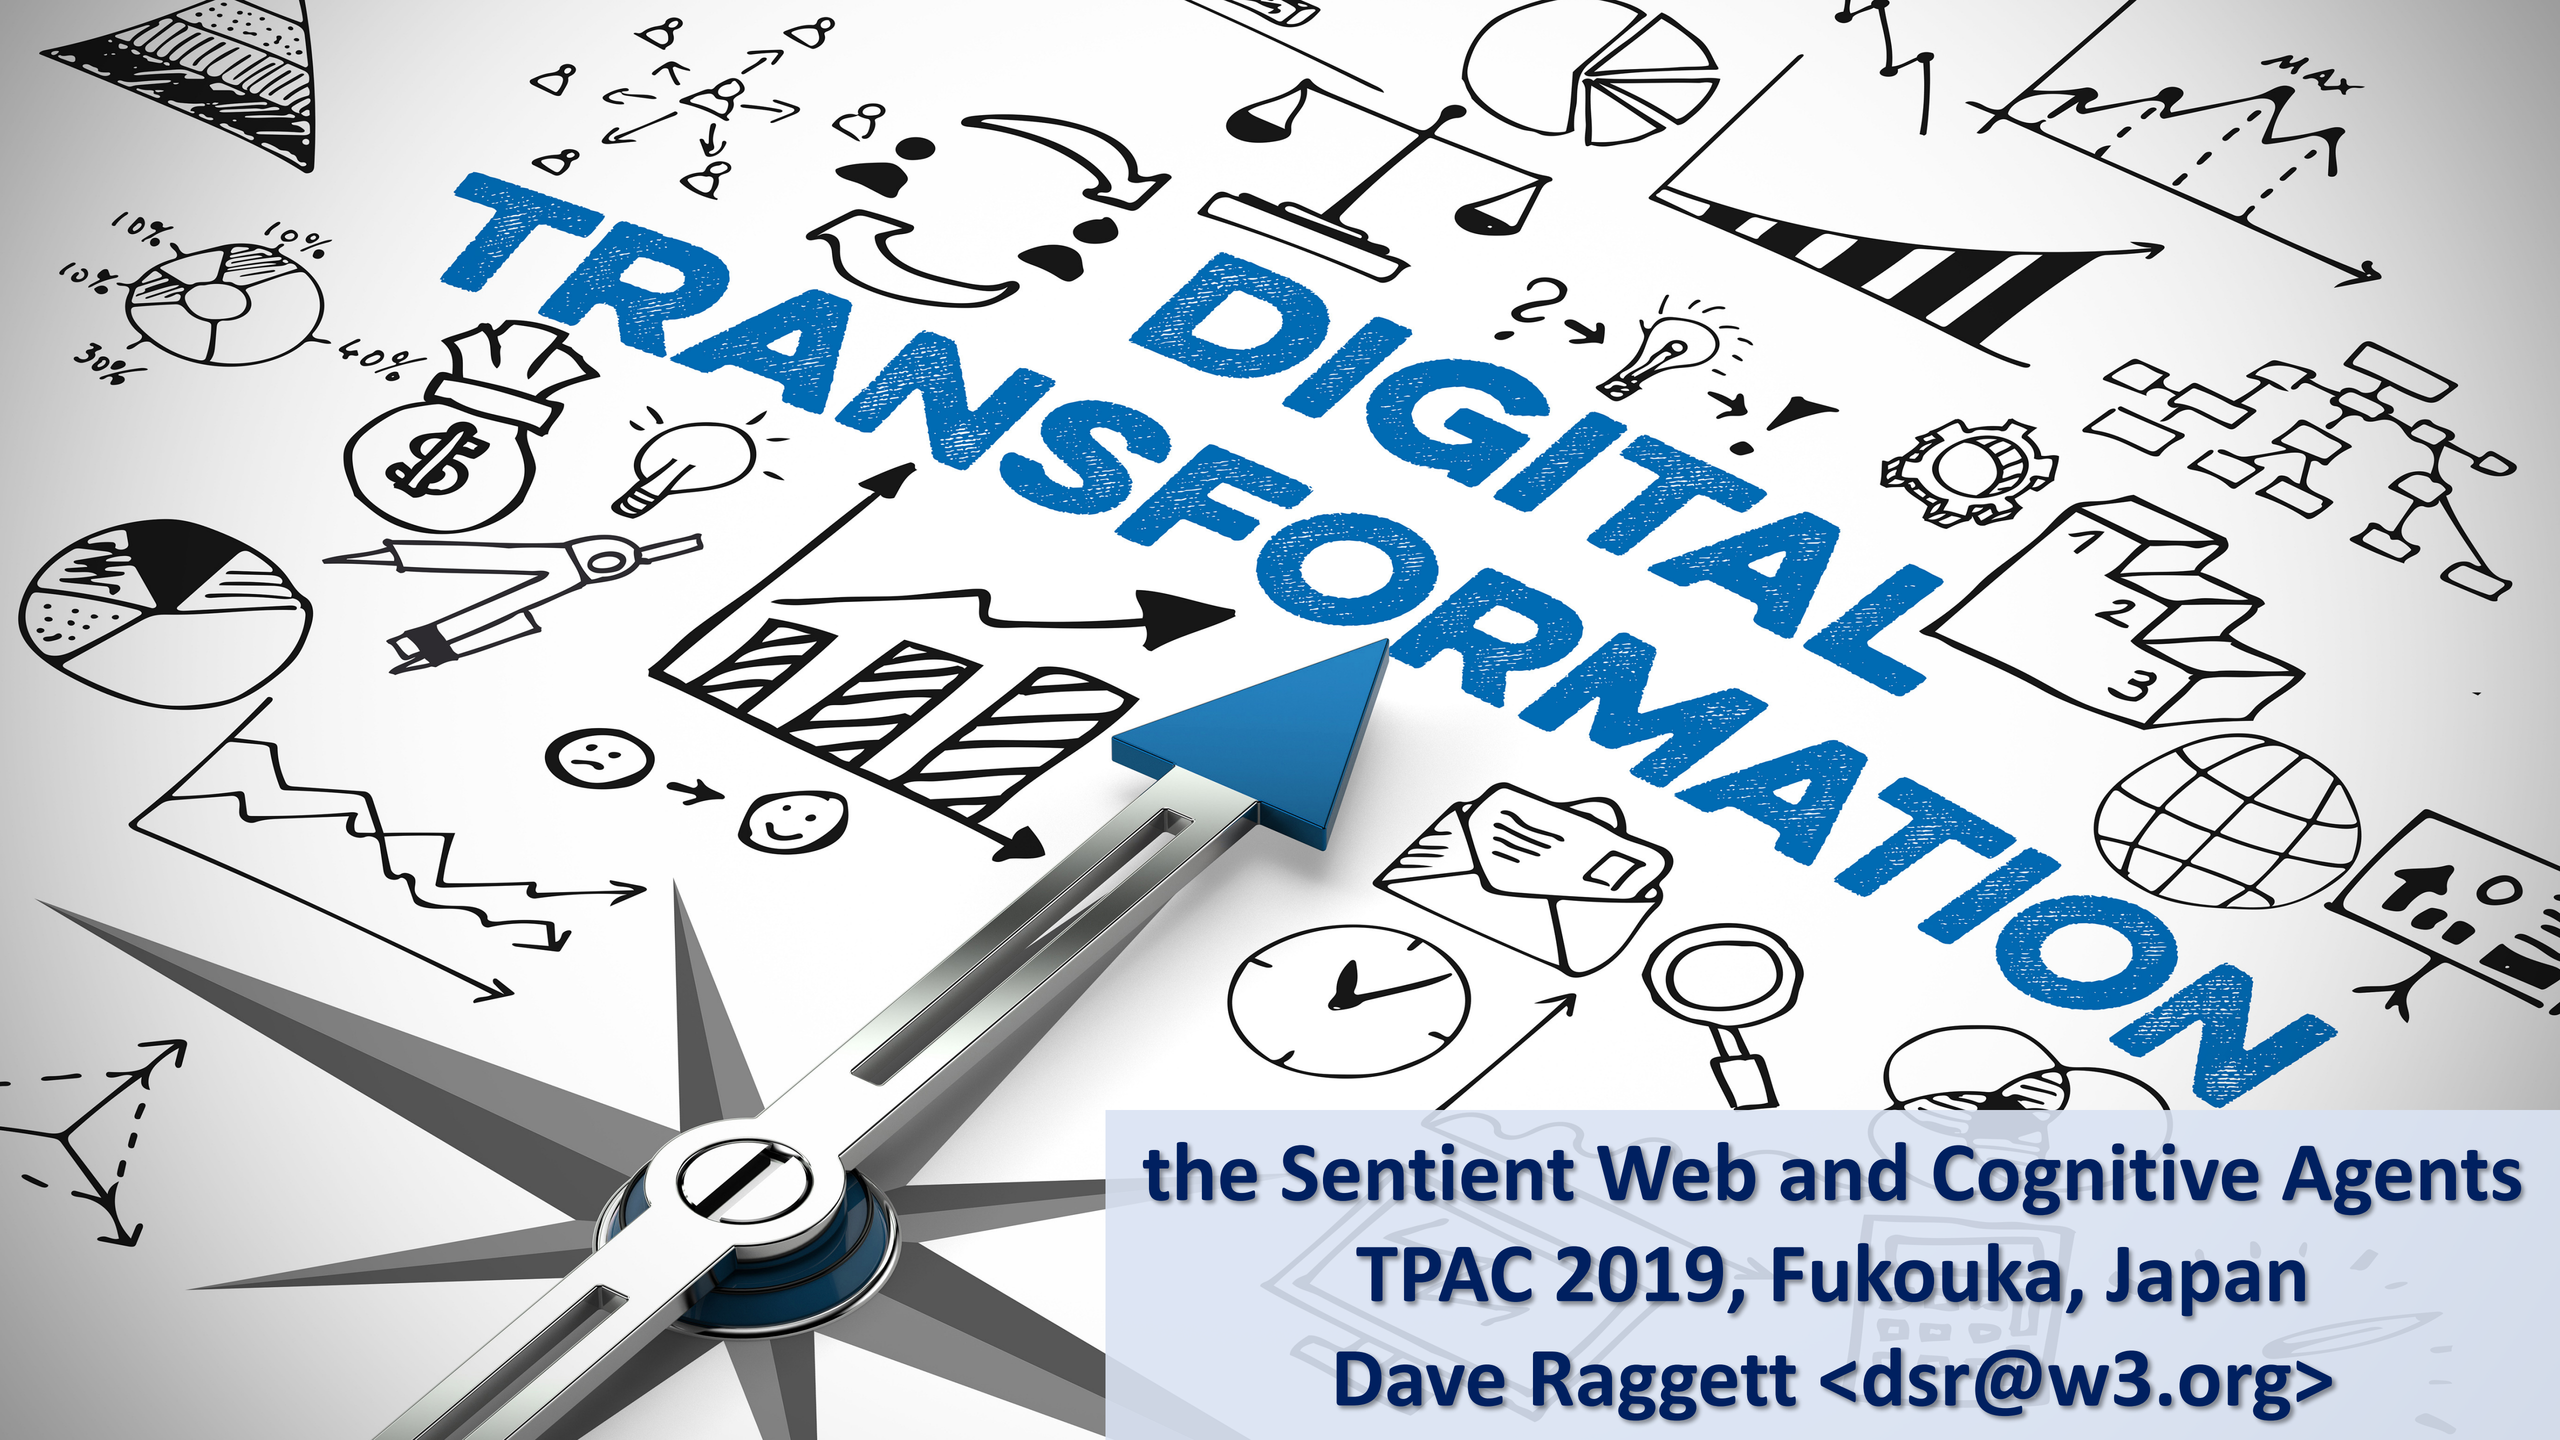

the Sentient Web and Cognitive Agents

TPAC 2019, Fukouka, Japan

Dave Raggett <dsr@w3.org>

Summary

Digital transformation can be defined as the adoption by organisations of advanced digital technologies as a means to achieve increased efficiency, improved flexibility and greater speed at adapting to changes in business conditions. This emphasises the strategic importance of data and the need for enterprise-wide data management and governance. The business challenges for digital transformation include the need to capture knowledge about information and processes in a machine interpretable form, and to bring people along rather than to impose solutions from above and risk costly failures.

From a technical perspective, the integration of information services across organisations is greatly eased with the use of graph models for concepts and relationships, moreover, when it comes to implementing digital twins for devices, along with open markets of information services, the Web of Things solves the fragmentation of the IoT.

