IoT Week 2022

# Combining Digital Twins with Cognitive Agents and Plausible Reasoning

Dave Raggett <dsr@w3.org> W3C/ERCIM

13/10/2022

This work has been supported through funding for the TERMINET project from the European Union's Horizon 2020 research and innovation programme under grant agreement No 957406.

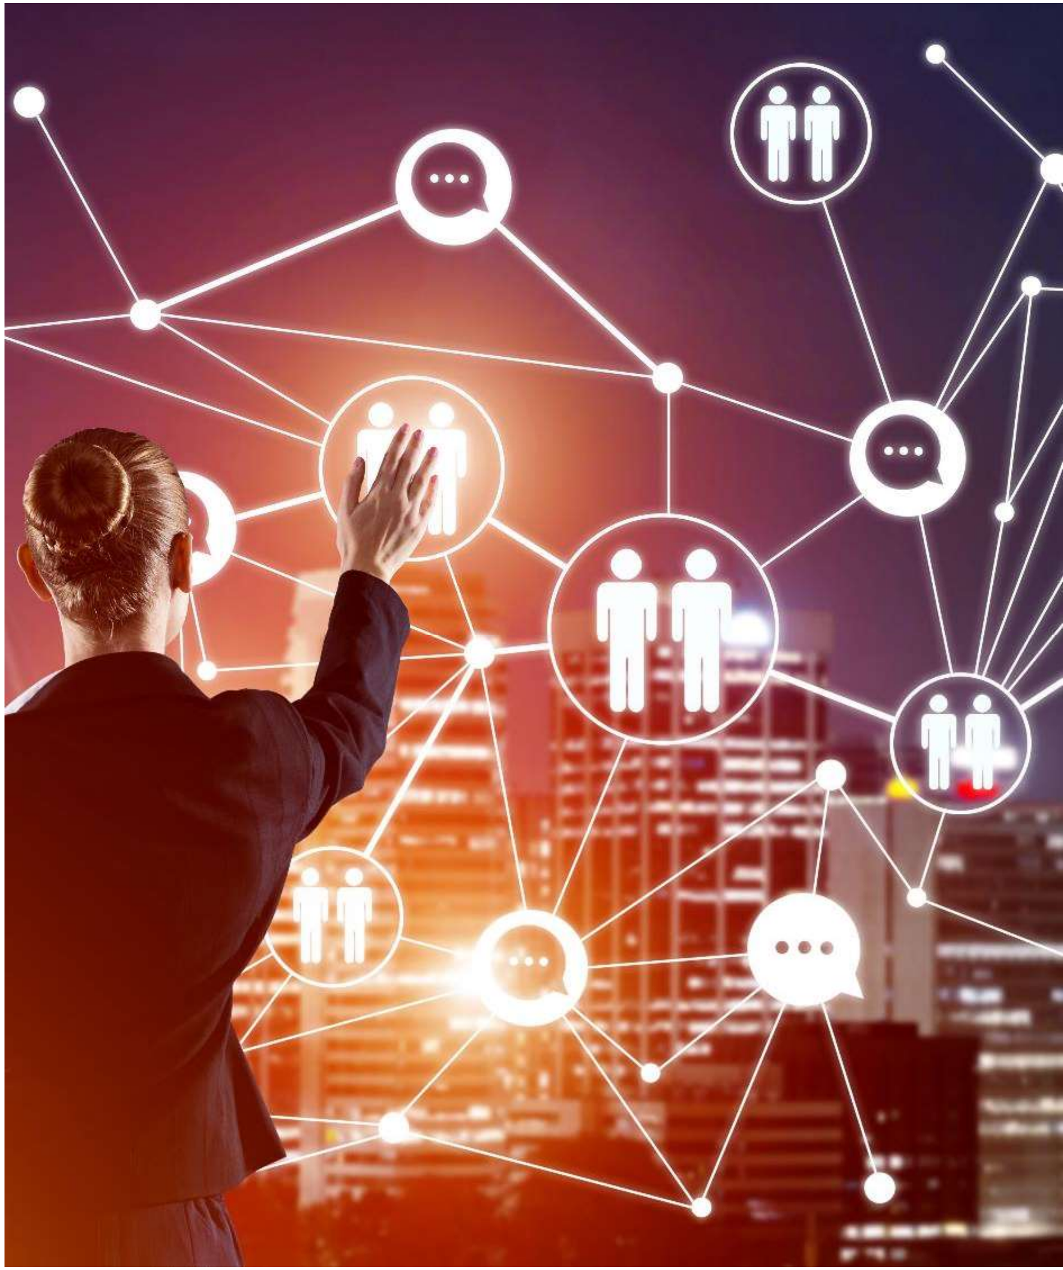

# TERMINET: next generation Internet of Things

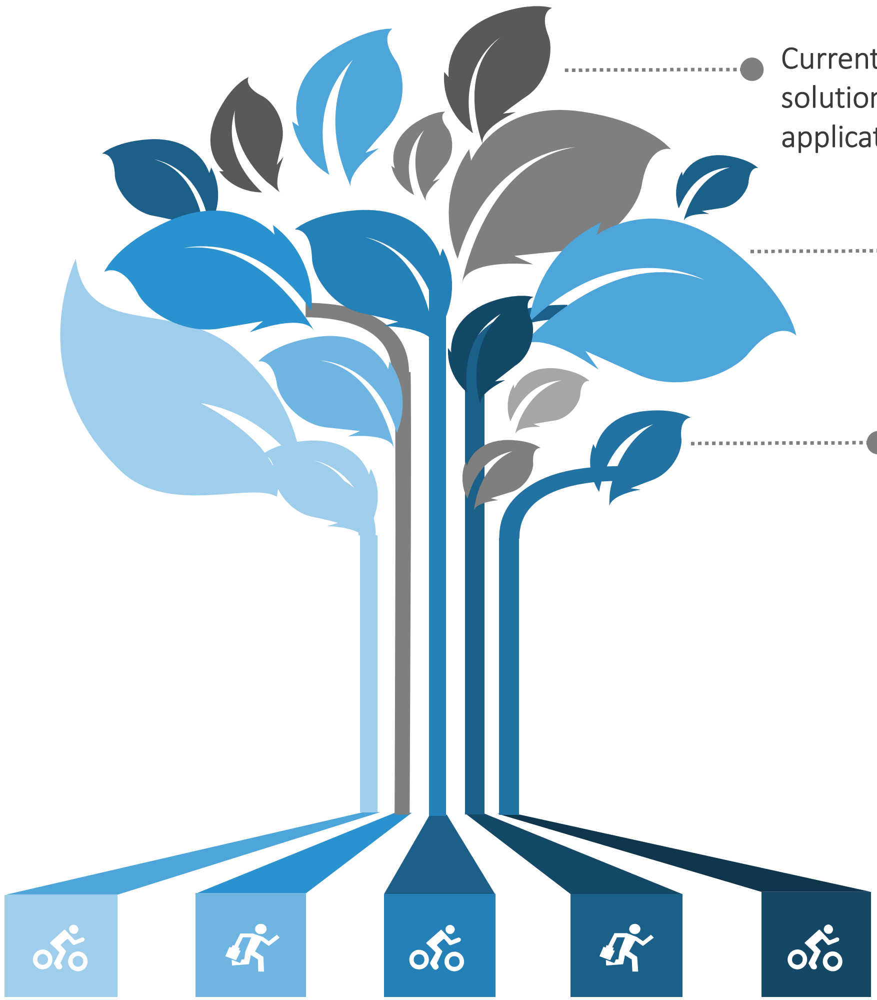

Current, traditional cloud computing and IoT solutions are not able to support real time applications

A new cost effective approach is needed
- New IoT systems could be closer to the data source
- Low latency services and applications are viable
- Data privacy could be increased

Connecting, configuring, and managing all these devices in the traditional way, e.g., manually, statically and per IoT domain is no longer viable

Combination of Smart Technologies

## Need for enhancing IoT

Smart solutions with embedded intelligence, connectivity and processing capabilities for edge devices, and support for real-time processing at the edge of the IoT network – near the end user

