CASR: Generating Complex Sequences with Autoregressive Self-Boost Refinement

Hongwei Han · Mengyu Zhou · Shi Han · Xiu Li · Dongmei Zhang
Tsinghua University & Microsoft Research Asia

Code: https://github.com/RalphHan/CASR

Email: hhw20@mails.tsinghua.edu.cn/ralphhan666@gmail.com

Multi-hop Dependency

What is dependency?

h: I like cooking
i: so I make dinner every day
h<--i

j: I adopt a dog
k: that is because I like animals
j-->k

l: We can speculate that the valence of O in H2O is -2
m: that is because O has 6 outermost electrons
n: which helps O steal an electron from H
o: and makes O display a negative valence.
l-->m<--n<--o

What is multi-hop dependency?

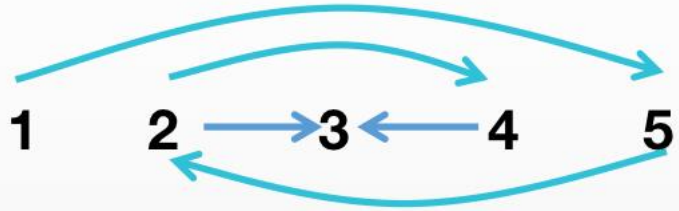

Best Order to Generate

The best order is to decouple the dependency.

A 4x4 Sudoku Example. White cells denote blanks, and the green numbers in them denote the ground truth. The game solving order of a human is from 'a' to 'g', rather than row-by-row.

There are many sequence generation tasks where the best order to generate the target sequence is not from-left-to-right. Humans often solve complex problems based on the intrinsic order of the problem. For example,

1) the order to write code or formula is often bottom-up or top-down.

2) the order to generate logical natural language is to first analyze and then draw conclusions.

3) the order to solve Sudoku game is often from the easy parts to the difficult parts, as the difficult parts become easier when the simple parts are correctly solved.

Does the language model know the complete answer at the beginning of generation, or does it gradually recognize the complete answer as it generates text from left to right?

