

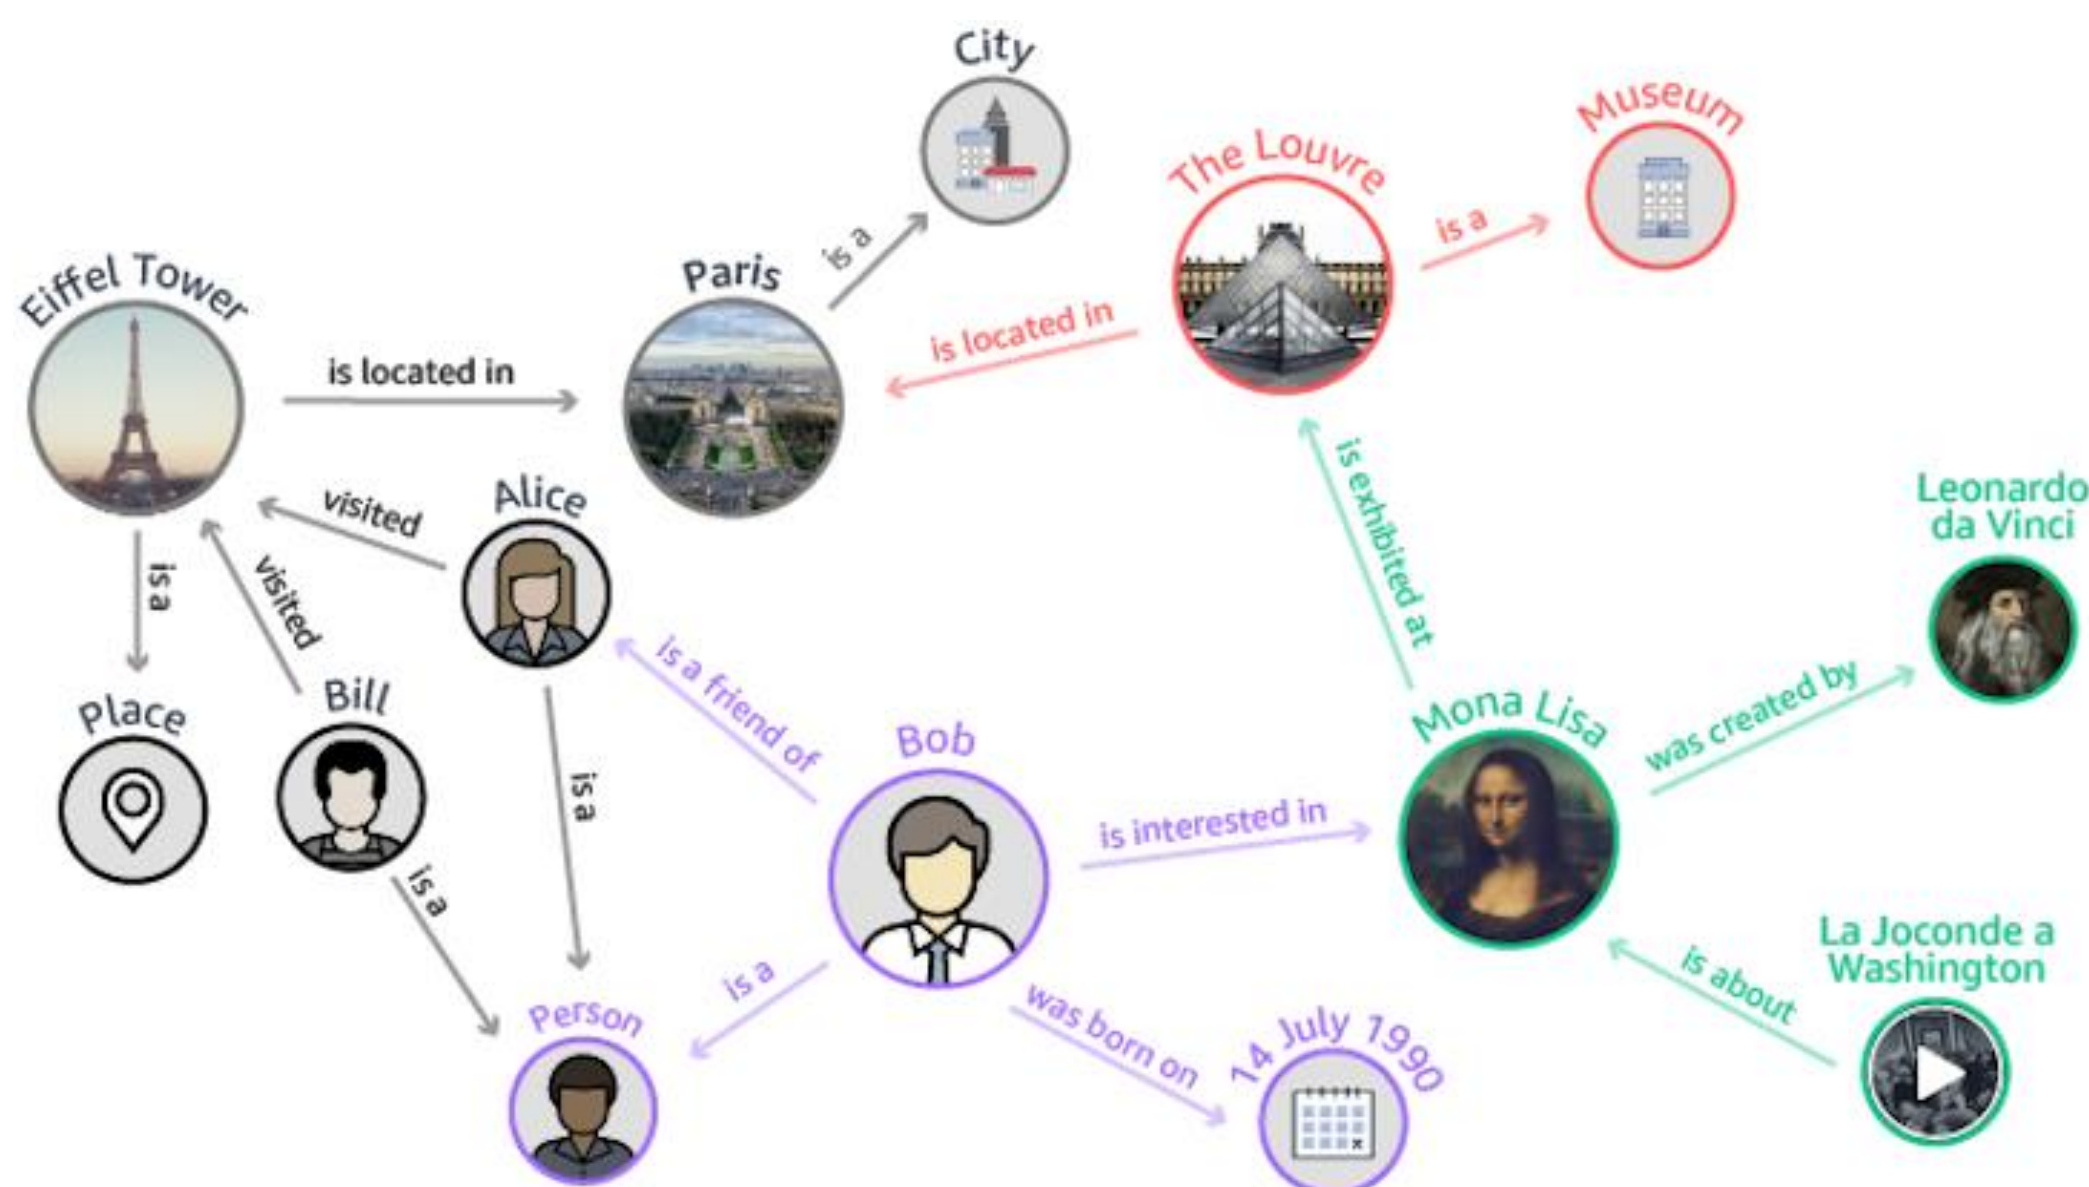
Latent Execution-Guided Reasoning for Multi-Hop Question Answering on Knowledge Graphs

