

The application of qualitative metadata to analogical reasoning and *artificial minds*

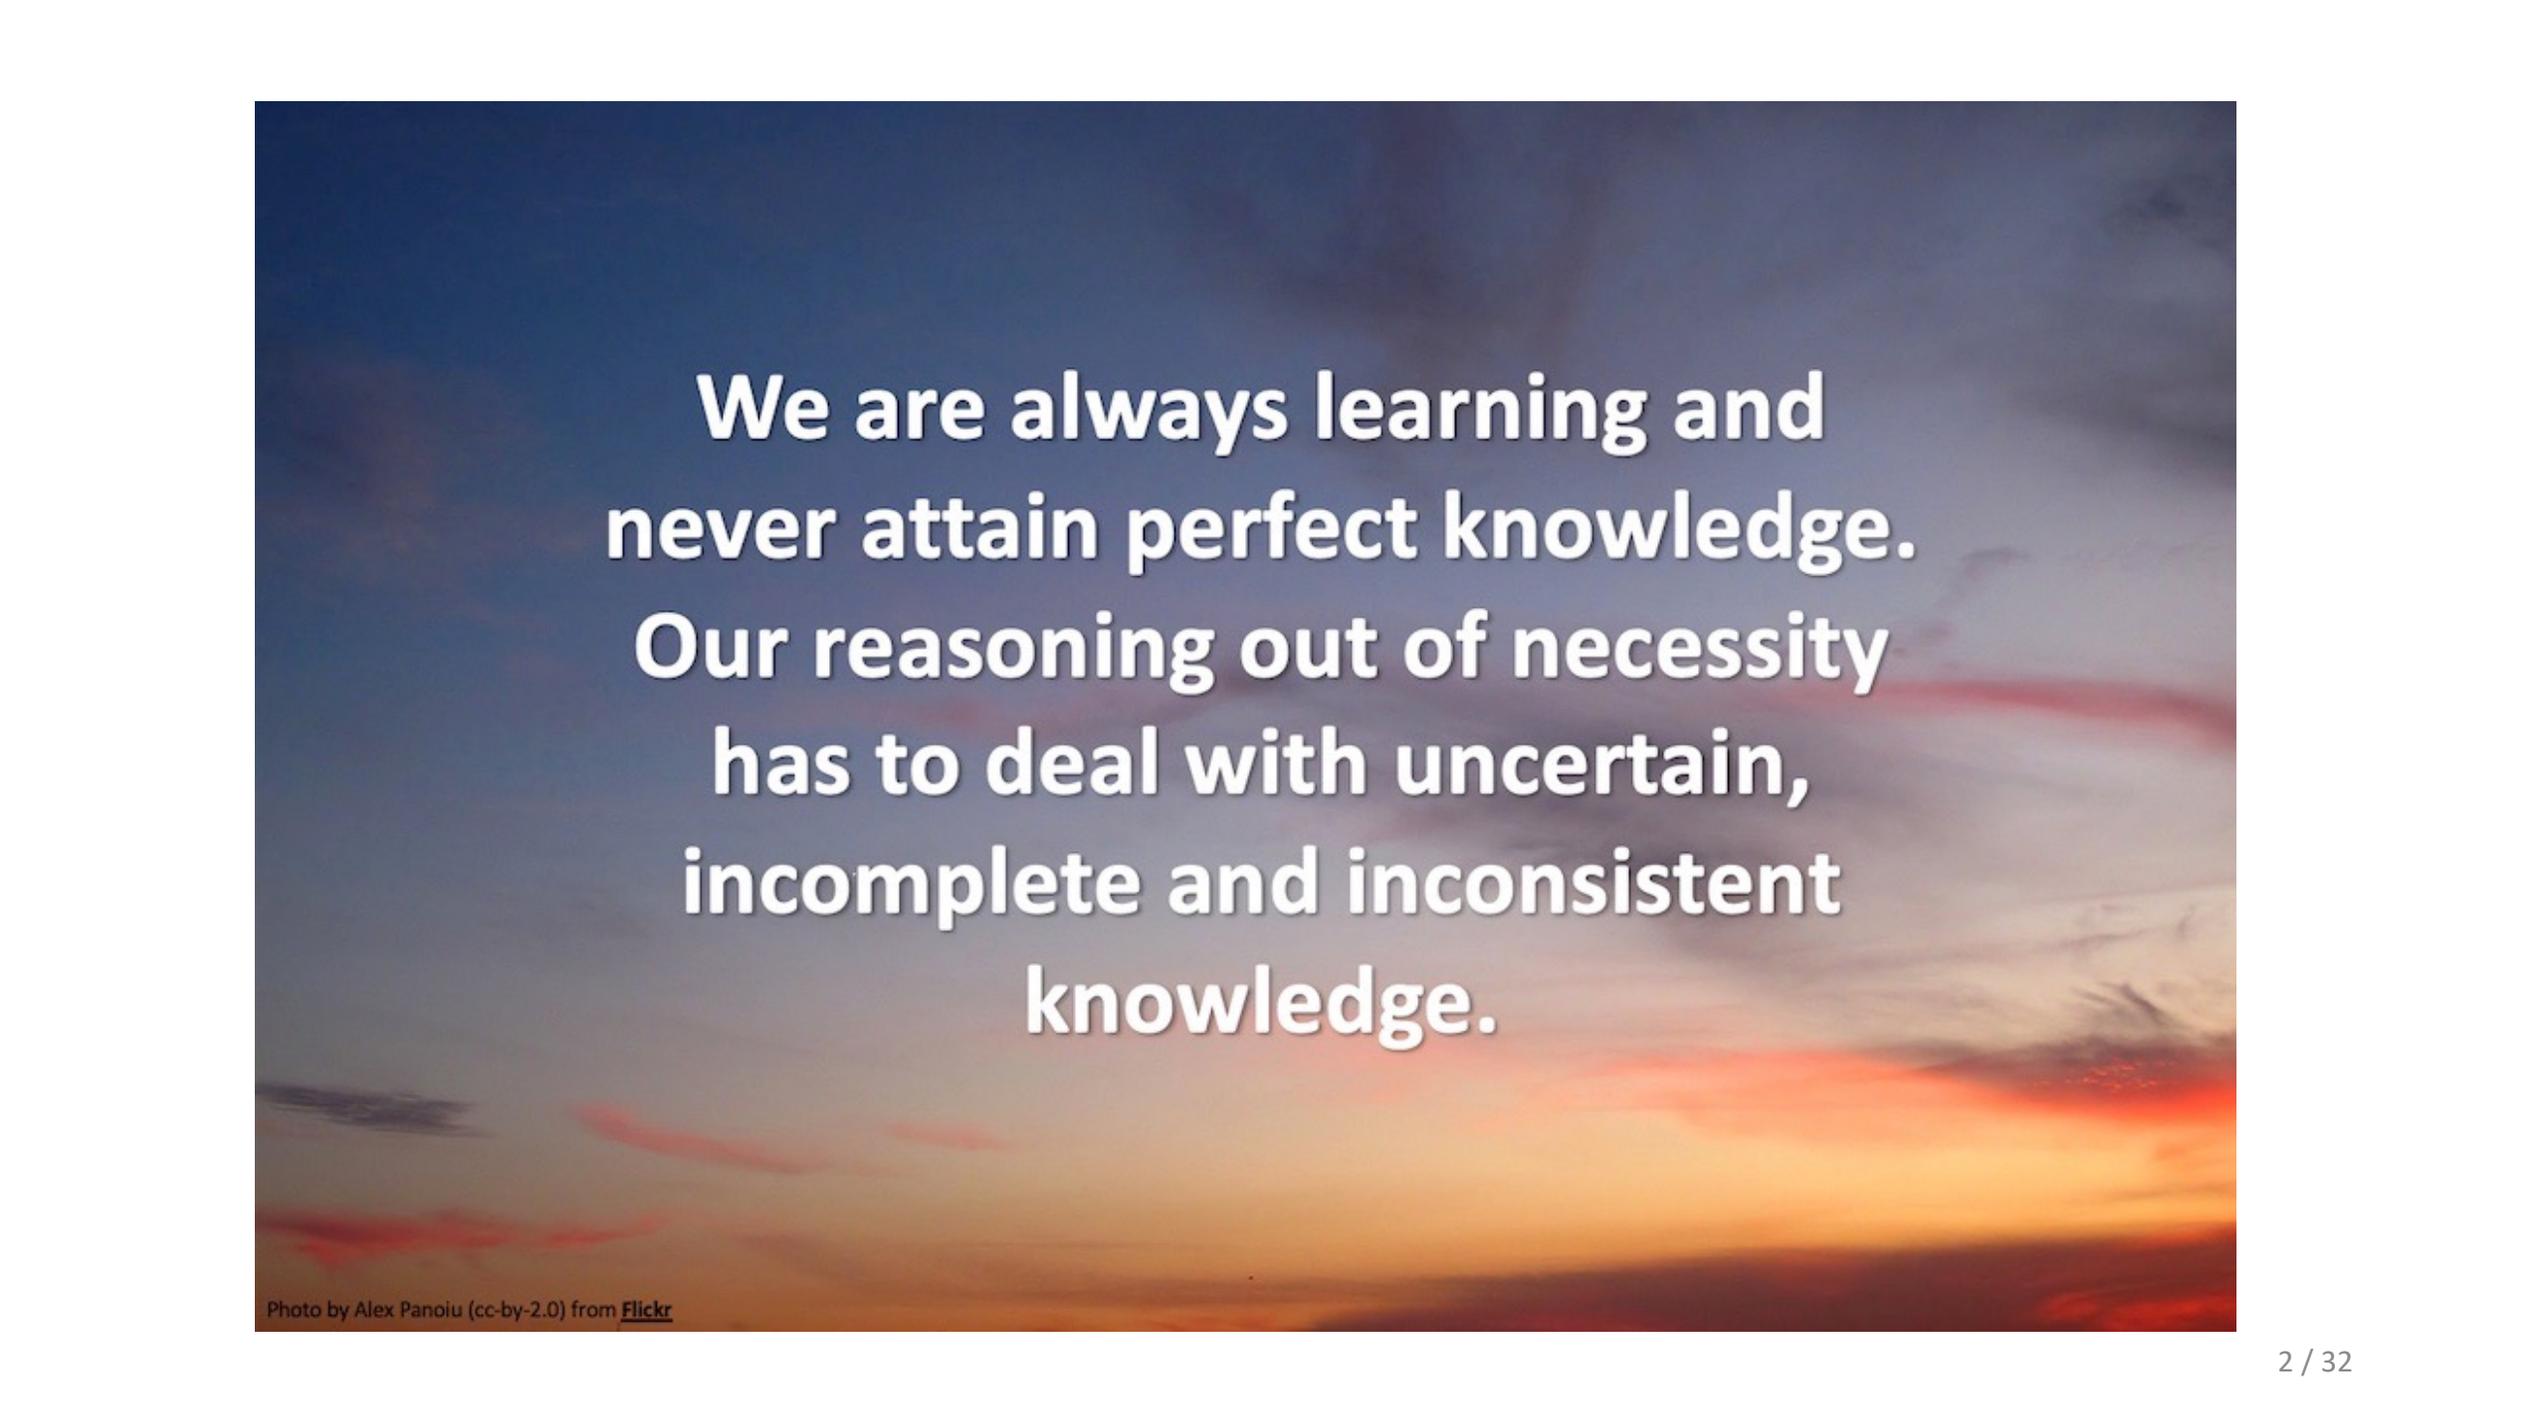
Workshop on Analogies: from Theory to Applications
12 September 2022, as part of ICCBR 2022



