Global-to-local Memory Pointer Networks for Task-Oriented Dialogue

Chien-Sheng Wu*, Caiming Xiong, Richard Socher
Salesforce Research
*Hong Kong University of Science & Technology

2018.12.07

Conversational Agents

- **Chit-Chat Dialogue Systems**
  - No specific goal, focus on conversation flow (engagement, fluency, consistent, etc) and expect more conversational turns
  - Work using variants of seq2seq model:
    - Seq2Seq models →
    - Seq2Seq + conversational context →
    - Knowledge-grounded Seq2Seq models
- **Task-Oriented Dialogue Systems**
  - Personal assistant, achieve a certain task and expect less conversational turns
  - Often combined rules and statistical components
    - Single domain, pipeline approaches →
    - Multi-domain, contextual, pipeline/end-to-end approaches →
    - Massively multi-domain, end-to-end approaches

Source: https://sites.google.com/view/deepdial/

Task-oriented Dialogue Systems

- **Description**:
  - Achieve specific user goals within a limited dialogue turns via natural language.
- **Challenges**:
  - language understanding (LU), dialogue management (DM), knowledge base (KB) understanding, language generation (LG), etc.

Point of interest (poi)	Distance	Traffic info	Poi type	Address
The Westin	5 miles	moderate traffic	rest stop	329 El Camino Real
Round Table	4 miles	no traffic	pizza restaurant	113 Anton Ct
Mandarin Roots	5 miles	no traffic	chinese restaurant	271 Springer Street
Palo Alto Cafe	4 miles	moderate traffic	coffee or tea place	436 Alger Dr
Dominos	6 miles	heavy traffic	pizza restaurant	776 Arastradero Rd
Stanford Express Care	6 miles	no traffic	hospital	214 El Camino Real
Hotel Keen	2 miles	heavy traffic	rest stop	578 Arbol Dr

End-to-End Approaches: Advantages & Challenges

- **Input**:
  - Dialogue history and knowledge base
- **Output**:
  - System response with real slot values
- **Advantages**:
  - No labels of belief states, slots, dialogue actions, intention, etc.
  - Free from learning dependency between modules
- **Challenges**:
  - How to incorporate large, dynamic KB into learning frameworks?
    - Entity selection, KB reasoning, etc.
  - How to interpret dialogue systems?
    - Belief states, slot-filling, etc
  - How to overcome rare data issue?

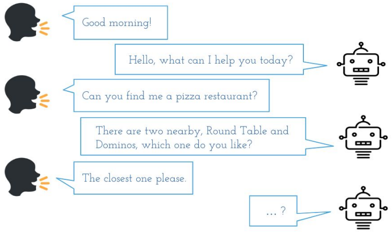

Point of interest (poi)	Distance	Traffic info	Poi type	Address
The Westin	5 miles	moderate traffic	rest stop	329 El Camino Real
Round Table	4 miles	no traffic	pizza restaurant	113 Anton Ct
Mandarin Roots	5 miles	no traffic	chinese restaurant	271 Springer Street
Palo Alto Cafe	4 miles	moderate traffic	coffee or tea place	436 Alger Dr
Dominos	6 miles	heavy traffic	pizza restaurant	776 Arastradero Rd
Stanford Express Care	6 miles	no traffic	hospital	214 El Camino Real
Hotel Keen	2 miles	heavy traffic	rest stop	578 Arbol Dr

Global-to-local Memory Pointer Networks (GLMP): Block Diagram

Point of interest (poi)	Distance	Traffic info	Poi type	Address
The Westin	5 miles	moderate traffic	rest stop	329 El Camino Real
Round Table	4 miles	no traffic	pizza restaurant	113 Anton Ct
Mandarin Roots	5 miles	no traffic	chinese restaurant	271 Springer Street
Palo Alto Cafe	4 miles	moderate traffic	coffee or tea place	436 Alger Dr
Dominos	6 miles	heavy traffic	pizza restaurant	776 Arastradero Rd
Stanford Express Care	6 miles	no traffic	hospital	214 El Camino Real
Hotel Keen	2 miles	heavy traffic	rest stop	578 Arbol Dr

