
Learning to Recover Meaning from Unannotated Conversational Interactions

Yoav Artzi

Computer Science & Engineering
University of Washington
Seattle, WA 98195
yoav@cs.washington.edu

Luke Zettlemoyer

Computer Science & Engineering
University of Washington
Seattle, WA 98195
lsz@cs.washington.edu

Introduction Conversational interactions provide significant opportunities for autonomous learning. This abstract describes ongoing work [1] towards building self-improving dialog systems that learn from their mistakes, with little or no human intervention. When a system fails to understand the meaning of the user’s utterance, it can ask for help or clarification. The feedback it receives not only helps to further the current interaction, it also provides a powerful signal for learning to avoid making similar mistakes in the future.

For example, consider applying this learning paradigm to develop algorithms for improved semantic understanding. Figure 1 presents a conversation where the system failed to understand the user’s initial request. To complete the task, the system must use a remediation strategy. Here, it takes the initiative by asking for and confirming each flight constraint in turn. This strategy produces an unnatural conversation but provides supervision for learning the meaning of the original utterance. The ability to automatically learn from this indirect feedback will enable the rapid development of simple self-improving dialog systems that learn through interaction and continually acquire better understanding of language.

SYSTEM: how can i help you?
USER: i would like to fly from atlanta georgia to london england on september twenty fourth in the early evening i would like to return on october first departing from london in the late morning
SYSTEM: leaving what city?
USER: atlanta georgia
SYSTEM: leaving atlanta. going to which city?
USER: london
SYSTEM: arriving in london england. what date would you like to depart atlanta?
USER: september twenty fourth in the early evening
[conversation continues]

Figure 1: Conversation excerpt from a travel-planning dialog. Conversations of this type provide the training data to learn to recover the context-dependent meaning of user utterances.

Learning from this weak signal is challenging. In any specific conversation, system remediations can fail to recover aspects of the original user meaning and can introduce spurious constraints, for example when users change their goals mid conversation. To learn effectively, the algorithm must accumulate evidence from many interactions.

Dialog Semantics We will learn a semantic parser to recover meaning from user utterances. For example, in a flight booking system, the user’s response to the question *Where do you want to fly?*

Response: Seattle on Friday
LF: $(\overset{REQUEST}{x \rightarrow flight1} : \lambda x.to(x, SEA) \wedge date(x, FRI))$

can be mapped to the meaning representation (MR) above: a speech act wrapping a lambda-calculus expression defining the set of flights that match the user’s given constraints. The speech act, *REQUEST* in our example, reflects how the utterance is used within the conversation, rather

than the literal meaning which is captured by the lambda-calculus expression. The speech act includes a mapping of variables from the logical expression to objects discussed in the conversation ($x \rightarrow flight1$), in our case: the flight planned.

One unique challenge for building these representations is that utterance meanings are often not independent from the rest of the interaction. For example, the above user utterance doesn't explicitly state what flight property the noun *Seattle* is describing. The semantic parser has to use the context provided by the current interaction to understand that *Seattle* is the destination. Incorporating this type of reasoning into a semantic parser is a central focus of our work.

Learning Signal To learn from interactions we assume access to logs of conversations with no manual labeling. However, we assume that the dialog system can provide a potentially incorrect, formal representation of the dialog state. This state includes both the meaning of generated system utterances and a representation of the context in which they appear. For example, the second system utterance in Figure 1

SYSTEM: leaving what city?
MR: (ASK: $\lambda x.from(fl, x)$)

may have the automatically generated meaning representation (MR) above: a speech act wrapping a lambda-calculus expression. We also assume the system can provide a set of conversational objects under discussion. In the example above, there is a single conversational object, the *fl* constant, which represents the flight currently being planned.

The information recorded from the series of system utterances following a complex user utterance can be used as indirect supervision for learning. The key challenge is to automatically correlate the incomplete and potentially incorrect information recorded by the system with potential meaning representations of user statements, in such a way as to provide effective signal for learning a high-quality semantic parser for the evolving dialog system.

Related Work There has been significant work on supervised learning for inducing semantic parsers. Earlier work focused on learning for complex, English database queries [23, 19, 24], while more recent approaches developed techniques that can learn to analyze many different languages [21, 22, 15, 11], and that can cope with spontaneous, unedited text [10, 25, 12]. Many current efforts are investigating learning with alternate forms of supervision. Clarke et al. [7] and Liang et al. [14] describe approaches for learning semantic parsers from questions paired with database answers. Chen et al. [6] and Goldwasser et al. [9] take advantage of natural language sentences paired with demonstrations of actions. While Goldwasser et al. [8] presents work on unsupervised learning. Our approach provides an alternative method of supervision that could complement these approaches. Additionally, there has been significant recent work on learning to do other, related, natural language semantic analysis tasks from context-dependent database queries [16, 26], grounded event streams [13, 5, 2], environment interactions [3, 4, 20], and even unannotated text [17, 18].

Initial Progress In recent work [1], we demonstrated that learning from conversations as we described above is possible. To do this we focused on learning to analyze the subset of context-independent sentences, whose meaning does not depend on their position in the conversation. To learn in this space, we presented an approach driven by a loss function that measures how a given meaning representation (a) aligns with system understanding (b) fits the domain of the dialog system. We showed that our approach is able to successfully learn a semantic parser without any manual annotation. This approach significantly outperforms a baseline lacking the conversational signal.

