Human-Like AI

Practical roadmap for realising general purpose AI

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

17th March 2021

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
- I have participated in many Horizon projects with European partners

Tim Berners-Lee at CERN, birthplace of the Web

Photo courtesy of CERN

What is Human-like AI?

- Artificial Intelligence (AI) lacks a precise agreed definition, but loosely speaking, it is about replicating intelligent behaviour, including perception, reasoning and action
- Human-like AI is general purpose AI inspired by over 500 million years of evolution of neural systems and many decades of progress across the cognitive sciences
- Emulation of the human brain at a functional level
  - The phenomenological requirements are essentially independent of how they are realised, e.g. as explicit graphs, vector spaces or pulsed neural networks, see David Marr's three levels of analysis

Comparison with other branches of AI

- Good old fashioned AI (e.g. expert systems and the semantic web) is based upon symbolic representations, but not statistics
- Deep Learning embodies statistics, but not symbols, where statistics act as a weak surrogate for semantics
- Human-like AI embodies both symbols and statistics with an explicit representation of semantics
- Old fashioned AI requires manual development of knowledge which acts as a straitjacket when it comes to scaling up. In addition, it has difficulties with uncertainties and inconsistencies
- Deep Learning has been very successful, and is good for scaling, but learns very slowly compared to humans, requiring large training sets, and suffers from a lack of explainability, a lack of understanding of salience, a lack of generality and an inability to support higher level reasoning
- Human-like AI, by contrast, seeks to mimic human abilities, using prior knowledge and past experience to speed learning, symbolic representations for reasoning and explanations, and metacognition for generality
- Traditional approaches to meaning have been based on the Aristotelian tradition of logic, formal semantics and model theory. This deals with what is provably true given a set of assumptions and inference rules. This approach underpins the Semantic Web and ontologies based upon OWL
- It can be contrasted with studies of human reasoning. Philip Johnson-Laird for instance, notes that humans don't rely on the laws of logic and probability, but rather by thinking about what is possible
- In other words, rather than reasoning based on mathematical proof, we consider examples and how they fit the problem under consideration. As work in cognitive linguistics has shown, we often use metaphors and analogies as a basis for thought
- We also use simple rules of inference in situations where such rules have been found to be appropriate. Such rules are part of the mix of declarative and procedural knowledge we bring to bear on particular tasks
- Human-like AI can be deployed cheaply on conventional computer hardware

If Human-like AI mimics human thought, does it make the same kinds of mistakes as humans?

- Yes, and just like people, this is a matter of training
- A well trained cognitive agent will do well in respect to the tasks it was designed for
- Unlike people, trained cognitive agents can be trivially cloned as needed!
- Systems based upon logical deduction and formal semantics offer mathematic proof given assumptions and inference rules, but are intolerant of inconsistencies and unable to exploit the statistics of past experience
- Human-like AI, inspired by human reasoning, and a combination of graphs and statistics, focuses on what is useful based upon prior knowledge and past experience in the face of uncertainty and inconsistency

Cognitive Architecture
for general purpose human-like AI

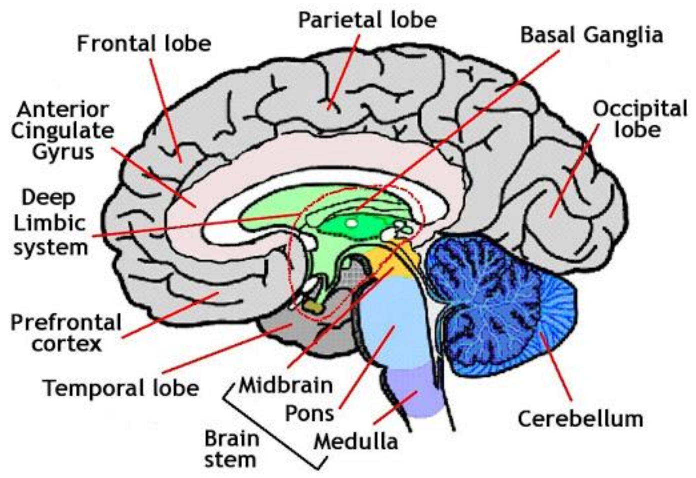

Courtesy of Clipart Library

Cognitive Architecture with multiple cognitive circuits loosely equivalent to shared blackboard

