



Cognitive AI

Dave Raggett <dsr@w3.org>

25 November 2020

Webinar on behalf of the

[**Centre for Artificial Intelligence, Robotics and Human-Machine Systems \(IROHMS\)**](#), Cardiff University

[**IROHMS Research themes**](#): Human-like AI, Ethical and explainable AI,
Human-centred technologies and society, Humans and robots

The Cambrian Explosion




Left credit: [Smithsonian Institution](#)


The Cambrian period occurred over 500 million years ago and included the biggest evolutionary explosion in Earth's history linked to higher concentrations of oxygen and warmer seas.

This included predators with the emergence of nervous systems supporting complex behaviour, and followed on from the emergence of multicellular organisms in the Ediacaran period, immediately after "Snowball" Earth 635 million years ago.


Bottom: 520 million year old fossilised ventral nerve cord of [*Chengjiangocaris kunmingensis*](#).

Credit: Top: Jie Yang, Bottom: Yu Liu






The cortex is highly convoluted to maximise its area in a limited volume – the human cortex has an area of roughly 2400 square centimetres



Cross Section of the Brain


The cortex on the outside, other systems on the inside



Courtesy of healthpages.org

The human cortex occupies about 76% of the brain's volume.

The outer part is the neocortex, some 2-4 mm thick and containing six horizontal layers, arranged into vertical cortical columns of approximately 0.5 mm diameter.



Cognitive AI and the Sentient Web

- **Cognitive AI aims to mimic human memory, reasoning and learning**, inspired by advances in the cognitive sciences and over 500 million years of neural evolution
- We can mimic the brain at a functional level using conventional computer technology rather than explicitly using artificial neurons
- Cognitive agents at the network edge or in the cloud
- There are many potential applications of cognitive agents for human-machine collaboration
- The **Sentient Web** as the web of interconnected cognitive agents featuring perception, cognition and action – subsuming both the IoT and Semantic Web
 - *Perception*: sensors and dynamic access to information services
 - *Cognition*: data processing, reasoning and learning
 - *Action*: control of actuators and responses to queries
 - *Sentient*: self awareness in respect to current state, goals and actions
- The future for Web search: intelligent, collaborative, value added services


Cognitive AI – Long Term Mission

is to enable smart cognitive agents that:

- Are knowledgeable, general purpose, collaborative, empathic and trustworthy
- Can apply metacognition and past experience to reason about new situations
- Support continuous learning based upon curiosity about the unexpected
- Have a level of self awareness in respect to current state, goals and actions
- Have an awareness of others in respect to their beliefs, desires and intents
- Are multilingual and can interact with people using their own language



A grand mission to boldly explore the universe of the mind



Modelling the Cortex with Cognitive Databases

- The human cortex is functionally equivalent to a set of specialised cognitive databases and associated algorithms
- A **cognitive database holds chunks**: collections of properties with values as literals or references to other chunks
- Chunks are associated with statistical information reflecting prior knowledge and past experience
- Cognitive databases have the potential to store vast amounts of information similar to the human cortex
- **Cognitive databases can be local or remote**, and **shared** with multiple cognitive agents, subject to **access control** policies


- Memory retrieval fits Web architecture
 - Remote invocation of graph algorithms in request/response pattern that sits nicely on top of HTTP
 - Analogous to Web search engines where results are computed based upon what is likely to be most relevant to the user – impractical and inappropriate to try to return complete set of matches
- Cognitive databases support a variety of algorithms that are executed local to the data
 - Scalable to efficiently handling Big Data
 - Distributed algorithms across multiple cognitive databases to mimic operation of the anterior temporal lobe*
- The algorithms depend on the intended function of the database, e.g.
 - Basic storage and recall
 - Specialised algorithms for natural language, spatial and temporal reasoning
 - Algorithms for data analytics



* [Sharon Thompson-Schill](#) on a hub and spoke model for how the anterior temporal lobe integrates unimodal information from different cortical regions.
Related to Numenta's "[thousand minds theory of intelligence](#)" involving non-hierarchical long range connections across cortical regions.

