香港中文大學
The Chinese University of Hong Kong

salesforce

Explicit Memory Tracker with Coarse-to-Fine Reasoning for Conversational Machine Reading

Yifan Gao[1], Chien-Sheng Wu[2], Shafiq Joty[2], Caiming Xiong[2], Richard Socher[2], Irwin King[1], Michael R. Lyu[1], Steven C.H. Hoi[2]

1. The Chinese University of Hong Kong 2. Salesforce Research

Code & Models: https://github.com/Yifan-Gao/explicit_memory_tracker

Machine Comprehension

SQuAD: 100,000+ Questions for Machine Comprehension of Text (Rajpurkar et al., 2016)

The current Chief Executive is **Carrie Lam**, who was selected on 26 March 2017, appointed by the Central People's Government with the State Council Decree signed by Premier Li Keqiang, on 11 April 2017 and took office on 1 July 2017.

✓ Literal Answer

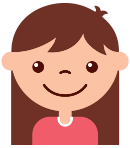

Q: Who is the chief executive of Hong Kong?

A: Carrie Lam

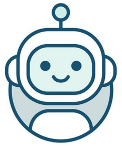

Conversational Question Answering

CoQA: A Conversational Question Answering Challenge (Reddy et al., 2018)

Incumbent **Democratic** President Bill Clinton was ineligible to serve a **third term** due to **term limitations** in the 22nd Amendment of the Constitution, and Vice President Gore was able to secure the Democratic nomination with relative ease.

- ✓ Literal Answer
- ✓ Dialog Understanding

Q: What political party is Clinton a member of?

A: Democratic

Q: What was he ineligible to serve?

A: third term

Q: Why?

A: term limitations

However...

Interpreting Natural Language Rules

The text to read may not contains the literal answer, but it contains a **recipe** to derive it.

You'll carry on paying National Insurance for the first 52 weeks you're abroad if you're working for an employer outside the EEA.

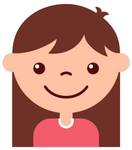

Q: I am working for an employer in Canada. Do I need to carry on paying UK National Insurance?

???

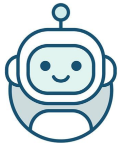

Background

Interpreting Natural Language Rules

You'll carry on paying National Insurance **for the first 52 weeks you're abroad** if you're working for an employer outside the EEA.

Q: I am working for an employer in Canada. Do I need to carry on paying UK National Insurance?

A: Have you been working abroad 52 weeks or less?

Yes, I have.

Yes, you need to carry on paying that.

- ✓ Non-literal Answer
- ✓ Dialog Understanding
- ✓ Proactive Interaction

ShARC: Shaping Answers with Rules through Conversation

(Saeidi et al., 2018)

- A user post her <u>scenario</u> and asks an <u>initial question</u> about her final pay
- Without knowing the rule text, the initial question is usually underspecified
- A machine need to read the relevant <u>rule text</u>, and ask a series of <u>clarification questions</u> until it can conclude with a certain answer

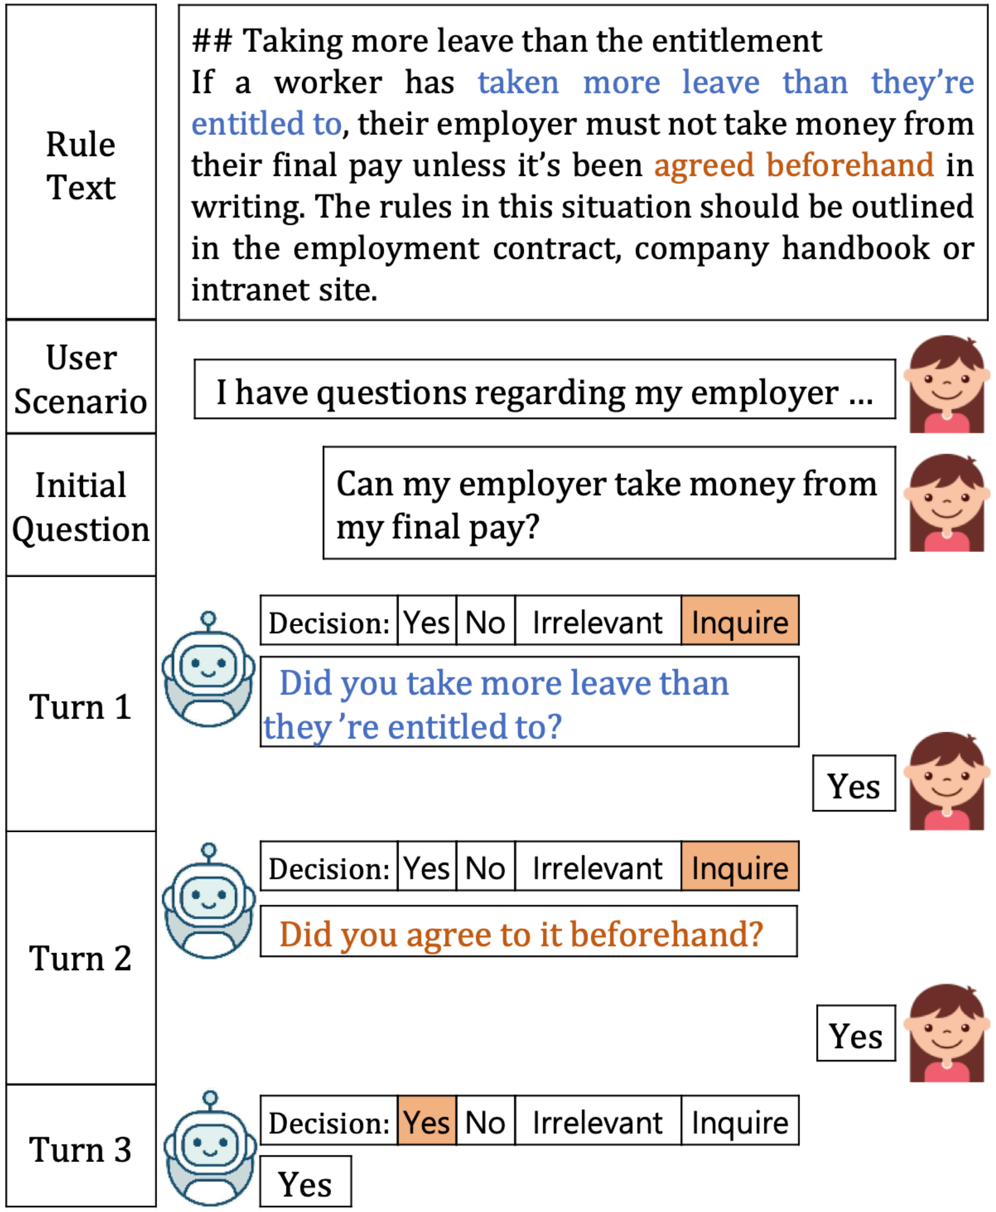

ShARC: Shaping Answers with Rules through Conversation

Task Definition

Input x = (q, h, r, s)

- r: rule text
- s: scenario
- q: underspecified question
- h: dialog history (QA1, QA2, ...)