Dr. Dave Raggett
W3C/ERCIM



This work has been supported through funding for the [GATEKEEPER project](#) from the European Union's Horizon 2020 research and innovation programme under grant agreement No 857223.



**We are always learning and
never attain perfect knowledge.
Our reasoning out of necessity
has to deal with uncertain,
incomplete and inconsistent
knowledge.**

Human-like AI

falling down the rabbit hole into a new world!

General purpose Human-like AI will dramatically change how we work, how we communicate, and how we see and understand ourselves!

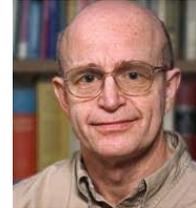


First, some acknowledgements

We owe these pioneers a huge debt for their hard won insights



Daniel Kahneman, Nobel prize winning psychologist who studied System 1 & 2 thinking, along with cognitive biases in “Thinking Fast and Slow”.



John R. Anderson, cognitive scientist renowned for his work on the ACT-R cognitive architecture for sequential cognition (System 2).



Philip Johnson-Laird, cognitive scientist renowned for his work on how humans reason in terms of mental models rather than using logic and statistics.



Allan M. Collins, cognitive scientist renowned for his work on plausible reasoning and intelligent tutoring systems.



Dedre Gentner, cognitive and developmental psychologist renowned for her work on analogical reasoning.



George Lakoff, cognitive linguist and philosopher renowned for his work on conceptual metaphors in language and cognition.



Allen Newell, researcher in computer science and cognitive psychology renowned for his work on AI and the Soar cognitive architecture.



Lotfi Zadeh, mathematician and computer scientist renowned for his work on fuzzy reasoning and control.

Why do this?

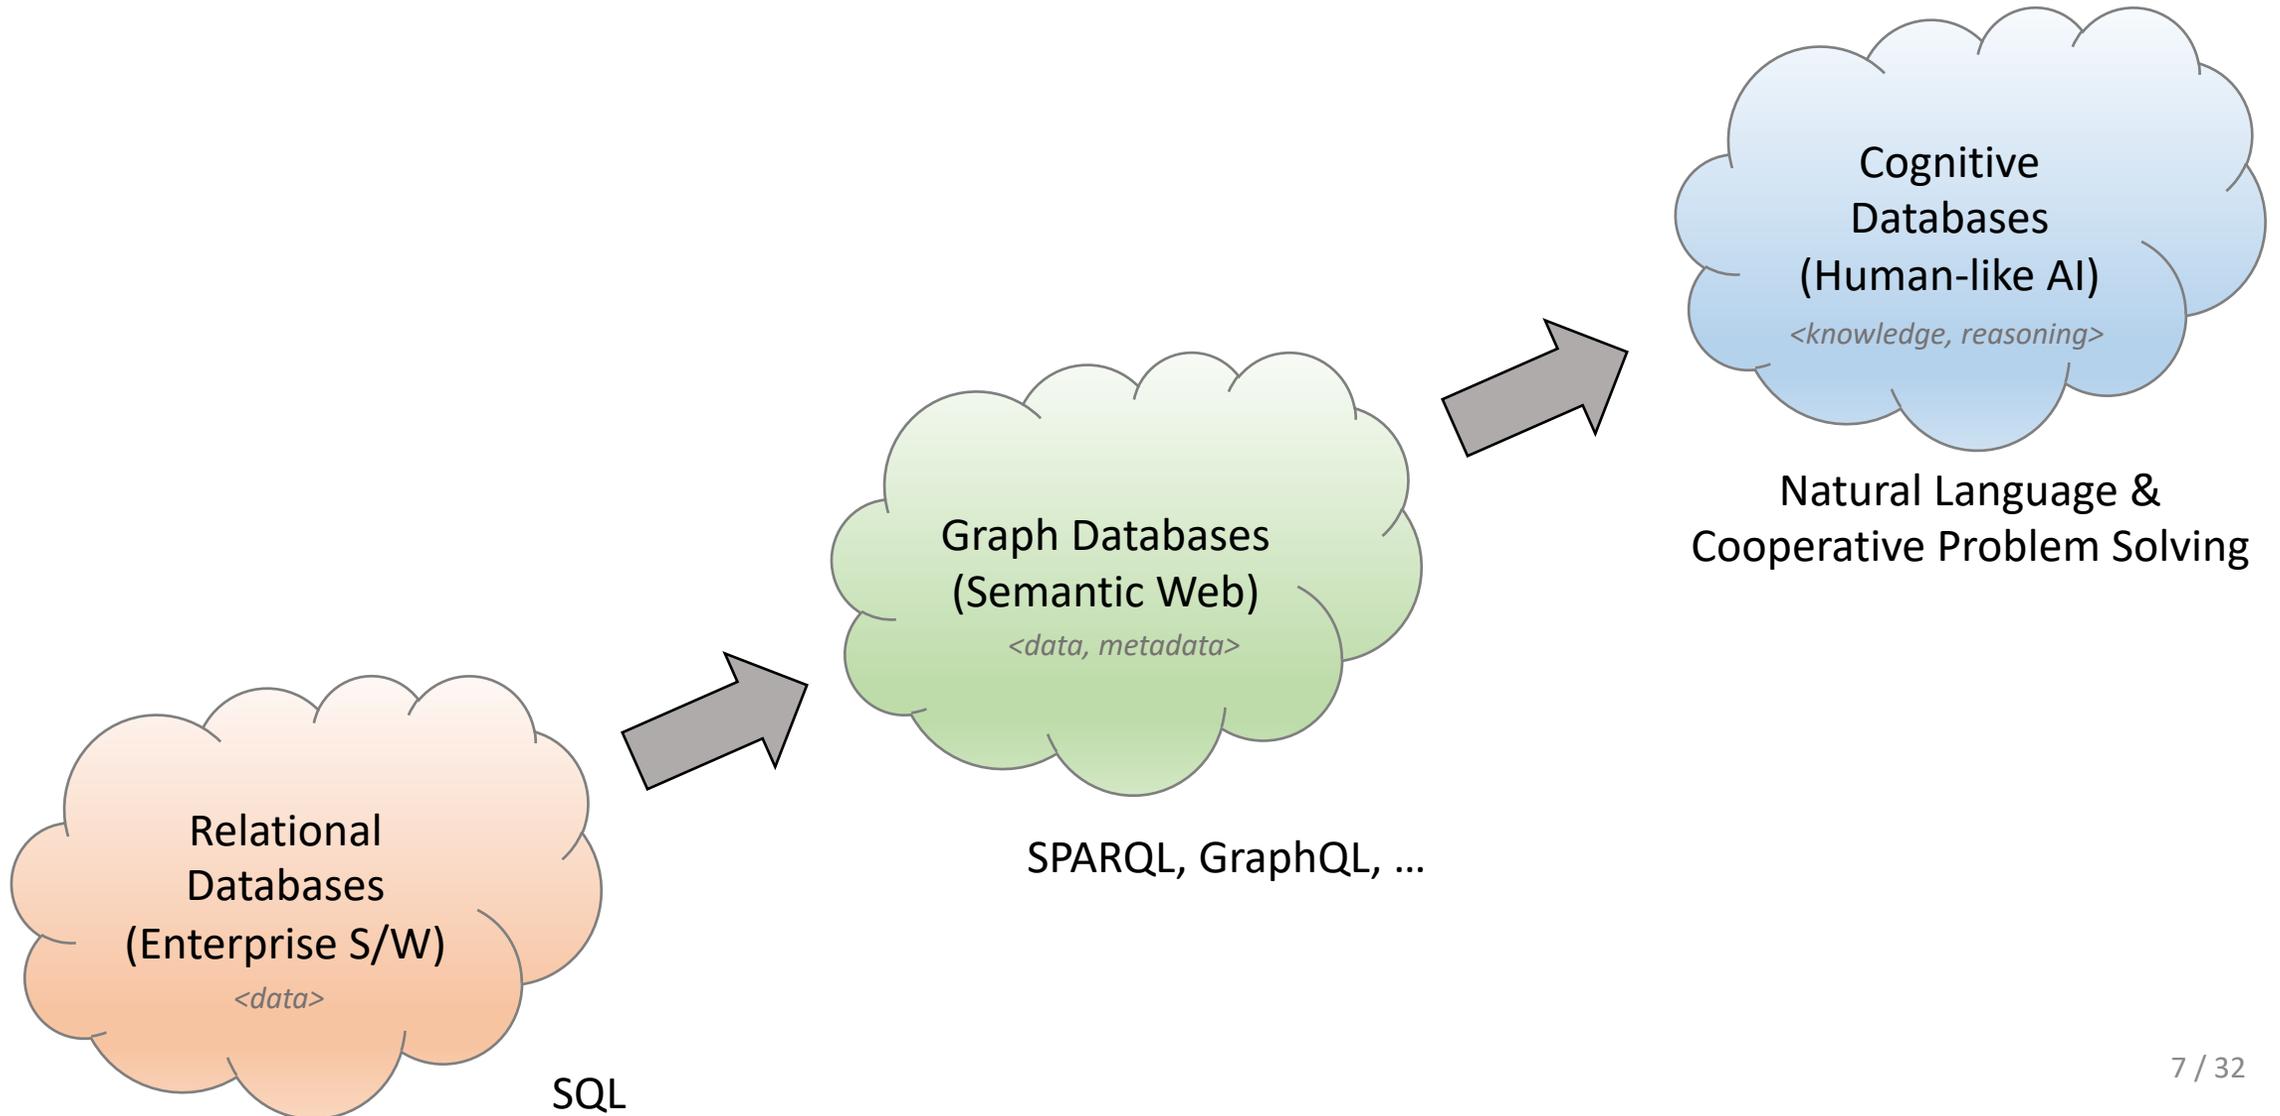
Why work on developing artificial minds?