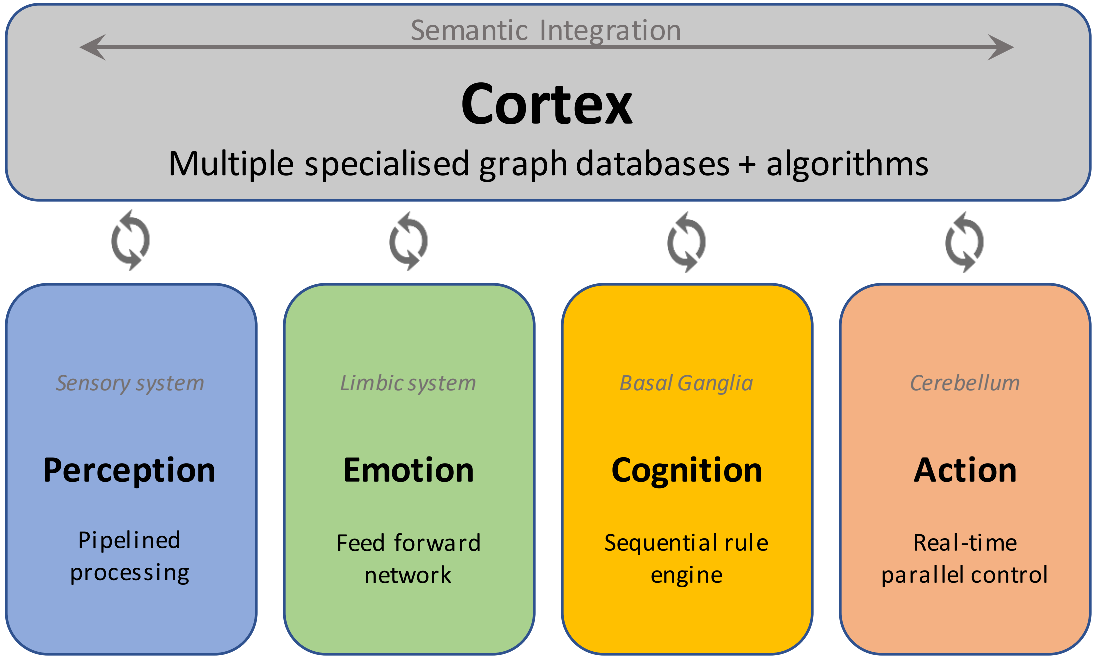

- **Perception** interprets sensory data and places the resulting models into the cortex. Cognitive rules can set the context for perception, and direct attention as needed. Events are signalled by queuing chunks* to cognitive buffers to trigger rules describing the appropriate behaviour. A prioritised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

