CMSC 373 Artificial Intelligence
Fall 2023
09-Knowledge Representation

Deepak Kumar
Bryn Mawr College

1

Meaning Representation Systems

- Logic
- Semantic Networks
- Frames
- Conceptual Dependency
- many many others…
- Japan’s Fifth Generation Project, and several others
- Successes of KR&R
- Limitations of KR&R

2

2

Dichotomies of AI: Neats versus Scruffies

- **Neats**

  Representations that had the characteristic of mathematical precision and formalisms. Example, First-Order Predicate Logic and other logics.

- **Scruffies**

  Modeling knowledge in representations that were intuitive and psychologically meaningful to humans.

3

3

Frames

- An **attribute/value data structure** (aka **slot/filler**) that captures commonsense knowledge about stereotypical situations (e.g. a birthday party, eating at a restaurant, etc.). Introduced by Marvin Minsky (1971) and others.

  Example:

  1. John went to a restaurant.
  2. He asked the server for a hamburger.
  3. He paid, tipped, and left.

4

4

Frames

- An attribute/value data structure that captures commonsense knowledge about stereotypical situations (e.g. a birthday party, eating at a restaurant, etc.). Introduced by Marvin Minsky (1971) and others.

Example:

1. John went to a restaurant.
2. He asked the server for a hamburger.
3. He paid, tipped, and left.

John went to a restaurant.
He sat down.
He read a menu.
The server came to take his order.
He ordered a hamburger.
He ate the hamburger.
He paid the check.
He left a tip.
He left the restaurant.

5

5

Frames

Restaurant
Name: Gulliftys
Address: 1149 Lancaster Ave, Bryn Mawr, PA 19010
Phone: (610) 525-1851
Website: http://gulliftys.com
Cuisine Type: American
Seating Capacity: 120
…etc…

Menu
Salads: …
Main Dishes:
Entrée: Chicken Cacciatore, $23.00
Entrée: Hamburger, $12.00
Entrée: etc.
Desserts:
Dessert: Rice Pudding, $8.00
Dessert: Crème Caramel, $8.00
Beverages:
Soda: Coke, $3.00
Hot Beverage: Tea, $3.50

John went to Gulliftys.
He sat down.
He read a menu.
The server came to take his order.
He ordered a hamburger.
He ate the hamburger.
He paid the check.
He left a tip.
He left the restaurant.

6

6

Frames – Organized in a Hierarchy/Network

Restaurant
Name: Gulliftys
Address: 1149 Lancaster Ave, Bryn Mawr, PA 19010
Phone: (610) 525-1851
Website: http://gulliftys.com
Cuisine Type: American
Seating Capacity: 120
...etc...

isa → Restaurant

has ↓

Menu
Salads: ...
Main Dishes:
Entrée: Chicken Cacciatore, $23.00
Entrée: Hamburger, $12.00
Entrée: etc.
Desserts:
Dessert: Rice Pudding, $8.00
Dessert: Crème Caramel, $8.00
Beverages:
Soda: Coke, $3.00
Hot Beverage: Tea, $3.50

John went to Gulliftys.
He sat down.
He read a menu.
The server came to take his order.
He ordered a hamburger.
He ate the hamburger.
He paid the check.
He left a tip.
He left the restaurant.

7

Frames – Example Taxonomic Hierarchy

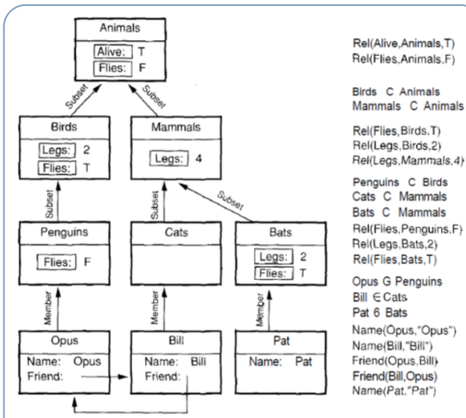

((a) A frame-based knowledge base (b) Translation into first-order logic

From: https://ui.kris.uniza.sk/ccat/knowledge-representation-methods/frame-based-representation/

8

Frames – Hey! This looks very familiar!!

