

New Frontiers of Knowledge Graph Reasoning: Recent Advances and Future Trends

Part IV: Neural Reasoning Beyond Entities and Relations

Yangqiu Song and Jiaxin Bai

Department of CSE, HKUST



KnowComp Group
Understanding the World by Computational Knowledge

Roadmap

Part I: Knowledge Graph Reasoning: Basic Concepts and Techniques

Part II: Recent Advance #1: Neural Reasoning for Natural Language Queries

Part III: Recent Advance #2: Neural Reasoning for Logical Queries

Part IV: Recent Advance #3: Neural Reasoning Beyond Entities and Relations

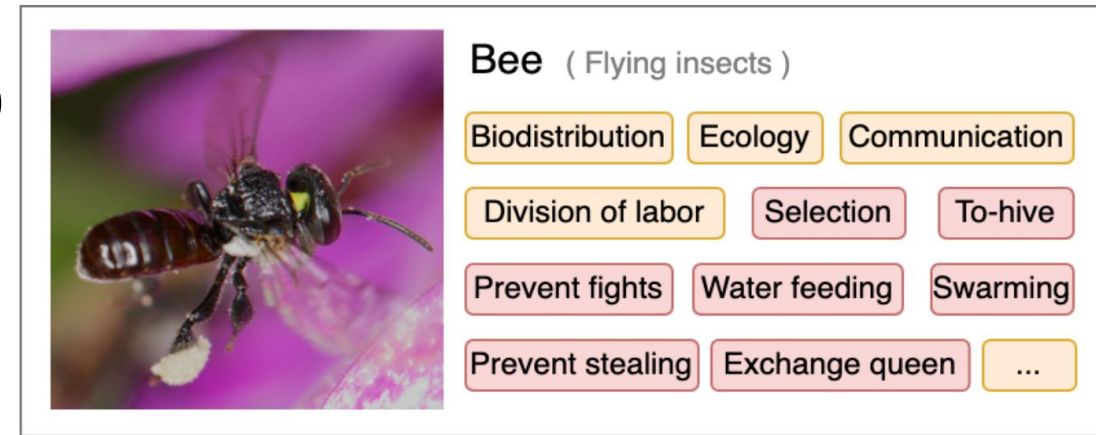
Part V: Recent Advance #4: LLM+KGR

Part VI: Open Challenges and Future Directions



Scope of Knowledge

- **Awareness of facts**
 - “Knowing that”: the truth of propositions
 - E.g., the population of Singapore is 5.92 million (2023)
- A possession of **practical skills**
 - “Knowing how”: understanding how to perform certain actions; procedural knowledge
 - E.g., How to hive bees?
- An experiential acquaintance: **familiarity with individuals and situations**
 - “Knowing by acquaintance”: directly perceiving an object, being familiar with it, or otherwise coming into contact with it.
 - E.g., by eating Durian, one becomes acquainted with the taste of it



The aspects in orange boxes are aspects that convey the knowledge of “know-what” or “know-why”, while those in red boxes convey the knowledge of “know-how”.

Source: Kuaipedia

Commonsense Knowledge in AI

- "Commonsense knowledge includes the basic facts about **events** (including actions) and **their effects**, ... **how it is obtained**, **facts about beliefs and desires**. It also includes the basic **facts about objects and properties**. " – John McCarthy
- "While to the average person the term 'commonsense' is regarded as synonymous with '**good judgement**', the AI community it is used in a technical sense to refer to the **millions of basic facts and understandings possessed by most people**." --ConceptNet
 - "Such knowledge is typically omitted from social communications", e.g.,
 - **If you forget someone's birthday, they may be unhappy with you.**
- Meanwhile, it is not invariably true
 - "a person is larger than a dog"



Source: Zorba - The World's Largest Dog Ever Lived

The Three-dimensional Development of Knowledge Graphs in Computer Science

The development of the Web

In the past years the development of knowledge graph has been largely changed (or revolutionized) by deep learning and NLP

Web 3.0

Web 2.0

Web 1.0

G#1: Entity-
Based KGs

G#2: Text-
Rich KGs

G#3: Dual
Neural KGs

Generations of KGs

Xin Luna Dong. Generations of Knowledge Graphs: The crazy ideas and the business impact. *VLDB*, 2023. Invited paper for **VLDB Women in Database Research Award**. [[Paper](#)][[Talk](#)]

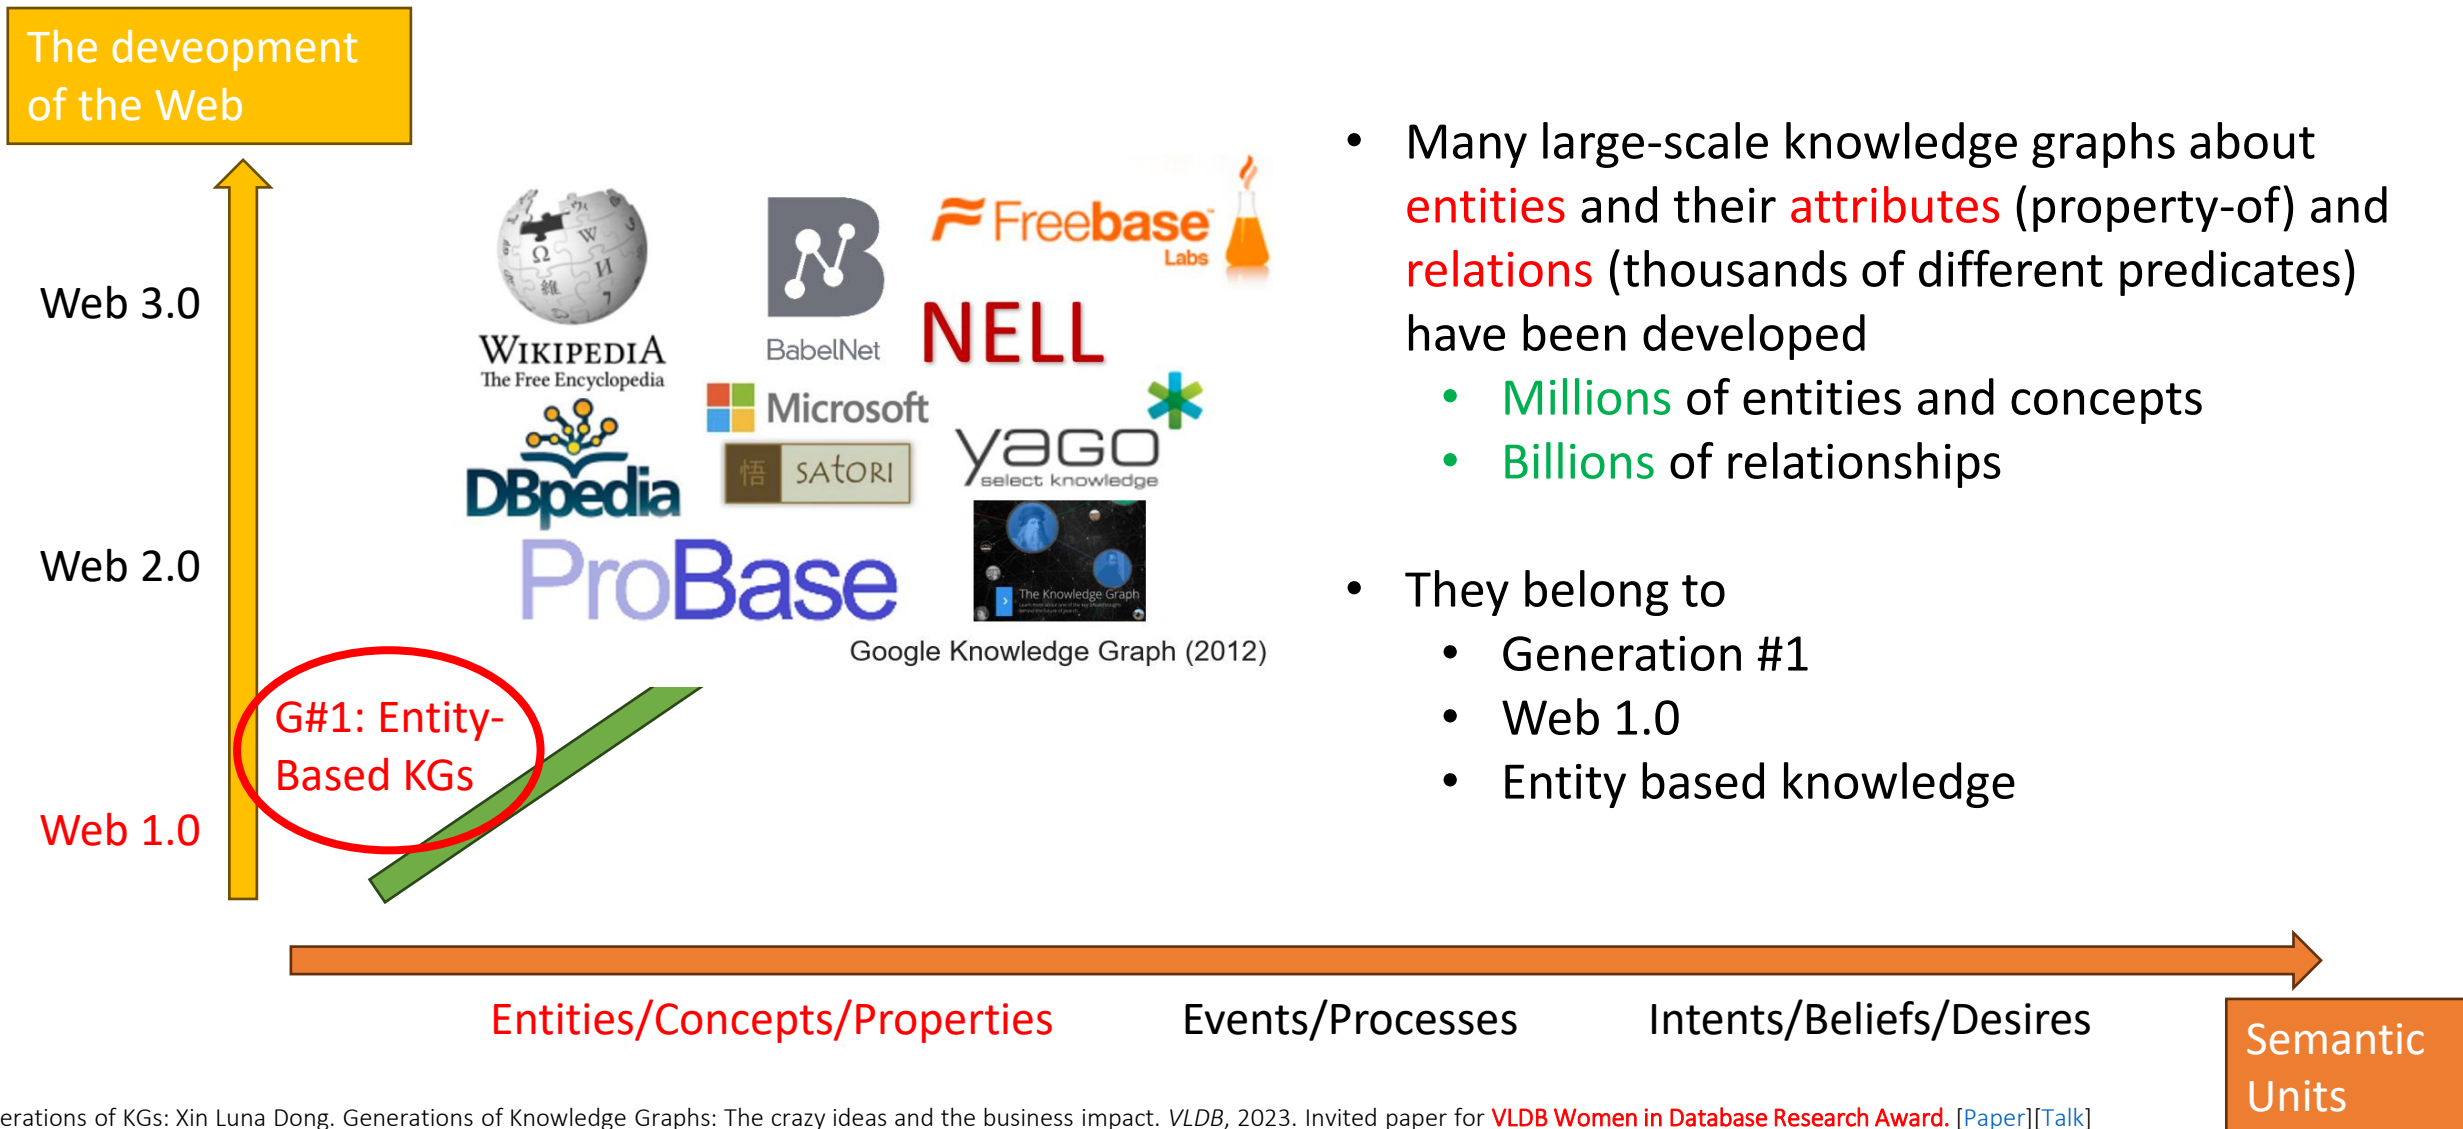
Entities/Concepts/Properties

Events/Processes

Intents/Beliefs/Desires

Semantic
Units

The Three-dimensional Development of Knowledge Graphs in Computer Science





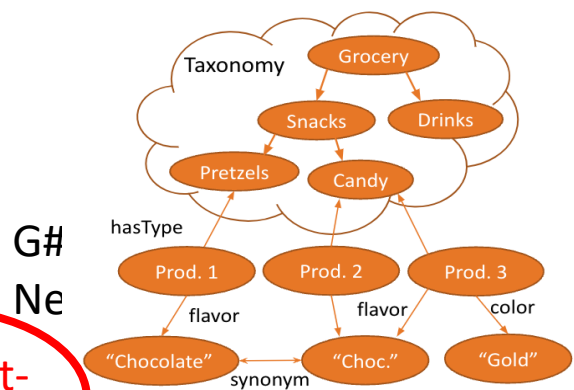
The Three-dimensional Development of Knowledge Graphs in Computer Science



Web 3.0

Web 2.0

Web 1.0

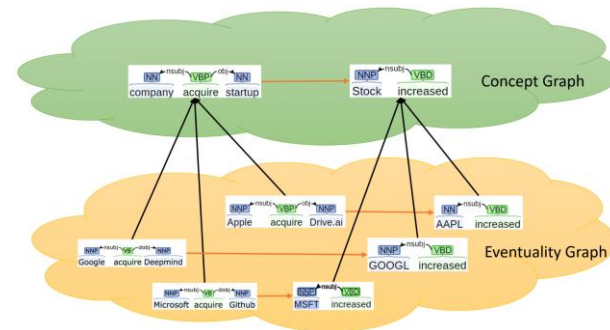


G#2: Text-Rich KGs

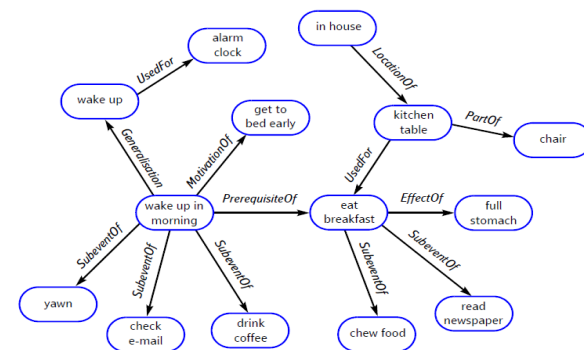
G#1: Entity-Based KGs

Product Graph

- Generation #2
- Web 1.0
- Entity (product) based knowledge



- ASER/Knowlywood
 - Generation #2
 - Web 1.0
 - Event based knowledge



- ConceptNet
 - Generation #2
 - Web 1.0
 - Entity and event based knowledge

Entities/Concepts/Properties

Events/Processes

Intents/Beliefs/Desires

Semantic Units

The Three-dimensional Development of Knowledge Graphs in Computer Science

The development of the Web

Web 3.0

Web 2.0

Web 1.0

G#1: Entity-
Based KGs

G#2: Text-
Rich KGs

G#3: Dual
Neural KGs

Generations of KGs



“Language models as knowledge bases”
Neural KGs/Memory Networks/NGDBs


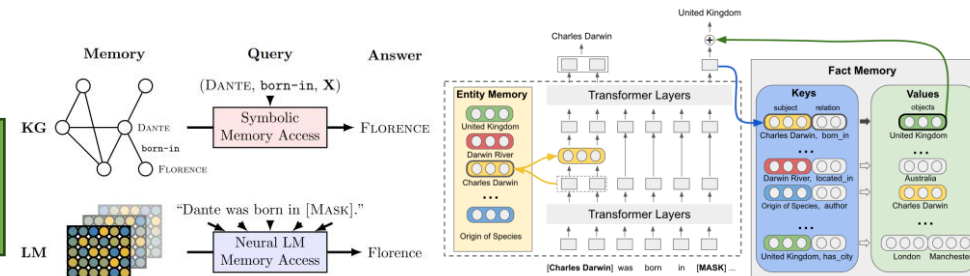


Figure of NGDBs

- Zihao Wang's part
 - Generation #3
 - Web 1.0
 - Entity (product) based knowledge



- Jiaxin Bai's part
 - Generation #3
 - Web 1.0
 - Event and Intention based knowledge