CASR framework

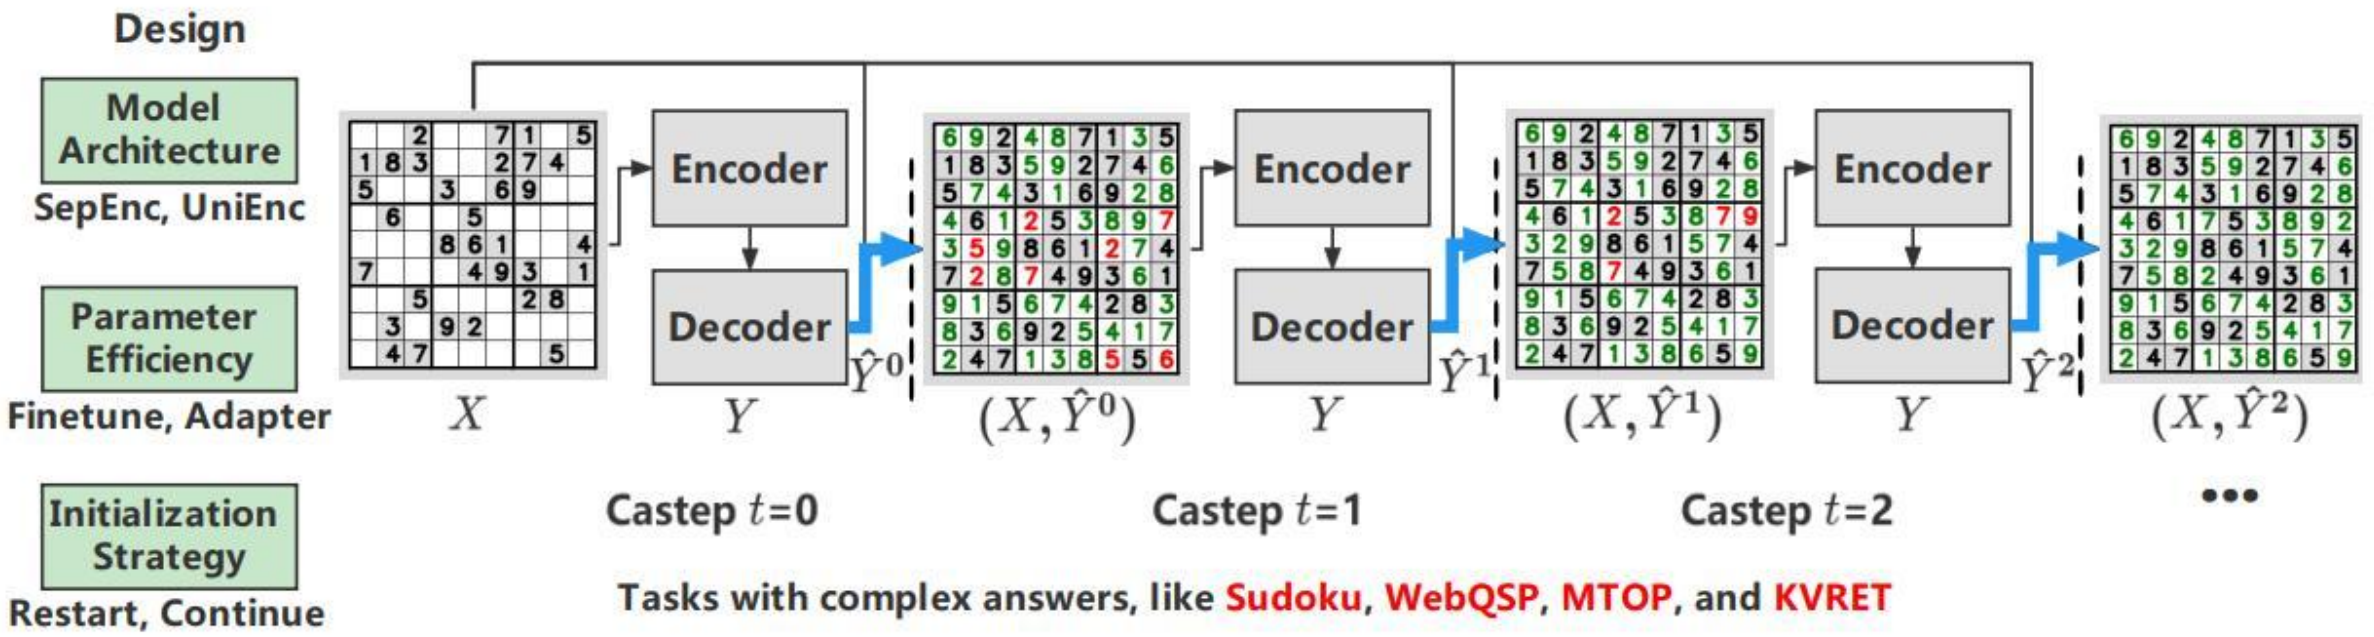

The Overview of CASR Framework. X, Y and \hat{Y} denote the input, ground truth and prediction, respectively. The blue arrows show how we iteratively added back the previous-step prediction \hat{Y}^{t-1} to the input for generating refined output \hat{Y}^ t. Through self-boost refinement, the model can refer to the results of the previous refinement at other positions (even on the right side) it depends on when deciding the token at a certain position.

