Digital Transformation

and the knowledge economy

Dave Raggett dsr@w3.org

W3C/ERCIM

Siemens IoT Conference

17 March 2021

This work has been supported by the European Union's Horizon 2020 research and innovation programme under grant agreement 957 406, project TERMINET on next generation IoT

SPONSORED BY

Next Generation Smart Interconnected IoT

- Reference architecture featuring software defined networking, multiple-access edge computing, virtualised IoT, stream processing, federated learning, AR/VR and support for low-latency use cases
- 26 partners and 6 pilots
  - smart farming
  - personalised healthcare
  - smart sustainable and efficient buildings
  - prediction and forecasting system for optimising the supply chain in dairy products
  - group training surgery Using VR-enabled IoT technologies
  - mixed reality and ML supported maintenance and fault prediction of IoT-based critical infrastructure

Three year Horizon 2020 project that started in November 2020, see https://terminet-h2020.eu/

My early involvement with the Web

- In 1991 I was a researcher at HP Labs, Bristol and part of the Knowledge-based Programming Department, working on applying AI to helping HP customers to order HP computer systems
- The “Protek” pilot field tested a combination of hypermedia and expert systems to simplify preparing quotations covering all of the components necessary for a fully functioning system
- I then looked around to find others with an interest in eCommerce over the Internet and learned about Tim Berners-Lee’s work in CERN, Switzerland
- I persuaded my boss (Martin Merry) to fly out with me to Geneva to meet with Tim during 1992, and from then on was directly involved in early work on the World Wide Web
- I drove efforts to define a richer version of HTML (HTML+), implemented an early Web browser and server, and subsequently launched the HTTP WG at the IETF, along with work on HTML forms and tables
- HP assigned me to work with Tim at MIT to help drive the work further, after Tim moved from Geneva to Boston
- Since then I have worked on many different Web technologies, including the Web of Things, organising a W3C Workshop in 2014 that was hosted by Siemens

Tim Berners-Lee at CERN, birthplace of the Web

Photo courtesy of CERN

The World Wide Web Consortium (W3C)

Team photo from November 2013

- W3C's mission is to lead the World Wide Web to its full potential by developing protocols and guidelines that ensure the long-term growth of the Web
  - Interoperability standards for Web Browsers and for the Web of Data and Web of Things
- W3C is an member funded international standards development organisation with over 400 member organisations
- W3C is hosted by MIT's Computer Science and Artificial Intelligence Laboratory in the United States, at ERCIM in France, at the Keio University in Japan and at the Beihang University in China
- W3C's Team has over 60 people working from locations around the World

Digital Transformation

- Businesses exist to provide a return on investment through generating value from products and services
- Digital Transformation is the adoption of digital technologies throughout the enterprise
- The aim is to boost efficiency, increase resilience, and agility for exploiting change

- Software defined networking for private networks as overlays on heterogeneous networks, 5G and the public Internet
- OT, IT networks and cybersecurity
- Digital Twins, Big Data and AI/ML for process optimisation
- General purpose human-like AI for supervisory control and human-machine collaboration

NEW

Enterprise Systems

- Enterprise Resource Planning
  - Integrated management of business processes
- Customer Relationship Management
  - Retaining customers and driving sales
- Supply Chain Management
  - Management of flow of goods and services
- Manufacturing Execution Systems
  - Tracking processing of raw materials to finished goods

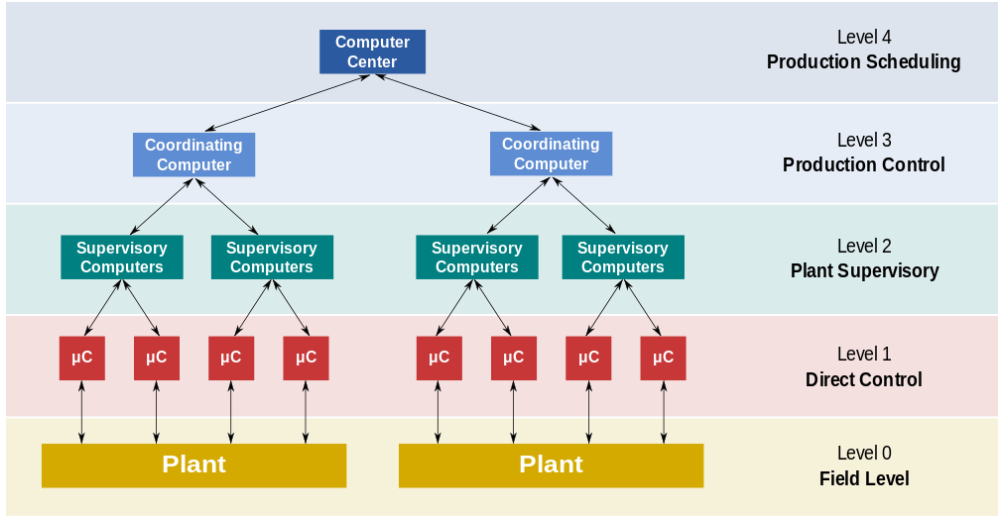

Courtesy of Wikipedia page on supervisory control

- Supervisory Control And Data Acquisition
  - Monitoring many individual controllers for an overall view of plant operation as part of MES
- Logical communication in a multi-agent system
  - Supervisors instruct workers
  - Workers keep supervisors informed
  - Supervisors monitor and decide on appropriate actions
  - Supervisors take on the role of workers in respect to their own managers
- Physical communication taking many forms
  - e.g. USB, Bluetooth, Modbus, OPC-UA, IP networks
- Machine understandable descriptions
  - Knowledge graphs and Data Governance
  - Accessible via searchable registries*
  - Expressed at both logical and physical levels
  - Use for discovery, planning and validation
  - Links between descriptions
  - Relationship to Web Architecture