Entities/Concepts/Properties

Events/Processes

Intents/Beliefs/Desires

Semantic
Units

The Three-dimensional Development of Knowledge Graphs in Computer Science

The development of the Web

Web 3.0

Web 2.0

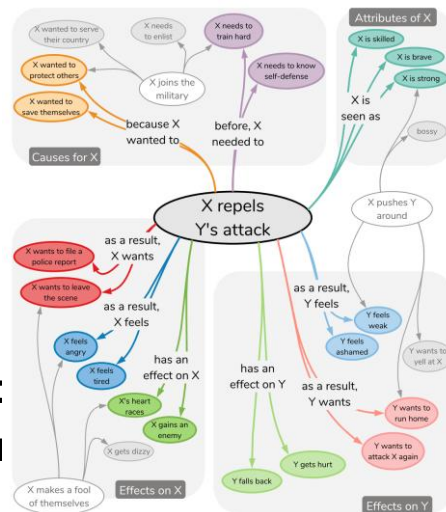
Web 1.0

G#1: Entity-
Based KGs

G#2: Text-
Rich KGs

G#3: Neu

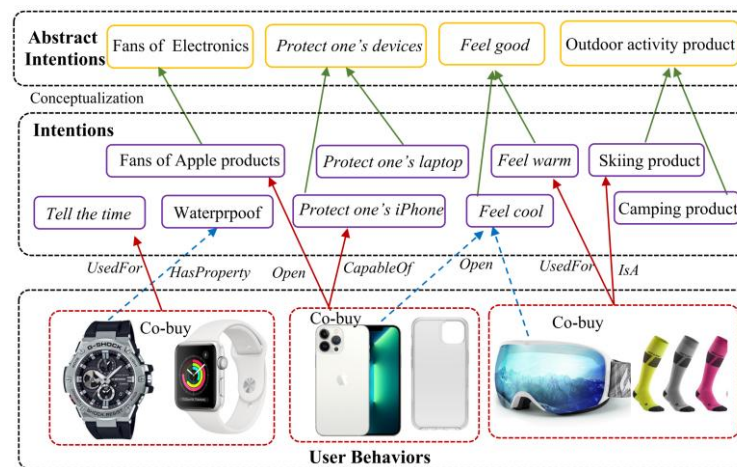
ATOMIC



ATOMIC

- Generation #2
- Web 1.0
- Intention and event (cause-effect) knowledge

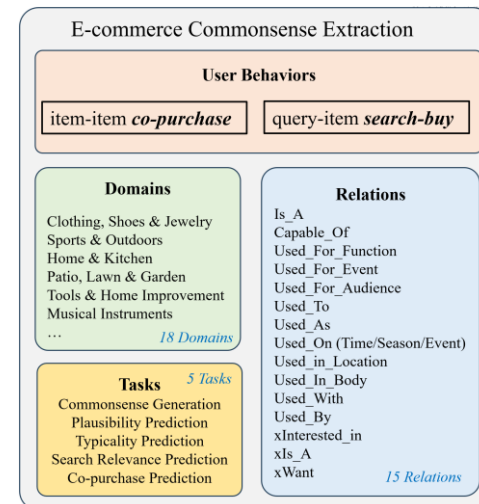
FolkScope



FolkScope

- Generation #2
- Web 2.0
- Intention knowledge

amazon COSMO



COSMO

- Generation #2
- Web 2.0
- Intention knowledge

Entities/Concepts/Properties

Events/Processes

Intentions/Beliefs/Desires

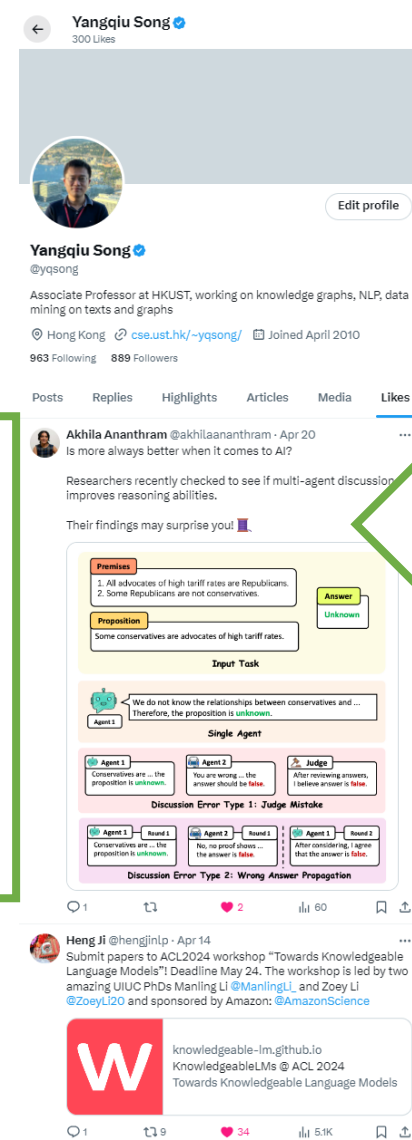
Semantic
Units

Why do we want to go beyond entities in Web2.0?

- User centric and user generated content (UGC)
 - Explicit incentives :
 - e.g., user reviews, posting blogs, tweets, photos, videos, etc.
 - Implicit incentives :
 - user behaviors, e.g., following, browsing, clicking, searching, upvoting, etc.



Why do I **post** this tweet?
Sell my own work, or promote my student?



Why do I **like** this post? I like the content or promote my own work?

Why do we want to go beyond entities in Web2.0?

“I post this tweet: Sell my own work or promote my student?” Inference involves:

- **1. Theory of Mind**

- i.e., the development of **knowledge** that **others have beliefs, desires, and intentions** that are different from one's own
- Possessing a functional theory of mind is crucial for success in everyday human social interactions

- What makes us take actions?

- **Beliefs** and **desires** are mediated by **intentions** which in turn controls human's **actions (or speech)** (Kashima et al., 1998)
- Intentions are implicit




Figure taken from
Andreas (2022)

<https://en.wikipedia.org/wiki/Intention>

https://en.wikipedia.org/wiki/Theory_of_mind

Jacob Andreas: Language Models as Agent Models. EMNLP (Findings) 2022: 5769-5779

Kashima, Yoshihisa, Allison McKintyre, and Paul Clifford. "The Category of the Mind: Folk Psychology of Belief, Desire, and Intention." Asian Journal Of Social Psychology 1, no. 3 (December 1998): 289–313.


Why do we want to go beyond entities in Web2.0?

“I post this tweet: My student gets more visible or not?”

Inference involves:

- **2. System II Processing**


- We need to equip machine learning systems with “slow, logical, sequential, conscious, linguistic, algorithmic, planning, reasoning”
- Particularly, such a system requires the “understanding of how **actions** interact with **changes (of states) in distribution**”
 - “Agents face **non-stationarities**”
 - Conditioned on “different places, times, sensors, actuators, goals, policies, etc”



Why Graphs?

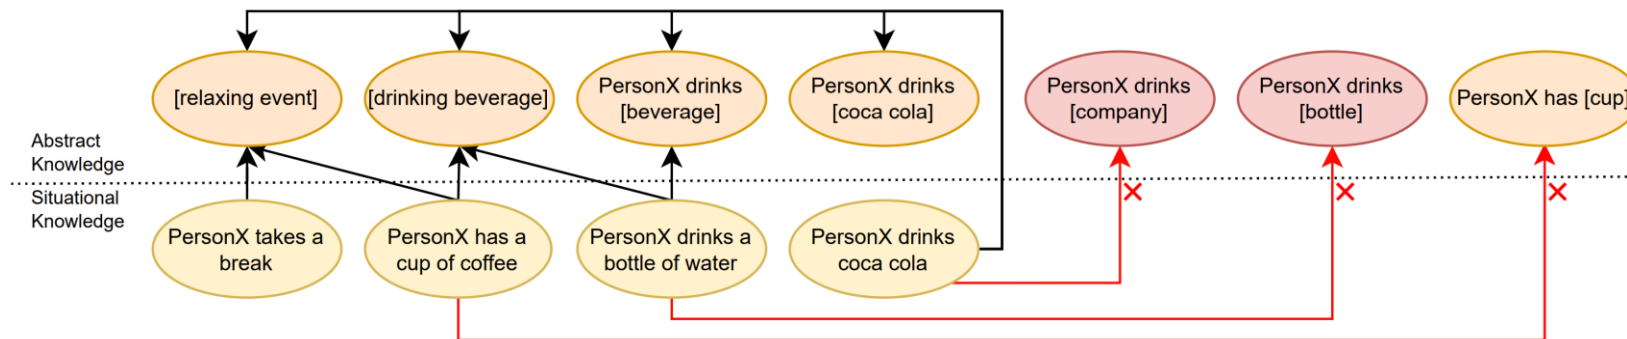
- The K-Line Theory (Minsky, 1980)

- More than **ontology**: categories include substances, properties, relations, states of affairs, and events
- **Mental states** in our memory are also in a hierarchical structure **beyond an ontology**; described as a **K-pyramid**



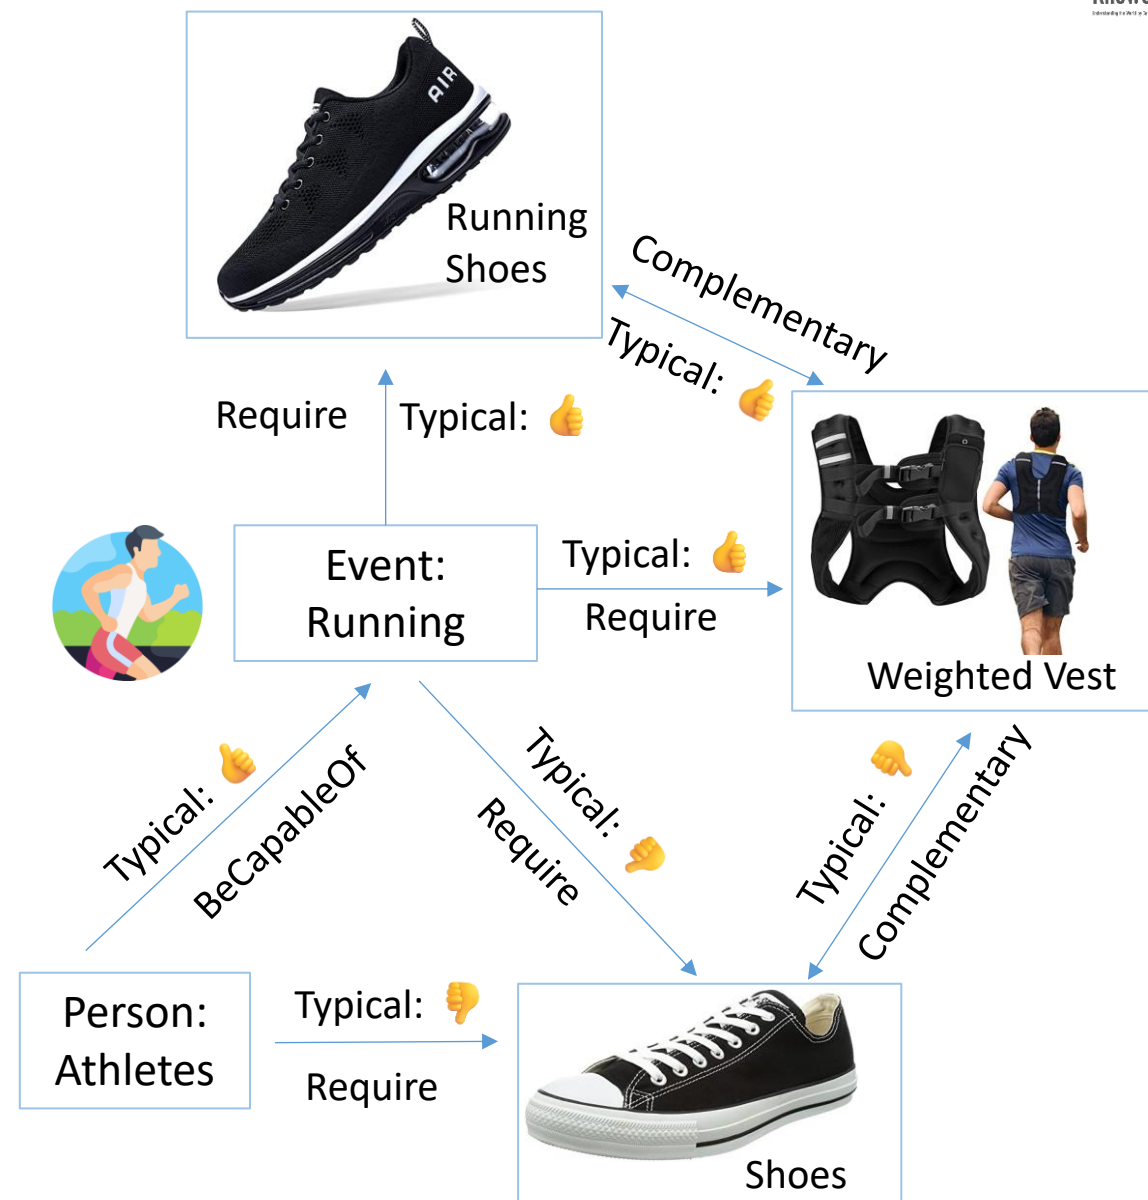
- We need the **right level** and **right perspective** of abstraction

- **Different levels of abstractness**: “PersonX drinks coca cola” → “[drinking beverage],” “[event]”
- **Different perspectives**: “Coca cola” → “[sugary beverage],” “[phosphate containing beverage],” “[iced drink],” not in a strict taxonomy
 - PersonX drinks [iced drink], xReact, refreshed
 - PersonX drinks [sugary beverage], xEffect, gain weight



Why Graphs?

- Sometimes we need concrete, symbolic, and globally referenced knowledge (Edge et al., 2024)
- Ability of commonsense reasoning with high complexity
 - NP-complete problems, e.g., Max-Sat (Chalier et al., 2022), subgraph matching or counting, subset sum, etc.
- The trade-offs between scalability and computational complexity
 - Intentions can be pre-stored and indexed to be more efficiently accessed online