Future Work Our long term goal is to build self-improving dialog systems that continuously learn from their mistakes, with little or no human intervention. While our recent work proves that learning semantics using a conversational signal is feasible, it does so by focusing on learning a context-independent parser. In future work, we aim to build a semantic parser that incorporates contextual information into the parsing process. Such a parser will be able to fully parse user utterances, by inducing implicit meaning indicated by the state of the dialog. This state includes two components that provide significant clues for parsing: the system request that triggered the user's response and the system understanding of the interaction so far. An important part of such work is to model and recover user speech acts, which will reflect how the user utterance is used within the conversation. In the long run we hope that such a learning approach will enable the rapid creation of dialog systems that will be able to automatically improve over time, through interaction with users.

References

- [1] Y. Artzi and L.S. Zettlemoyer. Bootstrapping semantic parsers from conversations. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2011.
- [2] A. Bordes, N. Usunier, R. Collobert, and J. Weston. Towards understanding situated natural language. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2010.
- [3] SRK Branavan, H. Chen, L.S. Zettlemoyer, and R. Barzilay. Reinforcement learning for mapping instructions to actions. In *Proceedings of the Joint Conference of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, 2009.
- [4] SRK Branavan, L.S. Zettlemoyer, and R. Barzilay. Reading between the lines: learning to map high-level instructions to commands. In *Proceedings of the Association for Computational Linguistics*, 2010.
- [5] D.L. Chen, J. Kim, and R.J. Mooney. Training a multilingual sportscaster: using perceptual context to learn language. *Journal of Artificial Intelligence Research*, 37(1):397–436, 2010.
- [6] D.L. Chen and R.J. Mooney. Learning to interpret natural language navigation instructions from observations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2011.
- [7] J. Clarke, D. Goldwasser, M. Chang, and D. Roth. Driving semantic parsing from the world’s response. In *Proceedings of the Conference on Computational Natural Language Learning*, 2010.
- [8] D. Goldwasser, R. Reichart, J. Clarke, and D. Roth. Confidence driven unsupervised semantic parsing. In *Proceedings of the Association of Computational Linguistics*, 2011.
- [9] D. Goldwasser and D. Roth. Learning from natural instructions. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2011.
- [10] Y. He and S. Young. Semantic processing using the hidden vector state model. *Computer Speech and Language*, 19(1):85–106, 2005.
- [11] T. Kwiatkowski, L.S. Zettlemoyer, S. Goldwater, and M. Steedman. Inducing probabilistic CCG grammars from logical form with higher-order unification. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2010.
- [12] T. Kwiatkowski, L.S. Zettlemoyer, S. Goldwater, and M. Steedman. Lexical Generalization in CCG Grammar Induction for Semantic Parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2011.
- [13] P. Liang, M.I. Jordan, and D. Klein. Learning semantic correspondences with less supervision. In *Proceedings of the Joint Conference of the Association for Computational Linguistics the International Joint Conference on Natural Language Processing*, 2009.
- [14] P. Liang, M.I. Jordan, and D. Klein. Learning dependency-based compositional semantics. In *Proceedings of the Association for Computational Linguistics*, 2011.
- [15] W. Lu, H.T. Ng, W.S. Lee, and L.S. Zettlemoyer. A generative model for parsing natural language to meaning representations. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2008.
- [16] S. Miller, D. Stallard, R.J. Bobrow, and R.L. Schwartz. A fully statistical approach to natural language interfaces. In *Proceedings of the Association for Computational Linguistics*, 1996.
- [17] H. Poon and P. Domingos. Unsupervised semantic parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2009.
- [18] H. Poon and P. Domingos. Unsupervised ontology induction from text. In *Proceedings of the Association for Computational Linguistics*, 2010.
- [19] G.N. Ramaswamy and J. Kleindienst. Hierarchical feature-based translation for scalable natural language understanding. In *Proceedings of the International Conference on Spoken Language Processing*, 2000.
- [20] A. Vogel and D. Jurafsky. Learning to follow navigational directions. In *Proceedings of the Association for Computational Linguistics*, 2010.

- [21] Y.W. Wong and R.J. Mooney. Learning for semantic parsing with statistical machine translation. In *Proceedings of the Human Language Technology Conference of the North American Association for Computational Linguistics*, 2006.
- [22] Y.W. Wong and R.J. Mooney. Learning synchronous grammars for semantic parsing with lambda calculus. In *Proceedings of the Association for Computational Linguistics*, 2007.
- [23] J.M. Zelle and R.J. Mooney. Learning to parse database queries using inductive logic programming. In *Proceedings of the National Conference on Artificial Intelligence*, 1996.
- [24] L.S. Zettlemoyer and M. Collins. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 2005.
- [25] L.S. Zettlemoyer and M. Collins. Online learning of relaxed CCG grammars for parsing to logical form. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, 2007.
- [26] L.S. Zettlemoyer and M. Collins. Learning context-dependent mappings from sentences to logical form. In *Proceedings of the Joint Conference of the Association for Computational Linguistics and International Joint Conference on Natural Language Processing*, 2009.