Applications

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

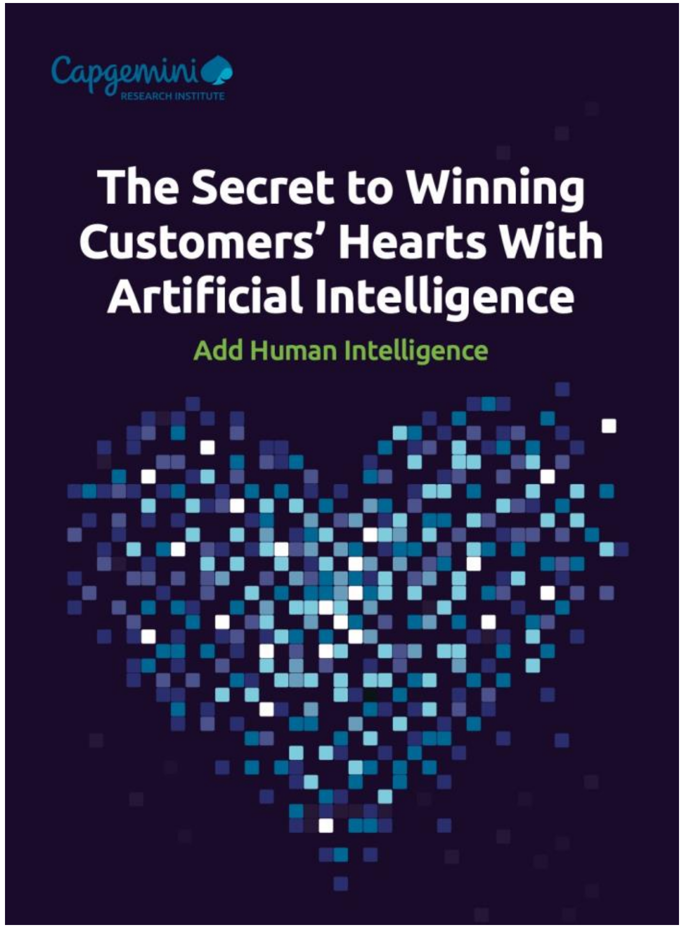

Survey conducted in 2018

Share of consumers who find human-like qualities compelling

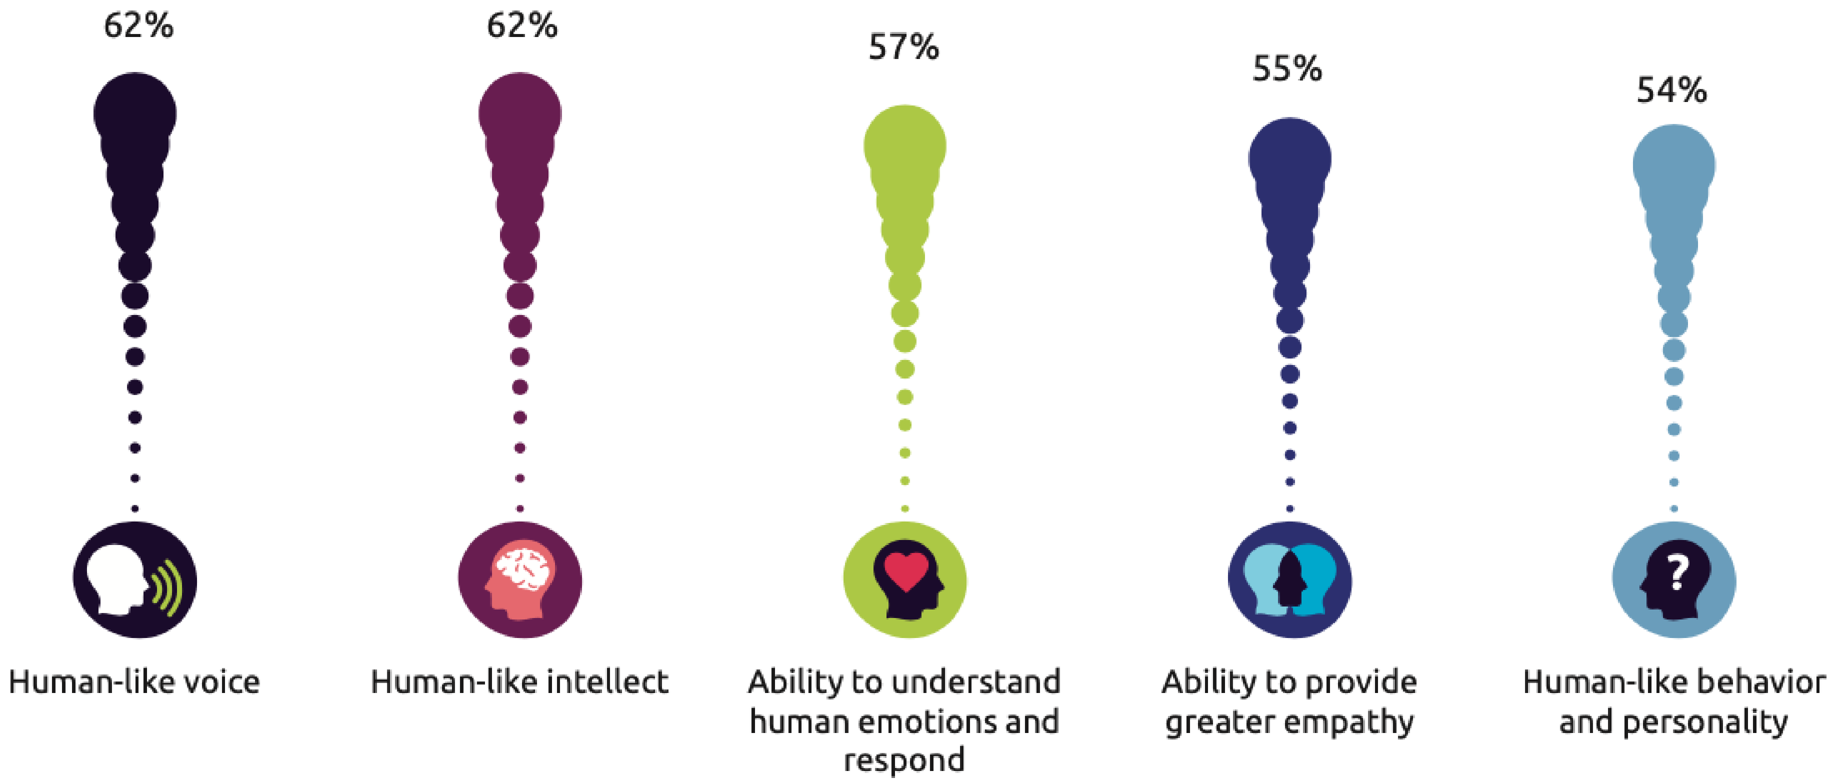

Source: Capgemini Research Institute, AI in CX Consumer Survey, May 2018, N=10,000 consumers.

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

Icons courtesy of flaticon, ultimatearm and geotatah

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

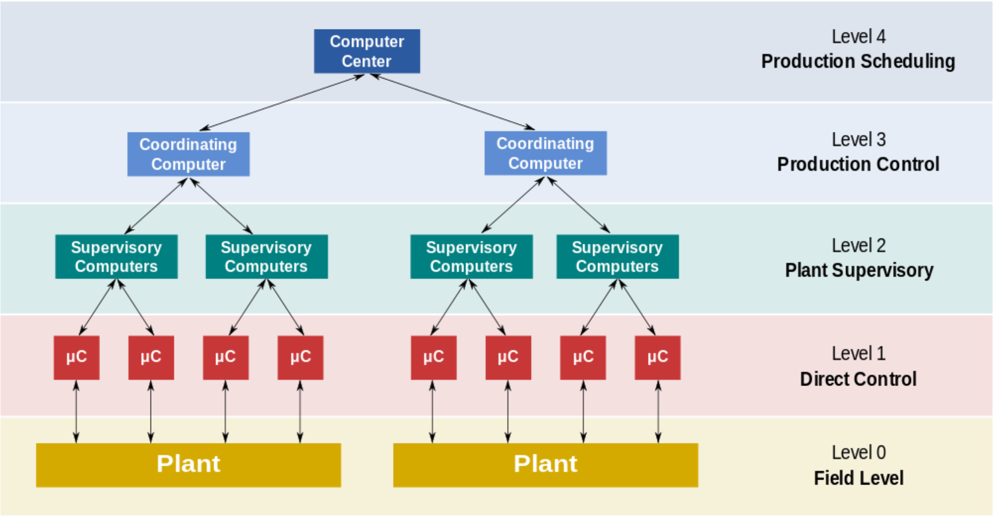

Courtesy of Wikipedia page on supervisory control

- Supervisory Control And Data Acquisition
  - Monitoring many individual controllers for an overall view of plant operation as part of MES
- Logical communication as 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