FolkScope and COSMO: User-centric Intention KGs

- AI generated knowledge graph construction framework

Actions on E-commerce Platform

Explainable recommendation

UserU bought [A] because

Query-item relevance

PersonX searches [Q] to

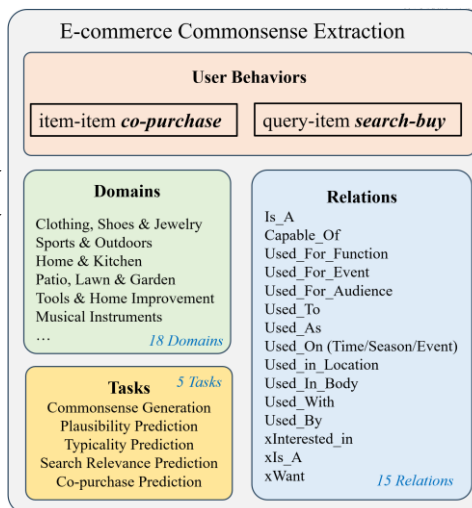
User behavior rationalization

PersonX bought [A] and [B] because

PersonX searched [A] and bought [B] because

Prompts

AI Generator



Human
Feedback

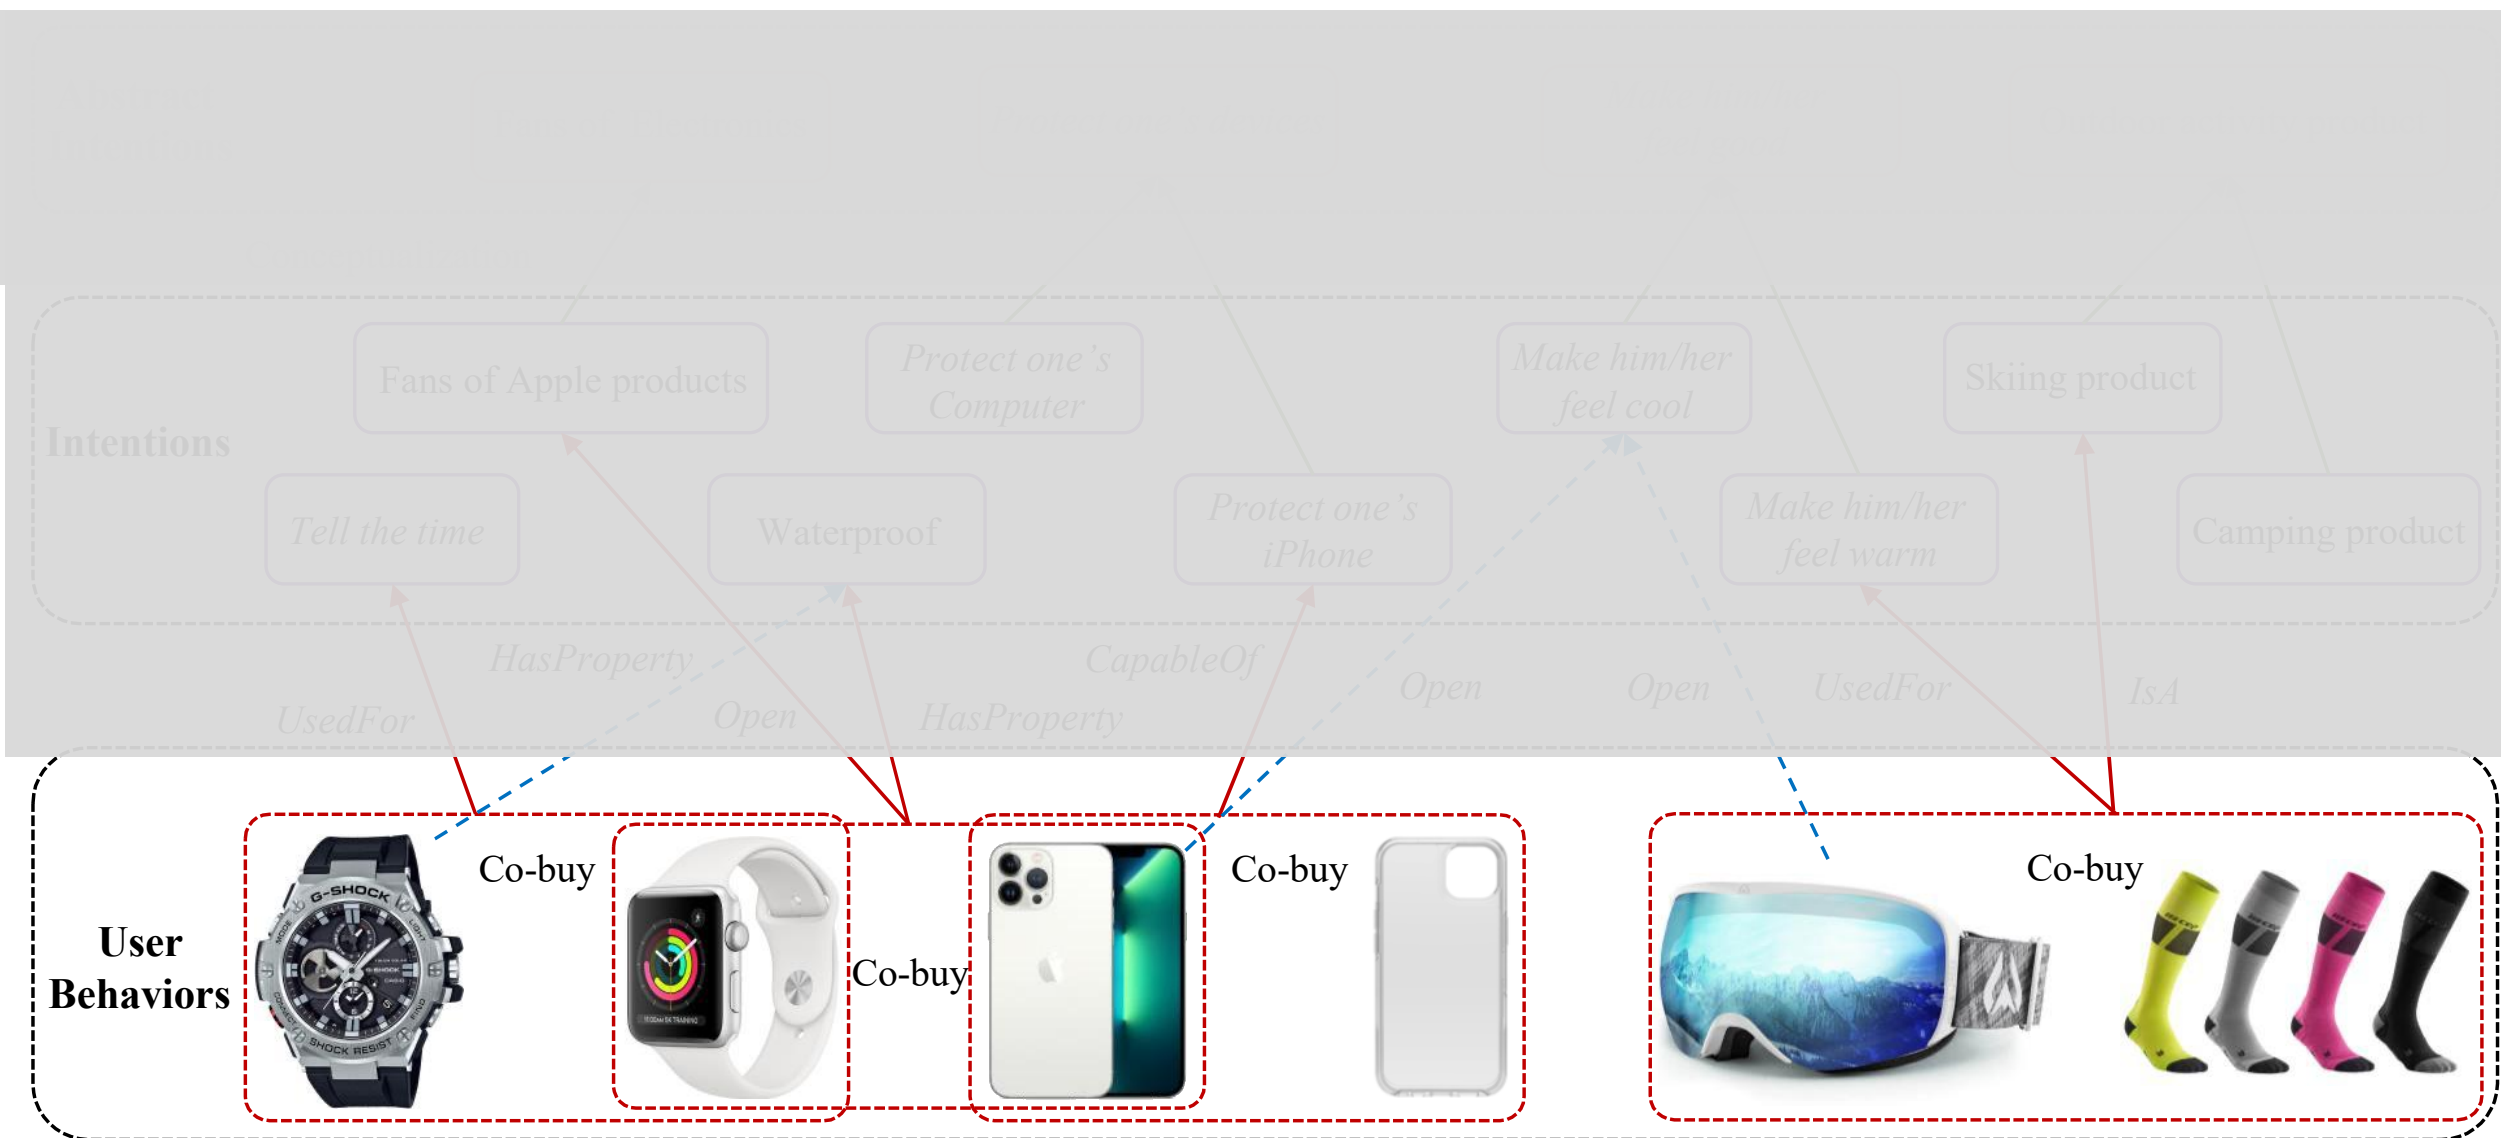
AIG-KG



Intention KG

Pattern mining, Filtering,
Normalization, Conceptualization

FolkScope: Collective Intention KG for Co-Purchases

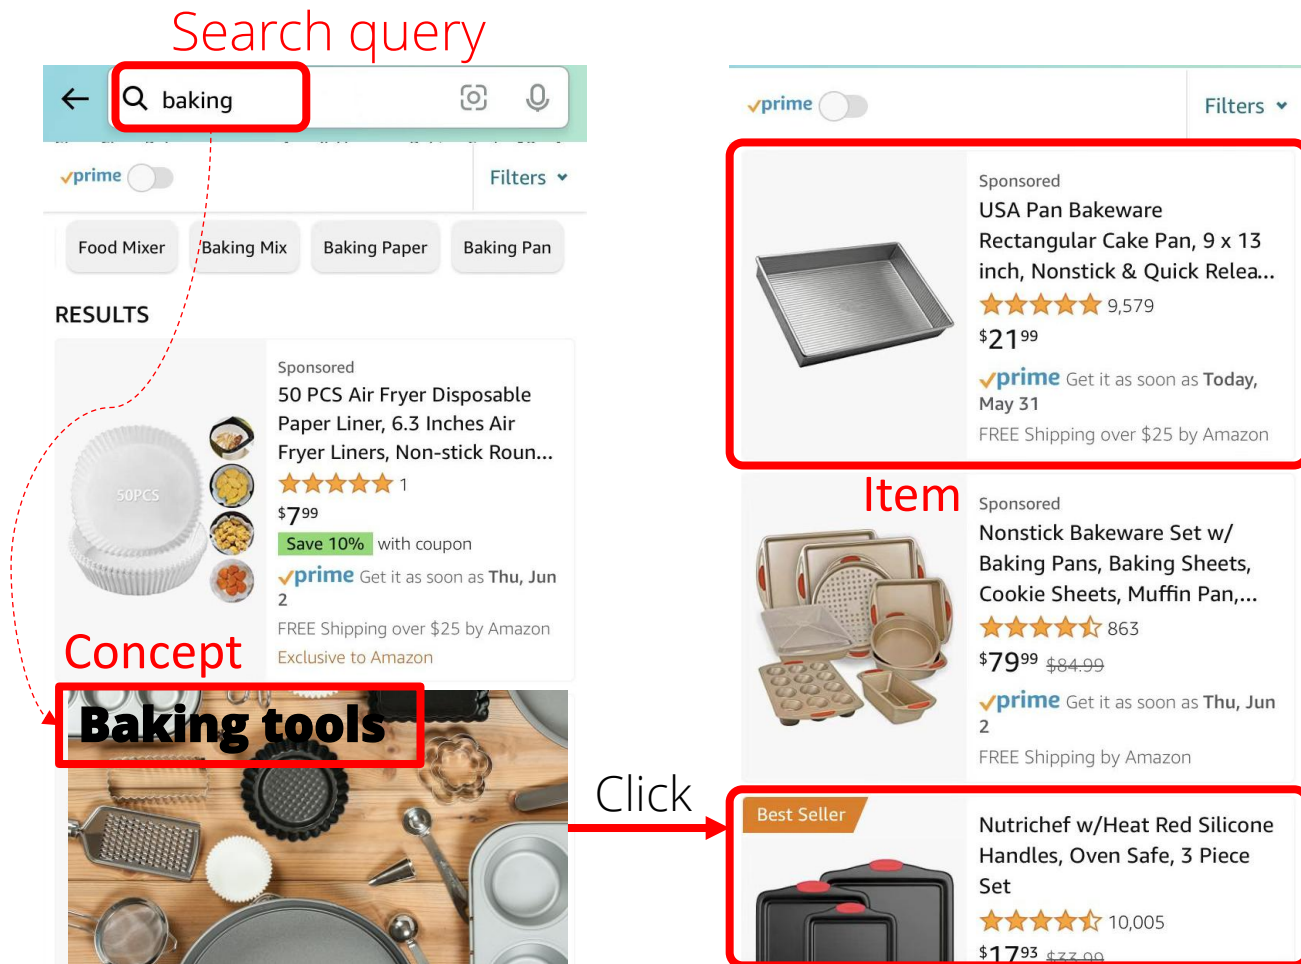


Deployments: COSMO, Intention KG for Search-buy

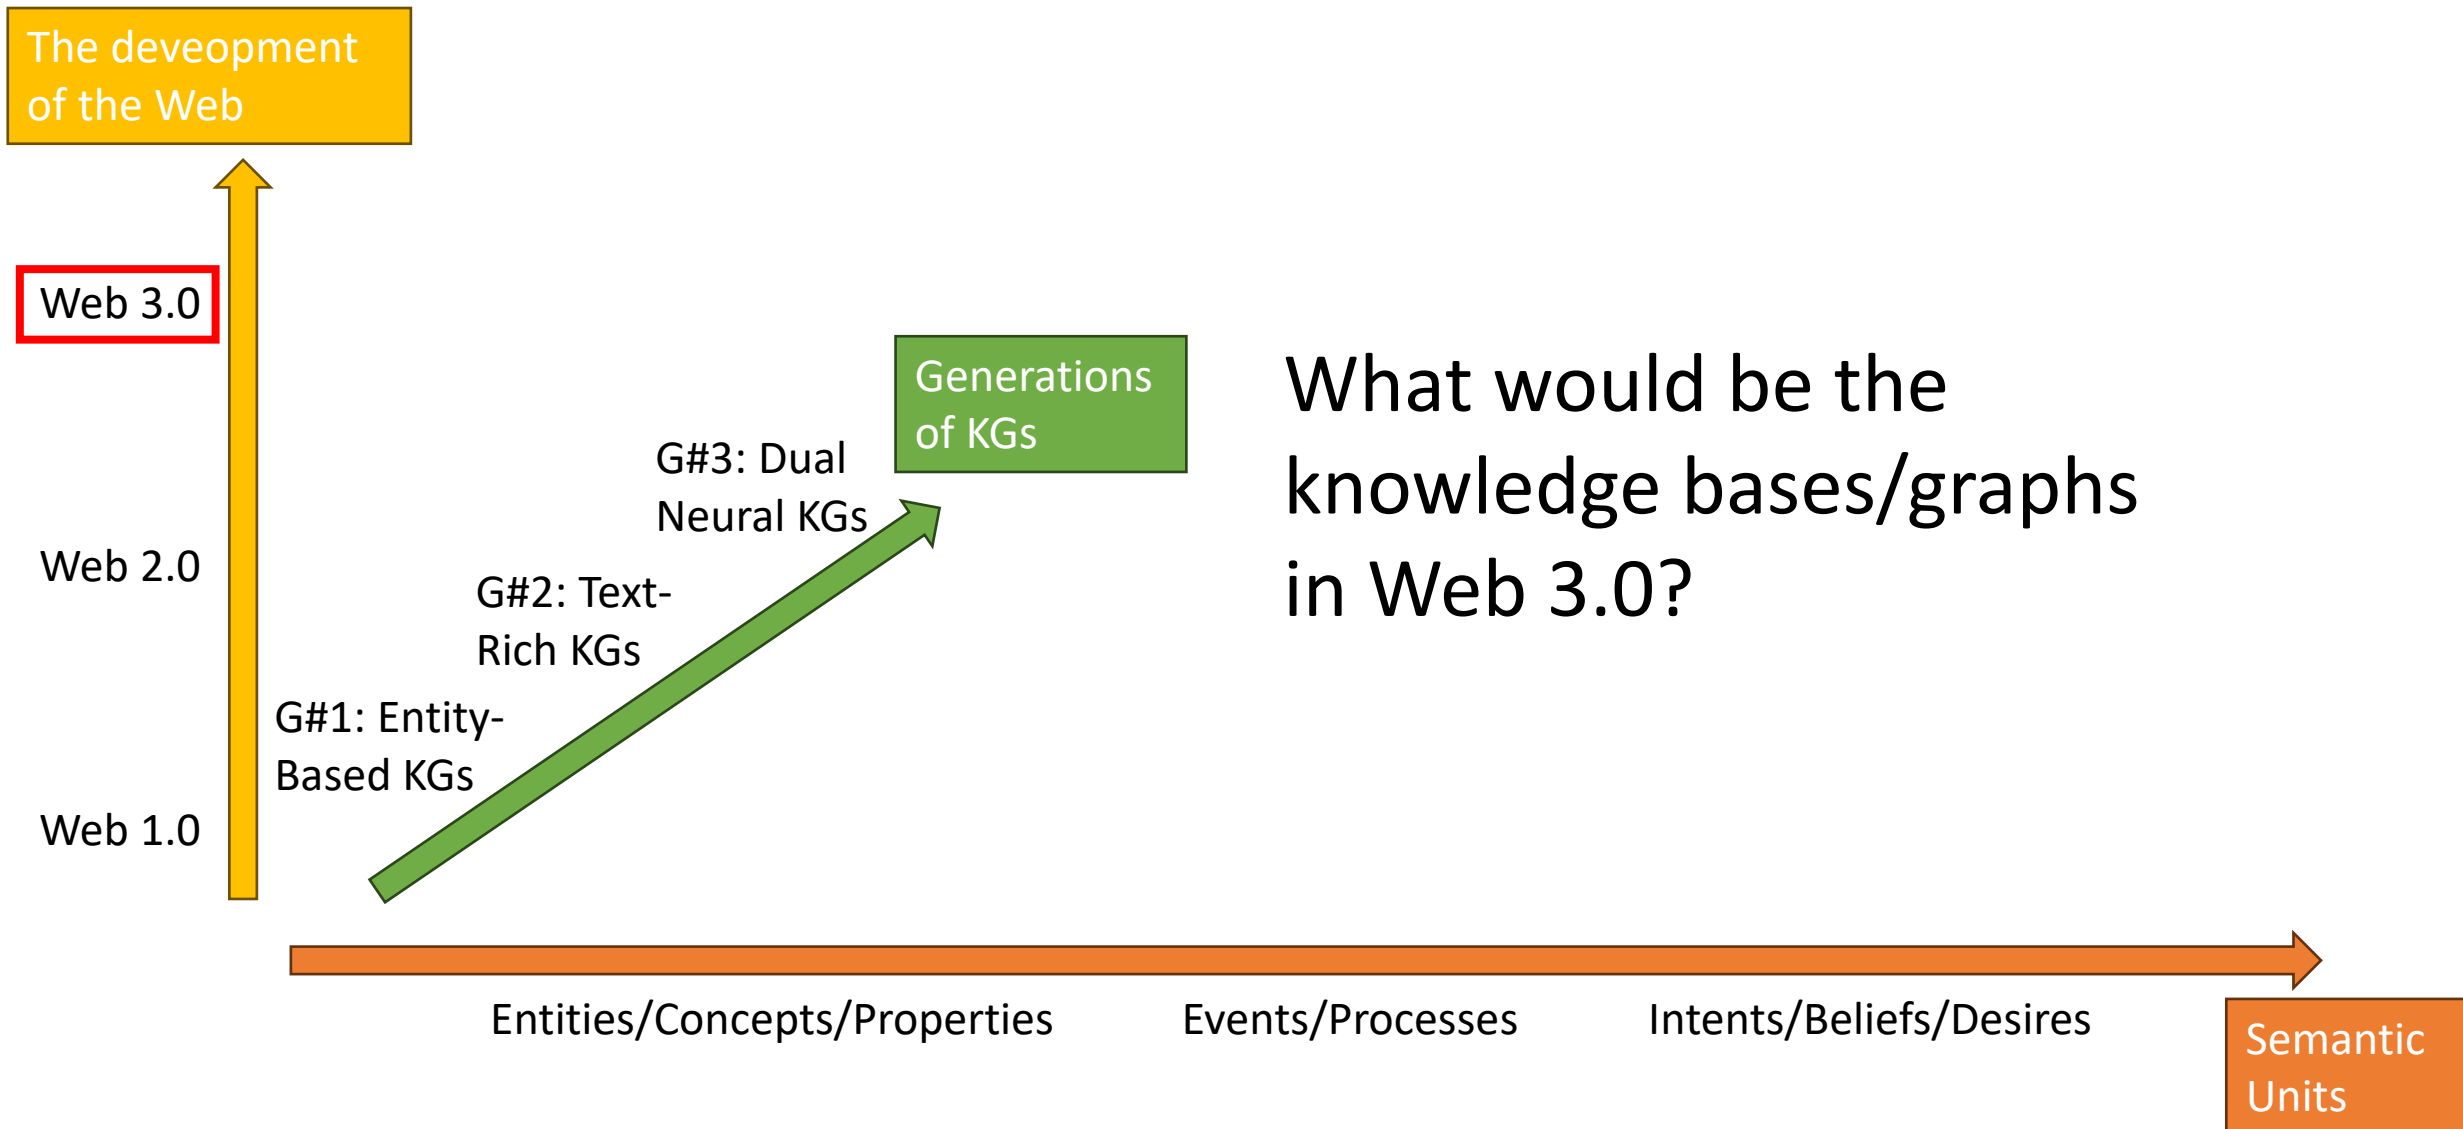
- Efficient **feature store** and asynchronous **cache store**
- Effectively meets Amazon's **restricted search latency** requirements while maintaining **storage costs** comparable to real-time serving for the **majority of traffic**

Search Query Navigation

- **A/B tests** carried out over **several months** in total
- approximately **10% of Amazon's U.S. traffic**
- a notable **0.7% relative increase** in product sales
- translating to **hundreds of millions of dollars** in annual revenue surge.