* DNS is a registry of host names, but weak on metadata

Digital Enterprises
as multi-agent systems

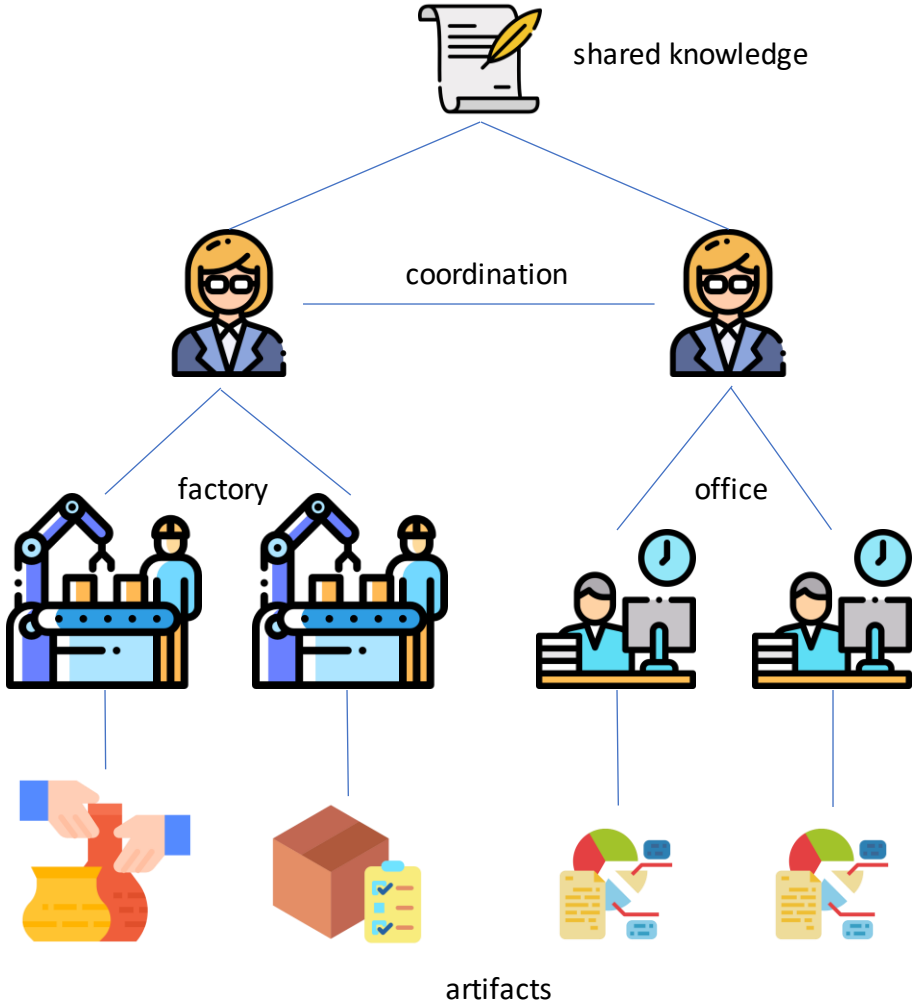

- Shared knowledge
  - **Databases**: read/write knowledge graphs, with access scoped according to need
  - **Ledgers**: signed non-erasable entries
- Supervisors
  - **General purpose human-like AI*** playing different roles at different levels of seniority, analogous to human teams, and working in collaboration with human colleagues
- Workers performing **specific** tasks
  - Robots, conveyor belts, AGVs, machining stations, PLCs, etc. with physical system + digital twin for zero defects and process optimisation
  - Virtual office workers as information services
- Physical and Virtual Artifacts
  - Materials and products under manufacture, with digital twins for service and usage histories
  - Information artifacts, e.g. certificates, contracts
  - Process flows of artifacts through production lines and business processes

* In contrast to narrow or weak AI

Example for a wine bottling plant

simple control of factory equipment

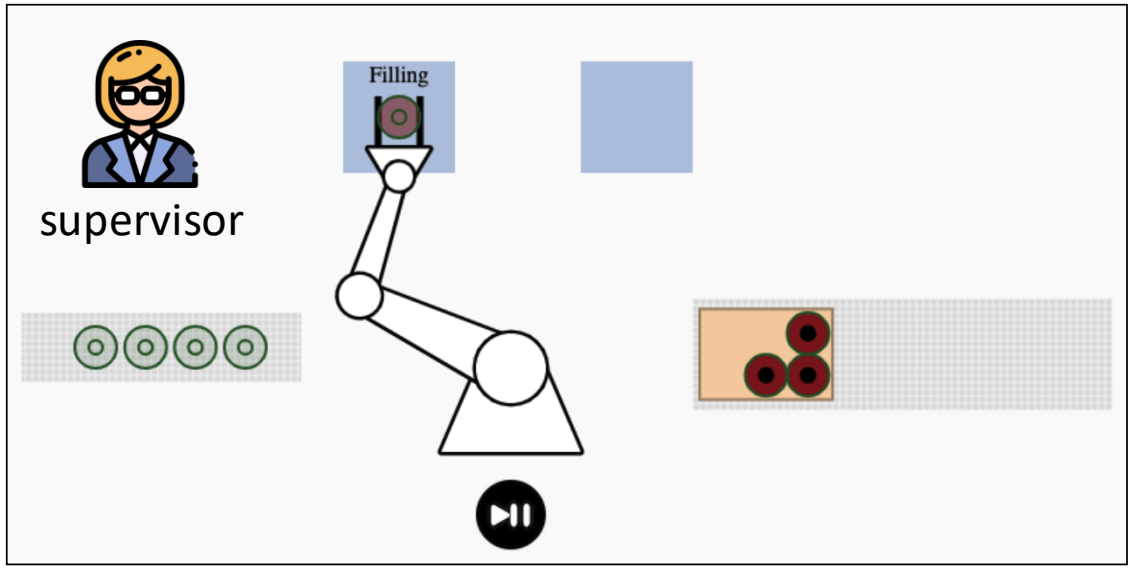

Demo: https://www.w3.org/Data/demos/chunks/robot/

Further technical details to follow in the second part of this talk

- Supervisor coordinates workers
  - Real-time concurrent control
  - Cognitive Rules + asynchronous messages
- Workers
  - Conveyor belts
  - Robot arm
  - Filling & Capping stations
- Artifacts
  - Batches of red wine
  - Empty bottles
  - Packing boxes
- Human-like AI is very scalable
  - Simple controllers on low cost platforms
  - Greater flexibility higher up the management chain
  - Phased deployment starting with simple controllers for SCADA and progressing to smarter supervisors for MES, SCM, ERP and CRM that work in collaboration with human colleagues

Data Ownership and What is it Worth*

agreements for data sharing

- Usage data from throughout the lifetime of a product can help in respect to timely servicing and improving product design
- Consumer resentment of 'digital slavery'†
  - Digital rights, copyright and GDPR
- Who owns the data: OEMs, service companies, consumers?
  - and where and how is the data stored?
- Clarifying which data we're talking about will help discussion on its use, e.g.
  - Vehicle performance data (OEMs)
  - Vehicle journeys (end user)
- Who owns sales, marketing and product data and how valuable is it?
- Shared data, subject to permissions, responsibilities and obligations
- Data storage considerations
  - Is access to your data at risk, and is it trapped in proprietary systems or third party repositories, you don't control?
- Data analytics as the process for making useful inferences from data along with actionable information
  - Analogous to human perception which takes place at staged levels of abstraction
- Data liability
  - Are you liable for misuse of data?
- Data marketability
  - Is your company data marketable?