New work is looking at ideas for simplifying the use of graphs for the average developer, and the increasing importance of machine learning to address scaling issues. Experimental work builds upon decades of progress in Cognitive Science, and embraces both RDF and Property Graphs. This exploits the latest models of how the human brain works, and some remarkable analogies with Web architecture. Help is sought for gathering use cases and data sets as a basis for evaluating alternative approaches and driving further progress.

This work is supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No 780732, project Boost 4.0

Part 1 – Digital Transformation

Digital Transformation

- The transformation of organisations based upon the adoption of advanced digital technologies to gain the following benefits
  - Greater efficiency
    - Greater productivity, reduced waste
  - Greater flexibility
    - e.g. bespoke products rather than limited choices
  - More nimble in respect to changing business needs
    - Keeping ahead of your competitors
  - A deeper engagement with your customers and suppliers

The Future is in Digital Integration

- **Vertically** from shop floor to office floor
  - Bridging different kinds of networks
- **Horizontally** across supply chain, enterprise value chain and post sales processes
  - Across organisational boundaries
- **Temporally** throughout the lifecycle for products and their instances
  - Learning how to make better products from customer experience

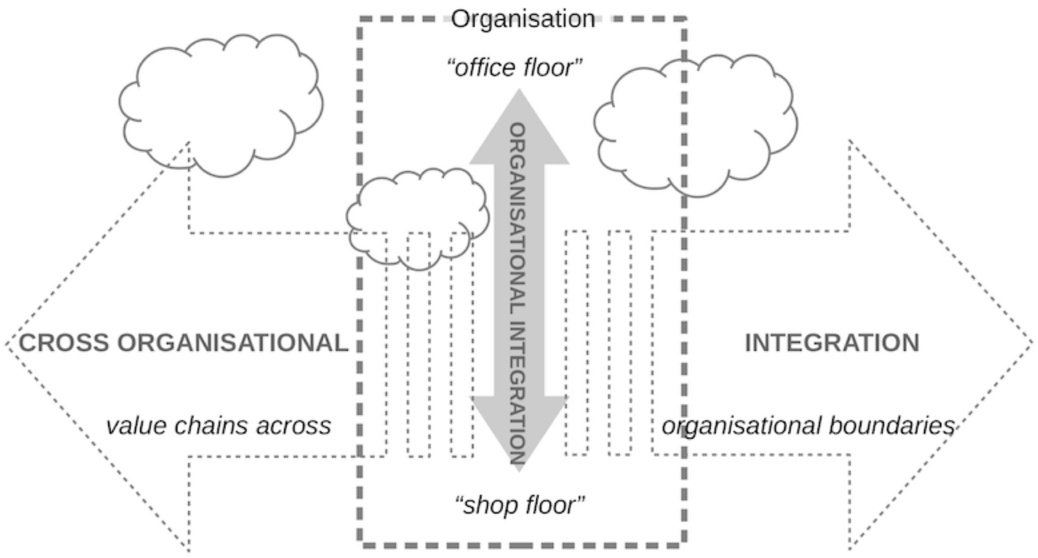

With product lifecycle as 3rd dimension

Some hurdles that will slow Integration

- Paper forms, spreadsheets and legacy database applications
- Knowledge about the meaning of data is locked inside people's heads, described in spec documents that are often outdated, or buried inside application code that no one is really sure about how to update
- People in different organisations/departments will have different ways of doing things along with different vocabularies that make it harder for collaborating across organisational boundaries

The Need for an Integration Framework based upon Open Standards

- Data is increasingly seen as a strategic asset, requiring enterprise wide management and governance
- Incremental approaches that integrate existing information systems and bring people along rather than imposing on them
- The challenges for capturing knowledge about data and processes
- This is important in respect to facilitating interoperability, management and governance

Data Management and Governance

- According to DAMA*
  - **Data Management** is a comprehensive collection of practices, concepts, procedures, processes, and a wide range of accompanying systems that allow for an organization to gain control of its data resources.
  - **Data Governance** is a collection of practices and processes which help to ensure the formal management of data assets within an organization

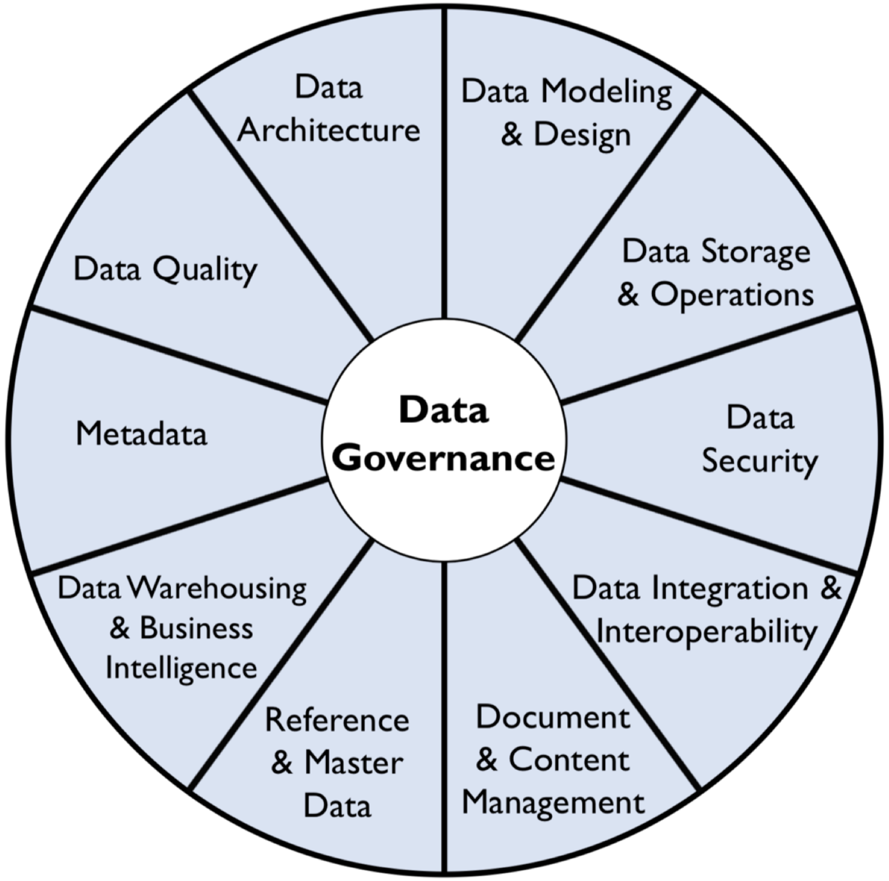

* **DAMA International** is a not-for-profit, vendor-independent, global association of technical and business professionals dedicated to advancing the concepts and practices of information and data management.

