# Intelligent Conversational Bot YUN-NUNG (VIVIAN) CHEN WWW.CSIE.NTU.EDU.TW/~YVCHEN/S105-ICB





renalation

Slides credit from Shawn

Language Generation May 2<sup>nd</sup>, 2017



### Task-Oriented Dialogue System (Young, 2000)

http://rsta.royalsocietypublishing.org/content/358/1769/1389.short



### Task-Oriented Dialogue System (Young, 2000)





### Language Modeling

□ Goal: estimate the probability of a word sequence  $P(w_1, \cdots, w_m)$ 

 Example task: determinate whether a sequence is grammatical or makes more sense



If P(recognize speech)

### N-Gram Language Modeling

- □ Goal: estimate the probability of a word sequence  $P(w_1, \cdots, w_m)$
- N-gram language model
  - Probability is conditioned on a window of (n-1) previous words  $P(w_1, \dots, w_m) = \prod_{i=1}^m P(w_i \mid w_1, \dots, w_{i-1}) \approx \prod_{i=1}^m P(w_i \mid w_{i-(n-1)}, \dots, w_{i-1})$ 
    - Estimate the probability based on the training data

 $P(\text{beach}|\text{nice}) = \frac{C(\text{nice beach})}{C(\text{nice})} \leftarrow \text{Count of "nice beach" in the training data}$ 

Issue: some sequences may not appear in the training data

### N-Gram Language Modeling

- Training data:
  - The dog ran .....
  - The cat jumped .....

P(jumped | dog) = 0.0001 P(ran | cat) = 0.0001

give some small probability → smoothing

- The probability is not accurate.
- The phenomenon happens because we cannot collect all the possible text in the world as training data.

### Neural Language Modeling

□ Idea: estimate  $P(w_i | w_{i-(n-1)}, \dots, w_{i-1})$  not from count, but from the NN prediction

P("wreck a nice beach") = P(wreck|START)P(a|wreck)P(nice|a)P(beach|nice)



### Neural Language Modeling

10



### Neural Language Modeling

11

The input layer (or hidden layer) of the related words are close



 If P(jump|dog) is large, P(jump|cat) increase accordingly (even there is not "... cat jump ..." in the data)

Smoothing is automatically done

### RNNLM

- Idea: condition the neural network on <u>all previous words</u> and tie the weights at each time step
- Assumption: temporal information matters



## <sup>13</sup> Natural Language Generation

**Traditional Approaches** 

### Natural Language Generation (NLG)

14

Mapping dialogue acts into natural language

inform(name=Seven\_Days, foodtype=Chinese)

Seven Days is a nice Chinese restaurant

#### **Template-Based NLG**

#### Define <u>a set of rules</u> to map frames to NL

Semantic Frame	Natural Language	
confirm()	"Please tell me more about the product your are looking for."	
confirm(area=\$V)	"Do you want somewhere in the \$V?"	
confirm(food=\$V)	"Do you want a \$V restaurant?"	
confirm(food=\$V,area=\$W)	"Do you want a \$V restaurant in the \$W."	

*Pros:* simple, error-free, easy to control *Cons:* time-consuming, rigid, poor scalability

#### Class-Based LM NLG (Oh and Rudnicky, 2000)

http://dl.acm.org/citation.cfm?id=1117568

□ Class-based language modeling  $P(X \mid c) = \sum_{t} \log p(x_t \mid x_0, x_1, \dots, x_{t-1}, c)$ □ NLG by decoding  $X^* = \arg \max_X P(X \mid c)$  $X^* = \arg \max_X P(X \mid c)$ 



*Pros:* easy to implement/ understand, simple rules *Cons:* computationally inefficient

16

#### Phrase-Based NLG (Mairesse et al, 2010)

http://dl.acm.org/citation.cfm?id=1858838 Charlie Chan Chinese Restaurant Cineworld in the centre is a near Phrase DBN Semantic DBN Charlie Chan Chinese Cineworld restaurant centre food name type near area inform inform inform inform inform