- ❑ Computers that can help us to do more with less effort, boosting productivity and prosperity
- ❑ Compensating for shortages of skilled workers
- ❑ Especially in countries with ageing or declining populations, as people have fewer children, later in life
- ❑ Cognitive systems that think like us, understand us and empathise with us
- ❑ Help us to reason better by shining a light on our many cognitive biases
- ❑ Work in our place in hostile environments, e.g. space exploration
- ❑ But in turn, we need to adapt our societies to avoid the unwanted effects of uncontrolled use of automation that serves the needs of the few instead of the many
- ❑ Freeing up and rewarding people to work on socially important jobs

Reasoning with Knowledge Graphs*

- ❑ Whilst today's knowledge graphs claim to capture knowledge there is very little attention currently to automated reasoning
 - One exception is inheritance down class hierarchies
- ❑ Application logic is instead embedded in application code
 - This makes it hard to understand and costly to update
 - Why do we put up with this?
- ❑ Knowledge presumes reasoning and is otherwise just information!
 - Information is structured labelled data, such as column names for tabular data
 - Knowledge is understanding how to *reason* with information
- ❑ It is high time to focus on automated reasoning for human-machine cooperative work
 - Boosting productivity and compensating for skills shortages

Evolution in action



Plausible Reasoning and Argumentation

- ❑ People have studied the principles for plausible arguments since the days of ancient Greece, e.g. Carneades and his guidelines for argumentation
- ❑ There has been a long line of philosophers working on this since then, including Locke, Bentham, Wigmore, Keynes, Wittgenstein, Pollock and many others
- ❑ Plausible reasoning is *everyday reasoning*, and the basis for legal, ethical and business discussions
- ❑ Researchers in the 20th century were sidetracked by the seductive purity of mathematical logic, and more recently, the amazing magic of deep learning
- ❑ It is now time to exploit plausible reasoning with imperfect knowledge for human-machine cooperative work using distributed knowledge graphs
- ❑ Enabling computers to analyse, explain, justify, expand upon, and argue in human-like ways

*This is a **major** step forward for AI*

Everyday Knowledge is Imperfect

- ❑ In the real world, knowledge is distributed and imperfect
- ❑ We are learning all the time, and revising our beliefs and understanding as we interact with others
- ❑ Imperfect in the sense of **uncertain, incomplete** and **inconsistent**
- ❑ Conventional logic fails to cope with this challenge
- ❑ The same is true for statistical approaches, e.g. Bayesian inference, due to difficulties in compiling the required statistics
- ❑ Evolution has equipped humans with the means to deal with this
- ❑ However, not everyone is rational, and some lack sound judgement*

** Moreover, all of us are subject to various kinds of cognitive biases*

Plausible Inferences

- Consider $A \Rightarrow B$, 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
- A more concrete example: *if it is raining then it is cloudy*
 - This can be used in both directions: Rain is more likely if it is cloudy, likewise, if it is not raining, then it might be sunny, so it is less likely that it is cloudy*
 - i.e. using our knowledge of weather
- In essence, plausible reasoning draws upon prior knowledge as well as on the role of analogies and consideration of examples, including *precedents*
- Mathematical proof is replaced by reasonable arguments, both for and against a premise, along with how these are assessed
 - In court cases: arguments are laid out by the Prosecution and Defence, the Judge decides what evidence is admissible, and guilt is assessed by the Jury