Restaurant
Name: Gulliftys
Address: 1149 Lancaster Ave, Bryn Mawr, PA 19010
Phone: (610) 525-1851
Website: http://gulliftys.com
Cuisine Type: American
Seating Capacity: 120
…etc…

isa → Restaurant

has ↓

Menu
Salads: …
Main Dishes:
Entrée: Chicken Cacciatore, $23.00
Entrée: Hamburger, $12.00
Entrée: etc.
Desserts:
Dessert: Rice Pudding, $8.00
Dessert: Crème Caramel, $8.00
Beverages:
Soda: Coke, $3.00
Hot Beverage: Tea, $3.50

John went to Gulliftys.
He sat down.
He read a menu.
The server came to take his order.
He ordered a hamburger.
He ate the hamburger.
He paid the check.
He left a tip.
He left the restaurant.

9

9

Frames – Hey! This looks like OOP!!

Restaurant
Name: Gulliftys
Address: 1149 Lancaster Ave, Bryn Mawr, PA 19010
Phone: (610) 525-1851
Website: http://gulliftys.com
Cuisine Type: American
Seating Capacity: 120
…etc…

isa → Restaurant

has ↓

Menu
Salads: …
Main Dishes:
Entrée: Chicken Cacciatore, $23.00
Entrée: Hamburger, $12.00
Entrée: etc.
Desserts:
Dessert: Rice Pudding, $8.00
Dessert: Crème Caramel, $8.00
Beverages:
Soda: Coke, $3.00
Hot Beverage: Tea, $3.50

John went to Gulliftys.
He sat down.
He read a menu.
The server came to take his order.
He ordered a hamburger.
He ate the hamburger.
He paid the check.
He left a tip.
He left the restaurant.

10

10

Frames & OOP

- While the idea of Frames came out of AI, the idea of OOP came out of Software Engineering. Some people claim the idea came from AI Frames (myth or reality??).
- Frames are similar to class hierarchies in OOP. But the focus is on intuitive representation of knowledge.
- OOP is primarily focused on encapsulation and information hiding.
- Both are modeling frameworks.

11

11

Frame Languages

- Conceptual Dependency, 1973
- KRL, 1977
- FRL, 1977
- OWL, 1977
- NETL, 1979
- KL-ONE, 1982
- KRYPTON, 1983
- FRAIL, 1983
- KODIAK, 1984
- UniFrame, 1984
- More recently, ontologies for Semantic Web…(another OWL, 2005)

12

12

Conceptual Dependency (Schank, 1969)

- A meaning representation theory that captures the content of linguistic communication.
- Events: physical events, mental events, intentions, physical causes, reasons (mental causes), etc.

  Every EVENT has
  - an ACTOR
  - an ACTION performed by ACTOR
  - an OBJECT that ACTION is performed upon
  - a DIRECTION in which ACTION is oriented
- Events are organized into **scripts** and **plans**.

13

13

Example CD Script

- Every EVENT has
  - an ACTOR
  - an ACTION performed by ACTOR
  - an OBJECT that ACTION is performed upon
  - a DIRECTION in which ACTION is oriented

John went to New York.	[ACTOR: John, ACTION: Go, OBJJECT: John?, DIRECTION: New York]
John likes to drink.	[ACTOR: John, ACTION: drink, OBJECT: (beverage)*, DIRECTION: (towards John's mouth)]
John fell.	[ACTOR: (gravity), ACTION: fell, OBJECT: John, DIRECTION: (towards ground)]

- Some slots must be filled or inferred (from existing knowledge)