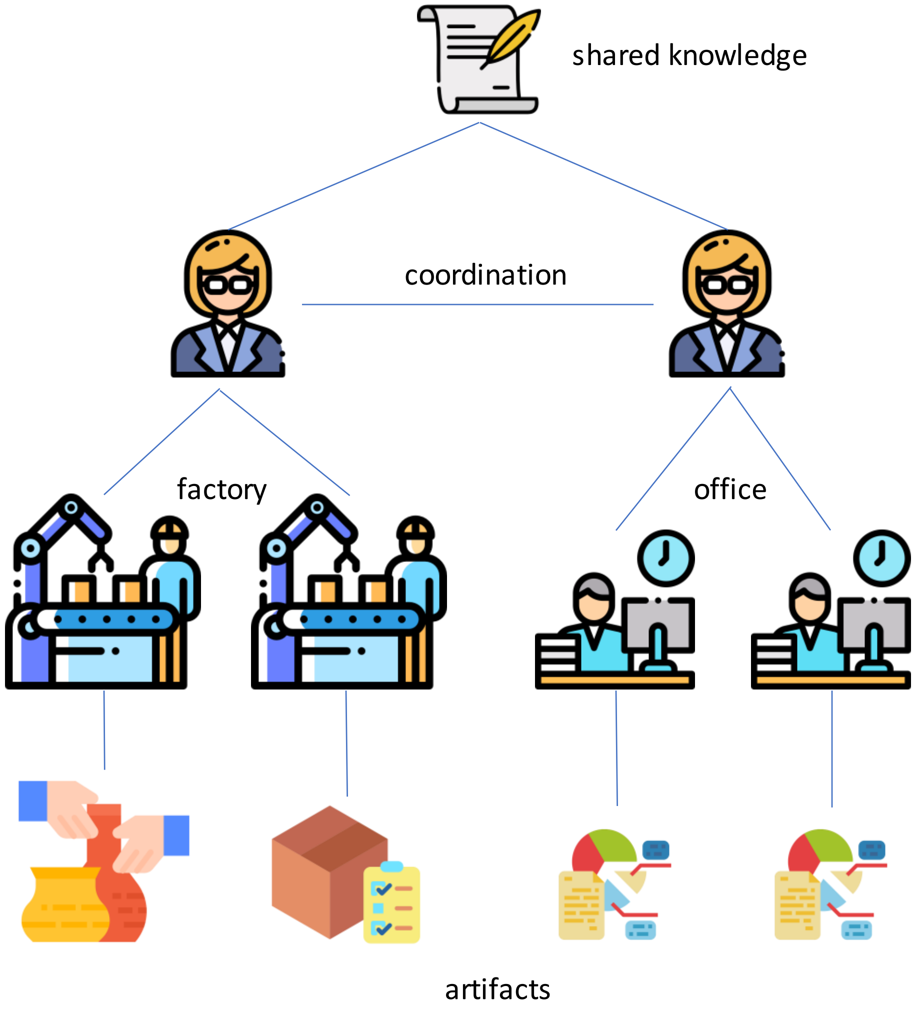

- Shared knowledge
  - **Databases**: read/write knowledge graphs, with access scoped according to need
  - **Ledgers**: signed non-erasable entries
- Supervisors
  - **General purpose human-like AI*** playing different roles at different levels of seniority, analogous to human teams, and working in 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