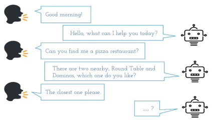

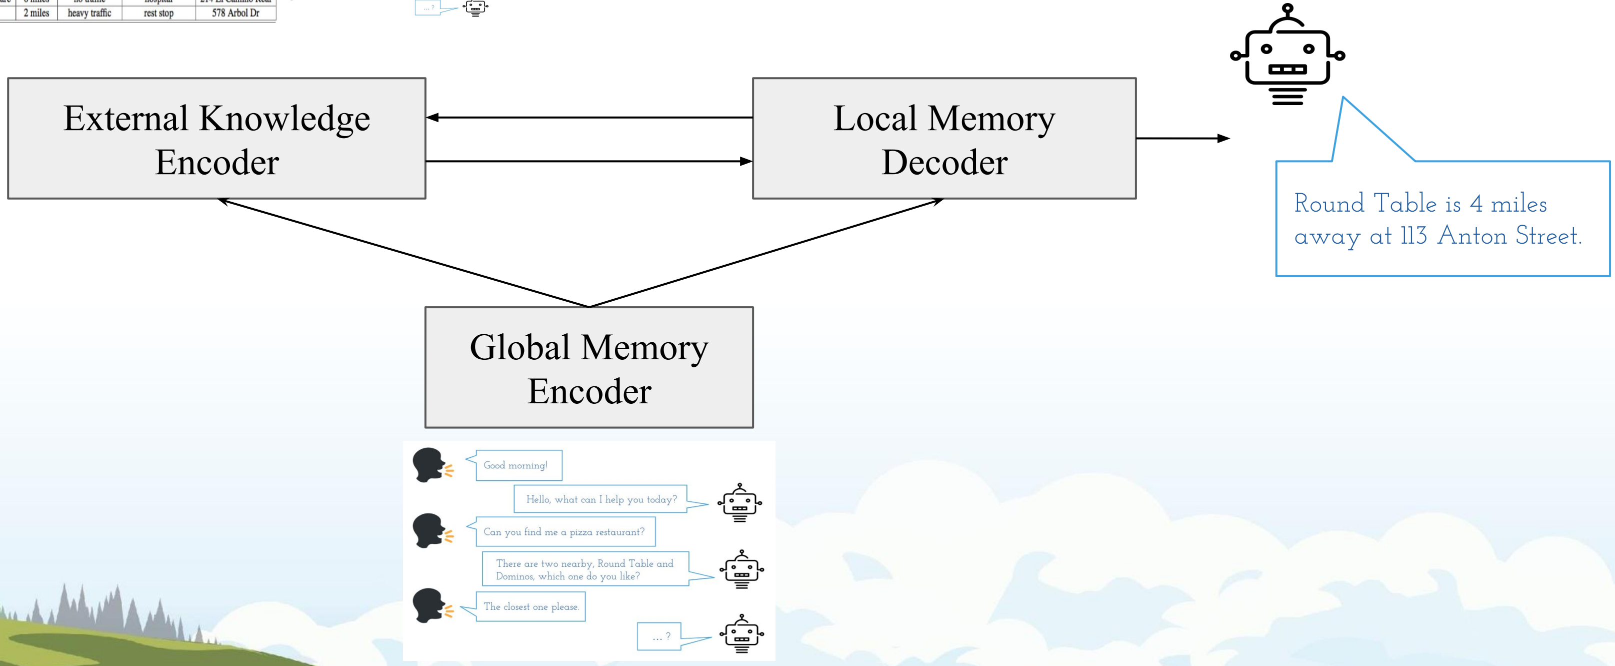

GLMP: External Knowledge

- **End-to-end Memory Networks**
  - A query vector
  - A set of trainable embeddings
  - Memory attention weights
  - Multiple hops reasoning

$$p_i^k = \text{Softmax}((q^k)^T c_i^k),$$

$$o^k = \sum_i p_i^k c_i^{k+1}, \quad q^{k+1} = q^k + o^k.$$

- **KB memory & Dialogue memory**
  - (Subject, Predicate, Object)
  - Copy Object word

