Multi-Turn Response Selection for Chatbots with Deep Attention Matching Network

Xiangyang Zhou[*] , *Lu Li*[*], *Daxiang Dong*, *Yi Liu*, *Ying Chen*,

Wayne Xin Zhao, *Dianhai Yu* and *Hua Wu*

Baidu – Natural Language Processing

Background

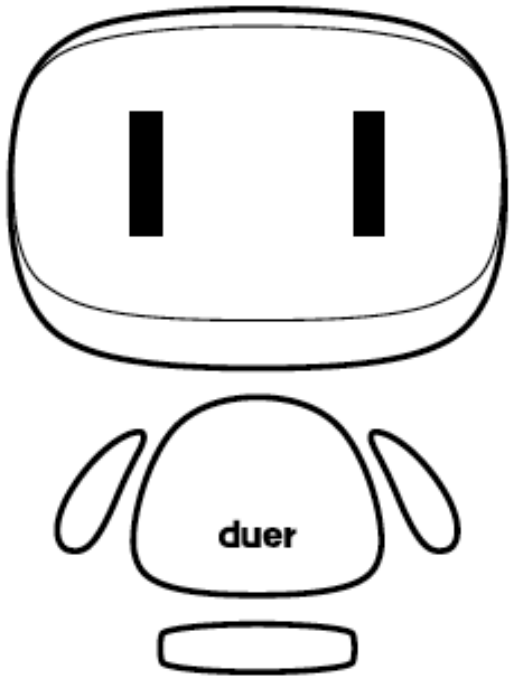

Chatbots

- Naturally and consistently converse with human-beings on open-domain topics.

- Data-driven
  - Retrieval-based method
  - Generation-based method
  - System ensemble

Context-Response Matching

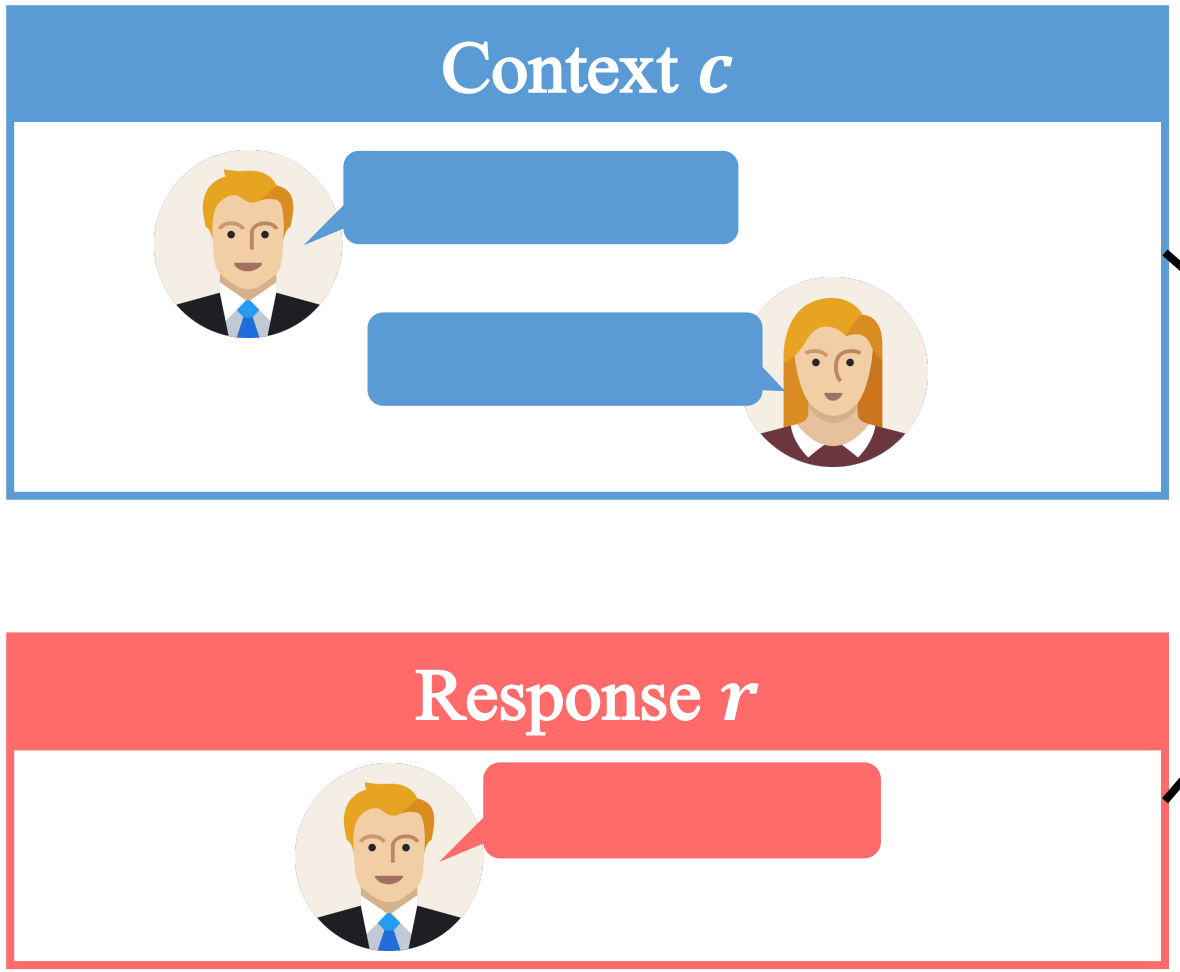

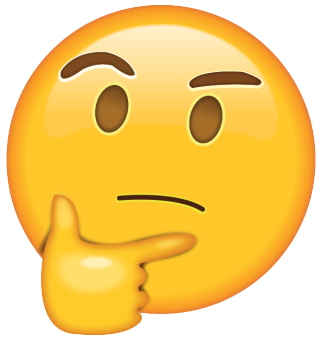

$$p(is\_good|c,r)$$

Retrieval-based Chatbot

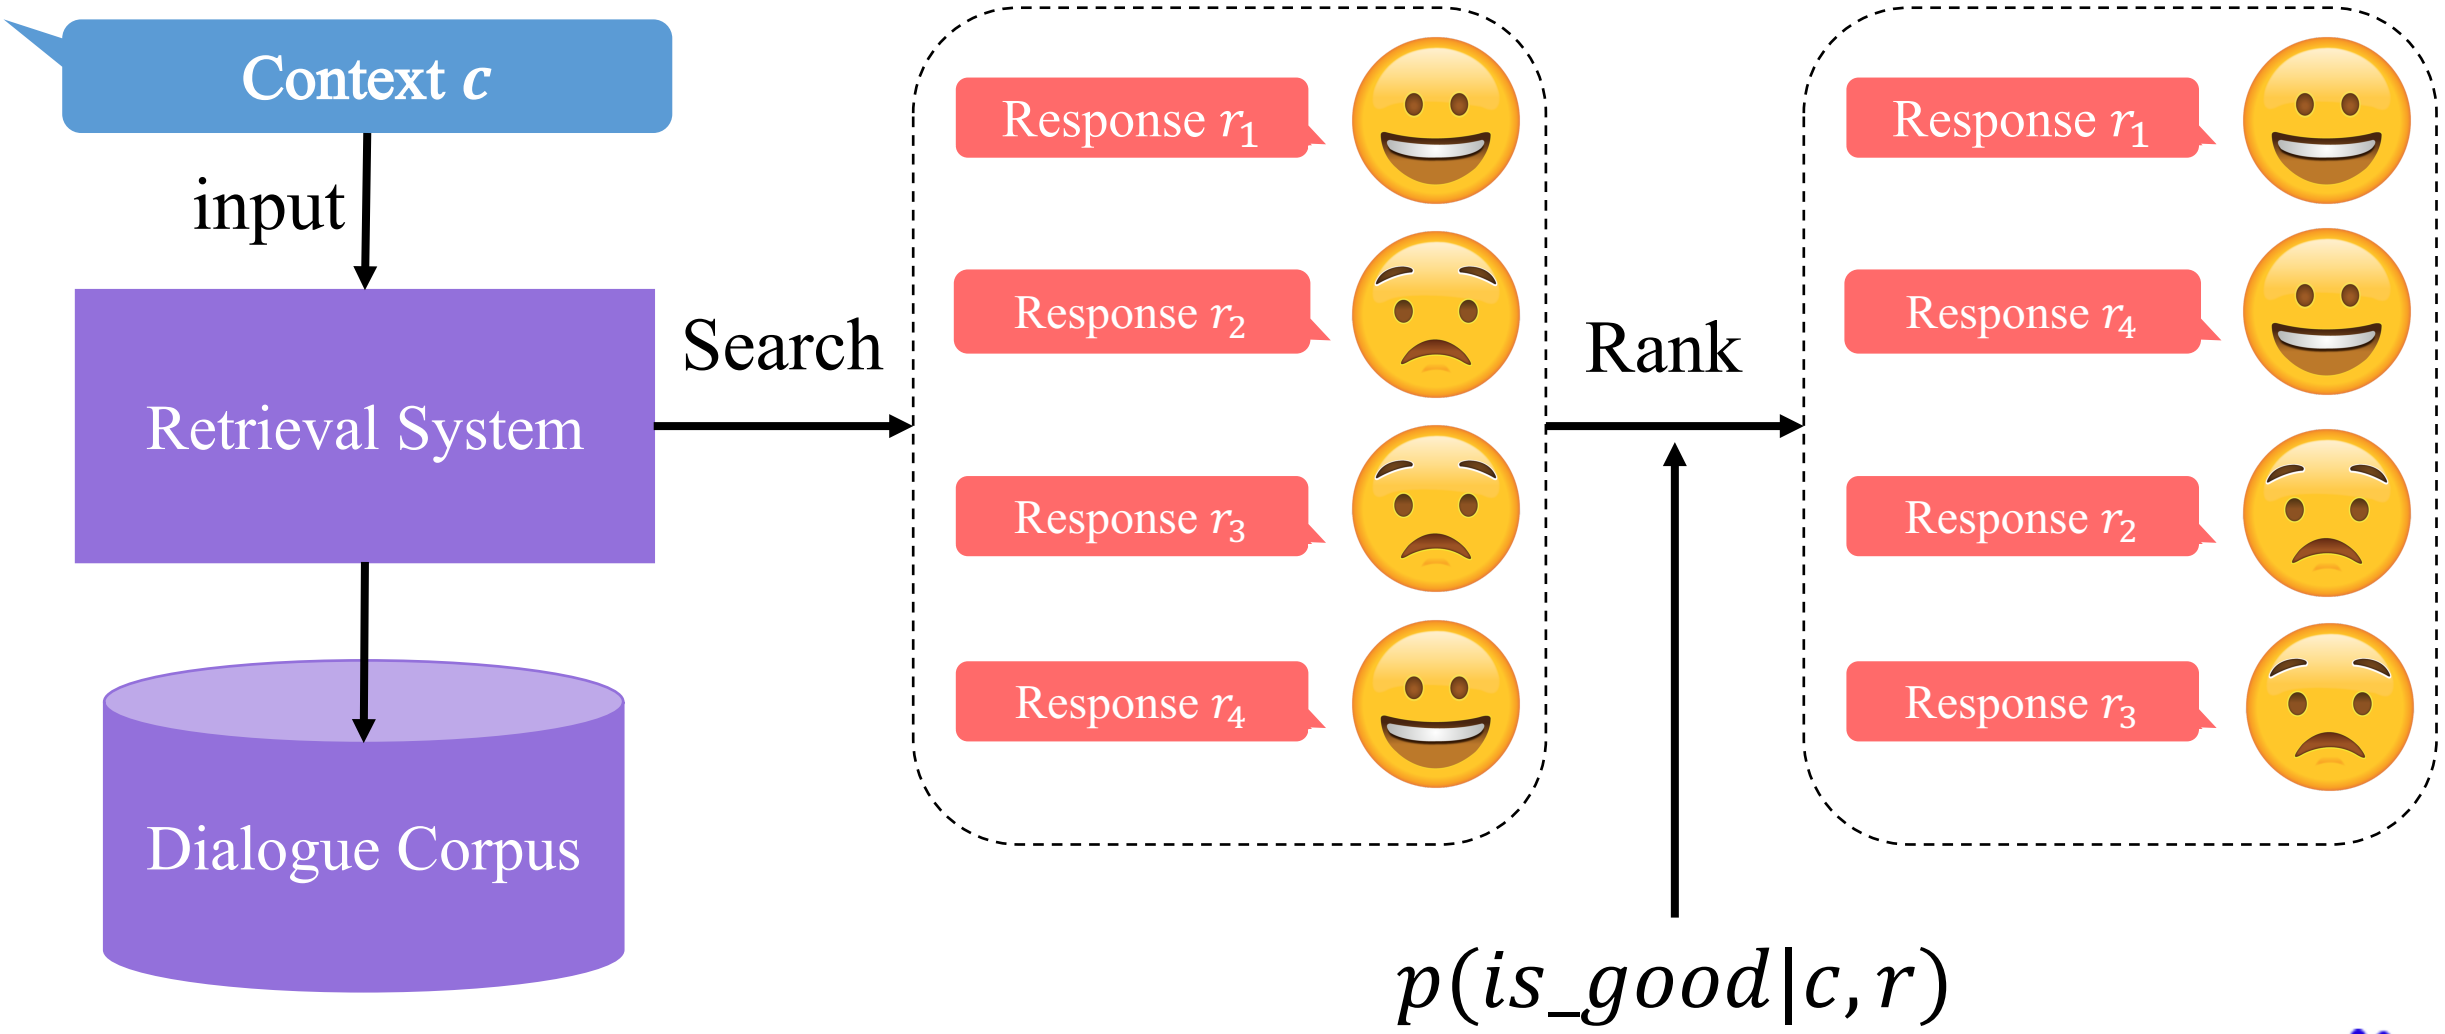

An information retrieval approach to short text conversation. Ji et al., 2014

Adversarial Dialogue Generation

Dialogue Corpus

Real Conversation

Sampled <c, r>

Generator $G(r^*|c)$

$p(is\_good|c,r)$

Discriminator

Adversarial Learning for Neural Dialogue Generation. Li et al., EMNLP-2017

Challenges

Challenges

- Detecting truly matched segment pairs across context and response.
- Segment pairs could be matched at different granularities.
- Segment pairs, across context and response, could be matched because of textual relevance or semantic dependency.

