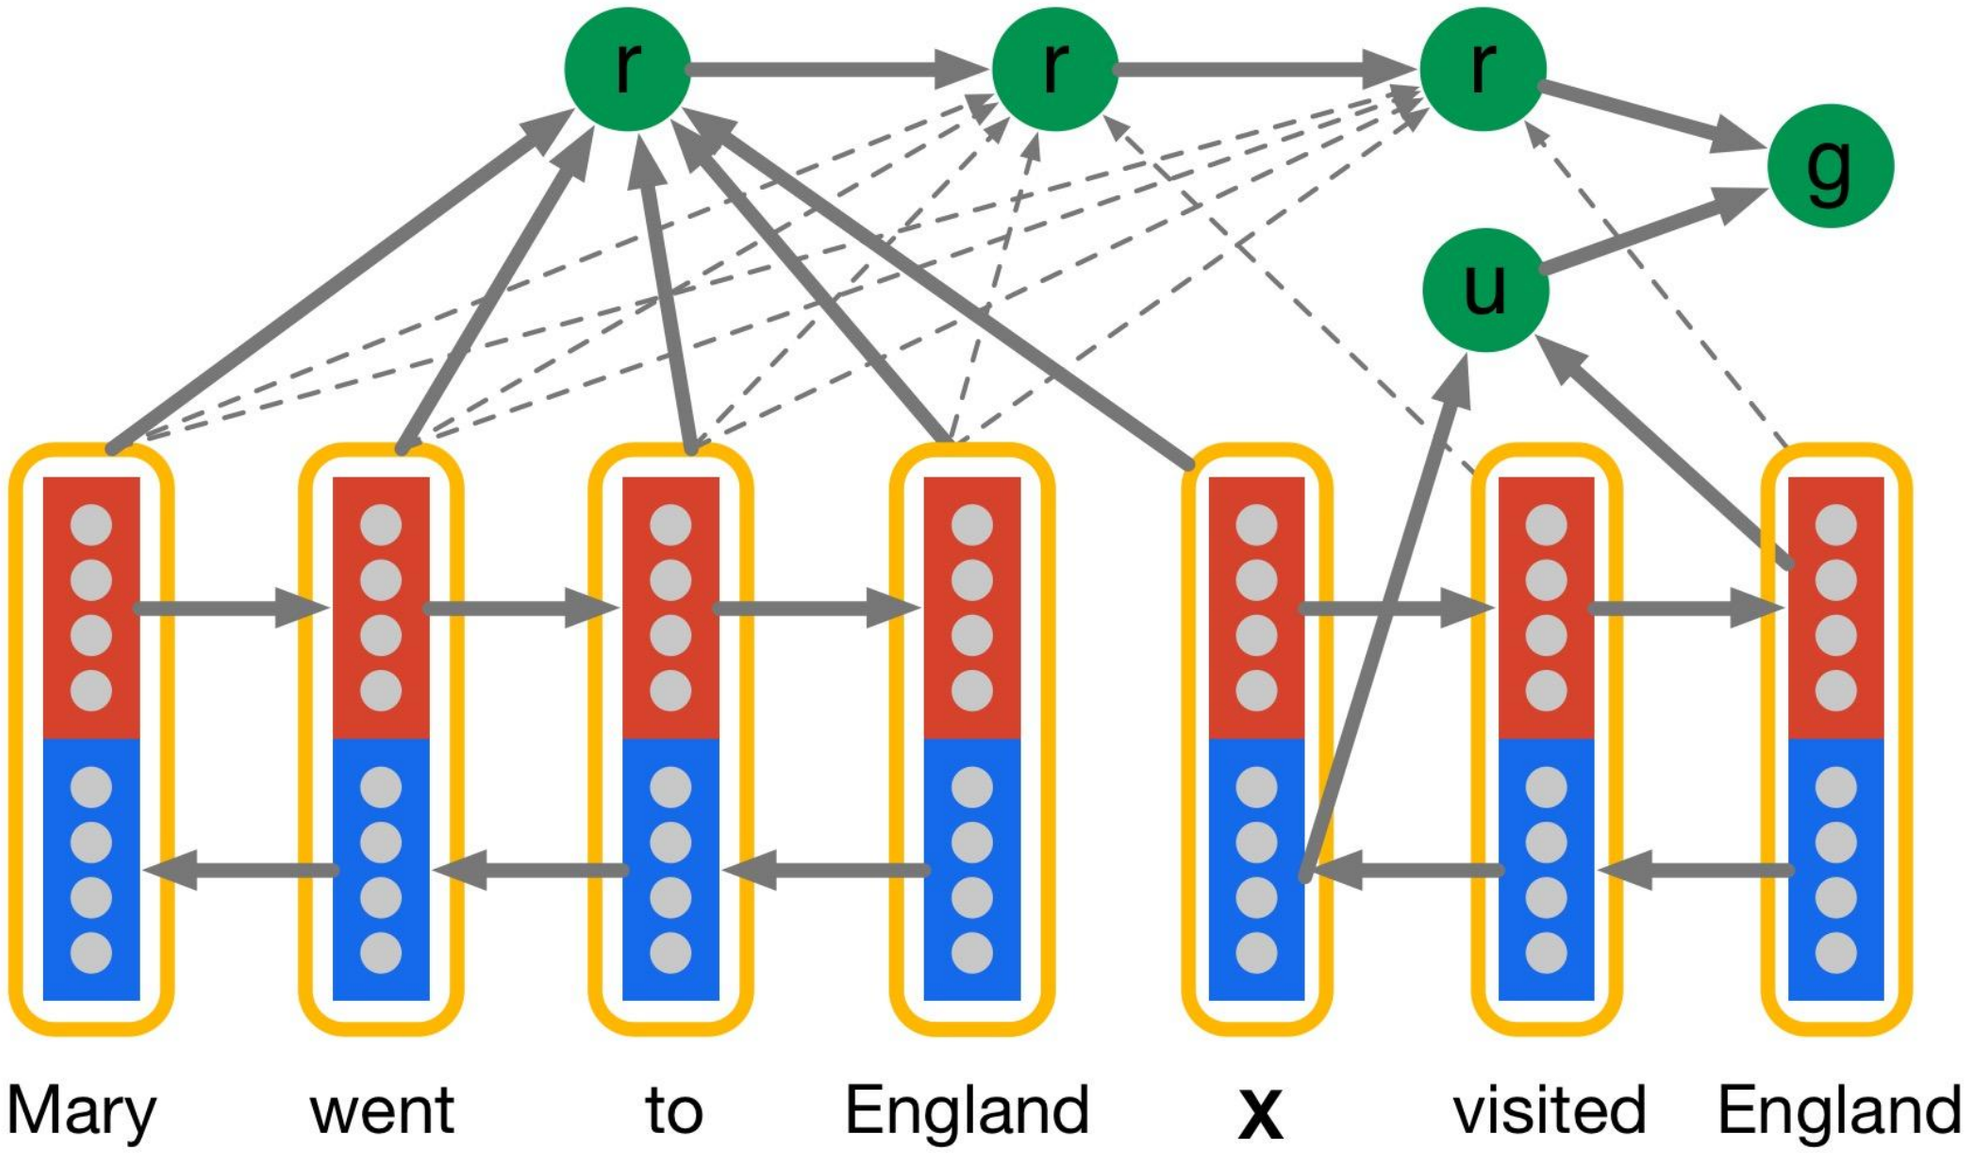

DeepMind

UNIVERSITY OF CAMBRIDGE

Learning to detect entailment

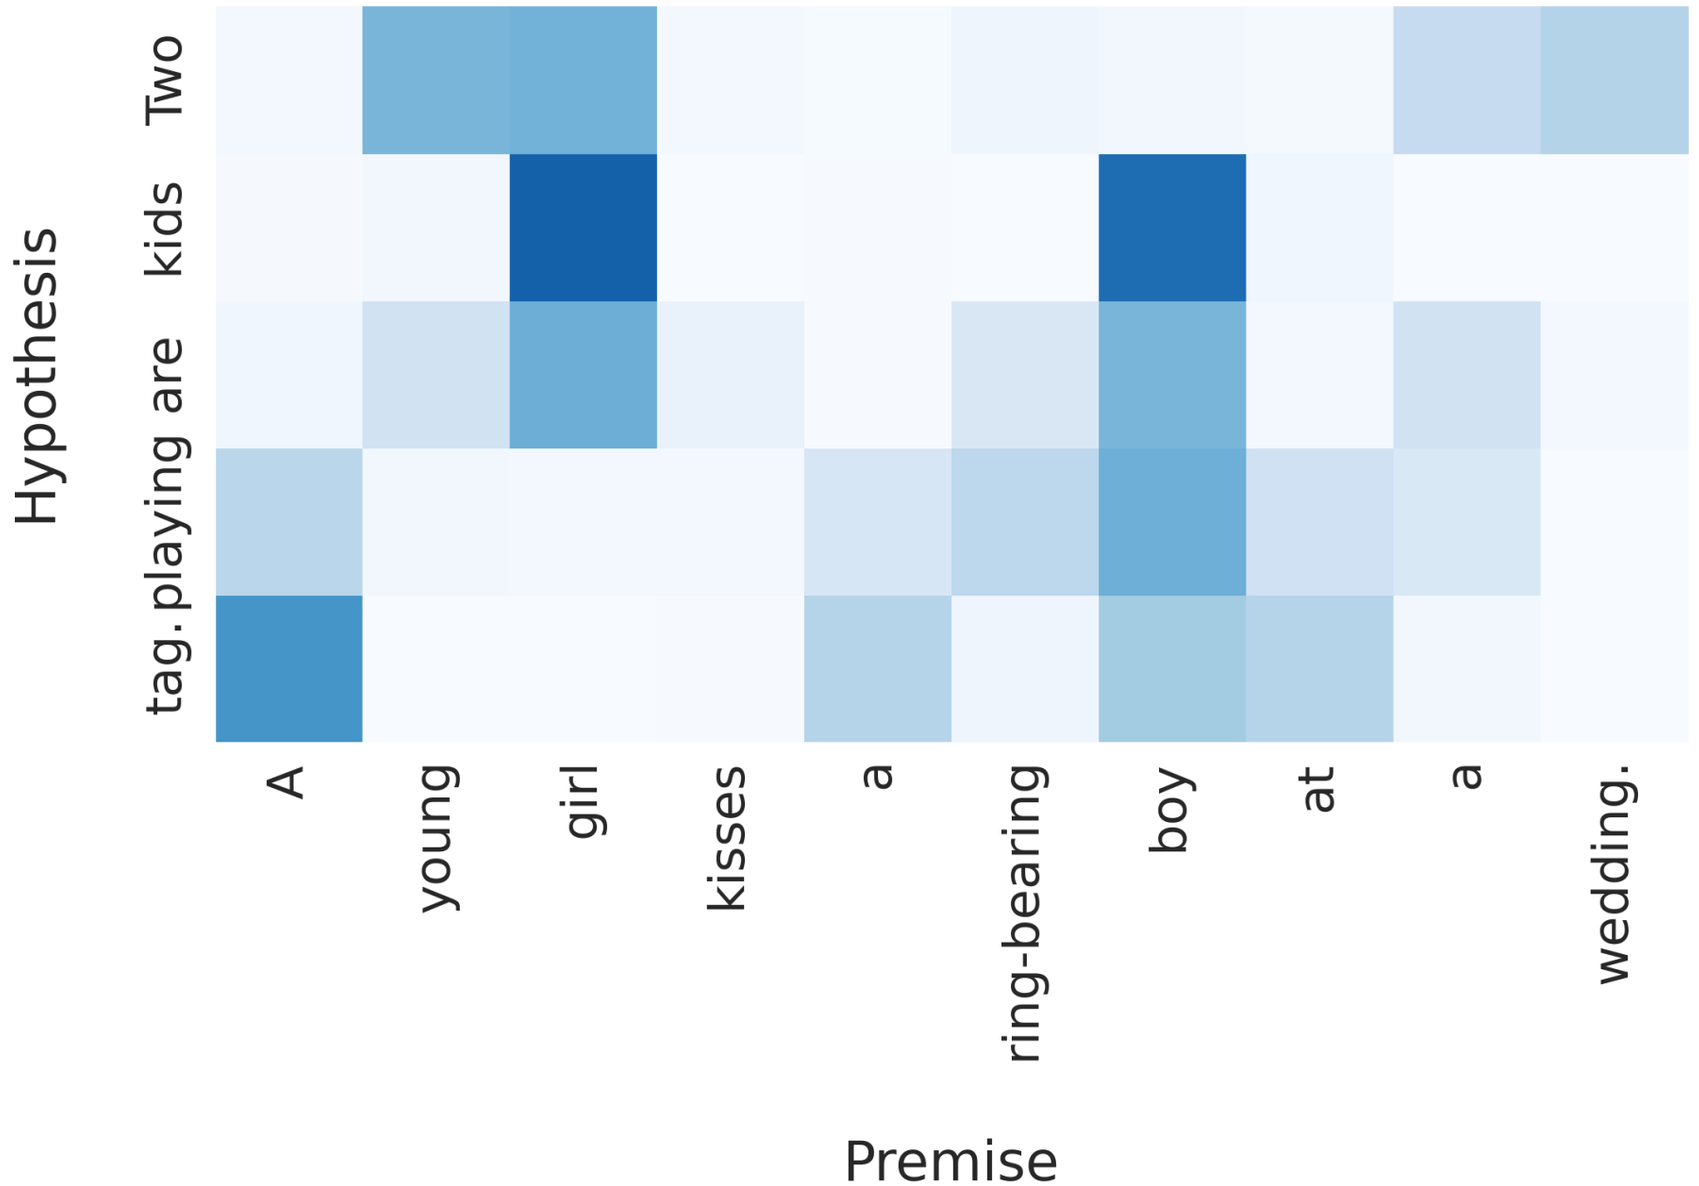

DeepMind

UNIVERSITY OF CAMBRIDGE

Conclusions

- Machine Comprehension can be tested using **the same sort of test we use on humans** (reading comprehension, reasoning exercises, etc).

- Tasks are reasonably well specified, but frequently imperfect, have **simplifying assumptions**, not enough data, **diversity**, etc.

- Many ways of approaching each problem (cf. literature), but **no winner-takes-all**.

- **Unsupervised learning is missing:** Should we be training on the same sort of task we are testing on? Do we learn to understand language by answering reading comprehension examinations?

DeepMind

UNIVERSITY OF CAMBRIDGE

Reading

Matthew Richardson, Christopher JC Burges, and Erin Renshaw. 2013. MCTest: A challenge dataset for the open-domain machine comprehension of text. In *Proceedings of EMNLP*.

Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proceedings of NIPS*.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of EMNLP*.

Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2017. The NarrativeQA Reading Comprehension Challenge. Upcoming in *TACL*.

Sam Bowman, Gabor Angeli, Chris Potts, and Christopher Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of ACL*.

Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, and Phil Blunsom. (2015). Reasoning about entailment with neural attention. In *Proceedings of ICLR*.