† Slavery is the ownership of a person by another; its extension to the digital self

* See 2019 blog by Lou Washington

Data and the Circular Economy

- We are effectively raping the planet that we live and rely on
  - Unsustainable exploitation and degradation of the environment
- We must make the world better for our children and everyone and everything else!
  - Phasing out “rape, make, discard”
- European Commission’s plan to accelerate the circular economy*
  - Making EU climate neutral by 2050
  - Reducing consumption of raw resources and doubling rate of reusing materials through repair and recycling

- Implications for different sectors
  - Construction, packaging, plastics, food and water
  - Less waste, more value
- Opportunities for knowledge-based approaches to digital transformation
  - Better and more suitable materials
    - e.g. replacing plastics, steel and concrete
  - Design for repair and recycling
    - Making this easier and more cost effective
    - Longer term aim to mimic how natural ecosystems recycle organic materials
  - Data collection as basis for improvements
  - The IoT’s role in providing actionable information

* European Commission’s March 2020 Circular Economy Action Plan

W3C Web of Things

https://www.w3.org/WoT

- Abstraction layer for digital twins to offset fragmentation of the IoT
  - Many different protocols and standards
- Thing Descriptions in JSON-LD format
  - Logical definition of objects with properties, actions and events
    - Exposed to apps via local scripting API
  - Semantic models and relationships
    - What kind of thing is it?
    - How does it relate to other things?
  - Communications and security metadata
    - Used by the client platform to pick the appropriate protocols and data formats
  - Retrieving Thing Descriptions via HTTP
- Object-oriented vs process oriented perspective

How does the Web of Things relate to the big picture for digital transformation?

Web of Things – usage patterns

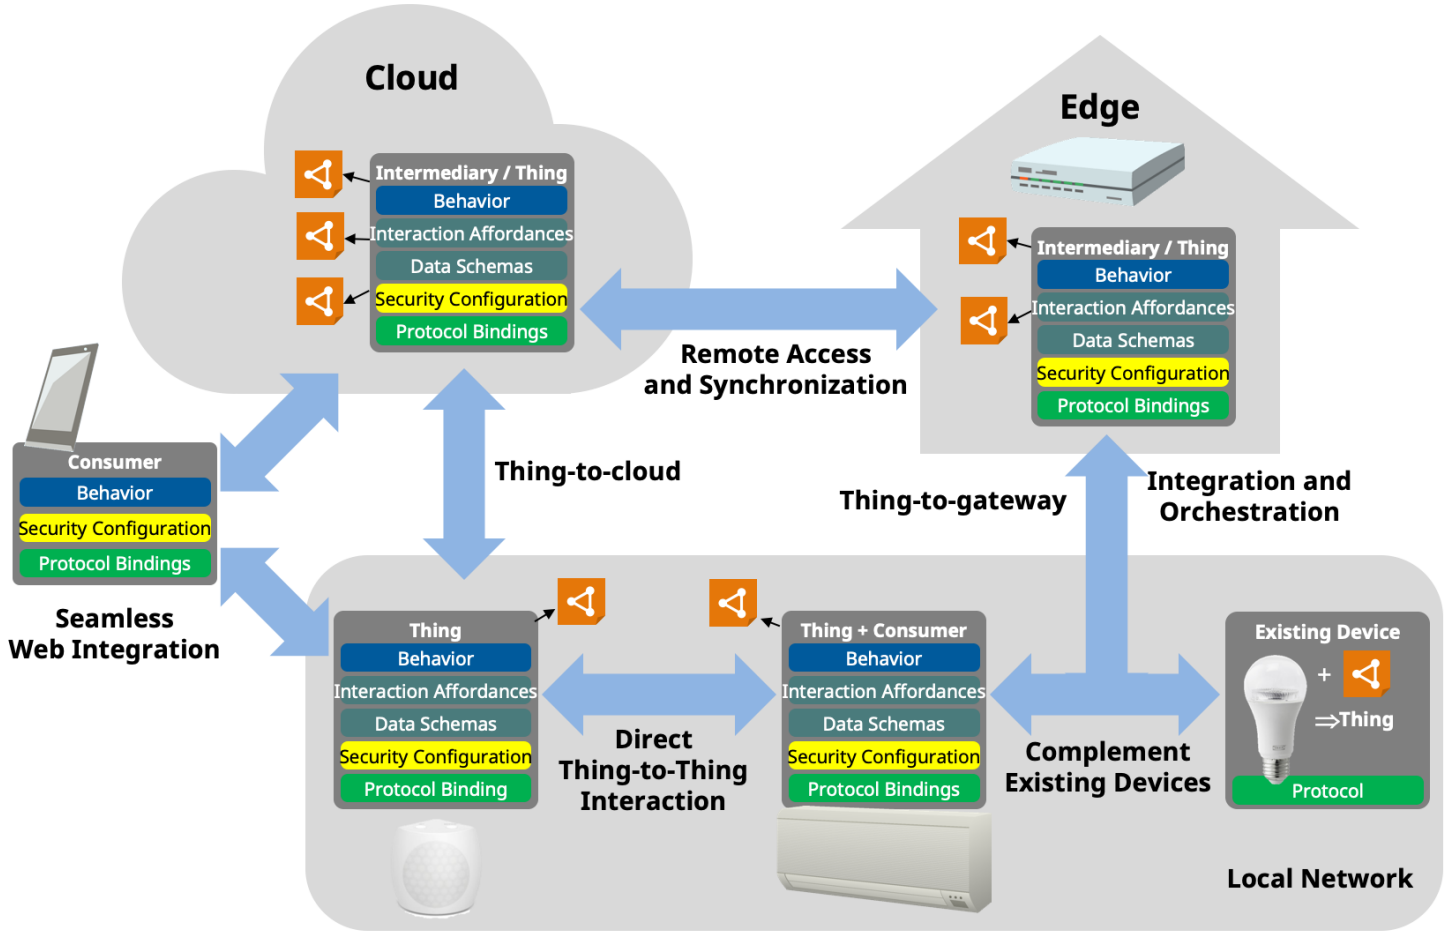

Contacts: Sebastian Kaebisch, *Siemens Corporate Technology* and Michael McCool, *Intel Technology Pathfinding*

Consumers want human-like approach, not human-like looks

- According to a recent survey by Capgemini
  - 64 % of consumers want AI to be more human-like
  - 62% are comfortable with human-like voice and intellect
  - 1 in 2 consumers say they are not comfortable with human-like physical features
  - 2 in 3 consumers want to know if they are interacting with an AI-enabled system or a human
  - 55% would prefer to have interactions enabled by a mix of AI and humans