Output (two subtasks):

- Make a decision $y \in$ {Yes, No, Irrelevant, Inquire}
- If *Inquire*, ask a follow-up question

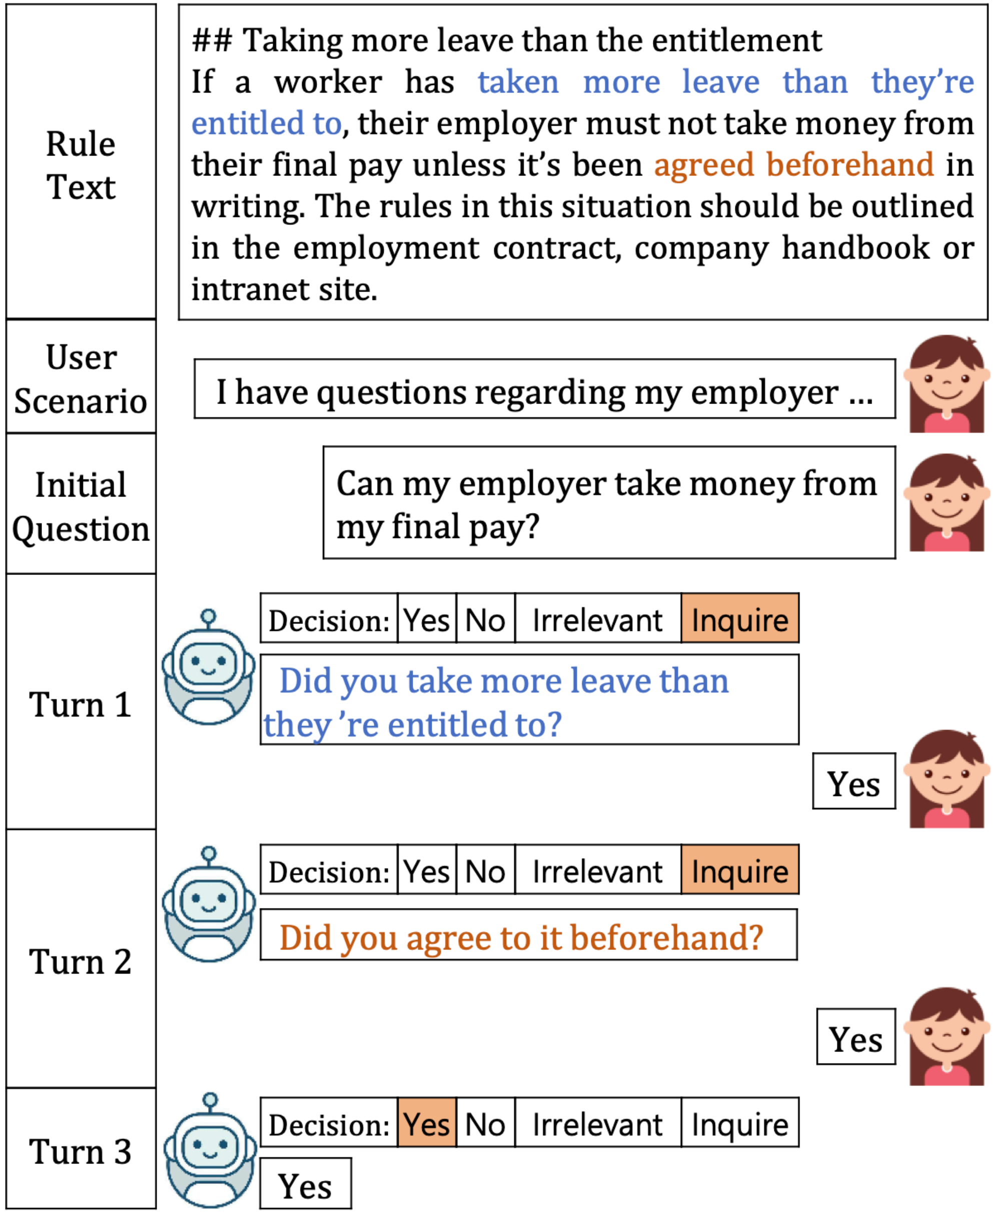

ACL 2020, Explicit Memory Tracker with Coarse-to-Fine Reasoning for Conversational Machine Reading

Proposed Solution

Contributions

- Explicit Memory Tracker (EMT)
  - Explicitly track whether conditions listed in the rule text have been fulfilled or not

- Coarse-to-fine (C2F) Reasoning
  - A coarse-to-fine approach to reason out which part of the rule text is underspecified, and ask a question accordingly

- Our proposed solution achieves new state-of-the-art results on the ShARC benchmark