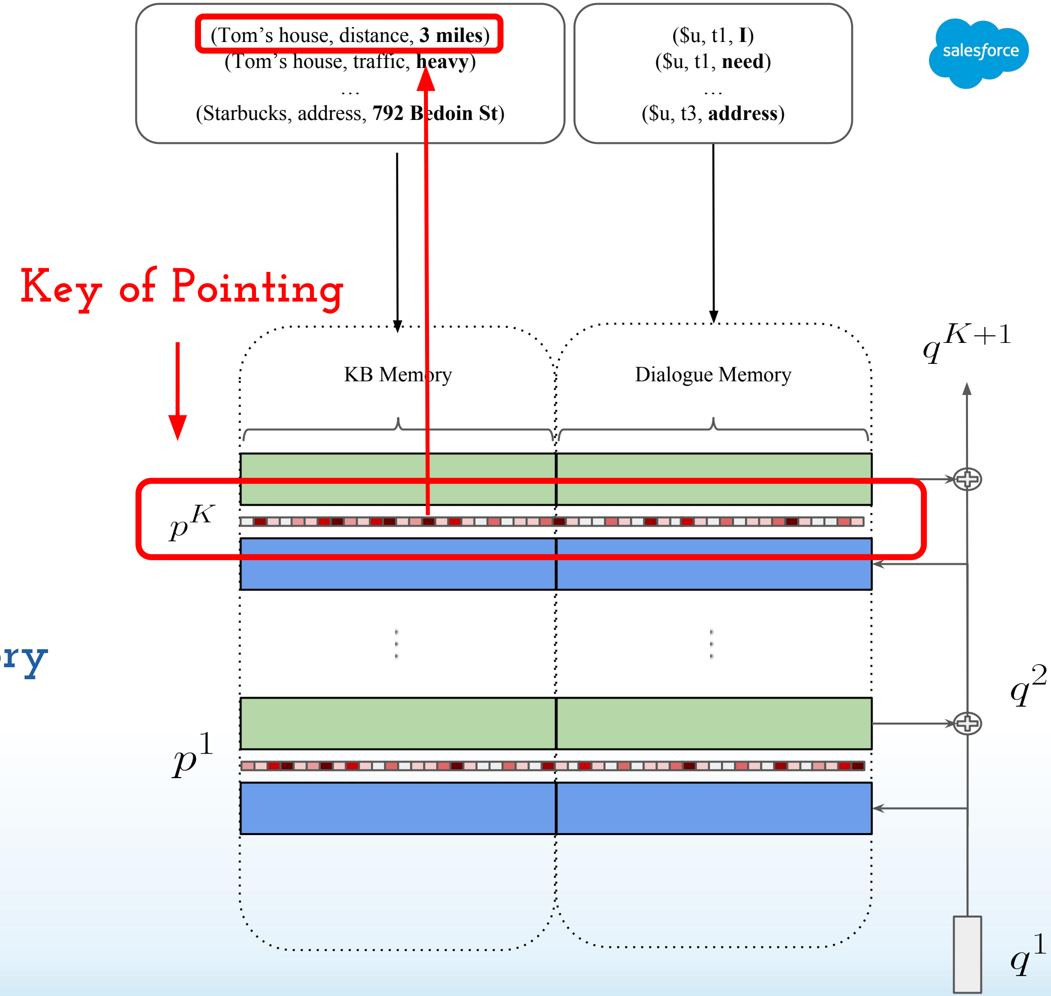

GLMP: Global Memory Encoder

- **Context RNN**
  - Encode plain text dialogue history
  - Query external knowledge
- **Contextual Dialogue History**
  - Write hidden states into dialogue memory module
  - Mitigate OOV copying problem
- **Global Memory Pointer**
  - Point to all the words that may appear in the system response.
  - Multi-label classification

$$g_i = \text{Sigmoid}((q^K)^T c_i^K)$$

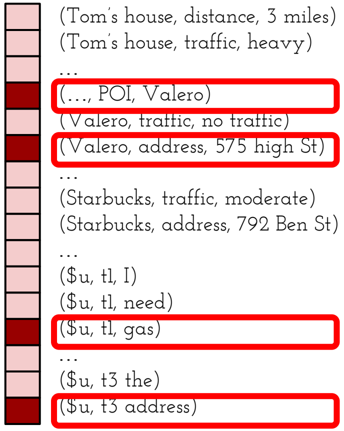

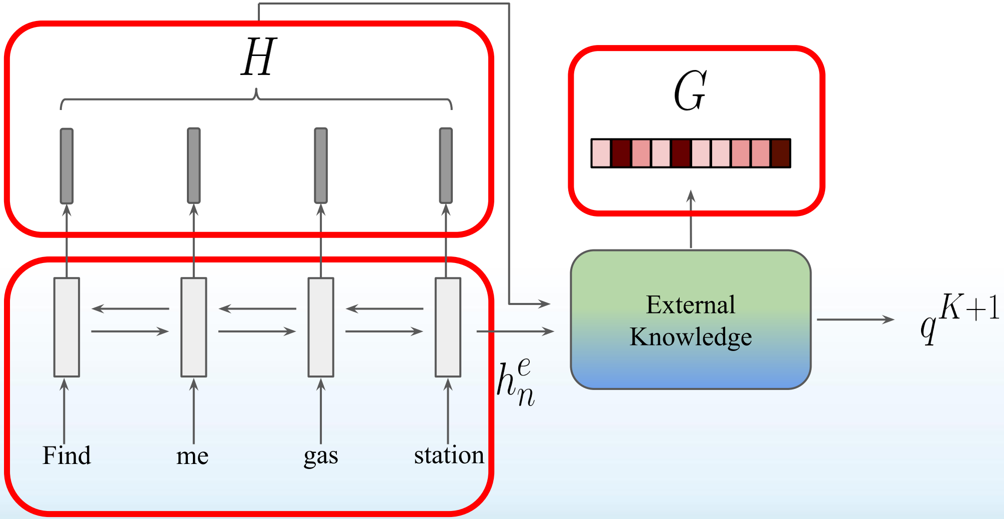

GLMP: Local Memory Decoder

$$h_t^d = \text{GRU}(C^1(\hat{y}_{t-1}^s), h_{t-1}^d), \quad P_t^{vocab} = \text{Softmax}(W h_t^d)$$

$$p_i^k = \text{Softmax}((q^k)^T c_i^k),$$

$$\hat{y}_t = \begin{cases} \arg\max(P_t^{vocab}) & \text{if } \arg\max(P_t^{vocab}) \notin ST, \\ Object(m_{\arg\max(L_t \odot R)}) & \text{otherwise,} \end{cases}$$