Hongyu Ren, Hanjun Dai, Bo Dai, Xinyun Chen, Jure Leskovec, Denny Zhou

Stanford University, Google Brain, UC Berkeley

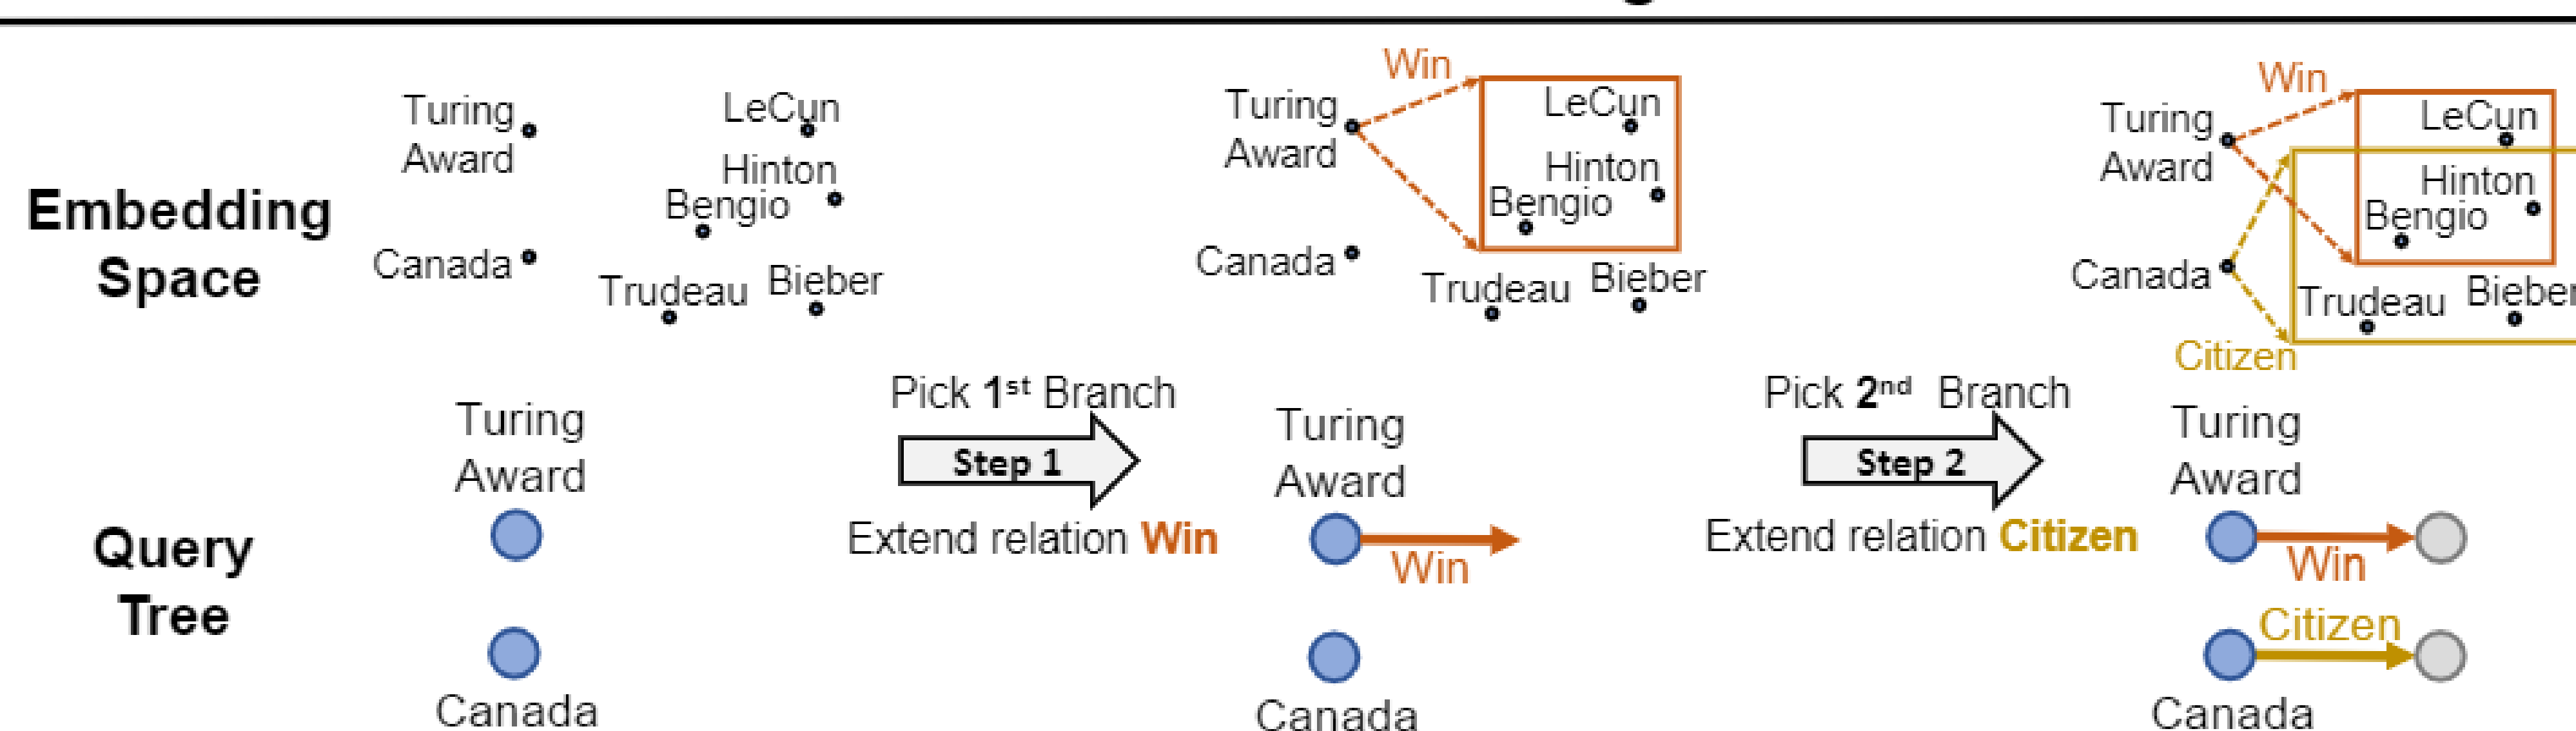


Multi-hop Reasoning on KGs



LEGO: Latent Execution-Guided Reasoning

Question: Who are Canadian Turing Award winners?

