Open Rule Induction

Wanyun Cui, Xingran Chen
Shanghai University of Finance and Economics

NeurIPS 2021

John McCarthy's criticism of some expert systems in 1980s

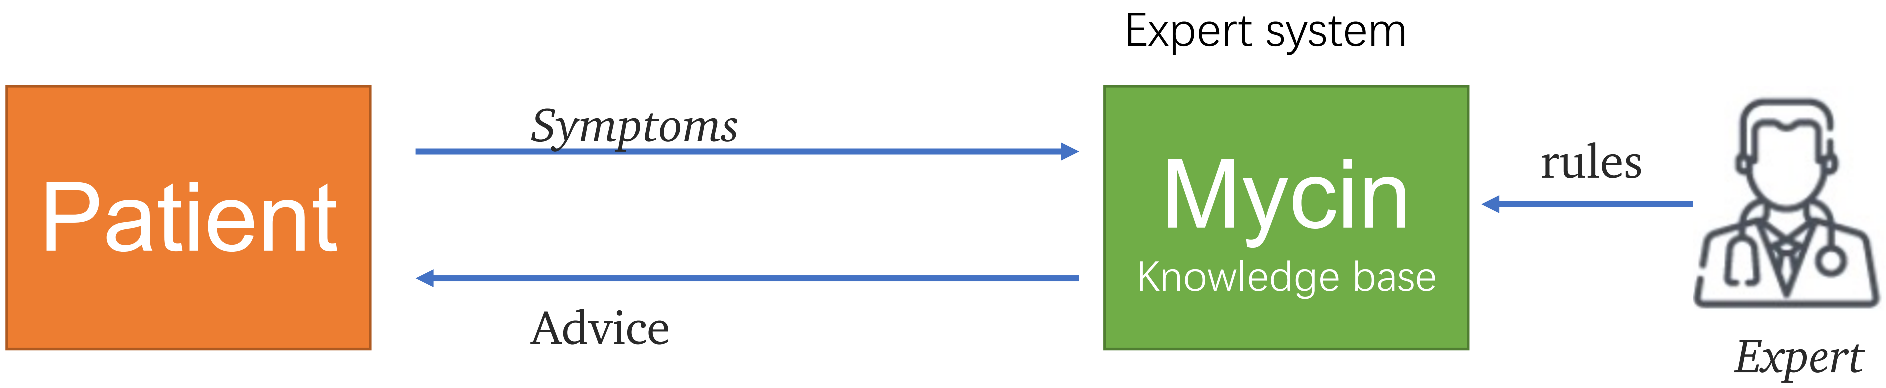

John McCarthy's criticism of some expert systems in 1980s

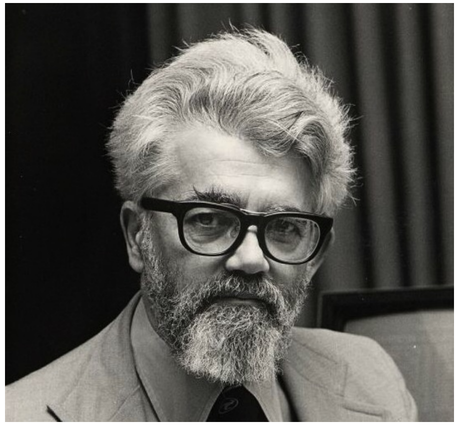

Cholerae Vibrio

two weeks of tetracycline

Expert system

Mycin

Knowledge base

rules

Expert

This would most likely kill off all the bacteria, but …
by then the patient would already be dead.