Cognitive AI Architecture with multiple cognitive circuits

Cortex is analogous to blackboard system



Perception


- Our senses
 - Smell, taste, touch, pain, heat, sound, vision, ...
 - **Perception creates short lived representations in the cortex**
 - The cortex can likewise direct sensory processing as needed
- Touch and pain are mapped to a homuncular model of our bodies
- Proprioception – sense of self-movement and body position
 - Limbs, joints, muscle load
 - Vestibular system (inner ear)
- Sound is fleeting
 - Processing word by word
 - Prosody & attentional cues
- Vision is much more complex
 - Two eyes for stereo depth perception
 - Each eye: high resolution narrow angle + low resolution wide angle
 - Saccades as eyes swivel to scan areas of interest
 - Good at recognizing many different kinds of things, including their structures & behaviours
 - Cognitive **context** determines what is interesting and relevant
 - **Alerts** signal relevant things in field of view (i.e. act as events)
 - Cognitive **focus** directs attention to specific things
 - Reinforcement learning from experience with neuro-evolution



Visual system



Hearing




Cortical homunculus



Cognition and conscious thought

- Symbolic (graphs) + sub-symbolic (statistics)
 - Chunk based symbolic representation of concepts and relationships
- Rule engine connected to many parts of the cortex
 - Rules represent reasoning & procedural knowledge
 - Learned from experience (hierarchical reinforcement learning)
- Sequential application of rules to cognitive buffers
 - Every 50 mS or longer
- Parallel processes for graph algorithms
 - Recall of memories
 - Selection of rules
- Autobiographical and episodic memories
- Reasoning at multiple levels of abstraction

Basal ganglia as a sequential rule engine
*and nature's solution to the scaling problem**



Each buffer holds a single chunk and represents the current state of a bundle of nerve fibres connecting to a particular cortical region. The rule conditions and actions operate over these buffers. Moreover, the buffers can be likened to HTTP clients, where the cortex is like a set of HTTP servers. This architecture originates in John Anderson's work on [ACT-R](#).

Chunks: a collection of properties that include references to other chunks, and statistical weights reflecting prior knowledge and past experience

Modules: specialised graph databases and algorithms, accessed via buffers that hold a single chunk

Rules: conditions ranging over module chunk buffers, and actions that either update the buffers or invoke graph algorithms

* High performance will entail emulating the Striatum, Pallidum and Thalamus, akin to the Rete algorithm



Action

- **Cortico cerebellar circuit**
- Handles actions devolved to it by conscious thought
- Real-time control with parallel processing
- Cerebellum's dense circuitry contains more than three times the number of neurons in the cortex*
- Cerebellum acts as a flight controller managing activation of myriad sets of muscles in coordination with perceptual input from the cortex
- Offloads processing from cortico basal-ganglia circuit thereby enabling higher level thought whilst actions are underway

- Performance degrades when conscious thought diverts visual attention, starving cerebellum of visual feedback
- Monitoring and signalling to cognition when higher level control is needed, e.g. to re-plan movement around obstacle
- Learning through experience, starting with conscious thought
- Implemented as suite of real-time continuous state machines
 - Exploiting Neural Ordinary Differential Equations ([Neural ODE](#)), Chen et al. 2018
- Examples: talking, walking, cycling and playing the piano

* The human cerebellum contains 70 billion nerves vs 20 billion for the cerebral cortex, see [Suzana Herculano-Houzel](#), 2010