Graphs as a key enabling technology

- Graphs are a flexible way of representing knowledge and data in terms of nodes, properties and relationships
- Graphs can be reasoned over with rules and processed with efficient graph algorithms
- Graphs make it easier to combine heterogeneous data sources compared to traditional approaches
  - Improved performance over large datasets
  - Better suited for situations where the data model is evolving
  - Better suited than noSQL stores when you have links within data
- Knowledge Graphs (ontologies) for describing concepts and relationships

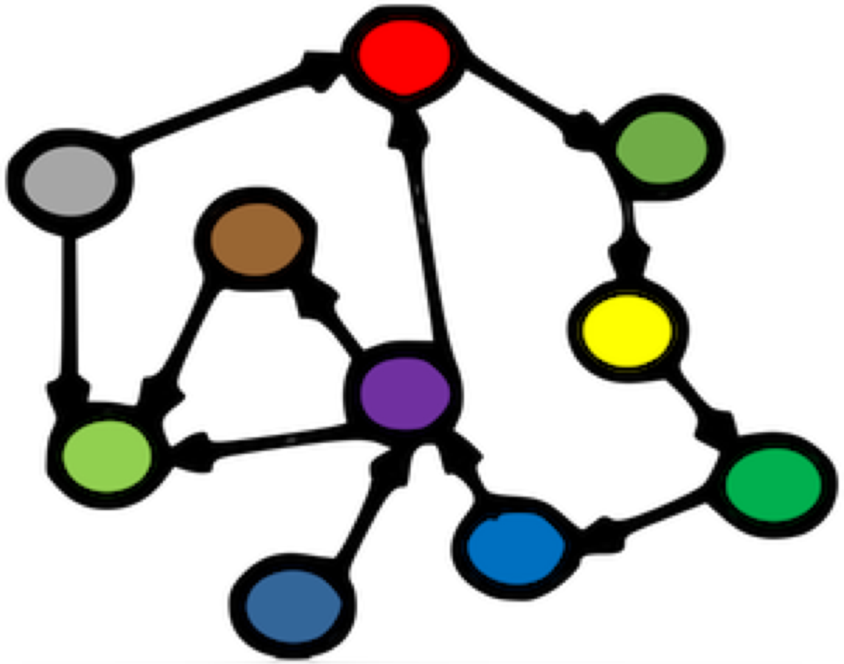

Many Graph Database Vendors, e.g.

amazon
ArangoDB
blazegraph™
CAMBRIDGE SEMANTICS

CRAY
DATASTAX
FRANZ INC.
GRAKN.AI

Graph Base
MEM GRAPH
Microsoft
neo4j

Objectivity
ontotext
OrientDB®
redislabs HOME OF REDIS

Stardog
TigerGraph
TROVARES seriously scalable
and many more …

See : https://www.bloorresearch.com/technology/graph-databases/

Sentient Web = Sensors + Graphs + AI/ML

- The Sentient Web brings together
  - **Awareness** based upon sensors and information services
    - Web of Things for digital twins and open markets of services
  - **Reasoning** based upon graphs, rules, algorithms and AI/ML
    - Federated across the enterprise

RDF as the foundations for data

- Binary named directed relationships between concepts
  - <subject, predicate, object>, e.g. Mary loves Jack
- Global identifiers for concepts and predicates
  - HTTP URLs and URNs e.g. urn:tdm:aws:property:switch
- HTTP based identifiers can be dereferenced for more information
  - HTTP content negotiation to select between formats
    - e.g. HTML, RDF/XML, Turtle, N3, JSON-LD, …
- Mature suite of standards, e.g. RDF core, OWL, SPARQL, Turtle, SHACL

W3C Graph Data Workshop (March 2019)

- Creating bridges across RDF, Property Graph and SQL communities
- Understanding requirements for interchange across graph databases
- Seeking alignment of graph query languages
  - SPARQL, GraphQL, Gremlin, AQL, Cypher, PGQL, SQL PGQ, …
- Discussion of extensions for RDF, SPARQL, JSON-LD, Turtle, N3
- Expected to lead to new standardisation work at W3C
- We're launching a **W3C Business Group on Graph Standardization**
  - Goals
    - Collect use cases and requirements across sectors
    - Guide technical standardisation work to match business needs
    - Liaisons with other SDOs and industry alliances
  - Now seeking attestations for our [White paper]() and help with drafting the Business Group Charter
  - Will proceed in parallel with work on incubating Easier RDF

See https://www.w3.org/Data/events/data-ws-2019/

Tips for Digital Transformation

Forester offers these 6 tips for successful digital transformation

1. **Start with buy-in and commitment from C-level leaders**
   - This will ensure that you have the authority to get the changes put in place
2. **Identify pathfinder projects to build momentum with small wins**
   - Early wins will convince fence sitters to support further work
3. **Communicate early and often with everyone – not just digital leaders**
   - This is essential to gain buy-in for digital projects
4. **Shift your firm's relationship with digital**
   - Understand where you are going, but don't necessarily try to get there very quickly
5. **Hire people who can help you embed agility into your organization**
   - You will need a mix of talents, e.g. including data scientists, sociologists
6. **Challenge the status quo to drive collaboration with technology teams**
   - Success depends on establishing an effective culture, which takes time to achieve

Some topics for discussion

- Anecdotal evidence, e.g.
  - UK NHS is still very much based on paper driven processes
  - Interoperability based on standard medical terminology
  - Top down centrally driven projects have been over-budget failures
  - Cybersecurity & ransom-ware attacks on old computers
- Getting the best from combination of top-down and bottom up
  - Bottom up for engagement by local staff who really know how things work
  - Top down for common framework for training, tooling and interoperability
  - Incremental, evolutionary steps to build confidence
- W3C Graph Standardisation Business Group
- What other ideas do you think are appropriate?