John McCarthy's criticism of some expert systems in 1980s

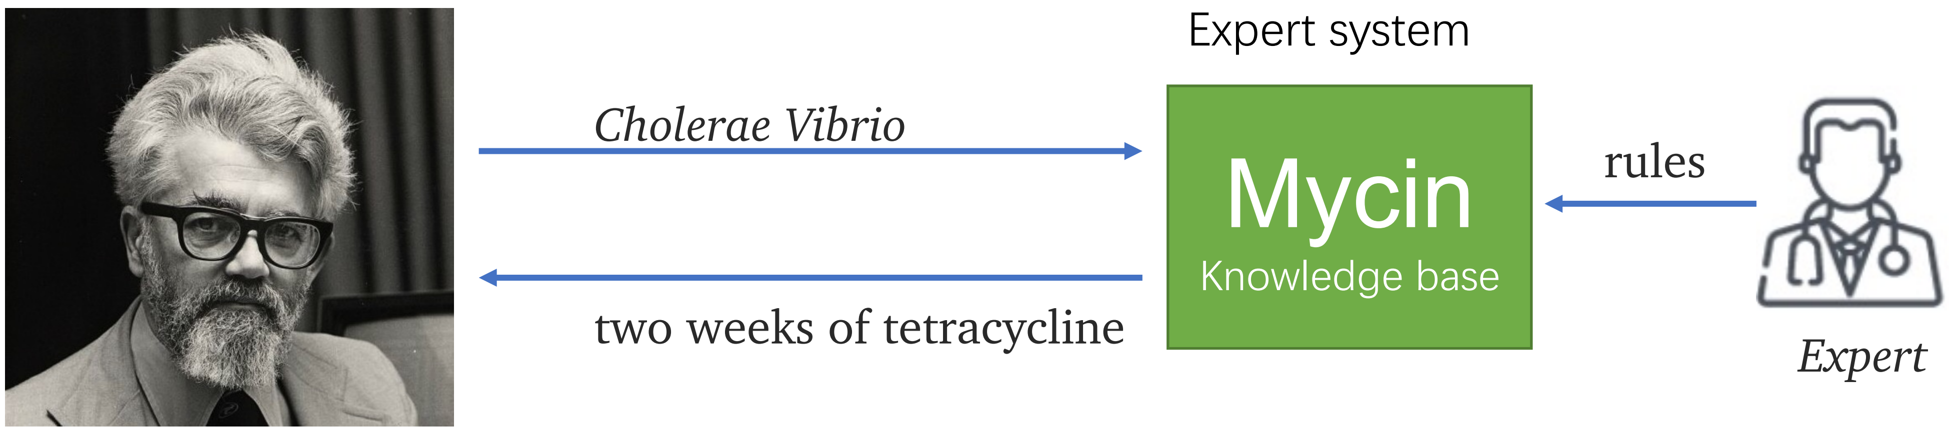

He point out that, this is because the expert rules lacked common sense and knowledge.

- The real-world knowledge is much more complex than the annotated rules.

How to make the rules more expressive?

- Mine rules from large-scale knowledge bases.

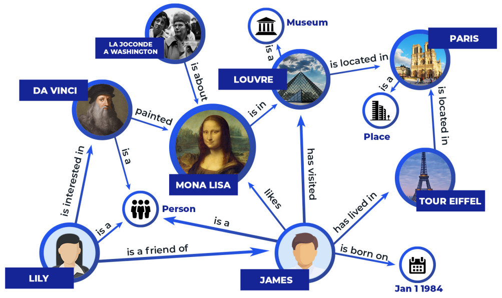

Find rules

?x livesIn ?y → ?x wasBornIn ?y

How to make the rules more expressive?

- Mine rules from large-scale knowledge bases.

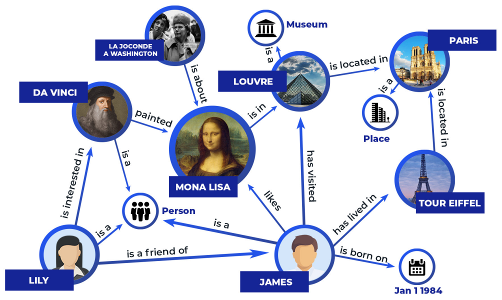

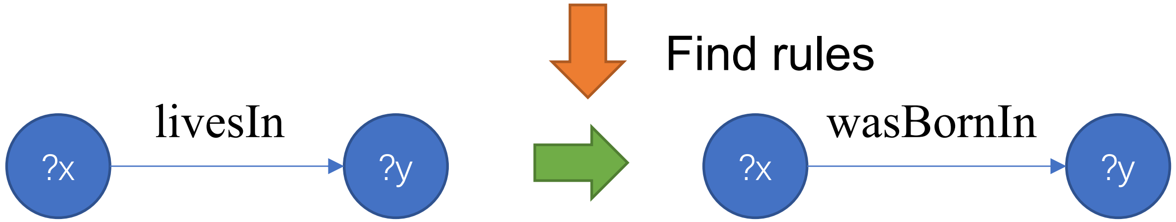

- Rules are limited to the existing entities and relations of the KB.

