

brmson (YodaQA)

A DeepQA-style Question Answering Pipeline

Petr Baudiš [⟨pasky@ucw.cz⟩](mailto:pasky@ucw.cz)

FEE CTU Prague; brmlab hackerspace

Summer 2014

Outline

- 1 Introduction
- 2 YodaQA Architecture
- 3 Current Performance
- 4 Review, Future Work

Petr Baudiš

First year **PhD student** at FEE CTU Prague (Petr Pošík),
Masters degree in AI from Charles University in Prague

Strong **software engineering** background: The original Git team,
GNU libc development, many open source projects, freelancing

Solid **AI, RL** background: Computer Go research
(MCTS software Pachi — top OSS program, ~4th worldwide)

Basic **ML, optimization** background: Algorithm portfolios, etc.

Newbie in Natural Language Processing!

brmson

A Question Answering system inspired by **IBM Watson** and its DeepQA pipeline architecture.

- Primary goals:**
- Practicality
 - Extensible design
 - Scientific rigor

Current aims: Open-domain factoid questions (TREC QA), generating (not choosing) the right answers.

Multiple implementations:
BlanQA (legacy), YodaQA (current)

brmson: BlanQA (legacy)

- **BlanQA:** *Legacy* pipeline based on CMU's **OAQA**
- Java, UIMA without CAS branching, **UIMA-ECD**
- Architecture based on OAQA **helloqa** prototype GitHub branch, but rewritten almost from scratch
- *enwiki* in **solr**, **Ephyra** answer type system, **Ephyra modules** provide the actual algorithms and rules
- Complete setup documentation, fairly clean code
- Interfaces: Interactive and chatbot (IRC)
- **Functional OAQA end-to-end pipeline!**

Second Thoughts on OAQA

Unhappy with OAQA architecture:

- “Single CAS for each phase” model does not fit UIMA
 - Prevents reuse of most third-party UIMA annotators
 - Prevents scalability
 - Ideally, each *real* subject of analysis should have its own UIMA CAS and subject-of-analysis
- UIMA-ECD source code is rather...opaque
- Unsatisfactory documentation wrt. the big picture
- Fixed answer type system unsatisfactory (though perhaps not essential in OAQA)

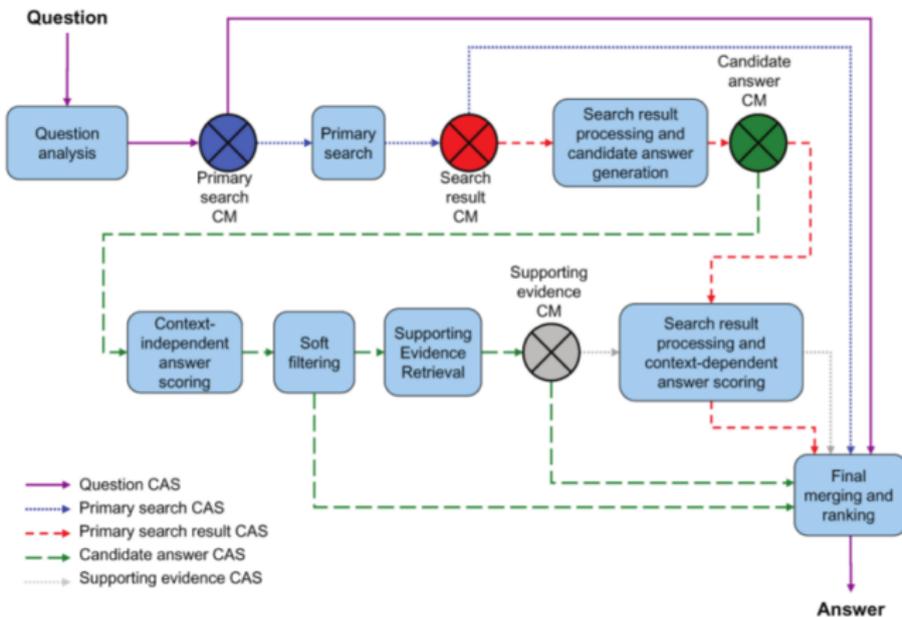
brmson: YodaQA (current)