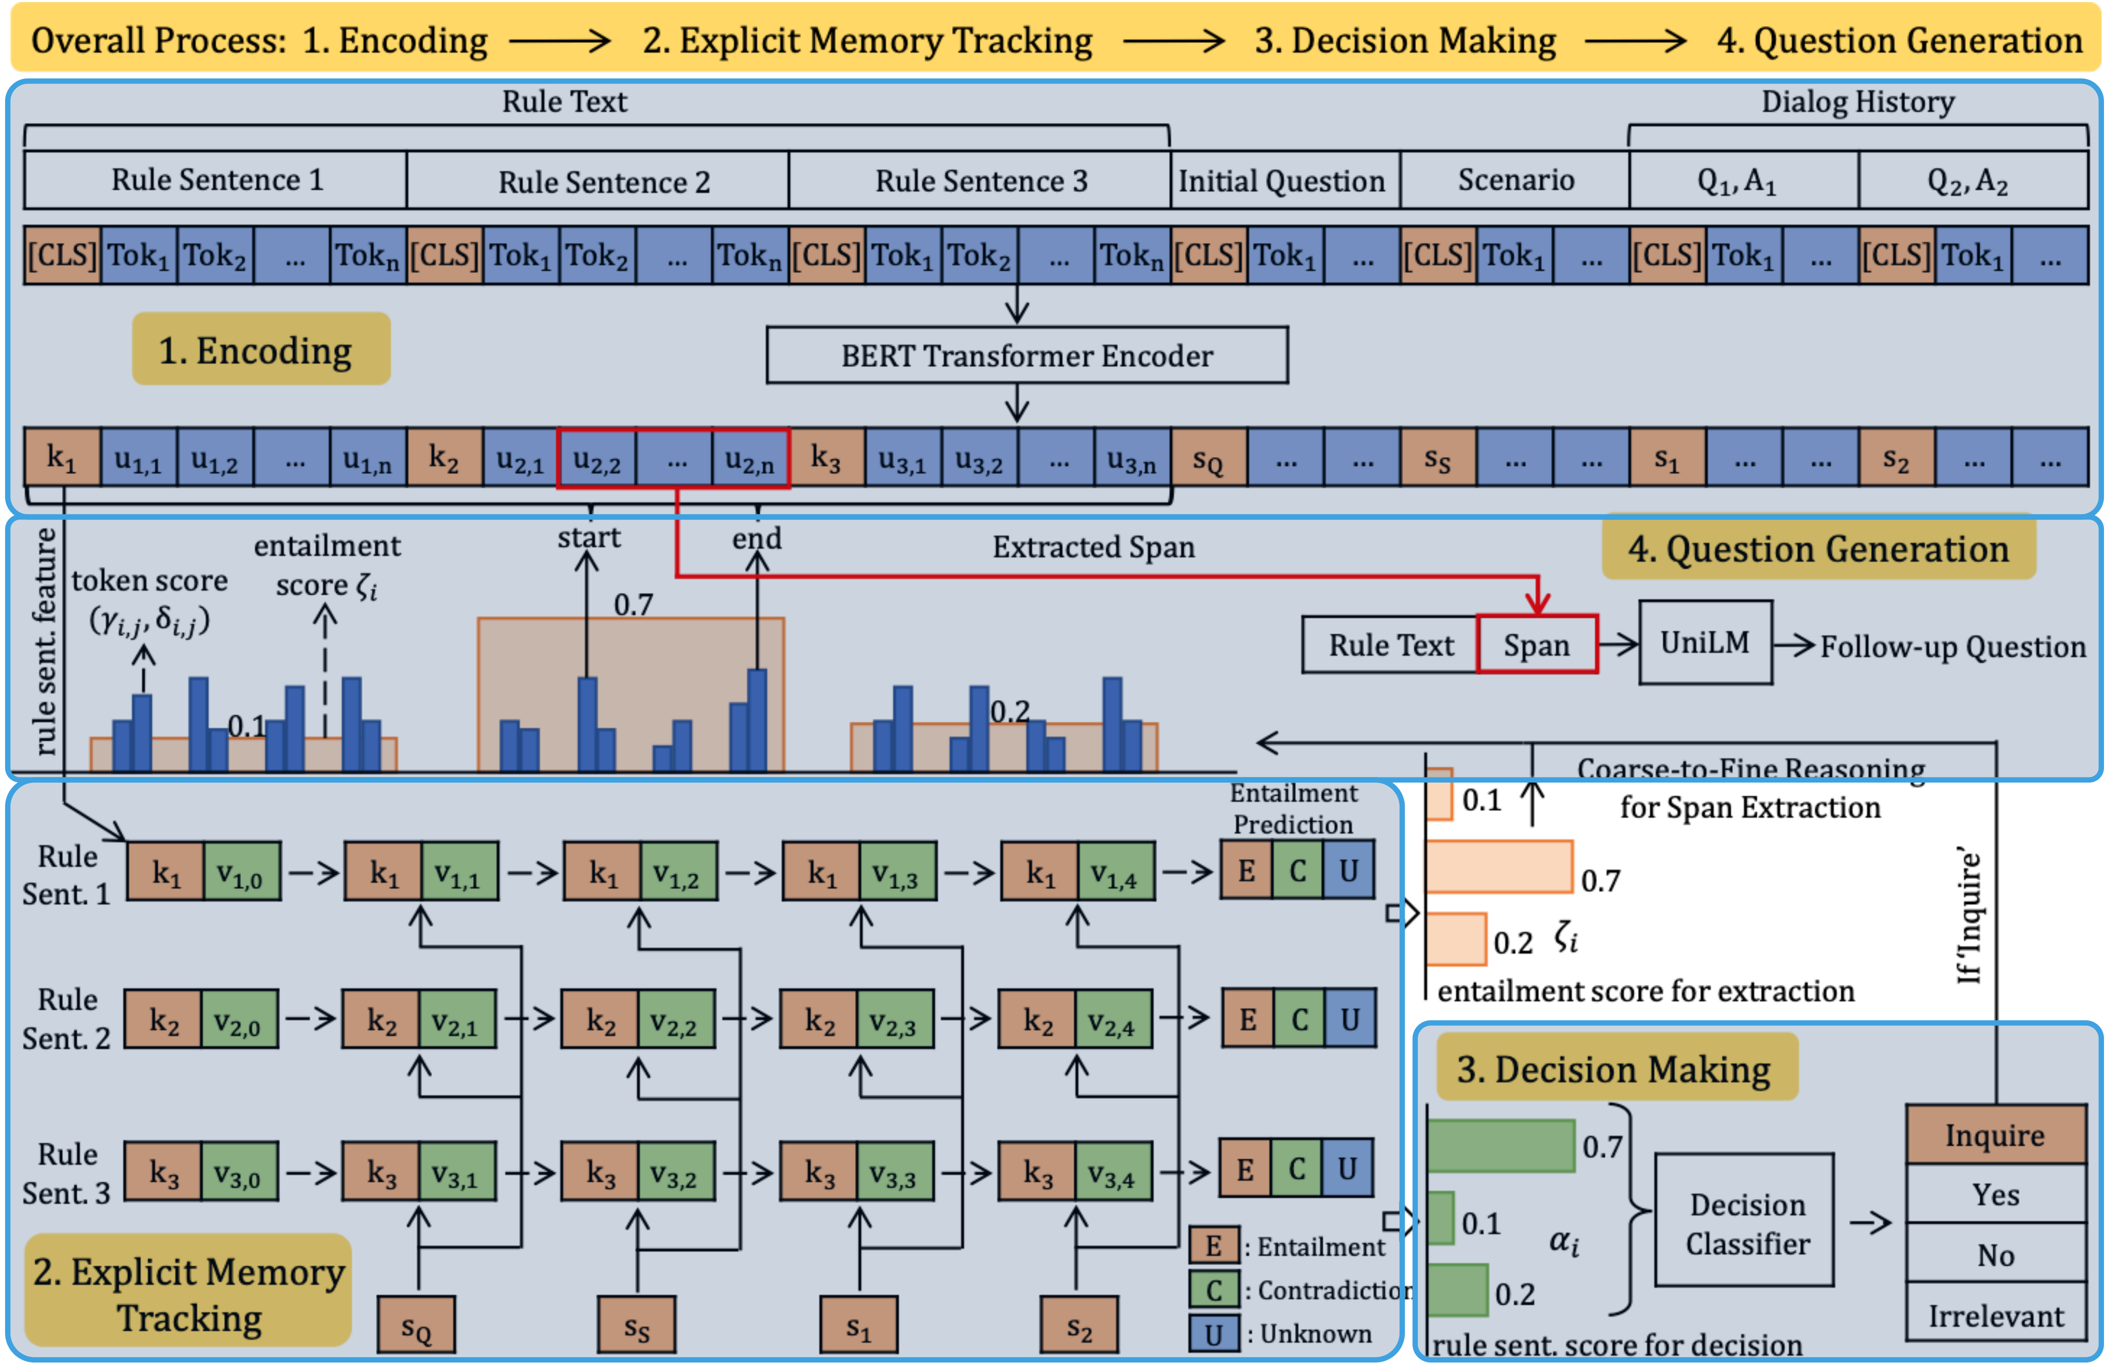

ACL 2020, Explicit Memory Tracker with Coarse-to-Fine Reasoning for Conversational Machine Reading

