

Natural Language Understanding

John F. Sowa & Arun K. Majumdar

Kyndi, Inc.

Data Analytics Summit, December 2015

Revised 15 June 2017

Outline

1. Why are natural languages so hard to analyze?

Computers process syntax and logic very well. Difficulties arise from the many ways of thinking and acting in and on a complex world.

2. Hybrid systems are necessary to support diversity

Flexibility and generality are key to intelligence. No single algorithm or paradigm can do everything or talk about everything.

3. Cognitive Computing


For any specific task, a computer simulation can often do as well or better than humans. But people are superior to any computer system in relating, integrating, and talking about all possible tasks.

4. Cycles of Learning and Reasoning

The cognitive cycle of induction, abduction, deduction, and testing.

For videos of the talks presented at the Data Analytics Summit (including this one), see <http://livestream.com/hulive/datasummit/videos/107034438>

Natural Language Processing



A classroom in 2000, as imagined in 1900 *

* <http://publicdomainreview.org/collections/france-in-the-year-2000-1899-1910/>

1. What Makes Language So Hard

Early hopes for artificial intelligence have not been realized. *

Language understanding is more difficult than anyone thought.

A three-year-old child is better able to learn, understand, and speak a language than any current computer system.

Tasks that are easy for many animals are impossible for the latest and greatest robots.

Questions:


- Have we been using the right theories, tools, and techniques?
- Why haven't these tools worked as well as we had hoped?
- What other methods might be more promising?
- What can research in neuroscience and psycholinguistics tell us?
- Can it suggest better ways of designing intelligent systems?

* See [The best AI still flunks 8th grade science](#), *Wired Magazine*.

Early Days of Artificial Intelligence

- 1960:** Hao Wang's theorem prover took 7 minutes to prove all 378 FOL theorems of *Principia Mathematica* on an IBM 704 – much faster than two brilliant logicians, Whitehead and Russell.
- 1960:** Emile Delavenay, in a book on machine translation:
“While a great deal remains to be done, it can be stated without hesitation that the essential has already been accomplished.”
- 1965:** Irving John Good, in speculations on the future of AI:
“It is more probable than not that, within the twentieth century, an ultraintelligent machine will be built and that it will be the last invention that man need make.”
- 1968:** Marvin Minsky, a technical adviser for the movie *2001*:
“The HAL 9000 is a *conservative estimate* of the level of artificial intelligence in 2001.”

HAL 9000 in 2001: A Space Odyssey

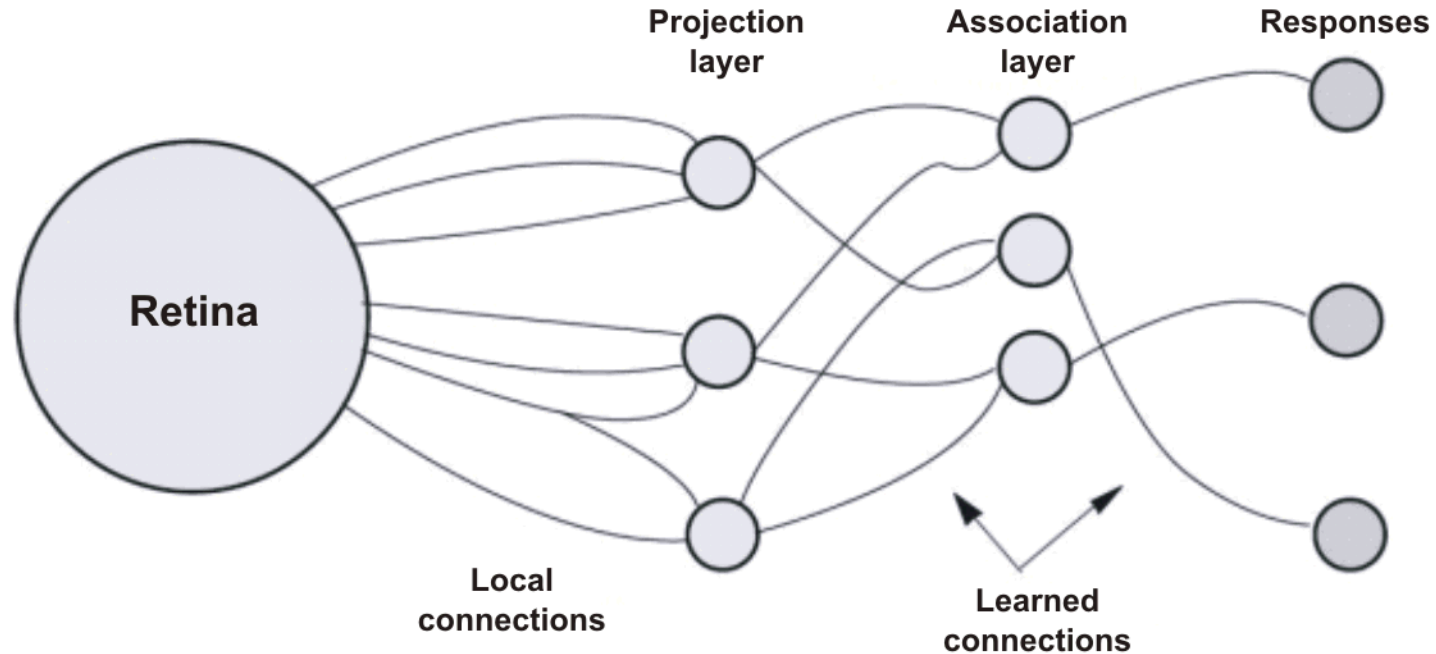


The advisers made two incorrect predictions:

- **Hardware technology developed faster than they expected.**
- **But software, including AI, developed much slower.**

Predicting a future invention is almost as hard as inventing it.

The Perceptron



One-layer neural network invented by Frank Rosenblatt (1957).

Mark I: a hardware version funded by the US Navy:

- Input: 400 photocells in a 20 x 20 array.
- Weights represented by potentiometers updated by electric motors.

The New York Times, after a press conference in 1958:

*The perceptron is “the embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.” **

* <http://query.nytimes.com/gst/abstract.html?res=9D03E4D91F3AE73ABC4B52DFB1668383649EDE>

A Breakthrough in Machine Learning



Program for playing checkers by Art Samuel in 1959:

- Ran on the IBM 704, later on the IBM 7090.
- The IBM 7090 was comparable in speed to the original IBM PC (1981), and its maximum RAM was only 144K bytes.