14

14

CD – Primitive Actions

NAME	Description	Example
ATRANS	Transfer of abstract relationship (ownership)	give, buy
PTRANS	Physical location transfer	go
MTRANS	Transfer of Information	read, tell, see, hear
PROPEL	Application of force	push, throw,
MBUILD	Input from some memory and output to memory	wonder, figure out
ATTEND	Action of attending on something	reading a book
SPEAK	An utterance	said, told
GRASP	Actor grasping an object	clutch
MOVE	Move a body part	kick, grab
INGEST	Ingesting something	drink, eat, breathe
EXPEL	Performing a bodily function	…

CD – Representations & Inferences

- *John went to New York.*

  ACTOR: John
  ACTION: PTRANS
  OBJECT: John
  DIRECTION TO: New York
  FROM: unknown

- John bought a book from Mary.

  ACTOR: John
  ACTION: ATRANS
  OBJECT: money
  DIRECTION TO: John
  FROM: Mary

- John read a book.

  ACTOR: John
  ACTION: ATTEND
  OBJECT: eyes
  DIRECTION TO: book
  FROM: unknown

- John drank a glass of milk.

  ACTOR: John
  ACTION: INGEST
  OBJECT: milk
  DIRECTION TO: mouth of John
  FROM: glass

Instrument:
  ACTOR: John
  ACTION: PTRANS
  OBJECT: glass containing milk
  DIRECTION TO: mouth of Johnk
  FROM: table

Instrument:
  ACTOR: John
  ACTION: MOVE
  OBJECT: hand of John
  DIRECTION TO: glass
  FROM: unknown

Instrument:
  ACTOR: John
  ACTION: GRASP
  OBJECT: glass of milk
  DIRECTION TO: hand of John
  FROM: unknown

Semantic Networks

- Represents semantic relations between concepts in a network.

  It is a **graph**
  vertices represent **concepts**
  edges represent **semantic relations**

- Also represented as a **triple**:

  <entity, relation, entity>
  <mammal, has, vertebra>
  <bear, is a, mammal>
  <fish, lives in, water>

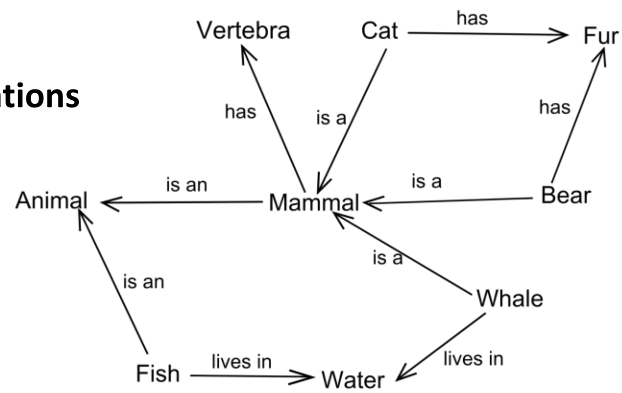

From: https://en.wikipedia.org/wiki/Semantic_network#/media/File:Semantic_Net.svg

17

17

SNePS – Semantic Network & Logic

- <u>**S**</u>emantic <u>**Ne**</u>twork
  <u>**P**</u>rocessing <u>**S**</u>ystem

- Represents:

  All humans are mortal.
  Socrates is a human.

- Can be used to perform logical inference.

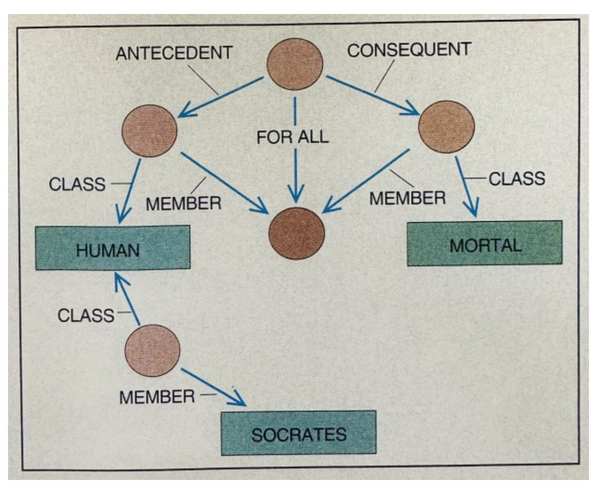

From: Paul Wallich, Silicon Babies, Scientific American, December 1991.

