The Application of Qualitative Metadata to Analogical Reasoning

Dave Raggett

W3C/ERCIM, Sophia Antipolis, France

Abstract

Analogical reasoning can be used for plausible inferences based upon direct similarities or structural mappings involving properties and relationships. This can be implemented on top of a combination of symbolic knowledge plus sub-symbolic qualitative metadata, with matching based upon structural or causal similarities, and noticing interesting differences, in essence, abstracting from similarities and dissimilarities, and will be applied to examples of the form "A is to B as C is to ?X". A further challenge is to support the use of literal and figurative analogies in natural language, e.g., comparing life to the wheel of fortune, when you want to highlight the role of chance. An easy-to-use syntax will be presented for expressing knowledge, along with a web-based proof of concept demonstrator, and a unifying cognitive architecture for human-like AI. This builds upon pioneering work by Alan Colins on plausible reasoning, and Dedre Gentner on analogies.

Keywords

Plausible reasoning, Human-like AI, analogies, Cognitive Architecture

1. Introduction

The paper starts with an introduction to plausible reasoning before moving on to analogical reasoning and how this can be supported as an extension of plausible reasoning. This very much work in progress, and part of a long term drive to realise human-like memory, reasoning and learning in cognitive agents.

2. Plausible Reasoning

We are learning all the time, and revising our beliefs and understanding as we interact with others. As such our knowledge is imperfect and subject to uncertainties, incompleteness and inconsistencies. This is challenging both for conventional mathematical logic, and for statistical approaches such as Bayesian inference due to the difficulties in obtaining the required statistics. Evolution has equipped humans with the means to deal with imperfect knowledge in a rational way based upon sound judgement, albeit subject to various kinds of cognitive biases, see, e.g., Daniel Kahneman [1].

ICCBR Analogies'22: Workshop on Analogies: from Theory to Applications at ICCBR-2022, September, 2022, Nancy, France

dsr@w3.org (D. Raggett)

People have studied the principles for plausible arguments since the days of Ancient Greece, e.g., Carneades and his guidelines for argumentation. This was developed further by a long line of philosophers, including Locke, Bentham, Wigmore, Keynes, Wittgenstein, Pollock and many others. Plausible reasoning is everyday reasoning, and the basis for legal, ethical and business discussions. It is now timely to exploit plausible reasoning with imperfect knowledge in support of human-machine cooperative work. This will enable computers to analyse, explain, justify, expand-upon and argue in human-like ways.

Consider $A \to B$, which means if A is true then B is true. If A is false then B may be true or false. If B is true, we still can't be sure that A is true, but if B is false then A must be false. We can go further with a little knowledge. Consider a more concrete example: if it is raining then it is cloudy. This can be used for inferences in both directions. Rain is more likely if it is cloudy, and likewise, if it is not raining, then it might be sunny, so it is less likely that it is cloudy. Such arguments draw upon qualitative terms in lieu of quantitative statistics.

In essence, plausible reasoning draws upon prior knowledge as well as on the role of analogies, and the consideration of examples as precedents. Mathematical proof is replaced by reasonable arguments, both for and against a premise, along with how these are assessed. In legal proceedings, for instance, arguments are laid out by the Prosecution and the Defence, the Judge decides what evidence is admissible, whilst guilt is assessed by the Jury.

During the 1980's Alan Collins and his co-workers developed a theory of plausible reasoning [2] based upon analysis of recordings of how people reasoned. They found that:

- There are several categories of inference rules that people commonly use to answer questions.
- People weigh the evidence that bears on a question, both for and against, rather like in court cases.
- People are more or less certain depending on the certainty of the premises, the certainty of the inferences, and whether different inferences lead to the same or oppositive conclusions.
- Facing a question for which there is an absence of directly applicable knowledge, people search for other knowledge that could help given potential inferences.

Plausible knowledge can be expressed using a combination of symbolic graphs and associated metadata. This paper introduces the plausible knowledge notation (PKN) as an easy-to-read extensible syntax accompanied with an implementation as a JavaScript library for use in web page demos for different kinds of plausible reasoning [3], and as part of work for the W3C Cognitive AI Community Group [4]. PKN supports a variety of different kinds of statements:

Properties

```
flowers of England includes daffodils, roses (certainty high)
```

where `flowers` is a property of the referent `England`, and the use of `includes` signifies that the property is an open set with values `daffodils` and `roses`. For a closed set, use `is` instead of `includes`. Trailing round brackets are used to list qualitative metadata, in this case declaring that the statement has a high certainty.