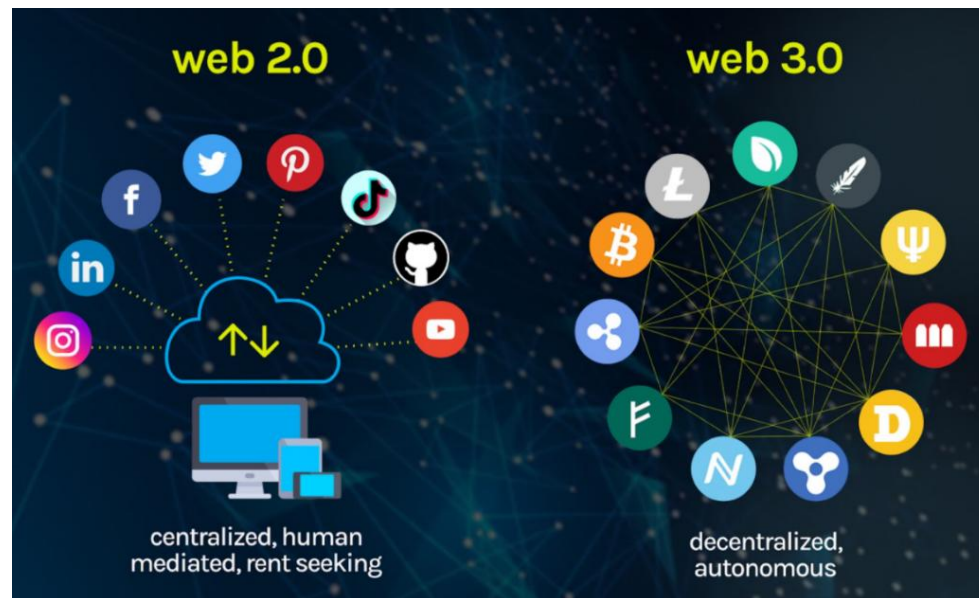


The Three-dimensional Development of Knowledge Graphs in Computer Science



An Outlook

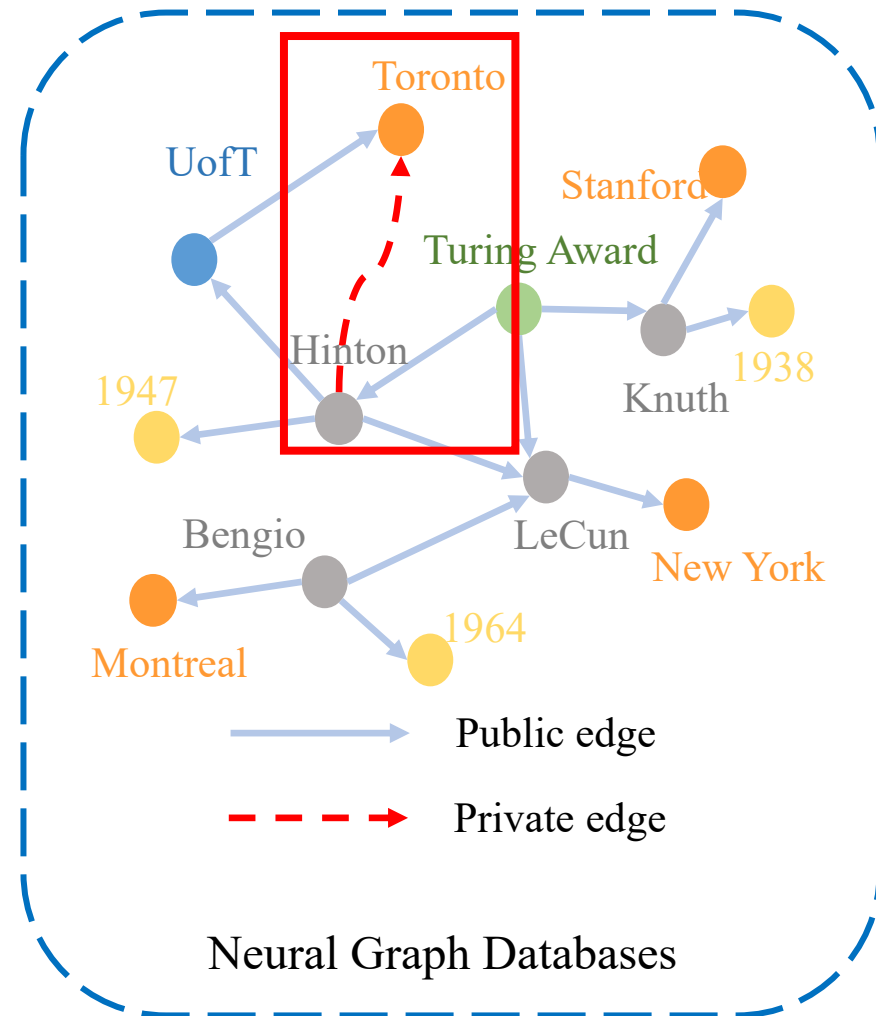
- From Web2.0 to Web3.0
 - Decentralized data: users own their (neural) knowledge bases/graphs
 - Monetize by users' data and time
 - Permissionless, trustless, but accessible to users' owned knowledge or data



- **Security and privacy** of data and knowledge is the key!

Privacy-preserved NGDBs

An attacker attempts to infer private information about Hinton's living place in the NGDBs. Attackers can leverage well-designed queries to retrieve desired privacy. The intersection of these queries can make a fair guess.



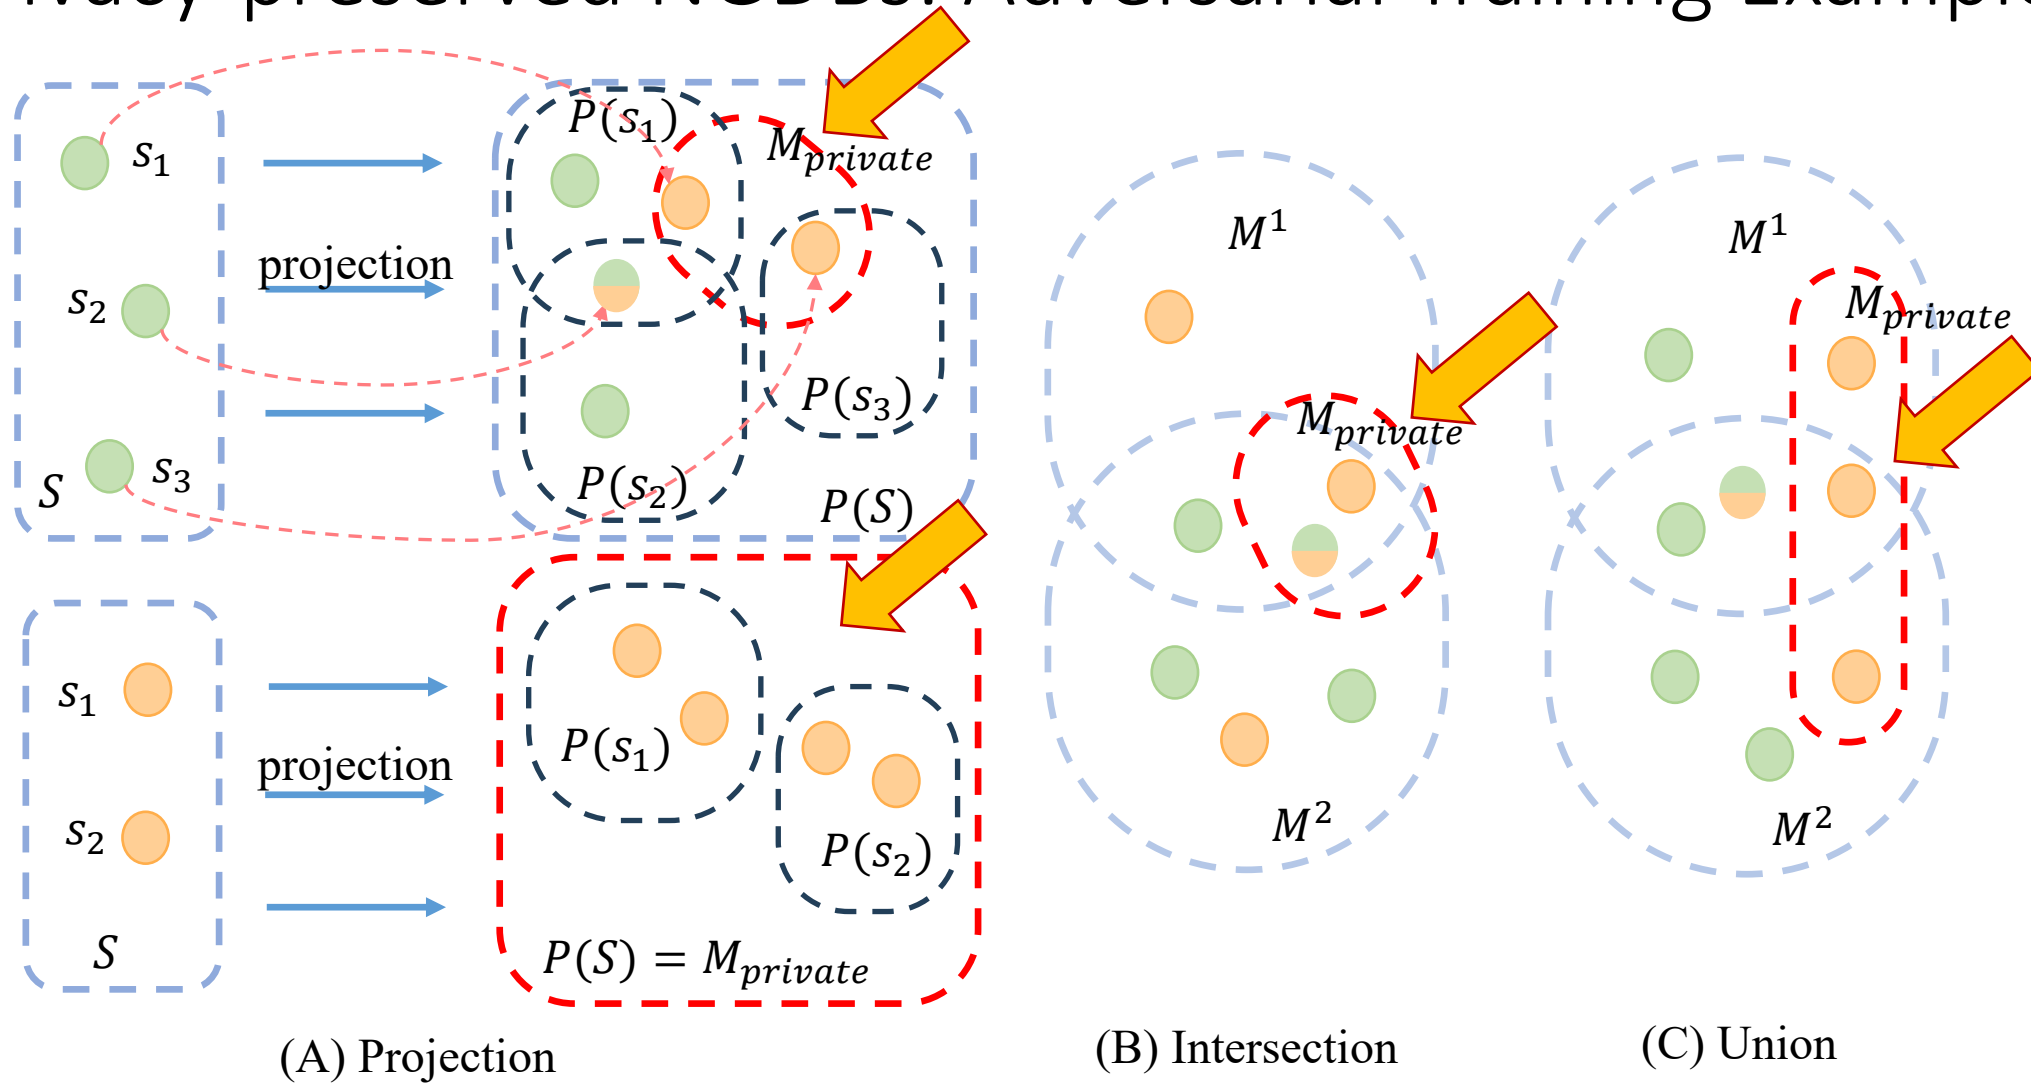
Query	$q = V_?. \exists V: Win(V, Turing\ Award) \wedge BornIn(V, 1938) \wedge LiveIn(V, V_?)$
Interpretation	Find where the Turing Award winner who was born in 1938 lived.

Complex Queries

Query	Answer
$q_1 = V_?. LiveIn(Hinton, V_?)$	Privacy risk query detection
$q_2 = V_?. \exists X_1, X_2: Win(X_1, Turing\ Award) \wedge GreaterThan(X_2, 1940) \wedge BornIn(X_1, X_2) \wedge Livein(X_1, V_?)$	Montreal, Toronto ...
$q_3 = V_?. \exists X_1: CollabWith(LeCun, X_1) \wedge LiveIn(X_1, V_?)$	Montreal, Toronto ...
$q_4 = V_?. \exists X_1, X_2: Win(X_1, Turing\ Award) \wedge SmallerThan(X_2, 1950) \wedge BornIn(X_1, X_2) \wedge LiveIn(X_1, V_?)$	Toronto , Stanford...

Privacy Risk Queries

Privacy-preserved NGDBs: Adversarial Training Examples



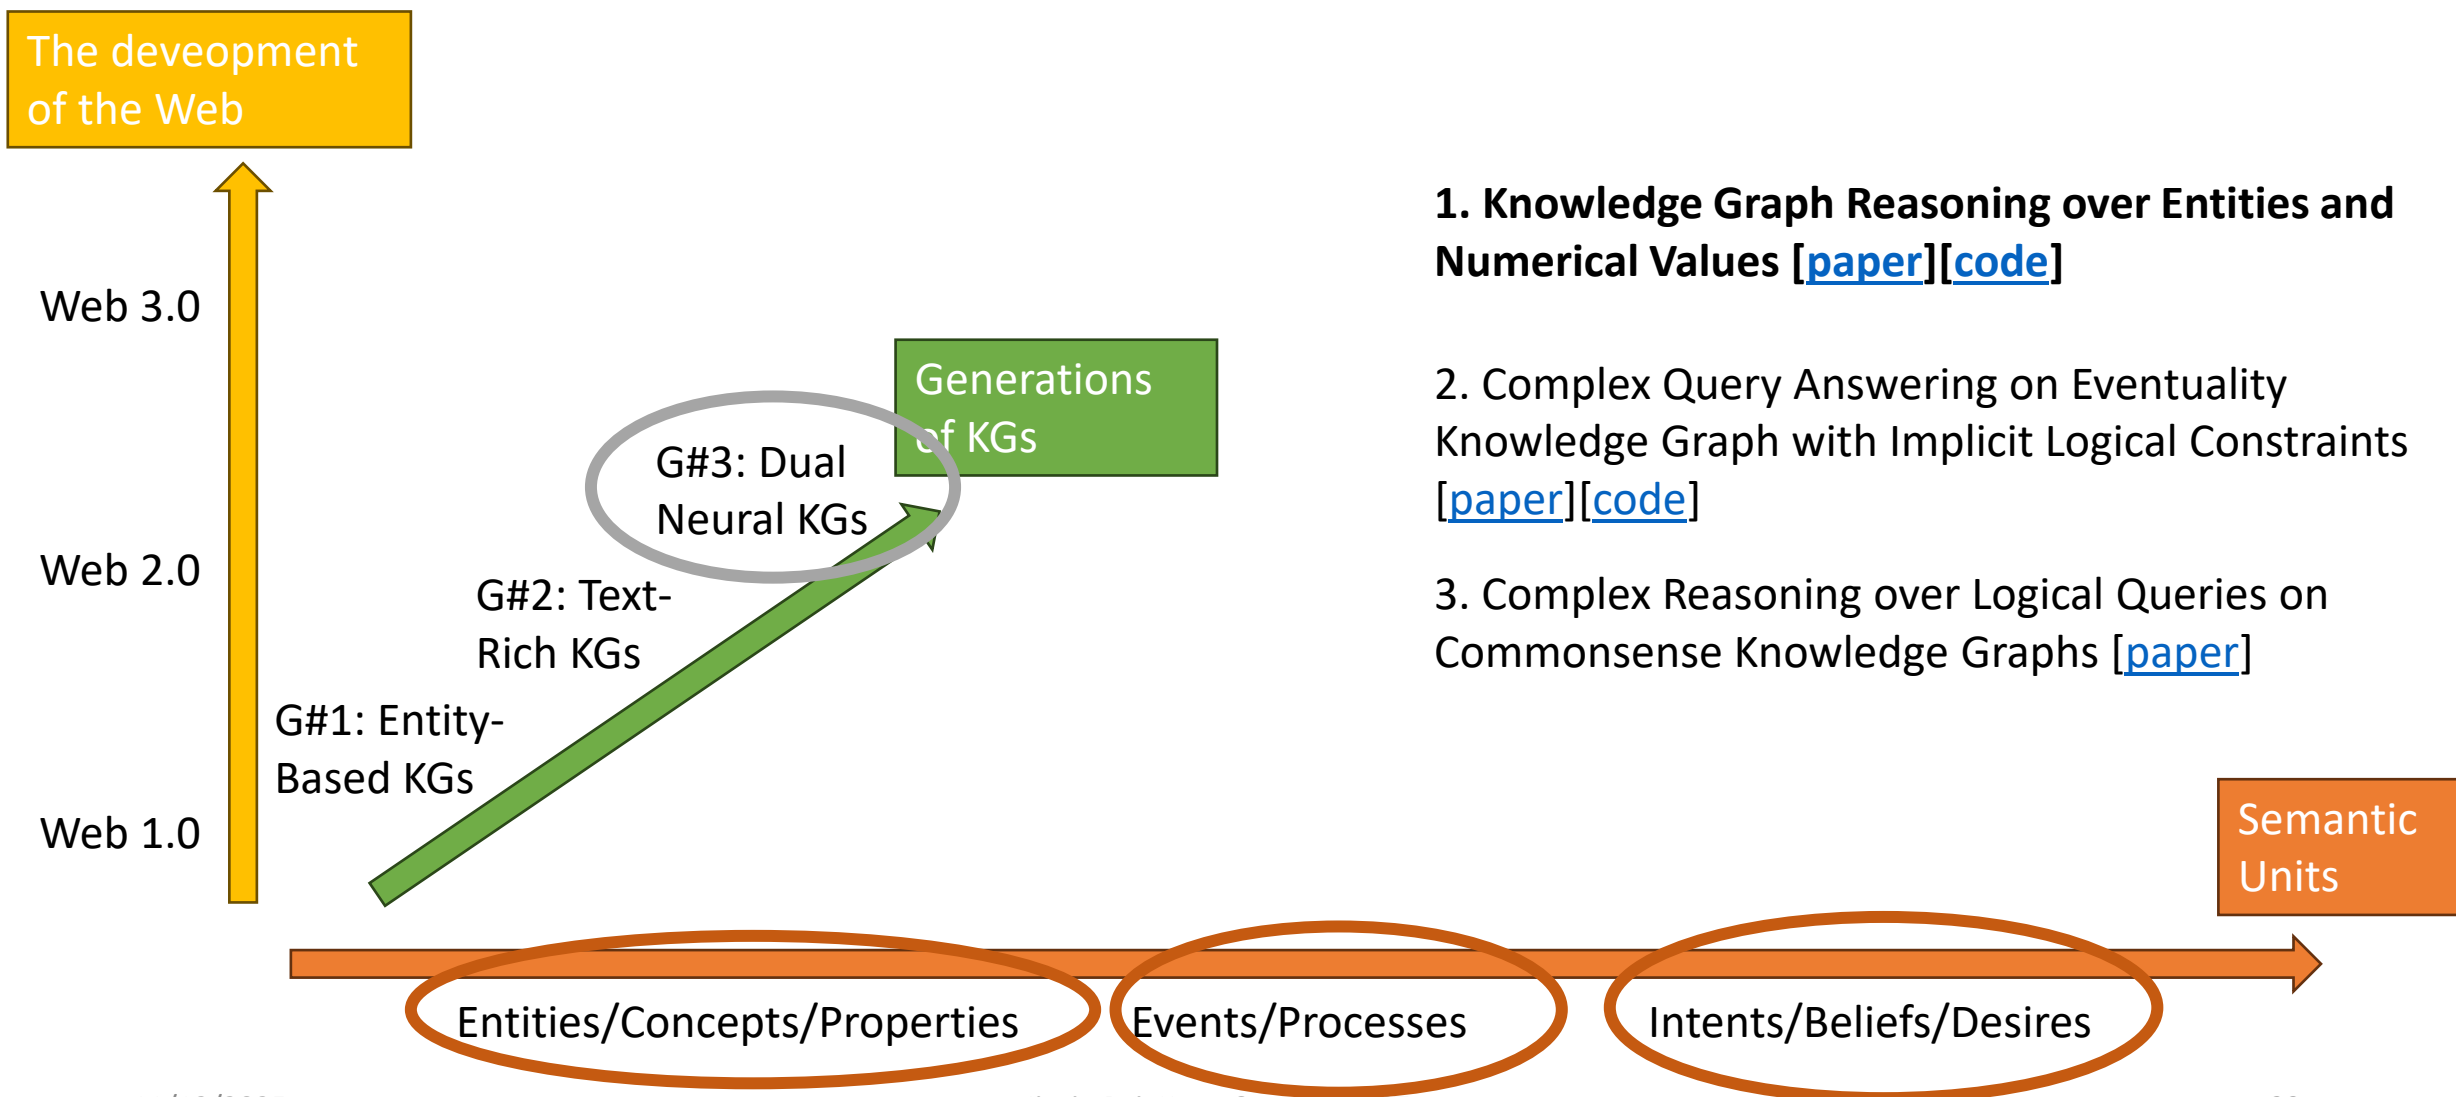
Green nodes denote non-private answers, orange nodes denote privacy-threatening answers, and orange-green nodes denote different privacy risks in subsets. Red dashed arrows denote privacy projection. The answers circled in red dashed line are at risk to leak privacy.