Samuel's program was a hybrid:

- A perceptron-like algorithm for learning to evaluate game positions.
- The alpha-beta algorithm for searching game trees.

Won a game against the Connecticut state checkers champion.

Bird Nest Problem

Robots can perform many tasks with great precision.

But they don't have the flexibility to handle unexpected shapes.

They can't wash dishes the way people do — with an open-ended variety of shapes and sizes.


And they can't build a nest in an irregular tree with irregular twigs, straw, and moss.

If a human guides a robot through a complex task with complex material, the robot can repeat the same task in the same way.

But it doesn't have the flexibility of a bird, a beaver, or a human.



Understanding Language



The syntax is easy: Parse the question and the answer.

Semantics is harder: Use background knowledge to

- Recognize the situation type and the roles of the two agents,
- Relate the word 'thing' to the picture and to the concept Car,
- Relate the verbs 'take' and 'move' to the situation,
- Apply the laws of physics to understand the answer.

Pragmatics is the hardest: Explain the irony and the humor.

* Search for 'moving' at <http://www.shoecomics.com/>

The Ultimate Understanding Engine

Sentences uttered by a child named Laura before the age of 3. *

Here's a seat. It must be mine if it's a little one.

I went to the aquarium and saw the fish.

I want this doll because she's big.

*When I was a little girl, I could go "geek geek" like that,
but now I can go "This is a chair."*

No computer system today can learn and use language as fast, as accurately, and as flexibly as a child.

Preschool children constantly ask "Why?"

Those questions get into the pragmatics. They are the hardest for parents and computer systems to answer.

* John Limber, The genesis of complex sentences.

http://pubpages.unh.edu/~jel/JLimber/Genesis_complex_sentences.pdf

Child Reasoning

A mother talking with her son, about the same age as Laura: *

Mother: *Which of your animal friends will come to school today?*

Son: *Big Bunny, because Bear and Platypus are eating.*

The mother looks in his room, where the stuffed bear and the platypus are sitting in a chair and “eating”.

The child relates the sentences to the situation:

- **The bear and the platypus are eating.**
- **Eating and going to school cannot be done at the same time.**
- **Big Bunny isn't doing anything else.**
- **Therefore, Big Bunny is available.**

This reasoning is more “logical” than anything that Siri says.

* Reported by the father, the psychologist Gary Marcus, in an interview with Will Knight (2015)
<http://www.technologyreview.com/featuredstory/544606/can-this-man-make-ai-more-human/#comments>

Mental Maps, Images, and Models

Quotation by the neuroscientist Antonio Damasio (2010):

“The distinctive feature of brains such as the one we own is their uncanny ability to create maps... But when brains make maps, they are also creating images, the main currency of our minds. Ultimately consciousness allows us to experience maps as images, to manipulate those images, and to apply reasoning to them.”

The maps and images form mental models of the real world or of the imaginary worlds in our hopes, fears, plans, and desires.

Words and phrases of language can be generated from them.

They provide a “model theoretic” semantics for language that uses perception and action for testing models against reality.

Like Tarski’s models, they determine the criteria for truth, but they are flexible, dynamic, and situated in the daily drama of life.

Role of Imagery in Mathematics

Paul Halmos, mathematician:

“Mathematics — this may surprise or shock some — is never deductive in its creation. The mathematician at work makes vague guesses, visualizes broad generalizations, and jumps to unwarranted conclusions. He arranges and rearranges his ideas, and becomes convinced of their truth long before he can write down a logical proof... the deductive stage, writing the results down, and writing its rigorous proof are relatively trivial once the real insight arrives; it is more the draftsman’s work not the architect’s.” *

Albert Einstein, physicist:

“The words or the language, as they are written or spoken, do not seem to play any role in my mechanism of thought. The psychical entities which seem to serve as elements in thought are certain signs and more or less clear images which can be *voluntarily* reproduced and combined... The above-mentioned elements are, in my case, of visual and some of muscular type. Conventional words or other signs have to be sought for laboriously only in a secondary stage, when the mentioned associative play is sufficiently established and can be reproduced at will.” **

* Halmos (1968). ** Quoted by Hadamard (1945). See also Lakoff & Núñez (2000).

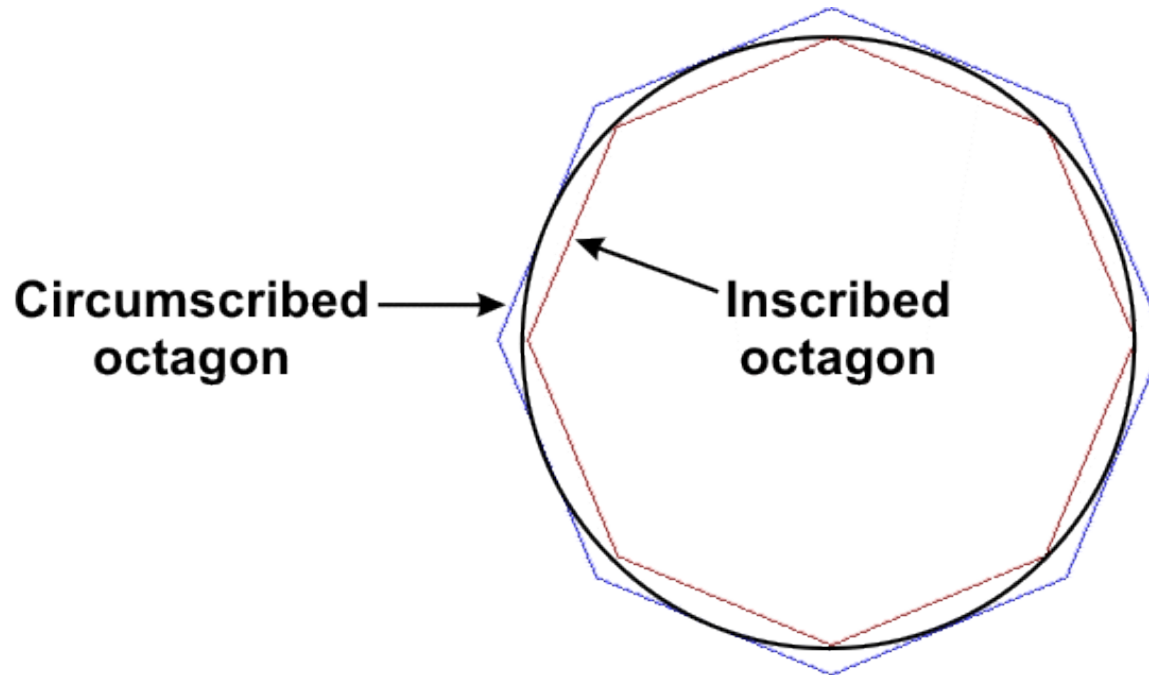
Archimedes' Eureka Moment



Insight: A submerged body displaces an equal volume of water.