How to make the rules more expressive?

- Mine rules from large-scale knowledge bases.

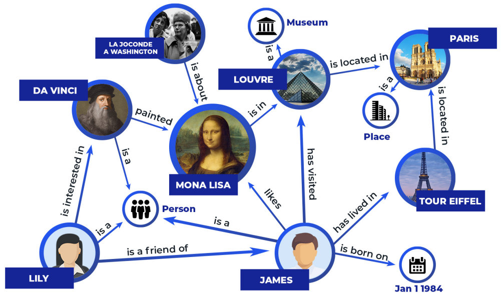

- Pre-trained language models as knowledge bases.

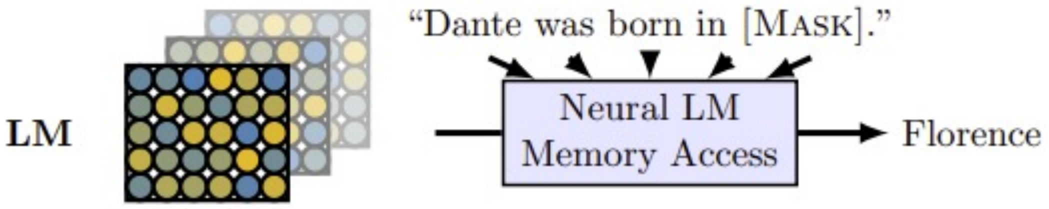

x worked at the brokerage Morgan Stanley for about 11 years until 2005, when he and some Morgan Stanley colleagues quit and later founded the hedge fund y

How to make the rules more expressive?

- Mine rules from large-scale knowledge bases.

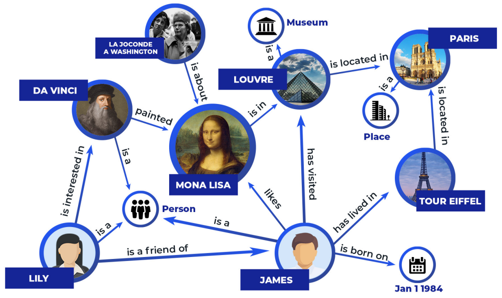

- Pre-trained language models as knowledge bases.

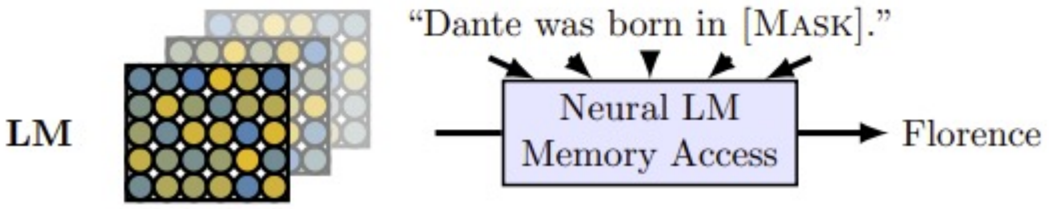

Find rules

x worked at the brokerage Morgan Stanley for about 11 years until 2005, when he and some Morgan Stanley colleagues quit and later founded the hedge fund y

...

