Response-based Learning for Grounded Machine Translation

Stefan Riezler, Patrick Simianer, Carolin Haas

Department of Computational Linguistics
Heidelberg University, Germany

Response-based Learning

- **Extract supervision signal from extrinsic response** to predicted structure.
- Prediction is **tried out** in extrinsic task:
  - **approved** as positive training example in case of positive task-based feedback,
  - **in addition to** or **instead of** learning from given gold standard annotations.

Response-based Learning

- **Extract supervision signal from extrinsic response** to predicted structure.
- Prediction is **tried out** in extrinsic task:
  - **approved** as positive training example in case of positive task-based feedback,
  - **in addition to** or **instead of** learning from given gold standard annotations.

Response-based Learning for MT

- **Try out** most probable translation in extrinsic task, and **approve as reference translation** in case of positive feedback.
- Advantages over learning from references only:
  - **Reproducability**: Multiple system translations can be converted into references.
  - **Reachability**: References are necessarily in decoder search space (compared to independently created human reference translations).

Response-based Learning for MT

- **Try out** most probable translation in extrinsic task, and **approve as reference translation** in case of positive feedback.
- Advantages over learning from references only:
  - **Reproducability**: Multiple system translations can be converted into references.
  - **Reachability**: References are necessarily in decoder search space (compared to independently created human reference translations).

Grounded Language Learning / Semantic Parsing

- **Grounded language learning**: Successful communication of meaning defined as successful interaction in a task ([Roy, 2002, Yu and Ballard, 2004, Yu and Siskind, 2013], *inter alia*).
- **Semantic parsing**: Successful execution of a meaning representation in a simulated world defined as returning the correct answer from a knowledge base (GEOQUERY, [Wong and Mooney, 2006]; ATIS [Zettlemoyer and Collins, 2009], FREEBASE [Cai and Yates, 2013], *inter alia*).

Grounded Language Learning / Semantic Parsing

- **Grounded language learning**: Successful communication of meaning defined as successful interaction in a task ([Roy, 2002, Yu and Ballard, 2004, Yu and Siskind, 2013], *inter alia*).
- **Semantic parsing**: Successful execution of a meaning representation in a simulated world defined as returning the correct answer from a knowledge base (GEOQUERY, [Wong and Mooney, 2006]; ATIS [Zettlemoyer and Collins, 2009], FREEBASE [Cai and Yates, 2013], *inter alia*).

Response-based Semantic Parsing

- Learn semantic parsers from question-answer pairs without recurring to annotated logical forms [Kwiatowski et al., 2013, Berant et al., 2013, Goldwasser and Roth, 2014].
- Term **response driven learning** coined by [Clarke et al., 2010].

Grounding SMT in Semantic Parsing

- QA-scenario:
  - **Question is translated successfully if correct answer is returned** based only on the translation of the question.
- Semantic parsing realization:
  - Translation quality defined by ability of semantic parser to **construct a meaning representation from the translated query**, which **returns correct answer** when executed against database.

Grounding SMT in Semantic Parsing

- QA-scenario:
  - **Question is translated successfully if correct answer is returned** based only on the translation of the question.
- Semantic parsing realization:
  - Translation quality defined by ability of semantic parser to **construct a meaning representation from the translated query**, which **returns correct answer** when executed against database.

Response-based Learning Cycle

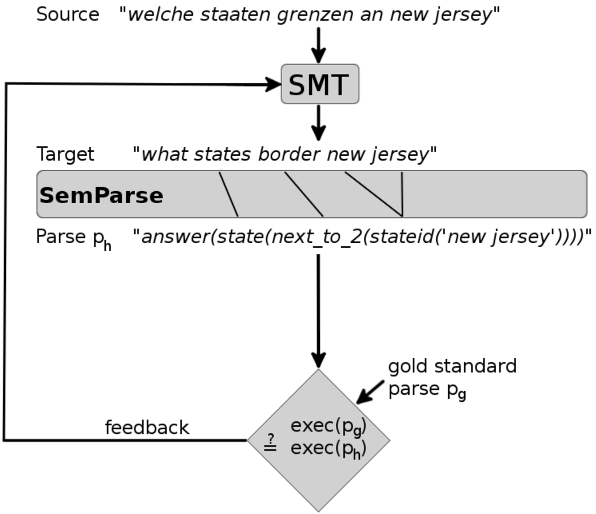

Response-based Learning Cycle

Advantages over learning from independent references:

- Task-approval of system translations **avoids problem of (un)reachability** of references by decoder.
- Structural and lexical variation of predicted and approved translations **broadens learning capabilities**,
- Task-approved supervision signal allows **learn optimally for task-specific aspects of translation quality**.

Response-based Learning Cycle

Advantages over learning from independent references:

- Task-approval of system translations **avoids problem of (un)reachability** of references by decoder.
- Structural and lexical variation of predicted and approved translations **broadens learning capabilities**,
- Task-approved supervision signal allows **learn optimally for task-specific aspects of translation quality**.

Response-based Learning Cycle

Advantages over learning from independent references:

- Task-approval of system translations **avoids problem of (un)reachability** of references by decoder.
- Structural and lexical variation of predicted and approved translations **broadens learning capabilities**,
- Task-approved supervision signal allows **learn optimally for task-specific aspects of translation quality**.

Example

German	Nenne prominente Erhebungen in den USA	
orig. query	Name prominent elevations in the USA	✔
sys. trans	Give prominent surveys in the US	–
sys. trans	Give prominent heights in the US	✔

Response-based Online Learning

- **Execution function** $\mathrm{e}(y) \in \{1, 0\}$ tests whether semantic parse for y receives same answer as gold standard.
- **Cost function** $\mathrm{c}(y^{(i)}, y) = (1 - \mathrm{BLEU}(y^{(i)}, y))$ based on sentence-level BLEU [Nakov et al., 2012].
- y^+ is **surrogate gold-standard translation** w/ positive feedback, high model score s, and low cost c:

$$y^+ = \arg\max_{y \in Y(x^{(i)}):\mathrm{e}(y)=1} \left(\mathrm{s}(x^{(i)}, y; w) - \mathrm{c}(y^{(i)}, y) \right).$$

- y^- **opposite**: negative feedback, high score and cost:

$$y^- = \arg\max_{y \in Y(x^{(i)}):\mathrm{e}(y)=0} \left(\mathrm{s}(x^{(i)}, y; w) + \mathrm{c}(y^{(i)}, y) \right).$$

Response-based Online Learning

- **Execution function** $\mathrm{e}(y) \in \{1, 0\}$ tests whether semantic parse for y receives same answer as gold standard.
- **Cost function** $\mathrm{c}(y^{(i)}, y) = (1 - \mathrm{BLEU}(y^{(i)}, y))$ based on sentence-level BLEU [Nakov et al., 2012].
- y^{+} is **surrogate gold-standard translation** w/ positive feedback, high model score s, and low cost c:

$$y^{+} = \underset{y \in Y(x^{(i)}):\mathrm{e}(y)=1}{\arg\max} \left(\mathrm{s}(x^{(i)}, y; w) - \mathrm{c}(y^{(i)}, y) \right).$$

- y^{-} **opposite**: negative feedback, high score and cost:

$$y^{-} = \underset{y \in Y(x^{(i)}):\mathrm{e}(y)=0}{\arg\max} \left(\mathrm{s}(x^{(i)}, y; w) + \mathrm{c}(y^{(i)}, y) \right).$$

