Building Effective Goal-Oriented Dialog Agents

Derek Chen
derekchen14@stanford.edu

Motivation

Recent progress has seen an explosion in dialog systems, including Voice activated bots (such as Siri), text based chat bots (such as those on Slack), and Email-based bots for scheduling meetings (such as X.ai). The prevalence of such personal assistants continues to grow, however far too often these so-called intelligent agents fall short of expectations. Consequently, the goal of this project is to study various dialog agents across a variety of tasks:

- *Craigslist* - buy or sell an item on Craigslist for highest price possible
- *Deal or No Deal* - negotiate for points from shared pool of items
- *Open Movies* - discuss anything about movies with another user
- *Mutual Friends* - two agents collaborate to identify a common friend

We find that rules-based methods often perform quite well and pinpoint strategic areas of improvement for neural-based goal-oriented systems.

Models

Mutual Friends

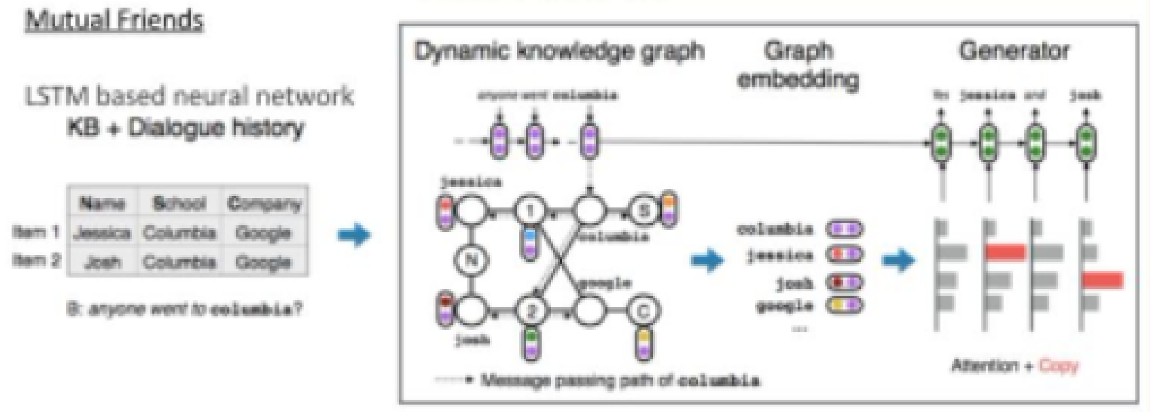

Deal or No Deal

Seq2Seq model with GRUs is used in addition with additional signal from REINFORCE policy with Dialog Rollouts. Reward Function:

$$R(x_t) = \sum_{x_t \in X^A} \gamma^{T-t}(r^A(o) - \mu)$$

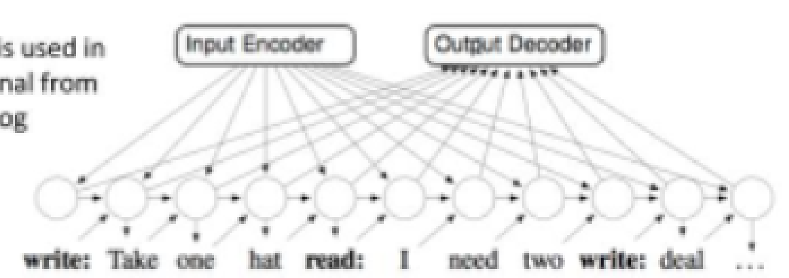

Craigslist & Open Movies

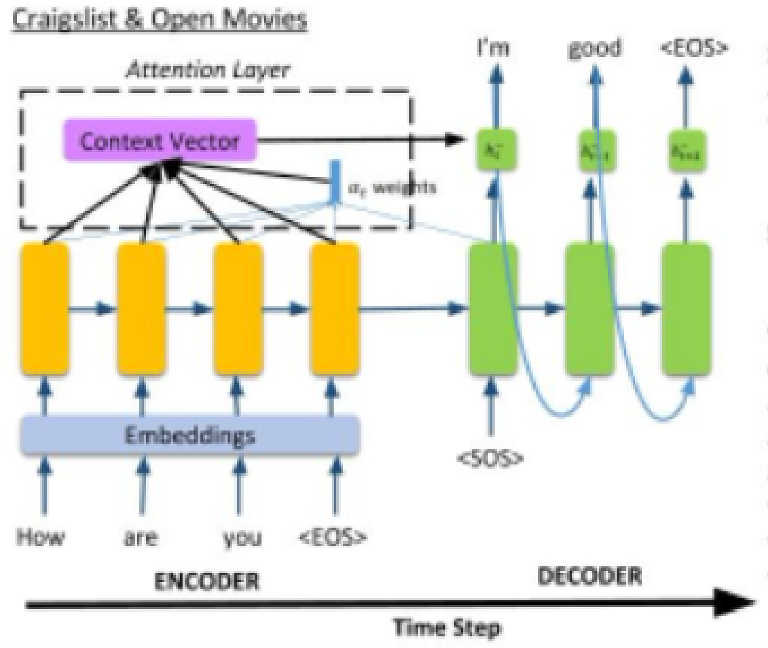

Seq2Seq model with LSTMs and an attention mechanism to better keep track of user belief states during discourse.

