Reading Wikipedia to Answer Open-Domain Questions

Danqi Chen, Adam Fisch, Jason Weston, Antoine Bordes
Stanford University & Facebook AI Research

Overview

Goal: build an end-to-end question answering system that can use full Wikipedia to answer any factoid question.

Large-scale QA + Machine comprehension of Text
"Machine Reading at Scale" (MRS)

Our system **DrQA:**

Q What is question answering? **A** a computer science discipline within the fields of information retrieval and natural language processing

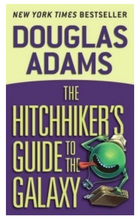

Q Who was the winning pitcher in the 1956 World Series? **A** Don Larsen

Q What is the answer to life, the universe, and everything? **A** 42

Try it out yourself!
https://github.com/facebookresearch/DrQA

Document Retriever + Document Reader

- Document retriever: finding relevant articles from 5 million Wikipedia articles
- Document reader (reading comprehension system): identifying the answer spans from those articles

Q: How many of Warsaw's inhabitants spoke Polish in 1933?

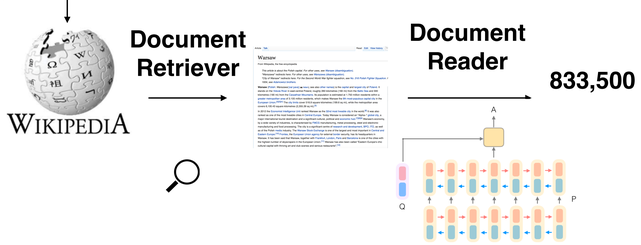

- Datasets:
  - SQuAD (Rajpurkar et al, 2016)
  - TREC (Baudiš and Šedivý, 2005)
  - WebQuestions Freebase (Berant et al, 2013)
  - WikiMovies (Miller et al, 2016)

Approach

Document Retriever

TF-IDF bag-of-words vectors + efficient bigram hashing (Weinberger et al., 2009)

Document Reader

Task: given paragraph P and question Q, the goal is to find a span A in the paragraph which answers the question.
Model: similar to AttentiveReader (Hermann et al, 2015; Chen et al, 2016). We aim to keep it **simple**!

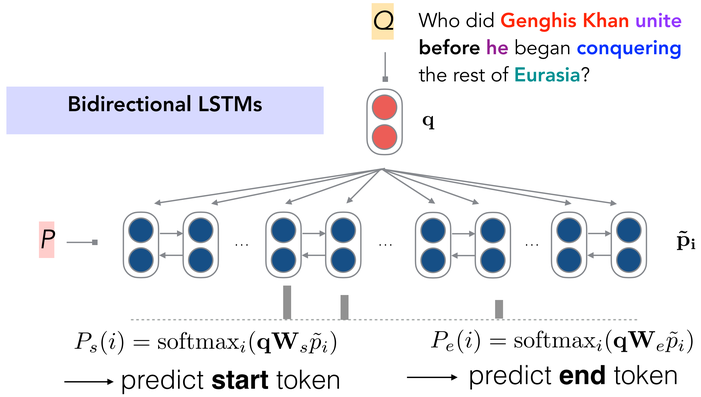

The input vectors consist of:

- Word embeddings
- Exact match features: whether the word appears in question
- Token features: POS, NER, term frequency
- Aligned question embedding

Data: SQuAD + **Distantly Supervised** Data

(Q, A) ⟶ (P, Q, A) if P is retrieved and A can be found in P

Q: What part of the atom did Chadwick discover? WebQuestions
A: neutron

Atom

From Wikipedia, the free encyclopedia

The atomic mass of these isotopes varied by integer amounts, called the whole number rule.[23] The explanation for these different isotopes awaited the discovery of the **neutron**, an uncharged particle with a mass similar to the proton, by the physicist **James Chadwick** in 1932. Isotopes were then explained as elements with the same number of protons, but different numbers of neutrons within the nucleus.

Results

Finding Relevant Articles

	Wiki Search	unigram	+bigram
SQuAD	62.7	76.1	**77.8**
TREC	81.0	85.2	**86.0**
WebQuestions	73.7	**75.5**	74.4
WikiMovies	61.7	54.4	**70.3**

70-86% of questions we have that the answer segment appears in the **top 5 articles**

Performance on SQuAD (single model, Feb 2017)

	EM	F1
Logistic regression	40.4	51.0
Fine-Grained Gating (Carnegie	62.5	73.3
Match-LSTM (Singapore	64.7	73.7
DCN (Salesforce)	66.2	75.9
BiDAF (UW & Allen Institute)	68.0	77.3
Ours	70.7	79.4
r-net (MSR Asia)	71.3	79.7
State-of-the-art (July 2017)	75.7	83.5
Human performance	82.3	91.2

Exact match features are important!

Features	F1
Full	78.8
No f_{token}	78.0 (-0.8)
No f_{exact_match}	77.3 (-1.5)
No $f_{aligned}$	77.3 (-1.5)
No $f_{aligned}$ and f_{exact_match}	59.4 (-19.4)

Full Results