

The Yahoo Query Treebank, V. 1.0

May 2016

Yuval Pinter
Yahoo Research
me@yuvalpinter.com

Roi Reichart
Yahoo Research & Technion IIT
roiri@ie.technion.ac.il

Idan Szpektor
Yahoo Research
idan@yahoo-inc.com

1 General

This dataset release accompanies Pinter et al. (2016) which describes the motivation and grammatical theory. Please cite that paper when referencing the dataset.

The dataset may be accessed via the Yahoo Webscope homepage¹ under **Linguistic Data** as dataset **L-28**. The description in Section 2 is included within the dataset as a Readme.

The dataset is sure to have annotation errors which are not covered by the special cases specified in this document. Please approach the first author for any corrections and they will appear in the next release. See Section 4 for known errors.

2 Dataset Description

User queries annotated for syntactic dependency parsing, version 1.0.

2.1 Full description

This dataset contains two files:

ydata-search-parsed-queries-dev-v1_0.txt 1,000 queries, 5,344 tokens
ydata-search-parsed-queries-test-v1_0.txt 4,000 queries, 26,015 tokens

These files differ in their level of annotation, but share the schema. They contain tab-delimited lines, each representing a single token in a Web query. The tokens in each query are given sequentially, and queries are given in order of an arbitrarily-selected numeric ID (with no empty lines between queries). The field schema is detailed in Table 1.

An example query is brought in Table 2. These lines represent a query, whose ID in the ‘test’ set is 183, and whose raw form is *charter school graduate early*. It is interpretable thus: the query is composed of two syntactic segments: [*charter school*]

¹<http://webscope.sandbox.yahoo.com>




Figure 1: Query #183 from the test set, tagged and parsed.

[*graduate early*]. In the first segment, the syntactic root is the noun *school*, and the noun *charter* modifies it in a nominal compound modifier relation. The second segment is rooted by the verb *graduate*, modified by the adverb *early* in an adverbial modifier relation. A dependency tree corresponding to this query is produced in Figure 1.

2.2 Linguistic pre-processing notes

All queries were tokenized using the ClearNLP tokenizer for English (Choi and McCallum, 2013)² and were not spell-corrected or filtered for adult content. For Excel-friendliness, initial quotes were replaced with backticks. When using off-the-shelf processing tools, we recommend to replace them in pre-processing.

3 Annotation Guidelines

The parse tree annotations followed the segmentation annotation. Both phases included supervision over automatic part-of-speech tagging. In general, parse tree annotators were instructed to adhere to the Stanford Dependency guidelines (De Marneffe and Manning, 2008)³ with the necessary caveats that come with non-standard text prone to errors.

Following are some selected issues which can be seen in the test dataset. Section 3.1 describes issues relating to segmentation tagging, and the subsequent subsections address issues of dependency

²Version 2.0.1, from www.clearnlp.com

³Version 3.5.2, from http://nlp.stanford.edu/software/dependencies_manual.pdf

Field #	Field content
0	query-id
1	token-in-query (starting with 1)
2	segmentation marker ('SEG' iff token starts a new segment, '-' otherwise)
3	token form
4	token part-of-speech †
5	index of syntactic head token, with 0 denoting root †
6	dependency relation of edge from head token to this token † ‡

Table 1: Fields marked by † are not populated for the dev set in V 1.0; fields with ‡ are not fully populated in the test set.

183	1	SEG	charter	NN	2	nn
183	2	-	school	NN	0	root
183	3	SEG	graduate	VB	0	root
183	4	-	early	RB	3	advmod

Table 2: A sample query (#183 in test set).

edge attachment.

3.1 Segmentation Ambiguity

Noun Strings The most difficult segmentation decisions have been in cases of long strings of nouns, both common and proper. The guideline is based on judgment, asking whether **the phrase can stand as a constituent in a coherent sentence**. For example, in query 116 words 1-3 are *big mac cost*, a phrase considered clunky as opposed to the much-preferred *cost of (a) big mac*, resulting in the decision to segment it as [*big mac*] [*cost*]. A more clear-cut case is the segmentation decision in q.284: [*first day of missed period*] [*symptoms*], where theoretically *day* could be tagged as a modifier phrase head of *symptoms*, but no imaginable well-formed sentence would use this formation as a constituent.

Sometimes reasonable semantics forced us to conclude towards a segmentation decision over an unlikely (but syntactically well-formed) single-segment constituent. For example, in q.271 [*fanfiction*] [*guest reviews*] we assume (and support this decision by running our own search) that the searcher was looking for guest reviews on fanfiction website platforms rather than guest reviews written about pieces of fanfiction. Likewise, in q.3013 [*questionnaire*] [*assisted suicide*] we treated *assisted* as an adjective relating to *suicide* rather than a main verb in a sentence describing an unlikely scenario.

Missing Auxiliary Verbs Another issue was certain constructions which may become a sentence by the addition of a copula or auxiliary verb. The guidelines stated that if the post-copular sequence describes an attribute of the semantic subject (pre-copula), it is considered the head of an adjunct phrase and shares its segment (e.g. q.2916 [*paragraph describing a family*], q.91.w.1-4: [*battery cables not tight*]). If it is construed as a sentence missing an auxiliary, it is segmented away from the subject (e.g. q.475 [*lab*] [*throwing up blood ?*]). In a way, this is an explicit application of the constituency test described earlier.

