Aligning English Strings with Abstract Meaning Representation Graphs

Nima Pourdamghani   Yang Gao   Ulf Hermjakob   Kevin Knight

Information Sciences Institute

University of Southern California

Overview

Abstract Meaning Representation (AMR) [1]:

- Logical meaning of sentences
- Directed acyclic graphs with labeled edges

The boy wants to go
(w / want-01
:arg0 (b / boy)
:arg1 (g / go-01
:arg0 b))

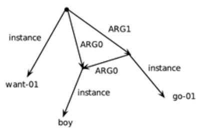

- Find alignment links between English tokens and AMR concepts
- Alignments are required for:
  - Semantic parsing
  - English generation

Approach

- Similar to Statistical Machine Translation
- Linearize AMR graph (not obvious how)
- Use string / string alignment
- Easier than SMT
  - AMR and English are highly cognate
- Harder
  - AMR is a graph with unordered nodes
  - Much less training data than in SMT

Corpus

- 13050 public AMR/English sentence pairs
- Hand Aligned 200
  - 100 dev, 100 test
- Ratio of aligned tokens in the gold data
  - English: ~3/4
  - AMR: ~ 1/2
  - AMR role tokens: ~ 1/4

	train	dev	test
Sent. pairs	13050	100	100
ENG tokens	248 K	2.3 K	1.7 K
AMR tokens	465 K	3.8 K	2.3 K
AMR role tokens	226 K	1.9 K	1.1 K

There   is   no   asbestos   in   our   products   now   .

(a / **asbestos**   **:polarity**   **"-"**   :time   (n / **now**)   **:location** (t / **thing**   **:ARG1-of**   (p2 / **produce-01**   **:ARG0** (w2 / **we**))))

The process

Preprocess

The boy wants to go
(w / want-01
:arg0 (b / boy)
:arg1 (g / go-01
:arg0 b))

- Linearize AMR:
  w / want-01 :arg0 b / boy :arg1 g / go-01 :arg0 b
- Remove stopwords
  - English: *boy wants go*
  - *AMR: want-01 boy go-01*
- Remove word sense indicator, etc. in AMR
  - *want boy go*
- Stem both English and AMR to first four letters
  - English: *boy want go*
  - AMR: *want boy go*

Extend Parallel Corpus

- Tokens that look the same after stemming

boy	boy
want	want
go	go

- English tokens that map to multiple AMR ones

higher	high :degree more
biggest	big :degree most
november	:month 11

Training

- Based on IBM word alignment models [2]
- Use EM to maximize likelihood:
- Generating AMR from English
  $$\theta_{A|E} = argmax\, L_{\theta_{A|E}}(A|E)$$
- *Or , generating English from AMR*
  $$\theta_{E|A} = argmax\, L_{\theta_{E|A}}(E|A)$$
- Decoding: get the most probable alignments given parameters using Viterbi algorithm

Symmetrized EM

- Word alignment is symmetric
- Training should be symmetric as well
- New objective:

$$\theta_{A,E} = argmax\left(L_{\theta_{A|E}}(A|E) + L_{\theta_{E|A}}(E|A)\right)$$

Subject to: $\theta_{A|E}\theta_E = \theta_{E|A}\theta_A = \theta_{A,E}$

- Approximate solution:
  1- optimize $\theta_{A|E} = argmax\, L_{\theta_{A|E}}(A|E)$
  2- satisfy constraint, initialize $\theta_{E|A} \propto \theta_{A|E}$
  3- optimize $\theta_{E|A} = argmax\, L_{\theta_{E|A}}(E|A)$
  4- satisfy constraint, initialize $\theta_{A|E} \propto \theta_{E|A}$
  5- Iterate
- Steps 1 and 3: EM (IBM models)
- Steps 2,4: simple initialization
- No extra code needed

Postprocess

- Goal: rebuild the **aligned** AMR graph
- Restore stopwords, change alignments
- Rebuild graph using recorded original structure

Experiments

Precision, Recall, F-measure

- We used Mgiza++ as implementation of IBM models
- Experiment setup (Model 4+):
  5 × Model 1 + 5 × HMM + 4 × symmetrized Model 4

	Model	precision	recall	F score
Dev	Model 1	70.9	71.1	71.0
	HMM	87.6	80.1	83.7
	Model 4	89.7	80.4	84.8
	Model 4+	94.1	80.0	86.5
Test	Model 1	74.8	71.8	73.2
	HMM	83.8	73.8	78.5
	Model 4	85.8	74.9	80.0
	Model 4+	92.4	75.6	83.1

- Test set was intrinsically harder
- Symmetrization increased F-measure by 1.7 and 3.1 points for dev and test sets on Model 4

Error Sources

performance breakdown for AMR role and non-role tokens

	token type	precision	recall	F score
Dev	role	77.1	48.7	59.7
	non-role	97.2	88.2	92.5
	all	94.1	80.0 (34%)	86.5
Test	role	71	37.8	49.3
	non-role	95.5	84.7	89.8
	all	92.4	75.6 (36%)	83.1

- Most of the error is on role tokens
  - role tokens don't have a specific translation in English
  - some hardly get aligned to any English word
    - :unit, :value, …
  - some can align to many different English words
    - :manner to most of the adverbs
  - they can match to part of an English word
    - :polarity to **un**popular
  - or the connection might be very implicit
    - (t / thank-01 **:arg0** (i / i) **:arg1** (y / you)) to thanks
- About 35% of recall loss is due to removing aligned stop words

Conclusions Future Work

We have presented:

- The first set of manually aligned English/AMR pairs (available in amr.isi.edu)
- The first system, and a strong baseline, for learning alignments between English sentences and AMR graphs
- The system is adaptable to any domain and any language
- First step for parsing AMR from English and generating English from AMR

References

[1] L. Banarescu, C. Bonial, S. Cai, M. Georgescu, K. Griffitt, U. Hermjakob, K. Knight, P. Koehn, M. Palmer, and N. Schneider. 2013. Abstract meaning representation for sembanking. Linguistic Annotation Workshop (LAW VII-ID), ACL.

[2] P. F. Brown, V. J. Della Pietra, S. A. Della Pietra, and R. L. Mercer. 1993. The mathematics of statistical machine translation: Parameter estimation. Computational linguistics, 19(2):263–311.

Acknowledgements

This work was supported by DARPA contracts HR0011-12-C-0014 and FA-8750-13-2-0045. The authors would like to thank David Chiang, Tomer Levinboim, and Ashish Vaswani for their comments and suggestions.

John Lennon Wall, Prague.
Thanks to: J. Flanigan, C. Wang, and Y. Zhang!