Generating Informative and Diverse Conversational Responses via Adversarial Information Maximization

Yizhe Zhang **Michel Galley** **Jianfeng Gao**
Zhe Gan **Xiujun Li** **Chris Brockett** **Bill Dolan**
Microsoft Research, Redmond, WA, USA
{yizzhang,mgalley,jfgao,zhgan,xiul,chrisbkt,billdol}@microsoft.com

NIPS18

Motivations

- The standard MLE average out the responses in the training data. (safe response problem)
- The problem is in fact twofold:
  - Diversity
  - Informativeness

Motivations

- Diverse but uninformative
  - I dont know, I haven’t a clue, I couldn’t tell you
- Informative but not diverse
  - “I like music” but never “I like Jazz”

Motivation

- To strike a right balance between informativeness and diversity.
- MMI attacked informativeness
- GAN tried, however, was explicitly not for either informativeness or diversity

Motivations

- Adversarial Information Maximization
  - GAN for diversity
  - Variational information Maximization Objective (VIMO) for informativeness

GAN

- Greedy sampling (argmax)
- Temperature trick for differentiability
- Compute cosine similarity between query and response
- WGAN-like objective
    - $$\mathcal{L}_{\mathrm{GAN}}(\theta,\psi) = -\mathbb{E}_{T,\tilde{T},S}\Big[f\Big(D_\psi(T,S) - D_\psi(\tilde{T},S)\Big)\Big]$$

VIMO

$$
\begin{aligned}
I_{p^e}(S,T) &\triangleq \mathbb{E}_{p^e(S,T)} \log \frac{p^e(S,T)}{p(S)p^e(T)} \\
=& H(S) + \mathbb{E}_{p^e(T)} D_{KL}(p^e(S|T), q_\phi(S|T)) + \mathbb{E}_{p^e(S,T)} \log q_\phi(S|T) \\
\geqslant& \mathbb{E}_{p(S)} \mathbb{E}_{p_\theta(T|S)} \log q_\phi(S|T) \triangleq \mathcal{L}_{\text{MI}}(\theta, \phi)\,,
\end{aligned}
$$

Dual Learning

$$\min_{\psi} \max_{\theta,\phi} \mathcal{L}_{\mathrm{DAIM}}$$
$$= -\mathbb{E}_{(T,\tilde{T},S)\sim p^e_\theta} f(D_\psi(S,T) - D_\psi(S,\tilde{T}))$$
$$- \mathbb{E}_{(T,\tilde{S},S)\sim p^d_\phi} f(D_\psi(S,T) - D_\psi(\tilde{S},T))$$
$$+ \lambda \cdot \mathbb{E}_{p(S)} \mathbb{E}_{p_\theta(T|S)} \log q_\phi(S|T)$$
$$+ \lambda \cdot \mathbb{E}_{p(T)} \mathbb{E}_{q_\phi(S|T)} \log p_\theta(T|S)\,,$$

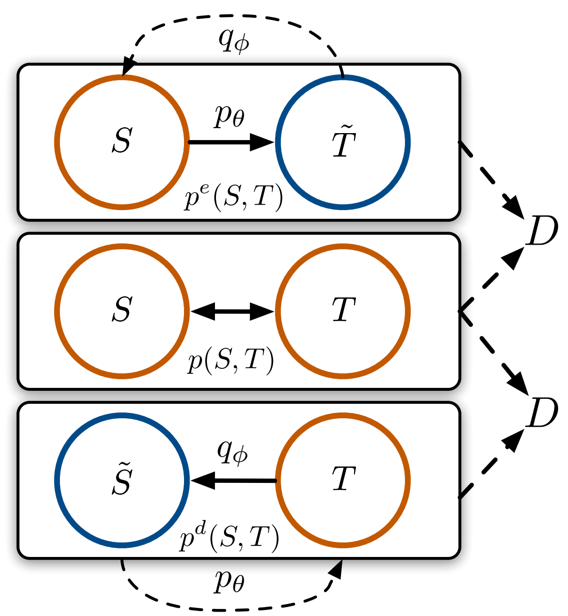

Experiments

Table 1: Quantitative evaluation on the Reddit dataset. (* is implemented based on [4].)