Relationships

```
robin kind-of songbird
duck similar-to goose for habitat
duck dissimilar-to goose for neck-length
```

where `robin` is declared as a subclass of `songbird`, and `duck` is declared as being similar to `goose` for `habitat` and dissimilar to `goose` in respect to `neck length`.

Dependencies

```
climate depends-on latitude
pressure decreases-with altitude
current increases-with voltage
```

where `climate` depends on `latitude` in some unspecified way, whilst `pressure` decreases with increasing `latitude` and `current` increases with increasing `voltage`.

Implications

```
temperature of ?place is warm &
rainfall of ?place is heavy
implies grain of ?place includes rice
```

Implications are a form of if-then rules where variables are prefixed with a question mark.

Metadata can be given with all kinds of PKN statements. Relationships, dependencies and implications can be used for inferences in both directions, subject to any associated metadata. Following Collins, PKN supports several kinds of statement metadata relevant to different kinds of inferences:

Typicality in respect to other group members, e.g., robins are typical song birds.

Similarity to peers, e.g., having a similar climate.

Strength as conditional likelihood, e.g., the strength of climate for determining which kinds of plants grow well. The forward and backward strengths may differ, e.g., rain is a strong indicator of cloudy weather, whilst cloudy weather is a weak indicator of rain.

Frequency as the proportion of children with a given property, e.g., most species of birds have the ability to fly.

Dominance as the relative importance in a given group, e.g., the size of a country's economy.

Multiplicity as the number of items in a given range, e.g., how many different kinds of flowers grow in England.

The web demonstrator [3] allows you to pick from an assortment of queries, and to then see a trace of the reasoning, proceeding from the facts to the premise. The inference engine itself works backwards from the premise to the facts, and the explanation is subsequently generated from the trace of execution. Here is an example of the reasoning associated with the query whether daffodils are grown in England:

```
Premise: flowers of England includes daffodils
Evidence supporting the premise:

  flowers of England includes temperate-flowers
  and daffodils kind-of temperate-flowers
  therefore flowers of England includes daffodils

  flowers of Netherlands includes daffodils, tulips
  and Netherlands similar-to England for flowers
  therefore flowers of England includes daffodils

Suggesting: flowers of England includes daffodils is likely
```

This develops two lines of argument in favour of the premise in the query. The first is based on recognising that daffodils are a sub-class of temperate flowers, which are known to grow in England. The second makes use of knowledge that England and the Netherlands are similar in respect to the flowers grown. The inference engine uses a fixed strategy for searching for and applying relevant inferences. This may involve the use of graph algorithms such as spreading activation to propose and prioritise potential inferences as suggested by Collins. Other algorithms are used to compute certainties of inferences based upon statement metadata, and for assessing and combining multiple lines of argument. Future work will explore a wide range of reasoning, including spatial, temporal, causal and social reasoning, along with metacognition for problem solving, and support for System 1 and 2 cognition [1].

3. Analogical Reasoning

What benefits are potentially possible for analogical reasoning by cognitive agents? A starting point is to distinguish between literal and figurative analogies. The former involves things that are really quite similar, whilst the latter are not obviously comparable at first glance. Analogies can help agents to generalise their knowledge based upon a few examples. This has potential applicability for the properties of things, understanding their behaviours, as well as for problem solving by drawing upon previous experience in similar situations.

Analogies are further related to similes and metaphors in language. Similes involve a comparison that explicitly emphasises some comparable characteristic, e.g., "his words were like a punch in the guts" as a way to establish the impact of the words on the listener, whilst metaphors involve an implicit comparison, e.g., "to get cold feet" is to have second thoughts about some proposed course of action. People commonly use similes and metaphors to communicate thoughts in ways that are more vivid and interesting, as well as to structure perceptions and understanding, see Lackoff and Johnson [5]. As such, this is expected to be an important aspect of human-machine communication, albeit one that is very challenging, at least in the near future.

Dedre Gentner [6] notes that analogies may involve matching based upon structural or causal similarities, and noticing interesting differences, in essence, abstracting from similarities and dissimilarities

Gentner cites the example of plumbing in that electrical circuits can be likened to a plumbing system for water, e.g., equating voltage to pressure, and electrical current to water flow. Causal relationships for the source can be used to suggest similar relationships for the target, e.g., higher voltage leads to greater current just as higher water pressure leads to greater water flow.