Emotions, Feelings and Moods

Towards strong empathic AI*

- **Cortico-Limbic system**
- Important from an evolutionary perspective
 - Avoidance of harm, fear of predators, interest in prey, courtship, care of young
- Enhanced for living in social groups
 - Emotional intelligence – awareness of what others are feeling, and signalling your own feelings
- Emotions are associated with a feeling and something they apply to
 - Valence describes whether feeling is positive, neutral or negative
 - Arousal describes whether feeling is calming or exciting
 - Moods are long lasting emotions that lack the cognitive element
- Triggered by
 - Perception (e.g. seeing a predator), reasoning about situations, recall of emotive memories
 - Functional implementation as a feed-forward classification network
- Effects
 - Instinctive behaviours and how these are regulated by cognitive control
 - Prioritising what you are thinking about and what feels important
 - Influences on recall, new memories, reinforcement of existing memories and reinforcement learning of behaviours
- Fast and instinctive vs slow and deliberate
 - Rapid instinctive appraisal and response, avoiding the delay incurred with conscious thought, but subject to errors of judgement due to lack of considered thought



* empathic: /ɛm'paðɪk/ adjective – showing an ability to understand and share the feelings of another

Many Ways to Think

- Many forms of thinking have to deal with uncertainties, e.g.
 - Induction: building models to explain regularities
 - Abduction: determining the most likely cause of some observations
 - Causal reasoning about plans for realising goals
 - Fuzzy reasoning involving blends of different states
- Mimicking human memory
 - In any large knowledgebase we only want to recall what is relevant to the current situation judging from past experience
 - Interference effect – new memories may block older memories
 - Ebbinghaus forgetting curve – our ability to recall information drops off over time unless boosted by repetition
 - Closely spaced repetitions have less effect – reduced novelty
 - Longer gaps means less salient features are forgotten
 - Spreading activation – concepts are easier to recall on account of their relationship with other concepts*
 - Stochastic selection of facts and rules, and the role of gaussian noise for recall of memories with similar strengths
 - Multiple implementations possible: chunk graphs (Anderson), holographic memory (Kelly), pulsed neural networks (Eliasmith)

Critical Thinking:


- Reasoning
- Analysing
- Evaluating
- Decision Making
- Problem Solving



Recommended reading:


- [Large-Scale Model of the Functioning Brain](#)
- [Holographic Declarative Memory](#)
- [Neural Network Model of Memory Retrieval](#)

* Nature figured out the “Page Rank” algorithm many millions of years ago!



Machine Learning

- Manual development of knowledge won't scale cost effectively
- We therefore need to rely on machine learning for declarative and procedural knowledge
 - Many existing algorithms to take advantage of
- Prior knowledge enables learning from small datasets
- Semi-supervised learning as human guided exploration with attention to salience
- Learning at multiple levels of abstraction
 - Case based reasoning to speed learning
 - Looking for meaningful explanations
- Active learning – continuous, surprise driven
 - Mimicking humans as prediction machines – we attend to novelty to improve our predictions
- Learning from experience with real or simulated environments and multiple cognitive agents
 - *Playground*
- Use with natural language for teaching skills to cognitive agents as solution to scaling problem
 - *Class room*



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 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 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




Nature vs Nurture

inborn ability and traits versus what is learned



- How much and what kinds of knowledge are needed to bootstrap learning?
 - Baseline of manually programmed facts, rules and graph algorithms including heuristics, etc.
- Three stages of learning and forgetting
 - Progression from declarative to procedural
 - Chunks as basis for facts and rules
 - Heuristics and metacognition
- Emotional control & cognitive regulation
 - Explored in Minsky's "The Emotion Machine"
- Recall of similar situations or instances using efficient graph algorithms
 - Simon (1992) Intuitive decision making as nothing more and nothing less than recognition
 - Klein (1986) Fire marshals tended to make rapid decisions by generating a single alternative, mentally simulating its outcome, and either making minor revisions or adapting the nearest closest alternative



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.

Contexts for episodic memory and statements about statements

- Beliefs, stories, reported speech, examples in lessons, abductive reasoning and even search query patterns involve the use of statements about statements. How can these be expressed as chunks and what else is needed?
- Here is an example borrowed from John Sowa's [Architectures for Intelligent Systems](#):
- Contexts are also useful for episodic memory when you want to describe facts that are just true in a given situation
- E.g. when visiting a restaurant for lunch, you sat by the window, you had soup for starters followed by mushroom risotto for the main course. A sequence of episodes can be modelled as relationships between contexts, and used for inductive learning.