3.2 Proper Names' Internal Structure

The dataset contains many instances of product names followed by numbers denoting model (e.g. query 2347: *how to replace crank sensor on 95 saab cs9000*). The guidelines in these cases were to place the phrase head as the last alphabetic word (here, *saab*) and depend the post-modifier (*cs9000*) on it as an *npadvmod*.

Long proper names with complex internal structures were left as their underlying structures. E.g. query 3252: *the₁ call of the₂ wild* is parsed with *call* as the phrase head with a *det* preceding and a *prep* following, further decomposed into *pobj(of,wild)* and *det(wild,the₂)*. However, more trivial internal structures were flattened into the common proper name representation: q.3260, *jeeves & wooster*: *nn(wooster,jeeves), nn(wooster,&)*. Nominal compounds were tagged as successive *NNPs*, e.g. q.3413, *what cameras do they use in planet_NNP earth_NNP*: *nn(earth,planet)*.

3.3 Truncated Sentences

Many of the queries in the dataset are in fact sentences aborted mid-way. Some because the query is a sentence-completion question (e.g. query 970

a major result of the european age of exploration was) and some for more opaque reasons (e.g. q.3828 *why are russians so*). In both these types, where the root of the sentence would normally lie in the missing complement, the root was assigned to the token nearest to it in the assumed full graph (*was* and *so*, respectively), with the other dependents collapsed into it. The same goes for any phrase which is truncated before its grammatical head (e.g. q.3840) or mandatory complement (e.g. q.1861).

3.4 Foreign Languages

The dataset contains several non-English queries that are treated as nonce (part-of-speech tag = ‘FW’, dependency relation = ‘dep’). Their parse tree is, as a rule, a flat tree headed by the final word (proper name convention). E.g. q.3299, q.3319. Other cases where foreign words are tagged as ‘FW’ is when they function meta-linguistically. E.g. q.3472, *what does baka_FW mean in japanese*.

3.5 Grammatical Errors

By the nature of Web Queries being written in real time by users possessing diverse proficiency and competence, the dataset contains many grammatical errors (as well as typos – see Section 3.6). These were not corrected during pre-processing in order to maintain the authenticity of the data. The guidelines in such cases, unless re-segmentation was in order, called to retain as much of the intended structure as possible. For example, in query 3274, word 4 *part* is meant to be *parts* and as such is tagged as a plural noun. In some cases, different parts of the sentence were fused together due to incorrect grammar. The solution was to represent the fused token by the head of the intended phrase if it is fully contained within (e.g. q.3233, *dogs* for *dog ’s*), or ignore a dependent if there is structural crossover (e.g. in q.3590, *my sister in laws husband*, the possessive *’s* was in effect excluded from the tree).

Another common case was auxiliary deletion common in ESL writing, e.g. query 3949 *why you study in university*. Here we opted again for the intended meaning as a sentence missing the auxiliary *do* rather than treating the full query as a constituent (adverbial clause, clausal complement, etc.) of a deleted governing predicate.

3.6 Typographical Errors

In general, obvious typos were treated as the intended word whether they resulted in a legal English word or not. E.g. q.3 w.13, *trianlgle_NN*, or q.3786 w.5 *if_VBZ* (for the intended *is*).

Sometimes, typos result in word merge, in which case they were either POS-tagged as *XX* (e.g. q.252 w.3) or by the head if they can be construed as a coherent phrase (e.g. q.3115 w.1 *search-where*, which is probably a mis-concatenation of a meta-linguistic *search* with the first intended term *where*, and so analyzed as if it were the latter alone).

Sometimes, extremely creative tokenization is employed by users. Behold the glory that is query 1814: *green chemistry.in day today life*. We parsed *day today* as if it were a noun phrase, and gave up representing the preposition *in* altogether.

3.7 BE-sentences

Sentences with forms of the verb BE are often ambiguous between attributive sentences (where the BE-verb acts as copula to the head of the following phrase) and proper essential statements where BE is the main verb. We tended to go for the former in case of ambiguity, excepting for very clear cases of essence (e.g. query 3264 *the free market is a myth: root(ROOT,is)*) and, of course, where the following can only be a clausal or prepositional complement (e.g. q.799 *trouble is you think you have time*, q.3593 *what is on sheldons shirt in season 1 episode 4*).

4 Known Annotation Errors

4.1 Segmentation Errors

Query ID	V 1.0	Correct
194	1,3	1
304	1,3	1
325	1,3	1
362	1,3	1
425	1,4	1
847	1,3	1
911	1,4	1
3779	1,5	1
1147	1	1,3
1812	1	1,4
2784	1	1,2,3
2883	1	1,7
2912	1	1,5,7
3348	1	1,2
3366	1	1,2,3,4

4.2 Attachment Errors

Query.Token	V 1.0	Correct
3153.6	2	7
3153.7	6	2

Credits

Segmentation tagged by **Bettina Bolla**, **Avihai Mejer**, **Yuval Pinter**, **Roi Reichart**, and **Idan Szpektor**. Parsing tagged by **Shir Givoni** and Yuval Pinter.

References

- [Choi and McCallum2013] Jinho D Choi and Andrew McCallum. 2013. Transition-based dependency parsing with selectional branching. In *ACL (1)*, pages 1052–1062.
- [De Marneffe and Manning2008] Marie-Catherine De Marneffe and Christopher D Manning. 2008. Stanford typed dependencies manual. Technical report, Technical report, Stanford University.
- [Pinter et al.2016] Yuval Pinter, Roi Reichart, and Idan Szpektor. 2016. Syntactic parsing of web queries with question intent. In *Proceedings of NAACL-HLT, San Diego, CA, June 12-17, 2016*.