Icons courtesy of flaticon, freepik, ultimatearm and icongeek

Example for a wine bottling plant

simple control of factory equipment

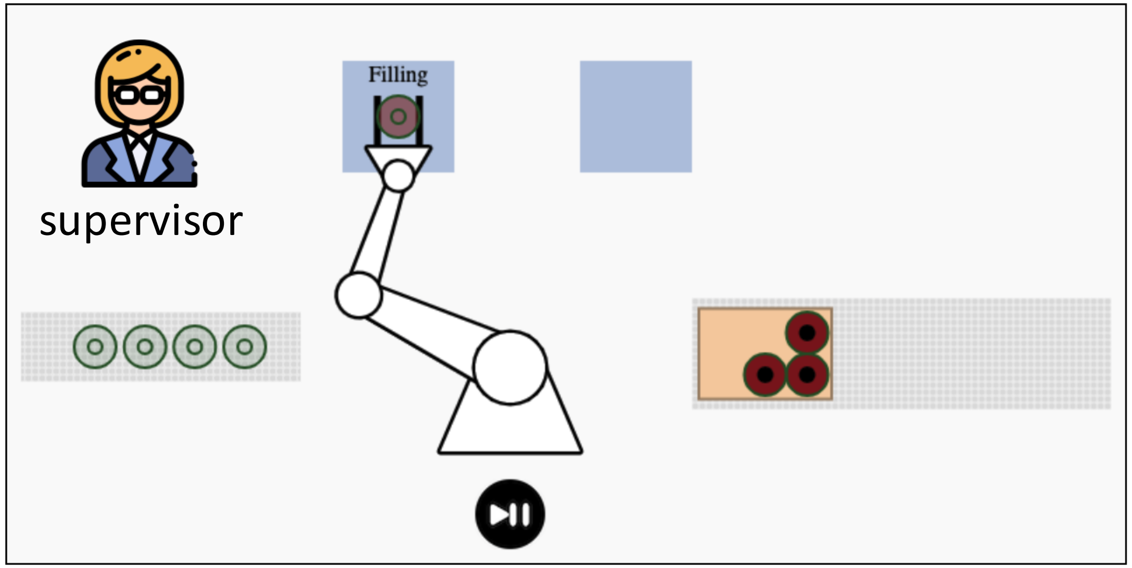

Demo: https://www.w3.org/Data/demos/chunks/robot/

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

- Supervisor coordinates workers
  - Real-time concurrent control
  - 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

Research

Research Challenges

To be worked on over many years

- Modelling cortico-basal ganglia circuit
  - *Sequential rule execution*
- Memory: learning and forgetting
  - *Chunk activation levels*
- Human-like reasoning
  - *Considering examples and analogies*
- Reinforcement learning of rules
  - *Hierarchical learning*
- Indexing for scalability
  - *Facts and rules*
- Semantic integration
  - *Spanning multiple lobes*
- Modelling the Limbic system
  - *How emotions are learned and applied*
- Mimicry as key to human behaviour
  - *Emotional and cognitive communication*
  - *Shared statistics for perception and action*
- Natural language processing
  - *Concurrent incremental processing across all stages in the NLU/NLP pipeline*
  - *Progressive understanding, no backtracking*
  - *Learning from experience*
- Human-like visual perception
  - *Scene & behavioural understanding*
- Modelling cortico-cerebellar circuit
  - *Real-time coordination (Neural ODE)*
  - *Interaction with cognition*
- Theory of mind and social interaction
  - *Human-machine, machine-machine*

Research Priorities

- Cortico-basal ganglia circuit
- Memory: learning and forgetting
- Natural language processing
- Indexing for scalability
- Thematic & taxonomic learning
- Learning natural languages
- Richer forms of reasoning
- Learning everyday skills
- The Limbic system

The rest of this talk focuses on the first three bullet points

Cognition

Cortico-Basal Ganglia Circuit

Cognition – Sequential Rule Engine

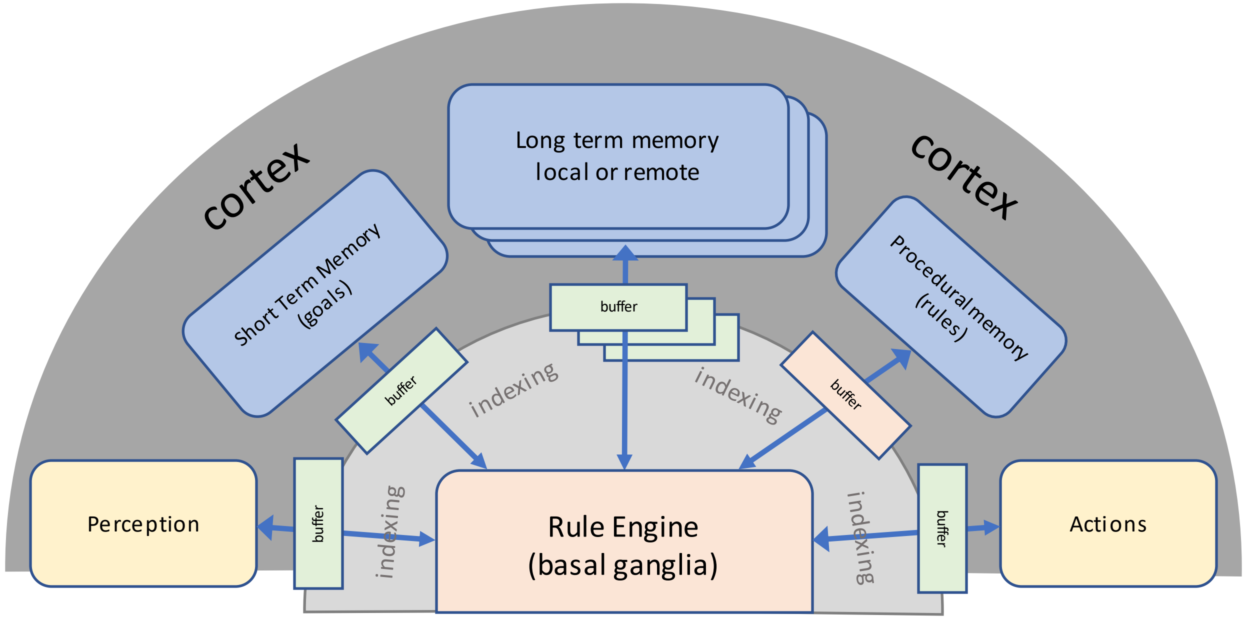

Cognitive Buffers hold single chunk
Analogy with HTTP client-server model

- Inspired by John R. Anderson's ACT-R
  - Semantic Networks + 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
  - Chunks map to *n*-ary relations in RDF
  - 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
  - Formal spec as draft W3C CG Report
- Growing suite of demos
  - counting, decision trees, industrial robots, smart homes, memory, natural language, self-driving cars, browser sandbox, chunks test suite, open source chunks library