18

18

SNePS – Semantic Networks & Logic

- **S**emantic **Ne**twork **P**rocessing **S**ystem
- Represents:

  All humans are mortal.
  Socrates is a human.

- Can be used to perform logical inference.

  Socrates is mortal.

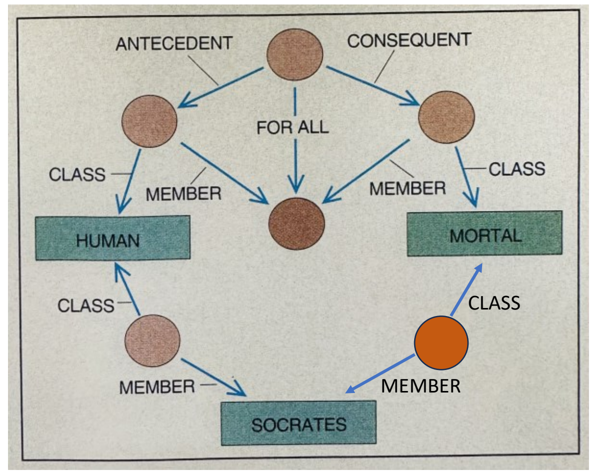

From: Paul Wallich, Silicon Babies, Scientific American, December 1991.

19

19

Example AI System (1993)

- Tell-Ask Systems

SNePS Semantic Network
+
Inference Engine

Tell

Blocks are supports. Red colored blocks are wooden. Picking up is a primitive action. Putting is a primitive action. Before picking up a block the block must be clear. After picking up a block the block is not clear and it is held. If a block is on a support then after picking up the block the support is clear and the block is not on the support. Before putting a block on a support the block must be held and the support must be clear. After putting the block on a support the block is not held, the block is clear, and the block is on the support. After putting a block on another block the latter block is not clear.

There is table. The table is a support. A, B, and C are blocks. Ais clear and it is on the table. B is clear and it is on the table. C is clear and it is on the table.

Grammar to parse all English sentences.
Semantic Interpretation.
Representing semantic interpretation as a semantic network (+ logic).
Grammar to generate English sentences from network representations.

20

20

Inference in Service of Acting (1993)

- Most KR Systems are designed to do reasoning.
- Acting is modeled by planning systems.
- Consider this:

  Red colored blocks are wooden.
  A is a block.

  Is A wooden?

Tell: All blocks are supports.
Tell: A is a block.

Ask: Is A a support?

I wonder if A is a support.
I wonder if A is a block.
I know A is a block.
Since A is a block and blocks are supports
I infer A is a support.

Tell: Pick up A.

I intend to do the act Pickup A.
I wonder if the act Pickup A has any preconditions.
I wonder if A is a block.
I know A is a block
and before picking up A it must be clear.
The act Pickup A has the following precondition:
A is clear.
I wonder if A is clear.
I know A is clear.
It is satisfied.
Now doing: Pickup A.

21

21

Inference in Service of Acting (1993)

Note: Explanations!

- Most KR Systems are designed to do reasoning.
- Acting is modeled by planning systems.
- Consider this:

  Red colored blocks are wooden.
  A is a block.

  Is A wooden?

Tell: All blocks are supports.
Tell: A is a block.

Ask: Is A a support?

I wonder if A is a support.
I wonder if A is a block.
I know A is a block.
Since A is a block and blocks are supports
I infer A is a support.

Do: Pick up A.