Proposed Solution

Encoding

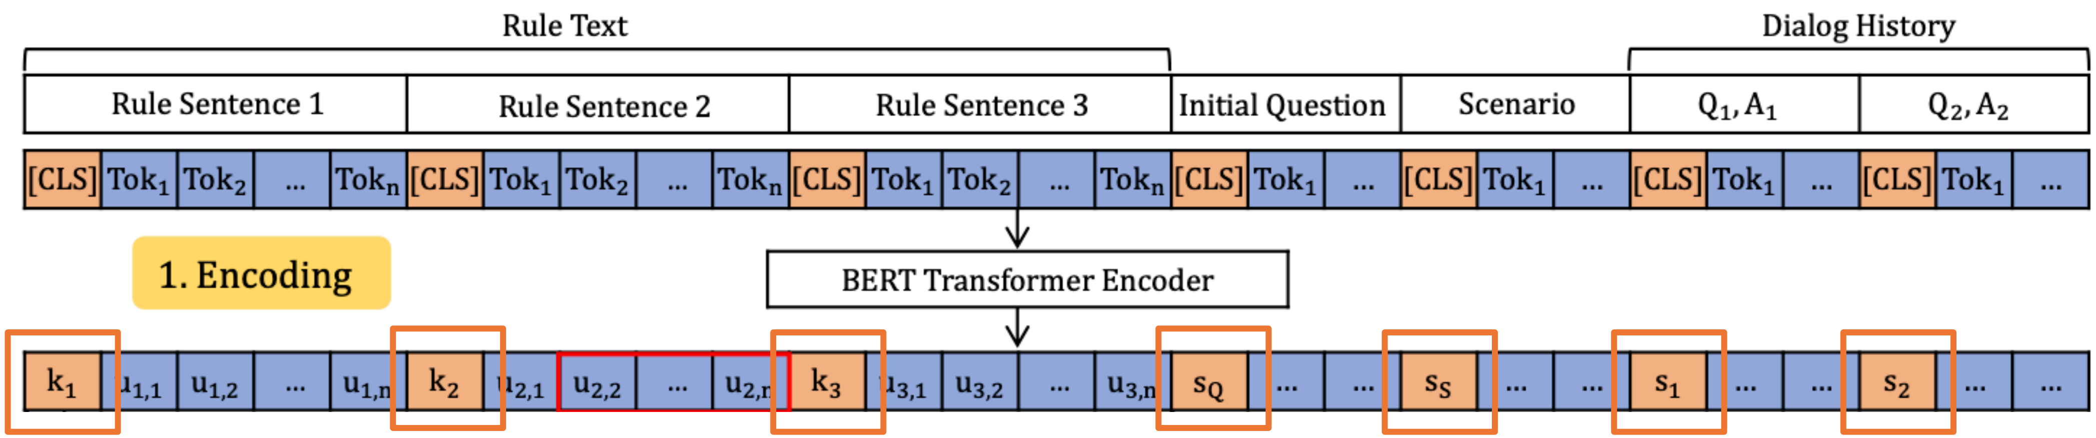

1. Parse the rule text into multiple rule sentences according to rules
2. Insert **[CLS]** token at the start of each rule sentence, initial question, scenario, and question-answer pairs in the dialog history
3. Concatenate all information and feed to BERT for encoding
4. **[CLS]** symbol is treated as the feature representation of the sentence that follows it

Proposed Solution

Explicit Memory Tracking

Rule sentences $\mathbf{k}_1, \mathbf{k}_2, \dots, \mathbf{k}_M$

Find Implications

User provided information:

- ❖ Initial question $\mathbf{s}_Q$
- ❖ Scenario $\mathbf{s}_S$
- ❖ Dialog history $\mathbf{s}_1, \dots, \mathbf{s}_P$

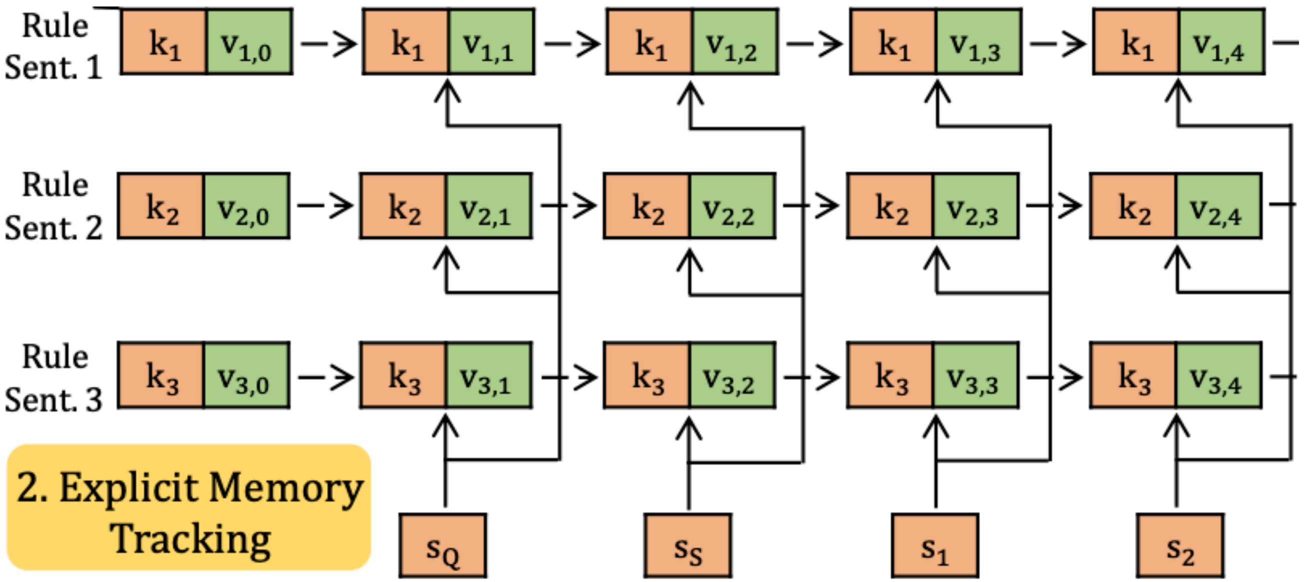

- ❖ We propose Explicit Memory Tracker (EMT), a gated recurrent memory-augmented neural network
- ❖ EMT explicitly tracks the states of rule sentences by sequentially reading the user provided information

Proposed Solution

Explicit Memory Tracking

EMT assigns a state $\mathbf{v}_i$ to each key $\mathbf{k}_i$, and sequentially reads user information

At time step t:

$$\tilde{\mathbf{v}}_{i,t} = \mathrm{ReLU}(\mathbf{W}_k\mathbf{k}_i + \mathbf{W}_v\mathbf{v}_{i,t} + \mathbf{W}_s\mathbf{s}_t),$$

$$g_i = \sigma(\mathbf{s}_t^\top\mathbf{k}_i + \mathbf{s}_t^\top\mathbf{v}_{i,t}) \in [0, 1],$$

$$\mathbf{v}_{i,t} = \mathbf{v}_{i,t} + g_i \odot \tilde{\mathbf{v}}_{i,t} \in \mathbb{R}^d, \mathbf{v}_{i,t} = \frac{\mathbf{v}_{i,t}}{\|\mathbf{v}_{i,t}\|}$$

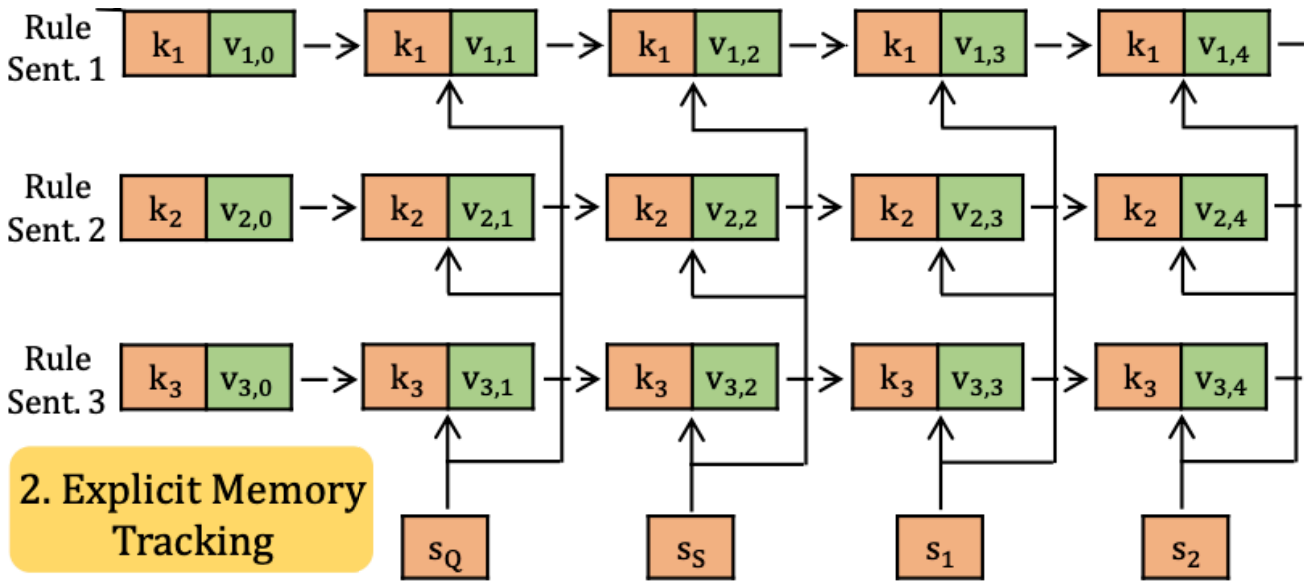

Keys and final states of rule sentences are denoted as $(\mathbf{k}_1, \mathbf{v}_1), \ldots, (\mathbf{k}_M, \mathbf{v}_M)$