Part 2 – We dive into the details

Some Barriers to Adoption

- The academic community is busy working on ontologies, but their use within industry for the most part is still at a very early stage
- It takes a lot of time to understand the ontologies and how to apply them in practical ways
  - Direct manual development will become too hard and excessively expensive as the number and size of vocabularies scales up and up
- Semantic technologies are seen as hard and something for the most skilled developers
- Lack of interoperability across graph database solutions
- To be widely adopted we need to make things **much easier** for the average developer

Easier RDF*

- The value of RDF has been well proven, but it has been categorised as a niche technology and W3C is at risk of being side-lined as companies switch to Property Graph databases
- It’s therefore time to do something proactive!
- Guiding Principles for Easier RDF
  - The goal is to make RDF – or some RDF-based successor – easy enough for *average* developers (middle 33%), who are new to RDF, to be consistently successful
  - Solutions may involve anything in the RDF ecosystem: standards, tools, guidance, etc. All options are on the table
  - Backward compatibility is highly desirable, but *less* important than ease of use
- See https://github.com/w3c/EasierRDF
- So far mostly a collection of comments
- This talk shows some ideas for taking that forward

*With thanks to David Booth for launching the Easier RDF initiative

Simpler Query Languages

- There are several existing query languages, e.g.
  - SPARQL for RDF
  - GraphQL for No-SQL DB based on JSON
  - Gremin based upon scripting APIs
  - Cypher (Neo4J) for Property Graphs
- What about simple ways to express queries in RDF itself?
  - Using RDF makes it easier to support machine learning
  - One demo based on patterns inspired by SPARQL
    - https://www.w3.org/Data/demos/chunks/patterns.html
  - Another demo based on graph traversal automata inspired by ATNs
    - https://www.w3.org/Data/demos/chunks/shapes.html

Two very different query languages

Both expressed In RDF using the Turtle notation

- Pattern based queries

```
@prefix sh: <http://www.w3.org/ns/shrl#> .
@prefix ex: <http://example.org/> .
@prefix foaf: <http://xmlns.com/foaf/0.1> .

:query sh:var :x , :name , :age ;
   sh:where {
      :x ex:name :name .
      :x ex:age :age .
   } ;
   sh:filter {
      :age sh:max 22 .
      :age sh:min 20 .
   } .
```

n.b. uses an extension to RDF for annotations on sets of triples, avoiding need for reification

- Automaton based queries*

```
@prefix sh: <http://www.w3.org/ns/shrl#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>.
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#>.
@prefix ex: <http://example.org/>.

:c1
   a sh:shape ;
   sh:node ex:fido ;
   sh:next :c2 .
:c2
   sh:rel rdf:type ;
   sh:any :c3 , :c4 .
:c3
   sh:rel rdfs:subClassOf ;
   sh:any :c4 , :c3 .
:c4
   sh:node ex:animal .
```

* See previous work on ATNs for RDF shapes:
https://www.w3.org/WoT/demos/shrl/test.html

Pause for demos and discussion

Chunks as an even easier way forward!

- An approach to Easier RDF inspired by decades of advances in Cognitive Psychology and Cognitive Neuroscience
- Chunks according to Wikipedia
  - *A chunk is a collection of basic familiar units that have been grouped together and stored in a person's memory. These chunks are able to be retrieved more easily due to their coherent familiarity*
- Each chunk has an ID, a chunk type and a set of properties
- Property values are one of the following data types
  - boolean, number, name, “quoted string” or a comma separated list thereof
- Links as a subclass of chunk
  - Chunk type is predicate, properties include: subject, object
- Related to Property Graphs
  - Nodes and Relationships map to Chunks
- Related to RDF via @rdfmap* and @prefix

Examples:

```
friend f34 {
    name Joan
}
friend {
    name Jenny
    likes f34
}
dog kindof mammal
cat kindof mammal
```

* @context is used for context chains for reasoning across multiple perspectives

Chunks are an even easier way forward!

- An approach to Easier RDF inspired by decades of advances in Cognitive Psychology and Cognitive Neuroscience
- Chunks according to Wikipedia
  - *A chunk is a collection of basic familiar units that have been grouped together and stored in a person's memory. These chunks are able to be retrieved more easily due to their coherent familiarity*
- Each chunk has an ID, a chunk type and a set of properties
- Property values are one of the following data types
  - boolean, number, name, “quoted string” or a comma separated list thereof
- Links as a subclass of chunk
  - Chunk type is predicate, properties include: subject, object
- Related to Property Graphs
  - Nodes and Relationships map to Chunks
- Related to RDF via @rdfmap* and @prefix

Chunk type   Optional Chunk ID

Examples:

```
friend f34 {
    name Joan
}
friend {
    name Jenny
    likes f34
}
dog kindof mammal
cat kindof mammal
```

Property values

Some Links
Syntactic sugar for full chunk

* @context is used for context chains for reasoning across multiple perspectives

Chunk Rules

- Simple rule language with conditions and actions expressed as chunks
- Inspired by ACT-R, one of the most popular cognitive architectures
  - Successfully used over the decades to model human behaviour, e.g. mental arithmetic and driving a car
  - Combines symbolic and subsymbolic approaches
  - Stochastic memory retrieval and reinforcement learning of rules
- A number of graph data modules: goals, facts, rules, input, output
- Each module is associated with a buffer that holds a single chunk
  - Mimics constraints of human short term memory
  - This turns out to be key to scalability
- Conditions match buffers
- Actions either directly update buffers or indirectly via exchange of messages with modules to invoke graph algorithms, e.g. queries and updates

Example rule:

```
rule r2 {
   @condition g2
   @action a4
}
count g2 {
   @module goal
   start ?num
   end ?num
   state counting
}
increment a4 {
   @module goal
   @action update
   state stop
}
```

Pause for Demos and Discussion

- Chunks and chunk rules in more detail
  - https://www.w3.org/Data/demos/chunks/chunks.html
- Existing demos
  - Counting 1, 2, 3, …
    - https://www.w3.org/Data/demos/chunks/counting/
  - Decision trees
    - https://www.w3.org/Data/demos/chunks/decision-tree/
- Other demos are on their way
  - Autonomous driving focusing on task management + attention + spatial/temporal reasoning
    - https://www.w3.org/Data/demos/chunks/driving/track.html
    - https://www.w3.org/Data/demos/chunks/driving/drive.html
  - Inductive learning of taxonomies from (noisy) examples
  - Reinforcement learning of rulesets for task acquisition
  - Reasoning within and across multiple contexts

