

Emergence of the Sentient Web and the revolutionary impact of Cognitive AI

Dave Raggett <dsr@w3.org>

08 June 2020

ISO/TC 211 50th Plenary meeting
WG4 Geospatial Services

Where we are now ...

- RDF is a successful framework for representing data and metadata
 - Large suite of standards including SPARQL, OWL, SHACL, Turtle, JSON-LD, ...
 - JSON-LD 1.1 is currently a W3C Proposed Recommendation
 - Data Catalog Vocabulary (DCAT) version 2, W3C Recommendation (Feb 2020)
 - Web of Things W3C Recommendation for abstraction layer for digital twins (April 2020)
 - Plenty of work on ontologies and wide deployment of *schema.org* vocabs
 - W3C/OGC standards for SSN/SOSA, ETSI SAREF suite for IoT
- Growing industry interest in Labelled Property Graphs (LPG)
 - Easier to work with n-ary relationships
 - No need for reification to annotate graph edges
- LPG is weak on interoperability across different vendors
 - ISO work on LPG extensions to SQL and new work item on GQL query language
- Easier RDF initiative seeking to make semantic technologies based upon RDF easier for the average developer (the middle 33%)
 - Desire for simpler solution for graphs and rules, including lists as arrays

Chunks as an amalgam of RDF and LPG

Inspired by work in Cognitive Psychology and Neuroscience

- Chunks as collection of properties whose values name other chunks
 - Values are names, numbers, true, false, quoted string literals, ISO8601 dates, or comma separated lists of values
 - Context chains for handling multiple perspectives
 - Simple mapping to RDF for integration with existing systems
- Simple syntax – simpler than JSON-LD

```
friend f34 {  
  name Joan  
}  
friend {  
  name Jenny  
  likes f34  
}
```

```
dog kindof mammal  
cat kindof mammal
```

is equivalent to

```
kindof {  
  subject dog  
  object mammal  
}  
kindof {  
  subject cat  
  object mammal  
}
```

- Where *friend* is a chunk type, *f34* is a chunk identifier, *name* and *likes* are property names, *Joan* and *Jenny* are also names.
- *likes f34* signifies that Jenny likes Joan via the link to the chunk for Joan.
- Missing chunk identifiers are automatically assigned when inserting a chunk into a graph
- Uses line break or semicolon as punctuation

Chunks correspond to activation across bundles of nerve fibres, see Chris Eliasmith's work on [semantic pointers](#)

