



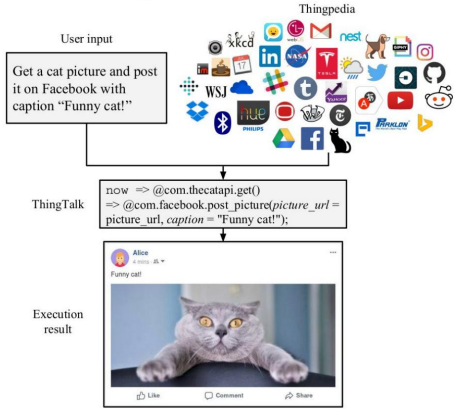
Confirmation Generation for Almond Virtual Assistant

Silei Xu silei@stanford.edu

Almond Virtual Assistant

Almond lets users issue compound commands in natural language and automatically translates them to programs in a formal language called ThingTalk.

Compound commands combine APIs from an open-source, crowdsourced Thingpedia.



Model

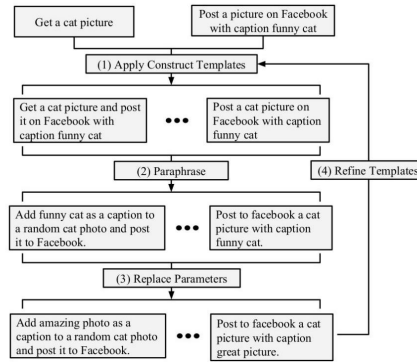
A Seq2Seq Model with attention

- Encoder: 1 embedding layer; 1 layer, bi-directional GRU
- Decoder:
 - 1 embedding layer
 - 1 layer GRU with attention (Luong et al 2015)
- Loss function: masked cross entropy
- Hyperparameters
 - Learning rate: 0.01
 - Batch size: 128
 - # of hidden units: 64
 - # of epochs: 100~300
 - Early stop after BLEU score stop improving within 10 epochs

Dataset

Dataset collected via Amazon Turk

- Primitives: 14,789 paraphrases (1,772 programs)
- Total: 9,777 paraphrases (1,638 programs)
- Total: 24,566 paraphrases (3,410 programs)



Confirmation Generation


It's a hard problem to translate from NL to ThingTalk, and different from Q&A, virtual assistant commands make side effects! We need confirmation!

Since there are infinite number of possible compound commands, it's impossible to provide manually created confirmation. We need machine learning!




Experimental results

Different learning rate:



Different size of neural network:



Discussions

The work is still very raw and incomplete at the current state:

- Training loss is still relatively high despite overfitting
- A couple insights gained from a smaller training set:
 - The model tend to append random tokens after a perfect confirmation, penalty for duplication & long output is added, but does not seem to improve much
 - The loss function is not a very good measurement of the accuracy, BLEU score is used for early stopping