x is the founder and CEO of y

x serves on the board of y

x is currently a venture capitalist at y

Previous rule mining methods for KBs/LMs

For knowledge bases

- Knowledge

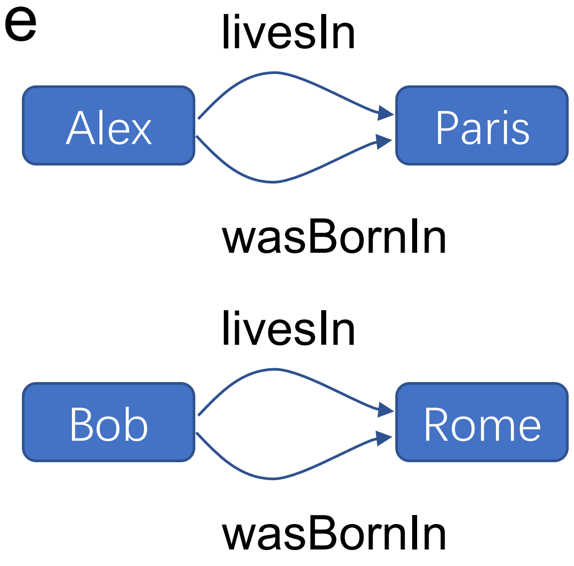

- Rule

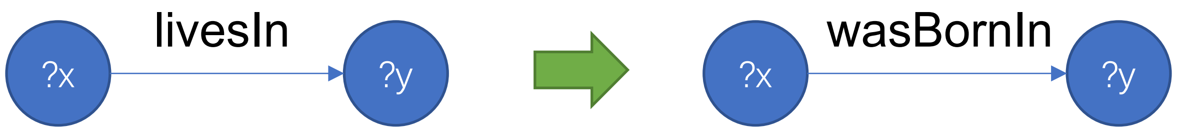

Discover commonalities of a group of entities from the data

Previous rule mining methods for KBs/LMs

For knowledge bases

- Knowledge

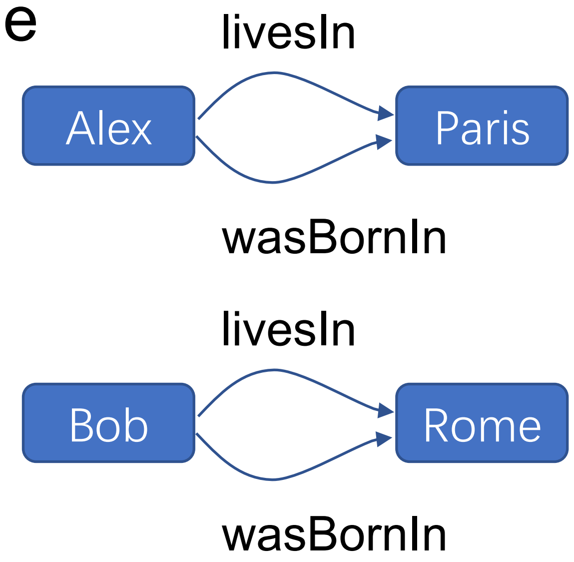

- Rule

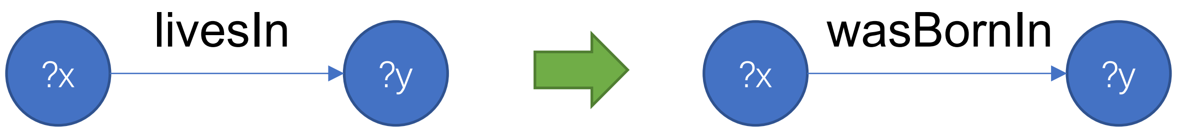

Discover commonalities of a group of entities from the data

For language models (e.g. COMET)

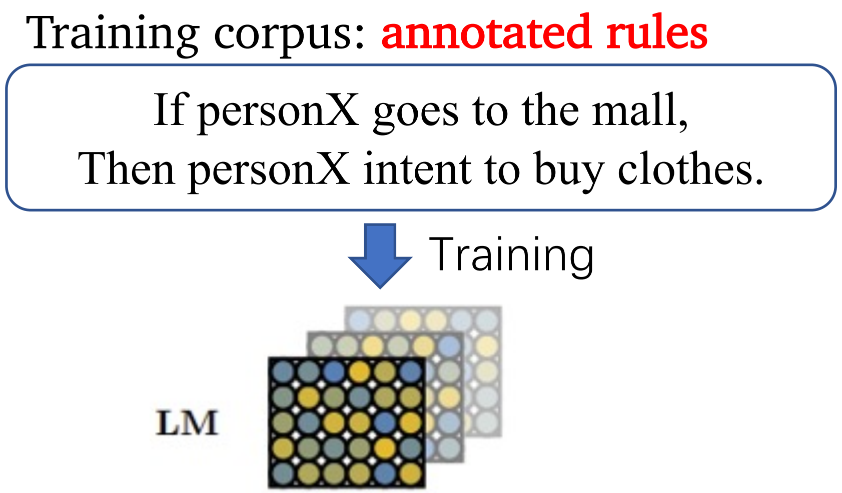

- The patterns of learned rules are constrained by the annotated rules.
- The limits the expressiveness of the generated rules.