Blending Symbols with Statistics

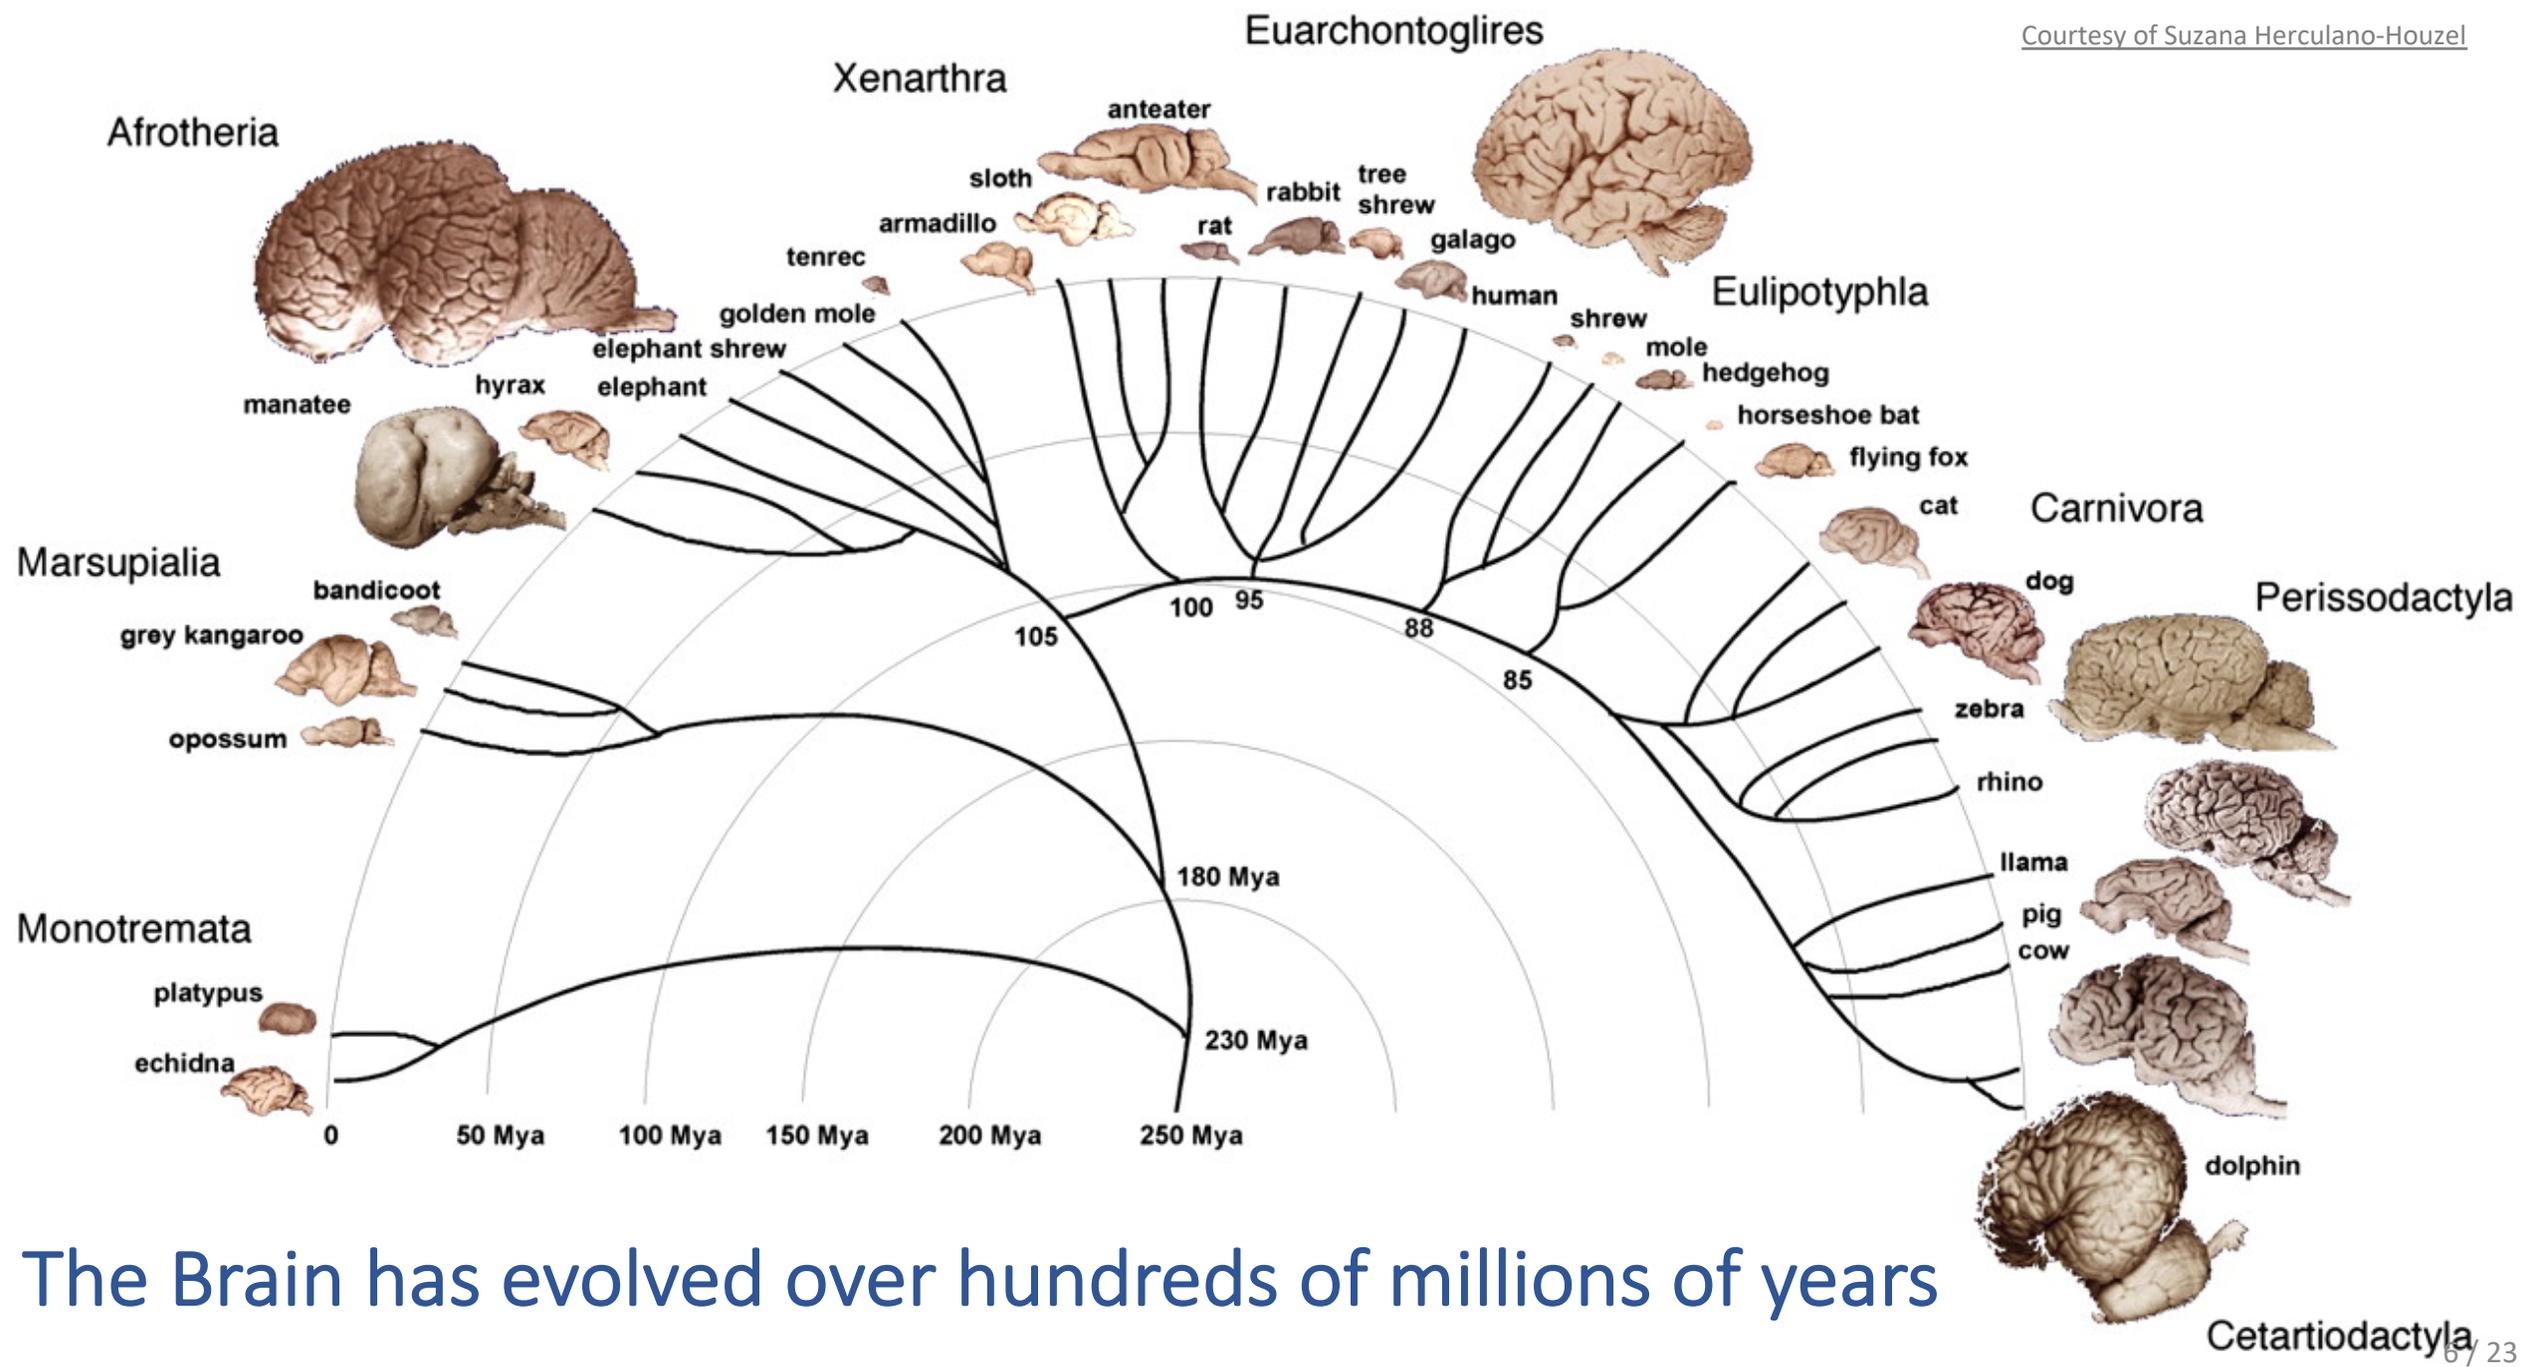


- Traditional approaches to handling data struggle in respect to the uncertainty, incompleteness and inconsistency commonly found in real-world situations
 - This exacerbates the cost for preparing and cleaning data prior to analysis, a major bugbear for data science
- Remembering what's important based upon prior knowledge and past experience
 - Data recall is like web search engines that determine which matches are most likely to be useful as distinct to all the rest
- Machine learning with relatively few examples, just like humans!
 - Unlike *Deep Learning* which starts from scratch, requiring huge numbers of training examples, and lacks salience, making it brittle and easy to fool
- Forms of reasoning that rely on statistical considerations
 - e.g. abduction which seeks explanations of observed behaviours
- Graph manipulation for operational semantics rather than formal semantics and logical proof
- Relational databases are now giving way to graph databases, and will in turn give way to cognitive databases that combine graph data, statistics, rules and graph algorithms
- **Sentient Web**: the combination of the IoT and Cognition to enable ecosystems of smart services
 - Sensing + reasoning federated across the Web

Cognitive AI



- In short, Artificial Intelligence inspired by advances in the cognitive sciences
- In other words, we would do well to borrow from nature when it comes to building AI systems
- We can mimic nature at a functional level using conventional computer technology without having to implement cognitive agents in terms of artificial neurons
- There are many potential applications of cognitive agents for human-machine collaboration