- It's a mathematical principle, a property of Euclidean space.
- Scientists and engineers have used it ever since.
- They don't prove it. They use it to define *incompressible fluid*.

Determining the Value of π




Archimedes had two creative insights, both inspired by images:

- **The circumference of the circle is greater than the perimeter of the inner polygon and less than that of the outer polygon.**
- **As the number of sides increases, the inner polygon expands, and the outer polygon shrinks. They converge to the circle.**

Given these insights, a good mathematician could compute π to any desired precision. Archimedes used 96-gons.

Euclid's Proposition 1



Euclid's statement, as translated by Thomas Heath:

- **On a given finite straight line, to draw an equilateral triangle.**

The creative insight is to draw two circles:

- **The circle with center at A has radii AB and AC.**
- **The circle with center at B has radii BA and BC.**
- **Since all radii of a circle have the same length, the three lines AB, AC, and BC form an equilateral triangle.**

Feelings and Emotions

Damasio and Carvalho (2013),

- **“Feelings are mental experiences of body states.”**
- **“They signify physiological need, tissue injury, optimal function, threats to the organism, or specific social interactions.”**
- **“Feelings constitute a crucial component of the mechanisms of life regulation, from simple to complex.”**
- **“Their neural substrates can be found at all levels of the nervous system, from individual neurons to subcortical nuclei and cortical regions.”**

Damasio (2014), *


- **“I’m ready to give the very teeny brain of an insect – provided it has the possibility of representing its body states – the possibility of having feelings.”**
- **“Of course, what flies don’t have is all the intellect around those feelings that could make use of them: to found a religious order, or develop an art form, or write a poem.”**

* Interview, <http://www.technologyreview.com/qa/528151/the-importance-of-feelings/>

Cognitive Learning

Areas of the cerebral cortex are highly specialized.

A study with fMRI scans shows which areas are active at various stages of cognitive learning. *




14 participants studied how four devices work: bathroom scale, fire extinguisher, automobile braking system, and trumpet.

Cognitive learning is more complex than DNN learning:

- 1. The visual cortex is most active in the initial perceptual stage.**
- 2. Parietal lobes are active while “imagining the components moving.”**
- 3. All lobes become active as participants are “generating causal hypotheses” about how the system works.**
- 4. Finally, the frontal cortex is active while “the person (probably oneself)” imagines what it would be like to “interact with the system.”**

* R. A. Mason & M. A. Just (2015) <http://www.sciencedirect.com/science/article/pii/S1053811915000841>

The Brain in Language Learning



Language learning increases connectivity among brain regions.

- 39 native English speakers studied Chinese for 6 weeks. *
- fMRI scans showed an increase in connectivity in successful learners compared to less successful learners.
- Those who learned the fastest had more connectivity at the start.


Learning a new language, natural or artificial, rewires the brain.

* Ping Li (2014) <http://medicalxpress.com/news/2014-11-languages-workout-brains-young.html>

Signatures of Consciousness

Stanislas Dehaene on introspection:

- People can't observe their own brains.
- But introspection is data about experience.
- At the right is an fMRI scan of Dehaene's own brain while he was reading concrete words.
- He calls the patterns of neural activity *signatures of consciousness*. *
- “These signatures are remarkably stable and can be observed in a great variety of visual, auditory, tactile, and cognitive stimulations.” (p. 13)



Example: Signatures for numbers and computation.

- Mathematicians treat numbers and arithmetic as a unified system.
- But people learn, use, and experience numbers in a variety of ways: verbal, computational, spatial (diagrams), and temporal (counting).
- By examining the neural signature, the experimenters can distinguish which version a subject is thinking about.

* Stanislas Dehaene (2014) *Consciousness and the Brain*, New York: Viking.

2. Hybrid Systems to Support Diversity

Flexibility and generality are key to intelligence.

- **The languages of our stone-age ancestors can be adapted to any subject: science, technology, business, law, finance, and the arts.**
- **When people invent anything, they find ways to describe it.**
- **When people in any culture adopt anything from another culture, they borrow or adapt words to describe it in their native language.**

Minsky's proposal: A society of heterogeneous agents:

“What magical trick makes us intelligent? The trick is that there is no trick. The power of intelligence stems from our vast diversity, not from any single, perfect principle. Our species has evolved many effective although imperfect methods, and each of us individually develops more on our own. Eventually, very few of our actions and decisions come to depend on any single mechanism. Instead, they emerge from conflicts and negotiations among societies of processes that constantly challenge one another.” *

* Marvin Minsky (1986) *The Society of Mind*, New York: Simon & Schuster, §30.8.
See also Push Singh & Marvin Minsky (2004) [An architecture for cognitive diversity](#).