- ❖ Decision Making Module
- ❖ Question Generation Module

Proposed Solution
Decision Making

Based on the most up-to-date key-value states of rule sentences $(\mathbf{k}_1, \mathbf{v}_1), \ldots, (\mathbf{k}_M, \mathbf{v}_M)$, EMT makes a decision among *Yes, No, Irrelevant, Inquire*

$$\alpha_i = \mathbf{w}_\alpha^\top [\mathbf{k}_i; \mathbf{v}_i] + b_\alpha \in \mathbb{R}^1$$

$$\tilde{\alpha}_i = \mathrm{softmax}(\alpha)_i \in [0, 1]$$

$$\mathbf{c} = \sum_i \tilde{\alpha}_i [\mathbf{k}_i; \mathbf{v}_i] \in \mathbb{R}^d$$

$$\mathbf{z} = \mathbf{W}_z \mathbf{c} + \mathbf{b}_z \in \mathbb{R}^4$$

The decision making module is trained with the following cross entropy loss:

$$\mathcal{L}_{\mathrm{dec}} = -\log\ \mathrm{softmax}(\mathbf{z})_l$$

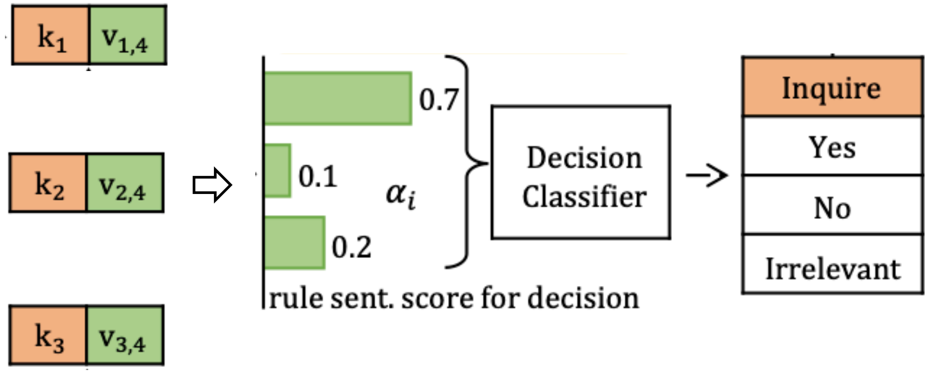

Proposed Solution

Subtask: Entailment State Prediction

- Explicitly track whether a condition listed in the rule has already been satisfied or not
- The possible entailment labels are:
  - Entailment (E)
  - Contradiction (C)
  - Unknown (U)

$$\mathbf{e}_i = \mathbf{W}_e[\mathbf{k}_i; \mathbf{v}_i] + \mathbf{b}_e \in \mathbb{R}^3$$

$$\mathcal{L}_{\text{entail}} = -\frac{1}{M}\sum_{i=1}^{M} \log\ \text{softmax}(\mathbf{e}_i)_r$$

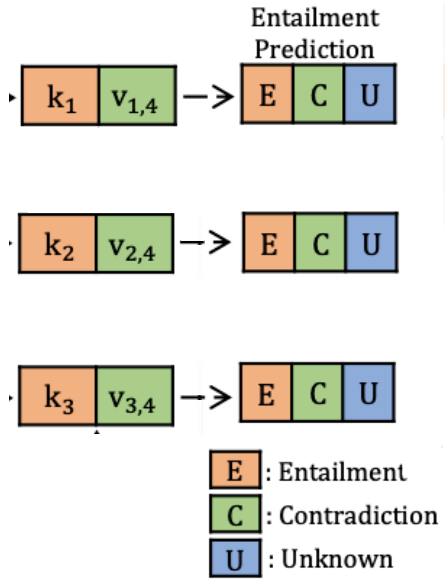

Proposed Solution

Follow-up Question Generation

When the decision is ‘Inquire’, a <u>follow-up question</u> is required for further clarification.

We adopt a two-step approach:

1. Extract a span inside the rule text which contains the underspecified user information
2. Rephrase the extracted span into a follow-up question

Proposed Solution

Follow-up Question Generation: Coarse-to-fine Underspecified Span Extraction

1. Coarse-to-fine Underspecified Span Extraction
   1) Identify underspecified rule sentence ζ_i

$$\zeta_i = \text{softmax}(\beta_{\text{unknown}})_i \in [0, 1]$$

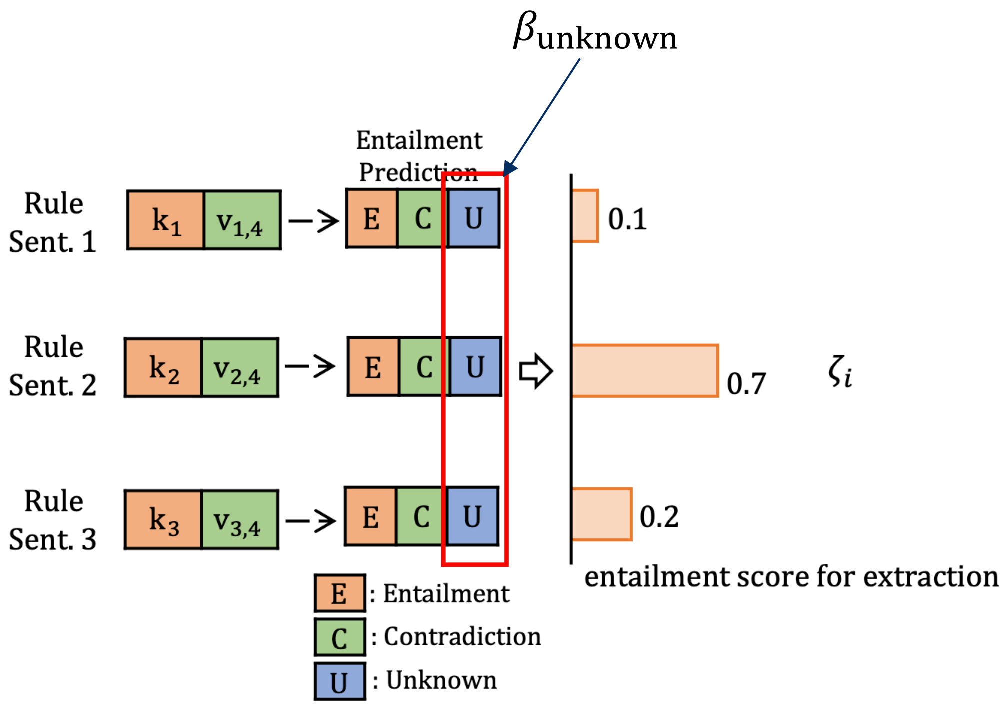

Proposed Solution

Follow-up Question Generation: Coarse-to-fine Underspecified Span Extraction

1. Coarse-to-fine Underspecified Span Extraction
   1) Identify underspecified rule sentence ζ_i
   2) Extract a span within each rule sentence $(\gamma_{i,j}, \delta_{i,j})$