Learning rules from rules.

Open Rule Induction: Overview

- Discover commonalities as traditional KB-based methods.
- Let the language model "speak" the commonalities without annotations.

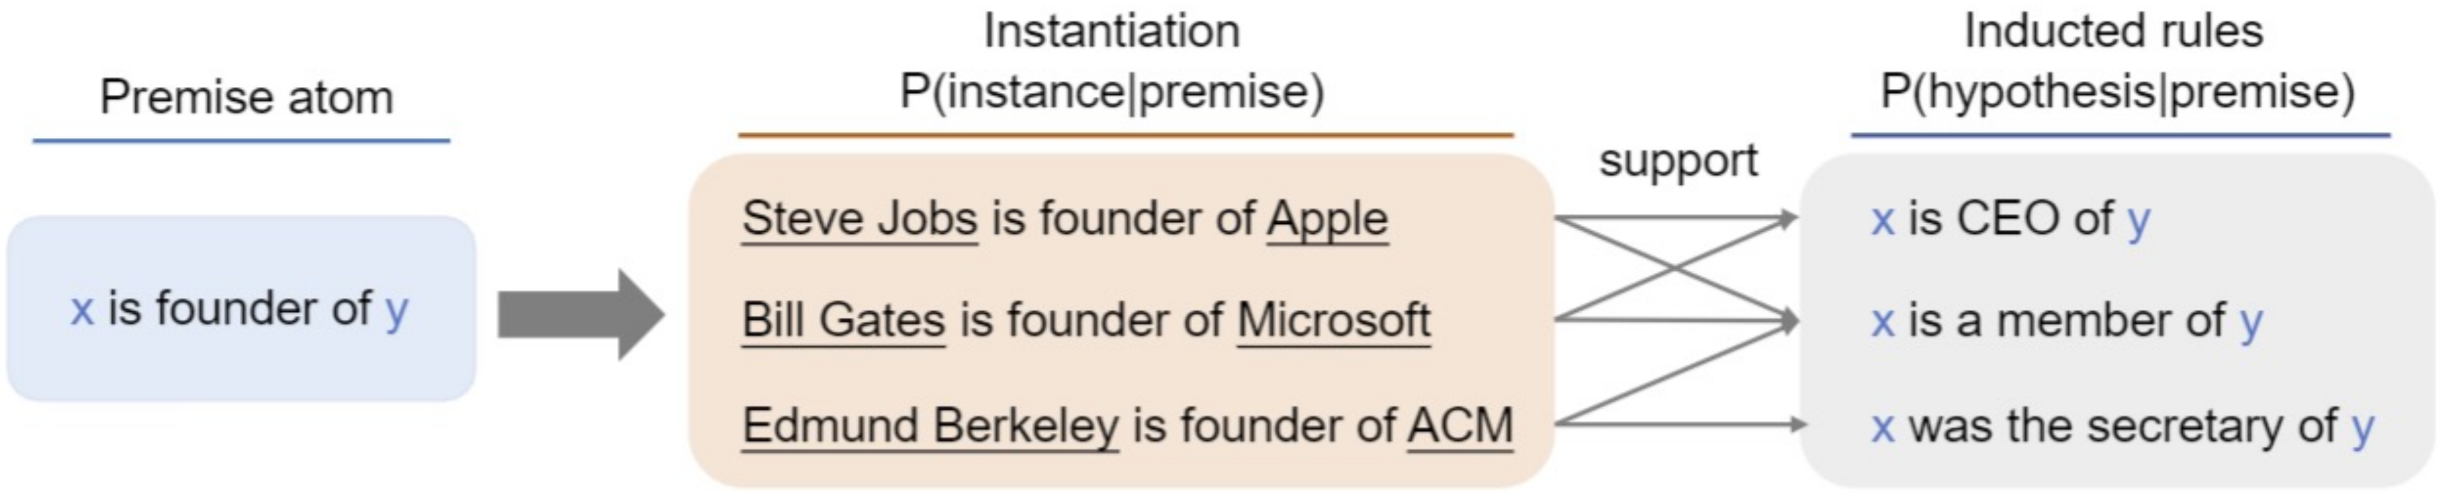

Open Rule Induction: Problem Formulation

- Given a premise atom (x, r_p, y) and k, find top-k of r_h, w.r.t. $P(r_h|r_p)$

$$P(r_h|r_p) = \sum_{ins} P(r_h|ins, r_p) P(ins|r_p)$$

- One key observation is that given ins, r_p and r_h are independent.