Inform(name=Charlie Chan, food=Chinese, type= restaurant, near=Cineworld, area=centre)

#### realization phrase semantic stack

$r_t$	St	$h_t$	$l_t$
<s></s>	START	START	START
The Rice Boat	inform(name(X))	X	inform(name)
is a	inform	inform	EMPTY
restaurant	inform(type(restaurant))	restaurant	inform(type)
in the	inform(area)	area	inform
riverside	inform(area(riverside))	riverside	inform(area)
area	inform(area)	area	inform
that	inform	inform	EMPTY
serves	inform(food)	food	inform
French	inform(food(French))	French	inform(food)
food	inform(food)	food	inform
8	END	END	END

**Pros:** efficient, good performance **Cons:** require semantic alignments

## <sup>18</sup> Natural Language Generation

Deep Learning Approaches

#### RNN-Based LM NLG (Wen et al., 2015)

19 http://www.anthology.aclweb.org/W/W15/W15-46.pdf#page=295 Input dialogue act 1-hot Inform(name=Din Tai Fung, food=Taiwanese) representation 0, 0, 1, 0, 0, ..., 1, 0, 0, ..., 1, 0, 0, 0, 0, 0... } SLOT NAME SLOT\_FOOD <EOS> serves conditioned on the dialogue act <BOS> SLOT\_NAME SLOT\_FOOD serves Output <BOS> Din Tai Fung Taiwanese serves delexicalisation Slot weight tying

#### Handling Semantic Repetition

- Issue: semantic repetition
  - Din Tai Fung is a great Taiwanese restaurant that serves Taiwanese.
  - Din Tai Fung is a child friendly restaurant, and also allows kids.
- Deficiency in either model or decoding (or both)
- Mitigation
  - Post-processing rules (Oh & Rudnicky, 2000)
  - Gating mechanism (Wen et al., 2015)
  - Attention (Mei et al., 2016; Wen et al., 2015)

### Visualization



### Semantic Conditioned LSTM (Wen et al., 2015)

- Original LSTM cell
  - $\mathbf{i}_{t} = \sigma(\mathbf{W}_{wi}\mathbf{x}_{t} + \mathbf{W}_{hi}\mathbf{h}_{t-1})$  $\mathbf{f}_{t} = \sigma(\mathbf{W}_{wf}\mathbf{x}_{t} + \mathbf{W}_{hf}\mathbf{h}_{t-1})$  $\mathbf{o}_{t} = \sigma(\mathbf{W}_{wo}\mathbf{x}_{t} + \mathbf{W}_{ho}\mathbf{h}_{t-1})$  $\hat{\mathbf{c}}_{t} = \tanh(\mathbf{W}_{wc}\mathbf{x}_{t} + \mathbf{W}_{hc}\mathbf{h}_{t-1})$  $\mathbf{c}_{t} = \mathbf{f}_{t}\odot\mathbf{c}_{t-1} + \mathbf{i}_{t}\odot\hat{\mathbf{c}}_{t}$
  - $\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t)$
- Dialogue act (DA) cell

$$\mathbf{r}_t = \sigma(\mathbf{W}_{wr}\mathbf{x}_t + \mathbf{W}_{hr}\mathbf{h}_{t-1})$$

 $\mathbf{d}_t = \mathbf{r}_t \odot \mathbf{d}_{t-1}$ 

Modify Ct

 $\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \hat{\mathbf{c}}_t + \tanh(\mathbf{W}_{dc} \mathbf{d}_t)$ 



### Attentive Encoder-Decoder for NLG

Slot & value embedding

 $\mathbf{z}_i = \mathbf{s}_i + \mathbf{v}_i$ 

Attentive meaning representation

 $e_{ti} = \mathbf{v}^{\mathrm{T}} \tanh(\mathbf{W}_{hm}\mathbf{h}_{t-1} + \mathbf{W}_{zm}\mathbf{z}_{i})$  $\alpha_{ti} = \mathrm{softmax}(e_{ti})$ 

$$\mathbf{d}_t = \mathbf{a} \oplus \sum_i \alpha_{ti} \mathbf{z}_i$$



### **Attention Heat Map**



id-5: skyCover(time=6-21, mode=50-75); id-10: precipChance(time=06-21, min=19, mean=32, max=73);

id-15: thunderChance(time=13-21, mode=SChc)

Figure 3: An example generation for a set of records from WEATHERGOV.