Conclusions for Part 4-1

- We have reviewed the frontier of recent development of knowledge graphs in terms of three dimensions
 - Three generations: entity, text, and neural
 - Semantic units: entity/attributes, events/processes, intent/desire/belief
 - **Web 1.0/2.0/3.0**
- More methods will be introduced in Part 4-2 by Jiaxin Bai
 - In terms of different types of semantic units



The Three-dimensional Development of Knowledge Graphs in Computer Science

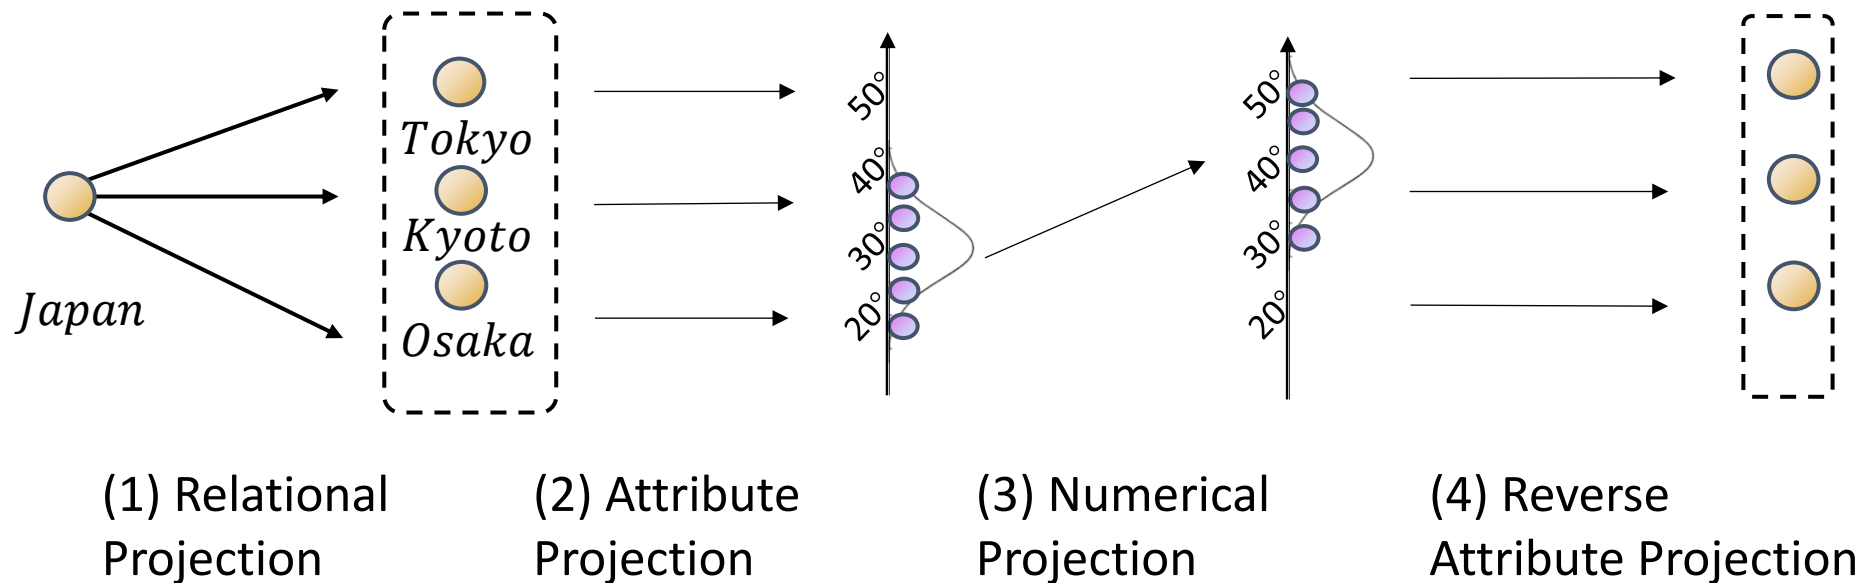



Numerical Complex Query Answering

Category	Complex Queries	Interpretations
Numerical CQA	$q_2 = V_? . \exists X_1, X_2: \text{Win}(V_?, \text{TuringAward})$ $\wedge \text{GreaterThan}(1927, X_2) \wedge \text{BornIn}(V_?, X_2)$	Find the Turing award winners that <u>is born before</u> the year of 1927.
Numerical CQA	$q_3 = V_? . \exists X_1, X_2: \text{LocatedIn}(V_?, \text{UnitedStates})$ $\wedge \text{HasLatitude}(V_?, X_1)$ $\wedge \text{GreaterThan}(X_1, X_2)$ $\wedge \text{HasLatitude}(\text{Beijing}, X_2)$	Find the states in US that have <u>a higher latitudes</u> than Beijing.
Numerical CQA	q_4 $= V_? . \exists X_1, X_2, X_3: \text{LocatedIn}(V_?, \text{UnitedStates})$ $\wedge \text{HasPopulation}(V_?, X_1)$ $\wedge \text{SmallerThan}(X_1, X_2) \wedge \text{TimesByTwo}(X_2, X_3)$ $\wedge \text{HasPopulation}(\text{California}, X_3)$	Find the states in US that have a <u>twice smaller population</u> than California?

Number Reasoning Network

Find the cities that have a higher latitudes than Japanese cities.

$$q = V_? . \exists V_1, X_1, X_2: \underline{HasLatitude}(V_?, X_2) \wedge \underline{GreaterThan}(X_2, X_1) \wedge \underline{HasLatitude}(V_1, X_1) \wedge \underline{LocatedIn}(V_1, Japan)$$


Number Reasoning Network

(1) Relational Projection (rp):

Query Embedding \rightarrow Entity Set

(2) Attribute Projection (ap):


Query Embedding \rightarrow Value Distribution

(3) Numerical Projection (np):

Value Distribution \rightarrow Value Distribution

(4) Reverse Attribute Projection (rap):

Value Distribution \rightarrow Query Embedding



Number Reasoning Network

(1) Relational Projection:
Adopted from the backbones:
GQE, Query2Box, Query2Particles.

(2) (3) (4) Other Projections: Gated Transitions

$$p_i = W_p^p q^i + b_p^p$$

Linear projection

$$z_i = \sigma(W_z^p e_a + U_z^p p_i + b_z^p)$$

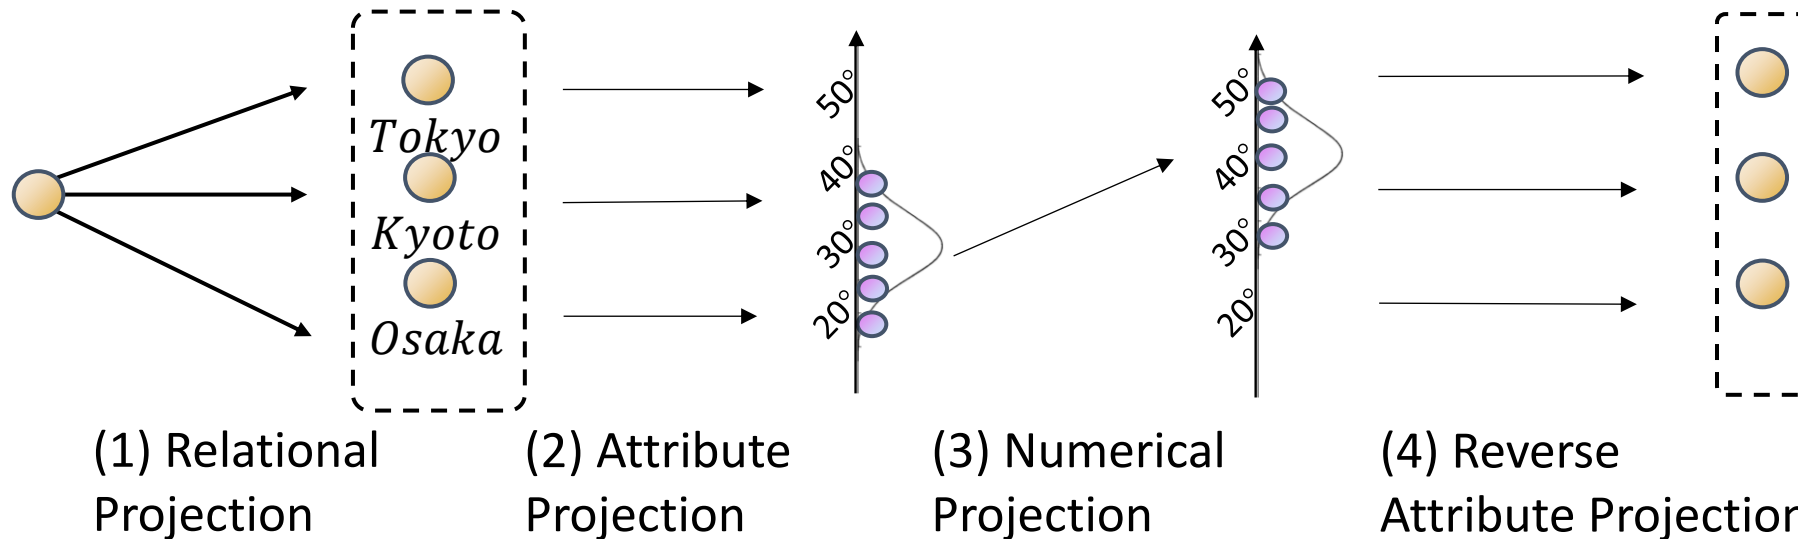
$$r_i = \sigma(W_r^p e_a + U_r^p p_i + b_r^p)$$

$$t_i = \varphi(W_h^p e_a + U_h^p (r_i \odot p_i) + b_h^p)$$

MLP

$$\theta_{i+1} = (1 - z_i) \odot p_i + z_i \odot t_i$$

Gate selection



Number Reasoning Network

Entity embeddings:

Adopted from the backbones:

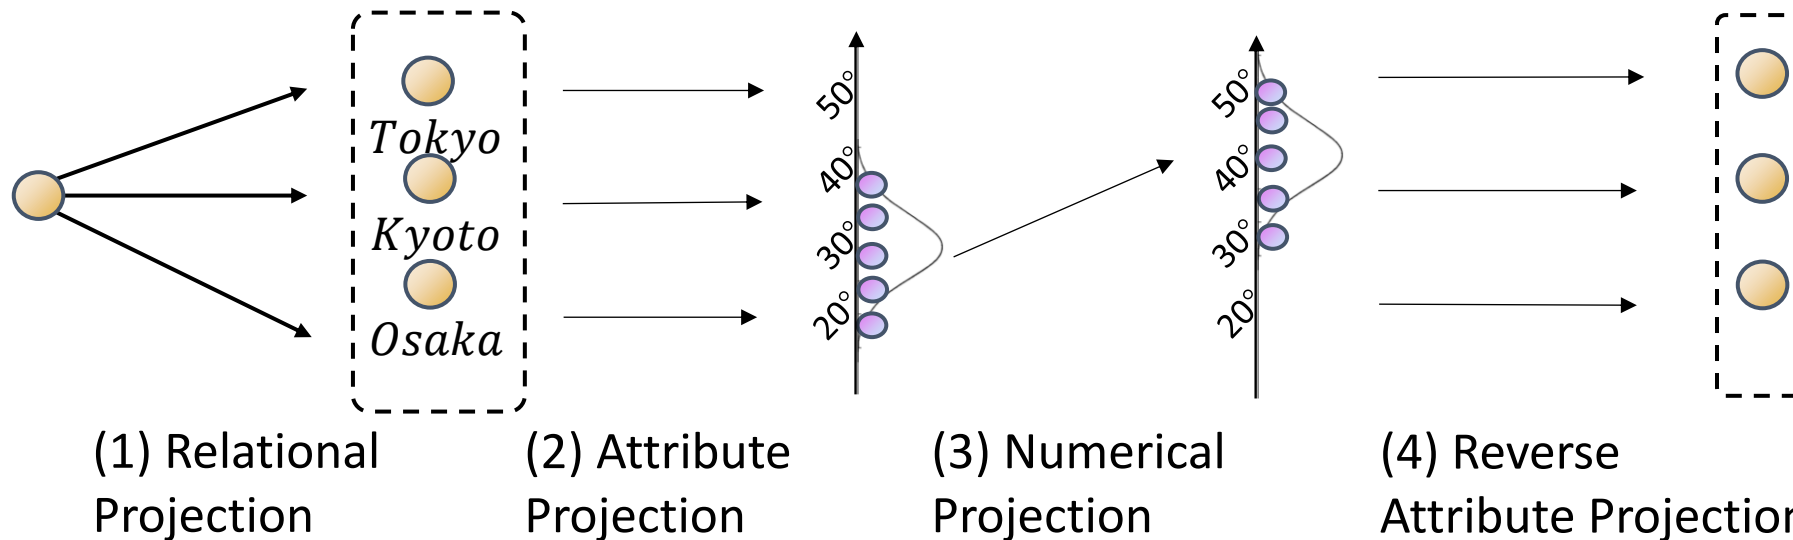
GQE, Query2Box, Query2Particles.