I intend to do the act Pickup A.
I wonder if the act Pickup A has any preconditions.
I wonder if A is a block.
I know A is a block
and before picking up A it must be clear.
The act Pickup A has the following precondition:
A is clear.
I wonder if A is clear.
I know A is clear.
It is satisfied.
Now doing: Pickup A.

22

22

Acting in Service of Inference (1993)

```
Tell: Red colored blocks are wooden.
Tell: A is a block.

Ask: Is A wooden?
```

I wonder if A is wooden.
I wonder if A is colored red.

I wonder if A is a block.
I know A is a block.

I wonder if A is red.
I don't know if A is red.

I don't know if A is wooden.

23

23

Acting in Service of Inference (and *vice versa)* (1993)

```
Tell: Red colored blocks are wooden.
Tell: A is a block.

Ask: Is A wooden?
```

I wonder if A is wooden.
I wonder if A is colored red.

I wonder if A is a block.
I know A is a block.

I wonder if A is red.
I don't know if A is red.

I don't know if A is wooden.

\+

```
Tell: Looking is a primitive action.
Tell: If you want to know the color
      of something, look at it.
```

24

24

Acting in Service of Inference (1993)

```
Tell: Red colored blocks are wooden.
Tell: A is a block.

Ask: Is A wooden?
```

I wonder if A is wooden.
I wonder if A is colored red.

I wonder if A is a block.
I know A is a block.

I wonder if A is red.
I don't know if A is red.

I don't know if A is wooden.

+

```
Tell: Looking is a primitive action.
Tell: If you want to know the color
      of something, look at it.
```

+

```
Tell: Red colored blocks are wooden.
Tell: A is a block.

Ask: Is A wooden?
```

I wonder if A is wooden.
I wonder if A is colored red.

I wonder if A is a block.
I know A is a block.

Since A is a block I infer
If you want to know the color of a block A look at it.

I intend to do the act look at A.
I wonder if act look at A has any preconditions.
…
Now doing: look at A.
Sensory-add: A is colored red.

Since A is a block and A is colored red
and all red colored blocks are wooden
I infer A is wooden.

25

25

Large-Scale AI/KR&R Efforts, 1980s-1990s

- Japanese Fifth Generation Project
- United States
  Strategic Computing Initiative (SGI)
  Microelectronics and Computer Technology Corporation (MCC)
- United Kingdom Alvey Project
- European Strategic Program for Research in Information Technology (ESPRIT)
- European Computer Industry Research Centre (ECRC)
  In collaboration with ICL (UK), Bull (France), Seimens (Germany)

26

26

Large-Scale AI/KR&R Efforts

- **Japanese Fifth Generation Project**

  10-year project started in 1982.
  Creating supercomputers for future AI development using PROLOG.
  Develop Knowledge Information Processing Systems.
  Spent ~$300 million over twelve years.
  (IBM's research expenditure in 1982 was $1.5 billion!)

- **Limited Results**

  Developed foundations for concurrent logic programming.
  Developed Kappa (a parallel DBMS)
  Developed an automated theorem prover, MGTP
  Attempt at applications in bioinformatics.

- **Did not meet commercial success**

  The arrival of SUN Workstations and x86 based machines far surpassed the abilities.
  Failed on attempts to develop concurrent KR&R systems.

- **Perhaps the ideas were far ahead of its time.**

  For example, current multi-core processors being used in current AI work.

27

27

CYC – The Ultimate Expert System

- Project started in 1984 by Doug Lenat at MCC (!!)
- Since 1995 has been under active development by Cycorp.

  Recognized at Knowledge-based AI projects, while having good early results, failed to scale-up.

- Hypothesis: The problem of general AI was primarily one of knowledge and that it could be solved by a suitable, large knowledge system by representing millions of pieces of human knowledge that is considered commonsense.