Machine Learning Applications

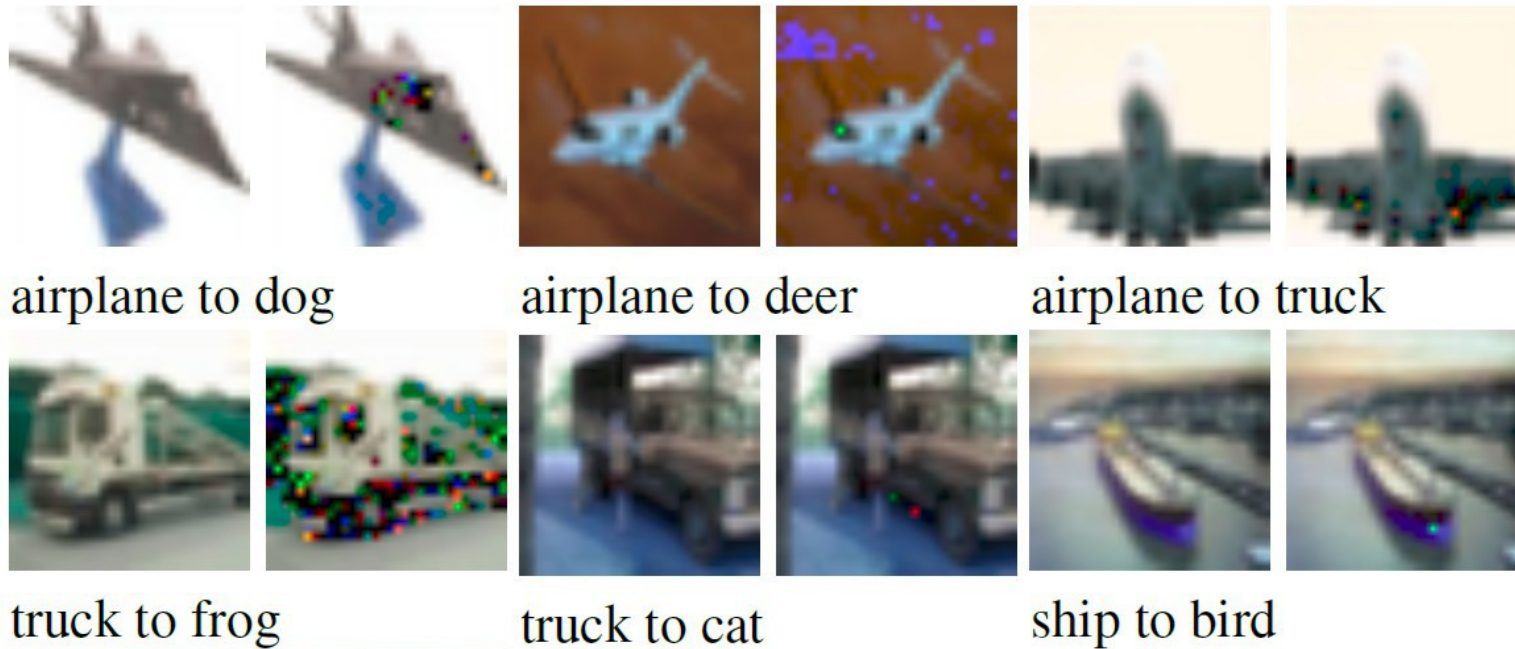
| INPUT A | RESPONSE B | APPLICATION |
|--|------------------------------------|------------------------|
| Picture | Are there human faces? (0 or 1) | Photo tagging |
| Loan application | Will they repay the loan? (0 or 1) | Loan approvals |
| Ad plus user information | Will user click on ad? (0 or 1) | Targeted online ads |
| Audio clip | Transcript of audio clip | Speech recognition |
| English sentence | French sentence | Language translation |
| Sensors from hard disk, plane engine, etc. | Is it about to fail? | Preventive maintenance |
| Car camera and other sensors | Position of other cars | Self-driving cars |

Deep neural nets can be very useful, but they're not magic. *

- **Observation by Andrew Ng: Current ML methods automate tasks that take less than one second of mental effort by humans.**
- **They only automate perception, not cognition (understanding).**
- **Cognition requires the ability to explain what was learned.**

* Andrew Ng (2016) <https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now>

Slight Perturbations



Random noise or perturbations may cause a misclassification.

- Six pairs of images show the originals and slightly modified versions.
- A DNN that correctly classified the first member of each pair made serious mistakes in classifying the modified versions. *
- Such mistakes in a self-driving car could cause a disaster.

People and other animals use cognition to correct perception.

- If they see something unexpected, they blink and look again.

* Huang et al. (2017) [Safety Verification of Deep Neural Networks](#)

Take Advantage of Available Tools

Sixty years of R & D in AI and machine translation.

Tools and resources for a wide variety of paradigms:

- **Parsers and translators for natural and artificial languages.**
- **Grammars, lexicons, ontologies, terminologies, corpora, Wikipedia, DBpedia, Linked Open Data, and the Semantic Web.**
- **Theorem provers and inference engines for formal logic and many kinds of informal and fuzzy reasoning.**
- **Qualitative, case-based, and analogical reasoning.**
- **Statistical, connectionist, and neural network methods.**
- **Pattern recognition, data mining, and graph data mining,**
- **Genetic algorithms and machine-learning methods.**
- **Thousands of implementations of all the above.**

But many systems are designed around a single paradigm – they do not take advantage of all the available resources.

Google Translate

Based on statistical methods for matching strings (N-grams).

String matching is good for short sentences:

English source: *The electrician is working.*

German: *Der Elektriker arbeitet.*

Polish: *Elektryk pracuje.*

English source: *The telephone is working.*

German: *Der Telefon funktioniert.*

Polish: *Telefon działa.*

But it can't keep track of long-distance connections: *

English source: *The electrician that came to fix the telephone is working.*

German: *Der Elektriker, der das Telefon zu beheben kam funktioniert.*

Polish: *Elektryk, który przyszedł naprawić telefon działa.*

English source: *The telephone on the desk is working.*

German: *Das Telefon auf dem Schreibtisch arbeitet.*

Polish: *Telefon na biurku pracuje.*

* Ernest Davis & Gary Marcus (2015) [Commonsense Reasoning and Knowledge in AI.](#)

Translating Latin to English

For Latin, word inflections are more important than word order.

If a text is very common, Google may find enough English-Latin copies for high quality statistical translation:

- Latin source: *E pluribus unum.*
- English: *Out of many, one.*

- Latin source: *Credo in unum Deum, Patrem omnipotentem, factorem coeli et terrae, visibilium omnium et invisibilium.*
- English: *I believe in one God, the Father, the Almighty, maker of heaven and earth, and of all things visible and invisible.*

Google sometimes uses, sometimes confuses punctuation:

- Latin source: *Gaudeamus, igitur, juvenes dum sumus.*
- English: *Let us, therefore, while we are young.*
- If the second comma is deleted, “*Gaudeamus, igitur ...*” is translated “*Us rejoice, therefore ...*”
- These translations were made in November 2016. The results may change with updates to the algorithms or the language data.