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg - get-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it for, its just reassurance as I don't know the debian package manager

Baidu 百度

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg - get-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it for, its just reassurance as I don't know the debian package manager

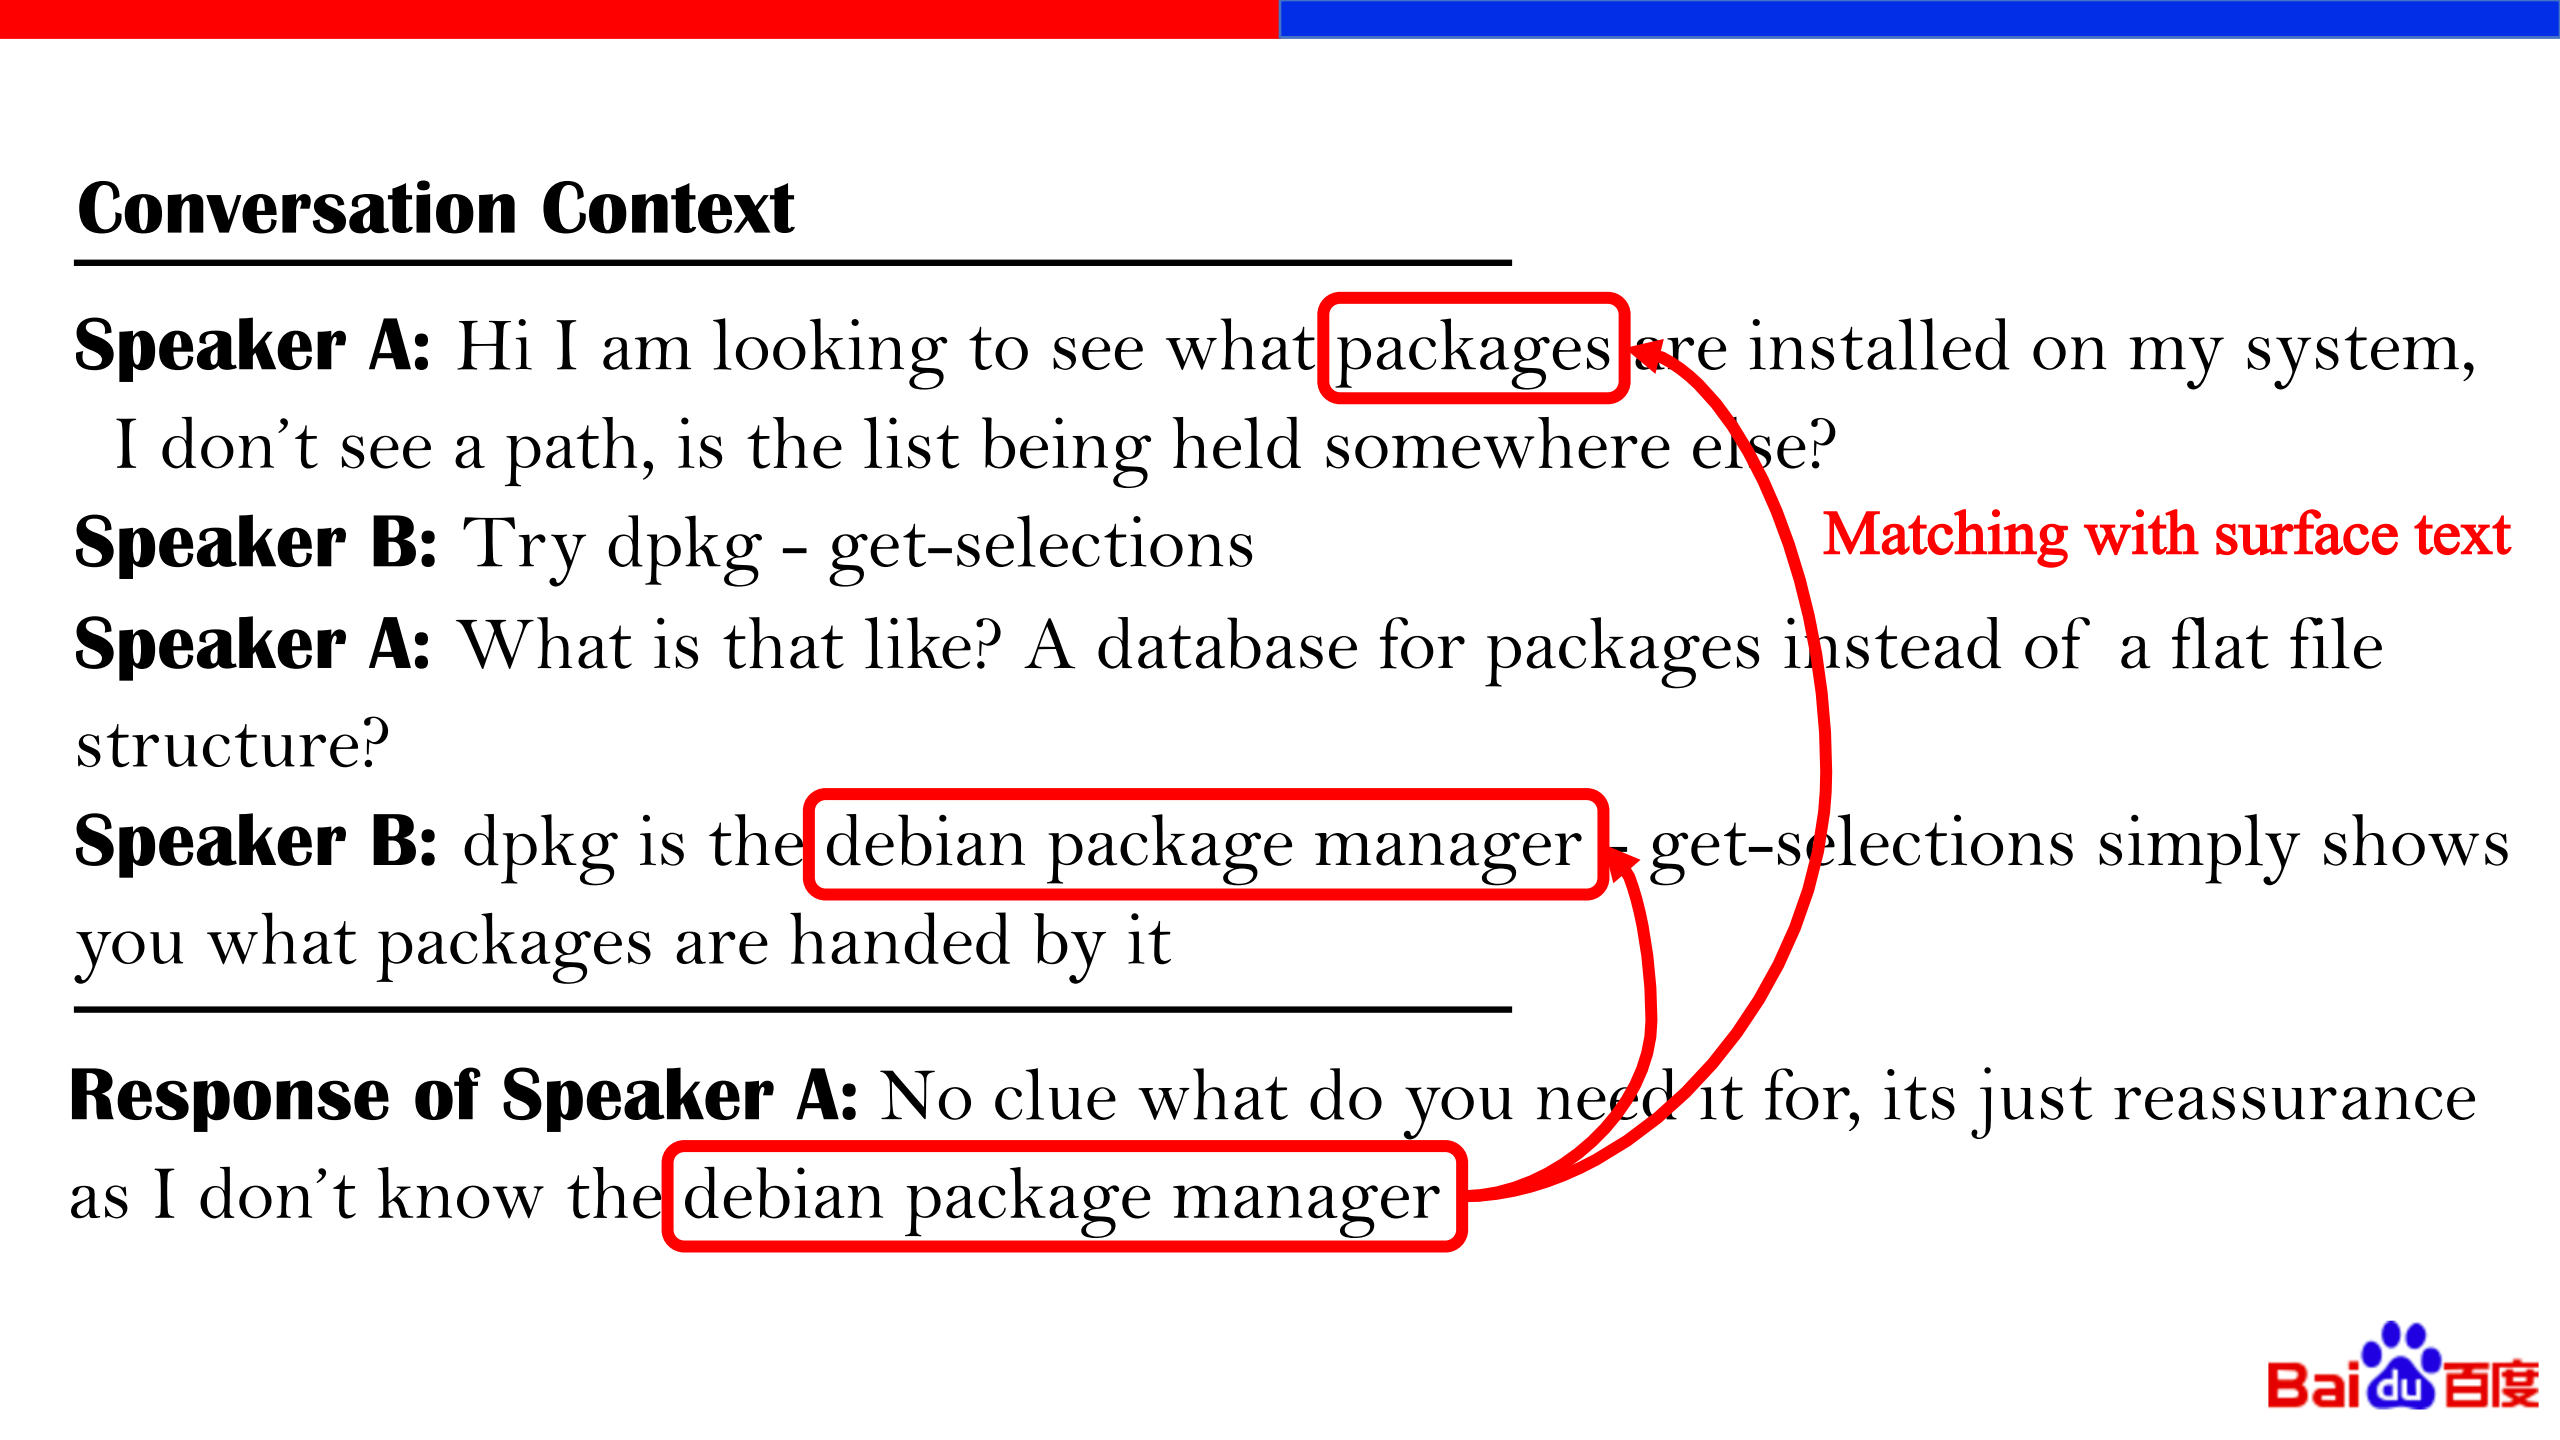

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg --get-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

Matching with dependency

Response of Speaker A: No clue what do you need it for, its just reassurance as I don't know the debian package manager

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg - get-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it for, its just reassurance as I don't know the debian package manager

Baidu 百度

Motivation

Previous works

Use GRU/LSTM to encode segments and match context with response only considering textual relevance.

Motivation

- ***Self-Attention***: Using intra-attention of utterance/response to gradually construct multi-grained semantic representations.
- ***Cross-Attention***: Using attention across context and response to match with dependency information.

Baidu 百度

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg - get-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it for, its just reassurance as I don't know the debian package manager

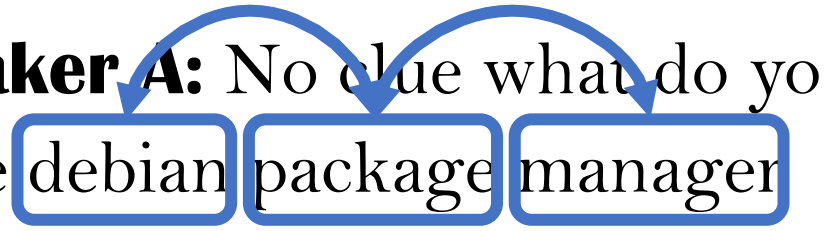

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg - get-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it for, its just reassurance as I don't know the debian package manager

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg - get-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it for its just reassurance as I don't know the debian package manager

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg - get-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it for its just reassurance as I don't know the debian package manager

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg - get-selections

Speaker A: What is that like? A dat for packages instead of a flat file structure?