28

28

CYC – The Ultimate Expert System

- Methodology

  Developed a representation language, CycL.
  Developed a set of representations (ontological engineering)
  Developed a massive knowledgebase comprising human knowledge.
  Connected CYC's knowledge to Wikidata (Wikipedia) and other large knowledgebases.
  An inference engine.

29

29

CYC – The Ultimate Expert System

- While it could successfully answer many questions, it failed on these types of questions in a demo in 1994:

  If Tom is 3 inches taller than Dick, and Dick is 2 inches taller than Harry, how much taller is Tom than Harry?

  Can two siblings each be taller that the other?

  Which is wetter, land or sea?

  Can a car drive backward? Can a plane fly backward?

  How long can people go without air?

30

30

CYC – The Ultimate Expert System

- While it could successfully answer many questions, it failed on these types of questions in a demo in 1994:

  If Tom is 3 inches taller than Dick, and Dick is 2 inches taller than Harry, how much taller is Tom than Harry?

  Can two siblings each be taller that the other?

  Which is wetter, land or sea?

  Can a car drive backward? Can a plane fly backward?

  How long can people go without air?

- Can ChatGPT answer any of these correctly? Would you then consider it intelligent??

31

31

CYC – The Ultimate Expert System

- In 2002 OpenCyc was relesed to the public. It contained 6000 concepts and 60,000 facts.
- A version of OpenCyc was released in 2012. It contained 239,000 concepts and 2,093,000 facts.
- There have been some successful applications of Cyc in pharmaceuticals, terrorism knowledgebases, 6^{th} grade level math.
- Overall considered a controversial project. Another AI idea “ahead of its time”.
- “…the sad fact is that Cyc’s main role in AI history is an extreme example of AI hype, which very publicly failed to live up to the grand predictions that were made for it.”
  -: Michael Wooldrdge

32

32

Knowledge-based AI – An Epitaph???

- "The project of using logical languages to represent commonsense knowledge for AI systems is now sixty years old. Though the literature is large, the fraction of commonsense knowledge covered is extremely small, and the impact of this approach on practical AI technology is small and grows steadily smaller as AI research becomes ever more dominated by corpus-based learning techniques…"

  -: Ernest Davis, 2017.

33

33

Knowledge-based AI – ~~An Epitaph???~~

- "The fact remains, however, that there are many forms of commonsense knowledge for which the *only* epistemically adequate representation known is some form of logical language,…[and cannot be carried out by any known formalism]. It therefore still remains likely that the study of logical representation of commonsense knowledge will ultimately play an important role in the development of human level AI."

  -: Ernest Davis, 2017.

- Knowledge-based systems can **intrinsically explain** the basis for all their inferences.

  Since Socrates in human and all humans are mortal,
  I infer Socrates is mortal.

34

34

1988-1990s Second AI Winter sets in…

35

35

The Seasons of AI

- **1950s – 1966 First AI Summer: Irrational Exuberance**

  Early successes in game playing, theorem proving, problem solving

- **1967 – 1977 First AI Winter**

  No useful deliverables led to loss of research funding and cancellation of AI programs. In UK *The Lighthill Report* (toy AI systems do not scale due to combinatorial explosion).

- **1978 – 1987 Second AI Summer/Spring**

  Rise of knowledge-based systems, success of Expert Systems. Boom times.

- **1988 – 1993 Second AI Winter**

  Failure of AI Hardware companies (Symbolics, LMI, Lisp Machines) and AI Companies (Teknowledge, Inference Corp. etc.) Commercial deployments of Expert Systems were discontinued.

- **1993 – 2011 Third AI Summer (Mostly academic advances)**

  Statistical approaches and extensions to logic (Bayesian Nets), Non-Monotonic Reasoning (in Logic), Fuzzy Logic, advances in Machine Learning (Decision Trees, Random Forests, Neural Nets), Cognitive Models, Logic Programming, Case-Based Reasoning, Genetic Algoritms, Agent-based approaches, etc.