Detailed Designs

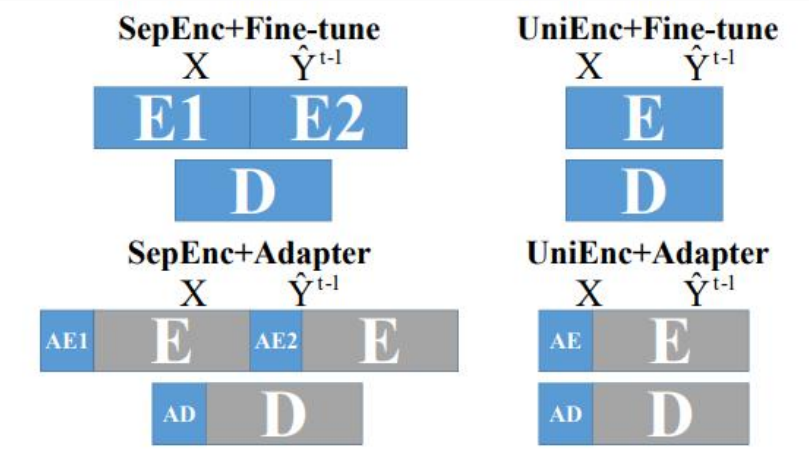

(a) Model Architectures and Parameter Efficiency Choices.

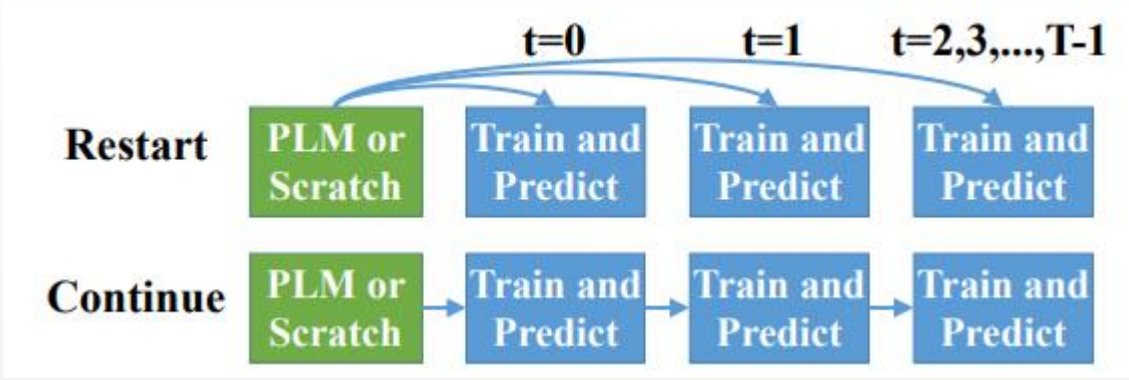

(b) Initialization Strategies.

SepEnc, $H = \text{Concat}(\text{Encoder1}(X), \text{Encoder2}(\hat{Y}^{t-1}))$。

UniEnc, $H = \text{Encoder}(\text{Concat}(X, \hat{Y}^{t-1}))$。

"**Restart**" means initializing each Mt with the same PLM weights (or from scratch). For SepEnc architecture, Both Encoder1 and Encoder2 can be initialized with the parameter weights of the PLM encoder.

"**Continue**" means initializing with the best checkpoint from the previous castep. In this way, M t inherits the knowledge from M t−1 to avoid the potential cold start issue at castep t.

Tasks & Datasets

Figure (a) is an example of **Sudoku**, which is an open dataset on Kaggle. The objective of the game is to correctly fill in the blanks, with no repetition of any two numbers in the same row, column, or 3x3 grid.

Figure (b) is an example of **WebQSP**. This is a classic knowledge base question-answering dataset, where the input consists of a knowledge graph and a natural language query, and the output is an executable s-Expression on the knowledge graph.

Figure (c) is an example of **MTOP**. Its input consists of a set of API calls and a natural language query, and the output is an executable tree-based task TOP Representation.

Figure (d) is an example of **KVRET**. This is a table-based dialogue task, where the input consists of a table and a natural language query, and the output is a natural language response.