- **YodaQA:** “Yet anOther Deep Answering pipeline”
- Designed and implemented from scratch — again
- Java, UIMA, DeepQA-style CAS branching, **UIMAFit**
- Architecture based on simplified **DeepQA** (as published)
- Every entity (question, retrieved document, answer) == CAS
- NLP analysis: Third-party UIMA annotators via **DKPro**
- Uses type coercion and parse trees instead of a fixed type system and regexs; no Ephyra components

Outline

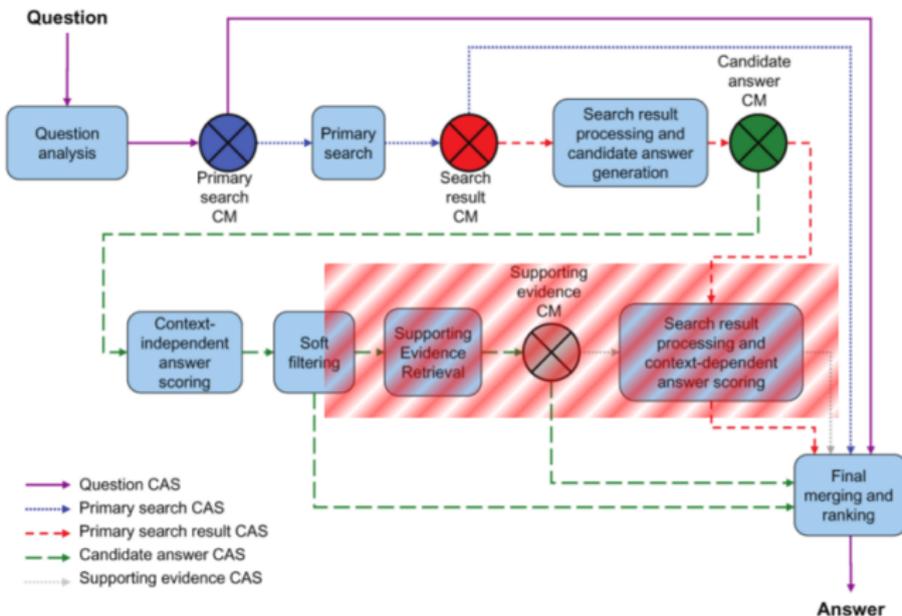
- 1 Introduction
- 2 YodaQA Architecture**
- 3 Current Performance
- 4 Review, Future Work

YodaQA Pipeline



DeepQA architecture (Epstein et al., Making Watson fast).
A series of CAS multipliers.

YodaQA Pipeline



Architecture inspired by DeepQA,
but many modules are obviously much simpler.

Question Analysis

- Full dependency parse
- **Focus** generation (hand-crafted dependency, pos rules)
 - What was the first **book** written by Terry Pratchett?
 - The **actor** starring in Moon?
- **LAT** (Lexical Answer Type) generation (from focus)
 - **Where** is Mount Olympus? **location**
- **Clues** (search keywords, keyphrases) generation:
 - POS and constituent whitelist
 - Selecting verb (hand-crafted rules)
 - Named entities
 - Focus and the NSUBJ constituent
 - enwiki article title exact match

Outcome: Set of Clue and LAT annotations in QuestionCAS

Answer Production

Two answer production pipelines run independently in parallel (custom flow controller developed).

- *SolrFull*: Passage-yielding search
 - *Fulltext*: Full-text + title search for clues, **passages containing clues** are considered
 - *Title-in-clue*: Title search for clues, **initial passage** is considered
 - Passages are parsed, NEs and NPs *not* containing clues are answer candidates
- *SolrDoc*: Full-text search for clues, **document titles** are answer candidates

Outcome: Set of CandidateAnswer CASes

Answer Analysis

Our focus has been maximizing the chance that the right candidate answer enters this phase, answer scoring is pretty naive for now.

- Each answer is POS-tagged and has dependency tree
- Focus and LAT (Lexical Answer Type) generated (dependency roots plus some hand-crafter rules)
- **Type coercion** of question + answer LAT: *Unspecificity* is path length in the **WordNet** (*hypernymy, hyponymy*) graph
- Naive Answer Score:

$$\sum (e^{-tyCorUnspec} \cdot passScore \cdot solrRelevance)$$

where $passScore = \sqrt{\#clueMatches}$ or 2

- N.B. no parameter tuning performed so far!