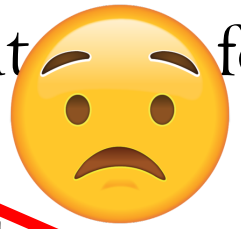

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it for, its just reassurance as I don't know the debian package manager

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: Try dpkg - get-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

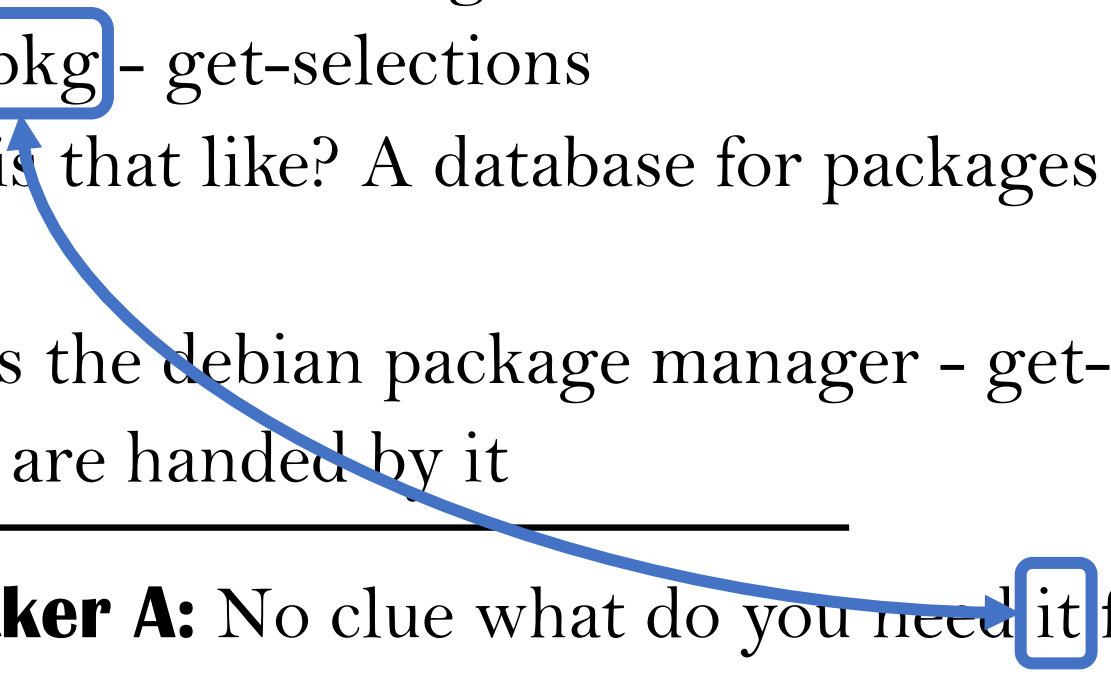

Response of Speaker A: No clue what do you need it for, its just reassurance as I don't know the debian package manager

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: dkpg + ``it" et-selections

Speaker A: What is that like? A database for packages instead of a flat file structure?

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it + ``dpkg" reassurance as I don't know the debian package manager

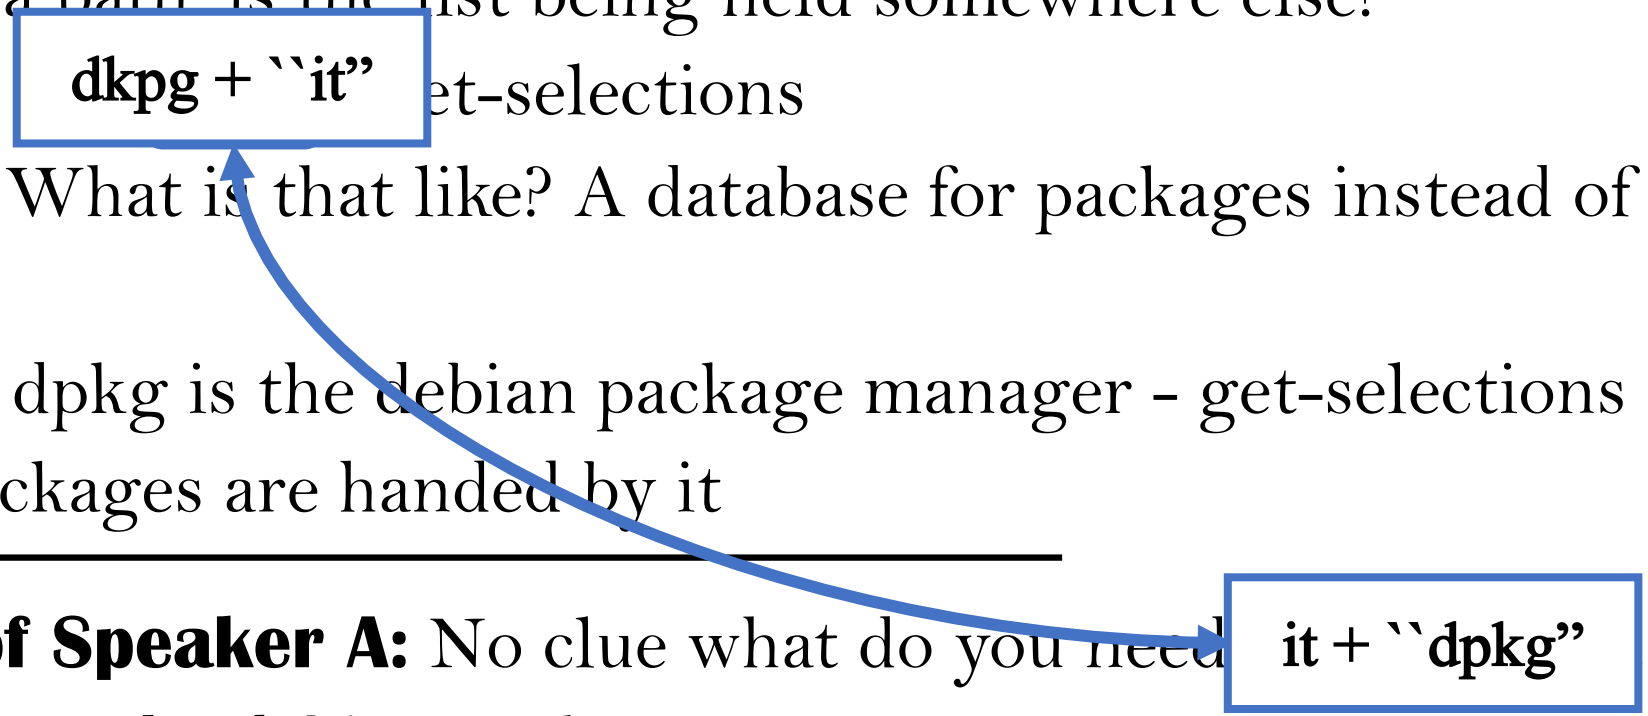

Conversation Context

Speaker A: Hi I am looking to see what packages are installed on my system, I don't see a path, is the list being held somewhere else?

Speaker B: dkpg + ``it" et-selections

Speaker A: What is that like? A database of packages instead of a flat file structure?

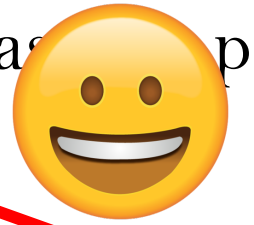

Speaker B: dpkg is the debian package manager - get-selections simply shows you what packages are handed by it

---

Response of Speaker A: No clue what do you need it + ``dpkg" reassurance as I don't know the debian package manager

Attentive Module

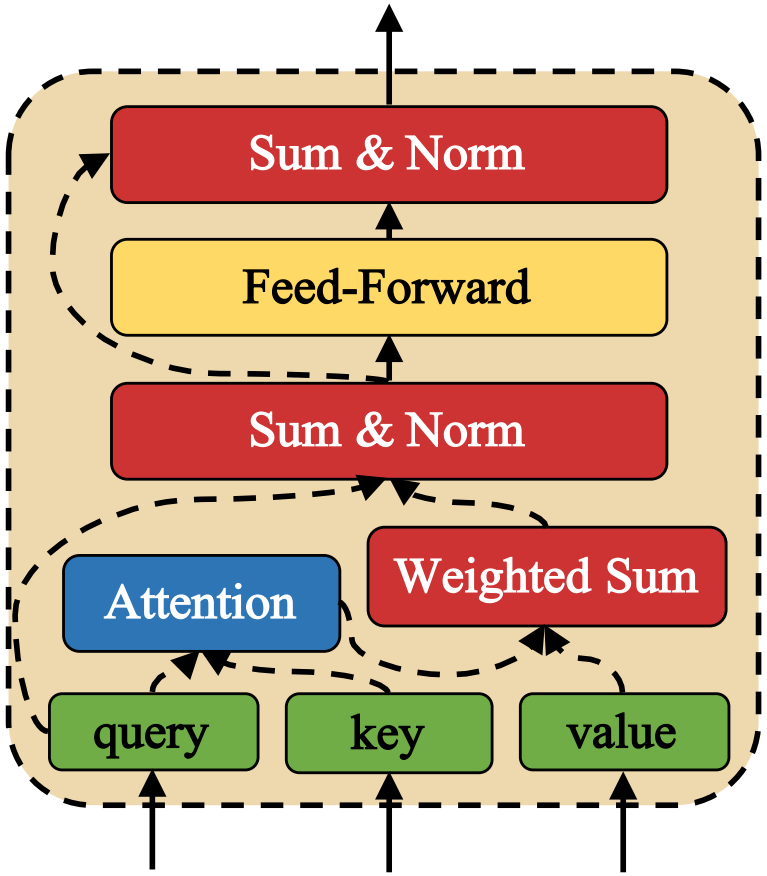

Attention is All You Need. Vaswani et al., NIPS-2017

Input

$$\mathcal{Q} = [e_0, ..., e_{n_\mathcal{Q}-1}] \quad \mathcal{K} = [e_0, ..., e_{n_\mathcal{K}-1}] \quad \mathcal{V} = [e_0, ..., e_{n_\mathcal{V}-1}]$$

AttentiveModule(Q,K,V)

- $\mathcal{V}_{att} = softmax(\frac{\mathcal{Q} \cdot \mathcal{K}^T}{\sqrt{d_k}})\mathcal{V}$
- $\mathcal{Q}_{att} = LayerNorm(\mathcal{V}_{att} + \mathcal{Q})$
- $FFN(\mathcal{Q}_{att})$
  $= max(0, \mathcal{Q}_{att}W_1 + b_1)W_2 + b_2$
- $LayerNorm(FFN(\mathcal{Q}_{att}) + \mathcal{Q}_{att})$

Summary

- Capture structures across Q and K-V
- Composite semantic representations of captured structures with input embedding