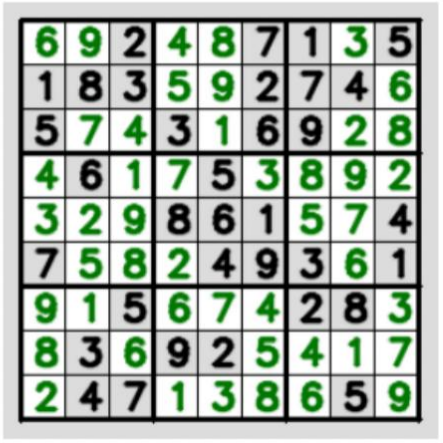

(a) Sudoku

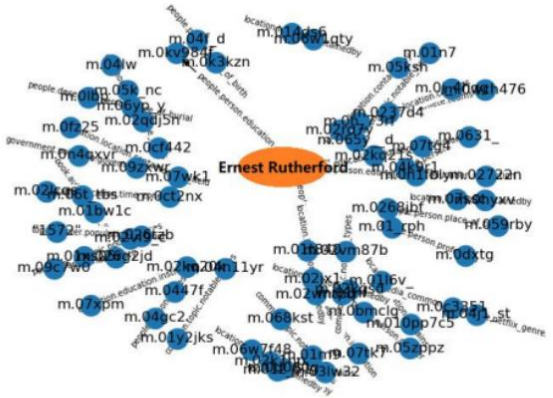

Request: what school did sir ernest rutherford go to?
Answer: (JOIN (R education.education.institution) (JOIN (R people.person.education) m.02m7r))

(b) WebQSP

IN:GET: MESSAGE, WEATHER, ALARM, INFO_RECIPES, STORIES_NEWS, REMINDER, RECIPES, EVENT, CALL_TIME, LIFE_EVENT, INFO_CONTACT, CONTACT, TIMER, REMINDER_DATE_TIME, AGE, SUNRISE, EMPLOYER, EDUCATION_TIME, JOB, AVAILABILITY, CATEGORY_EVENT, CALL, EMPLOYMENT_TIME, CALL_CONTACT, LOCATION, TRACK_INFO_MUSIC, SUNSET, MUTUAL_FRIENDS, UNDERGRAD, REMINDER_LOCATION, ATTENDEE_EVENT, MESSAGE_CONTACT, REMINDER_AMOUNT, DATE_TIME_EVENT, DETAILS_NEWS, EDUCATION_DEGREE, MAJOR, CONTACT_METHOD, LIFE_EVENT_TIME, LYRICS_MUSIC, AIRQUALITY, LANGUAGE, GENDER, GROUP
IN:SEND: MESSAGE
IN:SET: UNAVAILABLE, RSVP_YES, AVAILABLE, DEFAULT_PROVIDER_MUSIC, RSVP_INTERESTED, DEFAULT_PROVIDER_CALLING, RSVP_NO
IN:DELETE: REMINDER, ALARM, TIMER, PLAYLIST_MUSIC
.....
IN:DISPREFER:
IN:HELP: REMINDER
IN:FOLLOW: MUSIC

Request: does this recipe have dairy?
Answer: [IN:IS_TRUE_RECIPES [SL:METHOD_RECIPES recipe] [SL:RECIPES_INCLUDED_INGREDIENT dairy]]

(c) MTOP

poi	poi_type	address	distance	traffic_info
pizza my heart	pizza restaurant	528 anton ct	5 miles	moderate traffic
whole foods	grocery store	819 alma st	4 miles	heavy traffic
hotel keen	rest stop	578 arbol dr	3 miles	no traffic
safeway	grocery store	452 arcadia pl	4 miles	no traffic
midtown shopping center	shopping center	338 alester ave	3 miles	no traffic
round table	pizza restaurant	113 anton ct	4 miles	heavy traffic
mandarin roots	chinese restaurant	271 springer street	3 miles	moderate traffic

Request: where is the closest grocery_store?
Response: we are 4 miles away from whole_foods and from safeway: which one do you prefer?