- **Sketch RNN**
  - Initialize with encoded dialogue history and KB
  - Generate sketch response. Ex: @poi is @distance away.
  - Query external knowledge using its hidden states
- **Local Memory Pointer**
  - Filter external knowledge using global memory pointer
  - Copy one single word at each time step
- **Record Function**
  - Mask the copied words

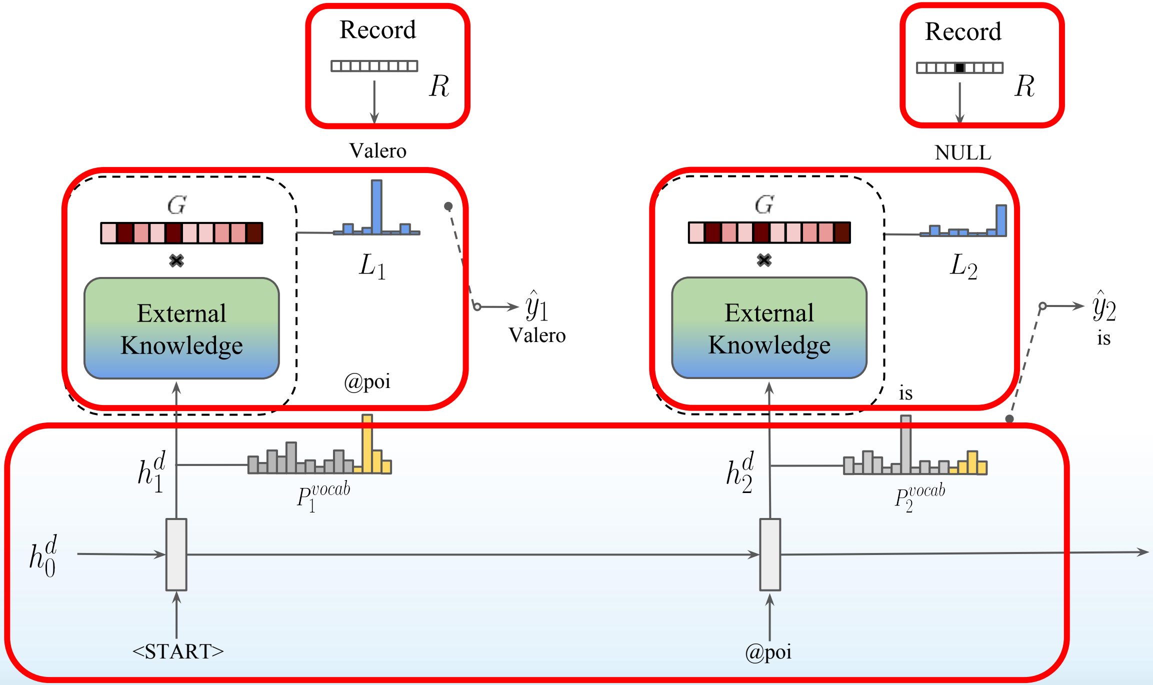

GLMP: Workflow

System Response: Valero is 3 miles away

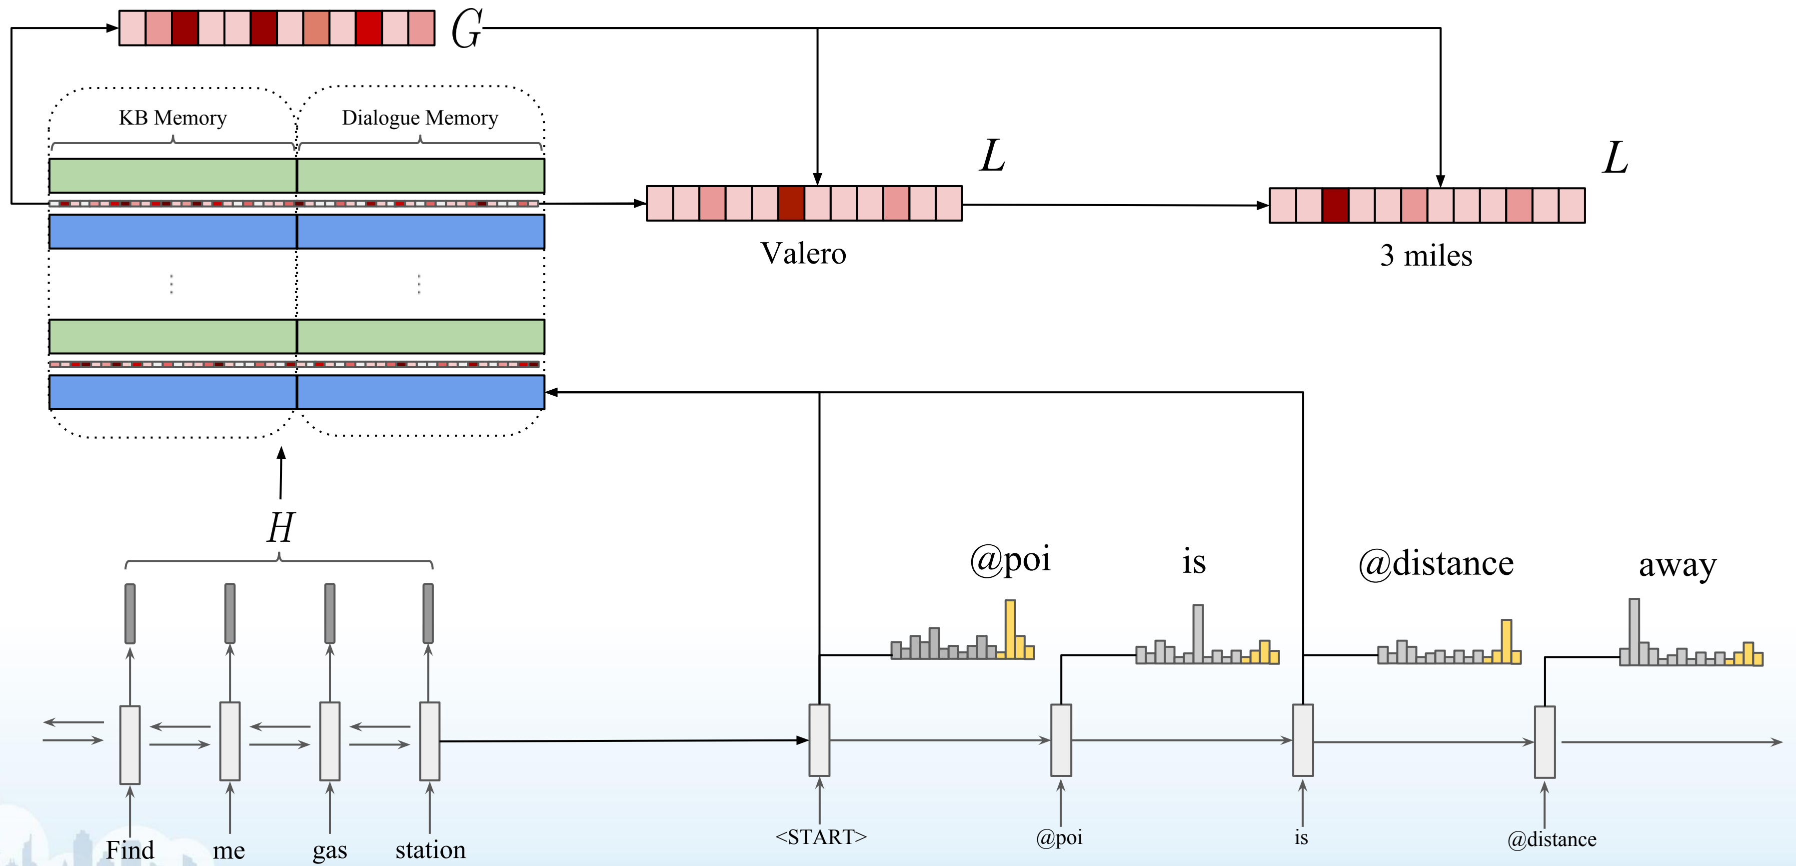

Experiment Setup

$$Loss_g = -\sum_{i=1}^{n+l}[g_i^l \times \log g_i + (1 - g_i^l) \times \log(1 - g_i)]$$

$$Loss_v = \sum_{t=1}^{m} -\log(P_t^{vocab}(y_t^s))$$

$$Loss_l = \sum_{t=1}^{m} -\log(L_t(L_t^{label})).$$

- **Datasets**:
  - **bAbI dialogue**:
    - Simulated dialogues on restaurant domain.
    - Include out-of-vocabulary setting for slot values.
    - Evaluation metrics: per-response accuracy, per-dialogue accuracy
  - **Stanford multi-domain** (SMD):
    - Human-human dialogue of a car assistant.
    - Three domains: calendar scheduling, weather retrieval, and point-of-interest navigation.
    - Evaluation metrics: BLEU, Entity F1, human evaluation
- **Objective functions**:
  - $Loss = \alpha Loss_g + \beta Loss_v + \gamma Loss_l$
  - Global memory pointer ($Loss\_g$): binary cross-entropy loss
  - Sketch RNN ($Loss\_v$): cross-entropy loss
  - Local memory pointer ($Loss\_l$): cross-entropy loss
- **Others**: Simple greedy decoding. Hyperparameter grid search over hidden size, number of hops, and dropout ratio. Adam optimizer. Without pre-trained embedding.