* Note use of qualitative terms in lieu of quantitative statistics

Plausible Reasoning

During the 80's Alan Collins and co-workers developed a theory of plausible reasoning* based upon 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 proceedings
- ❑ 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 opposite 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

* See: Collins & Michalski (1988) and subsequent extensions by Burstein, Collins and Baker (1991)

Plausible Knowledge Notation (PKN)

□ Properties

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

- where *certainty* is an example of qualitative metadata

□ Relationships

robin kind-of songbird
duck similar-to goose for habitat
duck dissimilar-to goose for neck-length
dingy is small for sailing-boat*

- where *for* lists properties that the relationship applies to

□ Ranges

range of guilt includes innocent, guilty
(domain closed, overlap none)

- used to describe referent domains

□ Dependencies

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

- describes a coupling between a pair of properties

□ Implications

climate of ?place includes hot and
rainfall of ?place includes heavy
implies crops of ?place includes rice

- a form of *if-then* rules

□ Qualitative metadata can be given with all PKN statements

- PKN queries will be introduced later in this talk

Relationships, implications and dependencies can be used for inferences in both directions

** The "is" relation is used for properties that lack descriptors.*

Qualitative Metadata

Used to estimate *certainty* for each plausible inference with algorithms to combine multiple sources of evidence

- ❑ *typicality* in respect to other group members
 - e.g. robins are typical song birds
- ❑ *similarity* to peers
 - e.g. having a similar climate
- ❑ *strength, inverse* – conditional likelihood in each direction
 - e.g. strength of climate for determining which kinds of plants grow well
- ❑ *frequency* – proportion of children with given property
 - e.g. most species of birds can fly
- ❑ *dominance* – relative importance in a given group
 - e.g. size of a country's economy
- ❑ *multiplicity* – number of items in a given range
 - e.g. how many different kinds of flowers grow in England

Plausible Inferences

- ❑ Let's start with something we want to find evidence for and against
 - flowers of England includes daffodils
and its inverse
 - flowers of England excludes daffodils
- ❑ We first check if this is a known fact and if not look for other ways to gather evidence
- ❑ We can either *generalise the **property value***
 - flowers of England includes ?flower
- ❑ We find a matching property statement
 - flowers of England includes temperate-flowers
- ❑ We then look for ways to relate daffodils to temperate flowers
 - daffodils kind-of temperate-flowers
- ❑ So we infer that daffodils grow in England
- ❑ Or we can *generalise the **property argument***
 - flowers of ?place includes daffodils
- ❑ We look for ways to relate England to a similar country
 - Netherlands similar-to England for flowers
- ❑ We then find a related property statement
 - flowers of Netherlands includes daffodils, tulips
- ❑ This also allows us to infer that daffodils grow in England
 - The certainty depends on the parameters, in this case *similarity*
- ❑ These examples use properties and relationships, but we can also look for implications and dependencies
 - e.g. a medium latitude implies a temperate climate, which in turn implies temperate flowers

We can prioritise inferences that seem more certain, and ignore those that are too weak

Plausible Reasoning

- ❑ Add the premise as a goal
 - It may include variables
- ❑ Iteratively select a goal and look for potential inferences
- ❑ Direct evidence using explicit property statements (facts)
- ❑ Indirect evidence via inferences
 - kind-of and part-of relations
 - similarities and dissimilarities
 - dependencies
 - implications
- ❑ Inferences may result in new goals for multi-step arguments
- ❑ Skip previously considered goals to avoid indefinite looping
 - mimicking human recall
- ❑ Record reasoning as backward chain of satisfied goals from the premise to the facts
- ❑ Compute certainty working from the facts to the premise
- ❑ Finally generate explanation

When to curtail search: e.g. after finding direct evidence, or on diminishing return on effort?

(analogous to when people get bored and lose interest)

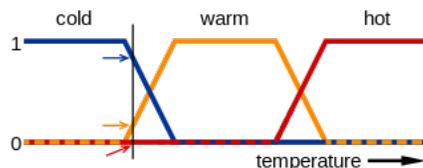
Computing the Degree of Certainty

- ❑ Conditional likelihood
 - Rain implies cloudy weather with a high degree of certainty, but there is only a low certainty of rain given cloudy weather
- ❑ Psychological experiments* show that people aren't good at compiling and using statistics – hence justifying use of qualitative rather than quantitative metadata
 - e.g. low, medium, high
 - Can be temporarily mapped to range 0 – 1 for internal calculations
- ❑ Start from the certainty of the facts and work towards the premise - i.e. the query posed to the reasoner
- ❑ For each inference we estimate its certainty given its metadata
- ❑ Given multiple matching facts, the more of them, the greater the certainty
- ❑ Gather property values from the matching facts using set union
 - e.g. the set of different kinds of flowers
- ❑ What to do when different values have different certainties?
 - Either track at the level of individual values
 - Or take an average and drop those that are below a threshold set relative to the average
- ❑ For implications (rules) with multiple antecedents, use the minimum
 - Corresponding to *conjunction* in fuzzy logic
- ❑ Metadata parameters strengthen, weaken or have no effect on certainty of an inference
 - Dependent on the type of inference

* e.g. the work by Daniel Kahneman

Relationship to Fuzzy Logic

- Lotfi Zadeh* describes fuzzy logic as a means to compute with words, e.g.
 - *small* can be multiplied by *a few* and added to *large*, or *colder* can be added to *warmer* to get something in between
- Fuzzy logic is often used for device automation
 - Scalar ranges are described as a blend of overlapping values, e.g. cold partially overlaps warm, and warm partially overlaps hot for temperatures
 - A given temperature can then be described with the fuzzy set, e.g. {cold 0.8, warm 0.2, hot 0.0}



- Control logic can then be described in terms of simple rules expressed using terms from the ranges
 - if cold then fan is stopped
 - if warm then fan is slow
 - if hot then fan is fast
- The fuzzy set for fan speed is then computed from the fuzzy set for the current temperature
 - {stopped 0.8, slow 0.2, fast 0.0}
- The selected speed is then derived from the fuzzy set
- Fuzzy sets are related to multiple lines of argument, e.g. that the fan speed is stopped, slow or fast