Response-based Online Learning

- **Execution function** $\mathrm{e}(y) \in \{1, 0\}$ tests whether semantic parse for y receives same answer as gold standard.
- **Cost function** $\mathrm{c}(y^{(i)}, y) = (1 - \mathrm{BLEU}(y^{(i)}, y))$ based on sentence-level BLEU [Nakov et al., 2012].
- y^+ is **surrogate gold-standard translation** w/ positive feedback, high model score s, and low cost c:

$$y^+ = \operatorname*{arg\,max}_{y \in Y(x^{(i)}):\mathrm{e}(y)=1} \left(\mathrm{s}(x^{(i)}, y; w) - \mathrm{c}(y^{(i)}, y) \right).$$

- y^- **opposite**: negative feedback, high score and cost:

$$y^- = \operatorname*{arg\,max}_{y \in Y(x^{(i)}):\mathrm{e}(y)=0} \left(\mathrm{s}(x^{(i)}, y; w) + \mathrm{c}(y^{(i)}, y) \right).$$

Response-based Online Learning

- **Ramp loss objective** [Gimpel and Smith, 2012]:

$$\min_{w} \Bigg(- \max_{y \in Y(x^{(i)}) : \mathrm{e}(y)=1} \Big(\mathrm{s}(x^{(i)}, y; w) - \mathrm{c}(y^{(i)}, y) \Big) + \max_{y \in Y(x^{(i)}) : \mathrm{e}(y)=0} \Big(\mathrm{s}(x^{(i)}, y; w) + \mathrm{c}(y^{(i)}, y) \Big) \Bigg).$$

- Stochastic (sub)gradient descent **(SSD) update** [McAllester and Keshet, 2011]:

$$w = w + \phi(x^{(i)}, y^{+}) - \phi(x^{(i)}, y^{-}).$$

Response-based Online Learning

- **Ramp loss objective** [Gimpel and Smith, 2012]:

$$\min_{w} \left(- \max_{y \in Y(x^{(i)}):\mathrm{e}(y)=1} \left(\mathrm{s}(x^{(i)}, y; w) - \mathrm{c}(y^{(i)}, y) \right) \right.$$
$$\left. + \max_{y \in Y(x^{(i)}):\mathrm{e}(y)=0} \left(\mathrm{s}(x^{(i)}, y; w) + \mathrm{c}(y^{(i)}, y) \right) \right).$$

- Stochastic (sub)gradient descent **(SSD) update** [McAllester and Keshet, 2011]:

$$w = w + \phi(x^{(i)}, y^{+}) - \phi(x^{(i)}, y^{-}).$$

Response-based Online Learning

Algorithm 1 Response-based Online Learning

repeat
    for $i = 1, \ldots, n$ **do**
        Receive input string $x^{(i)}$
        Predict translation $\hat{y}$
        Receive task feedback $\mathrm{e}(\hat{y}) \in \{1, 0\}$
        if $\mathrm{e}(\hat{y}) = 1$ **then**
            $y^{+} \leftarrow \hat{y}$
            Store $\hat{y}$ as reference $y^{(i)}$ for $x^{(i)}$
            Compute y^{-}
        else
            $y^{-} \leftarrow \hat{y}$
            Compute y^{+}
        end if
        $w \leftarrow w + \eta(\phi(x^{(i)}, y^{+}) - \phi(x^{(i)}, y^{-}))$
    end for
until Convergence

Experimental Setup

Data

- 880 English queries of GEOQUERY database, manually translated to German [Jones et al., 2012].
- Semantic parser [Andreas et al., 2013]:
  - Monolingual SMT system trained for full accuracy on 880 pairs of English queries and linearized logical forms (= extended parser).
  - Rationale: Translations that match original English query should be rewarded, however, no GEOQUERY test data used in SMT training!
  - Additional comparison with semantic parser trained only on 600 GEOQUERY training data (= standard parser).