Two situations can be identified as similar if they share some of the same properties, with the implication that you may be able to infer properties of the target from properties of the source. You may also be able to infer relationships, e.g., part/whole or cause/effect. More generally, the situations have different properties, that can however be mapped one to another (as in voltage to pressure). Such mappings have to be learned or guessed from matching relationships. Thus, if two situations/contexts have several properties or relationships in common, then we may consider them as analogical equivalents.

The notion of similarity introduced by Collins [2] supports inferences on shared property values at least in some given context, see the `similar-to` and `dissimilar-to` statements in PKN above. A generalisation is to relate pairs of different properties, e.g., voltage corresponds to pressure, and current to flow when making an analogy between electrical circuits and plumbing. Such pairings can be represented by adding a `corresponds-to` statement to PKN:

```
voltage corresponds-to pressure for circuit
current corresponds-to flow for circuit
flow increases-with pressure
# thus allowing us to infer
current increases-with voltage
```

We also need a way to describe that voltage and current are characteristics of electrical circuits, which are a sub-class of circuits, e.g.

```
electrical-circuit kind-of circuit
voltage property-of electrical-circuit
current property-of electrical-circuit
```

An open question is how people learn such knowledge from examples and being taught by others. That relates to the notion of syntagmatic and paradigmatic learning. Syntagmatic learning deals with learning co-occurrence patterns within episodes, whilst, paradigmatic learning involves identifying generalisations, and is believed to develop at a later age in childhood.

Analogies as part of critical thinking. It is easy to find web sites that propose the use of analogies for teaching purposes. These are based upon simple patterns, e.g., synonyms, antonyms, part/whole, cause/effect, etc. Here are some examples:

```
battery is-to torch as ?x is-to car # engine powers a car
itch is-to scratch as ?x is-to cold # virus causes a cold
wall is-to brick as bottle is-to ?x # a bottle is made of glass
```

Solving such queries involves identifying the pattern, and then applying background knowledge. The first step is to recognise the query as using an analogy. The next step is to use the pair that doesn't involve a variable to identify likely patterns, e.g., battery/torch in the first example. The knowledge base may contain plenty of facts and relationships, and it will be important

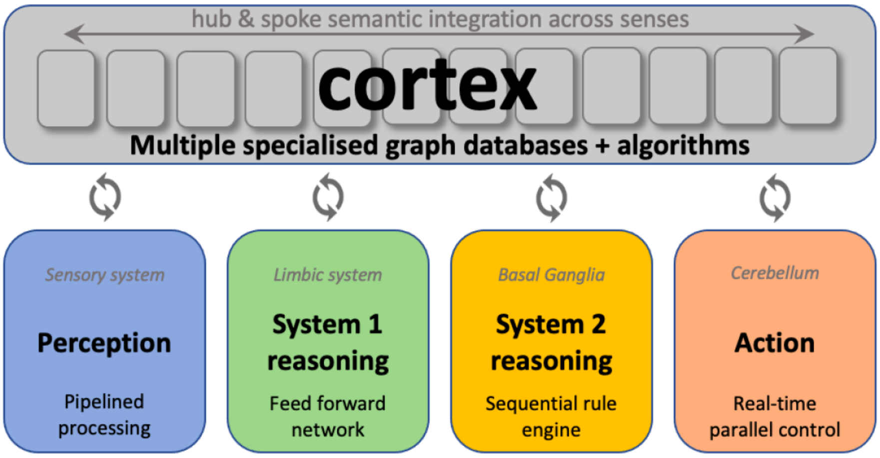

Figure 1: Cognitive architecture with multiple cognitive circuits loosely equivalent to shared blackboards – where semantic integration mimics the role of the Anterior Temporal Lobe as a hub for unimodal spokes.

to look for patterns that also occur for the pair with the variable. It may be the case that two pairs use different relationships, in which case, we need to find plausible evidence that they are comparable patterns.

Simple analogies are amenable to a fixed strategy plus associated graph algorithms. Qualitative metadata can be used to reason about certainty and to prioritise processing. What about more complicated analogies? The work by Jaime Carbonell on derivational analogies [7] is inspiring. The paper describes a problem solver that searches for analogies with previously solved problems, adapting the solution as needed based upon an analysis comparing the old and new problems.

4. Cognitive Architecture

The quest for realising human-like AI owes a huge debt to many pioneers over many decades. To mention just a few: Daniel Kahneman, a Nobel-prize willing psychologist who studied System 1 & 2 thinking along with cognitive biases [1]; Philip Johnson-Laird, a cognitive scientist renowned for his work on how humans reason in terms of mental models rather than logic and statistics [8]; John R. Anderson, a cognitive scientist renowned for his work on the ACT-R cognitive architecture for sequential cognition [9]; and Alan Collins, a cognitive scientist renowned for his work on plausible reasoning and intelligent tutoring systems [2].