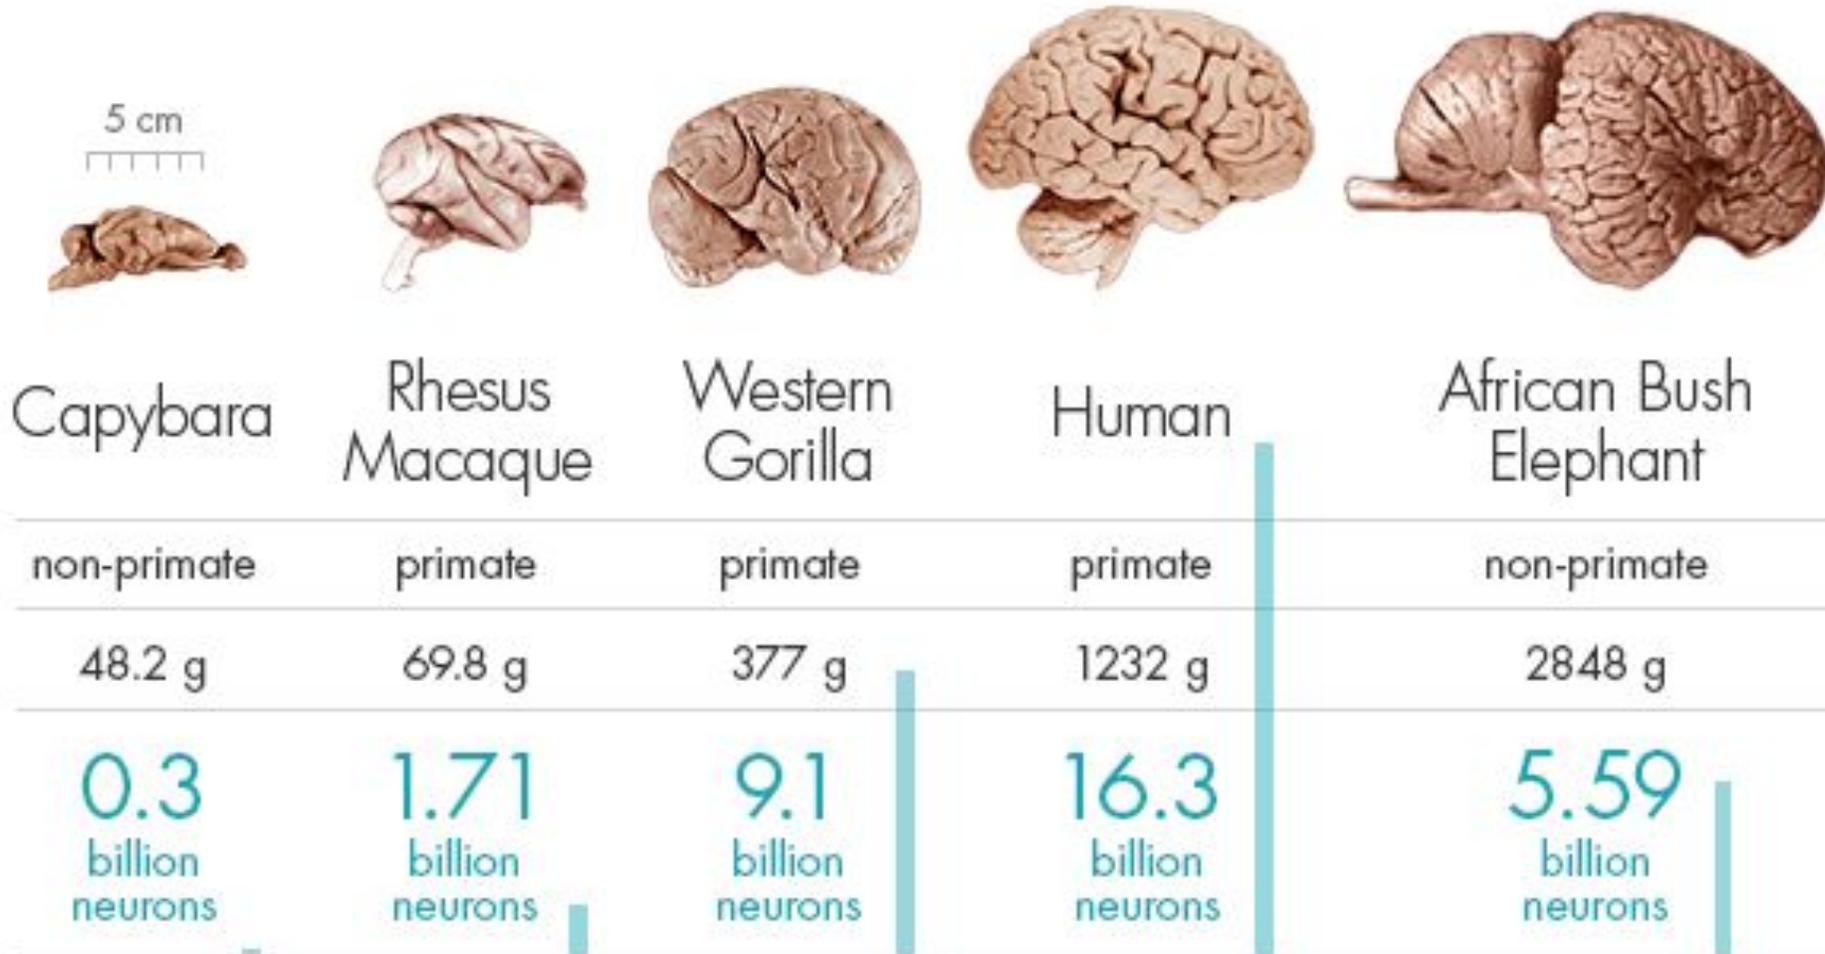


The Brain has evolved over hundreds of millions of years