Combined

For all tasks, the neural based models all employed some variant of an encoder to capture the user's intent and a decoder for generating the agent's final output. Similarly, all the rule-based models follow a 3-step framework of tokenization, parsing and template retrieval.

Data Collection

Amazon Mechanical Turk

Databases and Scraping

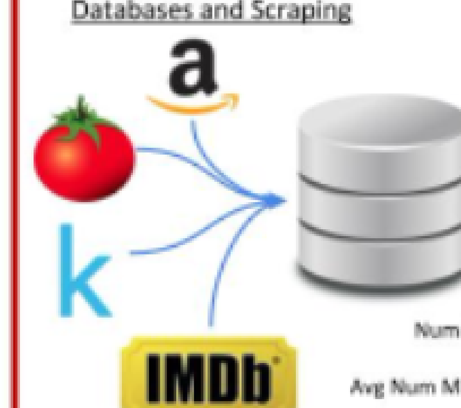

In addition to Mechanical Turk, some tasks required merging data from a variety of databases, scraping websites, and manual entry. For example, data on popular movies was an aggregation of data from IMDB, Kaggle, Rotten Tomatoes, and Amazon movie reviews.

	Craigslist Negotiation	Deal or No Deal	Open Movies	Mutual Friends
Num Dialogs	6682	12,234	1680	11,157
Avg Num Messages	7.39	4.98	10.58	11.41

Discussion

Unified Framework

Tokenization
The lexer tokenizes the raw utterance and identifies the logical forms in the dialog. For example, it finds the movie titles, genres, actors and actresses in the Open Movie Discussion task.

Parsing
The parser analyzes the dialog tokens to uncover the user's intent, such as *Greet*, *Inquire* or *Accept*. The intent can be thought of as a node within a state-based model.

Manager
Generates agent response by selecting from a set of templates labeled by the parser, based on the user intent. Suppose the user *Inquire* about the plot of a movie, so the agent *Inform*.

Results show rules-based systems following a standard framework is sufficient for performing well across a variety of dialog tasks. This is accomplished with a few strategic insights:

- Most discussions can be understood by tracking a limited number of key entities, so if the right items have been identified, then the agent is able to ignore most other items
- Difficulty of open domain discussion can be mitigated by driving conversations into "safe zones" where the agent has better understanding
- Context is critical to conversation, but can be tracked with very small number of states

Using a framework suggests that gains do *not* come from crafting copious amounts of rules.

Takeaways

Context and key entities are important for generating meaningful dialog, and thus improved neural based bots should work to be incorporate this knowledge when making decisions.

Results

Survey Responses

Task	*Rule-based*	*Neural-based*	*Human*	*Oracle*	*Statistic*
Craigslist	0.73	0.61	**1.34**	1.0	avg. sale price
Deal or No Deal	**8.7**	7.1	7.4	10	points
Open Movies	3.5	3.4	**4.8**	5.0	humanlikeness
Mutual Friends	90%	**96%**	82%	100%	success rate
Deal or No Deal	**90.9**	72.1	89.4	100.0	% agreement

Example Dialogue

Agent	Utterance	Intent	Entity
1	Hello, there. let's talk about movies!	Greet	---
2	What is your favorite comedy?	Inquire	Genre
1	I like Ghostbusters all of them	Inform	Title
2	Bill Murray is a legend	Inform	Actor
1	yes he is, i don't like the remakes though	Opinion	---
2	I don't either, I also like the "monty python" ones	Inform	Title
1	great idea, I should watch those again. Great chatting with you.	Closing	---

Future Plans

Incorporating Context: Given the analysis of results, an obvious next step is exploring what happens when the neural systems have a chance to directly incorporate past context. Historically, keeping track of explicit belief states meant modular training, but we might be able to work this now using RL.

Reinforcement Learning: Concretely, the value function should give high rewards correct interpretations of the user intent. In contrast, Facebook's model optimized for maximizing points rather than understanding the user. Further gains may be achieved by directly requesting feedback from the user.

Humanlikeness: All the tasks implicitly placed more value in agents that behave in a human like manner. However, what if real-life users don't mind dealing with a bot as long as the agent effectively solves the user's problems. Knowing how important humanlikeness is helps prioritize how bots are built.

References

1] A.H.Miller,W.Feng,A.Fisch,J.Lu,D.Batra, A.Bordes, D. Parikh, and J. Weston. Parlai: A dialog research software platform. *arXiv preprint arXiv:1705.06476*, 2017.
2] H. He, A. Balakrishnan, M. Eric, and P. Liang. Learn- ing symmetric collaborative dialogue agents with dynamic knowledge graph embeddings. *arXiv preprint arXiv:1704.07130*, 2017
3] J. McAuley and J. Leskovec. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. https://snap.stanford.edu/data/web-Movies.html, 2013.
4] M. Lewis, D. Yarats, Y. N. Dauphin, D. Parikh, and D. Batra. Deal or no deal? end-to-end learning for negotiation dialogues. *arXiv preprint arXiv:1706.05125*, 2017.
5] O. Vinyals and Q. Le. A neural conversational model. *arXiv preprint arXiv:1506.05869*, 2015.
6] T.-H. Wen, D. Vandyke, N. Mrksic, M. Gasic, L. M. Rojas-Barahona, P.-H. Su, S. Ultes, and S. Young. A network-based end-to-end trainable task-oriented dia- logue system. *arXiv preprint arXiv:1604.04562*, 2016.