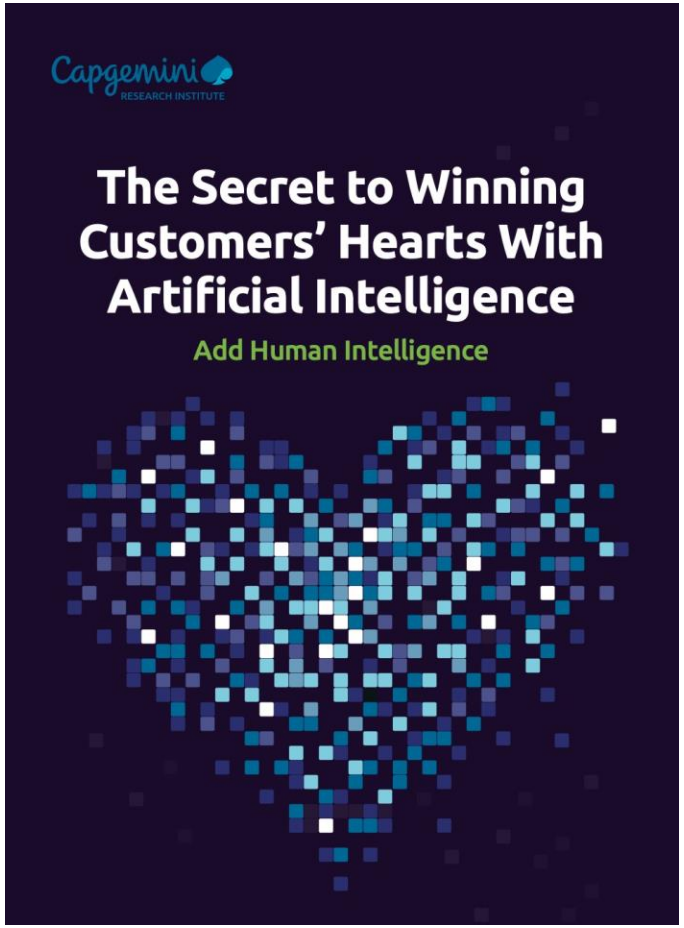

Survey conducted in 2018

Share of consumers who find human-like qualities compelling

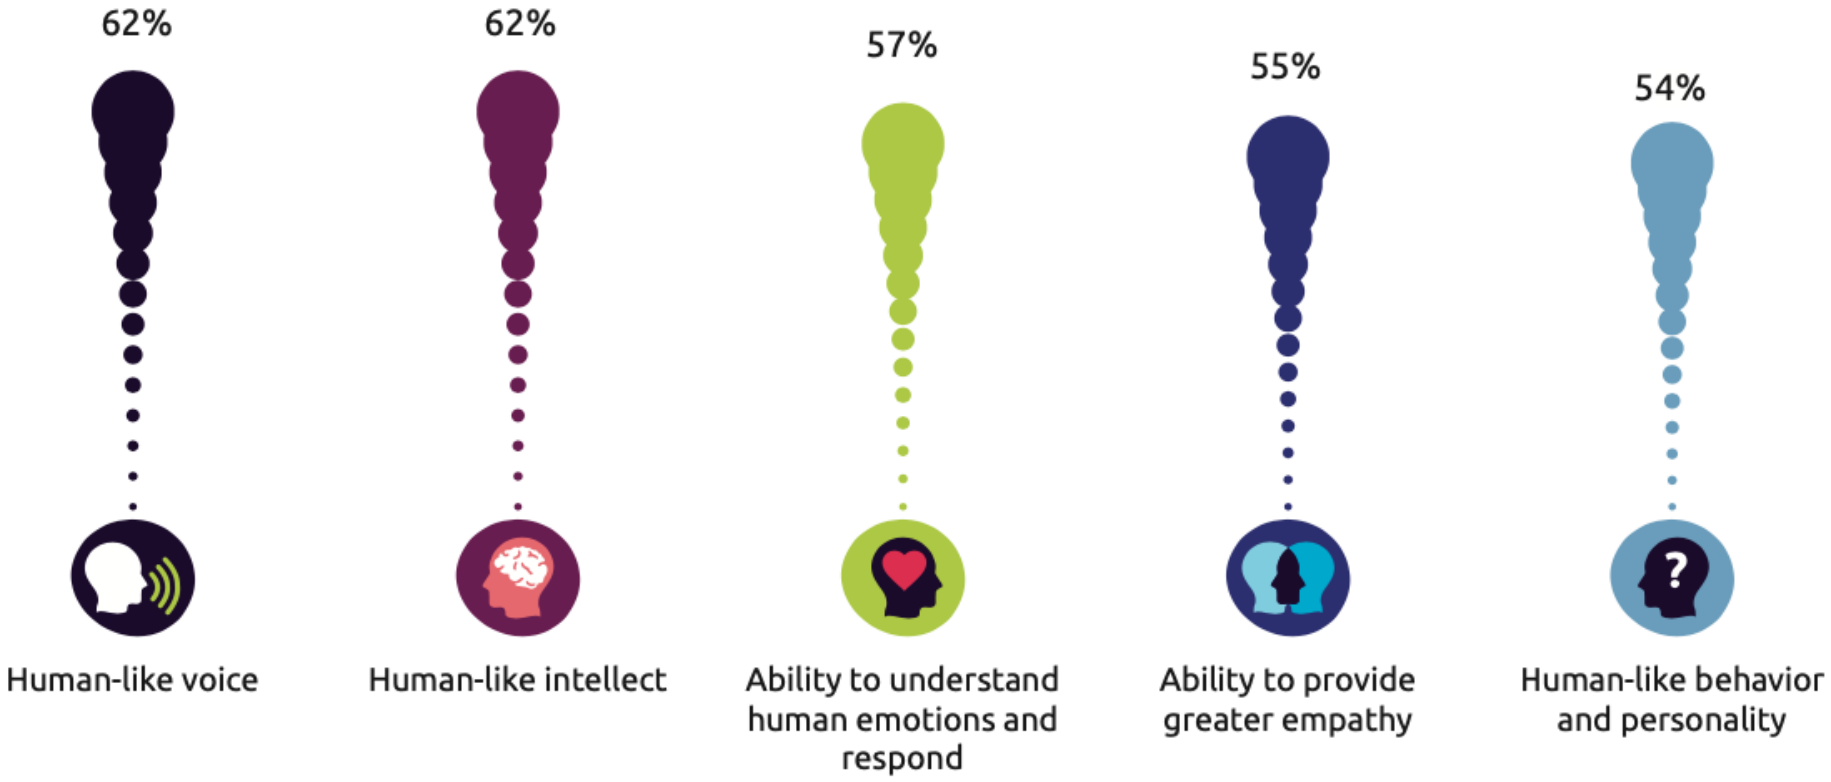

Source: Capgemini Research Institute, AI in CX Consumer Survey, May 2018, N=10,000 consumers.