$$P(r_h|r_p) = \sum_{ins} \underbrace{P(r_h|ins)}_{Appcalibility} \underbrace{P(ins|r_p)}_{Instantiation}$$

Open Rule Induction: Problem Formulation

- $P(r_h|r_p) = \sum_{ins} \underbrace{P(r_h|ins)}_{Appcalibility} \underbrace{P(ins|r_p)}_{Instantiation}$

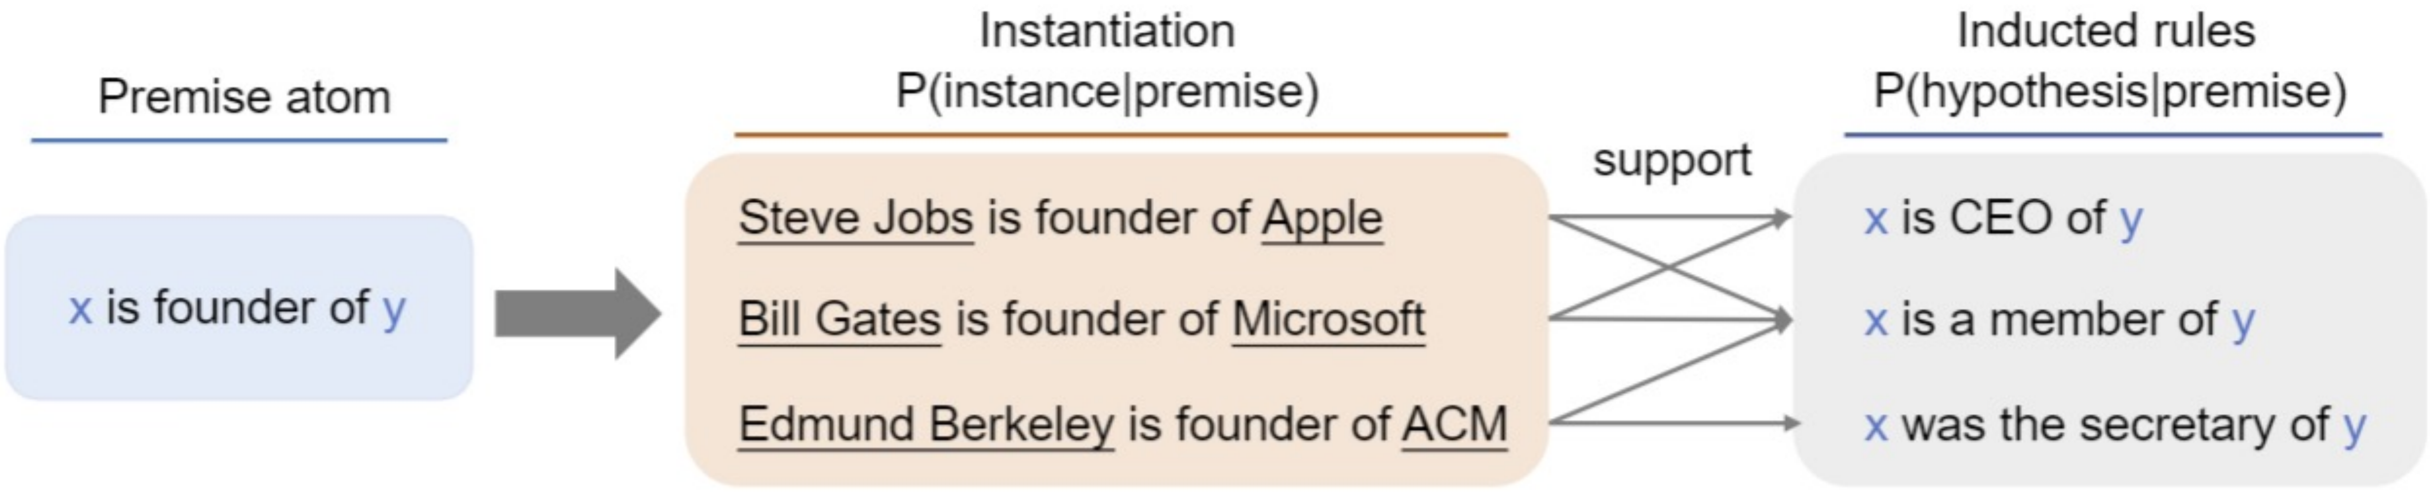

How to Compute $P(ins|r_p)$ and $P(r_h|ins)$ using LMs?

- With the language model, both probabilities can be computed through the masked language modeling task $P(w_{<mask>}|w_1, \ldots, w_n)$.
- We use different language model for $P(ins|r_p)$ and $P(r_h|ins)$.
- Following this strategy, we use Spacy to annotate entities and continue training language models on Wikipedia and BookCorpus.

x y

Steve Jobs is the founder of **Apple**.

Language Model

<*mask*> is the founder of <*mask*>.

$$P(ins = (x, y)|r_p)$$

r_h

Steve Jobs **is the founder of** Apple.

Language Model

Steve Jobs <mask> Apple.

$$P(r_h|ins = (x, y))$$

Experiments: OpenRule155

- Manual constructed dataset: OpenRule155
  - We collect 121 relations from all relations from 6 relationship extraction datasets: Google-RE, TREx, NYT10, WIKI80, FewRel, SemEval, and 34 relations from Yago2.

Our Dataset	BLEU-1	BLEU-2	BLEU-4	ROUGE-L	METEOR	self-BLEU-2
Prompt	17.77	3.65	0.48	18.65	12.94	86.63
Prompt (fine-tuned	20.95	7.58	0.86	22.37	17.24	82.13
Comet	21.58	8.15	1.04	23.45	5.44	90.78
Orion - STS	44.92	20.24	1.21	49.72	39.68	89.84