Baselines

- End-to-end Memory Network (**MN**)
  - [Sukhbaatar et al., NIPS 2015]
- Query Reduction Network (**QRN**)
  - [Seo et al., ICLR 2017]
- Gated Memory Network (**GMN**)
  - [Liu et al., EACL 2017]
- Sequence-to-sequence (**S2S**) + Attention
  - [Luong et al., EMNLP 2015]
- Pointer Network (**Ptr-Unk**)
  - [Gulcehre et al., ACL 2016]
- Memory-to-sequence (**Mem2Seq**)
  - [Madotto et al., ACL 2018]

Results: bAbI Dialogue

Retrieval

Generation

Task	QRN	MN	GMN	S2S+Attn	Ptr-Unk	Mem2Seq	GLMP K1	GLMP K3	GLMP K6
T1	99.4 (-)	99.9 (99.6)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)
T2	99.5 (-)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)
T3	74.8 (-)	74.9 (2.0)	74.9 (0)	74.8 (0)	85.1 (19.0)	94.7 (62.1)	**96.3 (75.6)**	96.0 (69.4)	96.0 (68.7)
T4	57.2 (-)	59.5 (3.0)	57.2 (0)	57.2 (0)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)
T5	**99.6 (-)**	96.1 (49.4)	96.3 (52.5)	98.4 (87.3)	99.4 (91.5)	97.9 (69.6)	99.2 (88.5)	99.0 (86.5)	99.2 (89.7)
T1 oov	83.1 (-)	72.3 (0)	82.4 (0)	81.7 (0)	92.5 (54.7)	94.0 (62.2)	**100 (100)**	**100 (100)**	99.3 (95.9)
T2 oov	78.9 (-)	78.9 (0)	78.9 (0)	78.9 (0)	83.2 (0)	86.5 (12.4)	**100 (100)**	**100 (100)**	99.4 (94.6)
T3 oov	75.2 (-)	74.4 (0)	75.3 (0)	75.3 (0)	82.9 (13.4)	90.3 (38.7)	95.5 (65.7)	**96.7 (72.9)**	95.9 (67.7)
T4 oov	56.9 (-)	57.6 (0)	57.0 (0)	57.0 (0)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)
T5 oov	67.8 (-)	65.5 (0)	66.7 (0)	65.7 (0)	73.6 (0)	84.5 (2.3)	**92.0 (21.7)**	91.0 (17.7)	91.8 (21.4)

without copy ability

with copy ability

Results: Stanford Multi-Domain

Automatic Evaluation

	Rule-Based*	KVR*	S2S	S2S + Attn	Ptr-Unk	Mem2Seq	GLMP H1	GLMP H3	GLMP H6
BLEU	6.6	13.2	8.4	9.3	8.3	12.6	13.83	**14.79**	12.37
Entity F1	43.8	48.0	10.3	19.9	22.7	33.4	57.25	**59.97**	53.54
Schedule F1	61.3	62.9	9.7	23.4	26.9	49.3	68.74	**69.56**	69.38
Weather F1	39.5	47.0	14.1	25.6	26.7	32.8	60.87	**62.58**	55.89
Navigation F1	40.4	41.3	7.0	10.8	14.9	20.0	48.62	**52.98**	43.08

Human Evaluation

	Mem2Seq	GLMP	Human
Appropriate	3.89	4.15	4.6
Humanlike	3.80	4.02	4.54

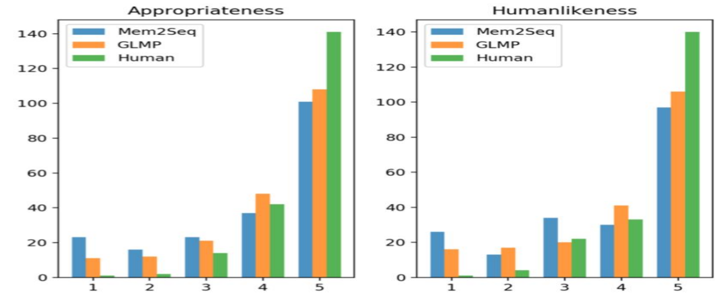

Results: Ablation Study

- **Ablation contextual dialogue history (w/o H)**
  - The hidden states of context RNN are not written into the external knowledge
  - The performance drop is serious in bAbI OOV scenario. (-5.3% in T5)
- **Ablation global memory pointer (w/o G)**
  - The external knowledge is not filtered by the global memory pointer
  - The performance drop is serious in SMD human-human scenario. (-11.47%)

	bAbI Dialogue OOV Per-response Accuracy					SMD Entity F1
	T1	T2	T3	T4	T5	All
GLMP	100 (-)	100 (-)	95.5 (-)	100 (-)	92.0 (-)	57.25 (-)
GLMP w/o H	90.4 (-9.6)	85.6 (-14.4)	95.4 (-0.1)	100 (-0)	86.2 (-5.3)	47.96 (-9.29)
GLMP w/o G	100 (-0)	91.7 (-8.3)	95.5 (-0)	100 (-0)	92.4 (+0.4)	45.78 (-11.47)