Input number embeddings

- DICE

- Sinusoidal

$$\psi(v)_d = \begin{cases} \sin^{d-1}(\alpha) \cos(\alpha) \\ \sin^D(\alpha) \end{cases} \quad \psi(v)_d = \begin{cases} \sin \frac{v}{v^{d/D}}, & d \equiv 0 \pmod{2} \\ \cos \frac{v}{v^{(d-1)/D}}, & d \equiv 1 \pmod{2} \end{cases}$$



Number Embeddings and Learning Objective

End-to-end training by Joint optimization of two losses:

$$L_A = \frac{1}{M} \sum_{j=1}^M \left(-\log p_{\theta_I^{(j)}} \left(\psi(v^{(j)}) \right) - \log \phi_{t^{(j)}}(\theta_I^{(j)}) \right)$$

The likelihood of the value

$v^{(j)}$ sampled from distribution of $\theta_I^{(j)}$

The likelihood of the distribution

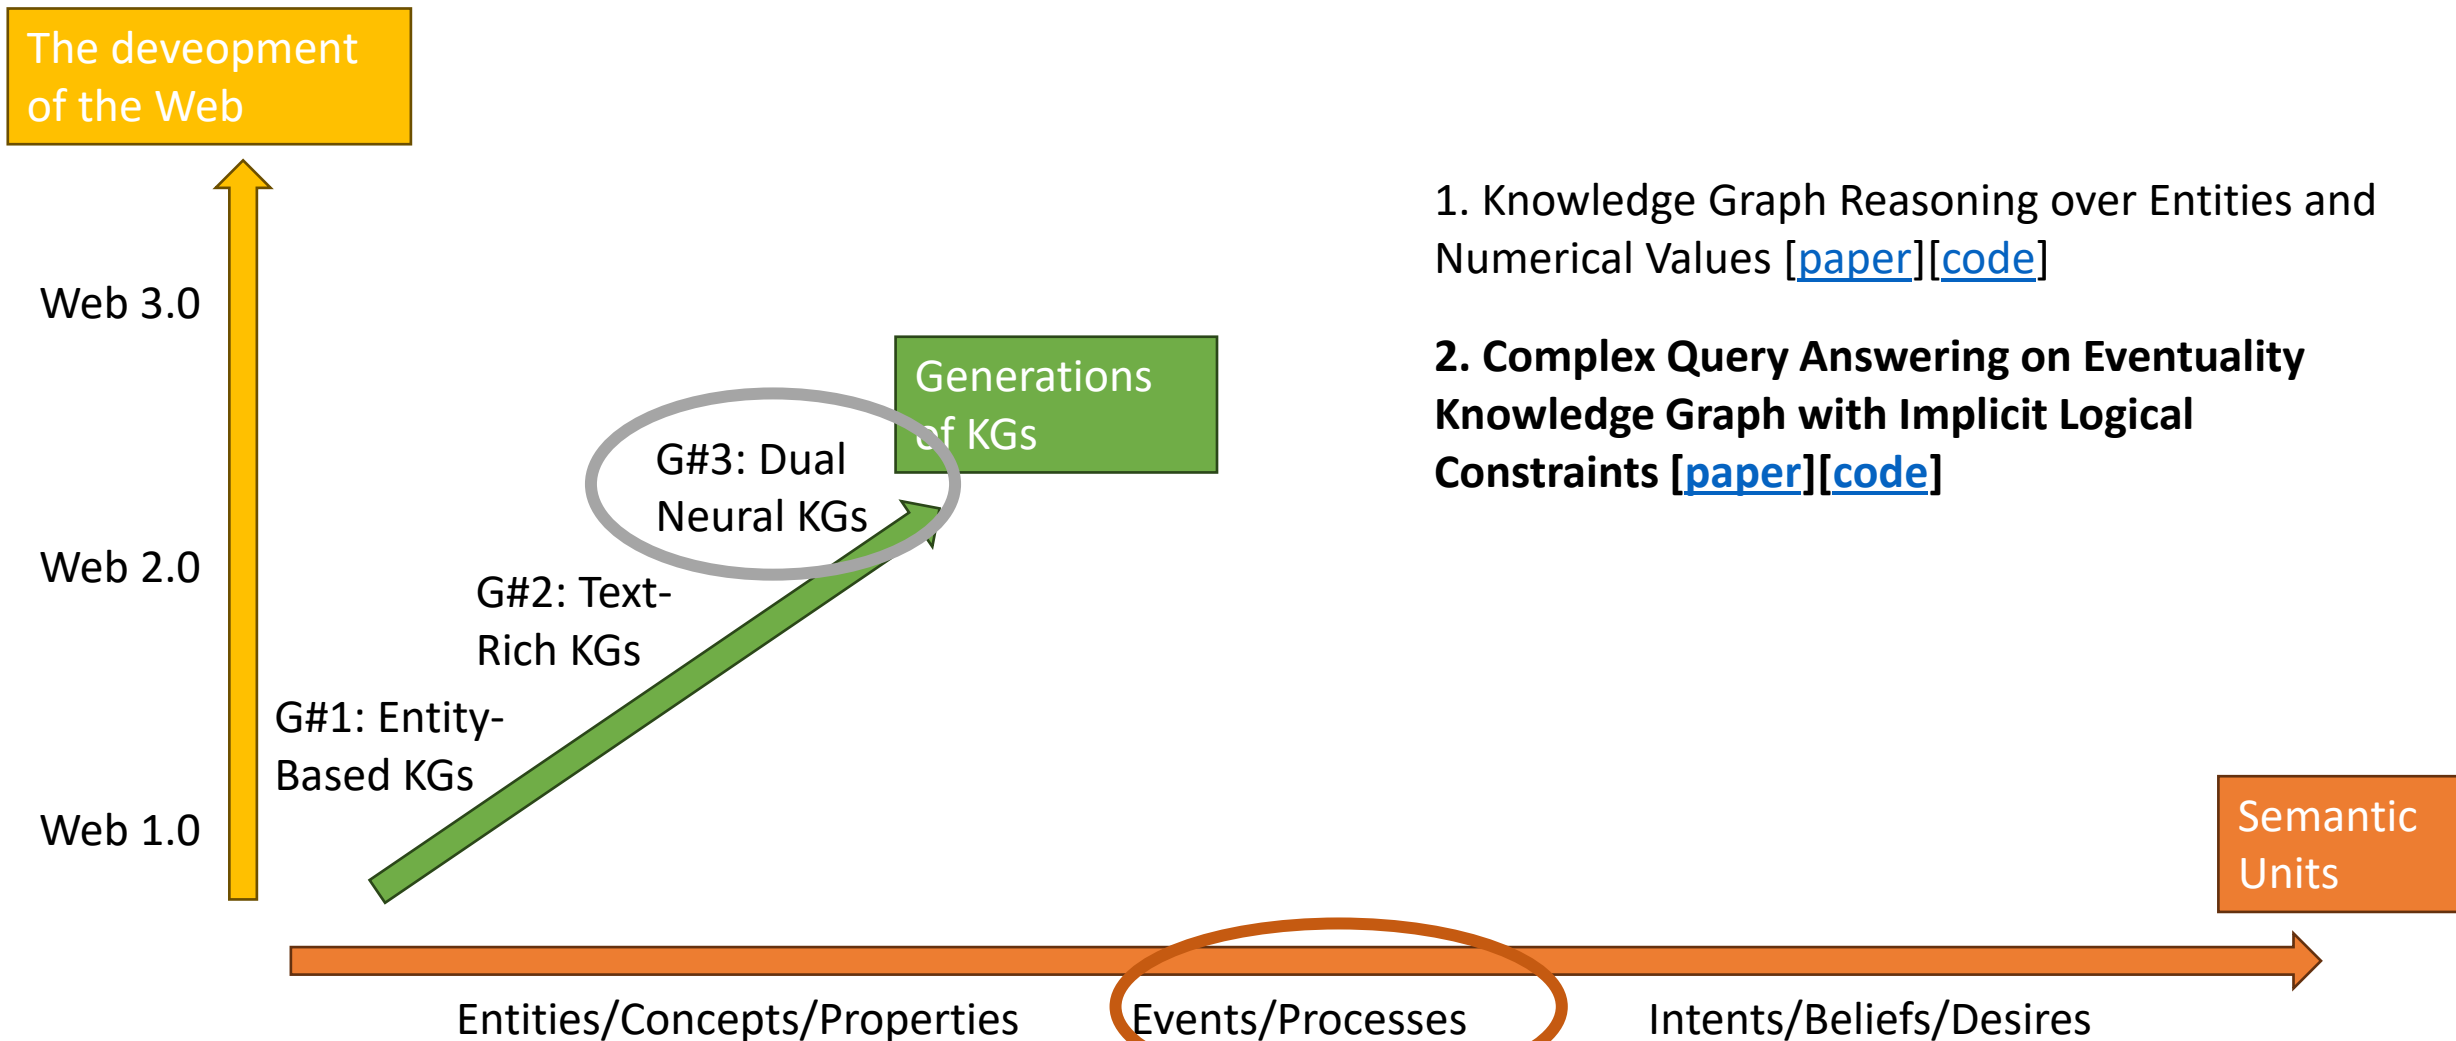
parameter $\theta_I^{(j)}$ is of type $t^{(j)}$

$$L_E = -\frac{1}{N} \sum_{j=1}^N \log p(q_I^{(j)}, v^{(j)})$$

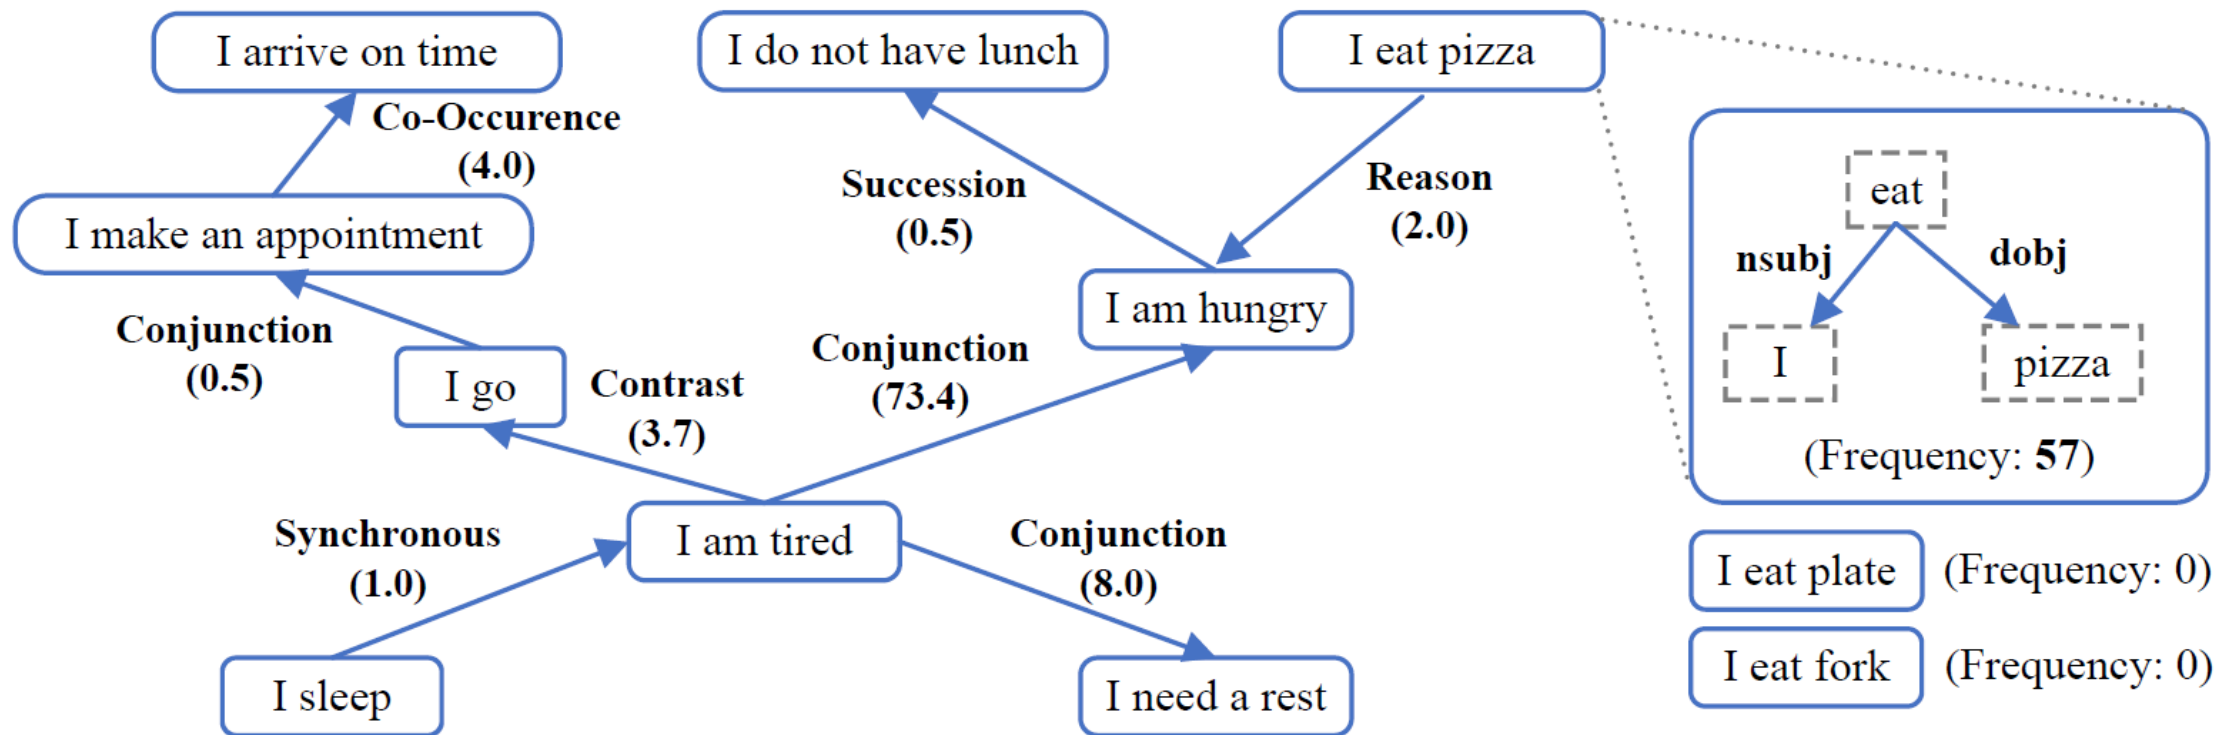
The likelihood of the entity $v^{(j)}$ is the answer of the query encoding $q_I^{(j)}$.

j is means the j -th sample, and I means the last step of distribution parameter encoding.

The Three-dimensional Development of Knowledge Graphs in Computer Science



ASER (Activities, States, Events, and their Relations)



Principle 1: Comparing semantic meanings by fixing grammar (Katz and Fodor, 1963)

Principle 2: The need of language inference based on ‘partial information’ (Wilks, 1975)

<https://github.com/HKUST-KnowComp/ASER>

Hongming Zhang, Xin Liu, Haojie Pan, Yangqiu Song, Cane Wing-Ki Leung: ASER: A Large-scale Eventuality Knowledge Graph. WWW 2020: 201-211

Katz, J. J., & Fodor, J. A. (1963). The structure of a semantic theory. Language, 39(2), 170-210.

Yorick Wilks. 1975. An intelligent analyzer and understander of English. Communications of the ACM, 18(5):264-274.

CQA on Eventuality Knowledge Graph

Complex query on eventuality graphs are **different** from the entity-relation graph

Whether and **when** the eventualities occur are important

Queries	Type	Interpretations
$q_1 = V_? . \exists V: \text{Interact}(V_?, V)$ $\wedge \text{Assoc}(V, \text{Alzheimer}) \wedge \text{Assoc}(V, \text{MadCow})$	Entity	Find the substances that interact with the proteins associated with Alzheimer's and Mad cow disease.
$q_2 = V_? . \text{Precedence}(\text{Food is bad}, \text{PersonX add soy sauce})$ $\wedge \text{Reason}(\text{Food is bad}, V_?)$	Eventuality	Food is bad before PersonX add soy sauce. What is the reason for food being bad?
$q_3 = V_? . \text{Precedence}(V_?, \text{PersonX go home})$ $\wedge \text{ChosenAlternative}(\text{PersonX go home}, \text{PersonX buy an umbrella})$	Eventuality	Instead of buying an umbrella, PersonX go home. What happened before PersonX go home?

Discourse Relations and Implicit Constraints

- PersonX did not eat anything because PersonX was full

Reason(PersonX did not eat anything, PersonX was full)

Occurrence
Constraint

$\eta(\text{PersonX did not eat anything}) \wedge \eta(\text{PersonX was full})$
 $\wedge \eta(\text{PersonX did not eat anything}) \leftarrow \eta(\text{PersonX was full})$

Temporal
Constraints

$\tau(\text{PersonX did not eat anything}) > \tau(\text{PersonX was full})$

$\eta(A) = 1$ if and only if it occurs
 $\tau(A) > \tau(B) : A$ happens after B

Discourse Relations and Implicit Constraints

- Food is bad before PersonX add soy sauce

Precedence(Food is bad, PersonX adds soy sauce)

Occurrence
Constraint

$\eta(\text{Food is bad}) \wedge \eta(\text{PersonX adds soy sauce})$

Temporal
Constraints

$\tau(\text{Food is bad}) < \tau(\text{PersonX adds soy sauce})$

$\tau(A) < \tau(B)$: A happens before B
 $\eta(A) = 1$ if and only if it occurs

Discourse Relations and Implicit Constraints

- Instead of buying an umbrella, PersonX go home

ChosenAlternative(buying an umbrella, *PersonX* go home)

Occurrence
Constraint

$\eta(\textit{PersonX go home}) \wedge \neg \eta(\textit{PersonX buy an umbrella})$

$\eta(A) = 1$ if and only if it occurs

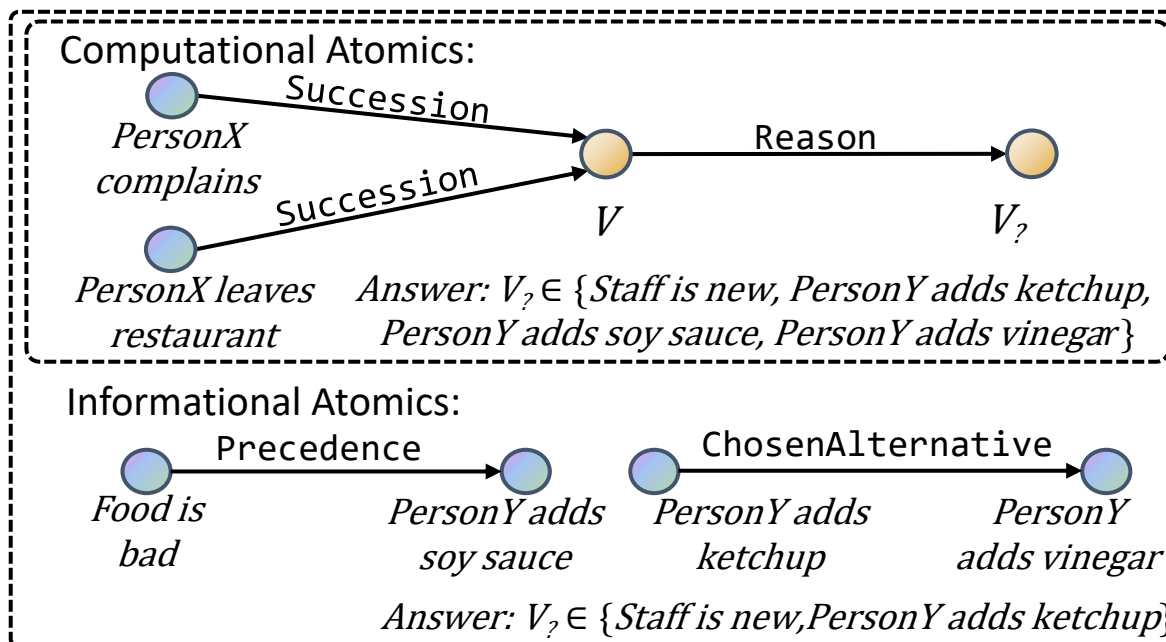
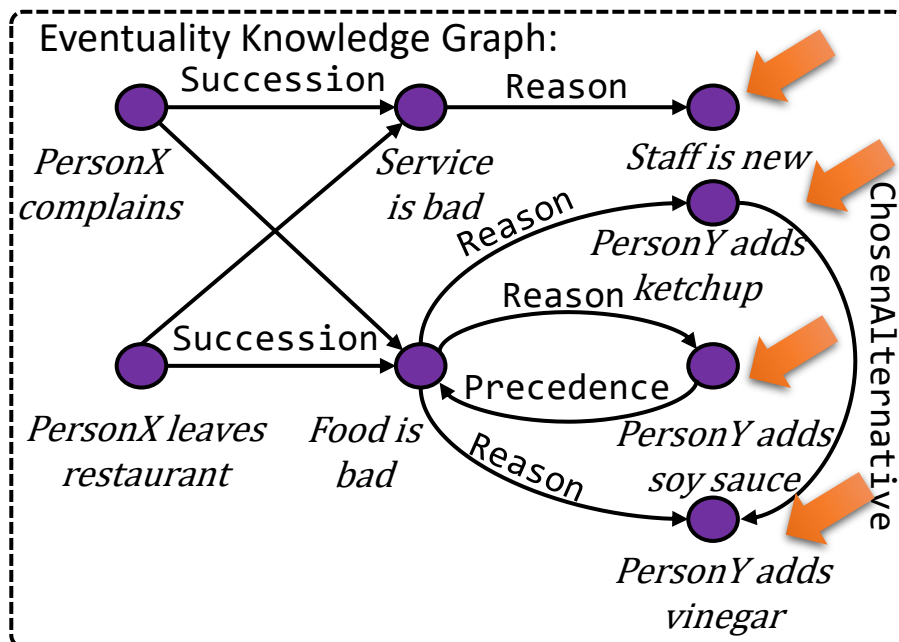
Logical Constraints behind Discourse Relations