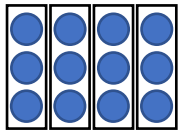

u_1

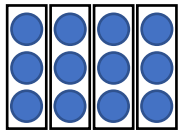

u_i

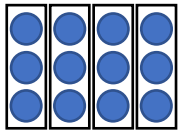

u_n

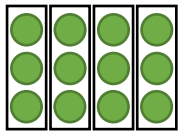

r

Word Embedding

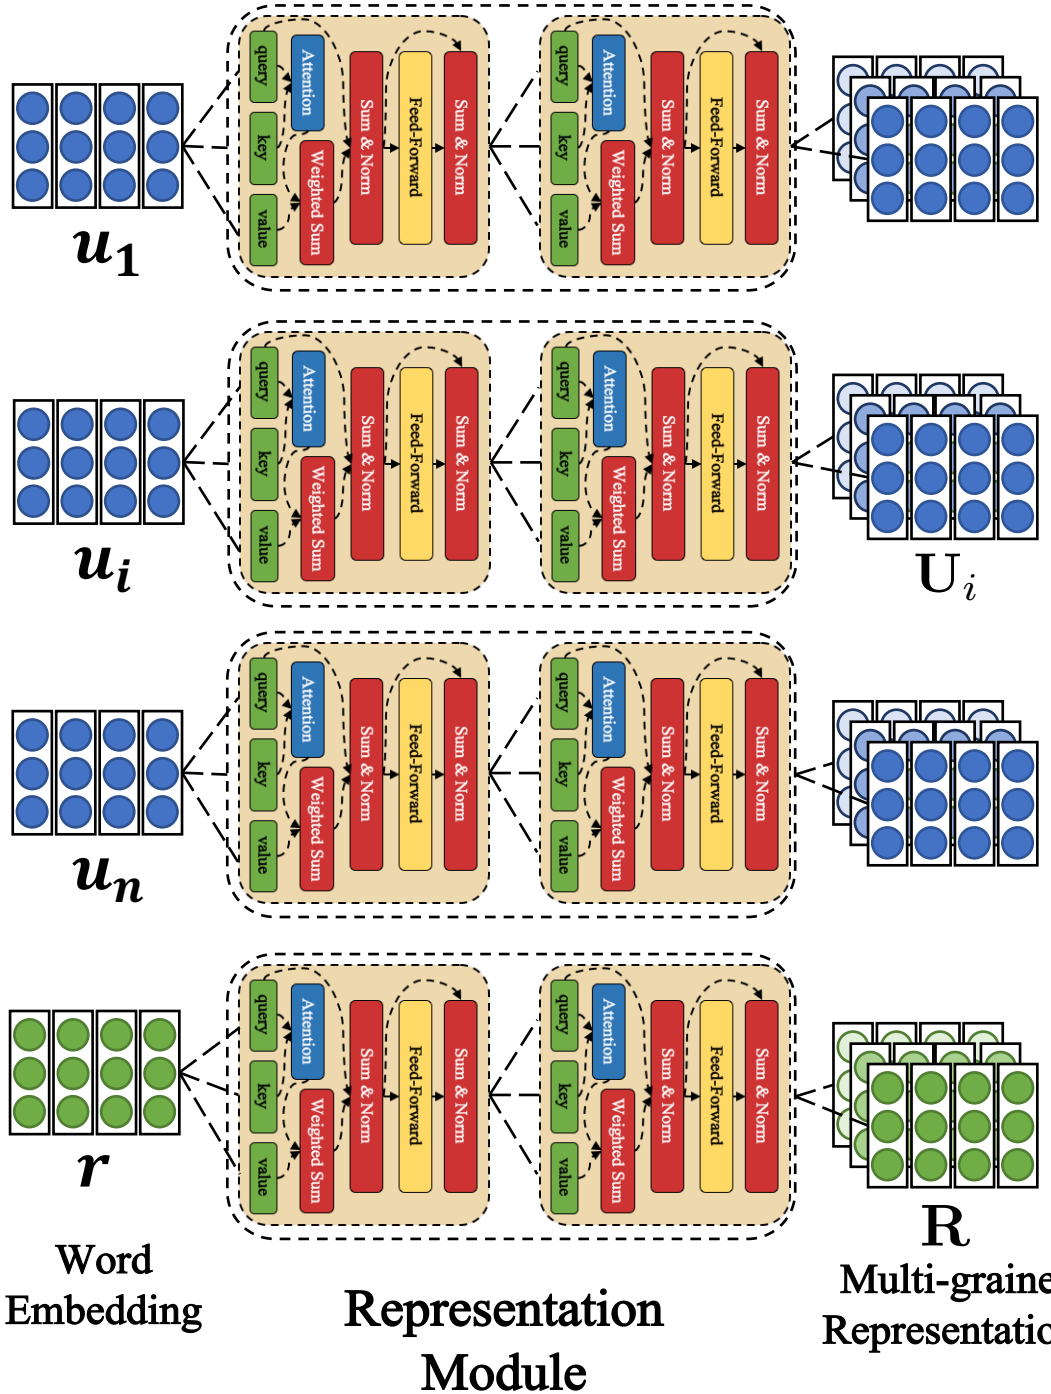

$$
\left.
\begin{aligned}
\mathbf{U}_i^0 &= \mathbf{u}_i, \\
&\dots \\
\mathbf{U}_i^l &= AttentiveModule(\mathbf{U}_i^{l-1}, \mathbf{U}_i^{l-1}, \mathbf{U}_i^{l-1}), \\
&\dots \\
\mathbf{U}_i^L &= AttentiveModule(\mathbf{U}_i^{L-1}, \mathbf{U}_i^{L-1}, \mathbf{U}_i^{L-1})
\end{aligned}
\right\}
$$

word-level

phrase-level

sent.-level

$$
\left\{
\begin{aligned}
\mathbf{R}^0 &= \mathbf{r}, \\
&\dots \\
\mathbf{R}^l &= AttentiveModule(\mathbf{R}^{l-1}, \mathbf{R}^{l-1}, \mathbf{R}^{l-1}), \\
&\dots \\
\mathbf{R}^L &= AttentiveModule(\mathbf{R}^{L-1}, \mathbf{R}^{L-1}, \mathbf{R}^{L-1})
\end{aligned}
\right\}
$$

Baidu 百度

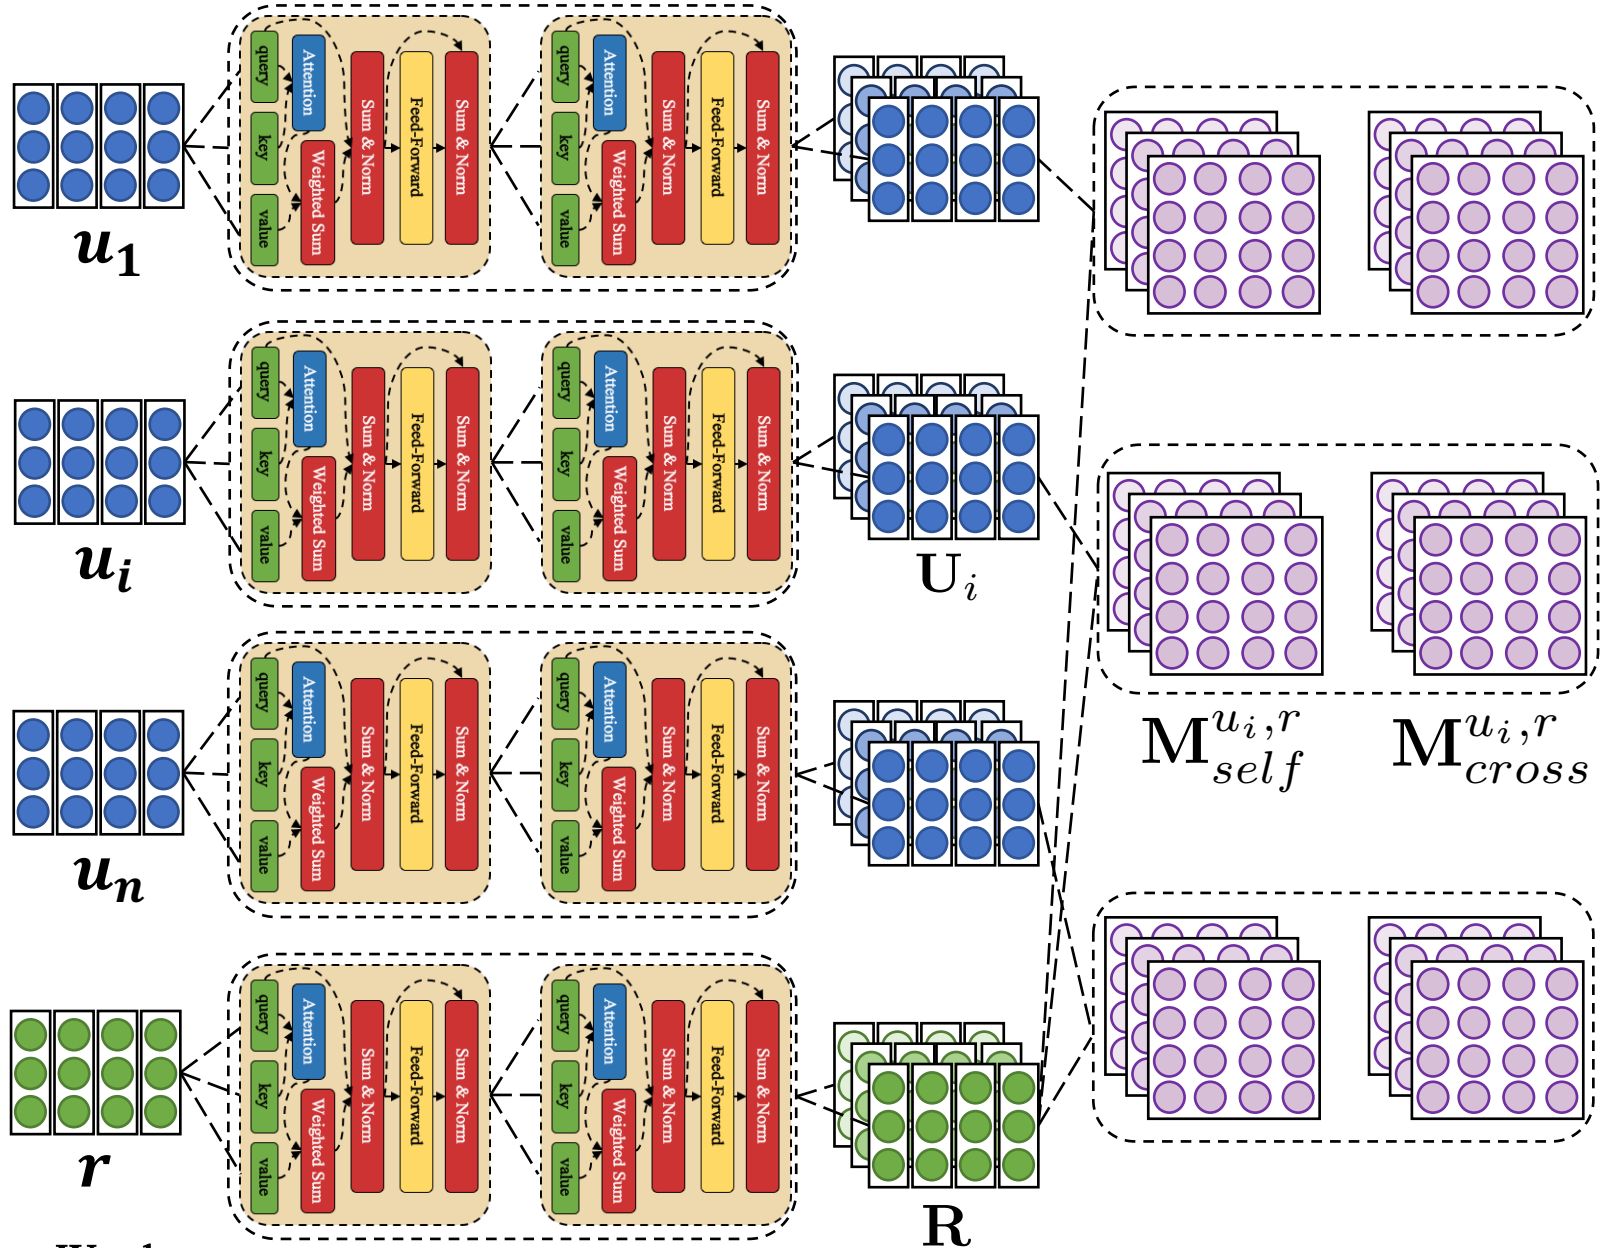
u_1
u_i
u_n
r
query
key
value
Attention
Weighted Sum
Sum & Norm
Feed-Forward
Sum & Norm
$\mathbf{U}_i$
$\mathbf{R}$
$\mathbf{M}_{self}^{u_i,r}$
$\mathbf{M}_{cross}^{u_i,r}$
Word Embedding
Representation Module
Multi-grained Representations
Word-word Matching with Cross-Attention