Results: Qualitative Study (1/2)

items	address	poi_type	traffic_info	distance
76	91_el_camino_real	gas_station	car_collision_nearby	5_miles
coupa	394_van_ness_ave	coffee_or_tea_place	road_block_nearby	6_miles
jing_jing	113_arbol_dr	chinese_restaurant	car_collision_nearby	3_miles
palo_alto_garage_r	481_amaranta_ave	parking_garage	no_traffic	5_miles
panda_express	842_arrowhead_way	chinese_restaurant	moderate_traffic	5_miles
peets_coffee	9981_archuleta_ave	coffee_or_tea_place	moderate_traffic	4_miles
stanford_childrens_health	899_ames_ct	hospital	moderate_traffic	5_miles
willows_market	409_bollard_st	grocery_store	car_collision_nearby	3_miles

Driver: Give me the directions to the closest grocery store
GLMP: The nearest grocery_store is willows_market, 3_miles away at 409_bollard_st
Driver: What is the next closest location?
GLMP: 409_bollard_st is the only 3_miles away
Driver: Ok then make a route to willows_market, avoid all traffic
GLMP: I sent the info on your screen, drive carefully .

Results: Qualitative Study

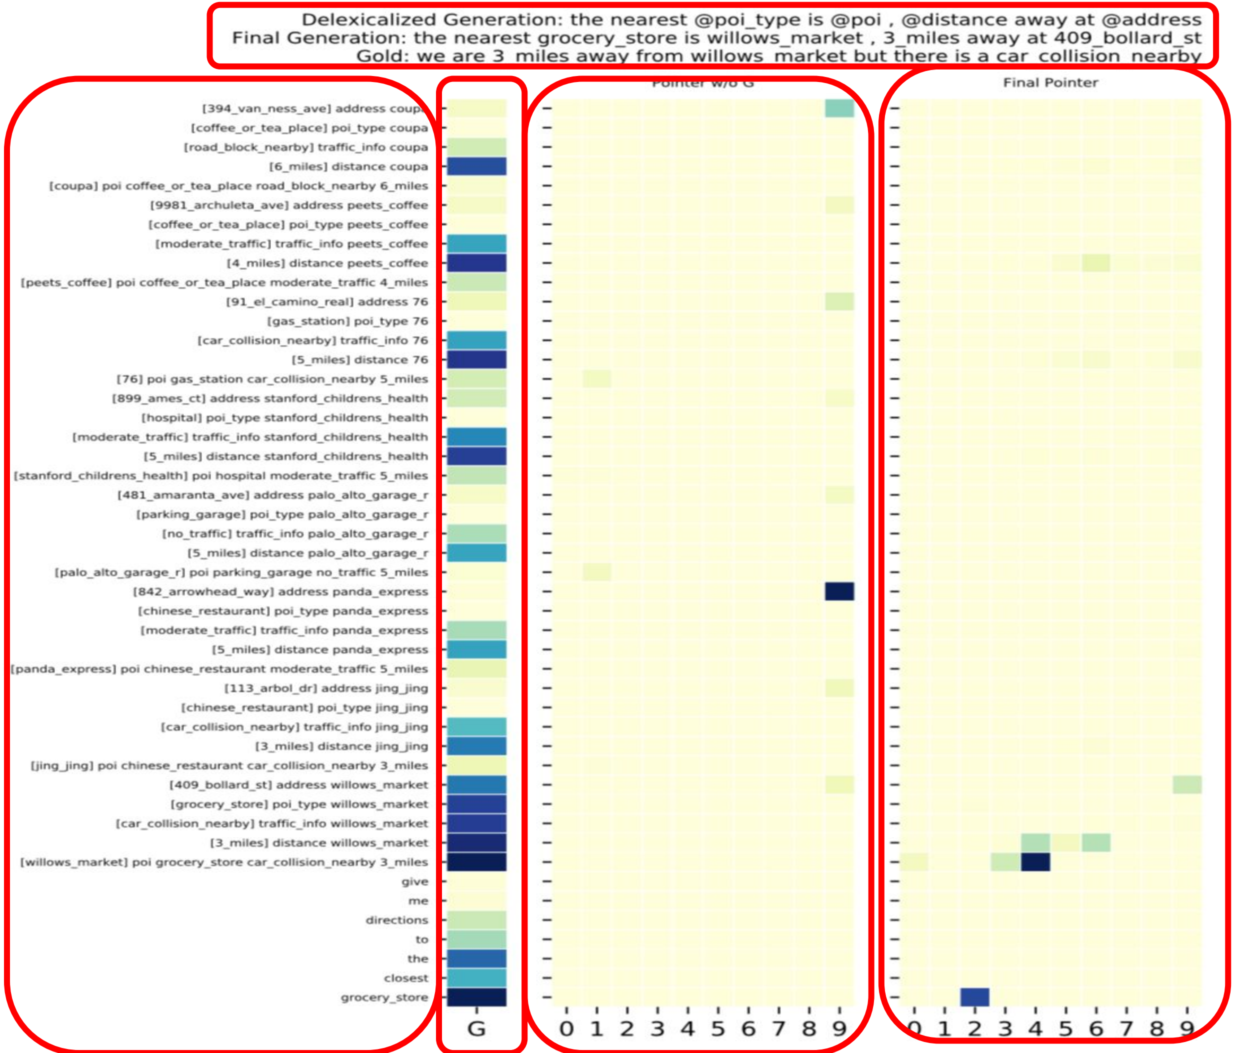

Results: Error Analysis

- **bAbI dialogue**
    - **Recommend restaurants** (Task 3)
        - Keep recommending some with higher scores but rejected by users
- **SMD dataset**
    - **KB understanding problem**
        - When several options matched are found in KB, model can barely recognize
            - Ex: Are you talking about the doctor appointment on Wednesday or Friday?
        - When the response requires checking multiple nodes
            - Ex: No, it will not rain for the next week in Palo Alto
    - **Copy mismatch**
        - Sketch tags mismatch with entity values
            - Ex: @address tag but copy "4 miles" out.
        - Generated sketch response includes missing slot values
            - Ex: Your @event is on @day with @party

Delexicalized Generation: your @event is on @date at @time with @party
Final Generation: your doctor is on the_5th at 6pm with alex
Gold: are you talking about the doctor_appointment on wednesday or the one on the_5th ?
Pointer w/o G
Final Pointer
[tom] party football_activity
[the_5th] date football_activity
[9am] time football_activity
[alex] party dentist_appointment
[sunday] date dentist_appointment
[2pm] time dentist_appointment
[discuss_the_merger] agenda conference
[conference_room_50] room conference
[management] party conference
[friday] date conference
[10am] time conference
[father] party doctor_appointment
[wednesday] date doctor_appointment