Discourse Relations	Semantics	Implicit Constraints	
		Occurrence Constraints	Temporal Constraints
Precedence(A, B)	A occurs before B.	$\eta(A) \wedge \eta(B)$	$\tau(A) < \tau(B)$
Succession(A, B)	A occurs after B happens.	$\eta(A) \wedge \eta(B)$	$\tau(A) > \tau(B)$
Synchronous(A, B)	A occurs at the same time as B.	$\eta(A) \wedge \eta(B)$	$\tau(A) = \tau(B)$
Reason(A, B)	A occurs because B.	$\eta(A) \wedge \eta(B) \wedge (\eta(A) \leftarrow \eta(B))$	$\tau(A) > \tau(B)$
Result(A, B)	A occurs as a result B.	$\eta(A) \wedge \eta(B) \wedge (\eta(A) \rightarrow \eta(B))$	$\tau(A) < \tau(B)$
Condition(A, B)	If B occurs, A.	$\eta(A) \rightarrow \eta(B)$	$\tau(A) > \tau(B)$
Concession(A, B)	B occurs, although A.	$\eta(A) \wedge \eta(B)$	-
Contrast(A, B)	B occurs, but A.	$\eta(A) \wedge \eta(B)$	-
Conjunction(A, B)	A and B both occur.	$\eta(A) \wedge \eta(B)$	-
Instantiation(A, B)	B is a more detailed description of A.	$\eta(A) \wedge \eta(B)$	-
Restatement(A, B)	A restates the semantics of B.	$\eta(A) \leftrightarrow \eta(B)$	-
Alternative(A, B)	A and B are alternative situations.	$\eta(A) \vee \eta(B)$	-
ChosenAlternative(A, B)	Instead of B occurs, A.	$\eta(A) \wedge \neg \eta(B)$	-
Exception(A, B)	A, except B.	$\neg \eta(A) \wedge \eta(B) \wedge (\neg \eta(B) \rightarrow \eta(A))$	-

Logical Query with Implicit Constraints

Question: Food is bad before PersonY adds soy sauce. Instead of adding vinegar, PersonY adds ketchup.
PersonX complains after V. PersonX leaves the restaurant after V. The reason V is V'?. What is V'?

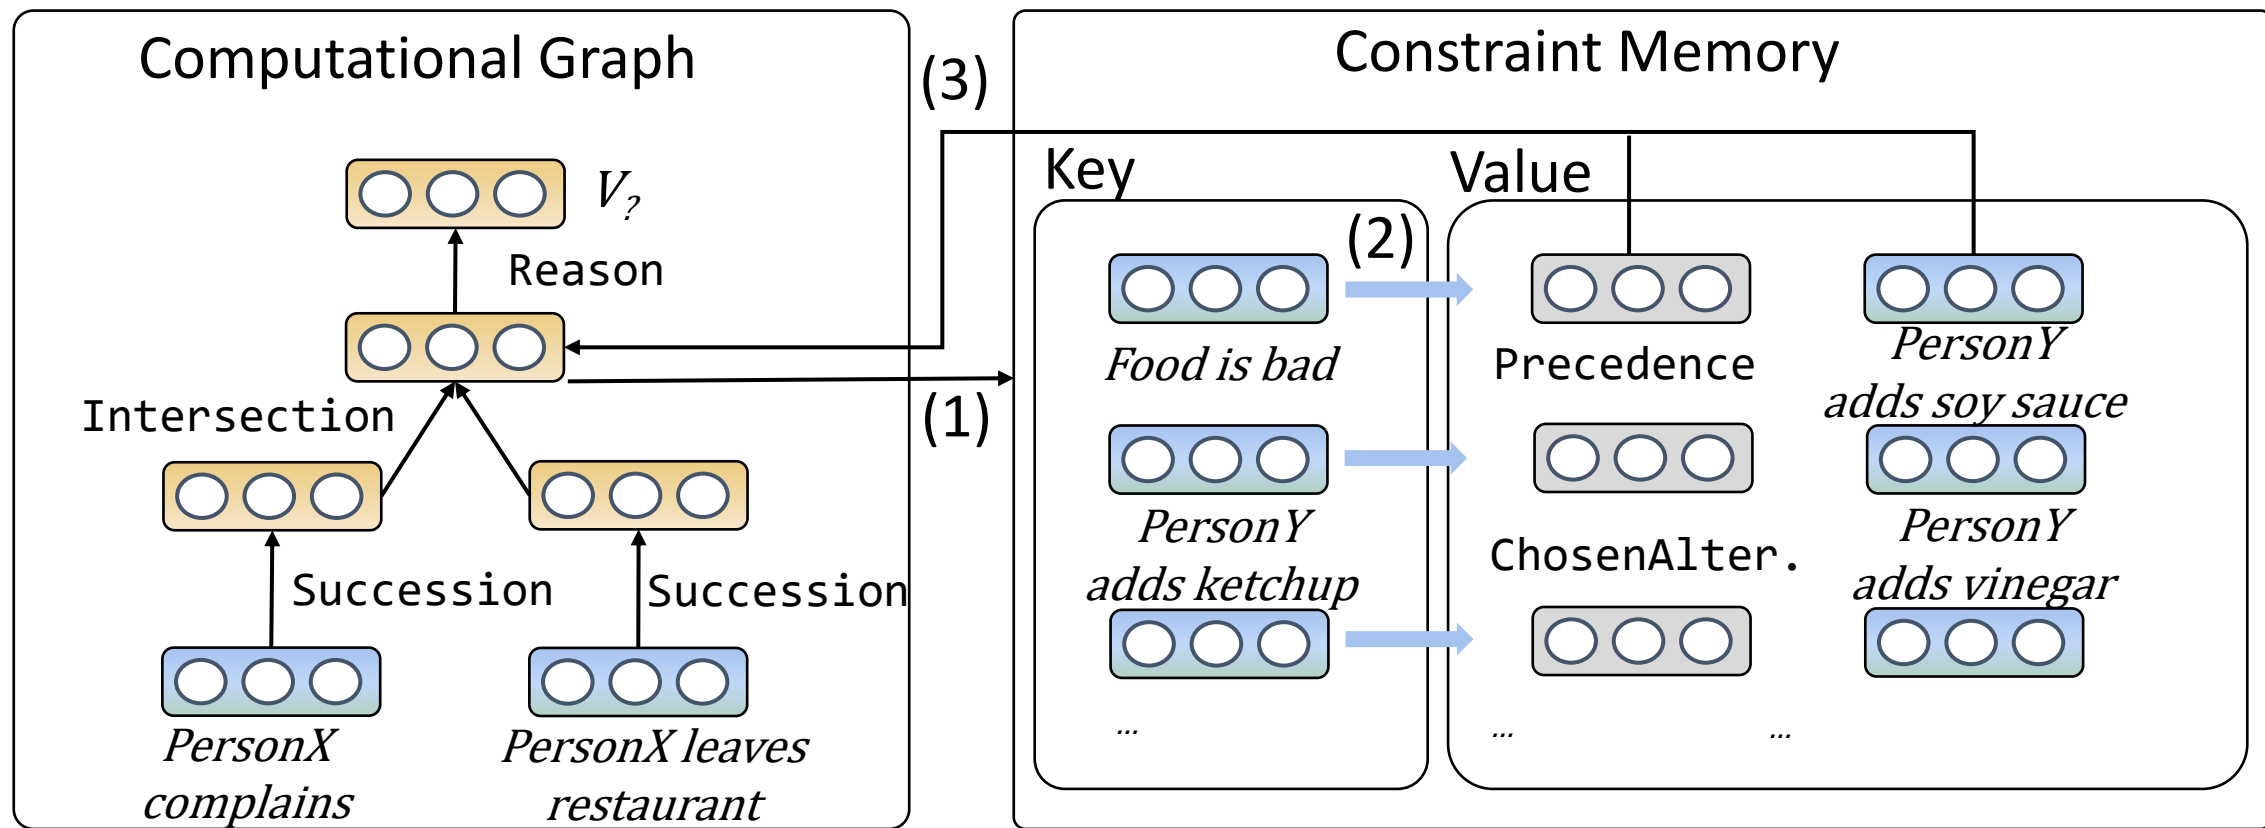
Query on Graph: $q = V_? . \exists V: \text{Succession}(\text{PersonX complains}, V) \wedge \text{Succession}(\text{PersonX leaves restaurant}, V) \wedge \text{Reason}(V, V_?) \wedge \text{Precedence}(\text{Food is bad}, \text{PersonY adds soy sauce}) \wedge \text{ChosenAlternative}(\text{PersonY adds ketchup}, \text{PersonY adds vinegar})$



Without context and its constraints:
4 answers

With Implicit Constraints:
Only 2 answers

Query Encoding with Constraint Memory



$$(1) s_{i,m} = \langle q_i, c_h^{(m)} \rangle$$

Computes the relevance of query embedding to the head of the memory key at position m .

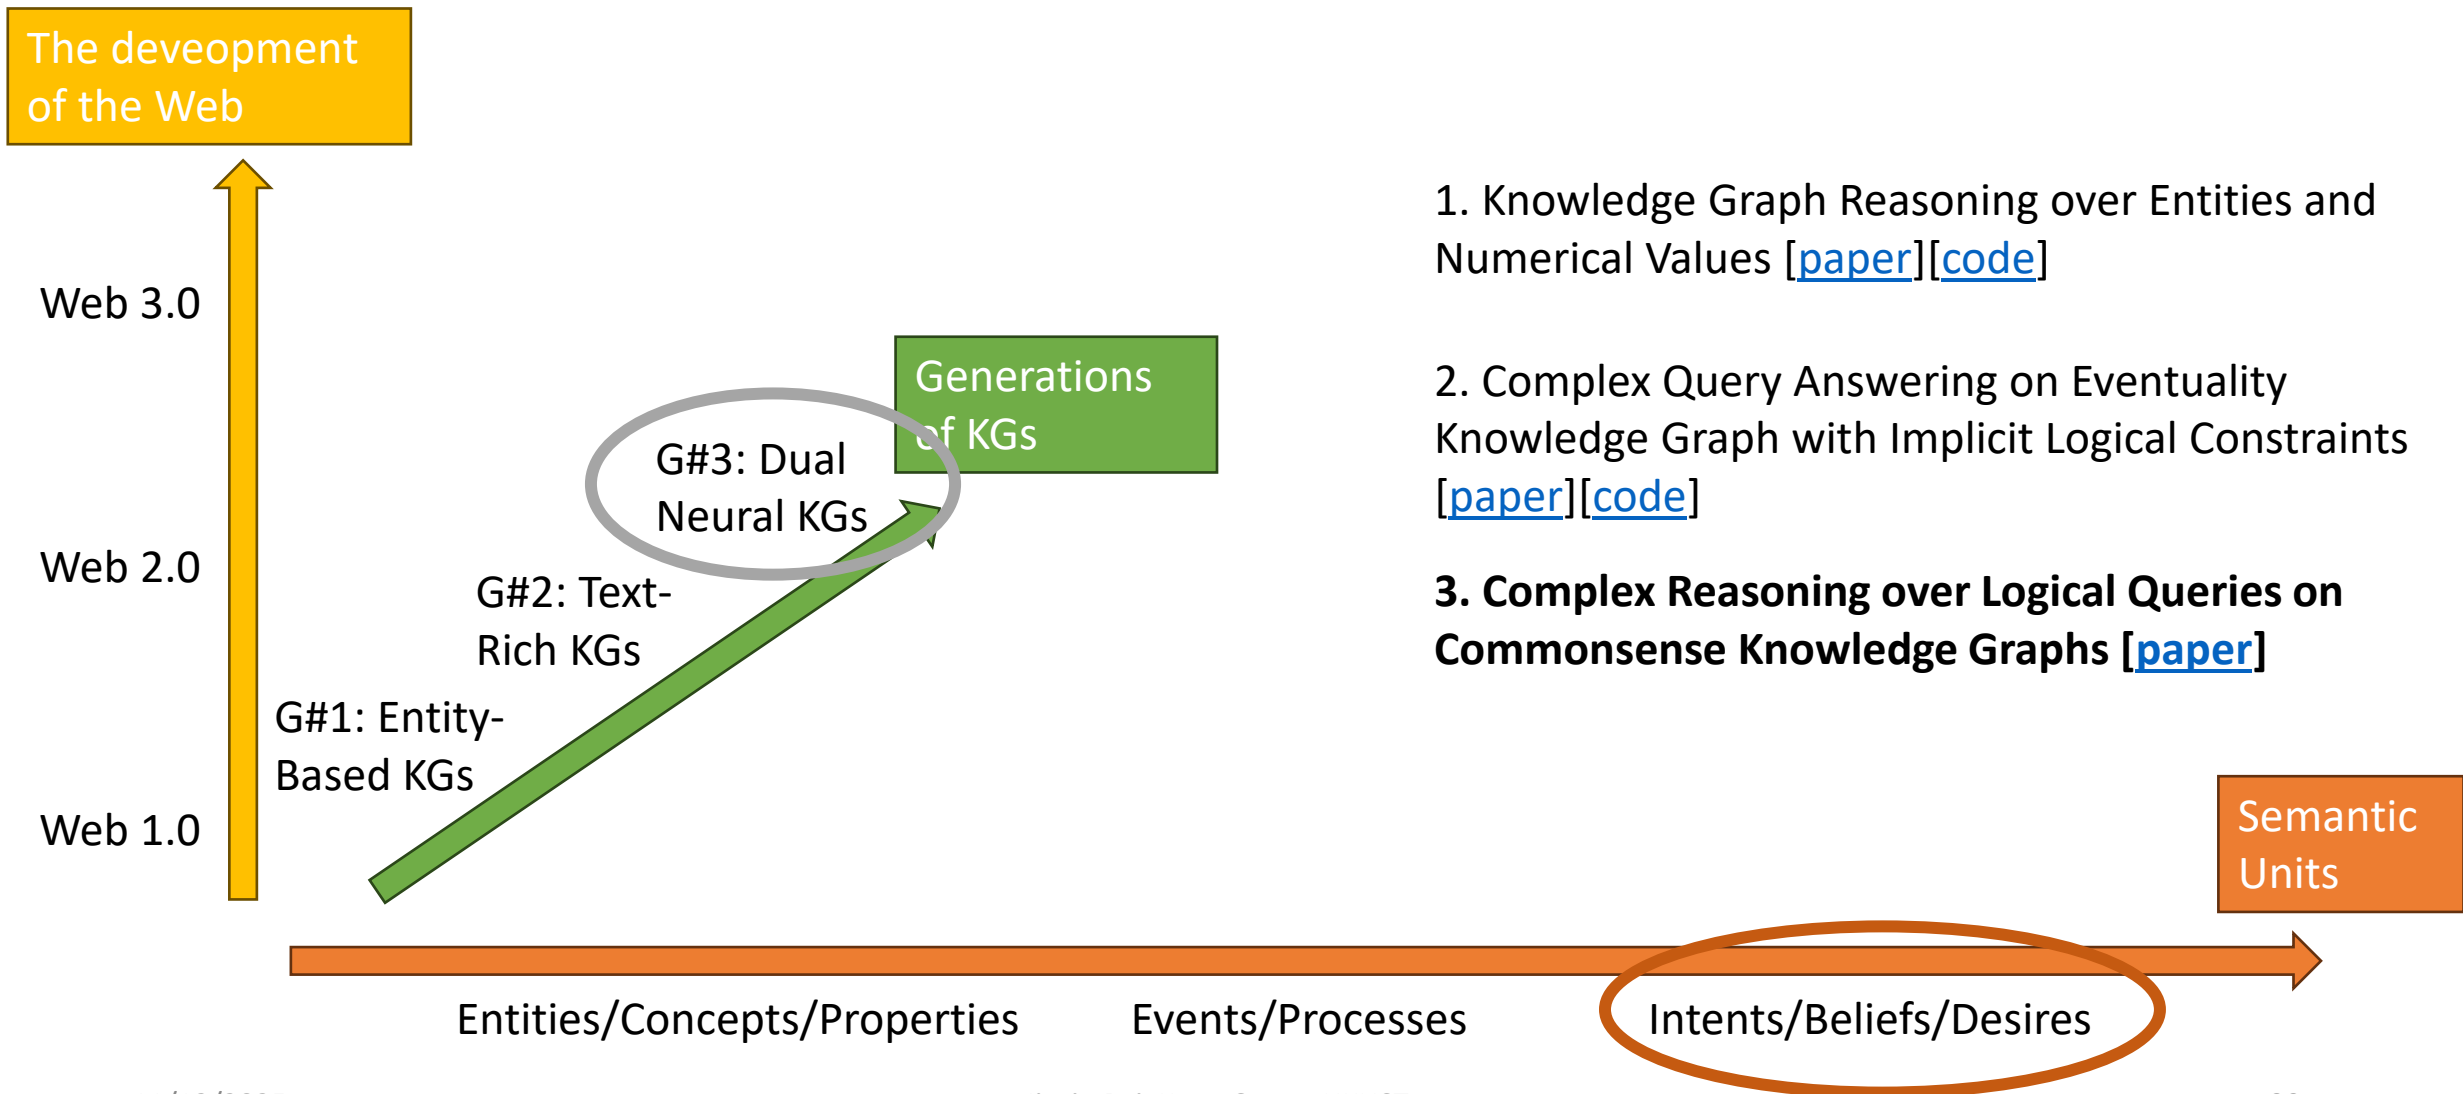
$$(2) v_i = \sum_{m=1}^M s_{i,m} (c_r^{(m)} + c_t^{(m)})$$

Computes the aggregated memory values across M memory cells with the importance weighted by relevance scores.

$$(3) q_i = q_i + \text{MLP}(v_i)$$

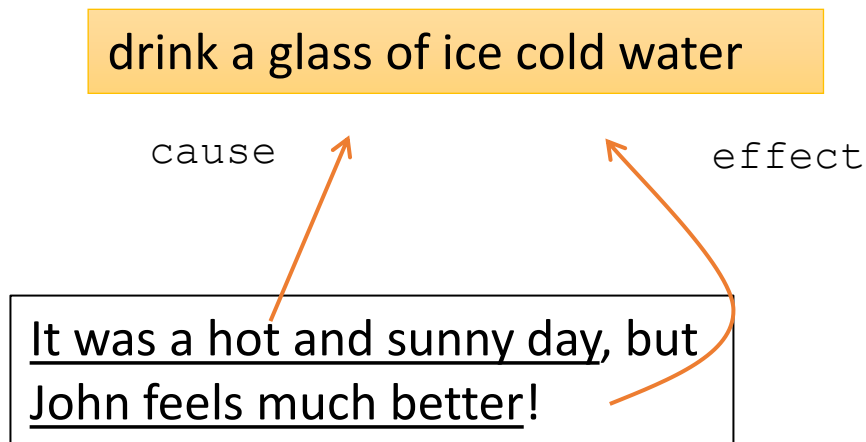
Computes the query embedding with memory values with the help of a MLP layer.