# TERMINET

https://terminet-h2020.eu

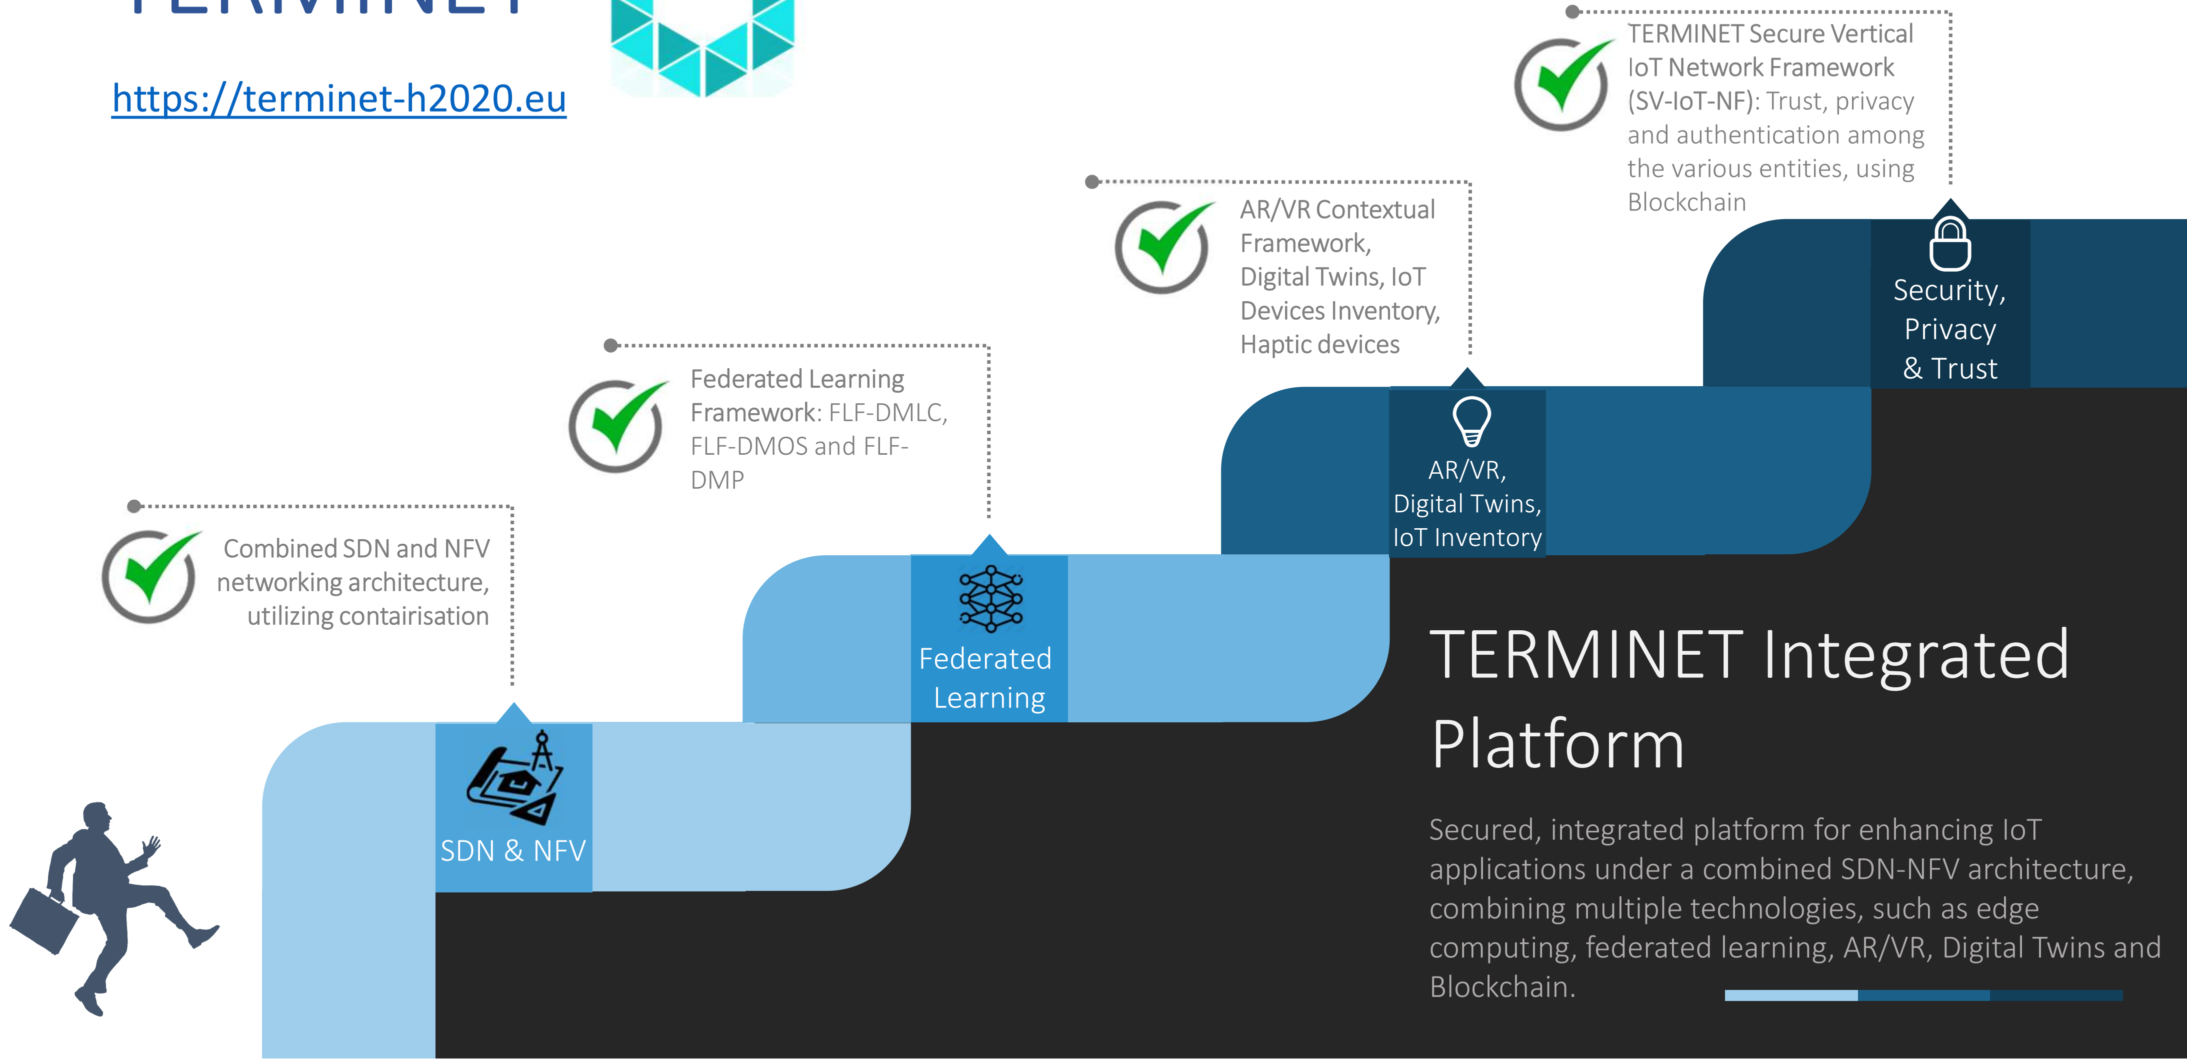

# Knowledge Fusion for Mastering Complexity

*Enabling situational awareness*

- Knowledge fusion seeks to provide an accurate and complete real-time view of a situation in a way that might be hard for a single human given the potentially very large amount of information and knowledge involved
- Consider disaster management, or efficiently managing a busy container port
- Knowledge fusion seeks to support people in their different roles to act as effectively as possible
- Digital twins should designed to integrate into a knowledge rich distributed environment
- This requires a systems perspective rather than thinking about sensors and actuators in isolation