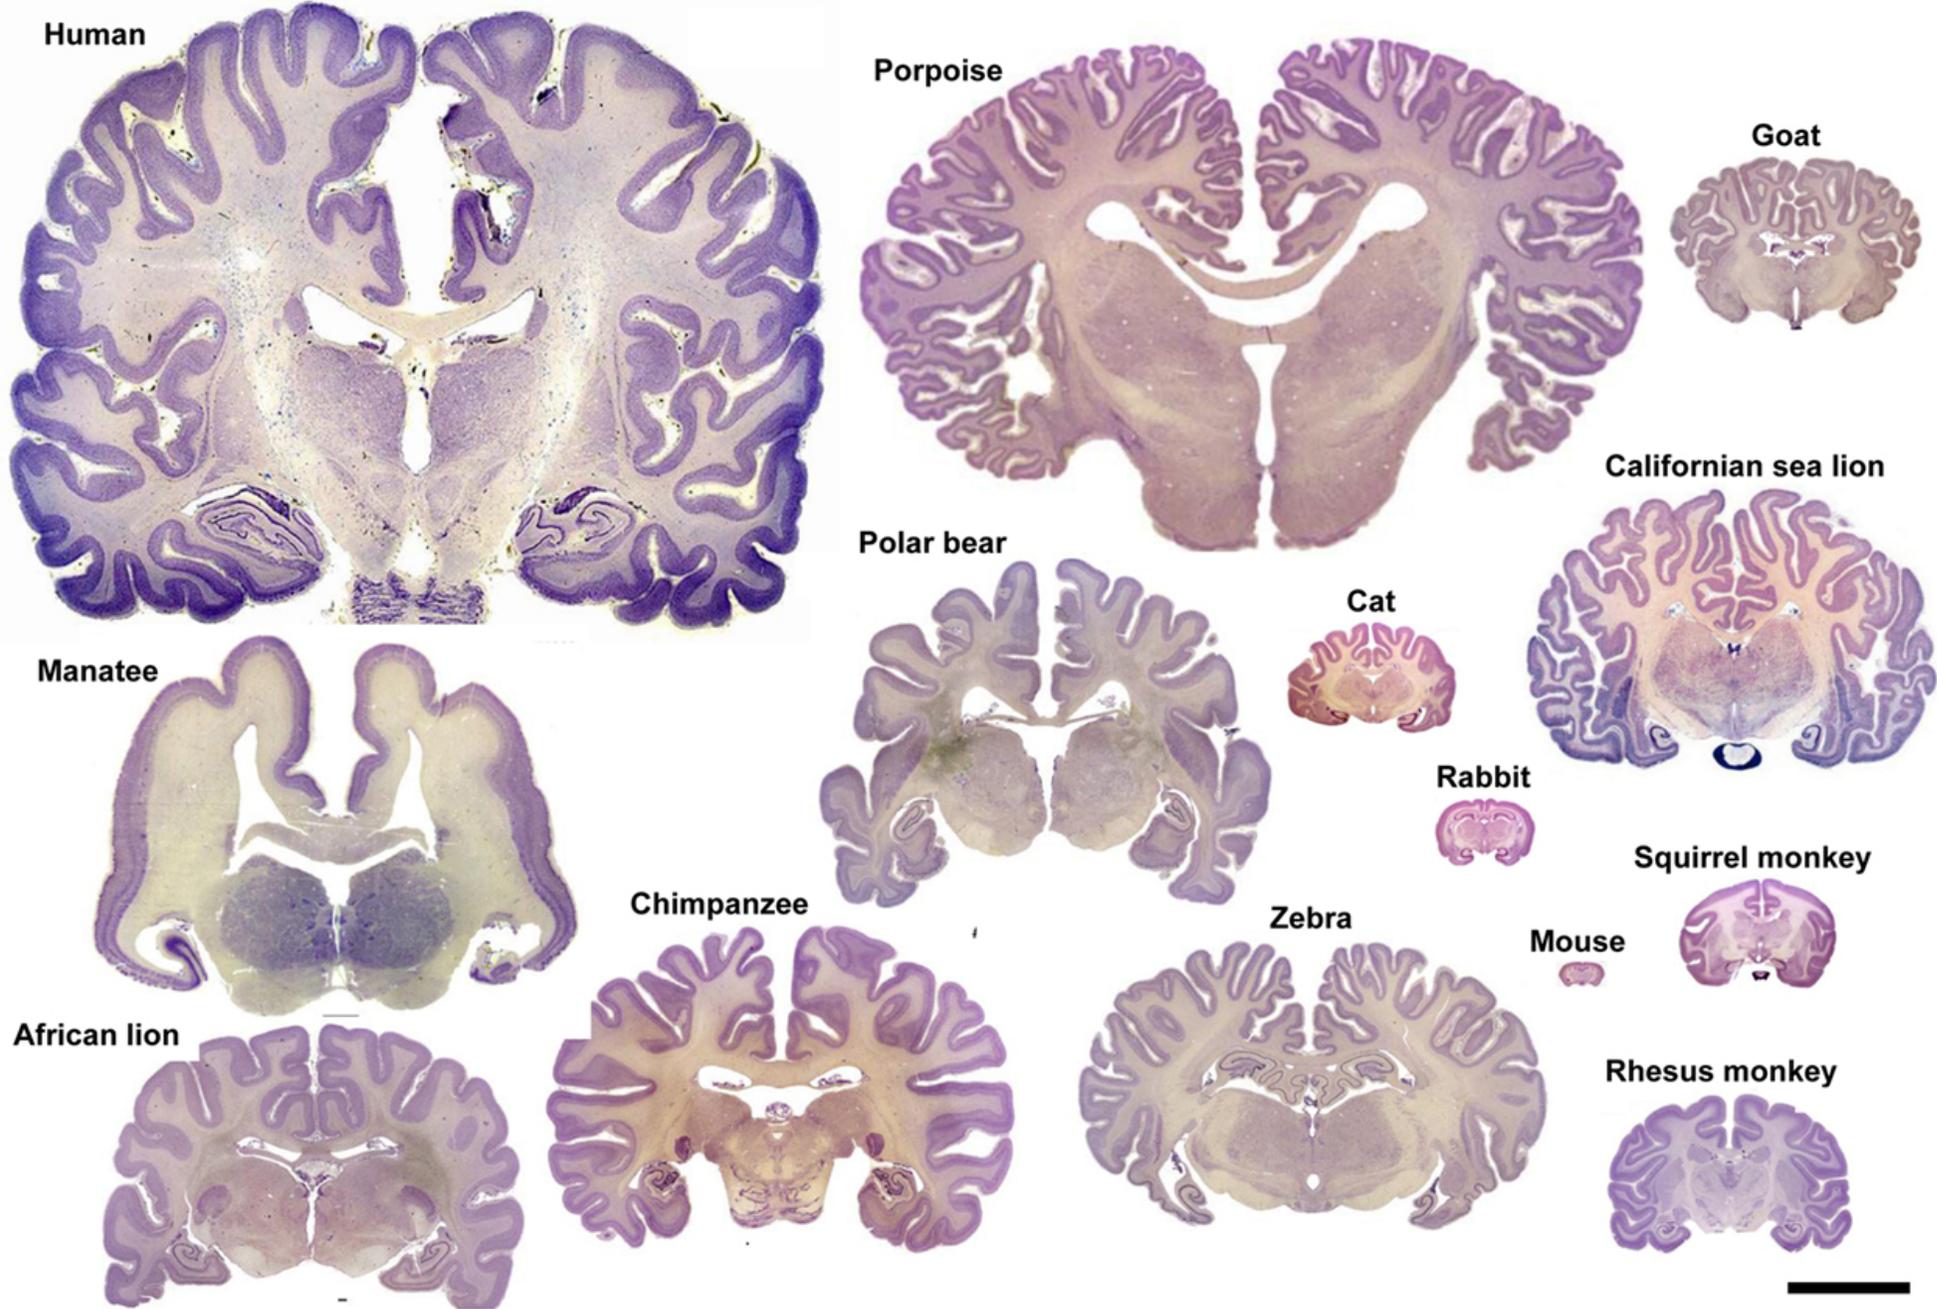
BRAIN SIZE AND NEURON COUNT

Cerebral cortex mass and neuron count for various mammals.