* The father of fuzzy logic

Reasoning By Analogy

- ❑ Plausible inferences based upon direct similarities or structural similarities
- ❑ If you suspect that something is analogous to something else, you can look for evidence of correspondences, as in:
 - circuit **analogous-to** plumbing
 - flow **increases-with** pressure **for** plumbing
 - current **increases-with** voltage **for** circuitallowing us to infer that voltage **corresponds-to** pressure and current **corresponds-to** flow
- ❑ You can solve problems like **leaf:tree::petal:?** by looking for structural similarities, e.g.
 - leaf **part-of** tree
 - petal **part-of** flower
- ❑ Similarly: **short:light::heavy:?**
 - short **less-than** long **for** size
 - light **less-than** heavy **for** weight
- ❑ And: **mansion:shack::yacht:?**
 - small **less-than** large **for** size
 - mansion **is** large **for** building
 - shack **is** small **for** building
 - yacht **is** large **for** sailing-boat
 - dingy **is** small **for** sailing-boat

More generally we can deal with softer matches and ranking which are more plausible than others

Proof of Concept

- ❑ Implementations are invaluable for testing understanding of previous work and for identifying challenges for new work
- ❑ A web-based demo that provides a novel simple notation and an inference engine inspired by the work of Allan Collins et al.
 - See: <https://www.w3.org/Data/demos/chunks/reasoning/>
- ❑ Collins distinguishes four kinds of plausible assertions
 - properties, relationships, implications and dependencies
- ❑ Inference involves qualitative parameters
 - certainty, typicality, similarity, frequency, dominance, conditional likelihood
- ❑ A collection of static strategies ...

“Machines that can use facts to present a convincing case could transform the way we make decisions – and help us understand our own rhetoric”

New Scientist, 7 September 2016 [[Link](#)]

Demo

Contact: Dave Raggett dsr@w3.org

PKN in relation to RDF and LPG

Existing graph database frameworks

Resource Description Framework (RDF)

- ❑ RDF uses triples for labelled directed graph edges <subject, predicate, object>
- ❑ PKN relationships and properties can easily be mapped to triples
 - If you ignore value lists and parameter lists, but RDF-star would help ...
- ❑ PKN statements correspond to sub-graphs in RDF
 - e.g. implication as a set of triples
- ❑ Considerations in respect to global names and ontologies
 - Role of lexicons for natural language

Labelled Property Graphs (LPG)*

- ❑ LPG allows name/value pairs on graph edges and vertices
- ❑ LPG nodes correspond to collections of PKN properties
 - If you ignore parameter lists
- ❑ LPG links correspond to PKN relationships
 - Link properties expressed as PKN properties that cite the relationship
- ❑ LPG lacks implications and dependencies
 - These could be modelled as sub-graphs

* e.g. Neo4J, Amazon Neptune, ...

Plausible reasoning is possible with RDF and LPG, but considerably easier with PKN

A Graph of Overlapping Graphs

and how to scale to much larger graphs

- ❑ Very large knowledge graphs are difficult to deal with
- ❑ Displays of graphs lack context when zoomed in, and are intimidatingly complex when zoomed out
- ❑ This is also a challenge for automated reasoning due to the lack of context
- ❑ Proposed solution is to map large graphs into overlapping smaller graphs that model sets of *named contexts*
- ❑ For digital transformation of an enterprise, use subgraphs for different business functions
 - Different views for different departments
 - Dependency tracking for old & new uses
 - Well defined business process for managing vocabularies at different levels of maturity
- ❑ Also very pertinent for natural language
 - “*John opened the bottle and poured the wine*”
 - i.e. a social situation with wine being transferred from bottle to guest’s glasses
 - How do humans find coherent explanations for natural language utterances so quickly?
- ❑ Exploit *spreading activation* to identify shared contexts and most plausible word senses*
 - Accounts for semantic priming
 - Activation wave weakened by high fan-out
 - Stochastic recall based upon adding noise
 - Forgetting curve and spacing effect
 - Semantic similarity for related terms
 - Easily implemented in software!
- ❑ Specify contexts as named sets of PKN statements, analogous to sets of triples
 - Statement metadata names contexts

* Collins et al. suggest spreading activation could play a role for guiding search for potential inferences

Metaphors and Similes

- ❑ Lackoff & Johnson have shown just how much humans rely on metaphor in everyday thought and language*
- ❑ **Simile** is a figure of speech using an explicit comparison to emphasise something, e.g.
 - *“What light through yonder window breaks? It is the east, and Juliet is like the sun!”*
- ❑ **Metaphor** is like simile, but leaves the comparison implicit, e.g.
 - *“to get cold feet”* about something
- ❑ Metaphors and similes are kinds of language based analogies
- ❑ We learn them from other people, and use them as an expected way to communicate, assuming and reinforcing a shared experience
 - Our peer group, the books we’ve read and the movies we’ve seen
- ❑ How do we recognise a metaphor we haven’t come across before?
- ❑ How do metaphors impact on plausible reasoning?

Richer Queries

- ❑ Quantifiers in queries and rules
 - Natural language is very flexible, e.g. *none, few, some, many, most* and *all*
 - Definite and indefinite references
 - Relative quantifiers*
 - Scoping to real or imagined contexts
- ❑ Example: are yellow roses found in England?
 - flower of England includes ?rose and ?rose kind-of rose and colour of ?rose includes yellow
- ❑ What about
 - Are *all* English roses red or white?
 - Are only a *few* roses yellow?
 - Are *most* people older than 20?
 - *Which* roses are red?
 - *How many* people are younger than 15?
- ❑ Comparisons, e.g. *smaller than*
- ❑ Transformations with modifiers
 - e.g. *very, slightly, smallest, largest, average*

Are *all* English roses red or white?
all ?x where colour of ?x includes red, white from ?x kind-of rose and flowers of England includes ?x

Are only a *few* roses yellow?
few ?x where colour of ?x includes yellow from ?x kind-of rose

Which English roses are yellow?
which ?x where colour of ?x includes yellow from ?x kind-of rose and flowers of England includes ?x

Are *most* people older than 20?
most ?x where age of ?x greater-than 20 from ?x isa person

Is *anyone* here younger than 15?
any ?x where age of ?x less-than 15 from ?x isa person

How many people are slightly younger than 15?
count ?x where ?x isa person and age of ?x slightly:younger-than 15

How many people are very old?
count ?x where ?x isa person and age of ?x includes very:old

* Note: *few* signifies that a given subset contains only small number of items relative to the set the selected items were taken from. Queries involving relative quantifiers like *few* need to specify the selection criteria and the set, the criteria applies to. Other queries such as *which* and *count* just specify a set of items. Modifiers refer to ranges, e.g. height, and may be used in combination, e.g. *very slightly smaller*, and involve a source which might be an individual or a set of individuals. In some cases, the source set may be implicit from the context.