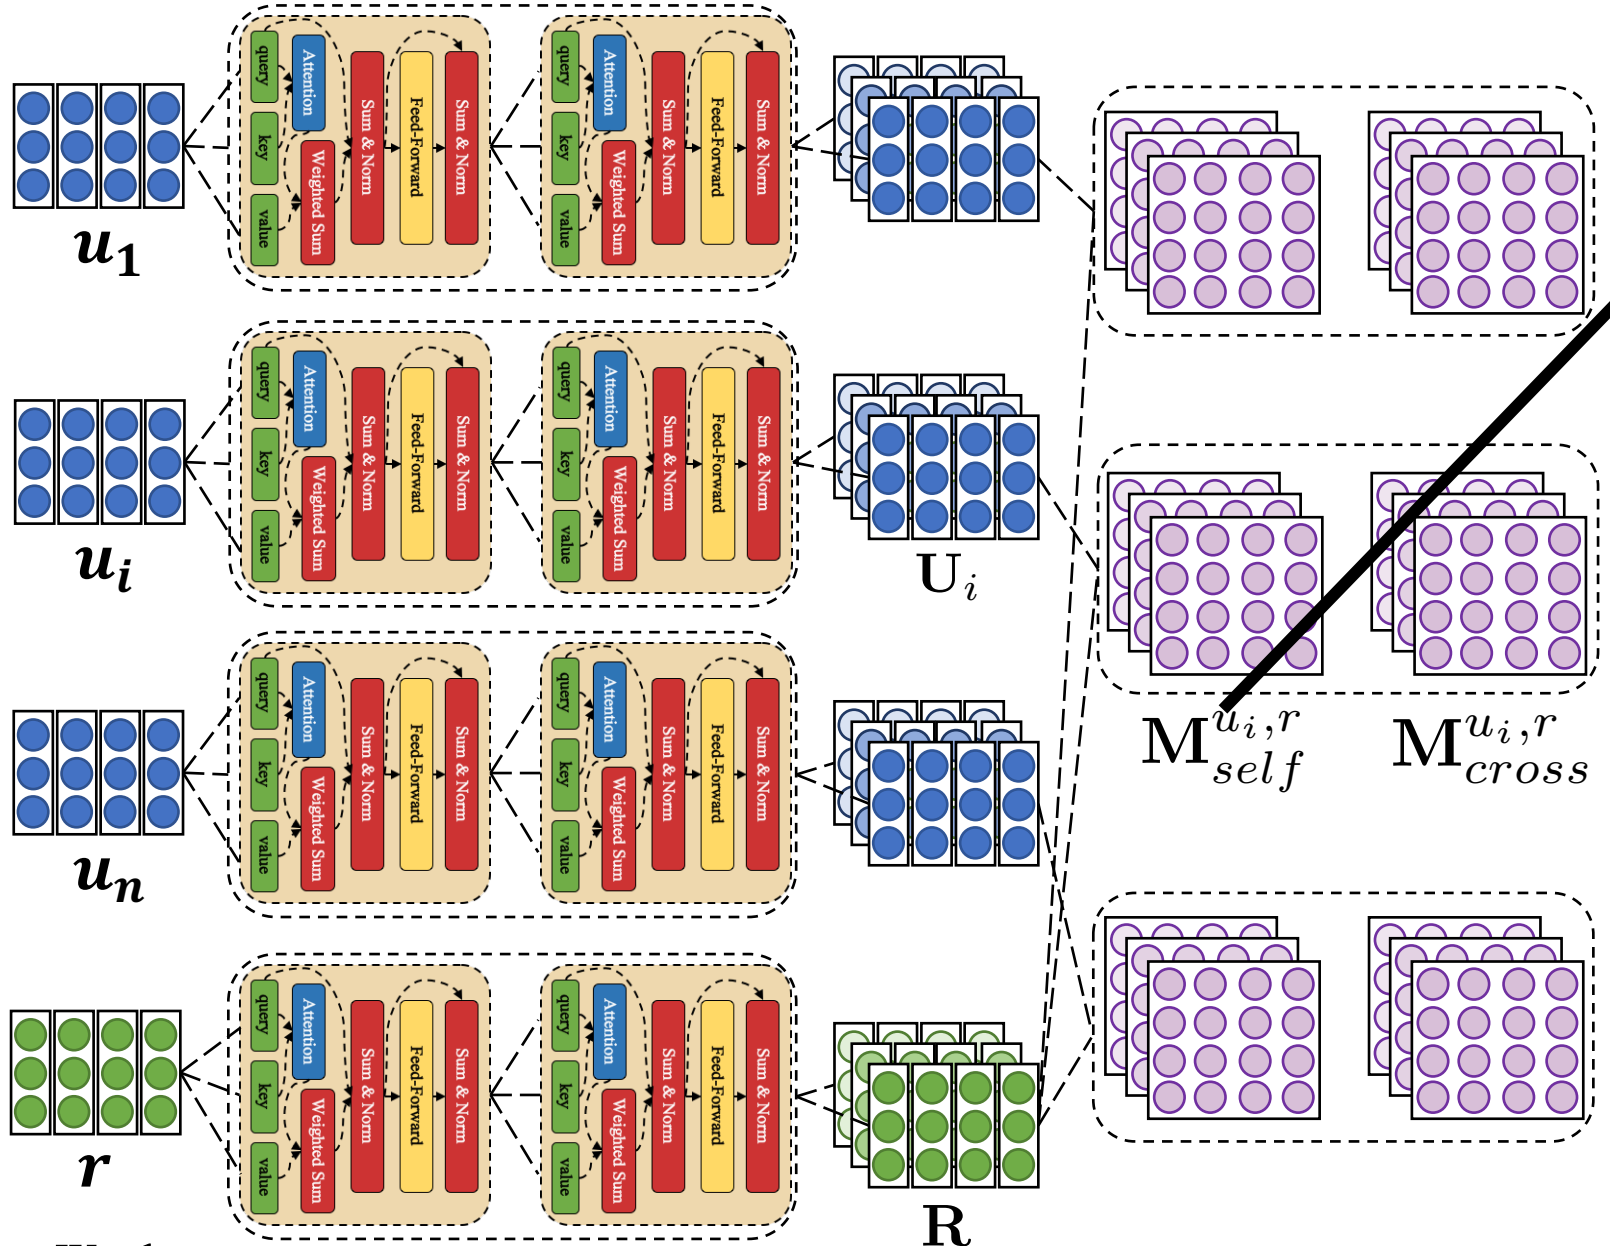

Self-Attention Match

$$\mathbf{M}_{self}^{u_i,r,l} = \{\mathbf{U}_i^l[k]^T \cdot \mathbf{R}^l[t]\}_{n_{u_i} \times n_r}$$

Match with textual information

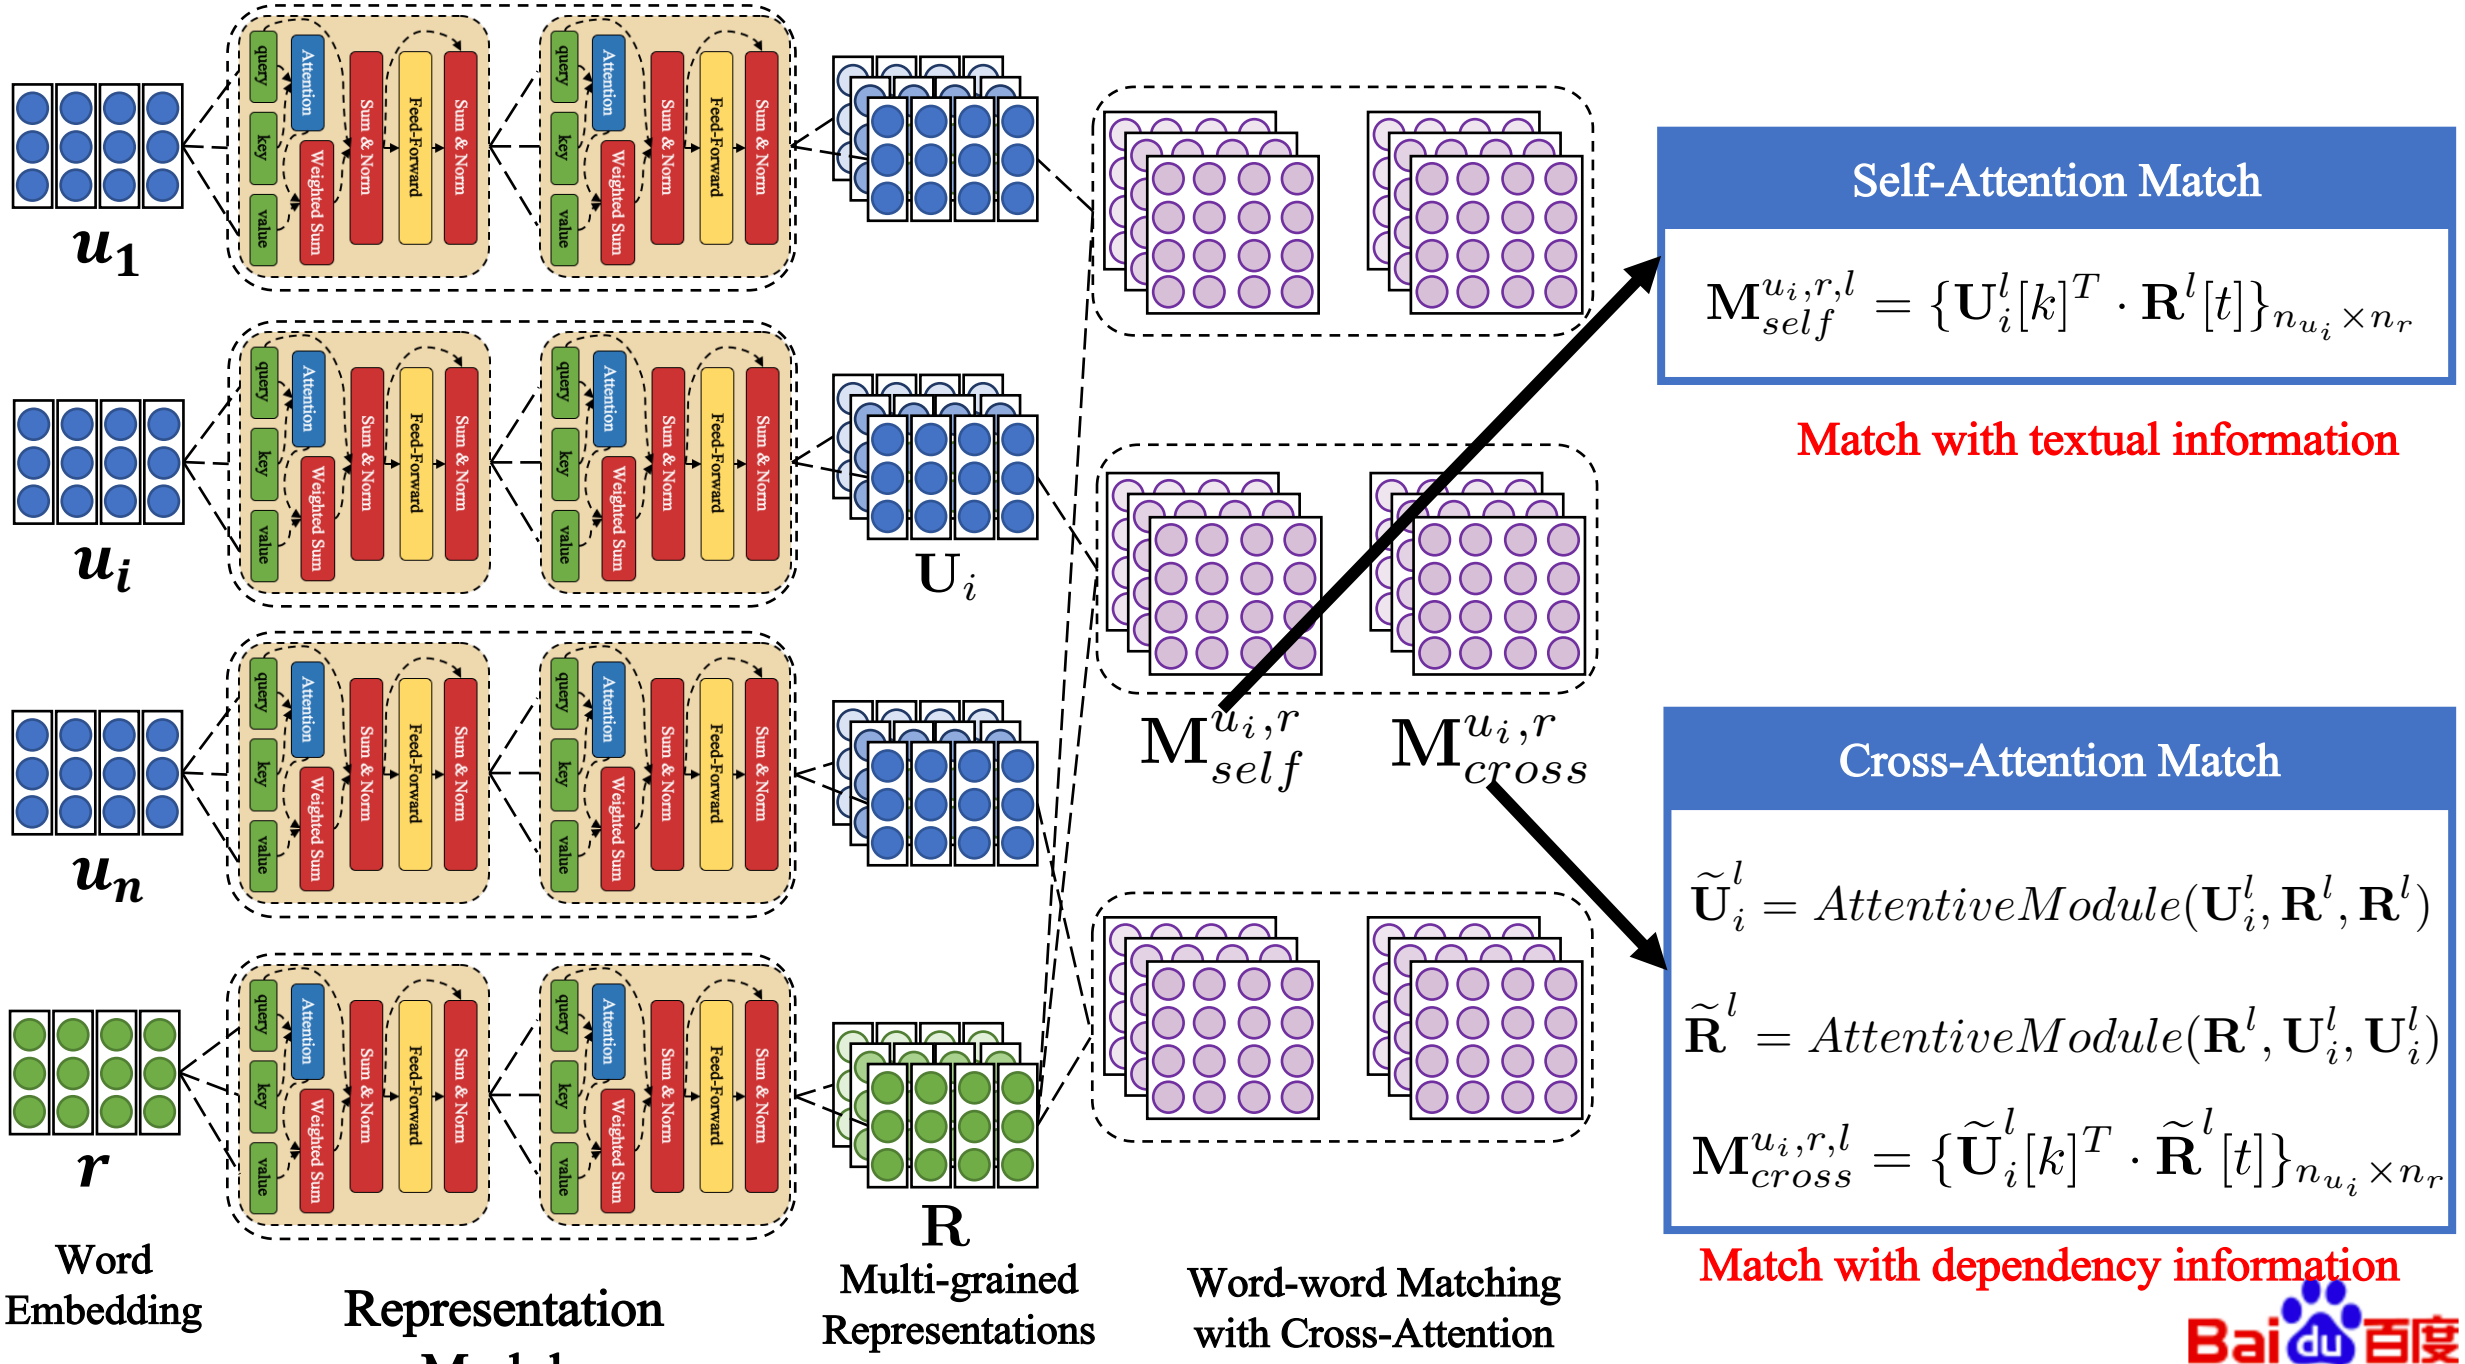
u_1
u_i
u_n
r
query
key
value
Attention
Weighted Sum
Sum & Norm
Feed-Forward
Sum & Norm
$\mathbf{U}_i$
$\mathbf{R}$
$\mathbf{M}_{self}^{u_i,r}$
$\mathbf{M}_{cross}^{u_i,r}$
Word Embedding
Representation Module
Multi-grained Representations
Word-word Matching with Cross-Attention
Self-Attention Match
$\mathbf{M}_{self}^{u_i,r,l} = \{\mathbf{U}_i^l[k]^T \cdot \mathbf{R}^l[t]\}_{n_{u_i} \times n_r}$
Match with textual information
Cross-Attention Match
$\widetilde{\mathbf{U}}_i^l = AttentiveModule(\mathbf{U}_i^l, \mathbf{R}^l, \mathbf{R}^l)$
$\widetilde{\mathbf{R}}^l = AttentiveModule(\mathbf{R}^l, \mathbf{U}_i^l, \mathbf{U}_i^l)$
$\mathbf{M}_{cross}^{u_i,r,l} = \{\widetilde{\mathbf{U}}_i^l[k]^T \cdot \widetilde{\mathbf{R}}^l[t]\}_{n_{u_i} \times n_r}$
Match with dependency information