# Evolution in action

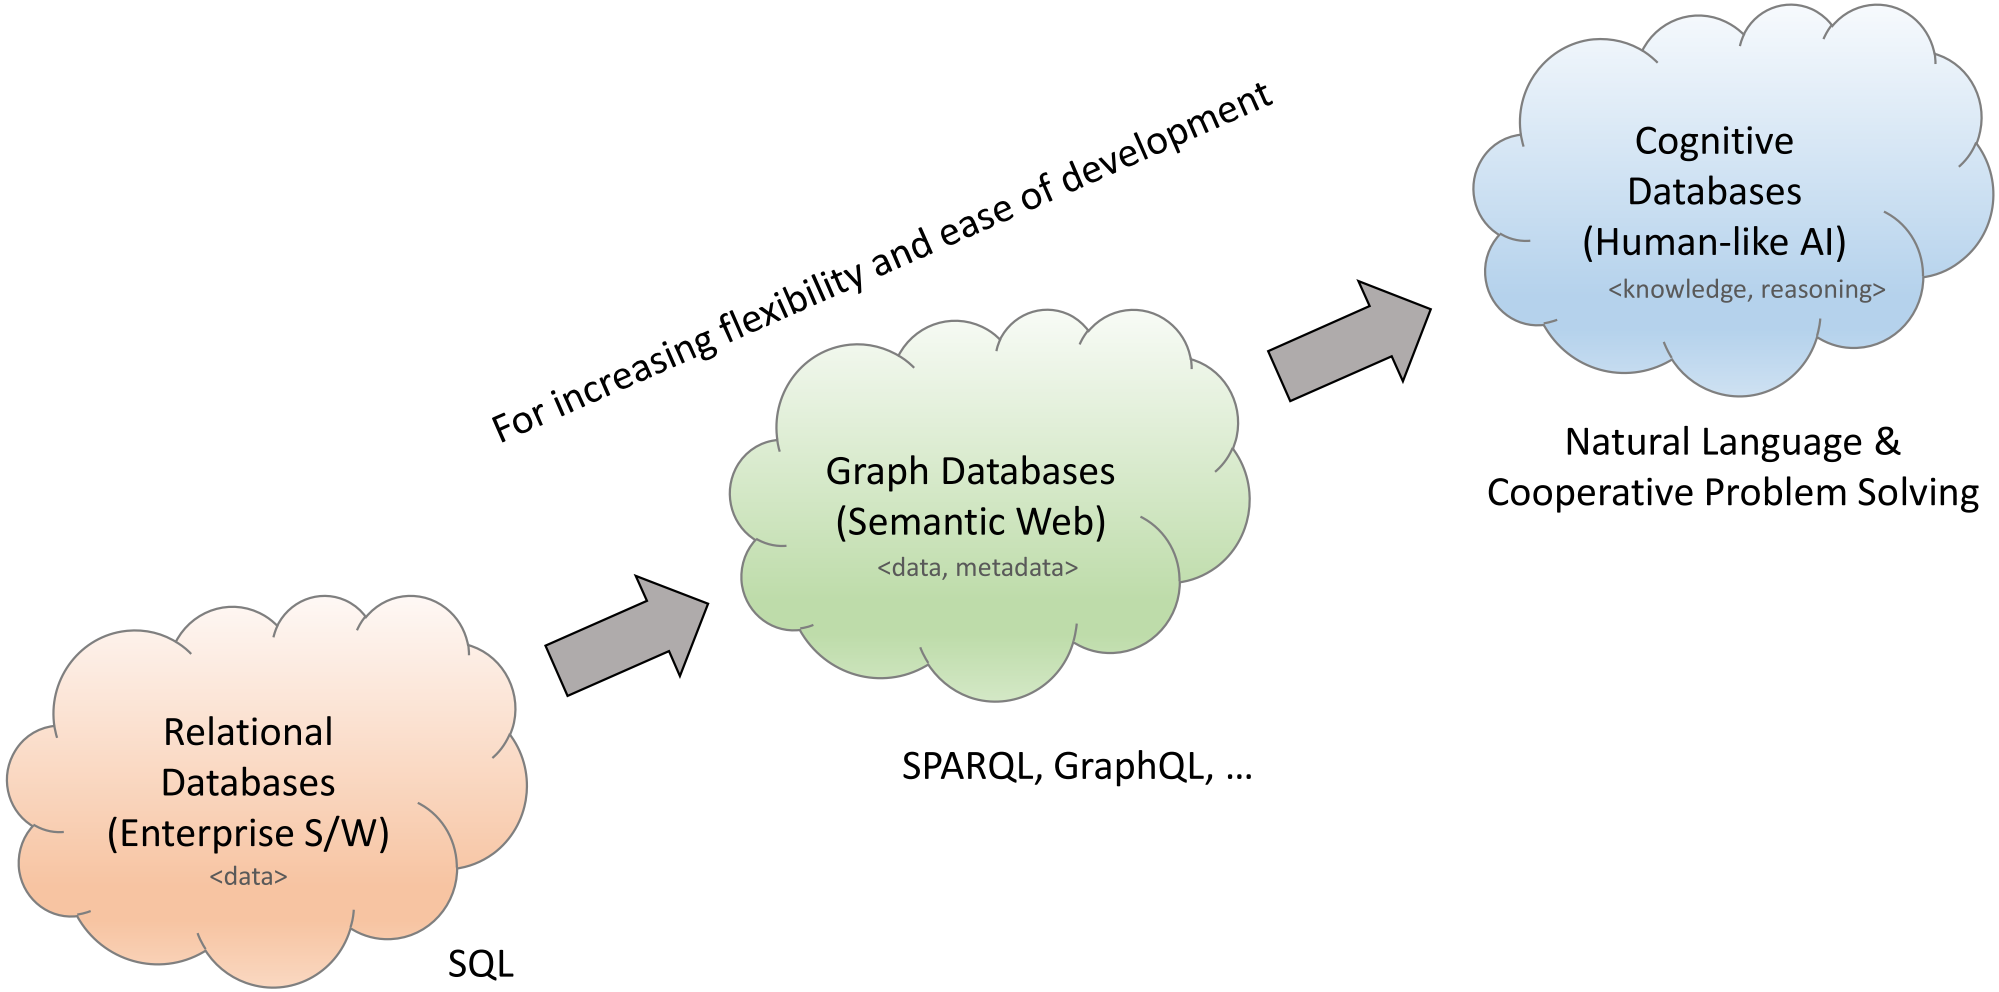

# Digital Twins and the Web of Things

- **Digital Twins** are a virtual manifestation of IoT sensors and actuators
- Simplify applications by abstracting IoT devices
- May support access to historical, current and (potential) future states: forecasting and planning
- Known as “Administration Shells” in Industry 4.0

- W3C’s **Web of Things** provides an abstraction for digital twins using RDF/JSON-LD
- Aimed at reducing cost for integrating heterogeneous IoT systems
- Thing Descriptions
  - Interaction affordances: properties, actions, events
  - Semantic descriptions - ontologies
  - Security & Protocol details

Data and metadata can be held in **Knowledge Graphs** for flexible representation

# The huge strides in AI of the last decade have been in machine perception, but *not* in machine reasoning

*Bradley Allen (2018)*

- Whilst today's knowledge graphs claim to capture knowledge there is very little attention currently to machine 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 machine reasoning for human-machine cooperative work
  - Boosting productivity and compensating for skills shortages

# Some acknowledgements

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

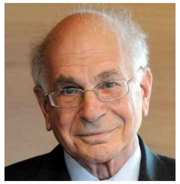

**Daniel Kahneman**, Nobel prize winning psychologist who studied System 1 & 2 thinking, along with cognitive biases in "Thinking Fast and Slow".

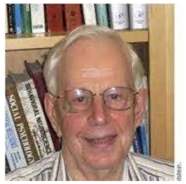

**Philip Johnson-Laird**, cognitive scientist renowned for his work on how humans reason in terms of mental models rather than using logic and statistics.

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

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

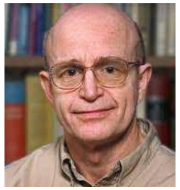

**John R. Anderson**, cognitive scientist renowned for his work on the ACT-R cognitive architecture for sequential cognition (System 2).

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

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

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

# 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 side-tracked 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 decide in human-like ways

*This is a **major** step forward for AI, complementing work on deep learning*



* See OpenAI's DALL•E 2 and DeepMind's AlphaFold

# Plausible Reasoning with Imperfect Knowledge

*Everyday knowledge is subject to uncertainty, incompleteness and inconsistencies: we're learning all the time, and consequently our current knowledge is imperfect*

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)*

# Web-based Proof of Concept Demo

The demo includes a variety of examples along with analogical reasoning and fuzzy quantifiers

- 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 (PKN*) 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 as metadata
  - certainty, typicality, similarity, frequency, dominance, conditional likelihood
- A collection of static reasoning strategies …
  - future work is planned on metacognition and continuous learning, including syntagmatic learning, paradigmatic learning and skill compilation

* Plausible knowledge notation for cognitive databases

# Plausible Inferences

- Let's start with something we want to find evidence for
  - flowers of England includes daffodils*