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
  - Variables pass information from conditions to actions
  - 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, fill a bottle, switch on a light, or to say "hello"
- Rule chunks use @module to name the module chunk 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}
```

* Deliberate use of HTTP methods to strengthen analogy with HTTP

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}
```

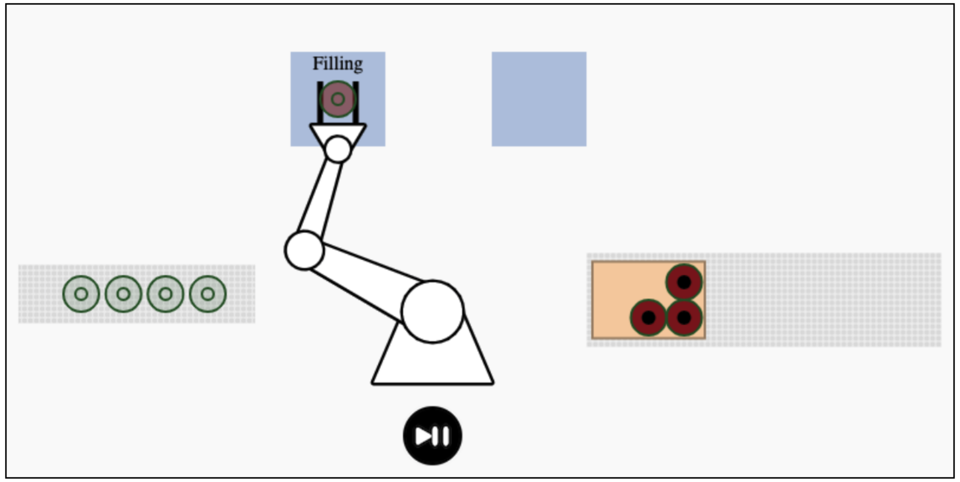

Log:

```
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
```

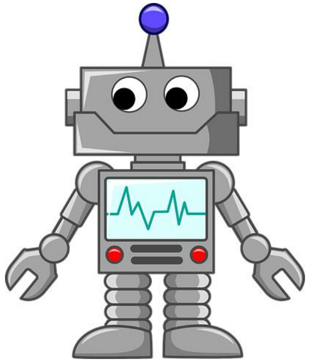
Courtesy of pixabay

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

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

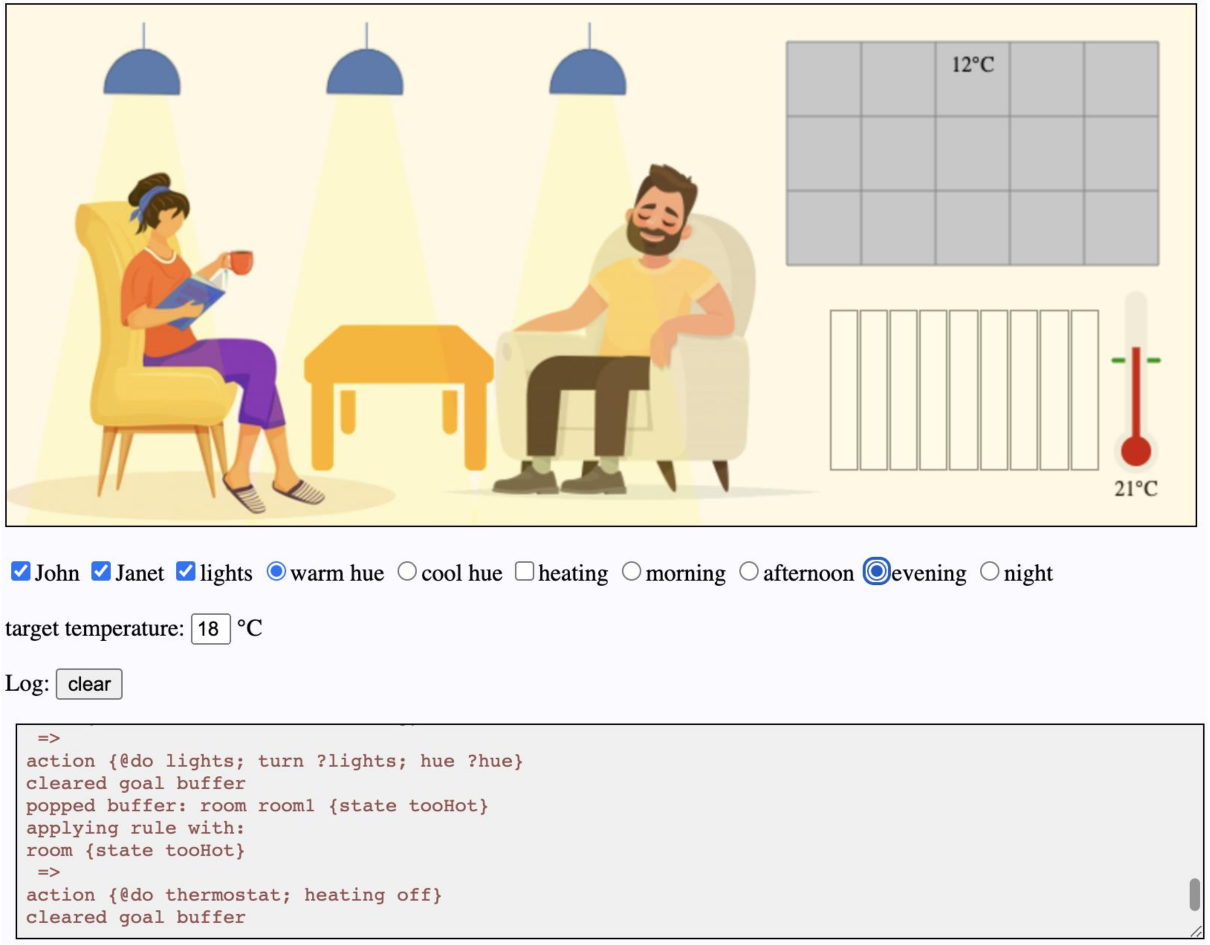

Other Web-based Demos

- 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

Memory

Memory Tests

- Study a set of items, then, sometime later, see how many you can remember
- The functional model of the cortico-basal ganglia circuit involves buffers that hold a single chunk
- To model the memory test we can
  a) either store successor relationships for the sequence of items, e.g. *apple* - *flower*, or
  b) we can iterate over matching chunks
- The second approach is more robust to memory failures
- How does this relate to interference effects?

Memory

- Forgetting as intrinsic memory decay or as interference from other memories, or some combination of the two?
- Exponential decay over time provides good fit to lab data
- But can also be ascribed to interference from new memories
- Priming effect on related memories as spreading activation
- We can recall memories from many years ago given the right cues
- New memories lack strong evidence for their lasting value
- Such evidence has to be acquired with experience
- What's the most effective model for all of these points?

Underwood (1957) showed that memory loss is largely attributable to interference with other memories. Memories can thus be recalled after an interval of many years provided that the interference is small. This reflects experience in selecting memories that have been more valuable.