Orion - train $P(ins\|r_p)$	15.85	3.11	0.00	32.91	13.19	90.29
Orion - train $P(r_h\|ins)$	19.17	3.05	0.07	34.99	10.30	**83.54**
Orion	**45.41**	**21.29**	**1.30**	**50.37**	**40.41**	90.94

More Examples of Generated Rules

Table 3: Effect of complex rule induction. Original sentence of **Case 1**: *[X]'s emergence from international isolation has been marked through improved and expanded relations with other nations such as [Y], France, Japan, Sweden, and India.* **Case 2**: *His guitar work on the title track is credited as what first drew [X] to him, who two years later invited allman to join him as part of [Y].*

Case	Orion	Comet
Case 1	[X] has a long history of military cooperation with [Y]. [X] is the largest exporter of oil to [Y]. [X]'s economy is heavily dependent on [Y]. [X]'s foreign policy is based on its close relationship with [Y]. [X] has been the largest exporter of uranium to [Y].	<Causes>: personx. <HasProperty>: happy. <MadeUpOf>: happy. <isAfter>: happy. <isBefore>: happy.
Case 2	[X], guitarist and singer of [Y]. [X] and his band [Y]. [X] has been a fan of [Y]. [X] was a fan of [Y]. [X] was a fan of the band [Y].	<Causes>: talented. <HasProperty>: talented. <MadeUpOf>: talented. <isAfter>: talented. <isBefore>: persony.

Application: Relation Extraction

- We evaluate inducted rules on relation extraction tasks.
- We use ExpBERT to add our explanation.