(d) KVRET

Experiment

Method	Objective
NAT (Gu et al., 2017)	$\prod_i P(y_i\|X)$
INAT (Lee et al., 2018)	$\prod_i P(y_i^t\|X, \hat{Y}^{t-1})$
Levenshtein (Gu et al., 2019)	Imitate an expert policy to delete and insert
L2R (Wu et al., 2016)	$\prod_i P(y_i\|X, y_{i-1}, ..., y_1)$
XLNet (Yang et al., 2019)	$\prod_i P(y_{z_i}\|X, y_{z_{i-1}}, ..., y_{z_1})$
Bidirectional (Zhang et al., 2018)	$\prod_i P(y_i\|X, y_{i+1}, ..., y_n) \cdot \prod_i P(y_i\|X, y_{i-1}, ..., y_1, C^{R2L})$
Progressive (Tan et al., 2021)	$\prod_i P(y_i^t\|X, Y^{t-1}, y_{i-1}^t, ..., y_1^t)$
Ours	$\prod_i P(y_i^t\|X, \hat{Y}^{t-1}, y_{i-1}^t, ..., y_1^t)$

(a) Objectives

		Methods	WebQSP F1	MTOP Acc.	MTOP EM	KVRET BLEU	KVRET Micro F1	AVG
Xa	NAR	INAT (Lee et al., 2018)	—.01—					0.00
Xb	NAR	Levenshtein (Gu et al., 2019)	22.54	—.01—				4.51
Xc	AR	CASR B / Dec	4.08	63.02	12.93	7.18	10.17	19.48
Xd	AR	Bidirectional (Zhang et al., 2018)	68.43	61.33	57.84	8.98	54.77	50.27
Xe	AR	Progressive (Tan et al., 2021)	72.05	80.80	77.04	15.49	64.54	61.98
B0	AR	Finetune	70.81	82.49	78.64	**18.33**	67.40	63.53
B2		CASR	**74.81**	**85.61**	**81.69**	**19.12**	**69.85**	**66.22**
L2		CASR-L	**77.99**	**87.73**	**84.54**	18.14	**68.80**	**67.44**

(b) CASR vs. Baselines

	t	CASR Variations: Param.E.	Init.	Arch.	WebQSP F1	MTOP Acc.	MTOP EM	KVRET BLEU	KVRET Micro F1	AVG
A0	0	Fine-tune	Restart	SepEnc	70.81	82.49	78.64	18.33	67.40	63.53
B0	0	Fine-tune	Continue	SepEnc						
C0	0	Fine-tune	Restart	UniEnc						
D0	0	Fine-tune	Continue	UniEnc						
E0	0	Adapter	Restart	SepEnc	70.20	77.34	73.10	16.38	65.59	60.52
A1	1	Fine-tune	Restart	SepEnc	73.03	82.92	78.93	18.17	67.65	64.14
B1	1	Fine-tune	Continue	SepEnc	**74.61**	85.07	81.19	18.55	**70.00***	**65.88**
C1	1	Fine-tune	Restart	UniEnc	73.14	82.83	78.80	**18.80**	67.88	64.29
D1	1	Fine-tune	Continue	UniEnc	74.16	84.86	81.24	18.06	69.12	65.49
E1	1	Adapter	Restart	SepEnc	67.92	77.43	73.21	16.60	65.86	60.20
A2	2	Fine-tune	Restart	SepEnc	73.09	82.92	78.93	18.26	67.70	64.18
B2	2	Fine-tune	Continue	SepEnc	**74.81**	**85.61**	**81.69**	**19.12**	**69.85**	**66.22**
C2	2	Fine-tune	Restart	UniEnc	73.14	82.83	78.77	18.78	68.06	64.32
D2	2	Fine-tune	Continue	UniEnc	74.32	**85.66**	**81.92**	18.34	68.80	65.81
E2	2	Adapter	Restart	SepEnc	68.13	77.45	73.23	16.63	65.66	60.22