- **2011 – Now Third AI Spring**

  Rise of Deep Learning, Neuro-symbolic AI, ChatGPT and other chatbots, generative AI.

36

36

Towards Usable Representations (1985-now)

- **Wordnet**, 1985 (Princeton U.)

  A lexical database of semantic relations between words.

  Over 150,000 organized in 207,000 word-sense pairs (eat-out, car-pool).

  Now available for multiple languages

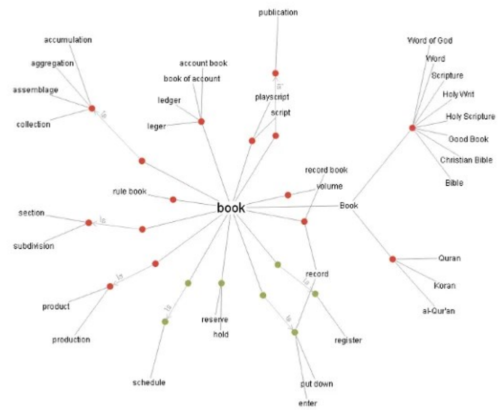

37

37

Wordnet Knowledge Structure

- dog, domestic dog, Canis familiaris canine, canid
  - carnivore
    - placental, placental mammal, eutherian, eutherian mammal
      - mammal
        - vertebrate, craniate
          - chordate
            - animal, animate being, beast, brute, creature, fauna
              - ...

38

38

Knowledge Graphs – The Mother of all Semantic Networks!

- Instead of hand coding a semantic network, can we extract knowledge from large corpora (text, as on the web) and structure it in a usable form? Currently, mostly from proprietary structured data or human coding in specific domains.
- Knowledge Graphs were introduced by Google in 2012
- Knowledge Graphs represent relationships between entities. It is an organized collection of interrelated facts or concepts.
- A relationship is represented as a collection of triples

  <*s*, *p*, *o*>: Object, *o* is related to subject, *s* via a relation *p* (predicate)

  e.g. <:deepak_kumar, foaf:name, 'Deepak Kumar'>

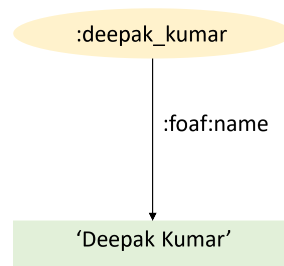

39

39

An Example KG

Concepts, relations etc. in a KG can be obtained from and linked to several disparate structured Sources of data:

ConceptNet
ATOMIC
Wikidata
WordNet
Roget
VerbNet
FrameNet
VisualGenome
ImageNet

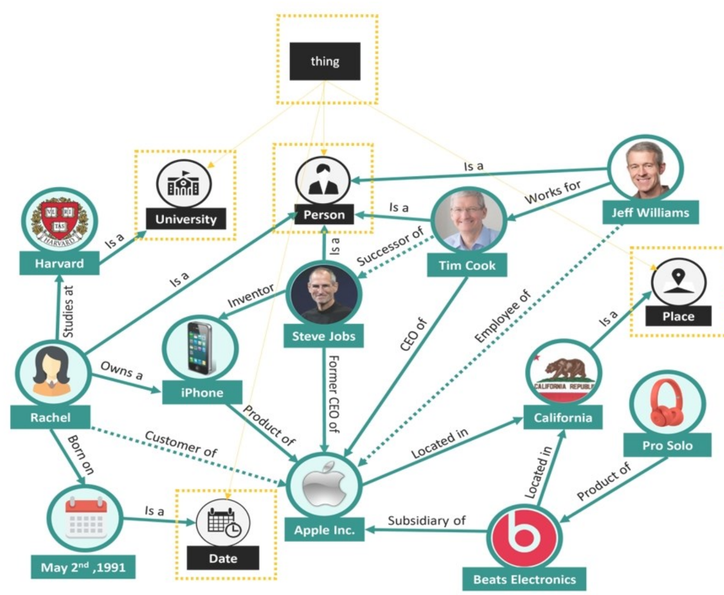

From: https://atonce.com/blog/knowledge-graph

40

40

Commonsense Knowledge

On stage, a woman takes a seat at the piano. She

1. sits on a bench as her sister plays with a doll.
2. smiles with someone as the music plays.
3. is in the crowd, watching the dancers.
4. nervously sets her fingers on the keys.