Outline

- ① Introduction
- ② YodaQA Architecture
- ③ Current Performance**
- ④ Review, Future Work

Testing Dataset

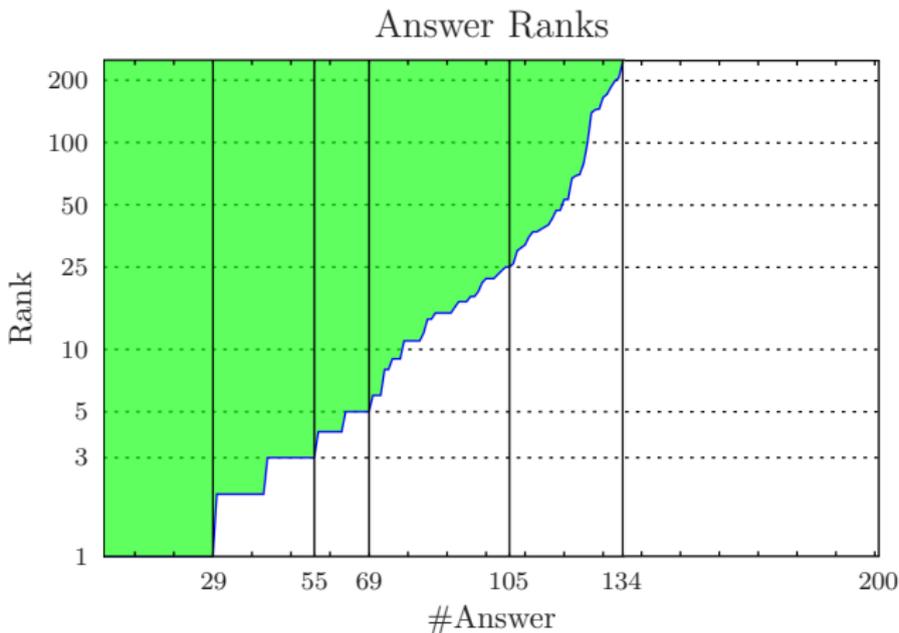
- **TREC QA 2002 + 2003** XML datasets converted to plaintext
- Filtered only questions with a single answer (avoiding some ambiguity)
- **First 200 questions** from this dataset are our first testing set

- Somewhat informative, but less than perfect
- Questions are tied to particular data sources, many ambiguities or imprecisions
- 200 is current practical limit for measurement turn-around (3-4 hour evaluation runs on my home computer)

Experimental Results

Candidate answer binary recall: 67.0% (134/200)

Final answer accuracy: 14.5% (29/200)



Analysis Tools

```
$ data/eval/trecnew-single200-measure.sh
```

```
...
```

```
$ data/eval/tsvout-stats.sh | head -n 5
```

```
3b46430 14-08-14 CluesMergeByText: Al... 29/134 14.5%/67.0% as 0.424
fdb239b 14-08-11 Revert "CluesToConce... 23/131 11.5%/65.5% as 0.395
4cd4a09 14-08-10 Clue: Add a label fe... 21/131 10.5%/65.5% as 0.390
e8ad387 14-08-10 SolrFullPrimarySearc... 24/131 12.0%/65.5% as 0.389
acf17eb 14-08-10 ClueConcept, CluesTo... 18/126 9.0%/63.0% as 0.372
```

```
$ data/eval/tsvout-compare.sh 01718ca 8fc9856
```

```
----- Gained answer to:
```

```
1424 Who wo... for best actor in 1970?      George C. Scott 0.00 1.00
```

```
----- Improved score for:
```

```
1417 Who wa... in less than four minutes?  Roger Bannister 0.57 0.77
```

```
----- Worsened score for:
```

```
1408 Which ... Lionel Jospin a member of?  Socialist          0.31 0.30
```

```
----- Lost answer to:
```

```
1427 What w... spaceship on the moon?      Eagle              0.04 0.00 $
```

Adapt Watson-style Analysis Snapshot

1407 When did the shootings at Columbine happen? | April 20\s?, 1999
#stopwords/#synsets Columbine High School massacre does not
appear; ignore "happen" or add synset that includes "occur"

1439 How deep is Crater Lake? | 1\s?, \s?932 feet
#wikipieces "Crater Lake" yields "crater lake" matches and
never the main article; also, it should be 1,943 feet

1666 What is the name of the US military base in Cuba? | Guantanamo
#abbrev US -> U.S., then should work (answer Guantánamo!)