Tom believes that Mary wants to marry a sailor
believes s1 {@subject tom; proposition s2}
wants s3 {@context s2; person mary; situation s4}
married-to s5 {@context s4; @subject mary; @object s6}
a s6 {@context s4; isa person; profession sailor}

In other words: Tom believes the proposition that Mary wants the situation in which she is married to someone who is a sailor by profession

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: °C

Log:

```
=>
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 (it may have already happened, or you may need to wait until it does)
 - 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

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, inspired by work on Cognitive Linguistics
- Incremental word by word concurrent syntactic and semantic processing without the need for any backtracking
 - Reducing load on limited working memory
- Use of statistical information to guide choices, e.g. for priming effect on word senses
 - Disambiguation on behalf of parser
- Offloading cognition with graph algorithms
 - Freeing up the cortico-basal ganglia circuit, and escaping the constraints it imposes
- 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
- Syntactic word classes (part of speech) allow utterances to be segmented
- These segments can be arranged in a properly nested spanning tree
- The choice of attachment depends on syntactic and semantic considerations which therefore need to be processed concurrently to avoid backtracking
- Semantic agreement overrides adjacency for the attachment of prepositional phrases
- Phrase structure as data stream for understanding meaning in terms of semantics and pragmatics
- Semantic mapping rules are designed or learned for the structures generated by the phrasal analysis
- Used for understanding and generation, with shared statistics




```
# 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}
```

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
- Parsing demo tests that parser is adequate for all dinner dialogue utterances
 - <https://www.w3.org/Data/demos/chunks/nlp/parsing/>
- Ongoing work on syntax-semantics mapping rules
 - And their use for both understanding and generation
 - Further work planned on learning mapping rules and meanings of previously unseen words


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


Relationship to other approaches to AI

- Deep Learning is inspired by models of neurons with back propagation algorithm for training across multiple layers
- Very successful, but increasingly obvious limitations
 - Lacks transparency of how it reaches its conclusions
 - Lacks salience, is brittle and readily fooled
 - Requires huge amounts of training data & compute power
 - Not general purpose, difficulties for transfer learning
- Yoshua Bengio, a Deep Learning pioneer, acknowledges this:

"We have machines that learn in a very narrow way," Bengio said in his keynote at NeurIPS in December 2019. "They need much more data to learn a task than human examples of intelligence, and they still make stupid mistakes."


Without question, deep learning is an imperfect model of intelligence. It cannot reason abstractly, does not understand causation and struggles with out-of-distribution generalization.

- Semantic Web is founded on Aristotelian tradition of logic, formal semantics and model theory
 - Predicate calculus and first order logic
- *What is provably true given assumptions and inference rules; intolerant of inconsistencies*
- By contrast Cognitive Science is the experimental study of the organizing principles of the mind
- Cognitive AI – mimicking human memory and reasoning with combination of graphs, statistics, rules and graph algorithms, inspired by evolution
- *What is useful based on prior knowledge and past experience in the face of uncertainty and inconsistency*



Aristotle Newell Anderson Minsky Johnson-Laird*


* Humans don't rely on the laws of logic and probability, but rather by thinking about what is possible



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

perception, cognition, and action federated across the Web

Cognitive AI – mimicking human memory, reasoning and learning




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
 - Complements Deep Learning
- W3C Cognitive AI Community Group
 - <https://www.w3.org/community/cogai/>
- Chunks as collection of properties referencing 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
 - JavaScript chunks library

Cognitive Architecture with multiple cognitive circuits loosely equivalent to shared blackboard




Cognition – Sequential Rule Engine



- 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

Natural Language Generation as 3 stage pipeline





Thanks for listening!