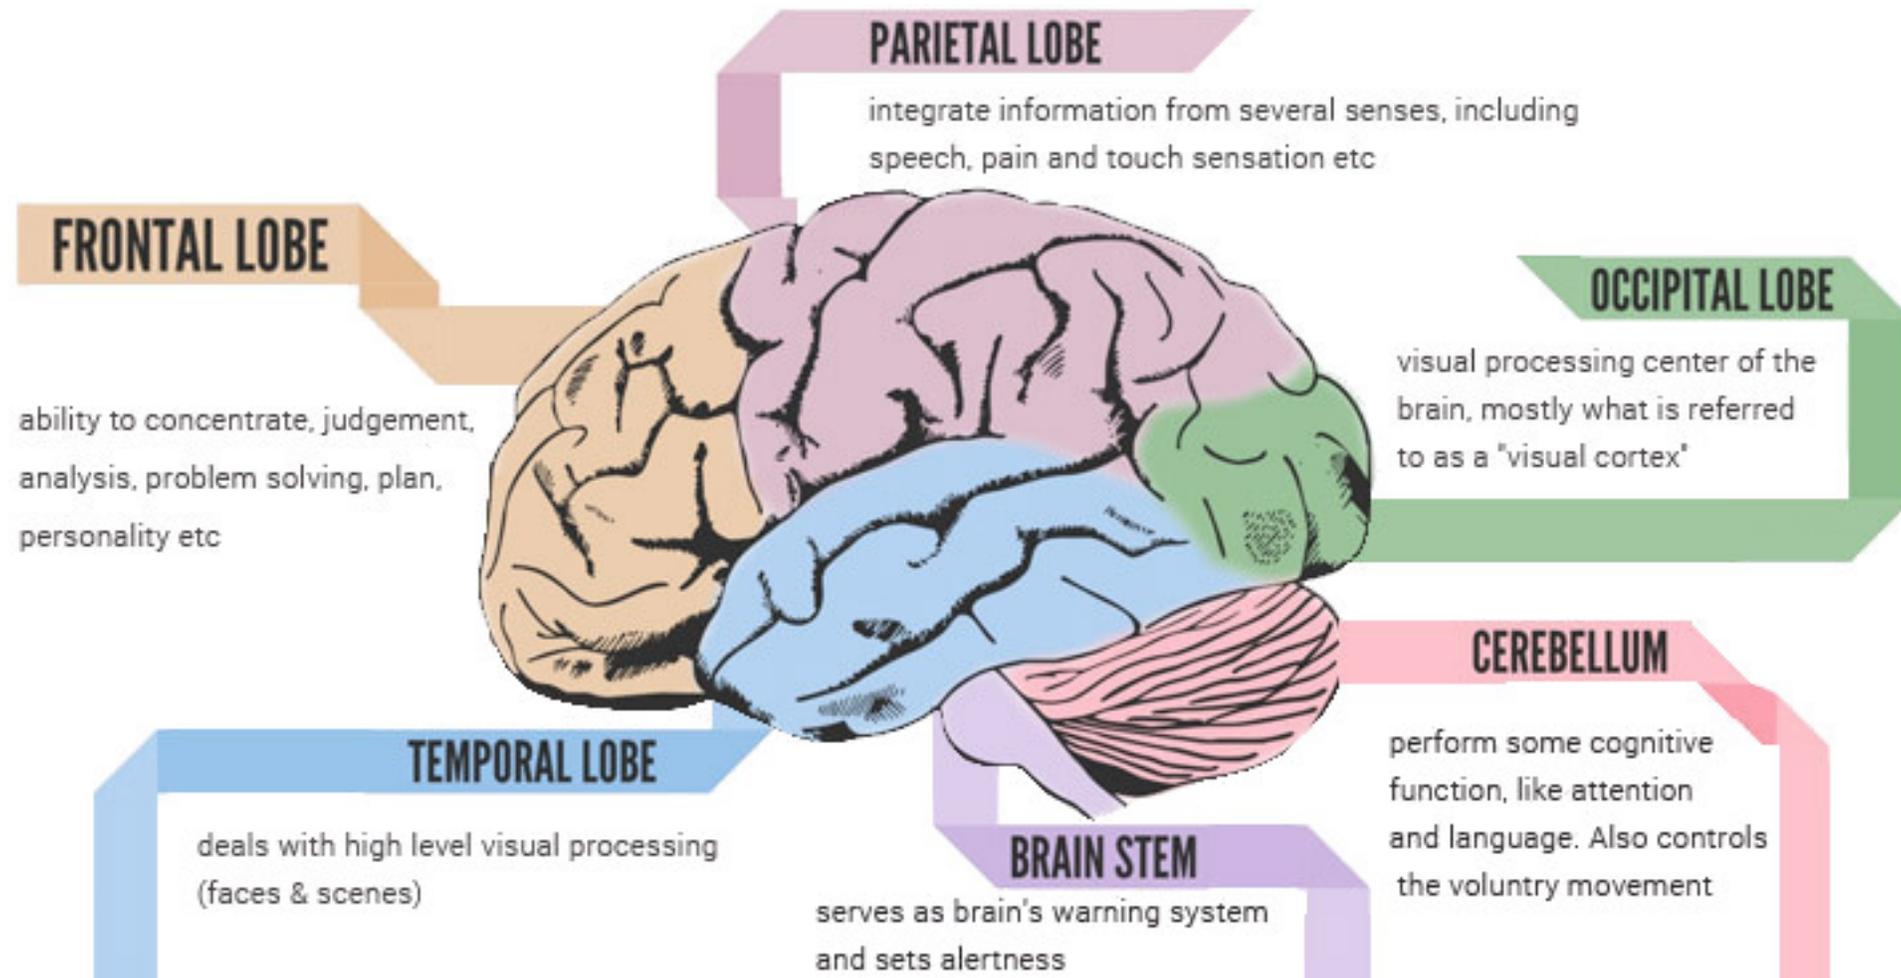
Courtesy of Quanta magazine



Note: across species, the cortex has under a third of the number of neurons in the cerebellum