Models	Relevance					Diversity		
	BLEU	ROUGE	Greedy	Average	Extreme	Dist-1	Dist-2	Ent-4
seq2seq	1.85	0.9	1.845	0.591	0.342	0.040	0.153	6.807
cGAN	1.83	0.9	1.872	0.604	0.357	0.052	0.199	7.864
AIM	**2.04**	**1.2**	**1.989**	**0.645**	0.362	0.050	0.205	8.014
DAIM	1.93	1.1	1.945	0.632	**0.366**	**0.054**	**0.220**	**8.128**
MMI*	1.87	1.1	1.864	0.596	0.353	0.046	0.127	7.142
Human	-	-	-	-	-	0.129	0.616	9.566

Table 3: Human evaluation results. Results of statistical significance are shown in bold.

Methods	Informativeness				Relevance			
	Method A		Method B		Method A		Method B	
MMI-AIM	MMI	0.496	AIM	0.504	MMI	0.501	AIM	0.499
MMI-cGAN	MMI	0.505	cGAN	0.495	MMI	**0.514**	cGAN	**0.486**
MMI-DAIM	MMI	**0.484**	DAIM	**0.516**	MMI	0.503	DAIM	0.497
MMI-seq2seq	MMI	0.510	seq2seq	0.490	MMI	**0.518**	seq2seq	**0.482**
seq2seq-cGAN	seq2seq	0.487	cGAN	0.513	seq2seq	0.492	cGAN	0.508
seq2seq-AIM	seq2seq	0.478	AIM	0.522	seq2seq	0.492	AIM	0.508
seq2seq-DAIM	seq2seq	**0.468**	DAIM	**0.532**	seq2seq	**0.475**	DAIM	**0.525**
Human-DAIM	Human	**0.615**	DAIM	**0.385**	Human	**0.600**	DAIM	**0.400**

Sequential Matching Network: A New Architecture for Multi-turn Response Selection in Retrieval-Based Chatbots

Yu Wu[†], Wei Wu[‡], Chen Xing[◇], Zhoujun Li[†*], Ming Zhou[‡]
[†]State Key Lab of Software Development Environment, Beihang University, Beijing, China
[◇]College of Computer and Control Engineering, Nankai University, Tianjin, China
[‡] Microsoft Research, Beijing, China
{wuyu,lizj}@buaa.edu.cn {wuwei,v-chxing,mingzhou}@microsoft.com