Translating Unfamiliar Texts

An epigram by Martial (Book VIII, No. LXI) and a human translation: *

*Livet Charinus, rumpitur, furit, plorat
Et quaerit altos, unde pendeat, ramos:
Non iam quod orbe cantor et legor toto,
Nec umbilicis quod decorus et cedro
Spargor per omnes Roma quas tenet gentes:
Sed quod sub urbe rus habemus aestivum
Vehimurque mulis non, ut ante, conductis.
Quid imprecabor, o Severe, liventi?
Hoc opto: mulas habeat et suburbanum.*

*Charinus is livid with envy, bursting, raging, weeping,
and he seeks high branches from which to hang himself,
not because I am acclaimed and read throughout the entire globe,
or because, ornamented with bosses and with cedar oil
I am spread over all the nations under Rome's sway,
But because in the suburbs I have a summer country place
and I am drawn by mules that are not hired, as before.
What curses, o Severus, shall I call down upon this envious man?
This I wish: that he have mules and a suburban house.*

Google produces unintelligible word hash:

*Livet health, bursting with rage, weep
He asks the high, which is bound up in the branches of:
And read that the world was no longer the whole,
Neither the muscles as well, and cedar
I will scatter through them all at Rome, which are bound to:
But the reason why we have the summer to the country under the city,
Vehimurque mules no longer, as before, they were hired.
What imprecabor o Severe, liventi?
This said, I would: mule, for it, and in the suburbs.*

* By Cora Sowa, <http://www.minervaclassics.com/quotat16.htm#Jul16>

Google Neural Machine Translation

Google upgraded their MT system with neural nets (GNMT). *

- In December 2016, they made a major improvement for 8 language pairs.
- For each pair of languages, neural nets match syntactic patterns better than the old word strings (N-grams).
- That upgrade improved the sentence-level translation.
- But the paragraph-level translation is just as bad.

Limitation: GNMT doesn't use semantics or pragmatics:

- People relate language to imagery, purpose, and feelings.
- They show their understanding by explaining what they know.
- 3-year-old children can explain what they say and why.
- But Google's neural networks cannot explain what they translate.
- They generate better syntax at the sentence level, but they cannot recognize and use paragraph-level context and intentions.

* Wu et al. (2016) [Google's neural machine translation system](#).

Lewis-Kraus (2016) [The great AI awakening](#).

Good Syntax. Bad Semantics.

A text translated by Google (24 March 2017):

From a novel: “She spoke earnestly, excitedly; eagerly he hung upon her words. Then her voice broke.” Satirical comment: “And down he went.” *

**Spanish: Ella habló con seriedad, con entusiasmo; ansiosamente que colgaba de sus palabras. Entonces su voz se quebró. Y abajo que iba.
Back to English: She spoke earnestly, with enthusiasm; Anxiously hanging from his words. Then his voice cracked. And down I was going.**

**German: Sie sprach sehr aufgeregt; Eifrig hing er auf ihre Worte. Dann brach ihre Stimme. Und er ging hinunter.
Back to English: She spoke very excitedly; Eagerly he hung his words. Then her voice broke. And he went downstairs.**

**Chinese: 她认真说着，激动地；他急切地挂在了她的话。
然后，她的声音哽咽了。上下他去。**

Back to English: She was so excited that he was eager to hang her words. Then, her voice choked. He went up and down.

Exercise: Try this example with other language pairs.


* Original text and comment from the *Altoona Tribune*, 29 December 1919, p. 8.

Multi-Layer Neural Nets

Deep Learning



Geoffrey Hinton: co-inventor of the backpropagation and contrastive divergence



Biomimicry of the human cortex


Deep neural nets (DNNs) are much better than earlier NNs.
But cognitive learning requires symbolic methods. *

* Wermter & Sun (2000) Hybrid Neural Systems, <http://www.cogsci.rpi.edu/~rsun/intro-w.pdf>

Learning to Play Games




Pong




Breakout



Space Invaders



Seaquest



Beamrider

Using DNNs to learn how to play games for the Atari 2600: *

- Seven games: Pong, Breakout, Space Invaders, Seaquest, Beamrider, Enduro, and Q*bert.
- No prior knowledge about objects, actions, features, or game rules.
- Bottom layer starts with pixels: 210 x 160 video and the game score.
- Each layer learns features, which represent the data at the next layer.
- Top layer determines which move to make at each step of the game.

Shows that DNNs can be used to learn time-varying patterns.

* Mnih et al. (2013) at DeepMind Technologies, <http://www.cs.toronto.edu/~vmnih/docs/dqn.pdf>

Developments at DeepMind

Good for perceptual learning, not cognitive learning.

Results on the Atari games:

- **Outperforms all other machine-learning methods on 6 of the 7 games.**
- **Better than a human expert on Breakout, Enduro, and Pong. Close to human performance on Beamrider.**
- **But far from human performance on Q*bert, Seaquest, and Space Invaders — because those games require long-term strategy.**



Comparison with Samuel's checker-playing system:

- **Modern DNNs are much more powerful than a perceptron.**
- **But lookahead methods are necessary for long-term strategy.**
- **Samuel's system was a hybrid that combined alpha-beta search with a learning method that was much simpler than a DNN.**
- **For chess, a hybrid that combined a DNN to learn the evaluation function with alpha-beta search reached the international master level. ***

* Michael Lai (2015) <http://arxiv.org/pdf/1509.01549.pdf>

Games of Go and Go-moku

Same syntax, but very different strategy:




Syntax defines legal moves, but not meaningful moves.
The meaning of any move is determined by its purpose.
In Go, the goal is to place stones that surround territory.
In Go-moku, the goal is to place five stones in a row.
Different goals change the strategy from the first move.

AlphaGo by DeepMind

Major breakthrough in learning to play Go.

- Computationally, the game of Go is far more challenging than chess.
- For both games, pattern recognition and tree search are important for high-level play.
- But Go has about 250 options at each step, and chess has about 37.
- The search strategies used for chess programs are inadequate for playing a good game of Go.



AlphaGo won 9 of 10 games with Go masters. *

- A hybrid system with DNNs and Monte Carlo Tree Search (MCTS).
- But the DNNs were trained on millions of games, far more than any human could play in a lifetime.
- And the 19x19 Go board is much simpler than the scenes and imagery that determine the semantics of natural languages.

* David Silver, et al. (2016) [Mastering the game of go with deep neural networks and tree search](#), *Nature*, vol. 529, pp. 484–489.

Cognitive Learning

With perceptual learning, ML systems can outperform humans.