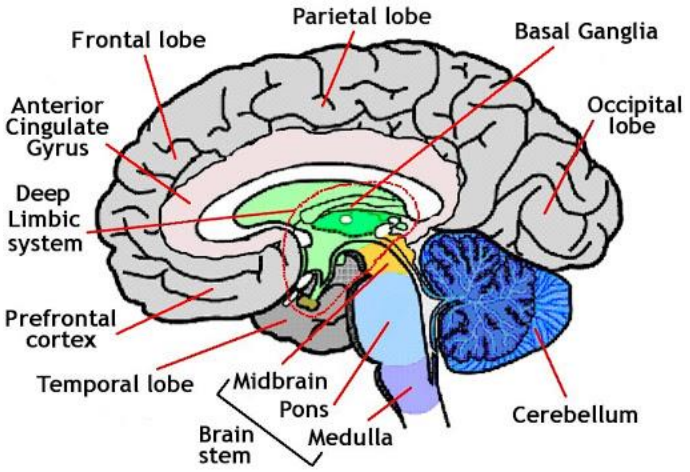

Courtesy of Clipart Library

Cognitive Architecture

for general purpose human-like AI

Cognitive Architecture with multiple cognitive circuits loosely equivalent to shared blackboard

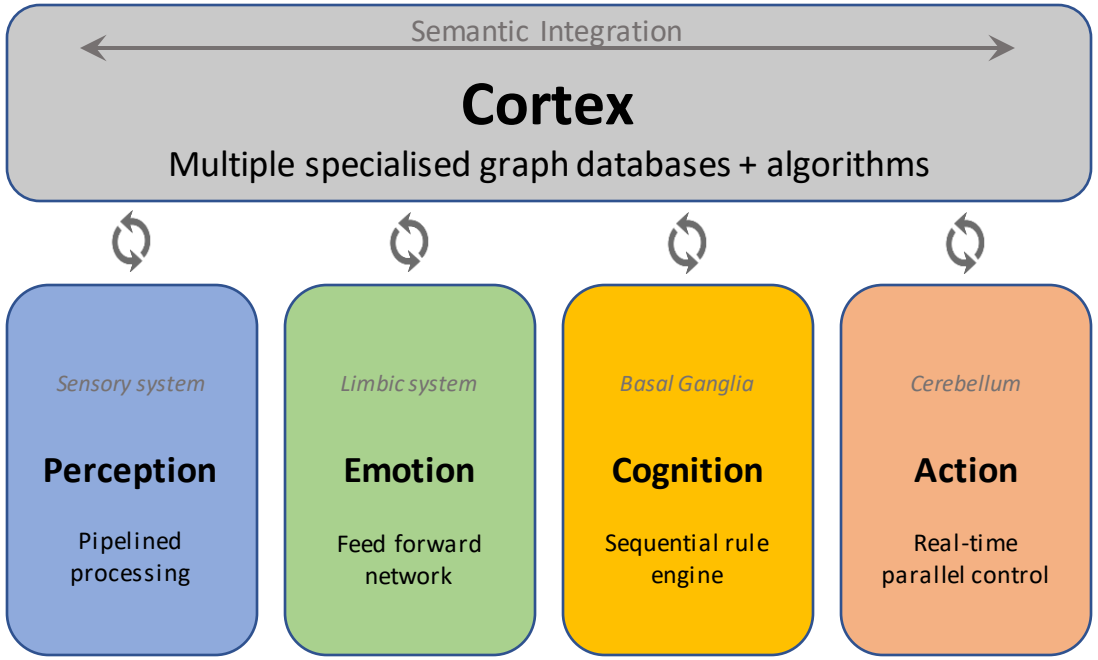

- **Perception** interprets sensory data and places the resulting models into the cortex. Cognitive rules can set the context for perception, and direct attention as needed. Events are signalled by queuing chunks* to cognitive buffers to trigger rules describing the appropriate behaviour. A prioritised first-in first-out queue is used to avoid missing closely spaced events.
- **Emotion** is about cognitive control and prioritising what's important. The limbic system provides rapid assessment of situations without the delays incurred in deliberative thought. This is sometimes referred to as System 1 vs System 2.
- **Cognition** 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.
- **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 coordinates muscle activation guided by perception.

* A chunk is a collection of properties that reference other chunks

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
- Key to prosperity of post-industrial societies as human populations shrink to a sustainable level
- Enabling us to safely exploit the resources of the solar system given the extremely harsh environment of outer space

Courtesy of Dave Lebow

Human-like AI

falling down the rabbit hole into a new world

- **Human-like in the sense of thinking like we do**
  - *Cognitive agents that are knowledgeable, general purpose, creative, collaborative, empathic, sociable and trustworthy*
  - *Metacognition and past experience to reason about new situations*
  - *Continuous learning based upon curiosity about the unexpected*
  - *Self aware in respect to current state, goals and actions*
  - *Awareness of others in respect to their beliefs, desires and intents*
  - *Multilingual, interacting with people using their own language*
- **Catalysing changes in how we live and work**
  - *Human-machine collaboration to boost productivity*
  - *Re-engineering capitalism in the post-industrial era*
- **Powering robots to help us in the physical world and beyond**
  - *Assisted living for people with cognitive or physical disabilities*
- **The Web 'verse* with distributed AR/VR as a place to meet, play, learn, do business, and much much more**
  - *Populated with avatars for humans and cognitive agents*
  - *Evolution of Web search with trusted personal agents*

* 1994 paper on VR Web

Technical Explanation and Demos

Cognitive AI FAQ

Human-agent collaboration and the **Sentient Web**

perception, cognition, and action federated across the Web

Human-like AI that mimics human memory, reasoning, feelings, and learning

(Democratising Artificial General Intelligence)

Web-based **Cognitive DB** chunks + algorithms

Smart data storage superseding property graphs and triple stores

- **Inspired by advances in the cognitive sciences and over 500 million years of neural evolution**
  - Functional models suitable for conventional computer hardware, complementing Deep Learning which is useful for big data, vision, etc.
- **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**
  - Each chunk is equivalent to concurrent firing of the bundle of nerve fibres connecting to a given cortical region
  - 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
  - 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

Cognitive Architecture with multiple cognitive circuits loosely equivalent to shared blackboard