  *and evidence against it using 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*



* Using the plausible knowledge notation (PKN)

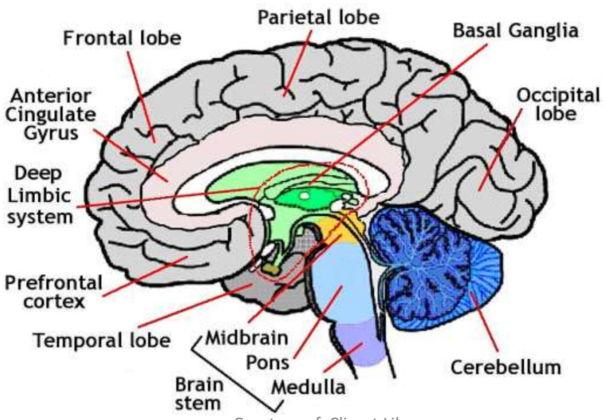

Courtesy of Clipart Library

Anterior temporal lobe as hub for integration across senses

# Cognitive Architecture for artificial minds

**Multiple cognitive circuits loosely equivalent to shared blackboard**

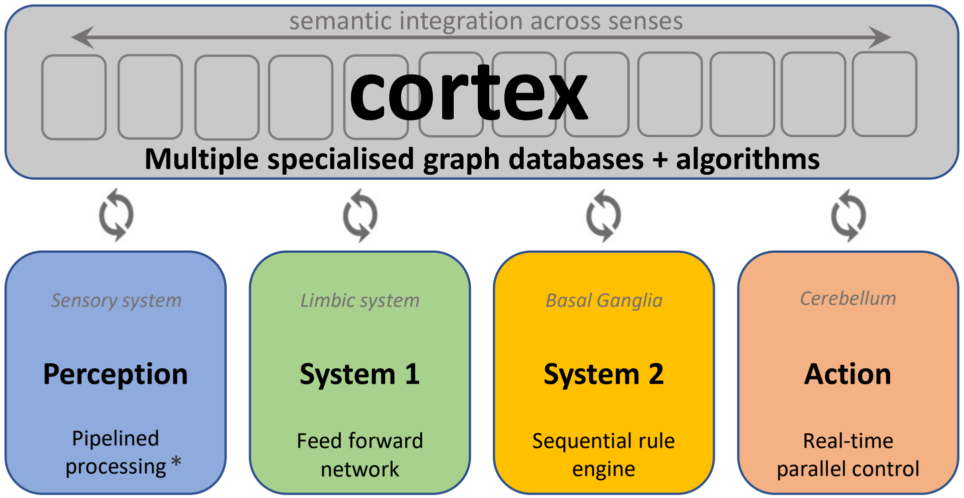


- **Perception** interprets sensory data and places the resulting models into the cortex. Cognitive rules can set the context for perception, and direct attention as needed. Events are signalled by queuing chunks to cognitive buffers to trigger rules describing the appropriate behaviour. A prioritised 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.

*e.g. based upon deep learning and artificial neural networks

# 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 (sets of name/value pairs)
  - My reworking of ACT-R in JavaScript
  - But unfortunately, still quite low level
- What about high level cognition?
  - Work underway on integrating natural language and everyday reasoning, along with combining System 1 and 2