- They can beat the world champions in chess, go, and poker.
- They can recognize signs and traffic patterns faster than human drivers.
- But they need to be trained on millions of examples – more than any human can experience in a lifetime.
- And pattern recognition must be supplemented with reasoning methods, such as search strategies (for chess and go) or statistics (for poker).

In cognitive learning, the agent understands what it learns.

- By analysis, people can learn new patterns from a single example.
- By analogy, they can relate patterns from different sensory modalities.
- Cognition uses “common sense”: it enables an agent to detect unusual situations, to correct errors in perception, and to explain what it learns.

Active ML is the first step toward cognitive learning. *

- The agent can choose the information from which it learns.
- Next step: Explain why it made a choice and what it learned.

* <http://burrsettles.com/pub/settles.activelearning.pdf>

The Role of Knowledge in Perception

Most animals are camouflaged with colors and features that blend with their native environment.

Where is the cat?

Prior knowledge enables faster, more accurate perception.

Neural networks, even DNNs, use bottom-up methods.



For better performance, they should be supplemented with top-down, knowledge-based methods.

The Role of Knowledge in Perception

The white oval directs attention to the cat.

But prior knowledge about cats is also important.

In science, new discoveries enable observers to “see” patterns they had overlooked.

Chromosomes, for example, were discovered and named in the 1880s. Before then, no drawings of cells showed chromosomes. Afterwards, they all showed chromosomes.



Stanford NLP Group

Developing statistical-symbolic hybrids. *

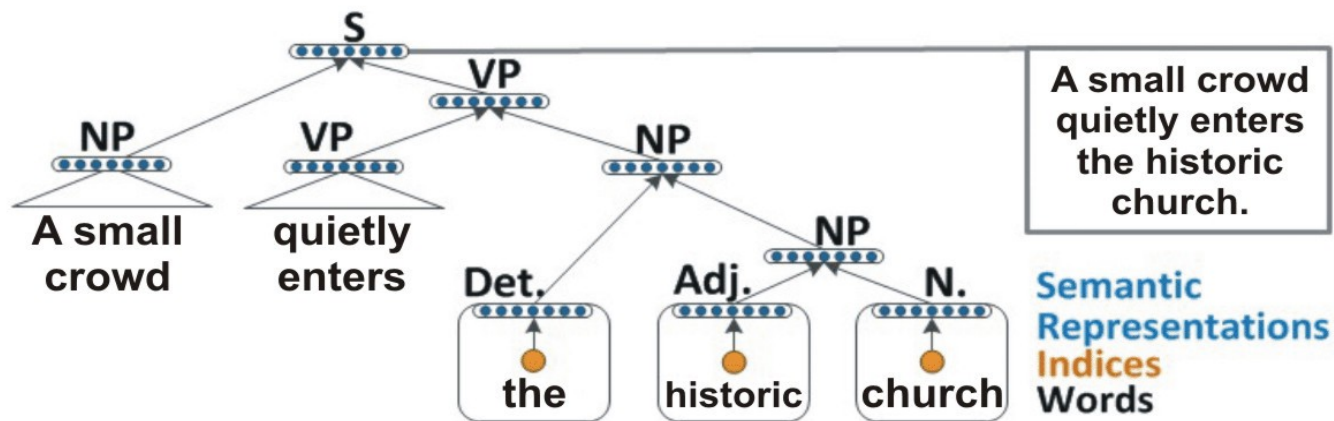
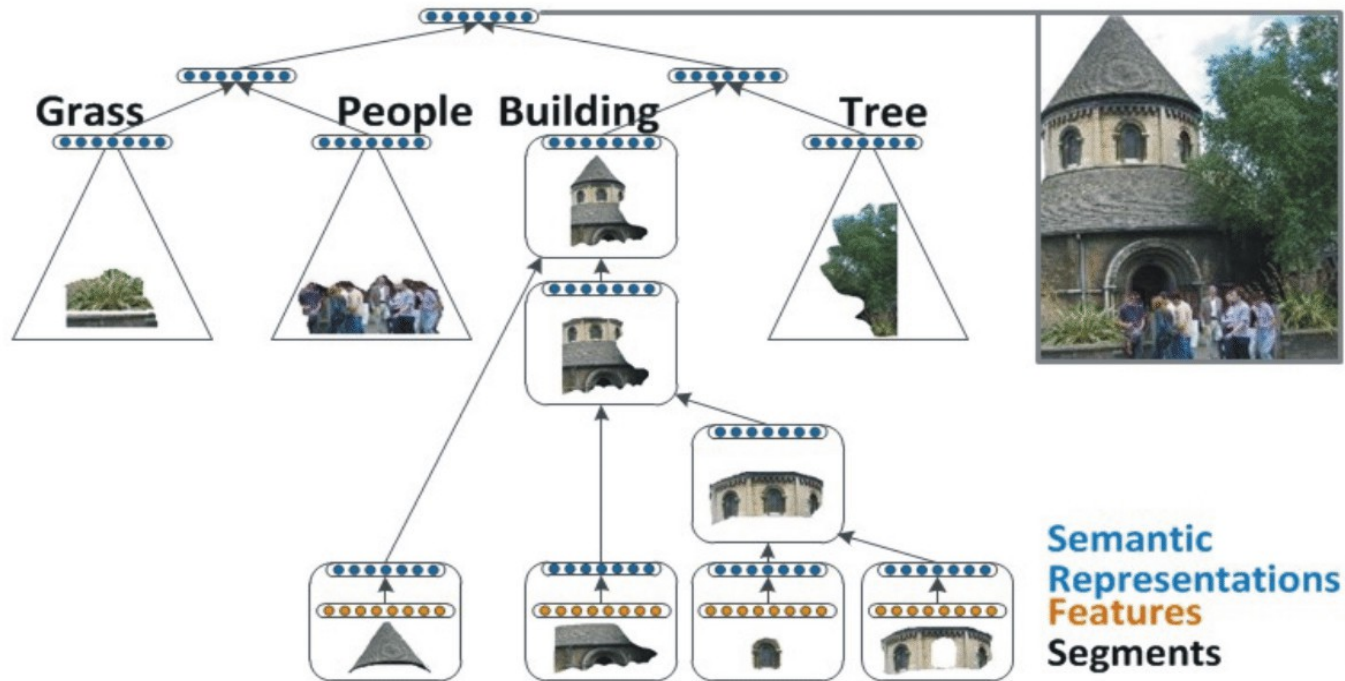
- Statistical methods for computing parse trees.
- DNNs for recognizing images, parse trees for images, and parse trees for language that describes the images.
- Statistical methods for representing word meaning in vectors and semantic graphs that relate the vectors.
- Logical inferences to derive the implications.
- Methods for translating texts to graphs, generating visual scenes from the graphs, and using the scenes for retrieving images.