The Three-dimensional Development of Knowledge Graphs in Computer Science

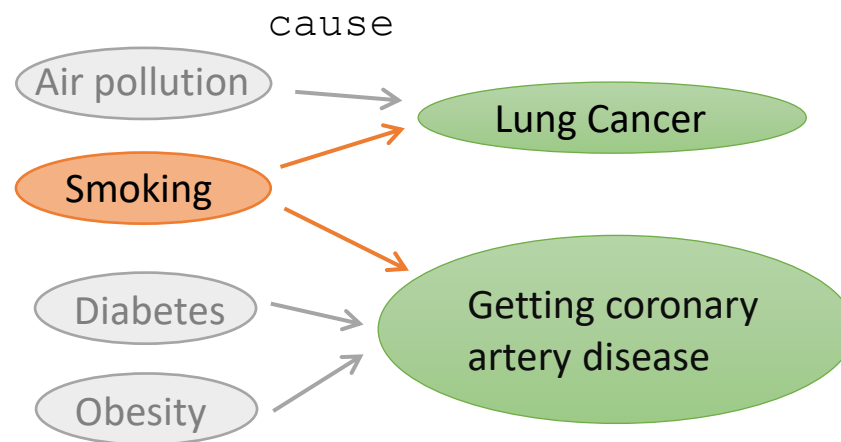


Moving towards intention based commonsense complex question answering

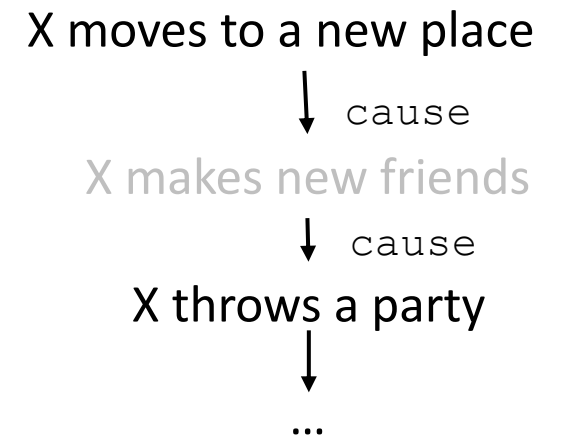
- Reasoning on real-world text and narratives requires complex reasoning over **multiple events**, and **inferring implicit context**.



Abduction



Common cause



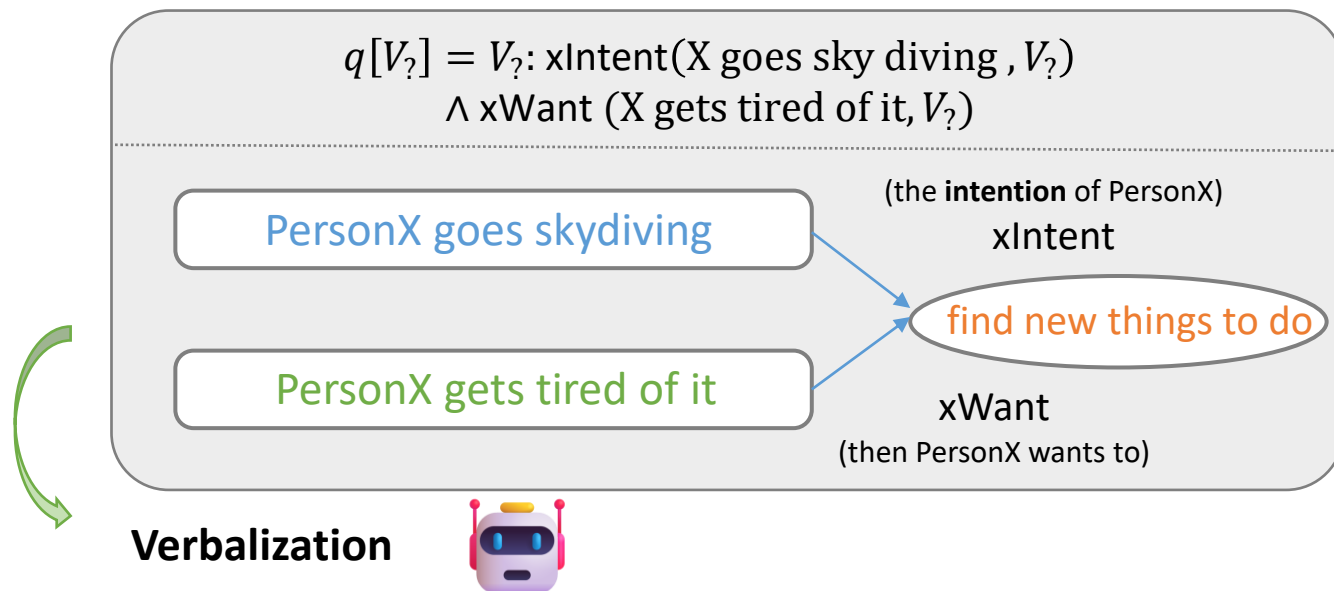
multi-hop effects

Complex Commonsense Reasoning

- Sampling conjunctive logical queries over existing CSKG

- Sampling synthetic query at scale
- Perform verbalization to make the query in natural language
- Define reasoning (question) based on the relations.

Base CSKG: ATOMIC2020



LLM-added
context

PersonX is living a boring life.

Rule-based
discourse

After getting tired of it
PersonX goes skydiving

Question:

What's both the intention of PersonX going skydiving
and what X wants to do after PersonX getting tired of it?

Answer:


find new things to do




Conjunctive Logical Queries

- Similar to CQA problems


Training Query Types




Unseen Query Types




2i: Common Attribution




2i-negative: Negated Common Cause



2p: 2nd order Effect



pi



Benchmarking

- Multiple-Choice QA:
 - Negative sampling: 2 are randomly sampled across CSKB, 2 are randomly sampled across one-hop answers (hard negative).
 - An additional “No answers are correct option”
- Generative complex commonsense reasoning
 - 53 annotators
 - Fleiss Kappa: 0.445; IAA: 78%
 - All disagreements are fixed by experts

#Train	790k
#Eval (annotated)	1,317
%wrong verbalization	4.5%
%sampled answers that are plausible	52.1%
%sampled negatives that are plausible	23.5%

LLMs Still Fall Short on This Task

Method	2i	2i-neg	3i	2p	ip	pi	All
API-based LLMs							
gpt-3.5-turbo-0613	33.56	43.12	42.01	38.66	38.05	28.40	37.74
- 1-shot	43.31	35.31	58.45	57.73	51.33	62.96	48.22
- 1-shot w/ CoT	45.80	36.43	54.34	57.73	50.44	66.67	48.75
- 8-shot (2i, 2p)	48.52	41.26	57.08	67.53	53.10	74.07	53.22
- 8-shot (2i, 2p) w/ CoT	52.61	46.10	60.27	59.79	52.21	65.43	54.37
gpt-4-1106-preview	44.67	46.47	52.05	32.47	40.71	53.08	44.64
- 1-shot	47.85	42.01	50.68	38.66	44.25	50.62	45.63
- 1-shot w/ CoT	48.97	46.46	52.96	49.48	52.21	58.02	50.04
- 8-shot (2i, 2p)	54.87	46.47	58.90	45.88	52.21	66.67	53.00
- 8-shot (2i, 2p) w/ CoT	57.82	49.07	62.56	61.34	52.21	66.67	57.40
Open-source (QA) Language Models							
HyKAS (Ma et al., 2021, zero-shot)	34.92	39.41	27.85	41.75	37.17	33.33	35.76
CAR (Wang et al., 2023a, zero-shot)	37.41	30.48	37.44	57.73	32.74	53.09	39.56
Llama2 (7B) (Touvron et al., 2023)	35.15	21.93	39.27	35.57	28.32	51.85	33.64
Vera (5B) (Liu et al., 2023)	47.62	27.51	40.18	66.49	52.21	58.02	46.09
UnifiedQA-v2 (Khashabi et al., 2022)	56.23	39.41	62.56	58.76	51.33	62.96	54.21
Flan-T5 (11B) (Chung et al., 2022)	58.28	47.21	65.30	76.29	56.64	79.01	60.97
Fine-tuned on COM²							
DeBERTa-v3-Large (+COM ²)	60.09	58.36	69.41	61.86	59.29	81.48	62.79
CAR-DeBERTa-v3-Large (+COM ²)	61.22	56.13	69.86	68.56	56.64	85.19	63.78

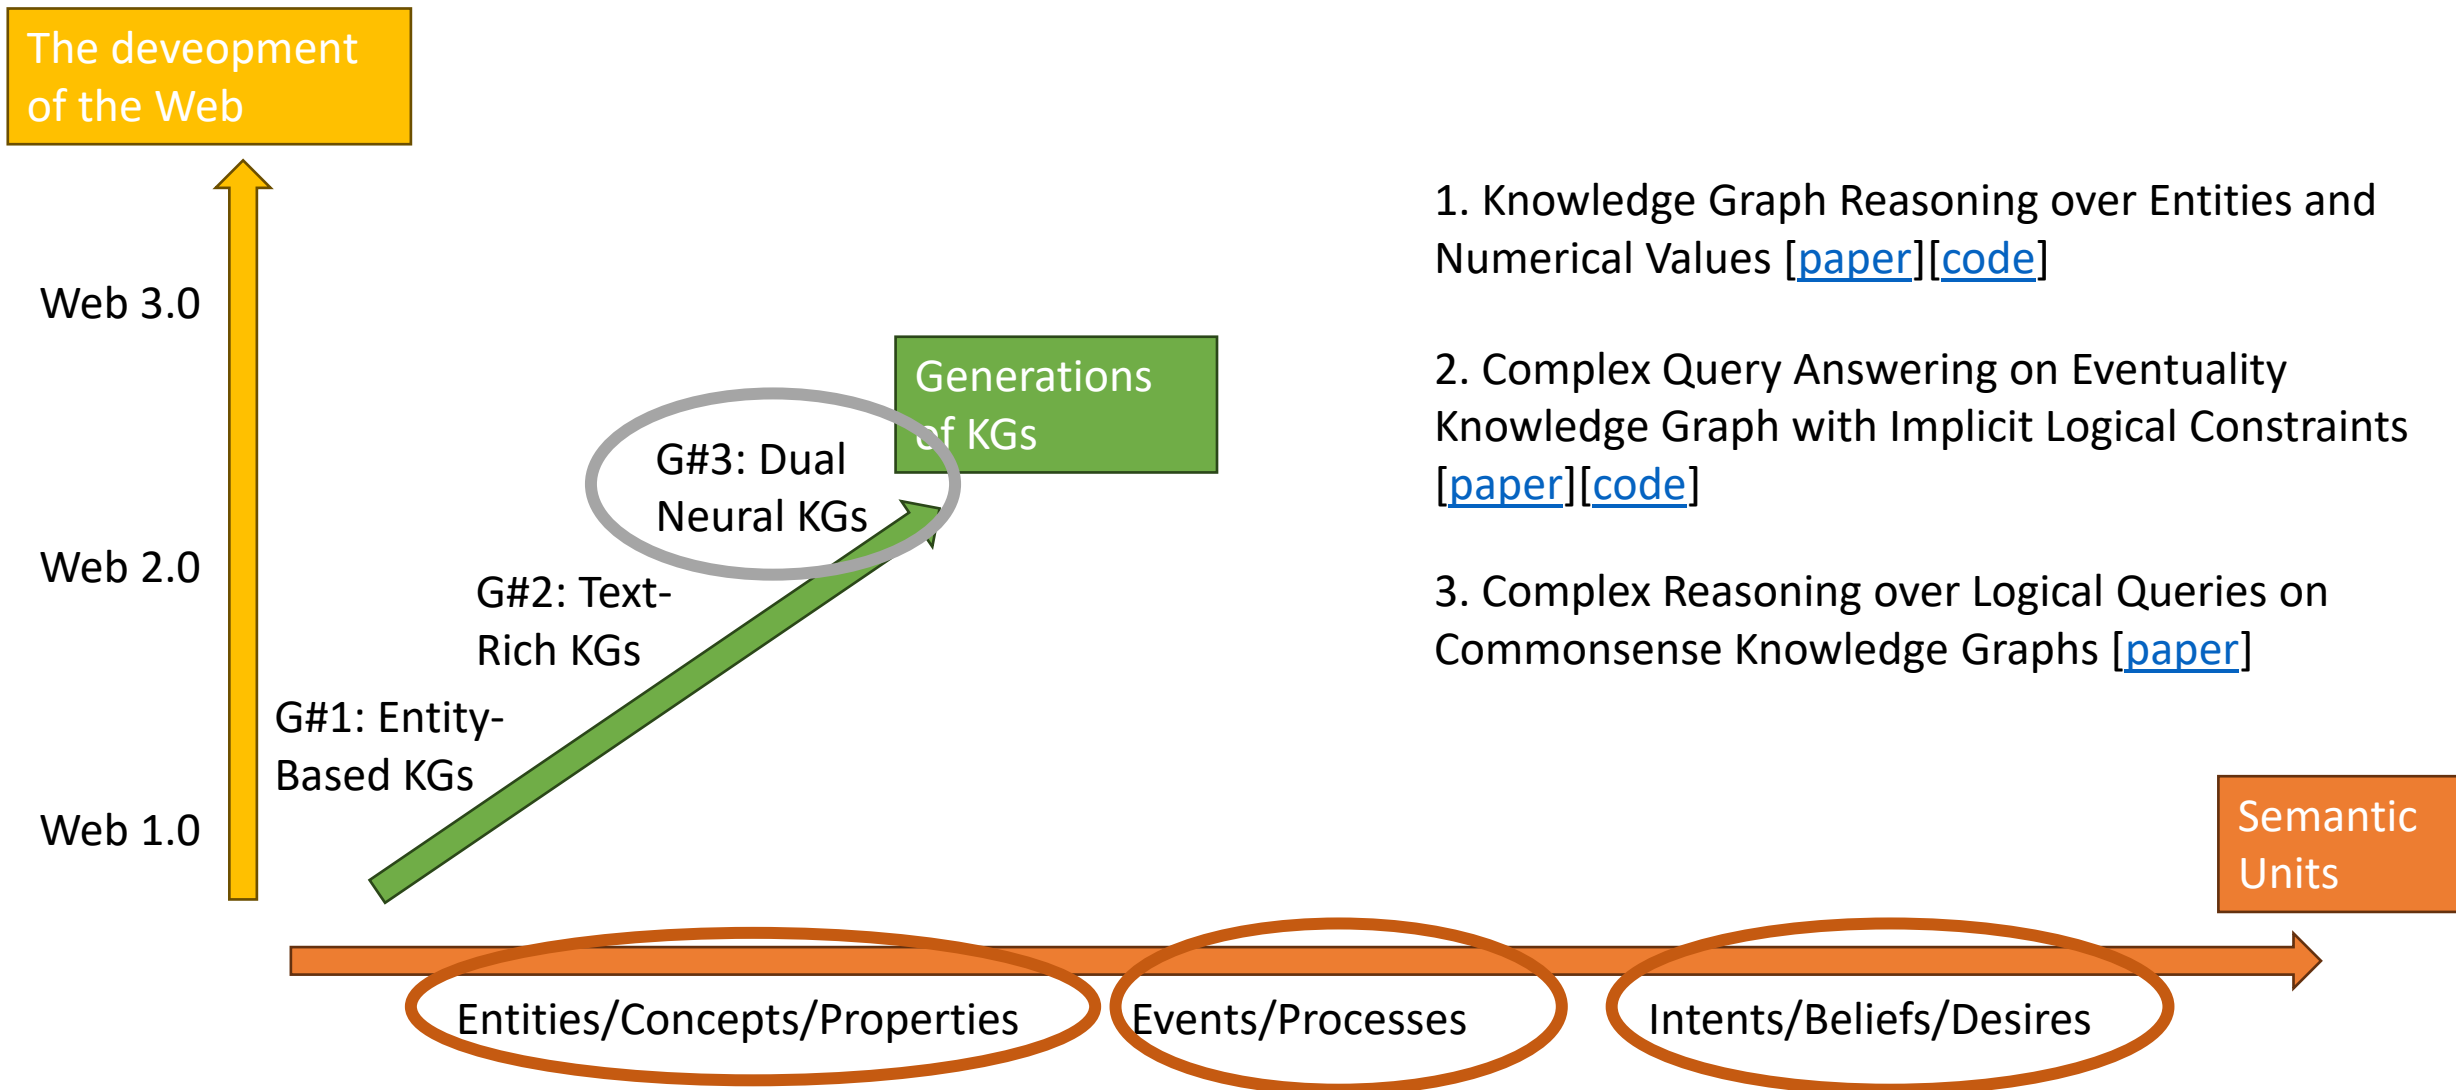
GPT:
37.74%~57.40%

Open LLMs:
33.64%~60.97%

Finetuned:
62.79%~63.78%

Table 1: Model performance (%) on the multiple-choice question answering evaluation set of COM².

The Three-dimensional Development of Knowledge Graphs in Computer Science



Conclusions for Part 4-2

- We have reviewed the frontier of recent development of knowledge graphs in terms methods of different semantic units:
 - Entity/attributes
 - Events/processes
 - Intent/desire/belief
- More recent advances in KG reasoning combined with LLMs will be introduced in Part 5 by Lihui Liu