### **Model Comparison**



#### Structural NLG (Dušek and Jurčíček, 2016)



inform name X-name inform eattype restaurant <GO> X is a restaurant

#### Contextual NLG (Dušek and Jurčíček, 2016)

- Goal: adapting users' way of speaking, providing contextaware responses
  - Context encoder
  - Seq2Seq model



https://www.aclweb.org/anthology/W/W16/W16-36.pdf#page=203



### Decoder Sampling Strategy

#### Decoding procedure



- Greedy search
- Beam search
- Random search

### **Greedy Search**

29

□ Select the next word with the *highest* probability



### **Beam Search**

30

 Select the next k-best words and keep a beam with width=k for following decoding





### **Random Search**

- Randomly select the next word
  - Higher diversity
  - Can follow a probability distribution





### **Chit-Chat Bot**

# Neural conversational model Non task-oriented

Human: what do you think about tesla ?
Machine: he 's a good conductor .
Human: what do you think about bill gates ?
Machine: he 's a good man .
Human: What do you think about messi ?
Machine: he 's a great player .
Human: what do you think about cleopatra ?
Machine: oh , she 's very regal .
Human: what do you think about england during the reign of elizabeth ?
Machine: it was a great place .

**Human:** what is the purpose of life? Machine: to serve the greater good. **Human:** what is the purpose of living? Machine: to live forever. **Human:** what is the purpose of existence ? Machine: to find out what happens when we get to the planet earth. Human: where are you now? Machine: *i* 'm in the middle of nowhere. **Human:** what is the purpose of dying? Machine: to have a life. **Human:** what is the purpose of being intelligent 2 Machine: to find out what it is. Human: what is the purpose of emotions? Machine: i don 't know

### Many-to-Many

- Both input and output are both sequences → <u>Sequence-to-</u> sequence learning
  - E.g. Machine Translation (machine learning→機器學習)



### A Neural Conversational Model

□ Seq2Seq [Vinyals and Le, 2015]



Human: what do you think about tesla ?
Machine: he 's a good conductor .
Human: what do you think about bill gates ?
Machine: he 's a good man .
Human: What do you think about messi ?
Machine: he 's a great player .
Human: what do you think about cleopatra ?
Machine: oh , she 's very regal .
Human: what do you think about england during the reign of elizabeth ?
Machine: it was a great place .

**Human:** what is the purpose of life? Machine: to serve the greater good. Human: what is the purpose of living? Machine: to live forever. **Human:** what is the purpose of existence ? Machine: to find out what happens when we get to the planet earth. Human: where are you now? Machine: *i* 'm in the middle of nowhere. **Human:** what is the purpose of dying? Machine: to have a life. **Human:** what is the purpose of being intelligent 2 Machine: to find out what it is . **Human:** what is the purpose of emotions? Machine: i don 't know.

### **Chit-Chat Bot**

36



LSTM Encoder

LSTM Decoder

電視影集 (~40,000 sentences)、美國總統大選辯論

### Sci-Fi Short Film - SUNSPRING

SUNSPRING

https://www.youtube.com/watch?v=LY7x2Ihqj

### **Concluding Remarks**

- The three pillars of deep learning for NLG
  - Distributed representation generalization
  - Recurrent connection long-term dependency
  - **Conditional RNN** flexibility/creativity
- Useful techniques in deep learning for NLG
  - Learnable gates
  - Attention mechanism
- Generating longer/complex sentences
- Phrase dialogue as conditional generation problem
  - Conditioning on raw input sentence → chit-chat bot
  - Conditioning on both structured and unstructured sources → task-completing dialogue system