* Large Language Models also do this, but in an opaque way

# Cortico-Basal Ganglia Circuit

## a functional model of System 2

**Cognition – Sequential Rule Engine**

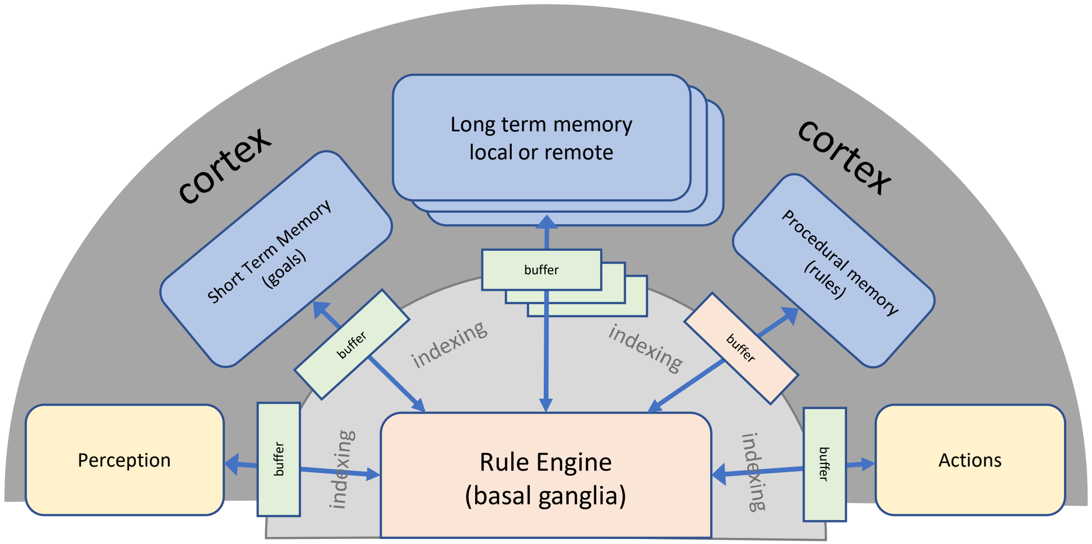


Cognitive Buffers hold single chunks
Analogy with HTTP client-server model

- 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

See W3C Cognitive AI Community Group for demos and specs

# Smart Factory Demo of a Cyber Physical System

https://www.w3.org/Data/demos/chunks/robot/

- **Cognitive AI demo that runs in a web page**
- Live simulation of bottling plant with robot, conveyor belts, filling and capping stations
- Real-time control by a cognitive agent

```
# add bottle when belt1 has space and wait afresh
space {thing belt1} =>
    action {@do addBottle; thing belt1},
    space {@do wait; thing belt1; space 30}

# stop belt1 when it is full and move arm
full {thing belt1} =>
    action {@do stop; thing belt1},
    action {@do move; x -120; y -75; angle -180; gap 40; step 1}

# move robot arm into position to grasp empty bottle
after {step 1} => robot {@do move; x -170; y -75; angle -180; gap 30; step 2}

# grasp bottle and move it to the filling station
after {step 2} =>  goal {@do clear}, robot {@do grasp},
    robot {@do move; x -80; y -240; angle -90; gap 30; step 3}
```

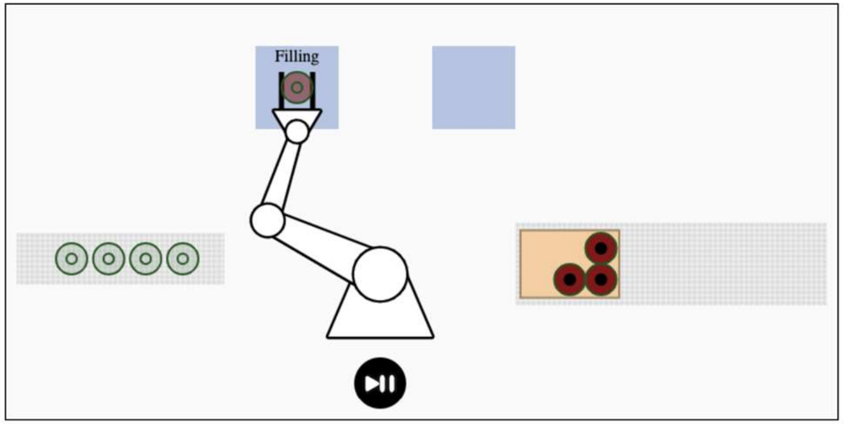


Log:

```
set goal to: after _:54 {step 1}
executed rule _:27 move
set goal to: after _:55 {step 2}
executed rule _:30 grasp
set goal to: after _:56 {step 3}
starting belt1
wait on filled
executed rule _:34 start
```

# Smart Home Demo

https://www.w3.org/Data/demos/chunks/home/

- Dynamic simulation of smart home
  - Live thermal model of heat flows between home and outside world
  - Control of lighting and heating
    - Manually
    - Automatically
  - Forms-based control of who is in the room, and the time of day
- Mix of declarative and procedural knowledge
  - Personal preferences and priorities
  - Example of default reasoning
- Web page with JavaScript library for Cognitive AI