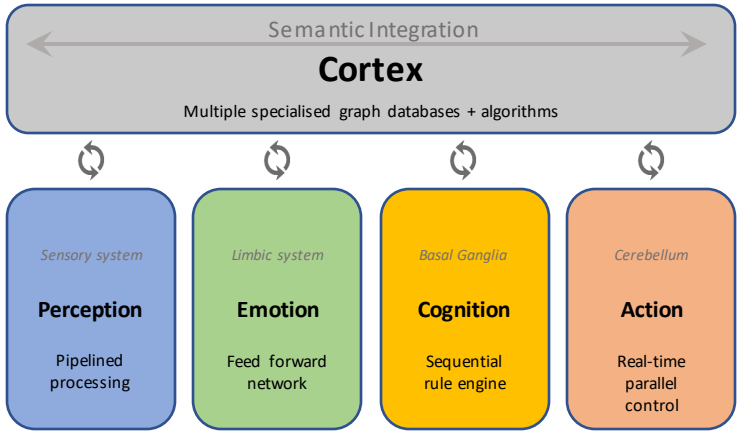

Cognition – Sequential Rule Engine

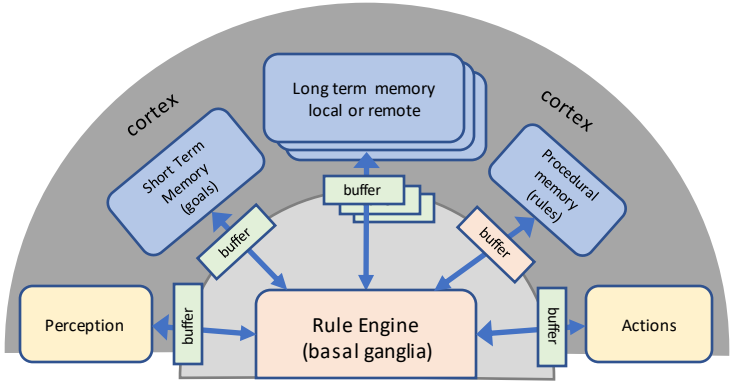

Cognitive Buffers hold single chunk
Analogy with HTTP client-server model

Contact: Dave Raggett <dsr@w3.org>

- **Natural language is key to human-agent collaboration as well as for teaching skills to bypass the manual programming bottleneck**
  - Human languages are complex yet easily learned by children, we need to emulate that for scalable AI
  - Explicitly represent semantics as chunk-based knowledge graphs in contrast to Computational Linguistics and Deep Learning which use large statistics as a weak surrogate
- **Syntax-semantics mapping rules and statistics shared between natural language understanding and generation**
  - Inductive generalisation from examples – incremental explanation-based continuous learning from experience
  - Informal task-related semantics fulfilling practical needs
  - Rule engine with concurrent asynchronous threads of execution that collaborate on refining interpretation
- **Lexicon, dialogue context, episodic, declarative and procedural memory are all represented with chunks with a simple syntax for rules, which act over chunk buffers and cortical algorithms**
- **The Sentient Web as federation of cognitive agents distributed across the Web with perception, reasoning and action**
  - Subsumes IoT, WoT and Semantic Web
  - Evolution of Web search: smarter & more personal
  - Pull-based ecommerce with trusted personal agents that works with other specialised agents on user's behalf
  - Personal agent collates rich personal information and shares directly relevant parts, subject to T&Cs, and based upon a model of your values & preferences as learned from your behaviour and those of others like you
  - Auction with 3rd parties to provide compelling offers
  - Integration with distributed AR/VR for the **Web 'verse**

Chunks

For details, see: https://github.com/w3c/cogai/blob/master/chunks-and-rules.md

Chunks is a simple amalgam of RDF and Property Graphs

Chunks correspond to concurrent firing patterns across bundles of nerves to a particular cortical region, see Chris Eliasmith's work on Semantic Pointers

Each chunk is a typed named collection of properties whose values are names or literals, e.g. numbers, booleans (i.e. true or false), dates, string literals or comma separated lists thereof*

A simple means is provided to map between chunks and RDF, mapping names to RDF URIs, and a short form syntax for chunks that denote single triples.

Here is an example of a chunk – you can use newline or semicolon as punctuation:

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

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

The chunk ID (e.g. *dog1)* is optional, and if missing, will be automatically assigned when adding the chunk to a graph. If the graph already has a chunk with the same ID, it will be replaced by this one.

You are free to use whitespace as you please, modulo the need for punctuation. String literals apart from URIs must be enclosed in double quote marks.

* A minimalist version of chunks limits properties to just the names of other chunks. *Type* is used for loose grouping and may imply an ontological constraint.

Chunk Rules

- Condition-action rules expressed as chunks with a convenient syntax for manual authoring when needed
- Updating any module buffer triggers rule engine to find and execute the best matching rule
- Stochastic selection of best rule from set of matching rules based upon their estimated utility according to past experience
- Conditions match content of module buffers
- Actions update buffers directly or invoke module operations, e.g. to recall a fact from memory, to assert a fact, or to invoke an external operation, e.g. to move a robot's arm, switch on a light, or to say "hello"
- Variables pass information from conditions to actions
- Rule chunks use @module to names the module it applies to, defaulting to "goal" module
- Module operations with @do
  - Built-in: *clear, update, queue, get, put, patch, delete, next, properties*
  - *Asynchronous except for clear, update and queue*
  - Applications can define additional module operations

```
# Given a goal like
#     count {state start; start 2; end 5}
# prepare to start counting using facts like
#     increment {number 1; successor 2}
count {state start; start ?num}
=>
    count {state counting},
    increment {@module facts; @do get; number ?num},
    console {@do show; value ?num}

# count up one at a time
count {state counting; start ?num1; end ~?num1},
   increment {@module facts; number ?num1; successor ?num3}
=>
   count {start ?num3},
   increment {@module facts; @do get; number ?num3},
   console {@do show; value ?num3}

# stop after last one
count {start ?num; end ?num; state counting}
=>
   count {@do update; state stop}
```

Autonomous Vehicles

- **Cognitive AI demo that runs in a web page**
- Mapping data for a small town was exported from Open Street Maps as XML (3.1MB) and transformed into chunks (637 KB or 128 KB compressed)
  - Points with latitude & longitude
  - Paths as sequence of points
  - Roads as collections of paths
- Graph algorithm for spatial indexing – constructs corresponding Quad Tree index with chunks
- Graph algorithm for route planning *("A star")*
- Visual model raises alerts that signal:
  - When approaching junction
  - When entering & leaving junction
  - When reaching the destination
- Cognitive rules as chunks for ease of learning
  - Start and stop turn indicator lights
  - Initiate braking or accelerating
  - Initiate lane tracking or turning
- Functional model of cortico-cerebellar circuit provides real-time control of brakes, acceleration and steering, as initiated by cognitive rules

Work in progress on integrating rule engine …

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

```
# 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}
```