[7pm] time doctor_appointment
[sister] party swimming_activity
[the_13th] date swimming_activity
[1pm] time swimming_activity
[alex] party doctor_appointment
[the_5th] date doctor_appointment
[6pm] time doctor_appointment
[jon] party dentist_appointment
[thursday] date dentist_appointment
[2pm] time dentist_appointment
[discuss_the_company_picnic] agenda meeting
[100_conference_room] room meeting
[management] party meeting
[the_3rd] date meeting
[2pm] time meeting
what
time
is
my
doctor
apartment
and
who
is
it
with
?
G
0 1 2 3 4 5 6 7 8
Delexicalized Generation: the nearest @poi_type is @poi , @distance away at @address
Final Generation: the nearest parking_garage is civic_center_garage , 4_miles away at 5_miles
Gold: the closest parking_garage is civic_center_garage , located 4_miles away at 270_altaire_walk
[783_arcadia_pl] address chevron
[gas_station] poi_type chevron
[moderate_traffic] traffic_info chevron
[3_miles] distance chevron
[chevron] poi gas_station moderate_traffic 3_miles
[271_springer_street] address mandarin_roots
[chinese_restaurant] poi_type mandarin_roots
[moderate_traffic] traffic_info mandarin_roots
[4_miles] distance mandarin_roots
[mandarin_roots] poi chinese_restaurant moderate_traffic 4_miles
[408_university_ave] address trader_joes
[grocery_store] poi_type trader_joes
[no_traffic] traffic_info trader_joes
[5_miles] distance trader_joes
[trader_joes] poi grocery_store no_traffic 5_miles
[638_amherst_st] address sigona_farmers_market
[grocery_store] poi_type sigona_farmers_market
[no_traffic] traffic_info sigona_farmers_market
[4_miles] distance sigona_farmers_market
[sigona_farmers_market] poi grocery_store no_traffic 4_miles
[347_alta_mesa_ave] address jills_house
[friends_house] poi_type jills_house
[heavy_traffic] traffic_info jills_house
[4_miles] distance jills_house
[jills_house] poi friends_house heavy_traffic 4_miles
[270_altaire_walk] address civic_center_garage
[parking_garage] poi_type civic_center_garage
[no_traffic] traffic_info civic_center_garage
[4_miles] distance civic_center_garage
[civic_center_garage] poi parking_garage no_traffic 4_miles
[434_arastradero_rd] address ravenswood_shopping_center
[shopping_center] poi_type ravenswood_shopping_center
[heavy_traffic] traffic_info ravenswood_shopping_center
[4_miles] distance ravenswood_shopping_center
[ravenswood_shopping_center] poi shopping_center heavy_traffic 4_miles
what
are
the
directions
to
the
closest
parking_garage
G
0 1 2 3 4 5 6 7 8 9


thank you
Q/A