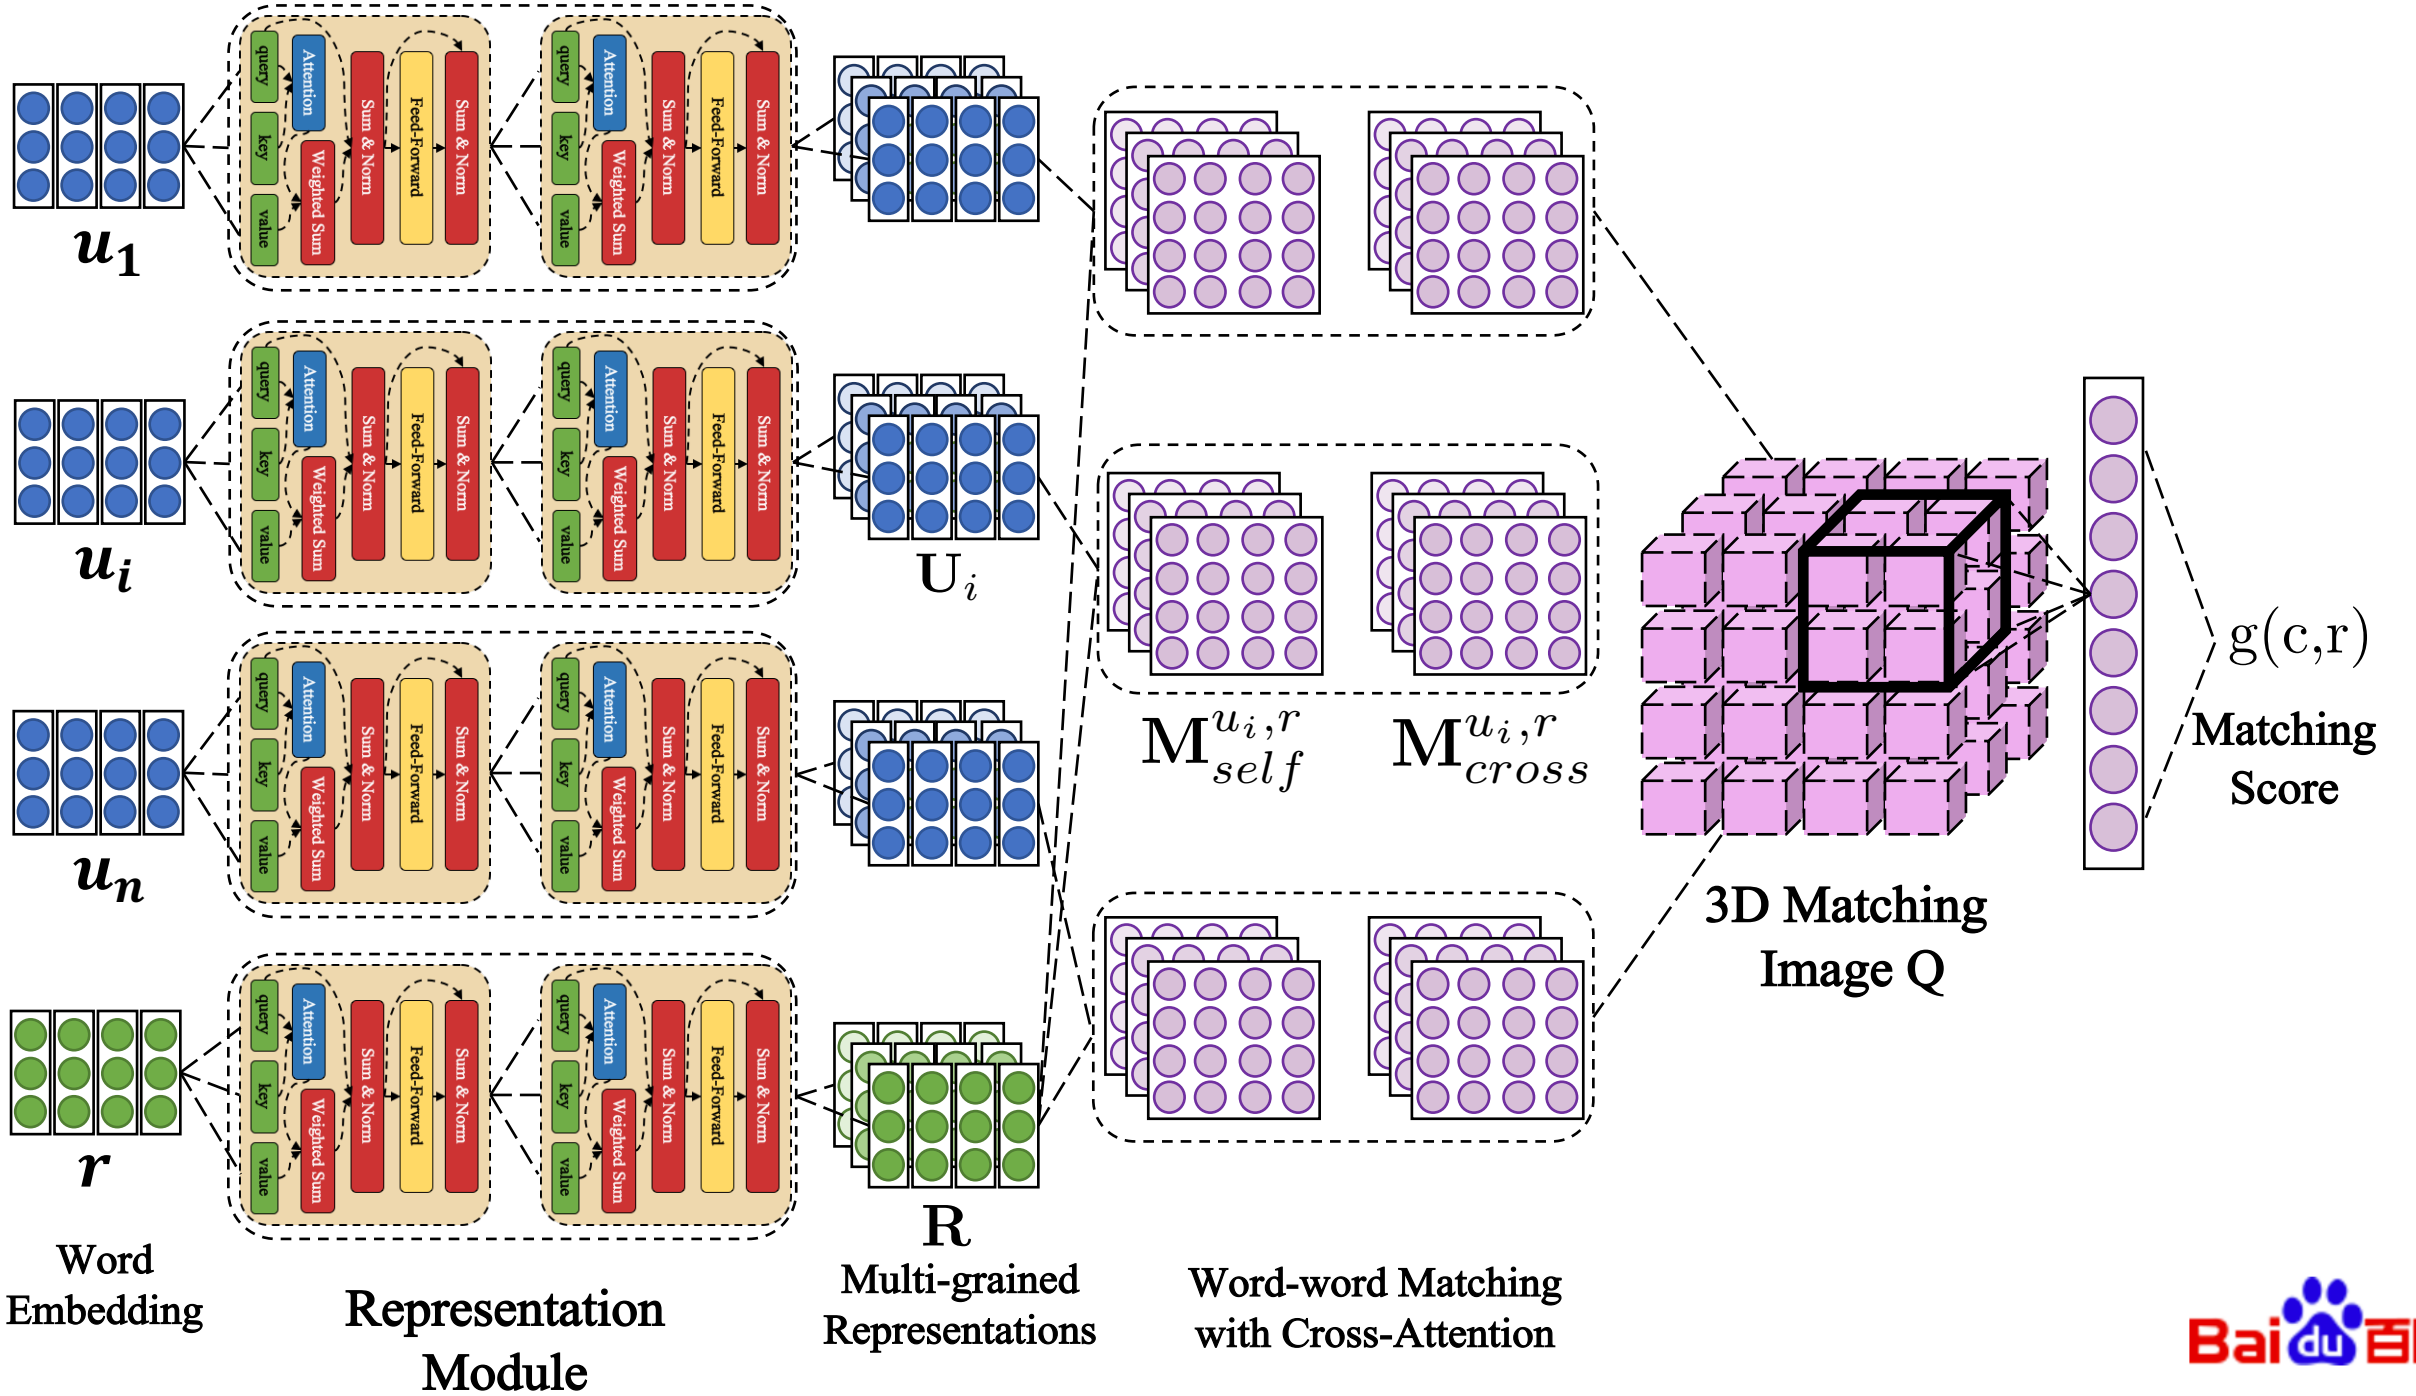
u_1
u_i
u_n
r
query
key
value
Attention
Weighted Sum
Sum & Norm
Feed-Forward
Sum & Norm
U_i
R
$M_{self}^{u_i,r}$
$M_{cross}^{u_i,r}$
3D Matching Image Q
g(c,r)
Matching Score
Word Embedding
Representation Module
Multi-grained Representations
Word-word Matching with Cross-Attention