For ACT-R, the decay of activation is only one component of the activation equation. There is also a context component to activation which works to increase the activation of items based on the current context. Thus, even chunks which have decayed significantly over time can have activations above the threshold if they are strongly related to the current context.

Learning and Forgetting

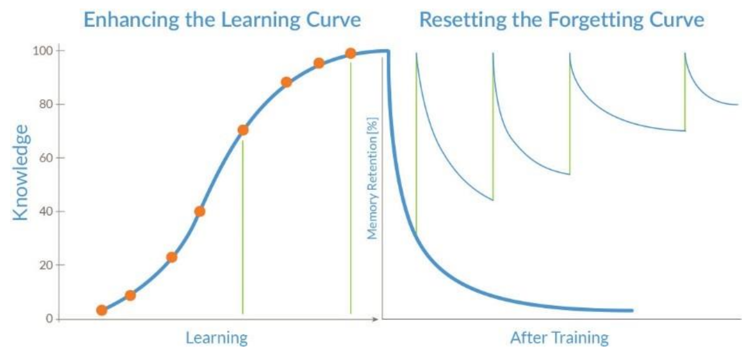

- Experiments that model how human memory works; how practice overcomes forgetfulness; replicating the forgetting curve and spacing effect
- You might ask why we would want to make computers forgetful, given how faithful they are at remembering
- The answer is that in everyday situations you want to recall just what is most important based upon past experience
- This is similar to web search engines which seek to provide the results that are likely to be most relevant given the words given in the search query

- Chunks have parameters for an activation level and a timestamp
- Activation decays over time like a leaky capacity losing its charge
- Recalling or updating a chunk boosts its activation level
- Boost is weaker for closely spaced rehearsals – *aka* the spacing effect*
- Decaying wave spreads through linked chunks to boost related concepts
- Stochastic recall - chunks with higher activation levels are more likely to be recalled, but sometimes weaker chunks are recalled in place of stronger chunks

* Spacing effect modelled with the Logistic function, whilst stochastic recall uses Gaussian noise

Spreading Activation

- Why is it easier to remember items in a group for groups with fewer items?
- A wave of spreading activation provides one possible explanation
- Activation of one item in the group spreads to other items in the same group following property links in both directions
- The amount of wave activation for each item is inversely related to the number of items in the group
- What is the underlying computational model for pulsed neural networks?

```
# items belonging to group animals
item {word dog; group animals}
item {word horse; group animals}
item {word cat; group animals}
```

- Remembering the item for *dog* boosts the chunk for the group (*animals*) and spreads out to boost the other items in that group
- Does this depend on the property (in this case *group*) being the same?
- How can we implement this efficiently on conventional computers?