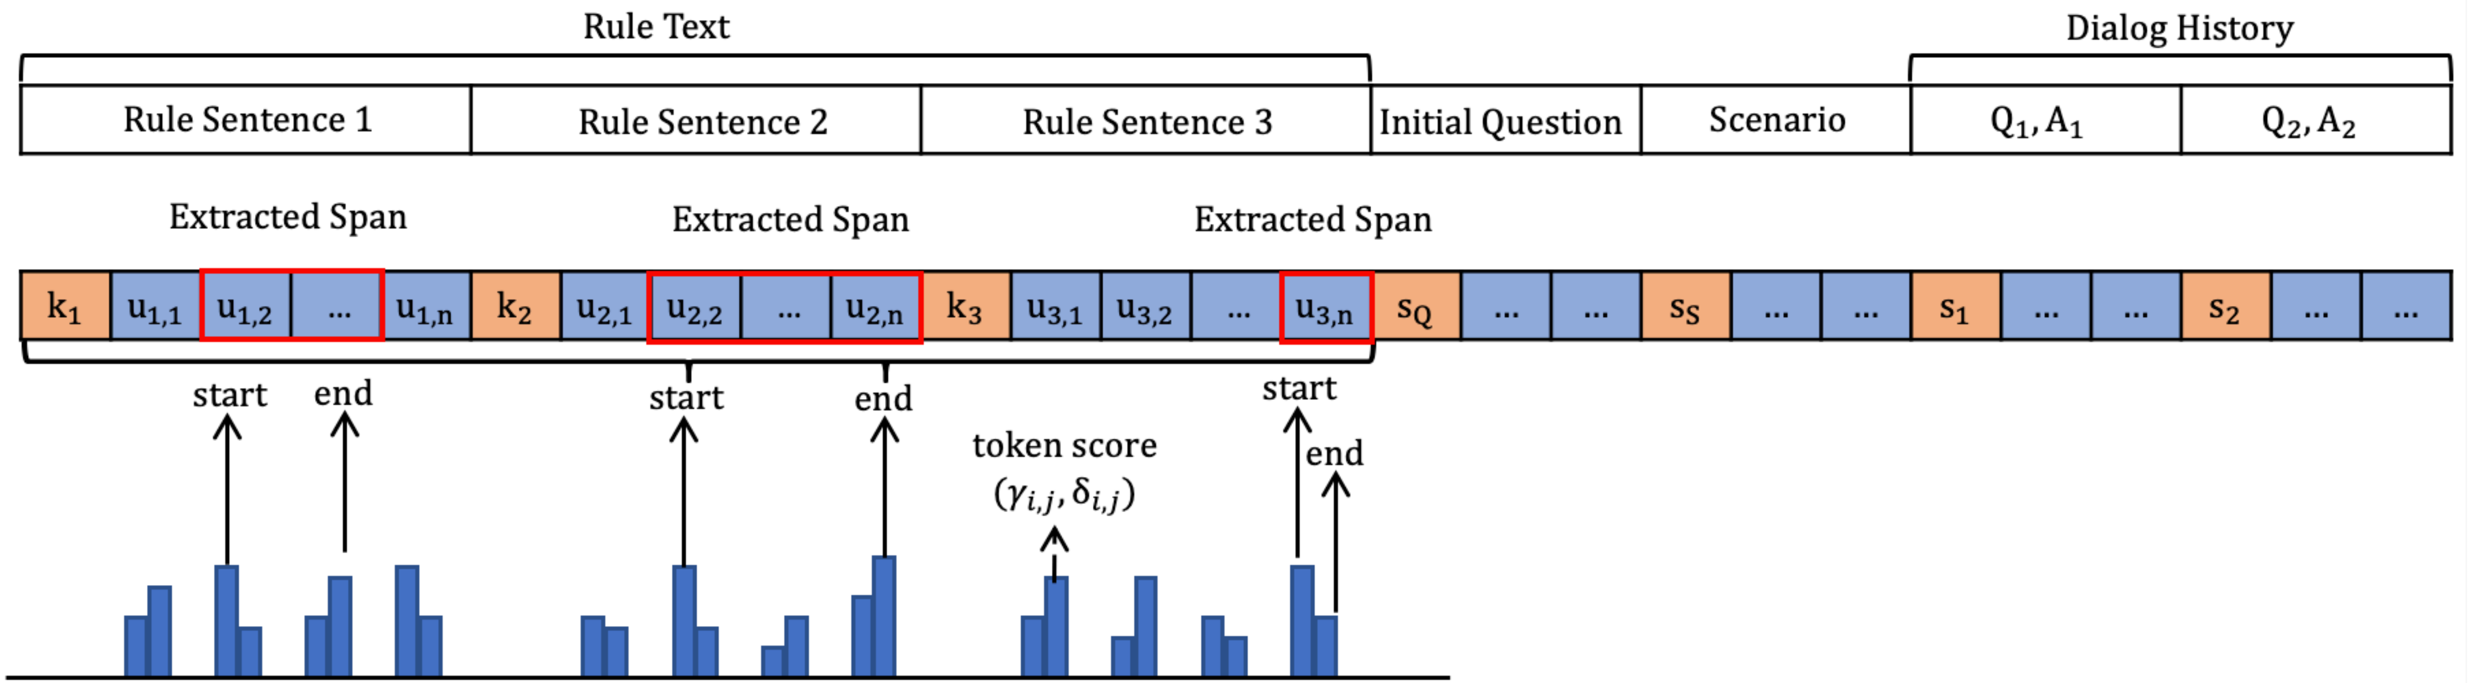

Proposed Solution

Follow-up Question Generation: Coarse-to-fine Underspecified Span Extraction

1. Coarse-to-fine Underspecified Span Extraction
    1) Identify underspecified rule sentence ζ_i
    2) Extract a span within each rule sentence $(\gamma_{i,j}, \delta_{i,j})$
    3) Select the span with the highest span score $\zeta_i * (\gamma_{i,j}, \delta_{i,j})$