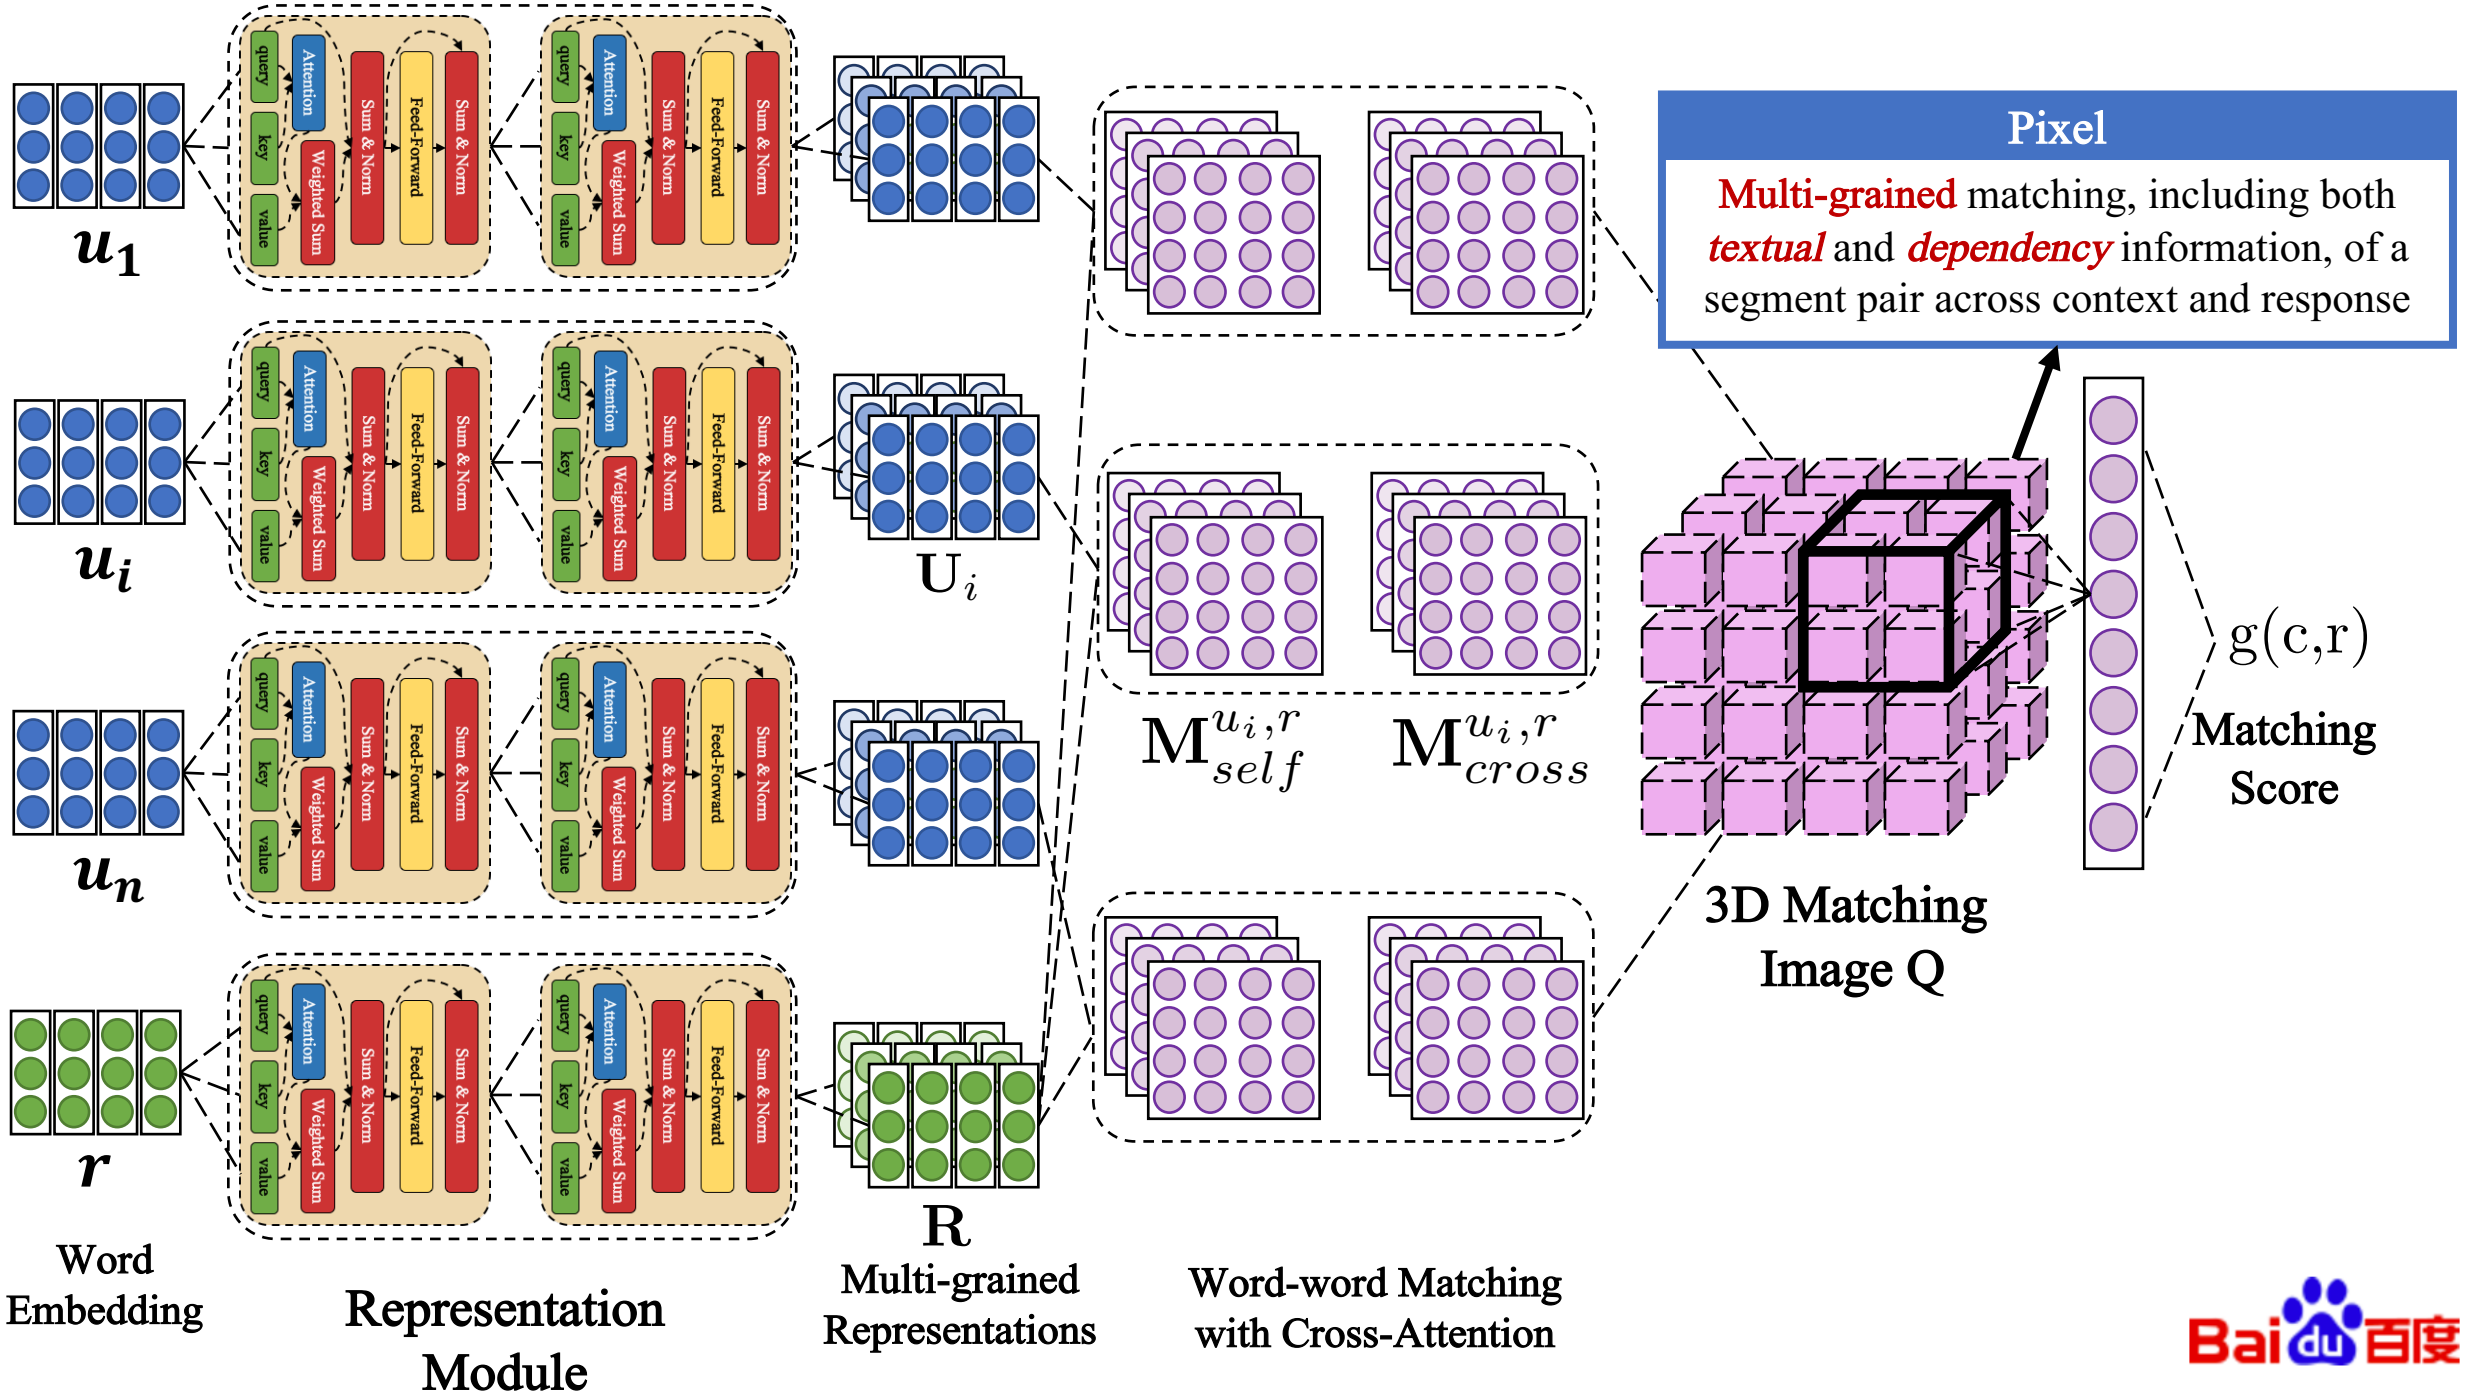
u_1
u_i
u_n
r
query
key
value
Attention
Weighted Sum
Sum & Norm
Feed-Forward
Sum & Norm
U_i
R
$M_{self}^{u_i,r}$
$M_{cross}^{u_i,r}$
Pixel
Multi-grained matching, including both *textual* and *dependency* information, of a segment pair across context and response
3D Matching Image Q
g(c,r)
Matching Score
Word Embedding
Representation Module
Multi-grained Representations
Word-word Matching with Cross-Attention

Experiment

- Ubuntu Corpus
  - One-one multi-turn conversation
  - Ubuntu troubleshooting
- Douban Corpus
  - One-one multi-turn conversation
  - Open domain topics
- Task
  - Given multi-turn context and serval response candidates
  - Select the best candidate based on matching score

	Ubuntu Corpus V1			Douban Conversation Corpus		
	Train	Dev	Test	Train	Dev	Test
# context-response-pairs	1M	50k	50k	1M	50k	10k
# candidates per context	2	10	10	2	2	10
# positive candidates per context	1			1	1	1.18
Min. # turns per context	3			3	3	3
Max. # turns per context	19			98	91	45
Avg. # turns per context	7.71			6.69	6.75	6.45
Avg. #words per utterance	10.34			18.56	18.50	20.74

Experiment

- DAM setup
  - Test stacking 3-7 self-attention layers

$$L(\cdot) = -\sum_{(c,r,y)\in\mathcal{D}} g(c,r)y + (1-g(c,r))(1-y)$$

- Comparison
  - ***Sequential Matching Network (SMN)*** (Wu et al., ACL-2017), ***Multi-view Matching*** (Zhou et al., EMNLP-2016), ***DL2R*** (Yan et al., SIGIR-2016), ***DualEncoder*** (Rowe et al., SigDial-2015)
- Ablation
  - DAM_{first} : without stacked self-attention
  - DAM_{last} : only using the last layer of stacked self-attention
  - DAM_{self} : only using self-attention-match
  - DAM_{cross} : only using cross-attention-match

Baidu 百度

Evaluation Results

	Ubuntu Corpus				Douban Conversation Corpus					
	$\mathbf{R_2@1}$	$\mathbf{R_{10}@1}$	$\mathbf{R_{10}@2}$	$\mathbf{R_{10}@5}$	MAP	MRR	P@1	$\mathbf{R_{10}@1}$	$\mathbf{R_{10}@2}$	$\mathbf{R_{10}@5}$
DualEncoder$_{lstm}$	0.901	0.638	0.784	0.949	0.485	0.527	0.320	0.187	0.343	0.720
DualEncoder$_{bilstm}$	0.895	0.630	0.780	0.944	0.479	0.514	0.313	0.184	0.330	0.716
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
Multiview	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
SMN$_{dynamic}$	0.926	0.726	0.847	0.961	0.529	0.569	0.397	0.233	0.396	0.724
DAM	**0.938**	**0.767**	**0.874**	**0.969**	**0.550**	**0.601**	**0.427**	**0.254**	**0.410**	**0.757**
DAM$_{first}$	0.927	0.736	0.854	0.962	0.528	0.579	0.400	0.229	0.396	0.741
DAM$_{last}$	0.932	0.752	0.861	0.965	0.539	0.583	0.408	0.242	0.407	0.748
DAM$_{self}$	0.931	0.741	0.859	0.964	0.527	0.574	0.382	0.221	0.403	0.750
DAM$_{cross}$	0.932	0.749	0.863	0.966	0.535	0.585	0.400	0.234	0.411	0.733