Which one?

41

41

Commonsense Knowledge

On stage, a woman takes a seat at the piano. She

1. sits on a bench as her sister plays with a doll.
2. smiles with someone as the music plays.
3. is in the crowd, watching the dancers.
4. nervously sets her fingers on the keys.

Which one?

Answering the question requires knowledge that humans possess and apply, but machines cannot distill from the communication.

Also, remember Winograd schemas?

42

42

Capturing Commonsense Knowledge

On stage, a woman takes a seat at the piano. She

1. sits on a bench as her sister plays with a doll.
2. smiles with someone as the music plays.
3. is in the crowd, watching the dancers.
4. nervously sets her fingers on the keys.

Lexical Nodes
Synset Nodes
Frame Nodes

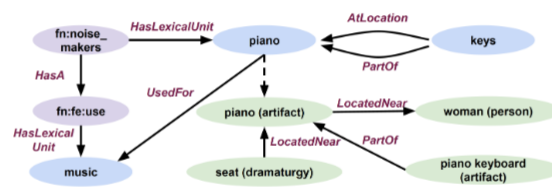

43

43

Some KG Applications

- Google's display of related topics in search results
  See: https://blog.google/products/search/introducing-knowledge-graph-things-not/
- Airbnb uses them to suggest popular areas or neighborhoods
  See: https://medium.com/airbnb-engineering/contextualizing-airbnb-by-building-knowledge-graph-b7077e268d5a
- Amazon's Product Graph
  See: https://www.aboutamazon.com/news/innovation-at-amazon/making-search-easier
- Many uses in machine learning

44

44

Some KG Applications

- Google's display of related topics in search results
  See: https://blog.google/products/search/introducing-knowledge-graph-things-not/
- Airbnb uses them to suggest popular areas or neighborhoods
  See: https://medium.com/airbnb-engineering/contextualizing-airbnb-by-building-knowledge-graph-b7077e268d5a
- Amazon's Product Graph
  See: https://www.aboutamazon.com/news/innovation-at-amazon/making-search-easier
- Many uses in machine learning

KR is not dead!!

45

45

References

- Stuart C. Shapiro & William J. Rapaport, *SNePS Considered as a fully Intensional Propositional Semantic Network*, in Proceedings of AAAi-86. AAAI Press, 1986.
- Roger Schank & Christopher Riesbeck, *Inside Computer Understanding*, Lawrence Erlbaum & Associates, 1981.
- Paul Wallich, *Silicon Babies*. Scientific American, December 1991.
- Anthony Maida, *Frames*, entry in Encyclopedia of AI, Volume 1, edited by Stuart C. Shapiro, John Wiley, 1992.
- Deepak Kumar, *Acting in Service of Inference (and vice versa)*, In Proceedings of FLAIRS-94, Douglas Dankel (editor), 1994.
- Ernest Davis, *Logical Formalizations of Commonsense Reasoning: A Survey*, Journal of AI Research, Number 59, pages 651-723, 2017.
- M. Wooldridge: *A Brief History of Artificial Intelligence*. Flatiron Books, 2020.
- Stuart Russell & Peter Norvig, *Artificial Intelligence: A Modern Approach*, Pearson, 2021.
- Filip Ilievski, Pedro Szkely, Bin Zhang, *CSKG: The Commonsense Knowledge Graph*, arXiv:2012.11490, 2021.
- Mayank Kejriwal, Craig A. Knoblock, PedroSzekely, *Knowledge Graphs*, MIT Press, 2021.

46

46