1606 What is the boiling point of water? | 212 degrees Fahrenheit
100 °C

...

#wikipieces (7) - for all NPs/NEs/nouns in question, include same-
titled wikipedia articles in primary search;
furthermore, do not split such to sub-clues?

#synsets (6) - include synsets instead of words

#abbrev (4) - acronym generation / expansion; e.g. PC = P.C. =
Personal Computer; in expansion, try using #redirects?

Outline

- 1 Introduction
- 2 YodaQA Architecture
- 3 Current Performance
- 4 Review, Future Work**

brmson: YodaQA vs. Primary goals

Practicality

- Detailed setup **instructions** (including data sources setup!)
- Detailed design **documentation**
- Interactive user interface
- Open source (**ASL2** licence), clean code and build system

Extensible design

- UIMA + DeepQA structure: Easy pipeline **branching** and addition of **new modules**
- DKPro: Third-party UIMA annotators (tokenizers, parsers, etc.) are **freely replaceable**
- Internal UIMA components are as **fine-grained** as possible

Scientific rigor

- Gold Standard interface, **TREC QA** based dataset
- All datasets, evaluation tools and measurements **published**
- **AdaptWatson** methodology for performance analysis driving development
- TODO: Cleaned up datasets, use larger test set

brmson: YodaQA vs. Primary goals

Practicality

- Detailed setup **instructions** (including data sources setup!)
- Detailed design **documentation**
- Interactive user interface
- Open source (**ASL2** licence), clean code and build system

Extensible design

- UIMA + DeepQA structure: Easy pipeline **branching** and addition of **new modules**
- DKPro: Third-party UIMA annotators (tokenizers, parsers, etc.) are **freely replaceable**
- Internal UIMA components are as **fine-grained** as possible

Scientific rigor

- Gold Standard interface, **TREC QA** based dataset
- All datasets, evaluation tools and measurements **published**
- **AdaptWatson** methodology for performance analysis driving development
- TODO: Cleaned up datasets, use larger test set

brmson: YodaQA vs. Primary goals

Practicality

- Detailed setup **instructions** (including data sources setup!)
- Detailed design **documentation**
- Interactive user interface
- Open source (**ASL2** licence), clean code and build system

Extensible design

- UIMA + DeepQA structure: Easy pipeline **branching** and addition of **new modules**
- DKPro: Third-party UIMA annotators (tokenizers, parsers, etc.) are **freely replaceable**
- Internal UIMA components are as **fine-grained** as possible

Scientific rigor

- Gold Standard interface, **TREC QA** based dataset
- All datasets, evaluation tools and measurements **published**
- **AdaptWatson** methodology for performance analysis driving development
- TODO: Cleaned up datasets, use larger test set

YodaQA: Future Work

My personal plan:

- More aggressive usage of Wordnet synsets, and other **tweaks** based on performance analysis
- **Parameter optimization** to boost answer scoring precision
- Basic **evidence gathering** for top candidate answers
- Declare and **publish baseline experimental testbed**

With more contributors:

- Cleaned up testing dataset
- UIMA component **unit tests**
- Verification dataset runs with **human judges**
- Scale-out, parallelization and memory usage **optimizations**
- **Apply** to some real-world projects and domains

Long-term Plans and Goals

- Post-YodaQA architecture reformulation as IE problem:

Latent knowledge graph paradigm

(QA pipeline as on-demand population of semantic network;
answer retrieved by path search, scored by edge coercion)

- brmson-based **startup**: Looking for good business cases
- Disembodied autonomous agent: QA with deduction + goal-setting + planning (maybe in 15 years)
- Personal: Internship at **NII Tokyo** in 1st quarter 2015 (answering of Physics questions in university entry exams)

Conclusion

- Practical, open source QA system
 - Clean architecture and development methodology
 - Reasonably documented!
 - Clear path forward, towards reference experimental testbed
 - Immediate tasks: Improve basic answer scoring, add evidence gathering
- Interested in any sort of **collaboration** with people at CMU (or elsewhere)

pasky@ucw.cz
petr.baudis@gmail.com

Thank you for your attention!