Experimental Setup

SMT

- Baseline SMT system CDEC [Dyer et al., 2010] trained on COMMON CRAWL [Smith et al., 2013] web data.
- Discriminative SMT learners:
  - Based on sparse features (rule ids, bigrams in rule source and target, rule shapes) [Simianer et al., 2012].
  - Training for 10 epochs on 10,000-best lists of translations of 600 GEOQUERY training examples.
  - Testing done offline on 280 unseen GEOQUERY test data.

Experimental Setup

Compared variants of discriminative SMT leaners:

- REBOL: Task feedback and cost w.r.t. references.
- EXEC: No cost or manual references, only task feedback.
- RAMPION: No task feedback, only manual references (SGD version of [Gimpel and Smith, 2012]).

Evaluation metrics:

- Precision = percentage of examples with correct answer out of parsed examples; Recall = percentage of total examples answered correctly, F1 = harmonic mean.
- BLEU measured against original English queries.

Experimental Setup

Compared variants of discriminative SMT leaners:

- REBOL: Task feedback and cost w.r.t. references.
- EXEC: No cost or manual references, only task feedback.
- RAMPION: No task feedback, only manual references (SGD version of [Gimpel and Smith, 2012]).

Evaluation metrics:

- Precision = percentage of examples with correct answer out of parsed examples; Recall = percentage of total examples answered correctly, F1 = harmonic mean.
- BLEU measured against original English queries.

Experimental Results w/ Extended Parser

	method	precision	recall	F1	BLEU
1	CDEC	63.67	58.21	60.82	46.53
2	EXEC	70.36	63.57	66.79[1]	48.00[1]
3	RAMPION	75.58	69.64	72.49[12]	**56.64**[12]
4	REBOL	**81.15**	**75.36**	**78.15**[123]	55.66[12]

- REBOL clear winner w.r.t. F1 on correct answers, at non-significant loss in BLEU.
- RAMPION wins w.r.t. BLEU, but far worse F1 than REBOL.
- EXEC improves F1 over CDEC, but far behind others.

Experimental Results w/ Standard Parser

	method	precision	recall	F1	BLEU
1	CDEC	65.59	57.86	61.48	46.53
2	EXEC	66.54	61.79	64.07	46.00
3	RAMPION	67.68	63.57	65.56	**55.67**^{12}
4	REBOL	**70.68**	**67.14**	**68.86**^{12}	**55.67**^{12}

- Training parser on 600 GEOQUERY gives same system ranking as extended parser.
- Statistically significant F1 result differences only for REBOL over EXEC and CDEC.
- BLEU differences not statistically significant between REBOL and RAMPION and between EXEC and CDEC.

Error Analysis

Structural variation in REBOL translations versus reference:

sys	how many inhabitants has new york
ref	how many people live in new york
sys	how big is the population of texas
ref	how many people live in texas
sys	which are the cities of the state with the highest elevation
ref	what are the cities of the state with the highest point
sys	what state borders california
ref	what is the adjacent state of california
sys	what rivers go through states with the least cities
ref	which rivers run through states with fewest cities

Error Analysis

REBOL winning against RAMPION:

reference	RAMPION	REBOL
what is the biggest capital city in the us	what is the largest city in the usa	what is the largest capital in the usa
what state borders new york	what states limits of new york	what states border new york
which states border the state with the smallest area	what states boundaries of the state with the smallest surface area	what states border the state with the smallest surface area

Error Analysis

RAMPION winning against REBOL:

reference	RAMPION	REBOL
how tall is mount mckinley	how high is mount mckinley	what is mount mckinley
what states does the mississippi river run through	through which states runs the mississippi	through which states is the mississippi
which is the highest peak not in alaska	how is the highest peaks of not in alaska is	what is the highest peak in alaska is

Conclusion

Response-based Learning for SMT

- New framework for structured learning in SMT from **weak supervision of task-based response**.
- **Broadening of learning capabilities** by task-approval of structural and lexical variants.
- **Translations still grammatical** due to additional use of cost function w.r.t. human references.