Cognitive Databases

Memory modules are Cognitive Databases

- Stores for chunks
  - Indexed by Chunk type and ID
  - Further indexes based on patterns of usage
- Accessed via buffers
- Stochastic recall imitating human memory
  - Probability of successful recall based upon estimated utility given prior knowledge and past experience
  - Analogous to Web search engines
- Scalable queries and graph algorithms
  - Invoked by clients in standard ways
  - Extensible via server-side apps

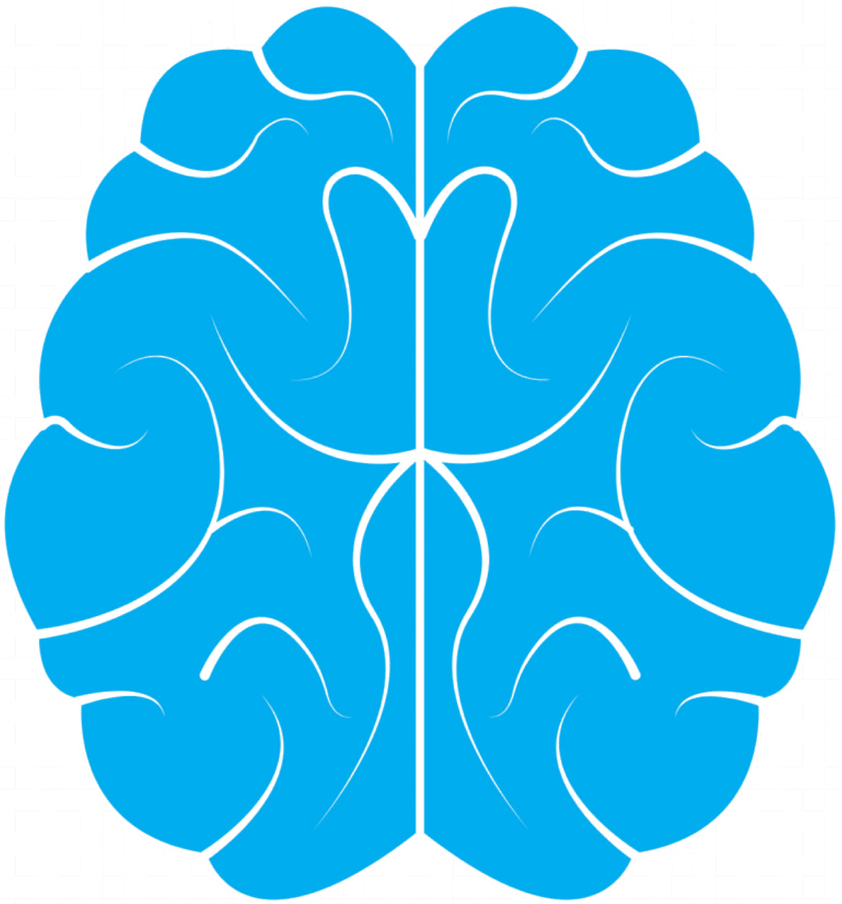

with thanks to the geoffrey w @ Neuroscience practice institute

Cognitive Architecture for Chunk Rules

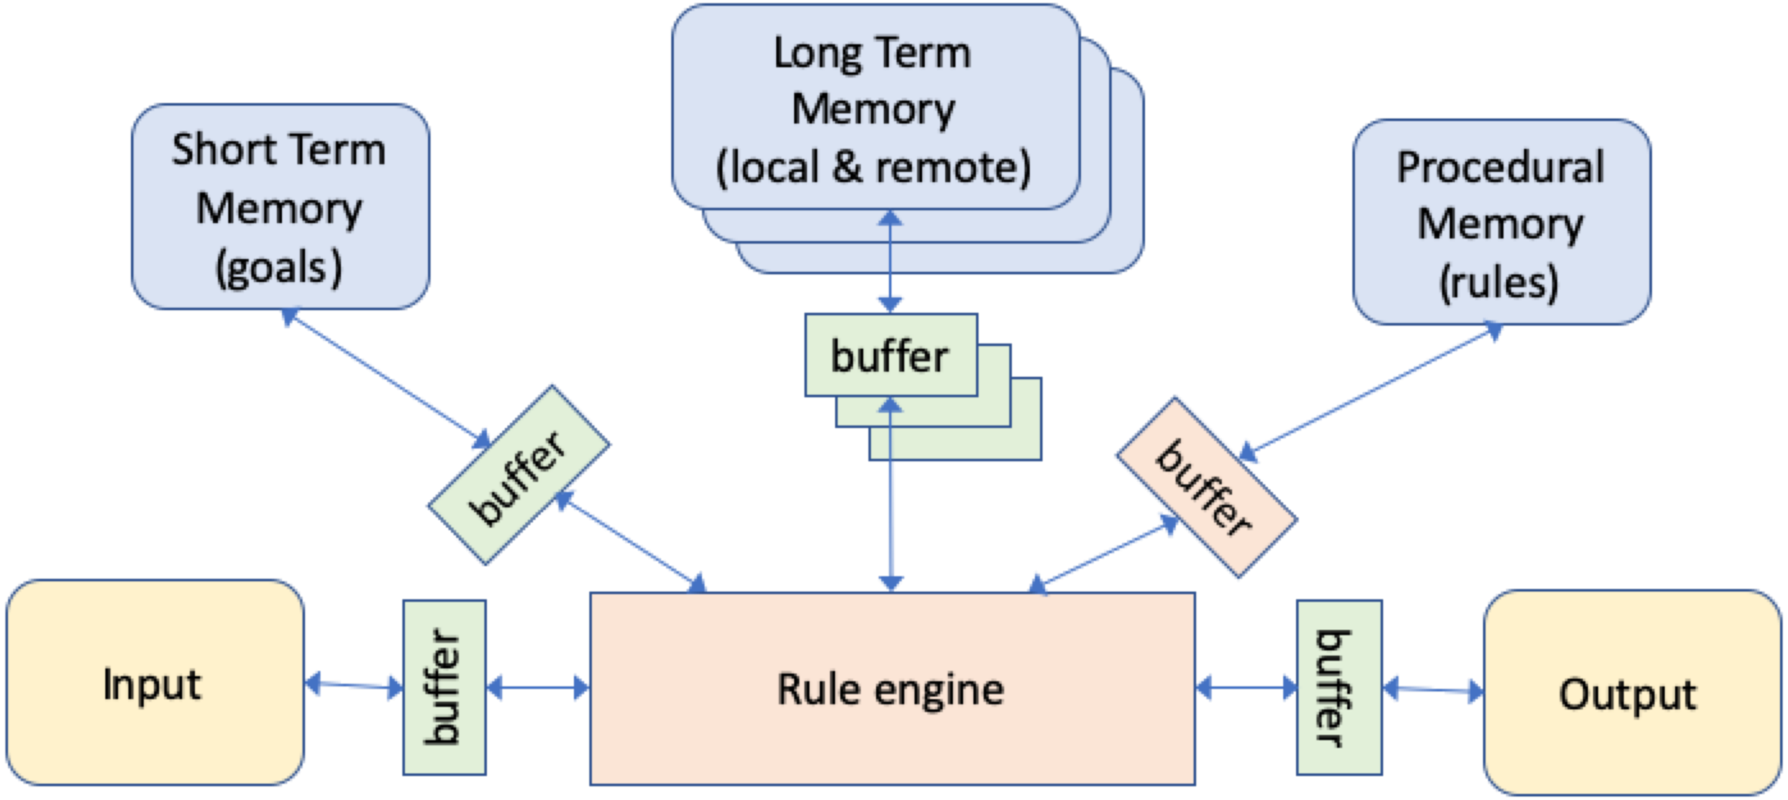

Module buffers hold a single chunk, Rule buffer holds a single rule

Machine Learning combining symbolic and sub-symbolic (statistical) approaches

- Stochastic database queries based upon prior knowledge and past experience
  - Closer to Web search engines than traditional databases
  - ACT-R's persistent strengths and dynamic activation levels that decay over time
- Enabling cognitive agents to learn from experience
  - Inductive learning of taxonomies and decisions trees from examples
    - Coping with noisy examples
    - Full, partial, or no supervision
  - Reinforcement learning of rulesets across trials
    - Heuristics propose new/updated rules