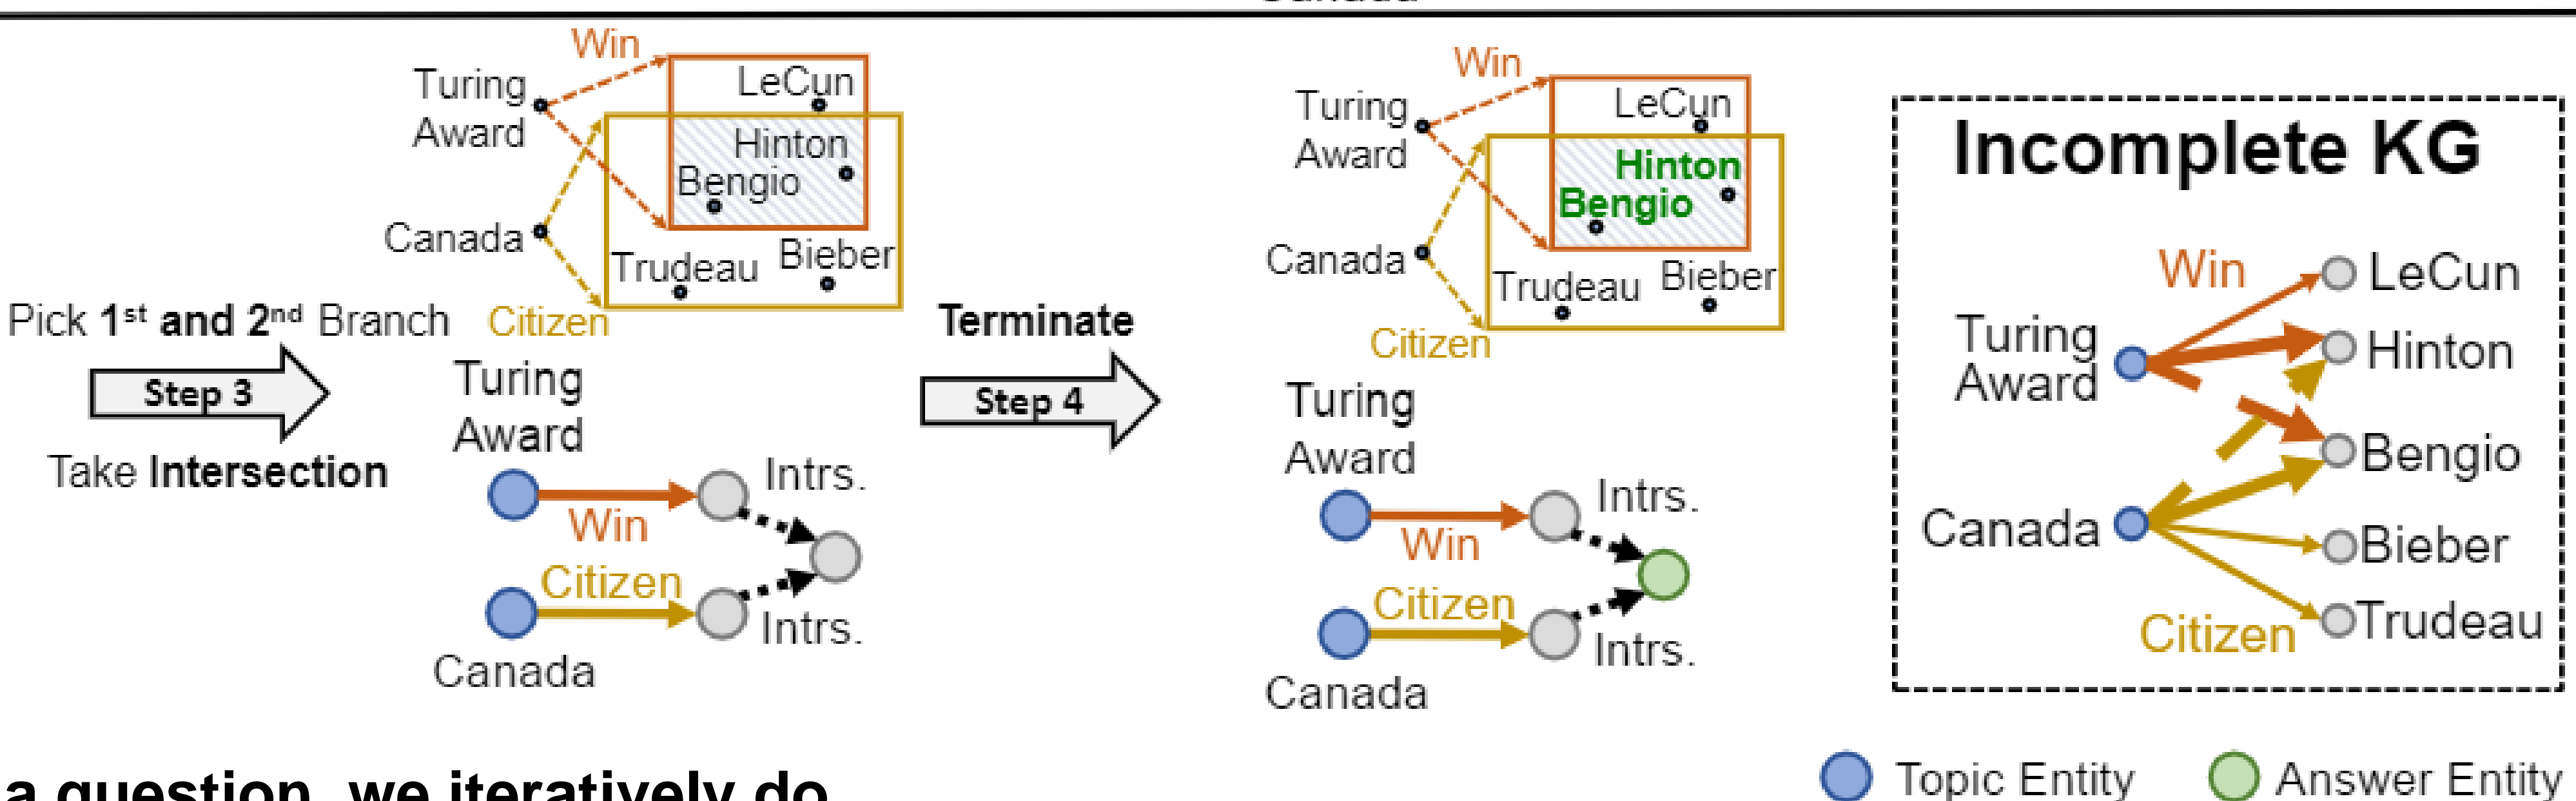
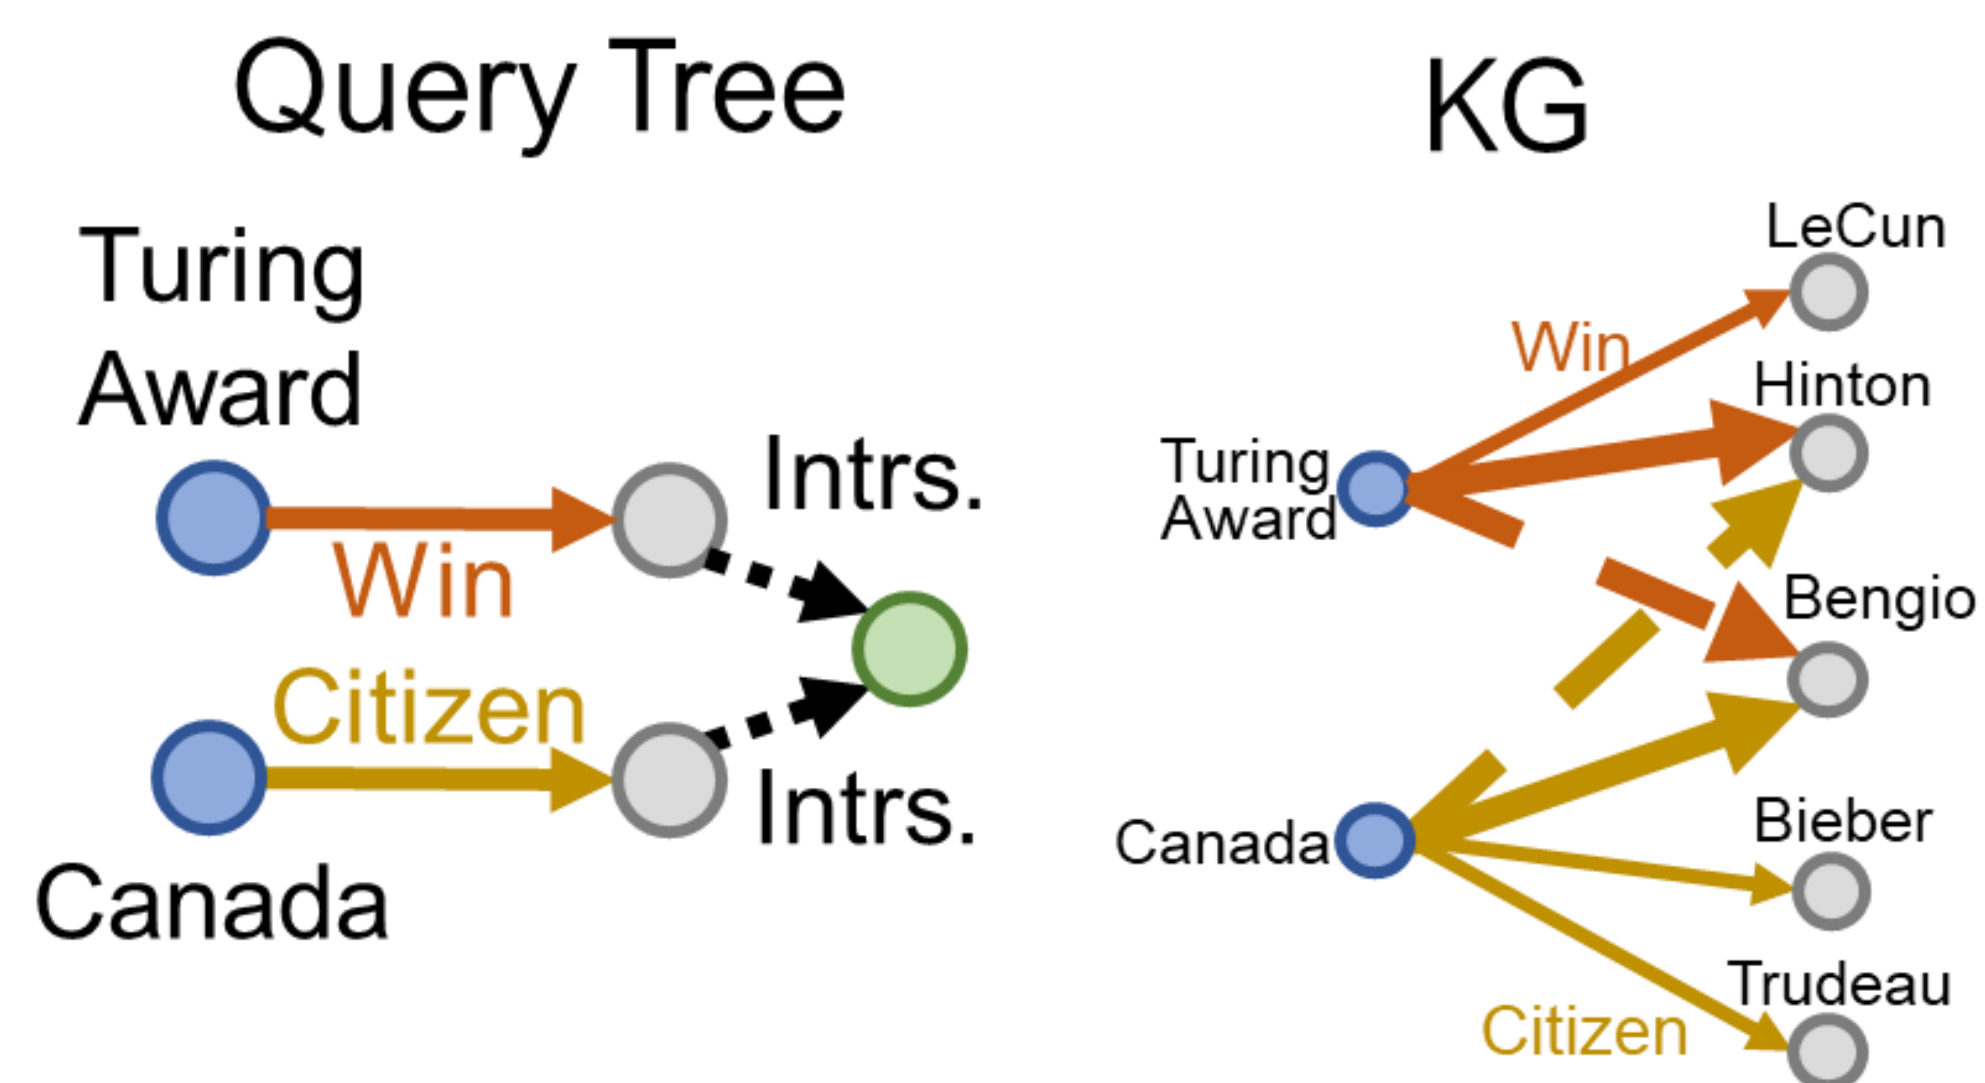


Answering complex questions on KGs

Who are Canadian Turing Award winners?

Latent query tree (**unobserved**): $q = V_?$

$Win(TuringAward, V_?) \wedge Citizen(WorldCup, V_?)$



Given a question, we iteratively do

Execution-guided synthesis: synthesize the query tree based on the current box embedding

Embedding-based query execution: execute the new reasoning step in the embedding space

Until we have a valid query tree, and the synthesis terminates.

Challenges:

1. Query tree is unobserved.
2. KGs are incomplete with missing edges.

Hits@1	1-hop	2-hop	3-hop	All
Semantic Parsing (GT)	63.3	45.8	45.3	51.5
Query2box (GT)	70.8	62.1	66.4	66.6
-----	-----	-----	-----	-----
PullNet	65.1	52.1	59.7	59.2
EmbedKGQA	70.6	54.3	53.5	60.2
Lego (ours)	69.3	57.8	63.8	63.8