Ongoing and Future Work

- Similar system rankings achieved on Free917 data [Cai and Yates, 2013].
- Scaling semantic parsers to larger coverage ongoing, but difficult.
- Extension to human feedback loop planned.

Thanks for your attention!

Andreas, J., Vlachos, A., and Clark, S. (2013).
Semantic parsing as machine translation.
In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL'13)*, Sofia, Bulgaria.

Berant, J., Chou, A., Frostig, R., and Liang, P. (2013).
Semantic parsing on freebase from question-answer pairs.
In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP'13)*, Seattle, WA.

Cai, Q. and Yates, A. (2013).
Large-scale semantic parsing via schema matching and lexicon extenstion.

In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL'13)*, Sofia, Bulgaria.

Clarke, J., Goldwasser, D., Chang, W.-W., and Roth, D. (2010).
Driving semantic parsing from the world's response.
In *Proceedings of the 14th Conference on Natural Language Learning (CoNLL'10)*, pages 18–27, Uppsala, Sweden.

Dyer, C., Lopez, A., Ganitkevitch, J., Weese, J., Ture, F., Blunsom, P., Setiawan, H., Eidelman, V., and Resnik, P. (2010).
cdec: A decoder, alignment, and learning framework for finite-state and context-free translation models.

In *Proceedings of the ACL 2010 System Demonstrations*, Uppsala, Sweden.

Gimpel, K. and Smith, N. A. (2012).
Structured ramp loss minimization for machine translation.
In *Proceedings of 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2012)*, Montreal, Canada.

Goldwasser, D. and Roth, D. (2014).
Learning from natural instructions.
Machine Learning, 94(2):205–232.

Jones, B. K., Johnson, M., and Goldwater, S. (2012).
Semantic parsing with bayesion tree transducers.

In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL'12)*, Jeju Island, Korea.

Kwiatowski, T., Choi, E., Artzi, Y., and Zettlemoyer, L. (2013).
Scaling semantic parsers with on-the-fly ontology matching.
In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP'13)*, Seattle, WA.

McAllester, D. and Keshet, J. (2011).
Generalization bounds and consistency for latent structural probit and ramp loss.

In *Proceedings of the 25th Annual Conference on Neural Information Processing Sytems (NIPS 2011)*, Granada, Spain.

Nakov, P., Guzmán, F., and Vogel, S. (2012).
Optimizing for sentence-level bleu+1 yields short translations.
In *Proceedings of the 24th International Conference on Computational Linguistics (COLING 2012)*, Bombay, India.

Roy, D. K. (2002).
Learning visually-grounded words and syntax for a scence description task.
Computer Speech and Language, 16:353–385.

Simianer, P., Riezler, S., and Dyer, C. (2012).
Joint feature selection in distributed stochastic learning for large-scale discriminative training in SMT.

In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL 2012)*, Jeju, Korea.

Smith, J. R., Saint-Amand, H., Plamada, M., Koehn, P., Callison-Burch, C., and Lopez, A. (2013).
Dirt cheap web-scale parallel text from the common crawl.
In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL'13)*, Sofia, Bulgaria.

Wong, Y. W. and Mooney, R. J. (2006).
Learning for semantic parsing with statistical machine translation.
In *Proceedings of the Human Language Technology Conference of the North American Chapter of the*

Association for Computational Linguistics (HLT/NAACL'06), New York City, NY.

Yu, C. and Ballard, D. H. (2004).
On the integration of grounding language and learning objects.
In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI'04)*, pages 488–493, San Jose, CA.

Yu, H. and Siskind, J. M. (2013).
Grounded langauge learning from video described with sentences.
In *Proceedings of the 51st Annual Meeting of the Association for computational Linguistics (ACL'13)*, pages 53–63, Sofia, Bulgaria.

Zettlemoyer, L. S. and Collins, M. (2009).

Learning context-dependent mappings from sentences to logical form.
In *Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics (ACL-IJCNLP'09)*, Singapore.