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

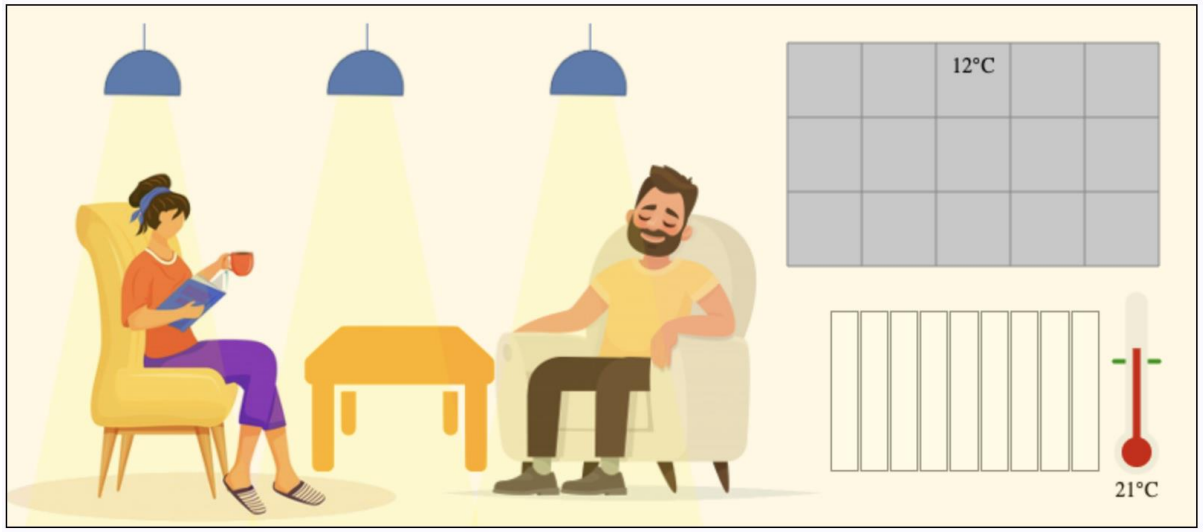

☑ John ☑ Janet ☑ lights ◉ warm hue ○ cool hue ☐ heating ○ morning ○ afternoon ◉ evening ○ night

target temperature: 18 °C

Log: clear

```
 =>
action {@do lights; turn ?lights; hue ?hue}
cleared goal buffer
popped buffer: room room1 {state tooHot}
applying rule with:
room {state tooHot}
 =>
action {@do thermostat; heating off}
cleared goal buffer
```

Smart Factory Demo

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}
```

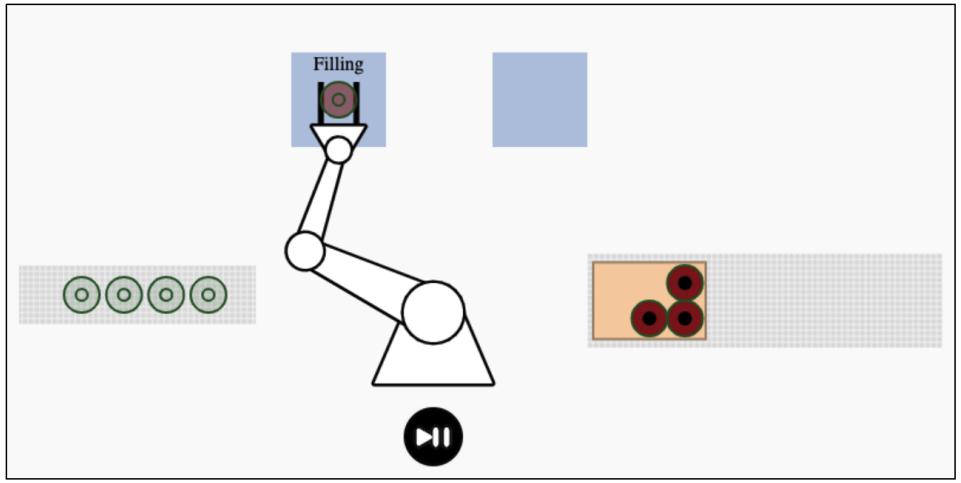

Log:

```
executed rule _:19 stop
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
```

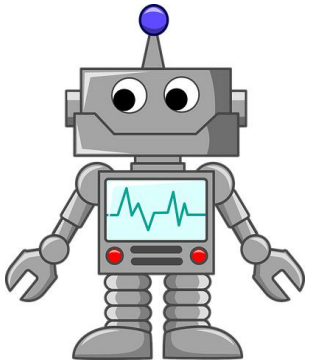

Courtesy of pixabay

Robot Control

- Perception
  - State models dynamically updated in declarative memory (cortical modules)
  - Events modelled as goals that trigger rule sets to handle them
  - Goal queue to avoid missing closely spaced events
- Actions
  - Concurrent asynchronous execution of actions analogous to HTTP request/response pairs
  - Execution is delegated to functional model of cortico-cerebellar circuit leaving the rule engine free to keep running and responding to other events
- Robot arm, conveyor belts, filling and bottling stations are all modelled as functions of time
  - Using high precision timer and plenty of trigonometric calculations
  - Robot has 3 rotational joints and a gripper, these are all smoothly accelerated and decelerated according to their individual capabilities
- Threaded Control with continuations
  - Doing something when something else has happened (no need to wait if it has already happened)
  - Waiting for space to free up at the start of the conveyor belt
  - Waiting for a robot motion to complete
  - Waiting for a bottle to be filled
- Integration with an existing robot
  - Robot exposes network API
  - Cognitive AI for high level control
- Robot demo with lightweight ontology
  - Validation of rules against available actions
  - Planning as basis for reconfiguring production
  - Meta-reasoning for resilience when needed
- Opportunities for richer human-machine interaction
  - Natural language and emotional intelligence

Other Web-based Demos

- Web-based demos that allow everyone to try things out themselves
  - No software installation required!
- JavaScript library for chunks and rules
- Easy to use from web page scripts
- Further technical work is planned on
  - Spreading activation & stochastic recall
  - Compiling rules into discrimination network akin to RETE algorithm
  - Reinforcement learning of rule sets
  - Exploration of Web Assembly and hardware acceleration
  - Exploration of holographic memory

- Counting 1, 2, 3, …
  - Ported from ACT-R tutorial
- Simple decision trees
  - How's the weather today? Is it suitable for a round of golf?
- Test suite for Chunks and Rules
  - In support of the formal spec
- Sandbox for getting started with Chunks and Rules
  - Edit, save and single step chunk facts and rules from within a web page
- Natural language demos
  - And ongoing work on NLP

Courtesy of snappygoat.com

Social Mimicry
it makes us human

- Mimicking others
  - Babies learning to smile from interacting with their mothers*
  - Children copying speech sounds of their peers (regional accents)
  - Learning how to say complex utterances by listening to others
  - Imitating dance movements of others on the dance floor or TV
  - Playing some music on a piano or guitar after listening to it
  - Choosing the same styles of clothes as your friends
- Socially driven
  - Emotionally satisfying, a feeling of belonging
- A common cognitive architecture
  - First, an internal model has to be learned from lower level sensory data, via increasing levels of abstraction, across multiple modalities
  - Second, you have learn how to map this internal model to a lower level model for motor control, via decreasing levels of abstraction, for execution by the cerebellum
  - Statistics for recognition of patterns is shared with their generation, e.g. shared across natural language understanding and generation
  - Incremental learning involving only weak supervision, and evolving effective models from many potential alternatives

* Smiling: there are 42 facial muscles and 6 muscles for moving each eye

Natural Language as social communication

And for teaching skills to cognitive agents as solution to scaling

- Cognitively plausible processing model for understanding and generation of natural language
  - Aim: to learn language like children do
- Incremental word by word concurrent syntactic and semantic processing without the need for any backtracking
  - Inspired by eye tracking data when reading text
- Production Line Metaphor – each stage in the line progressively elaborates and transforms information
  - Phonology, Morphology, Words, Phrase structure, Semantics, Pragmatics
- Use of statistical information to guide choices, e.g. for priming effect on word senses
  - Shared statistics for understanding and generation
- Simple robust shift-reduce parsing with implicit grammar and small set of word classes for part of speech
  - Parse tree and lexicon expressed with chunks

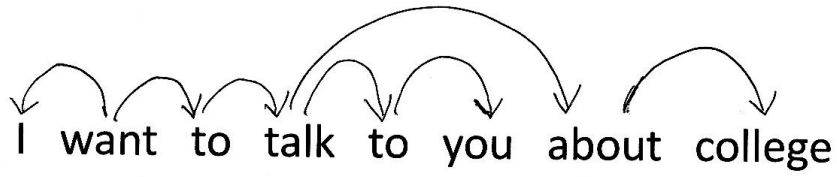