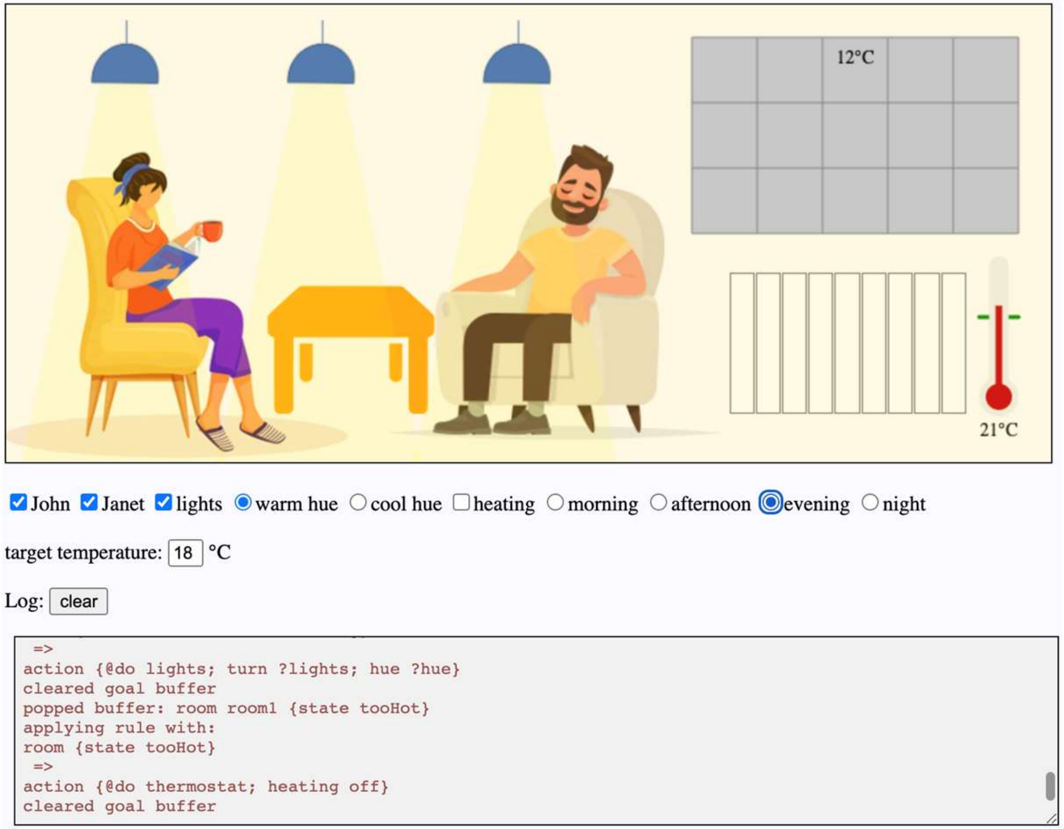

# Applying Cognitive Agents to Efficient Resource Management & Orchestration

Across the Computing Continuum: IoT, Edge and Cloud

Cognitive agents working on behalf of their owners

**Auctions** as established economic model for matching consumers to providers

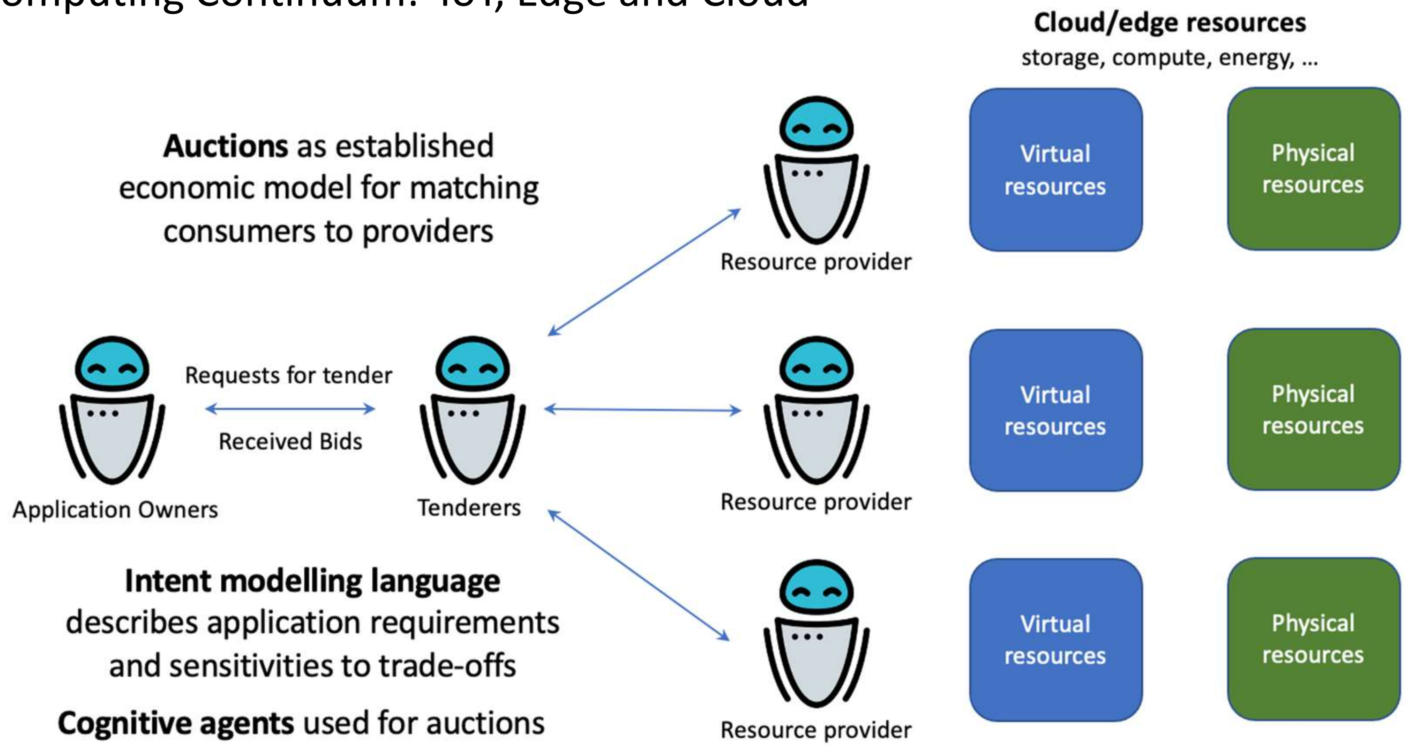


**Intent modelling language** describes application requirements and sensitivities to trade-offs

**Cognitive agents** used for auctions and throughout service life cycle

Robot icon courtesy of Pixel Perfect


Courtesy of Dave Lebow

# 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!

## Discussion time – your chance to ask questions!