Baidu 百度

Self-Attention Match Visualization

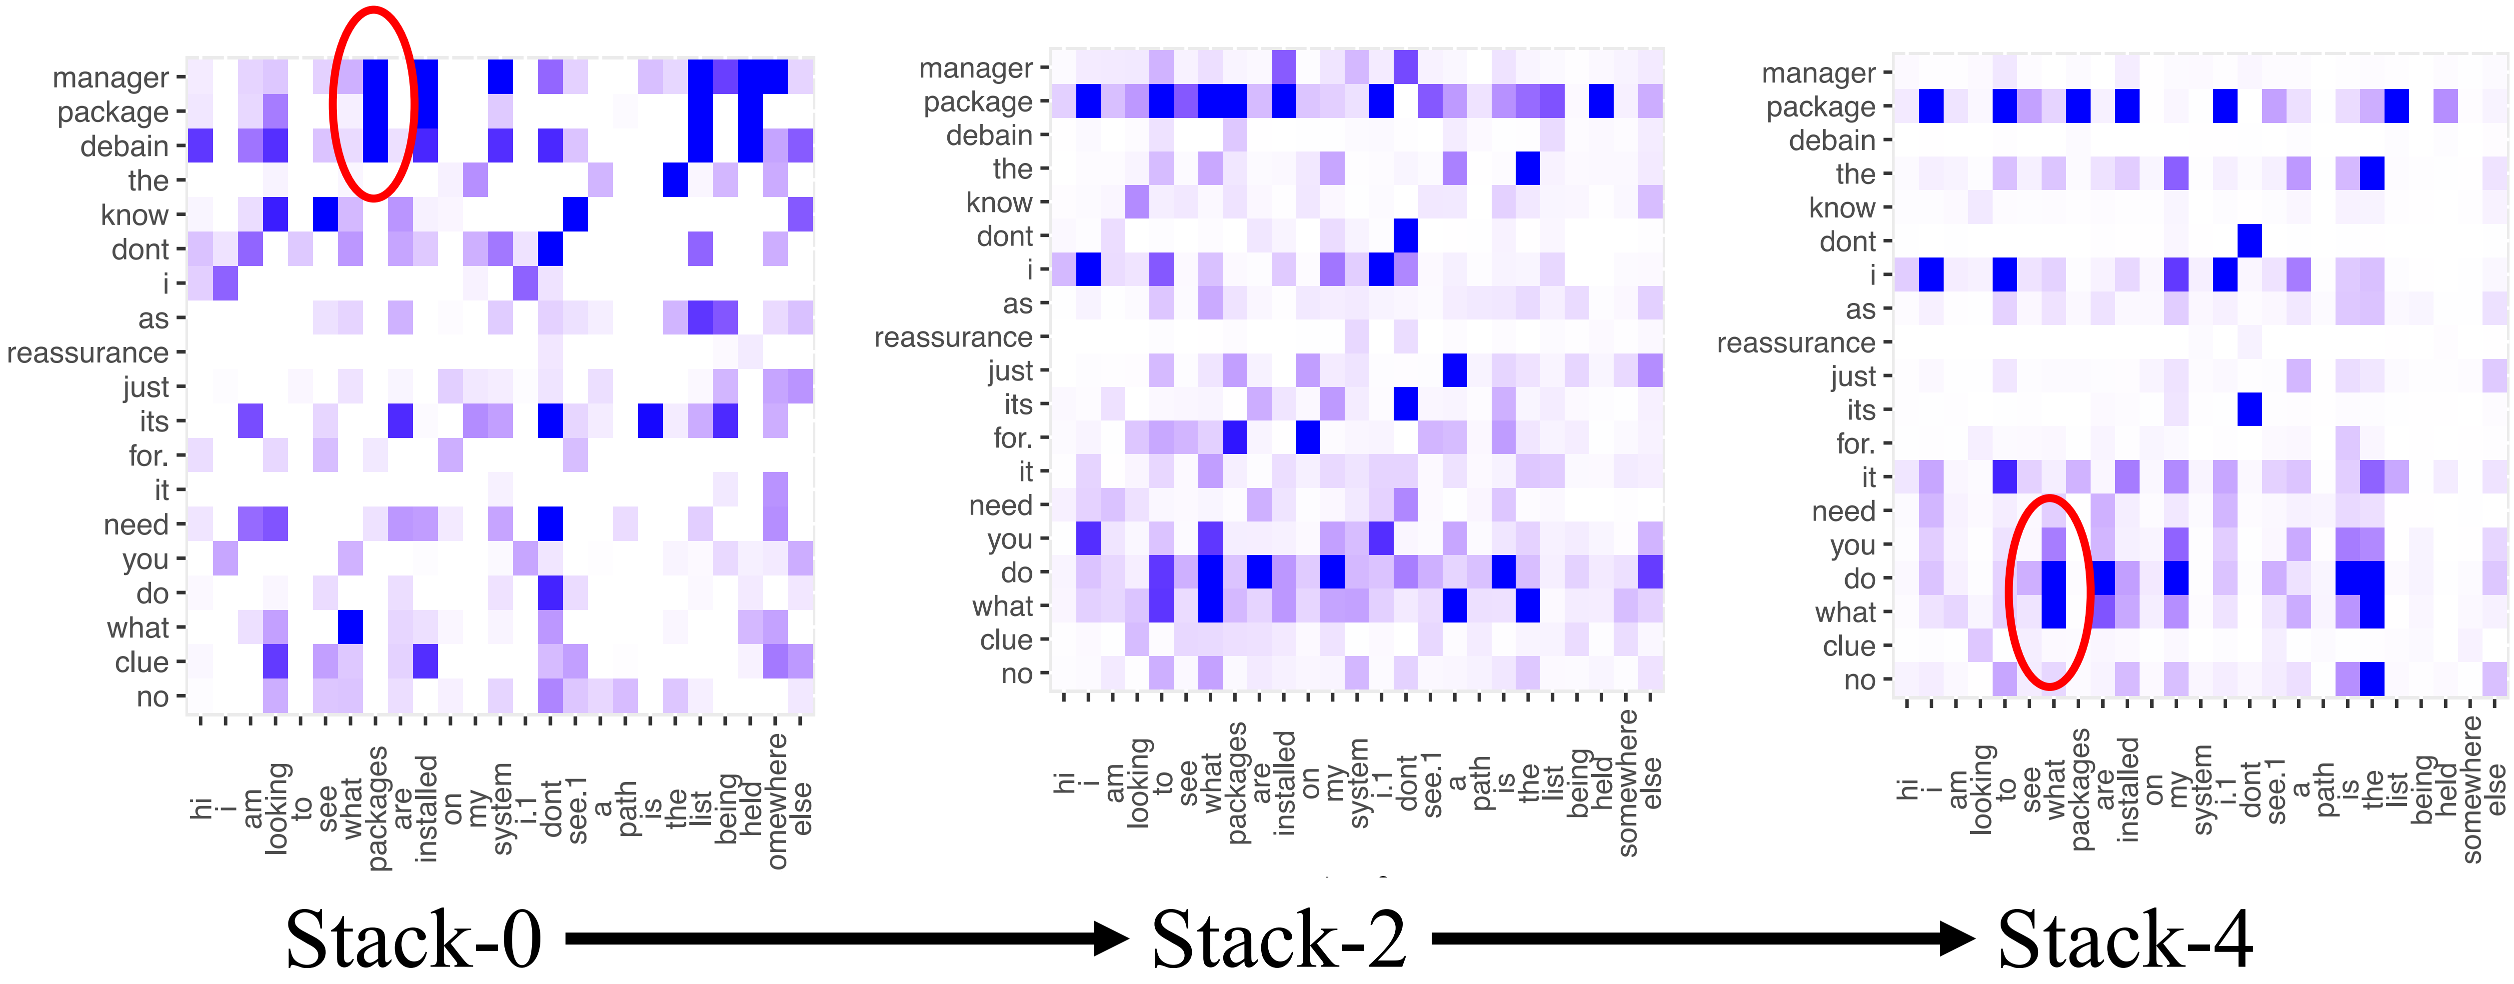

Cross-Attention Match Visualization

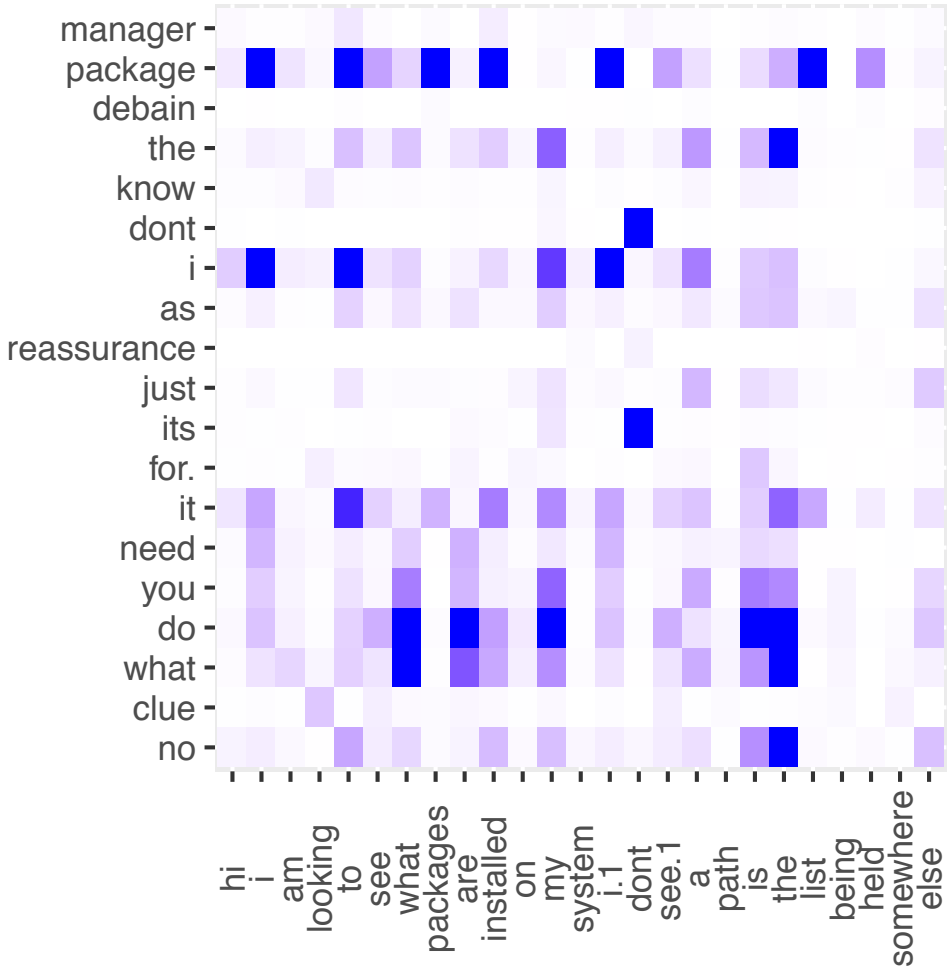

Self-Attention Match

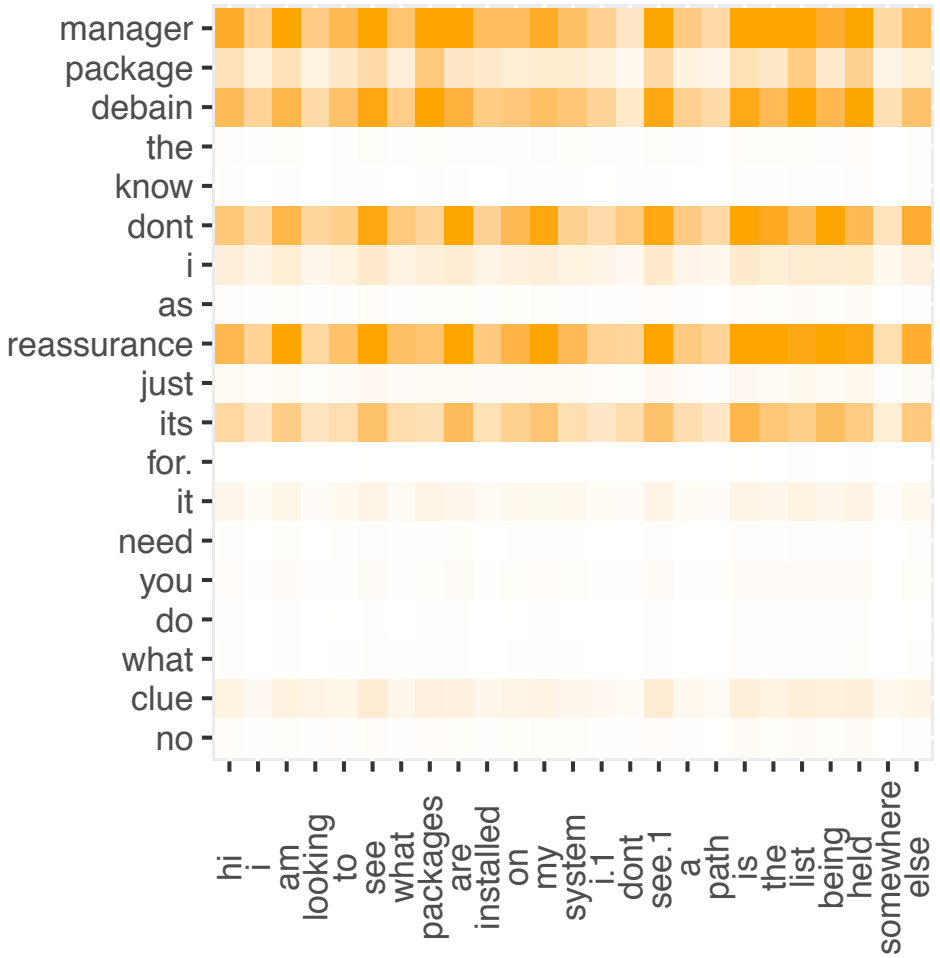

Cross-Attention Match

Summary

- We propose a novel deep attention matching network for multi-turn response selection that entirely based on attention.
- We use stacked self-attention to construct multi-grained semantic representations.
- We use cross-attention to match context with its candidate response considering both textual and dependency information

Thanks

- Code & Data: https://github.com/baidu/Dialogue/DAM
- Our Authors