

Deliverable 3

Claire Jaja, Andrea Kahn and Clara Gordon

Most important: a new
name...

QuAILS

Question Answering Integrated Linguistic System



Improving Runtimes

- Indexing - no longer using pymur, now using Indri directly
 - before: 6 hours, now: 15 minutes
- Parallelization for question pipeline
 - using Python multiprocessing module (no Condor blocking, you're welcome)
 - note: indexing with stopword list also has big improvement for runtime of question pipeline
 - before: ~3 hours, now: ~3 minutes



Improving Results

System	Strict	Lenient
D2	0.0051	0.0289
D3	0.1451	0.2639

2745% increase in strict
813% increase in lenient



Web Boosting

- Our most successful addition
- We scrape Ask.com results for the query + target
- Results are cached in a text file for later use
- No problems with throttling
- Treat web snippets as passages alongside Indri-returned passages
- Can be identified as web snippets because have document ID of “None”



Web Boosting

- After extensive testing, found best results with:
 - 3 pages of cached web results
 - Weighting web snippets very highly
 - Require all answers be in at least 1 AQUAINT doc
 - Require all answers be in 10 passages - but count each web snippet as 10 passages

System	Strict	Lenient
without web boosting	0.0051	0.0289
with web boosting	0.0742	0.1257



Question Classification

- Perform regular expression matching on original question text to determine whether we're looking for looking for answer that is a Person, Organization, Location, Time Expression, Number, etc. (using wh- words, etc.)
- Pass answer types and weights to answer-processing module
- Use regular expressions on possible answers
- Multiply answer score by weight for type identified - or by weight for other if no type identified

System	Strict	Lenient
without question classification	0.0742	0.1257
with question classification	0.0614	0.1330




Query Expansion

- Lin thesaurus-based synonym expansion
- NLTK returns synonyms in three POS categories: (N, V, Adj) along with similarity scores
- We select the top n synonyms for query terms after removing named entities
- After experimenting with weighting and n parameters, our best results...

System	Strict	Lenient
without Lin synonyms	0.0683	0.1279
with Lin synonyms	0.0331	0.0943

- :(
- Idea for next time: POS-tag query terms and return c with correct POS tag



Guess the Question!

142.4 D3 NYT19981201.0052 2000

142.4 D3 NYT19990717.0171 Ladies Professional Golf

142.4 D3 APW19991012.0204 longest-running women 's sports

142.4 D3 NYT19990717.0171 Professional Golf

142.4 D3 XIE20000111.0231 Commissioner 's Award

142.4 D3 NYT19990717.0171 Ladies Professional

142.4 D3 APW19991012.0204 longest-running women

142.4 D3 NYT20000719.0034 Hall of Famer

142.4 D3 NYT20000113.0020 13

142.4 D3 APW19981106.0075 U.S. Women 's Open



Guess the Question!

142 target: “LPGA”

142.4 question: “When does the LPGA celebrate its 50th anniversary?”

142.4 D3 NYT19981201.0052 2000

142.4 D3 NYT19990717.0171 Ladies Professional Golf

142.4 D3 APW19991012.0204 longest-running women 's sports

142.4 D3 NYT19990717.0171 Professional Golf

142.4 D3 XIE20000111.0231 Commissioner 's Award

142.4 D3 NYT19990717.0171 Ladies Professional

142.4 D3 APW19991012.0204 longest-running women

142.4 D3 NYT20000719.0034 Hall of Famer

142.4 D3 NYT20000113.0020 13

142.4 D3 APW19981106.0075 U.S. Women 's Open



Guess the Question!

192.2 D3 APW19980627.0818 nearly 800
192.2 D3 XIE20000920.0036 Spain and southern France
192.2 D3 XIE20000920.0036 Spain and southern
192.2 D3 APW20000625.0138 30-year campaign
192.2 D3 APW19980918.0676 Spain and France
192.2 D3 APW19980627.0818 Homeland and Freedom
192.2 D3 APW19981028.0645 Minister Jose Maria Aznar
192.2 D3 APW20000625.0138 five
192.2 D3 APW19981028.0645 Prime Minister Jose
192.2 D3 APW19981028.0645 Prime Minister



Guess the Question!

192 target: “Basque ETA”

192.2 question: “Approximately how many people has ETA killed?”

192.2 D3 APW19980627.0818 nearly 800

192.2 D3 XIE20000920.0036 Spain and southern France

192.2 D3 XIE20000920.0036 Spain and southern

192.2 D3 APW20000625.0138 30-year campaign

192.2 D3 APW19980918.0676 Spain and France

192.2 D3 APW19980627.0818 Homeland and Freedom

192.2 D3 APW19981028.0645 Minister Jose Maria Azn:

192.2 D3 APW20000625.0138 five

192.2 D3 APW19981028.0645 Prime Minister Jose

192.2 D3 APW19981028.0645 Prime Minister



OK, we give in...

Fine, fine, we'll return passages instead of answer n-grams


- Maintain all our answer processing but also hold onto the AQUAINT passages the answers come from
- At the end, use the document ID that has the most passages containing that answer
- Pick one of the passages from that document
- Return 250 characters, centered around the answer

System	Strict	Lenient
return n-gram answers	0.0868	0.1499
return passages	0.1451	0.2639



Issues

- Keeping track of multiple experiments and parameters
 - # Indri passages
 - Passage window
 - # synonyms in query expansion
 - query weighting: synonyms, named entities
 - # pages of web results
 - minimum required passage number
 - stemmer
 - stoplist
 - → Google spreadsheets!
- Determining appropriate ranges for weights from different sources -- Lin similarity, Indri document weighting, query and synonym weights



Next Steps

- Experiment with different applications for web results
 - redundancy-based query expansion
- Experiment with Indri Query Language -- ordering, chunking
- Better question and answer classification
- Experiment with POS-tagging queries before expanding them (and other forms of constraining expansion) to see if we can get improved results with query expansion
- Try different methods for selecting which passage that highly ranked answers occur in to return



Quentin the Quail



Thank you!