- EMNLP17

Model

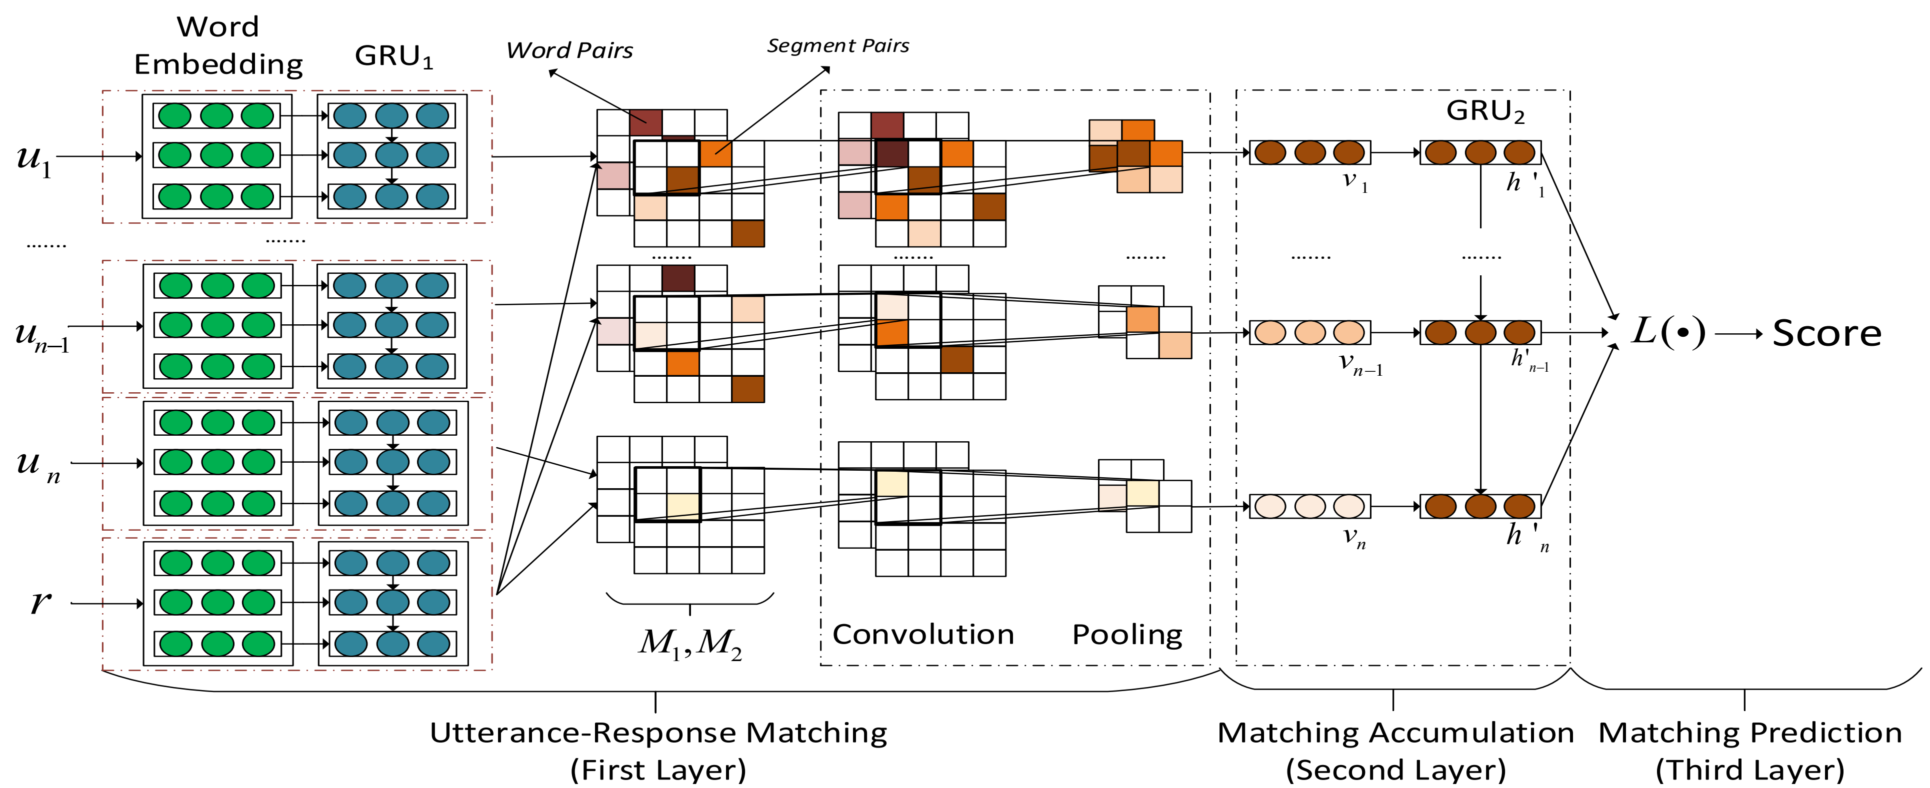

Word Embedding
GRU1
Word Pairs
Segment Pairs
GRU2
u1
un-1
un
r
M1, M2
Convolution
Pooling
v1
vn-1
vn
h'1
h'n-1
h'n
L(•)
Score
Utterance-Response Matching (First Layer)
Matching Accumulation (Second Layer)
Matching Prediction (Third Layer)

L?