```
# I want to talk to you about college
vp _:1 {verb want; subject _:2; to _:3}
np _:2 {pron i}
np _:3 {noun talk; to _:4; about _:5}
np _:4 {pron you}
np _:5 {noun college}
```

NLP as Concurrent Processing

John gives a book to Mary.

Word	Syntax	Semantics
John	Shift: np _:1 {noun John}	A named person
gives	Shift: vp _:2 {verb gives}	Action: to give
	Reduce: vp _:2 {verb gives; subject _:1}	argument *from*
a	Shift: np _:3 {det a}	
book	Extend: np _:4 {det a; noun book}	A book
to	Shift: pp _:5 {prep to}	
Mary	Shift: np _:6 {noun Mary}	A named person
	Reduce: pp {prep to; np _:6 }	
	Reduce: vp _:2 {verb gives; subject _:1; to _:6 }	argument *to*
	Reduce: vp _:2 {verb gives; subject _:1; object _:4; to _:6 }	argument *object*

- Table shows concurrent processing for each word in respect to phrase structure (syntax) and semantics
- The semantic graph describing the meaning is built step by step
- Nouns and Pronouns need to be mapped to what they refer to
- Verbs and auxiliaries are mapped to a model of when an action occurred and whether it is extended in time or a moment in time
- More complex examples involve syntactic and semantic ambiguities that can't be resolved immediately
- Examples include semantic priming of word senses, whether a word is part of compound noun, and whether a word is an object or an indirect object

Initial Experiments on NLP

- Use of text or speech to move discs in the towers of Hanoi game
  - https://www.w3.org/Data/demos/chunks/nlp/toh/
  - Initial proof of concept for shift-reduce parsing with chunks
- Dinner demo with two cognitive agents
  - https://www.w3.org/Data/demos/chunks/nlp/dinner/
  - Agents exchange chunks, whilst invoking speech API
  - Rules describe transitions between named tasks
  - Generalisation using plans and causal reasoning
- Parsing demo tests that parser is adequate for all dinner dialogue utterances
  - https://www.w3.org/Data/demos/chunks/nlp/parsing/
- Ongoing work on end to end communication of meaning
  - Modelling concurrent processing at different stages in the NLP pipeline for both understanding and generation
  - Future work on mimicking how children learn language

Towers of Hanoi

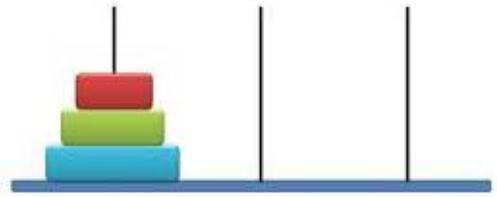

```
# move the red disc to the right peg
verb v1 {word move; subject p1; to p2}
phrase p1 {word disc; det the; adj red}
phrase p2 {word peg; det the; adj right}

# after application of ruleset
move m1 {disc disc3; to peg3}
```

Dinner Dialogue

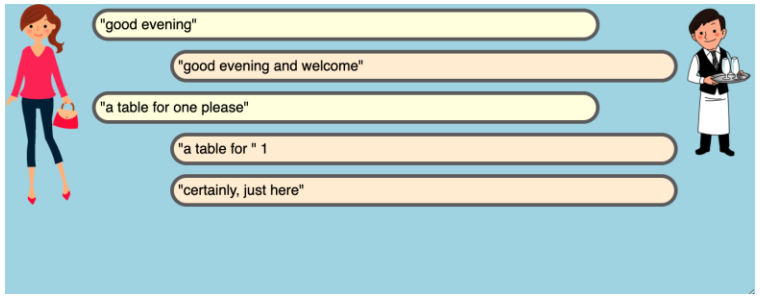

W3C Cognitive AI Community Group

See: https://www.w3.org/community/cogai/, https://github.com/w3c/cogai

- Participation is open to all, free of charge
- Focus on demonstrating the potential of Cognitive AI
  - A **roadmap for developing AI** that is general purpose, collaborative, empathic and trustworthy
- Collaboration on defining use cases, requirements and datasets for use in demonstrators
  - https://github.com/w3c/cogai/tree/master/demos
- Work on open source implementations and scaling experiments
- Work on identifying and analysing application areas, e.g.
  - Helping non-programmers to work with data (worth $21B by 2022 according to Forester)
  - Cognitive agents in support of customer services (worth $5.6B by 2023)
  - Smart chatbots for personal healthcare
  - Assistants for detecting and responding to cyberattacks
  - Teaching assistants for self-paced online learning
  - Autonomous vehicles
  - Smart manufacturing
- Outreach to explain the huge opportunities for Cognitive AI

Backup

Courtesy of Dave Lebow

Human-like AI

falling down the rabbit hole into a new world

- **Responsible AI** that learns and applies human values
  - Overcoming perils of prescribed rules of behaviour and inevitability of unforeseen effects, e.g. Azimov's 3 laws
- **Digital Self** as the evolution of your digital presence into a trusted personal agent that learns and applies your values, personality, memories and skills
  - Safeguards your privacy and personal data
  - Acts on your behalf when you are busy or offline
  - You are liable for actions by your digital self
  - Digital life after physical death – question of ownership
- **Abolition of digital slavery***
  - Slavery is the ownership of one person by another
  - Surveillance capitalism is the business of digital slavery and incompatible with democracy
  - We need to extend human rights to digital rights, and free our digital selves from slavery for others
  - Should also apply to synthetic humans as human-like AI

* With grateful acknowledgement to Aral Balkan's nature of the self in the digital age