(c) Controlled Experiments

Empirical Study

t	Task / Metric	WebQSP F1	MTOP Acc	MTOP Match	KVRET Bleu	KVRET Micro F1	AVG	Δ_t
0	Short	78.80	85.23	83.65	14.05	65.91	65.53	-
	Middle	83.07	84.83	80.51	21.93	71.16	68.30	-
	Long	50.51	77.31	71.62	15.08	65.69	56.04	-
1	Short	80.12	87.21	85.50	12.17	65.93	66.19	0.66
	Middle	86.41	87.19	83.01	19.81	71.51	69.59	1.29
	Long	57.52	80.71	74.95	16.33	69.48	59.80	3.76
2	Short	80.12	87.96	86.25	12.32	64.52	66.23	0.05
	Middle	86.41	87.53	83.34	19.77	71.79	69.77	0.18
	Long	58.13	81.26	75.36	17.00	69.18	60.19	0.39

(a) Result on Different Answer Length

t	AVG Difficulty	Ratio
0	82.80	63.54%
1	86.95	18.48%
2	88.60	8.50%
3	89.83	4.25%
4	90.86	2.51%
5 (failed)	92.21	2.72%

(b) The Average Difficulty and Ratio (sum to 1) of Each Correct-solving CASR step

Difficulty=$(r-1)\cdot(c-1)\cdot(h-1)$

Metric	WebQSP F1	MTOP Acc	MTOP Match	KVRET Bleu	KVRET Micro F1	AVG
Sparse	51.88	78.52	72.72	16.70	64.17	56.80
Middle	85.84	87.53	83.17	20.18	75.09	70.36
Dense	85.78	89.08	87.63	20.26	85.80	73.71

(c) The Relationship between Cross-Attention Density and the Quality of Attended Previous Predictions

t	Task	WebQSP	MTOP	KVRET	AVG
1	Input	0.10%	0.11%	0.40%	0.20%
	Previous Prediction	0.34%	0.33%	0.84%	0.50%
2	Input	0.09%↓	0.10%↓	0.39%↓	0.19%↓
	Previous Prediction	0.61%↑	0.43%↑	1.01%↑	0.68%↑