Hybrid methods have produced promising results.


But more research is needed to generalize and systematize the methods for relating heterogeneous paradigms.

* Chris Manning (2015) [Computational linguistics and deep learning](#).

A Hybrid with DNNs and Parse Trees



Learning Language as a Child



Syntagmatic-Paradigmatic Learner (SPL). *

- Learns to map language to and from situation descriptions.
- The box on the left represents the word *ball*, which was spoken in the two situations described by the boxes on the right.

* Designed by Barend Beekhuizen (2015), *Constructions Emerging*, [PhD dissertation](#).

Learning to Map Language to Situations


SPL learns to relate strings of words to discourse situations:

- The notation is based on Langacker's theory of cognitive grammar.
- SPL is trained with strings of words paired with diagrams of relevant situations (as in the diagram in the previous slide).
- In the early stages, SPL learns to map single words to appropriate nodes of a situation description.
- Later, it maps two-word and multi-word constructions of words to larger subgraphs of the situation descriptions.
- It also learns reverse mappings from situations to strings of words.
- Appropriate rewards train SPL to learn correct mappings.

But some human chose what features are significant and should be included in the description:


- That choice reflects human feelings, values, and intentions.
- Current DNNs do not recognize human intentions, but an infant very quickly learns to recognize and respond to them.
- A truly deep learning system must also recognize them.

Geometry Problem Solver (GeoS)



In the figure to the left, triangle ABC is inscribed in the circle with center O and diameter AC. If $AB=AO$, what is the degree measure of angle ABO?

- (A) 15°
- (B) 30°
- (C) 45°
- (D) 60°
- (E) 90°



In the figure to the left, CDE is an equilateral triangle and ABCE is a square with an area of 1. What is the perimeter of polygon ABCDE?

- (A) 4
- (B) 5
- (C) 6
- (D) 7
- (E) 8



GeoS solves typical problems on the Geometry SAT exam. *
It's a hybrid that relates imagery to language and mathematics.

* Developed by the [Allen Institute for AI and the University of Washington](#).

Cyc Project

The most ambitious attempt to build the HAL 9000:

- Cyc project founded by Doug Lenat in 1984.
- Starting goal: Implement the background knowledge of a typical high-school graduate.
- Ultimate goal: Learn new knowledge by reading textbooks.