    - Success or failure back propagated up the rule chain to estimate rule utility for solving tasks
    - Stochastic temperature to explore problem space
  - Case based reasoning to speed learning

Three Stages of skill retention

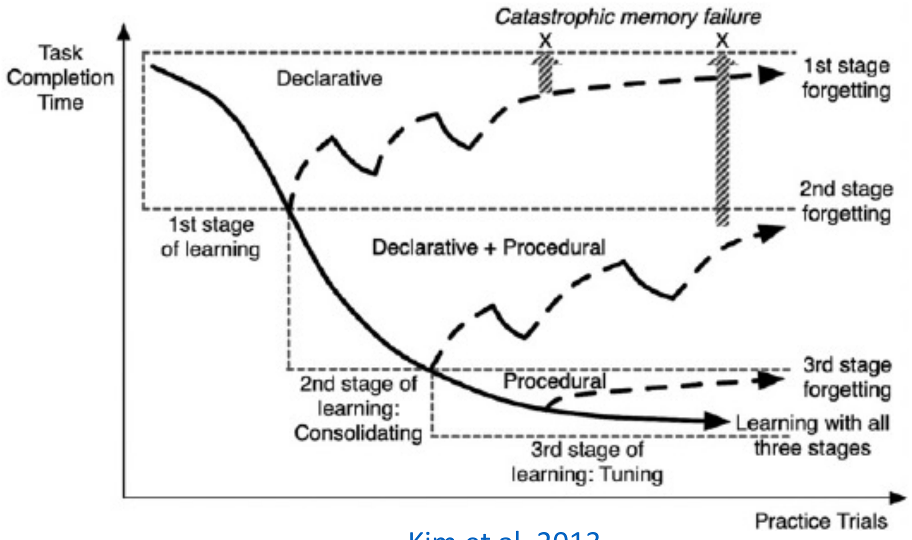

Kim et al. 2013

*ACT-R is a widely regarded cognitive architecture developed by John R. Anderson at CMU

Inspired by what we know of the brain

- The cortex functions as a collection of graph databases accessed via a request/response pattern
  - Analogy with Web architecture
- The basal ganglia and thalamus together function as a rule engine, transforming inputs to outputs
- Rules fire at approximately 50mS intervals
  - Consciousness is sequential, recall is highly parallel
- Combination of symbolic and subsymbolic (statistical) approaches
  - Essential for rational reasoning based upon prior knowledge and past experience
  - Analogy with Web search
- Layered on top of spiking neural networks
  - Chunks passed as spiking patterns over bundles of nerves
  - See Eliasmith 2015, "[Concepts as Semantic Pointers]"

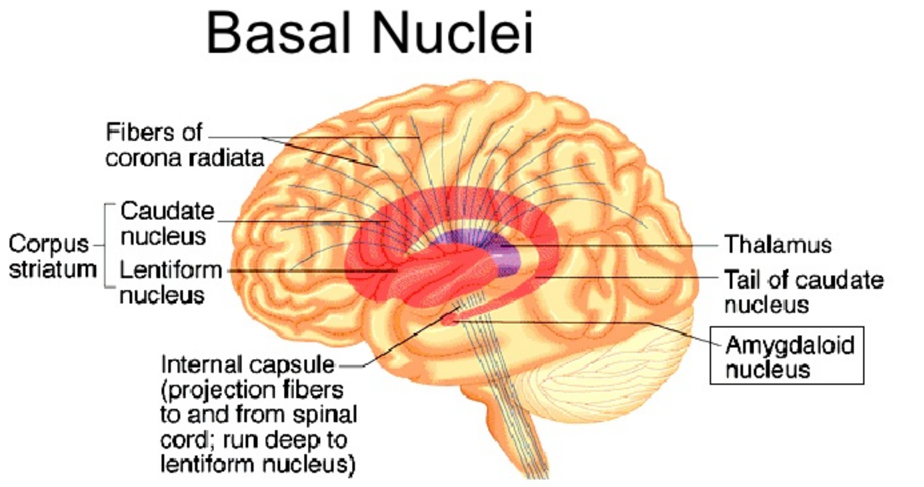

- The basal nuclei receive inputs from the entire cerebral cortex as well as from other subcortical nuclei
- Via relays through the thalamus, the basal nuclei project to the premotor and prefrontal cortices