Figure 1 illustrates a high-level cognitive architecture inspired by the structure and function of the human brain.

Memory is based on graph databases and associated graph algorithms. It combines symbolic graphs with sub-symbolic information, mimicking the human cortex, and defined at a conceptual level above that of RDF and Property Graphs (including NGSI-LD). Recall is stochastic reflecting prior knowledge and past experience. This involves activation boost/decay, spreading activation, the forgetting curve and spacing effect.

Perception interprets sensory data at progressively higher levels of abstraction, and places the resulting models into the cortex. Cognitive rules can set the context for perception, and direct attention as needed. Events are signalled by queuing chunks to cognitive buffers to trigger rules describing the appropriate behaviour. A prioritised first-in first-out queue is used to avoid missing closely spaced events.

System 1 is about intuitive/emotional thought, and prioritising what's important. The limbic system provides rapid automatic assessment of past, present and imagined situations without the delays incurred in deliberative thought. Emotions are perceived as positive or negative, and associated with passive or active responses, involving actual and perceived threats, goal-directed drives and soothing/nurturing behaviours.

System 2 is slower and more deliberate thought, involving sequential execution of rules to carry out particular tasks, including the means to invoke graph algorithms in the cortex, and to invoke operations involving other cognitive circuits. Thought can be expressed at many different levels of abstraction, and is subject to control through metacognition, emotional drives, internal and external threats.

Action is about carrying out actions initiated under conscious control, leaving the mind free to work on other things. An example is playing a musical instrument where muscle memory is needed to control your finger placements as thinking explicitly about each finger would be far too slow. The cerebellum provides real-time coordination of muscle activation actively guided by perception.

This architecture has been partially realised with a suite of web-based demos developed for the W3C Cognitive AI Community Group [4]. This includes the chunks and rules specification, and an implementation as a JavaScript library. Chunks are essentially collections of name/value pairs, where values are literals or references to other chunks, or lists thereof. Chunks are associated with decaying activation levels to mimic the characteristics of human-memory. Chunk rules support sequential reasoning (System 2).

Further work is underway to incrementally realise the requirements identified by Kahneman for System 1, and to understand how plausible reasoning, learning and metacognition can be layered on top of System 1 and 2. This will include the intuitive and deliberative reasoning involved in natural language processing, and the human ability to reason about the past, present, and imagined situations.

5. Conclusions

AI today can be broadly split into symbolic AI, statistical techniques, and approaches based upon deep learning and multi-layer artificial neural networks. Work in the cognitive sciences suggests a middle ground that combines symbols and sub-symbolic metadata, and is open to distributed representations (e.g., as vectors in noisy high dimensional spaces) where this would yield computational benefits. Traditional symbolic AI is hard to scale, relying on hand-coded knowledge, along with difficulties in dealing with imperfect knowledge, whilst deep learning scales well, but has challenges with reasoning and transparency. This paper draws attention to the potential for mimicking human-like memory, reasoning and learning, inspired by the wealth of research in the cognitive sciences.

References

[1] D. Kahneman, Thinking, fast and slow, Macmillan, 2011.
[2] A. M. Collins, R. S. Michalski, The logic of plausible reasoning: A core theory, Cogn. Sci. 13 (1989) 1–49. URL: https://doi.org/10.1207/s15516709cog1301_1. doi:10.1207/s15516709cog1301_1.
[3] D. Raggett, Plausible reasoning demo, visited on July 24, 2022. URL: https://www.w3.org/Data/demos/chunks/reasoning/.
[4] W3C Cognitive AI Community Group, visited on July 24, 2022. URL: https://github.com/w3c/cogai#readme.
[5] G. Lakoff, M. Johnson, Metaphors we live by, University of Chicago press, 1980.
[6] D. Gentner, Structure-mapping: A theoretical framework for analogy, Cogn. Sci. 7 (1983) 155–170. URL: https://doi.org/10.1207/s15516709cog0702_3. doi:10.1207/s15516709cog0702_3.
[7] J. G. Carbonell, M. Veloso, Integrating derivational analogy into a general problem solving architecture, in: Proceedings of the First Workshop on Case-Based Reasoning, Morgan Kaufmann Tampa, FL, 1988, pp. 104–124.
[8] P. Johnson-Laird, How we reason, Oxford University Press, 2006.
[9] J. R. Anderson, How can the human mind occur in the physical universe?, Oxford University Press, 2007.