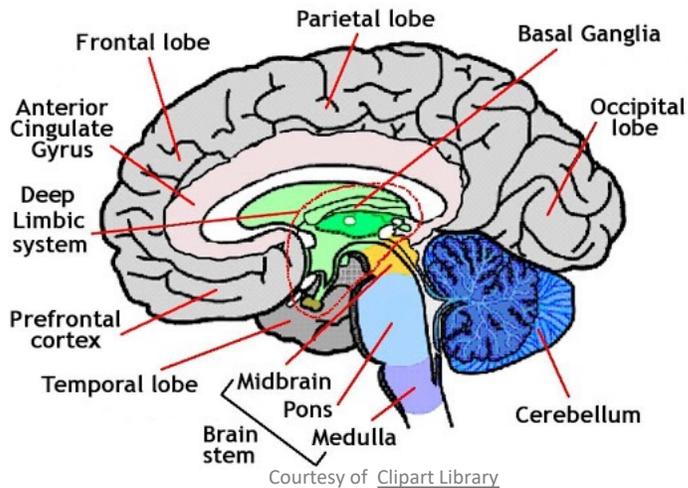
Enabling Natural Language

- ❑ Work on representing and reasoning with a broad range of natural language semantics
 - As a prerequisite for implementing natural language interaction
- ❑ How to respond in a human way
 - Taking the implicit context into account
 - *yes/no* answers to queries are often inappropriate
 - Long lists are likewise awkward, and it is better to just give a few pertinent examples – but how to select them from the longer list?
 - Determining when to use common metaphors and stock phrases
 - When to generate supporting graphics to communicate complex ideas
- ❑ The computer should be prepared to respond by asking questions that clarify what's wanted
- ❑ Role of natural language dialogue
 - including giving instructions, asking for advice, explanations, etc.
- ❑ Grice's maxims of conversation
 - quantity, quality, relation and manner

Cognition



- ❑ The plausible reasoning demo uses hard coded strategies for reasoning
- ❑ That's fine as a starting point, but needs to be replaced by an architecture that supports higher level cognition
 - Thinking about thinking
 - Which strategies make sense and when it is time to switch strategies?
 - Continuous learning using prior knowledge to get the most from limited data
 - Causality and intent
 - Higher level goals and drives
 - Theory of mind and of self
 - How are we doing in respect to our values and objectives?
- ❑ To realise artificial minds we would be wise to start with a cognitive architecture inspired by the human mind and brain



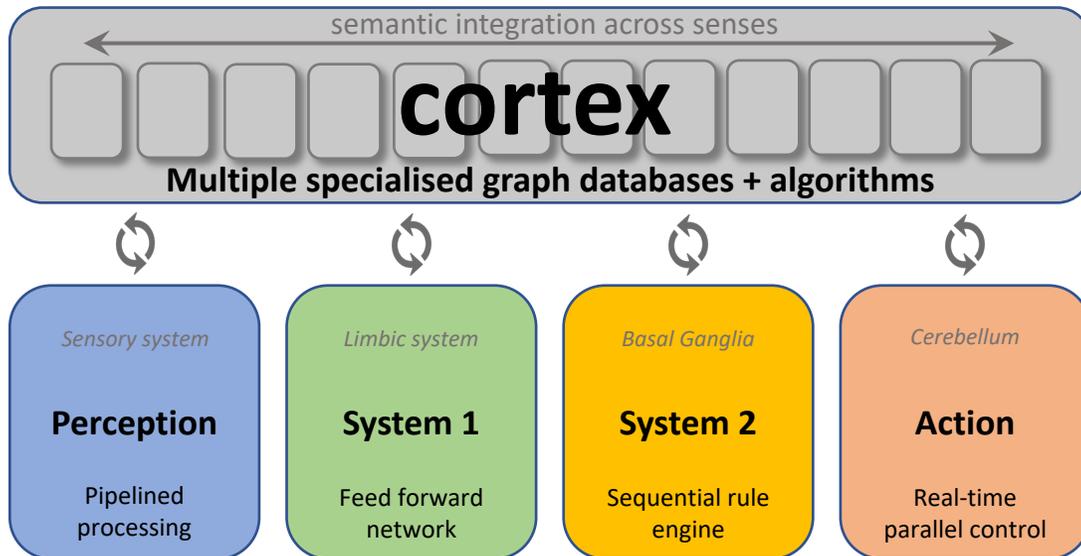
Courtesy of [Clipart Library](#)

Anterior temporal lobe as hub for integration across senses

Cognitive Architecture for artificial minds



Multiple cognitive circuits loosely equivalent to shared blackboard



- **Perception** interprets sensory data 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** covers intuitive/emotional thought, cognitive control and prioritising what's important. The limbic system provides rapid assessment of past, present and imagined situations. 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 systems. 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 guided by perception.