ACT-R provides accurate predictions for BOLD responses in brain scans of subjects performing the modelled task

Delegated Command and Control

- **The brain has two parallel systems**
- One system is responsible for consciousness and voluntary control
- Control for movement, speech and so forth is delegated to a second system
  - The cerebellum dynamically regulates movement using connections to the sensory systems, the spinal cord, and other parts of the brain

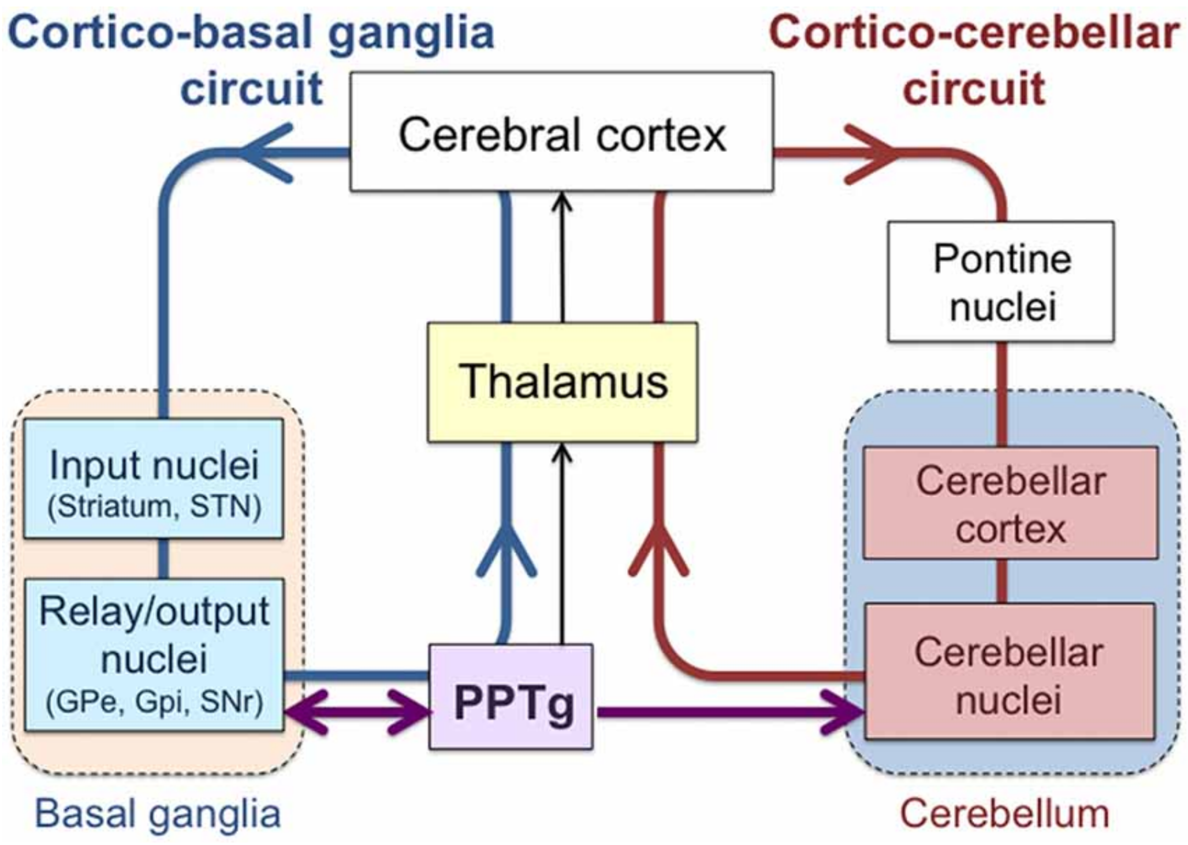

With thanks to Fumika Mori et al.

Cognitive Agents and the Sentient Web

Distributed architecture with many cognitive agents and information systems across the enterprise, designed for human-machine collaborative work

Local Components

- Input & Output (Web of Things)
- Goals module
- Local Rules module
- Short term memory module*
- Rule engine

Remote Components

- Long term memory modules with different locations and functions across the enterprise
  - Declarative & Procedural
  - Interfaces to legacy systems
- Delegated command & control†

* This mimics the hippocampus which focuses on short term memory, while the cortex is devoted to long term memory. The short term memory module supports rich queries on long term memory that return responses as graphs rather than as individual chunks.

† This mimics the cerebellum and manages coordinated control, analogous to walking.

The Sentient Web further out …

A new paradigm for AI – bringing a human touch

- The Semantic Web focuses on deductive logic
- This doesn't work with incomplete, uncertain and inconsistent data which is likely to include errors
  - In short, it doesn't work well with raw real world data
  - Work around involves cleaning data using rough rules of thumb
    - The dark underbelly of Data Science …
- We will need to make use of ***rational* forms of reasoning** exploiting prior knowledge and past experience
  - Many forms of reasoning: deductive, inductive, abductive, causal, counterfactual, temporal, spatial, social, emotional, humour, empathy, compassion, reflection (reasoning about reasoning) …
- Hence the need to **blend symbolic knowledge** (graph data and rules) **with computational statistics**
  - Cognitive Psychology has shown how this can be realised
  - This will have a very profound effect on society
- This will allow us to create the ***Sentient Web***
  - Ecosystems of services with **awareness** based upon sensors and information services, and **reasoning** based upon graph data and rules together with graph algorithms and machine learning
  - Sentient Web = digital twins + graphs + AI/ML

Where next?

- Continuing progress is critically dependent on identifying appropriate use cases and data sets
  - E.g. machine learning, different kinds of reasoning, everyday skills, natural language processing, integration with deep learning, …
  - **Let’s collaborate on compiling use cases, and seeing which kinds of approaches are the most effective!**
- This work is very much at the incubation phase, but builds on upon decades of advances in Cognitive Psychology and Cognitive Neuroscience
- Further opportunities to exploit ideas from other fields, e.g. Sociology, for social reasoning as part of human/machine collaboration
- Proprietary vs non-proprietary approaches to datasets, training materials and simulated environments for developing cognitive agents?
- Your help is very much welcomed!
- Please contact Dave Raggett <dsr@w3.org>