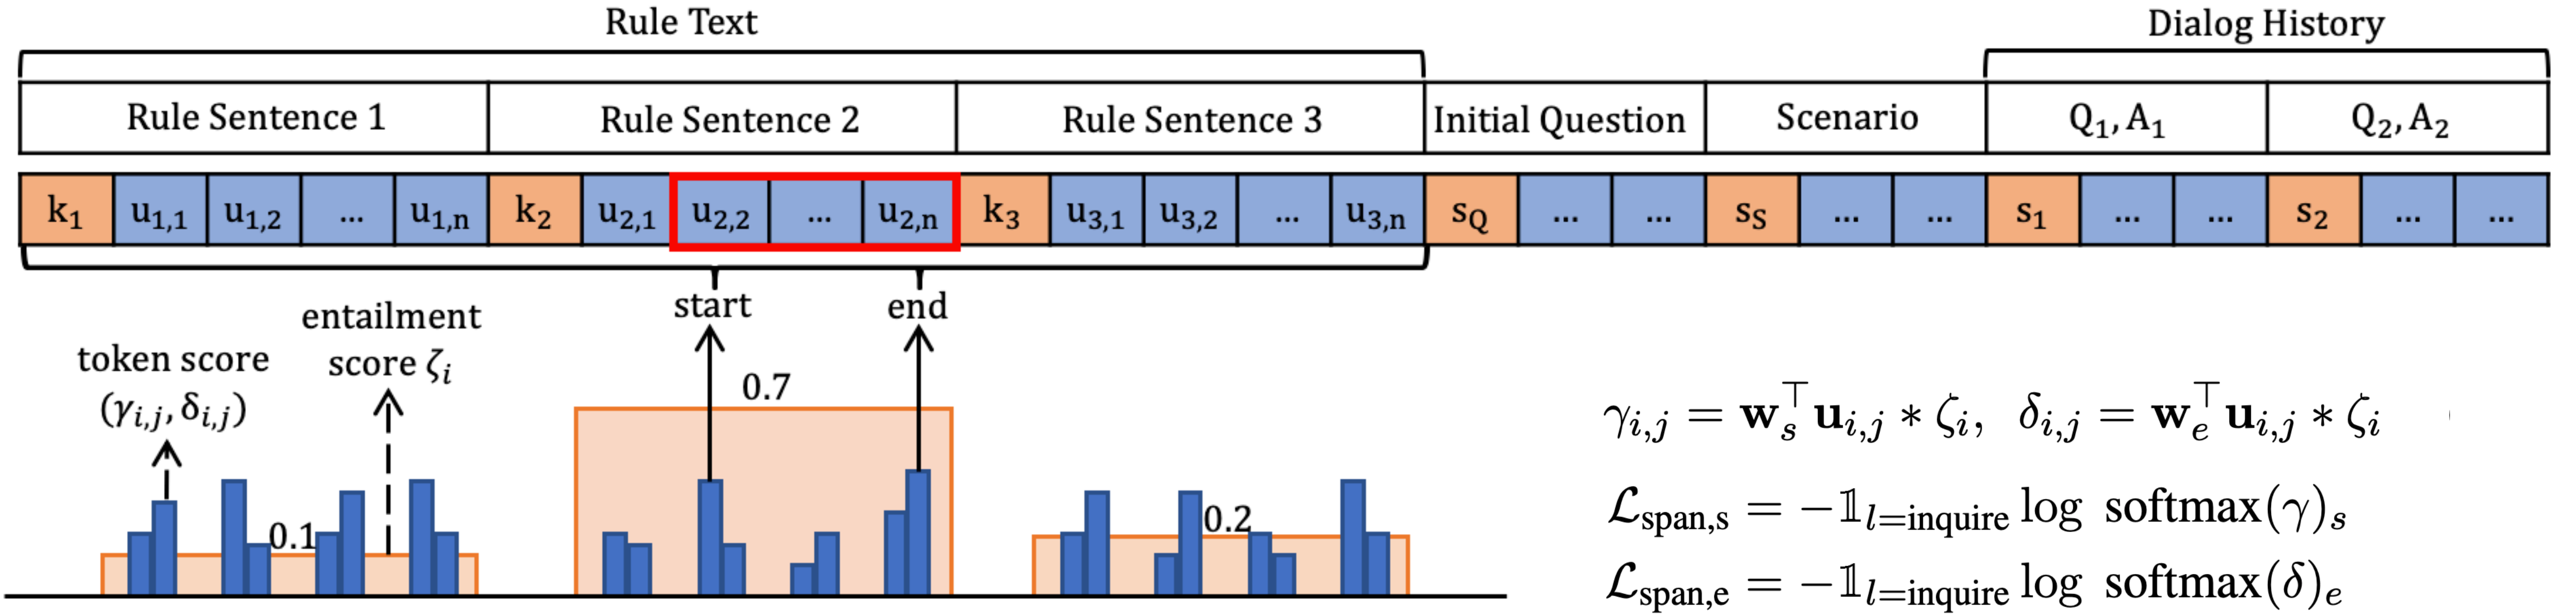

$$\gamma_{i,j} = \mathbf{w}_s^\top \mathbf{u}_{i,j} * \zeta_i, \quad \delta_{i,j} = \mathbf{w}_e^\top \mathbf{u}_{i,j} * \zeta_i$$

$$\mathcal{L}_{\text{span,s}} = -\mathbb{1}_{l=\text{inquire}} \log \ \text{softmax}(\gamma)_s$$

$$\mathcal{L}_{\text{span,e}} = -\mathbb{1}_{l=\text{inquire}} \log \ \text{softmax}(\delta)_e$$

Proposed Solution

Follow-up Question Generation

2. Question Rephrasing

1) Finetune UniLM (Dong et al, 2019), a pretrained language model
2) [CLS] rule text [SEP] span [SEP]