System 1 and 2

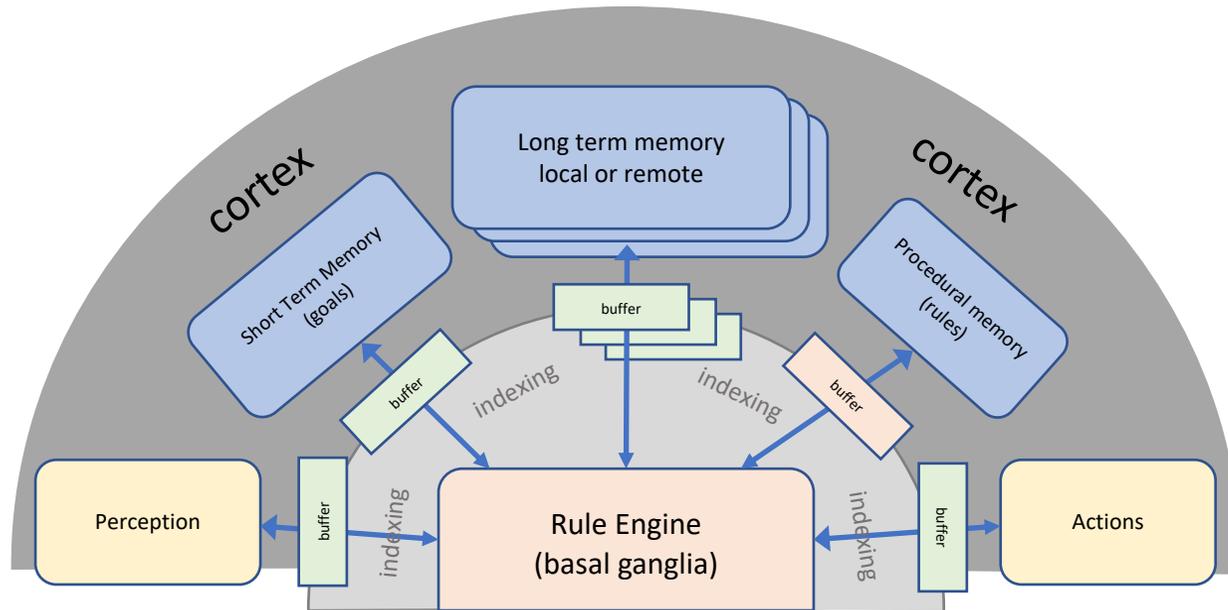
Popularised by Daniel Kahneman's "Thinking fast and slow"

- ❑ System 1 is fast, and apparently effortless, yet opaque
 - we aren't aware how we came to a conclusion
- ❑ System 1 is subject to many cognitive biases and often wrong
- ❑ Natural language is largely handled via System 1
 - We understand what people are saying, and construct a coherent explanation in real-time that hides the ambiguity of language
 - How do we do this?
- ❑ System 1 & 2 work in cooperation
- ❑ System 2 is accessible to introspection, and is much slower
- ❑ System 2 is effortful – making thinking hard work and quite exhausting!
- ❑ John Anderson's work on ACT-R
 - Sequential rule engine working on cognitive buffers that hold single chunks (a collection of name/value pairs)
 - My reworking of ACT-R in JavaScript
 - But unfortunately, still quite low level
- ❑ What about high level cognition?
 - Higher level rule language?

Cortico-Basal Ganglia Circuit

a functional model of System 2

Cognition – Sequential Rule Engine



Cognitive Buffers hold single chunk
Analogy with HTTP client-server model

- Inspired by John R. Anderson's ACT-R
 - Chunk graphs + Condition-Action rules
- Chunks as a collection of properties for literals and references to other chunks
 - Each chunk buffer is equivalent to the concurrent activity of a bundle of nerve fibres connecting to a given cortical region, see Chris Eliasmith's work on semantic pointers for pulsed neural networks
 - Rules operate over chunk buffers and invoke asynchronous operations on cortical modules that update the buffers
 - Stochastic selection from matching rules whenever buffer contents are updated

Chunks & Rules

- **Developed by W3C Cognitive AI Community Group**
 - [CogAI CG](#) is open to all, free of charge
 - [GitHub repository and documentation](#)
- **Chunks as collection of properties for literals and references to other chunks**
 - Chunks map to N -ary relations in RDF
 - Easier to work with than RDF
 - Formal spec as [draft CG Report](#)
- **Combination of symbolic + sub-symbolic approaches**
 - Graphs + statistics + rules + algorithms
 - Stochastic recall analogous to Web search, preferentially finding things that have proven useful in past experience
 - Mimics human memory – forgetting curve, spacing effect – using spreading activation and activation decay
 - Explainable AI/ML, learning with smaller datasets using prior knowledge and past experience
- **Growing Suite of web-based demos**
 - counting, decision trees, industrial robots, smart homes, natural language, self-driving cars, browser sandbox, chunks test suite, open source JavaScript chunks library

- Here are two examples of the same chunk – you can use newline or semicolon as punctuation:

```
dog dog1 {  
  name "fido"  
  age 4  
}  
  
dog dog1 {name "fido"; age 4}
```

- Here are some examples of chunk rules

```
# retrieve turn  
alert {@module goal; kind turn; turn ?id }  
=>  
turn {@module goal; @do recall; @id ?id}
```

```
# prepare for turn  
turn {@module goal; @id ?id; signal ?direction }  
=>  
action {@module car; @do brake; turn ?id},  
action {@module car; @do signal; signal ?direction},  
alert {@module goal; @do clear}
```

```
# start turn  
alert {@module goal; kind start-turn }  
=>  
action {@module car; @do steer; mode turn},  
action {@module car; @do cruise; speed 20},  
alert {@module goal; @do clear}
```

Longer Term Roadmap

- Plan to develop a suite of related demos to cover plausible reasoning in respect to induction, abduction, planning, belief revision, causal, social, emotional and other kinds of reasoning, including embracing fuzzy reasoning and qualitative reasoning
 - Not forgetting spatio-temporal reasoning and tenses
- Implementing two ways to support distributed knowledge:
 - Cognitive agents with shared access to remote cognitive databases, analogous to different lobes in the cerebral cortex – **hive minds** with a shared memory and accelerated communal learning
 - Web of collaborating agents and humans for ecosystems of services – relation to *data spaces*
- Potential for single framework encompassing plausible reasoning & Bayesian Inference
 - Exploiting mix of qualitative and quantitative metadata for better inferences
- Research on indexing, metacognition, and System 1 & 2 style cognition
 - Including techniques for mitigating different kinds of cognitive biases
- Work on integration with natural language interaction and natural language semantics
 - Along with continuous learning
 - Syntagmatic learning
 - Taxonomic learning
 - Skill compilation
- Work on developing common sense skills through teaching and learning through interaction in real or virtual environments – once trained, agents are easily cloned
- Commercialisation in respect to Digital Transformation and enabling non-programmers to work with information using collaborative multi-modal cognitive agents

Thank you for listening

Let's talk!

Contact: Dave Raggett dsr@w3.org