- Last: only use the last hidden state
- Static: static weights for hidden states
- Dynamic: dynamic weights for hidden states

Motivations

- Existing models: they first represent the whole context as a vector and then match the context vector with a response vector.

- Their approach: SMN matches a response with each utterance in the context at the beginning and encodes important information in each pair into a matching vector. The matching vectors are then accumulated in the utterances’ temporal order to model their relationships.

Experiments

	Ubuntu Corpus				Douban Conversation Corpus					
	R_2@1	R_{10}@1	R_{10}@2	R_{10}@5	MAP	MRR	P@1	R_{10}@1	R_{10}@2	R_{10}@5
TF-IDF	0.659	0.410	0.545	0.708	0.331	0.359	0.180	0.096	0.172	0.405
RNN	0.768	0.403	0.547	0.819	0.390	0.422	0.208	0.118	0.223	0.589
CNN	0.848	0.549	0.684	0.896	0.417	0.440	0.226	0.121	0.252	0.647
LSTM	0.901	0.638	0.784	0.949	0.485	0.527	0.320	0.187	0.343	0.720
BiLSTM	0.895	0.630	0.780	0.944	0.479	0.514	0.313	0.184	0.330	0.716
Multi-View	0.908	0.662	0.801	0.951	0.505	0.543	0.342	0.202	0.350	0.729
DL2R	0.899	0.626	0.783	0.944	0.488	0.527	0.330	0.193	0.342	0.705
MV-LSTM	0.906	0.653	0.804	0.946	0.498	0.538	0.348	0.202	0.351	0.710
Match-LSTM	0.904	0.653	0.799	0.944	0.500	0.537	0.345	0.202	0.348	0.720
Attentive-LSTM	0.903	0.633	0.789	0.943	0.495	0.523	0.331	0.192	0.328	0.718
Multi-Channel	0.904	0.656	0.809	0.942	0.506	0.543	0.349	0.203	0.351	0.709
Multi-Channel$_{exp}$	0.714	0.368	0.497	0.745	0.476	0.515	0.317	0.179	0.335	0.691
SMN$_{last}$	**0.923**	**0.723**	**0.842**	**0.956**	**0.526**	**0.571**	**0.393**	**0.236**	**0.387**	0.729
SMN$_{static}$	**0.927**	**0.725**	**0.838**	**0.962**	**0.523**	**0.572**	**0.387**	**0.228**	**0.387**	0.734
SMN$_{dynamic}$	**0.926**	**0.726**	**0.847**	**0.961**	**0.529**	**0.569**	**0.397**	**0.233**	**0.396**	0.724

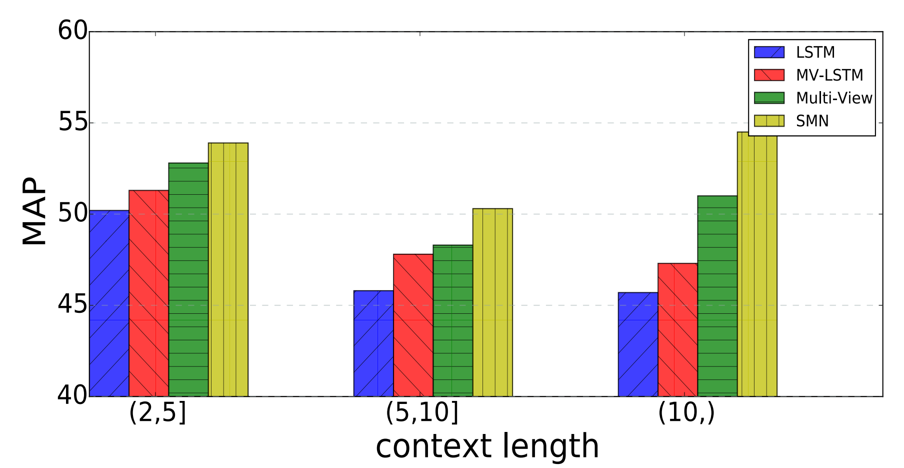

Figure 3: Comparison across context length