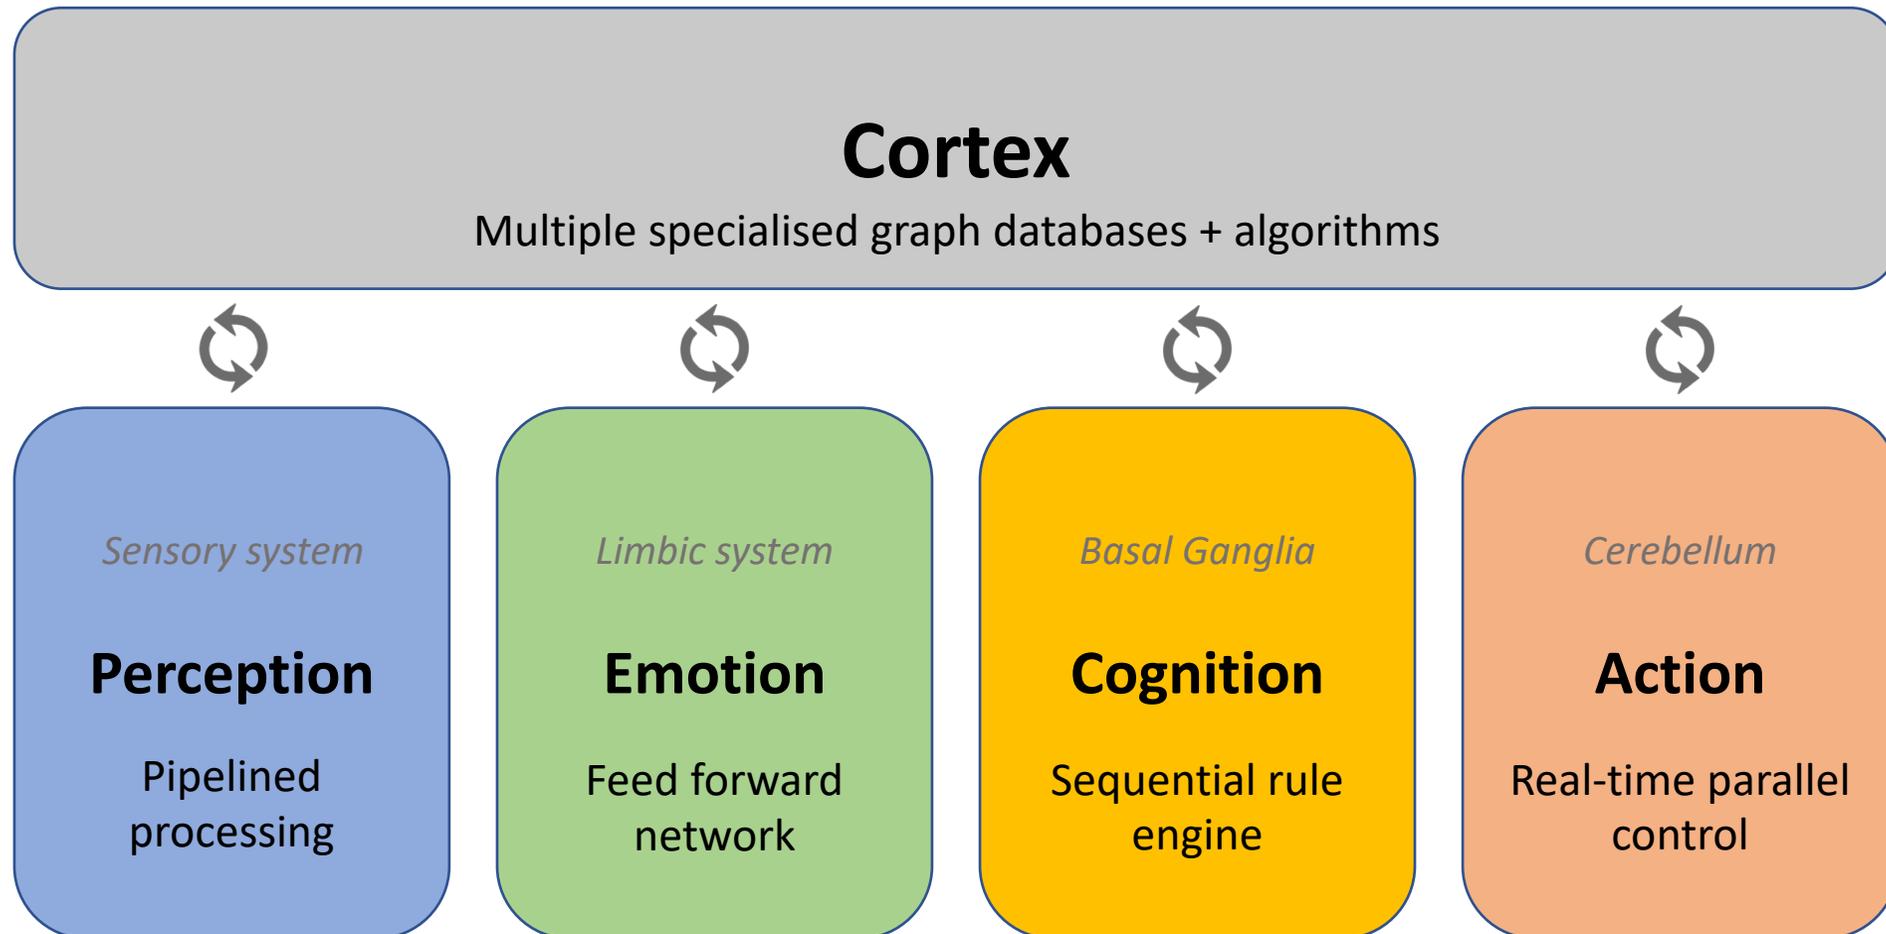


Brain function – many specialized areas



Cognitive AI Architecture

with multiple cognitive circuits





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 that include 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 rather like 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 handling Big Data
- 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

Sensory 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
 - Emotional cues

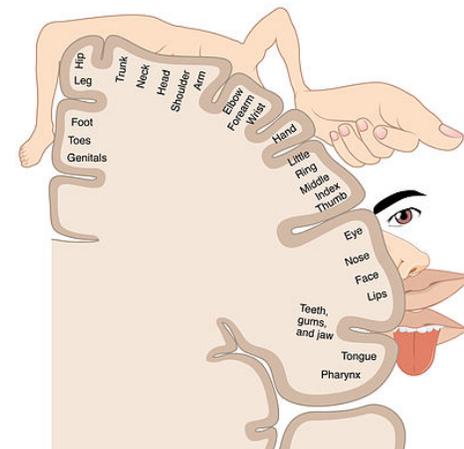
- Vision is much more complex
 - Two eyes for stereo depth perception
 - Each eye: high resolution narrow angle + low resolution wide angle
 - Saccades as eyes swivel to scan areas of interest
 - Good at recognizing many different kinds of things, including their structures & behaviours
 - Context determines what is interesting and relevant
 - Alerts signal relevant things in field of view
 - Focus directs attention to specific things
 - Reinforcement learning from experience



Visual system



Hearing



Cortical homunculus

Implementation as pipelined neural networks

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
- 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
 - Functional implementation as a feed-forward classification network

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

Cognition and Conscious Thought



- **Cortico basal-ganglia circuit**
 - The centre of conscious thought
- Symbolic (graphs) + sub-symbolic (statistics)
 - Chunk based symbolic representation of concepts and relationships
 - Statistical weights reflecting prior knowledge and past experience
- Rule engine connected to many parts of the cortex
 - Connections via buffers that hold single chunks
 - Rules represent reasoning & procedural knowledge
 - Learned from experience (hierarchical reinforcement learning)
- Sequential application of rules to cognitive buffers
 - Approximately 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

Chunks: a collection of properties that include references to other chunks

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

Action



Courtesy of Freepik

- **Cortico cerebellar circuit**
- Handles actions devolved to it by conscious thought
- Real-time control with parallel processing
- Contains more than three times the number of neurons in the cortex*
- Cerebellum acts as 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
- Learning through experience, starting with conscious thought
- Implemented as suite of real-time continuous state machines
- Examples: talking, walking and playing the piano

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

Application to 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 arriving to the destination
- Cognitive rules as chunks for ease of learning
 - Start and stop turn indicator lights
 - Initiate braking or accelerating
 - Initiate lane tracking and turning
- Functional model of cortico-cerebellar circuit provides real-time control of brakes, acceleration and steering, as initiated by cognitive rules

retrieve turn

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

prepare for turn

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

start turn

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

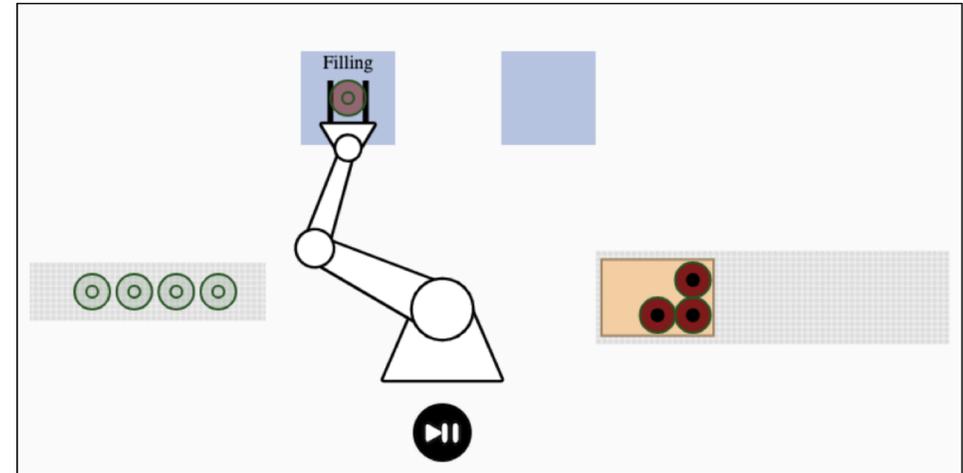
Application to Smart factories

- **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}
```

```
# add box when belt2 has space and wait afresh  
space {thing belt2} =>  
action {@do addBox; thing belt2},  
action {@do stop; thing belt2},  
space {@do wait; thing belt2; space 95}
```

```
# stop belt 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}
```



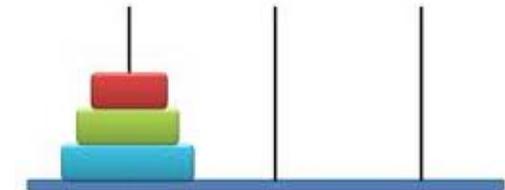
Log:

```
executed rule _:17 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
```

Natural Language Processing

- NLU as combination of pipelined processing + deliberative reasoning
 - Lexicon of words, their parts of speech, word senses and linguistic tags, e.g. person, number, case, gender
 - Spreading activation model for disambiguation based upon the context and statistical likelihood
 - Word by word generation of word dependency graph
 - Explicit reasoning to map word graph to semantics
 - Similar approach in reverse for NLG
- Work now underway on demos for NL understanding and generation
 - Start simple then iterative enrichment
 - Suggestions for scenarios and example dialogues welcomed!

Towers of Hanoi



move the red disc to the right peg

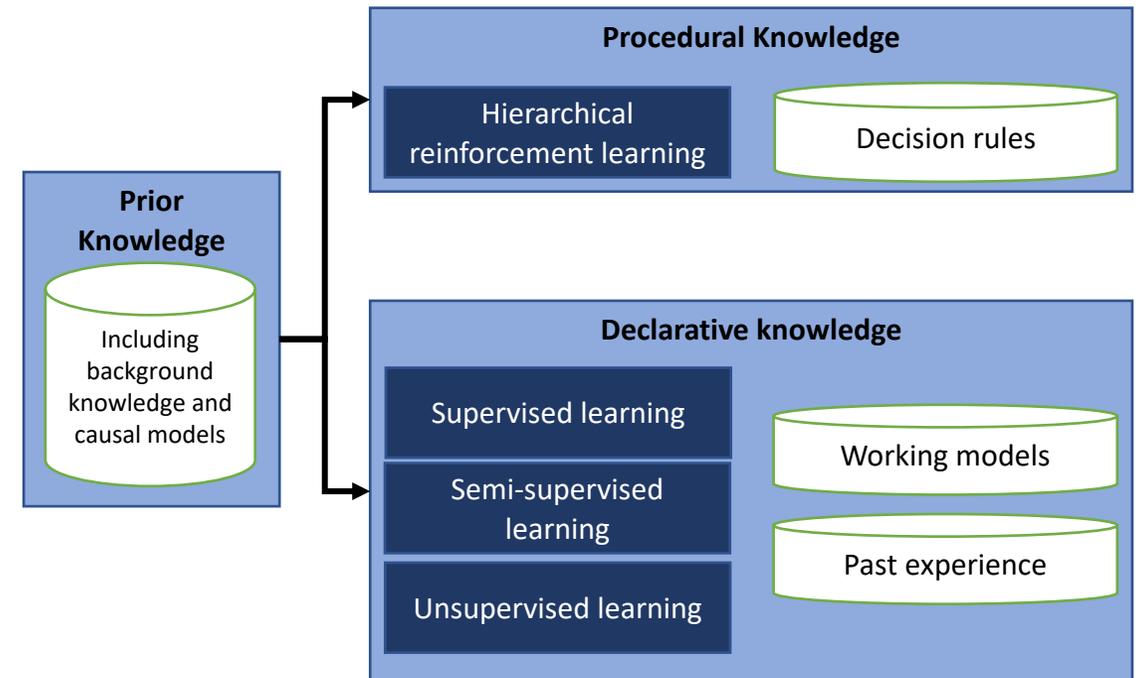
verb v1 {word move; subject p1; to p2}
phrase p1 {word disc; det the; adj red}
phrase p2 {word peg; det the; adj right}

after application of ruleset

move m1 {disc disc3; to peg3}

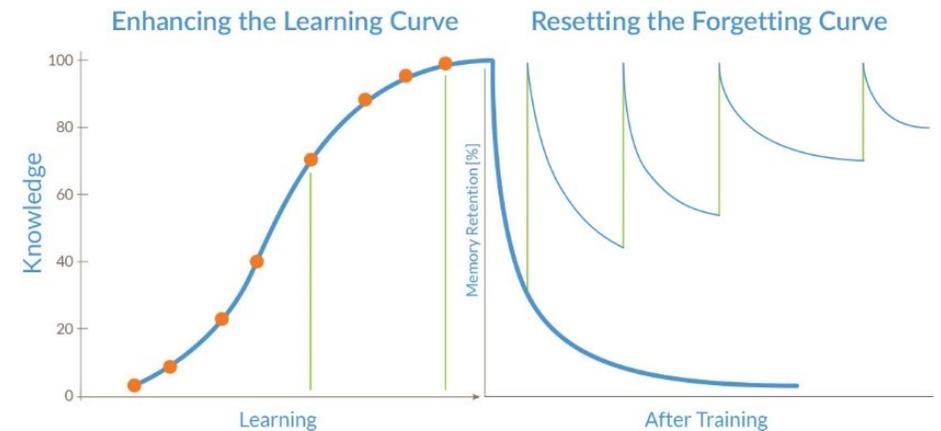
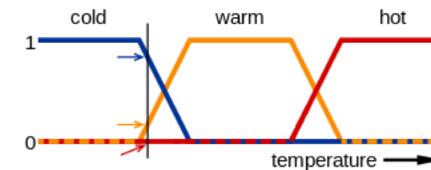
Machine Learning

- Manual development of knowledge won't scale cost effectively
- We therefore need to rely on machine learning for declarative and procedural knowledge
 - Many algorithms to take advantage of
- Prior knowledge enables learning from small datasets
- Semi-supervised learning as human guided exploration with attention to salience
- Active learning – continuous, surprise driven
 - Mimicking humans as prediction machines – we attend to novelty to improve our predictions
- Use with natural language for teaching skills to cognitive agents



Richer Ways to Reason

- Many forms of reasoning have to deal with uncertainties, e.g.
 - Induction: building models to explain regularities
 - Abduction: determining the most likely explanation of some observations
 - Causal reasoning about plans
 - 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 based upon past experience
 - Spreading activation – concepts are easier to recall on account of their relationship with other concepts*
 - Ebbinghaus forgetting curve – our ability to recall information drops off over time unless boosted by repetition
 - Closely space repetitions have less effect



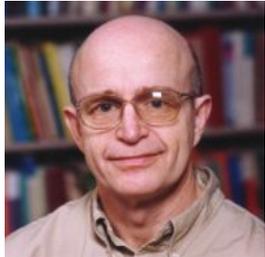
* Nature figured out the “Page Rank” algorithm many millions of years ago!

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 strong, 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

Acknowledgements



- Chunk rules are a form of production rules as introduced by [Alan Newell](#) in 1973 and later featured in his [SOAR](#) project
- [John Anderson's](#) work on human associative memory in 1973, later combined with production rules for ACT in 1976, maturing as ACT-R in 1993. [ACT-R](#) is a theory for simulating and understanding human cognition that has been widely applied to cognitive science experiments
- [Marvin Minsky](#) for his work on frames, metacognition, self-awareness and the importance of emotions for cognitive control
- European Commission for funding from the Horizon 2020 research and innovation programme under grant agreement No. 780732, [Boost 4.0](#) big data for factories



Cognitive AI

giving computing a human touch