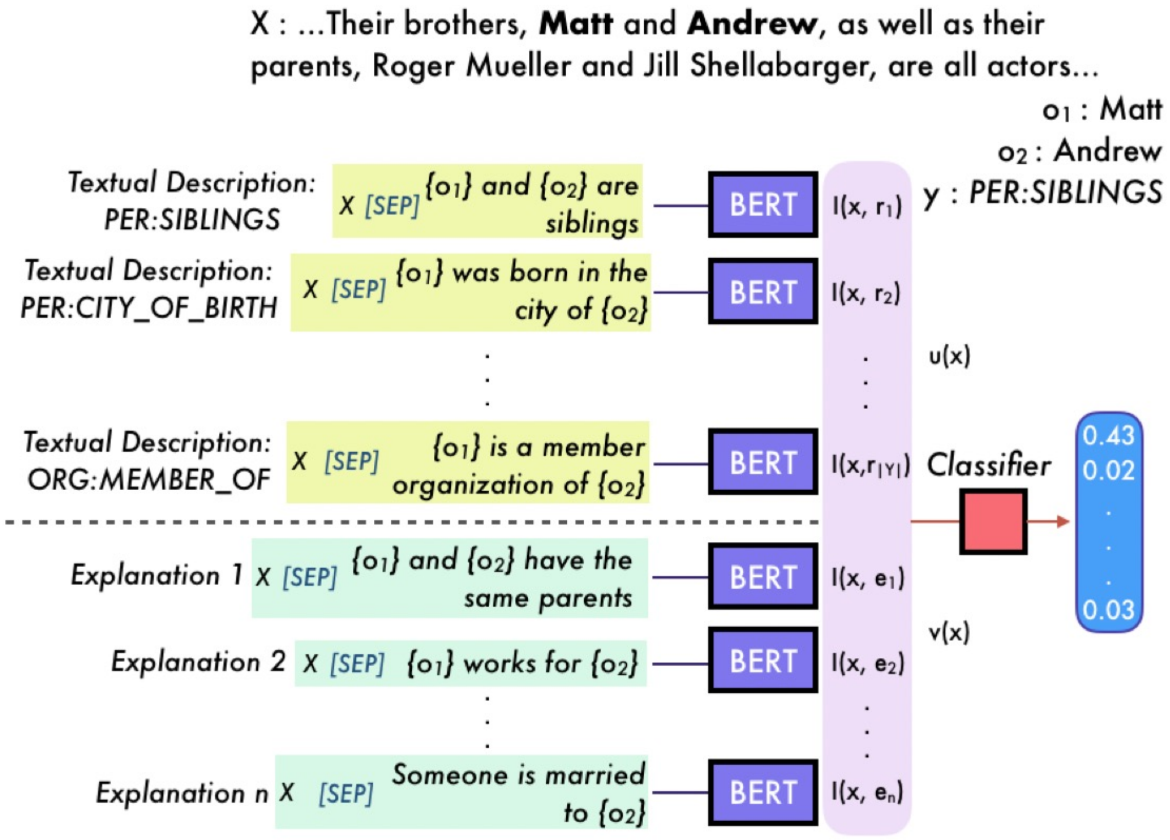

Figure 2: Overview of our approach. Explanations as well as textual descriptions of relations are interpreted using BERT for a given x to produce a representation which form inputs to our classifier.

Application: Relation Extraction

	Spouse	Disease
BERT	46.43 ± 0.84	40.20 ± 2.43
ExpBERT + annotated rules	76.04 ± 0.47	56.92 ± 0.82
ExpBERT + inducted open rules	$\mathbf{76.05} \pm \mathbf{0.52}$	$\mathbf{57.68} \pm \mathbf{1.34}$

- By adding these rules to BERT, the effect can be significantly improved.
  - Our rule induction method is unsupervised.

Application: Relation Extraction

	Spouse	Disease
BERT	46.43 ± 0.84	40.20 ± 2.43
ExpBERT + annotated rules	76.04 ± 0.47	56.92 ± 0.82
ExpBERT + inducted open rules	$\mathbf{76.05} \pm \mathbf{0.52}$	$\mathbf{57.68} \pm \mathbf{1.34}$

- By adding these rules to BERT, the effect can be significantly improved.
- These automatically inducted rules even slightly outperformed the manually annotated rules.

Application: Error Identification in LMs

- Some rules that defy human commonsense are incorrectly inducted.
- This is actually due to the bias of the language model.

Inducted rule: [X] is the politician of [Y]. → [X] was the founder and president of [Y].

Identified error: LMs assume that politician is always founder and president.

Reason: The training corpus description of politician has a disproportionate number of founder and president entities that general members.

Thank you!