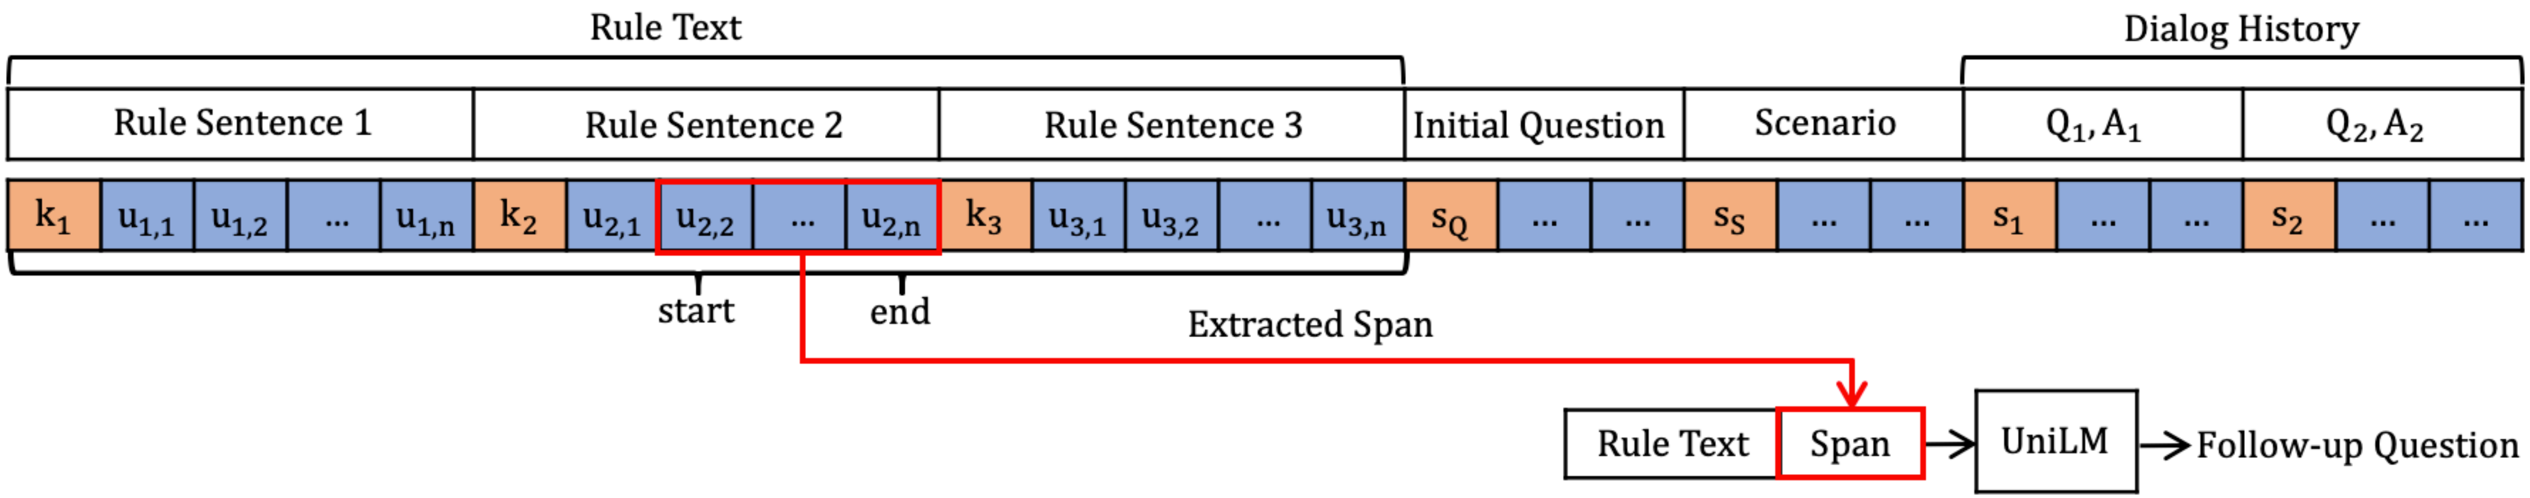

Proposed Solution

Overall Loss for EMT

The overall loss is the sum of the decision loss, entailment prediction loss and span extraction loss:

$$\mathcal{L} = \mathcal{L}_{\text{dec}} + \lambda_1 \mathcal{L}_{\text{entail}} + \lambda_2 \mathcal{L}_{\text{span}}$$

Experiments

Experimental Setup

- Dataset:
  - ShARC CMR dataset (Saeidi et a
  - Train/Dev/Test dataset sizes ar
  - Test set is not public.
  - Leaderboard: [https://sharc-dat]()

ShARC: End-to-end Task

#	Model / Reference	Affiliation	Date	Micro Accuracy[%]	Macro Accuracy[%]	BLEU-1	BLEU-4
1	[Anonymous]	[Anonymous]	May 2020	**73.2**	**78.3**	**64.0**	49.1
2	EMT	Salesforce Research & CUHK	Nov 2019	69.4	74.8	60.9	46.0
3	EMT + entailment	Salesforce Research & CUHK	Mar 2020	69.1	74.6	63.9	**49.5**
4	[Anonymous]	[Anonymous]	Dec 2019	69.0	74.6	56.7	42.0
5	E3	University of Washington	Feb 2019	67.6	73.3	54.1	38.7
6	BiSon (single model)	NEC Laboratories Europe	Aug 2019	66.9	71.6	58.8	44.3

Experiments

Experimental Setup

- Dataset:
  - ShARC CMR dataset (Saeidi et al. 2018)
  - Train/Dev/Test dataset sizes are 21980/2270/8276.
  - Test set is not public.
  - Leaderboard: https://sharc-data.github.io/leaderboard.html
- Evaluation Metrics
  - End-to-End Task
    - Macro/Micro Accuracy for the decision making task (Yes/No/Irrelevant/Inquire)
    - If both the ground truth decision and the predicted decision are *Inquire*, BLEU is used to evaluate the quality of the generated follow-up question

    (This evaluation metric makes the comparison unfair if two models have different *Inquire* predictions)

Experiments

Experimental Setup

- Dataset:
  - ShARC CMR dataset (Saeidi et al. 2018)
  - Train/Dev/Test dataset sizes are 21980/2270/8276.
  - Test set is not public.
  - Leaderboard: https://sharc-data.github.io/leaderboard.html
- Evaluation Metrics
  - End-to-End Task
    - Macro/Micro Accuracy for the decision making task (Yes/No/Irrelevant/Inquire)
    - If both the ground truth decision and the predicted decision are *Inquire*, BLEU is used to evaluate the quality of the generated follow-up question
  - Oracle Question Generation Task
    - We propose a new evaluation perspective.
    - We ask the models to generate follow-up questions *whenever* the ground truth decision is *Inquire*, and compute the BLEU score.

Experiments

Leaderboard Submission

Models	End-to-End Task (Leaderboard Performance)			
	Micro Acc.	Macro Acc.	BLEU1	BLEU4
Seq2Seq (Saeidi et al., 2018)	44.8	42.8	34.0	7.8
Pipeline (Saeidi et al., 2018)	61.9	68.9	54.4	34.4
BERTQA (Zhong and Zettlemoyer, 2019)	63.6	70.8	46.2	36.3
UrcaNet (Sharma et al., 2019)	65.1	71.2	60.5	46.1
BiSon (Lawrence et al., 2019)	66.9	71.6	58.8	44.3
E^3 (Zhong and Zettlemoyer, 2019)	67.6	73.3	54.1	38.7
EMT (our single model)	**69.1**	**74.6**	**63.9**	**49.5**

Table 1: Performance on the blind, held-out test set of ShARC end-to-end task.

Experiments

Class-wise Decision Prediction Accuracy

Models	Yes	No	Inquire	Irrelevant
BERTQA	61.2	61.0	62.6	96.4
E^3	65.9	70.6	60.5	96.4
UrcaNet*	63.3	68.4	58.9	95.7
EMT	**70.5**	**73.2**	**70.8**	**98.6**

Table 2: Class-wise decision prediction accuracy on the development set (*: reported in the paper).

Experiments

Oracle Question Generation Task

Models	Oracle Question Generation Task			
	Development Set		Cross Validation	
	BLEU1	BLEU4	BLEU1	BLEU4
E^3	52.79±2.87	37.31±2.35	51.75	35.94
E^3+UniLM	57.09±1.70	41.05±1.80	56.94	42.87
EMT	**62.32**±1.62	**47.89**±1.58	**64.48**	**52.40**

Table 3: Performance on Oracle Question Generation Task. We show both results on the development set and 10-fold cross validation. E^3+UniLM replaces the editor of E^3 to our finetuned UniLM.

Experiments

Interpretability

E : Entailment   C : Contradiction   U : Unknown

β_{unknown}

Regulation Text A (parsed into six rule sentences: S1 ~ S6)		Entailment States		
		Turn 1	Turn 2	Turn 3
S1	Statutory Maternity Pay	U (99.99)	U (99.99)	U (99.99)
S2	To qualify for smp you must:	U (99.99)	U (99.99)	U (99.99)
S3	* earn on average at least £113 a week	U (99.93)	E (99.91)	E (99.67)
S4	* give the correct notice	U (99.97)	U (99.61)	C (99.81)
S5	* give proof you're pregnant	U (99.98)	U (99.75)	U (99.94)
S6	* have worked for your employer...	U (99.98)	U (99.70)	U (99.96)

Scenario: I've been old enough to get my pension. My wife just reached pension age last year. Neither of us have applied for it yet.

Initial Question: Do I qualify for SMP?

	Decision:	Generated Question	Answer
Turn 1:	Inquire	Do you earn on average at least £113 a week?	Yes
Turn 2:	Inquire	Did you give the correct notice?	No
Turn 3:	No		

Conclusion

- We propose a new approach called Explicit Memory Tracker (EMT) for conversational machine reading.
- EMT achieved a new state-of-the-art result on the ShARC CMR challenge.
- EMT also gains interpretability by showing the entailment-oriented reasoning process as the conversation flows.

Thanks!

香港中文大學
The Chinese University of Hong Kong

salesforce

Yifan Gao

Chien-Sheng Wu

Shafiq Joty

Caiming Xiong

Richard Socher

Irwin King

Michael R. Lyu

Steven C.H. Hoi

Code & Models: https://github.com/Yifan-Gao/explicit_memory_tracker