Potential for human experiments to differentiate different models of spreading activation

Memory Experiments

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

- Grouping items into clusters makes them easier to recall
  - Random words that have no obvious relationship
  - Words that can be easily clustered
- Use of cognitive rules and graph algorithms to form clusters
- Thematic vs taxonomic relations
  - *dog – collar* vs *dog – bear*
- Thematic relations are based on co-occurrence in particular situations
- Taxonomic relations are abstractions over a history of situations (episodic memories)
- Future work on indexing to ease learning, to exploit analogies, metacognition and hierarchical reinforcement learning

Testing Data

Unrelated words

Nine	Swap	Cell	Ring	Lust
Plugs	Lamp	Apple	Table	Sway
Army	Bank	Fire	Hold	Worm
Clock	Horse	Color	Baby	Sword
Desk	Grab	Find	Bird	Rock

Taxonomically related

Horse	Cat	Dog	Fish	Bird
Orange	Yellow	Blue	Green	Black
Table	Chair	Desk	Bookcase	Bed
Teacher	School	Student	Homework	Class
Apple	Banana	Kiwi	Grape	Mango

Taxonomically and Thematically related

Animals	Food	Riding
Dog	Bone	Leash
Horse	Apple	Saddle
Mouse	Cheese	Trap

Natural Language


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

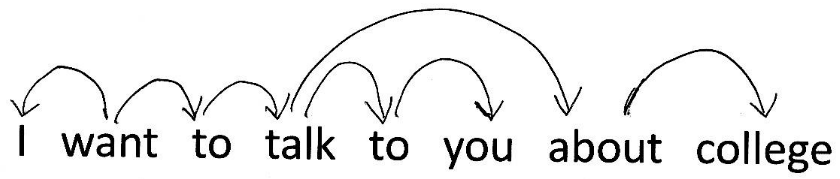


```
# I want to talk to you about college
vp _:1 {verb want; subject _:2; to _:3}
np _:2 {pron i}
np _:3 {noun talk; to _:4; about _:5}
np _:4 {pron you}
np _:5 {noun college}
```

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

I am working on a demo of **text → meaning → text** to show how this can work in practice for end to end communication of meaning

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

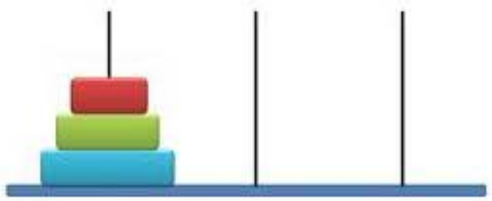

```
# move the red disc to the right peg
verb v1 {word move; subject p1; to p2}
phrase p1 {word disc; det the; adj red}
phrase p2 {word peg; det the; adj right}

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

Dinner Dialogue

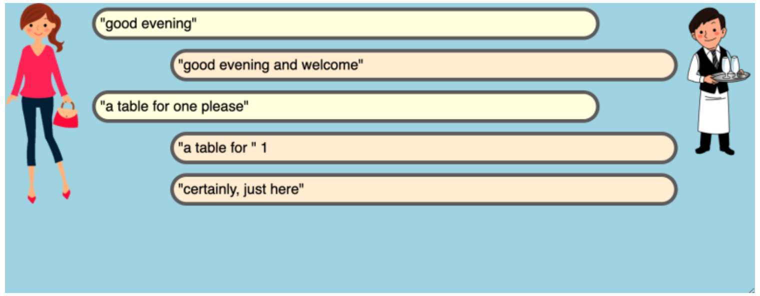

W3C Cognitive AI Community Group

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

Collaboration

Collaboration

- I am interested in a working partnership
  - Offering mutual benefit that can continue after I retire from W3C
- Mentoring given that I don’t have formal background in psychology
  - To help with design of experiments and outreach in scientific publications
- Exciting opportunities to break new ground for human-like processing of natural language
  - Modelling how children learn as a basis for multilingual agents and for escaping the straitjacket of manual knowledge engineering
- Some initial guidance on designing demos on modelling forgetting
- Opportunities for joint participation in Horizon Europe and UK funding calls

HORIZON-CL4-2021-HUMAN-01-03

European Network of AI Excellence Centres: Pillars of the European AI lighthouse (RIA)

- Research and Innovation Actions with an indicative budget of 9 million euros for the topic
- Start at TRL 2-3 and achieve TRL 4-5 by the end of the project
- Establishing a new pillar of the European AI lighthouse
- Reinforcing a leading unified European AI community
- Scientific progress in AI, addressing major challenges hampering its deployment

- Topics include, but not limited to
  - Technical robustness and safety, including methods for evaluating the resilience of systems, and standardized ways of quantifying robustness of AI
  - Privacy preserving techniques and infrastructures
  - Human agency and oversight in terms of system security and safety; including explainability in human-readable terms allowing to detect/prevent/mitigate/recover from harm and threats
- **Perfect fit for human-like AI**

Need to seek out potential partners, Digital innovation hubs and AI start-up initiatives