(d) Changes in Attention Density during the CASR steps

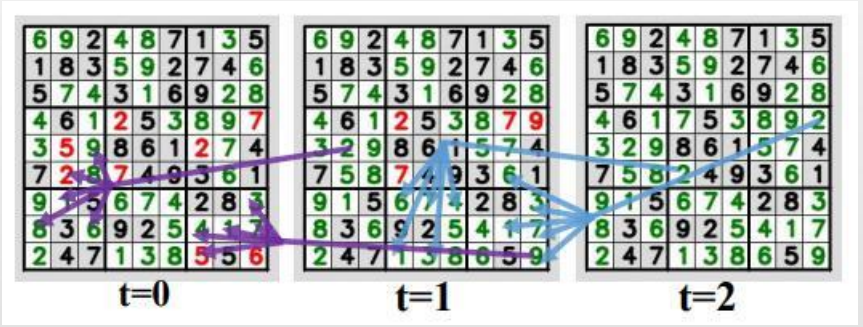

(e) Visualization of CASR's Cross-Attention

Case Study

Task&Input X	CASR Predictions $\hat{Y}^t$			Ground Truth Y
WebQSP input:task: webqsp ; what school did sir ernest rutherford go to? ; structured knowledge: Ernest Rutherford: m.02m7r | m.068kst common.topic.notable_types m.01m9 | (omitted to save space) | m.018y0w base.act	t=0 (AND (JOIN common.topic.notable_types m.01nf) (JOIN (R education.education.institution) (JOIN (R people.person.education) m.02m7r)))	t=1 (JOIN (R education.education.institution) (JOIN (R people.person.education) m.02m7r))	t=2 (JOIN (R education.education.institution) (JOIN (R people.person.education) m.02m7r))	(JOIN (R education.education.institution) (JOIN (R people.person.education) m.02m7r))
MTOP input:task: mtop ; does this recipe have dairy ; structured knowledge: IN:GET: MESSAGE, WEATHER,... | (omitted to save space) | IN:FOLLOW: MUSIC	t=0 [IN:IS_TRUE_RECIPES [SL:RECIPES_INCLUDED_INGREDIENT dairy]]	t=1 [IN:IS_TRUE_RECIPES [SL:METHOD_RECIPES recipe] [SL:RECIPES_INCLUDED_INGREDIENT dairy]]	t=2 [IN:IS_TRUE_RECIPES [SL:METHOD_RECIPES recipe] [SL:RECIPES_INCLUDED_INGREDIENT dairy]]	[IN:IS_TRUE_RECIPES [SL:METHOD_RECIPES recipe] [SL:RECIPES_INCLUDED_INGREDIENT dairy]]
KVRET input:task: kvret ; where is the closest grocery_store ; structured knowledge: col : poi | poi_type | address | distance | traffic_info row 1 : pizza my heart | pizza restaurant | (omitted to save space) | 3 miles | moderate traffic ; context:	t=0 whole_foods is 4 miles away at 819_alma_st.	t=1 the closest grocery_store is whole_foods at 819_alma_st.	t=2 the closest grocery_store is whole_foods which is 4 miles away at 819_alma_st.	we are 4 miles away from whole_foods and from safeway: which one do you prefer?
Sudoku	t=0	t=1	t=2,3,4	

Method	WebQSP	MTOP	KVRET
INAT	(JOIN (R education.education.institution) (AND (JOIN education.education.degree).02 (JOIN educationeducation.JO3.)www.wm.educationma02date)	and_IN: the_]ACT [ ] the for:]]IN for	error iserror__ theerror__ at is is_error___ is___.. would would is is_error is is
Levenshtein	(JOIN (R education.education.institution) (JOIN (R people.person.education) m.03xsv3))	_	what is for you.
CASR /Dec	(JOIN (R education.education.institution) (JOIN (R people.person.education) m.02mjmr))	[IN:GET_RECIPES [SL:RECIPES_DISH chicken]]	the nearest grocery_store is whole_foods at 819_alma_st. would you like directions there?
Bidirectional	(JOIN (R education.education.institution) (JOIN (R people.person.education) m.02m7r))	[IN:GET_RECIPES [SL:RECIPES_INCLUDED_INGREDIENT dairy]]	safeway is 4 miles away.
Progressive	(AND (JOIN common.topic.notable_types m.01nf) (JOIN (R education.education.institution) (JOIN (R people.person.education) m.02m7r)))	[IN:IS_TRUE_RECIPES [SL:RECIPES_INCLUDED_INGREDIENT dairy]]	the closest grocery_store is safeway.
CASR-L	(JOIN (R education.education.institution) (JOIN (R people.person.education) m.02m7r))	[IN:IS_TRUE_RECIPES [SL:RECIPES_INCLUDED_INGREDIENT dairy]]	the closest grocery_store is whole_foods at 819 alma st.

Connecting to OpenAI ChatGPT

Prompt:

KVRET is a benchmark for table conversation. The input consists of a table and an NL query, and the output is an NL response corresponding to the dialog.

For each query, you should give the answer without any explanation or any additional information.
When a suggested answer is given (may not be correct), you should repeat it if it's correct, or correct it if it's wrong.
For example,

Query: <query1>
Suggested Answer: None
Answer: <answer1>

Query: <query2>
Suggested Answer: <suggested_answer2>
Answer:

t	BLEU	Micro F1
0	0.43	25.66
1	0.66	29.41
2	1.35	28.37

one-shot learning
using gpt-3.5-turbo
with temperature=1.0
on KVRET