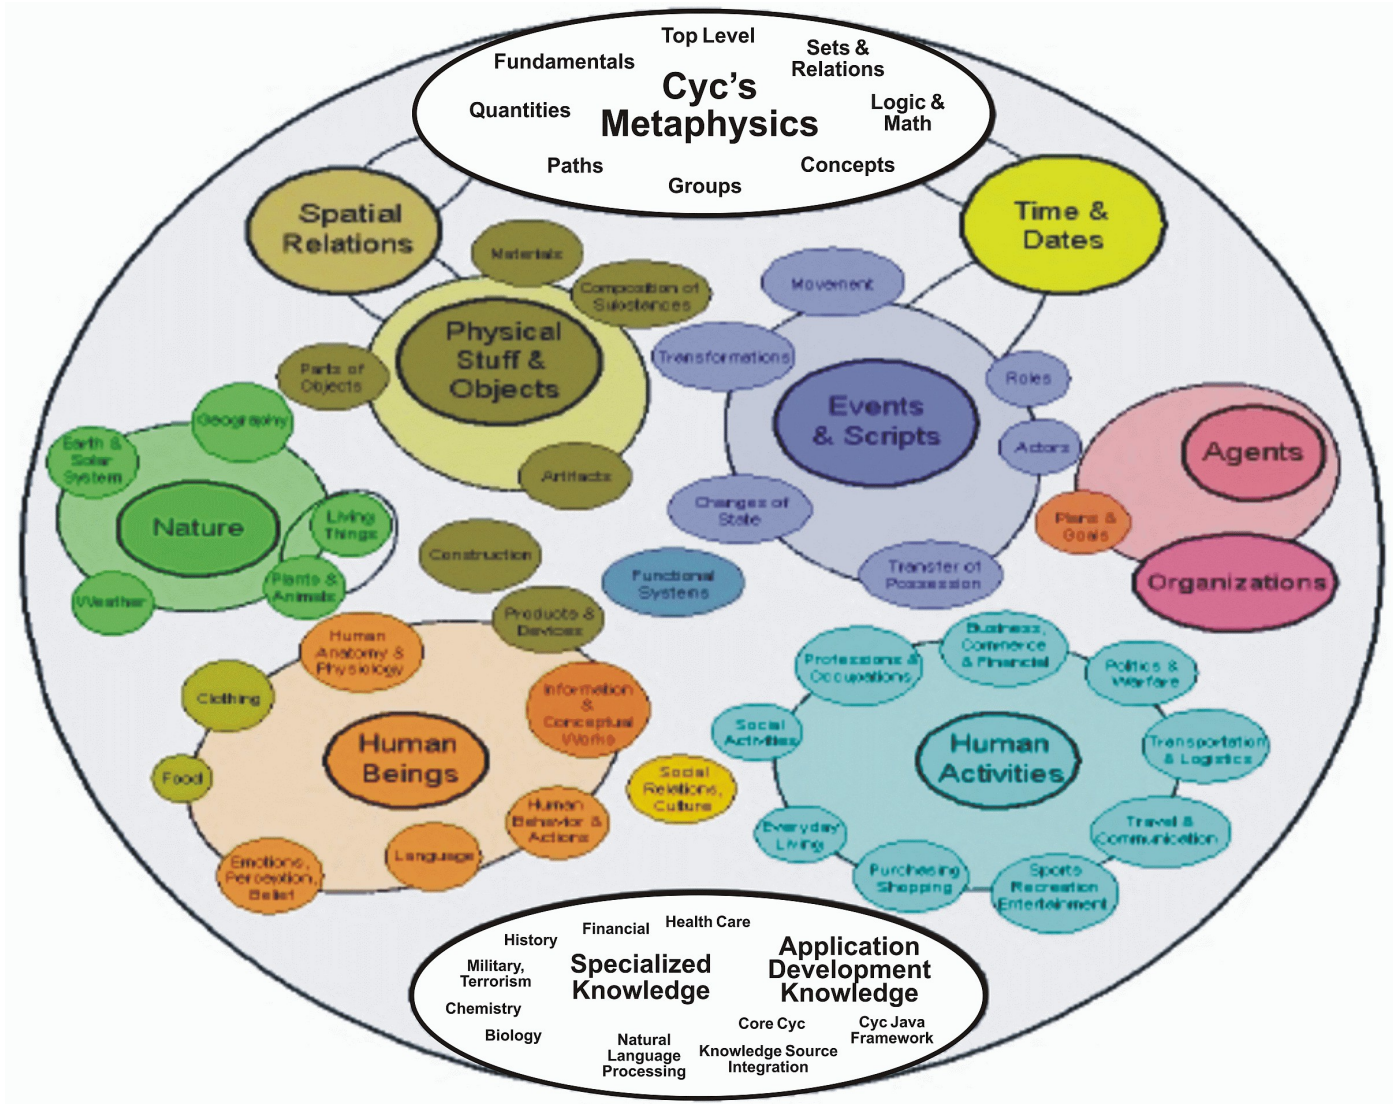
After the first 25 years,

- 100 million dollars and 1000 person-years of work,
- 600,000 concepts,
- Defined by 5,000,000 axioms,
- Organized in 6,000 microtheories.

Some good applications, but more needs to be done:

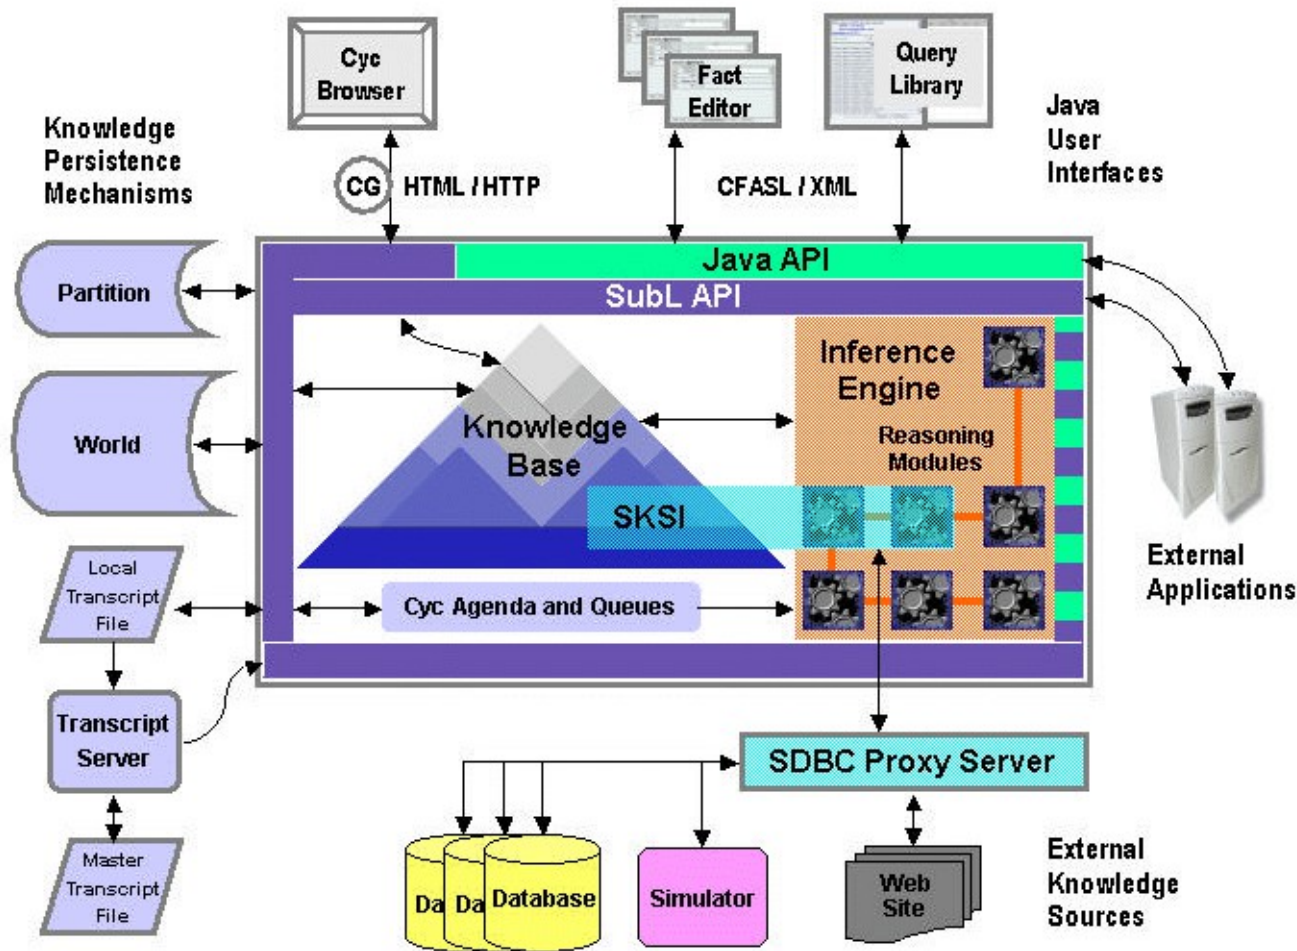
- Cyc cannot yet learn by reading a textbook.
- Cyc cannot understand language as well as a child.

Cyc Ontology



World's largest formal ontology – with controlled English as an option.
 See <http://www.cyc.com> and <http://opencyc.org>

Cyc System



A hybrid with a formal ontology, knowledge base, and NLP. *

- Logical foundation for ontology and multiple reasoning modules.
- Can use a variety of structured and unstructured data.

* See https://www.academia.edu/16911744/Common_Sense_Reasoning_From_Cyc_to_Intelligent_Assistant

Case Study: Cyc and IBM Watson

Why did IBM, not Cyc, beat the Jeopardy! champion?

Short answer: Cyc was not designed for game shows.

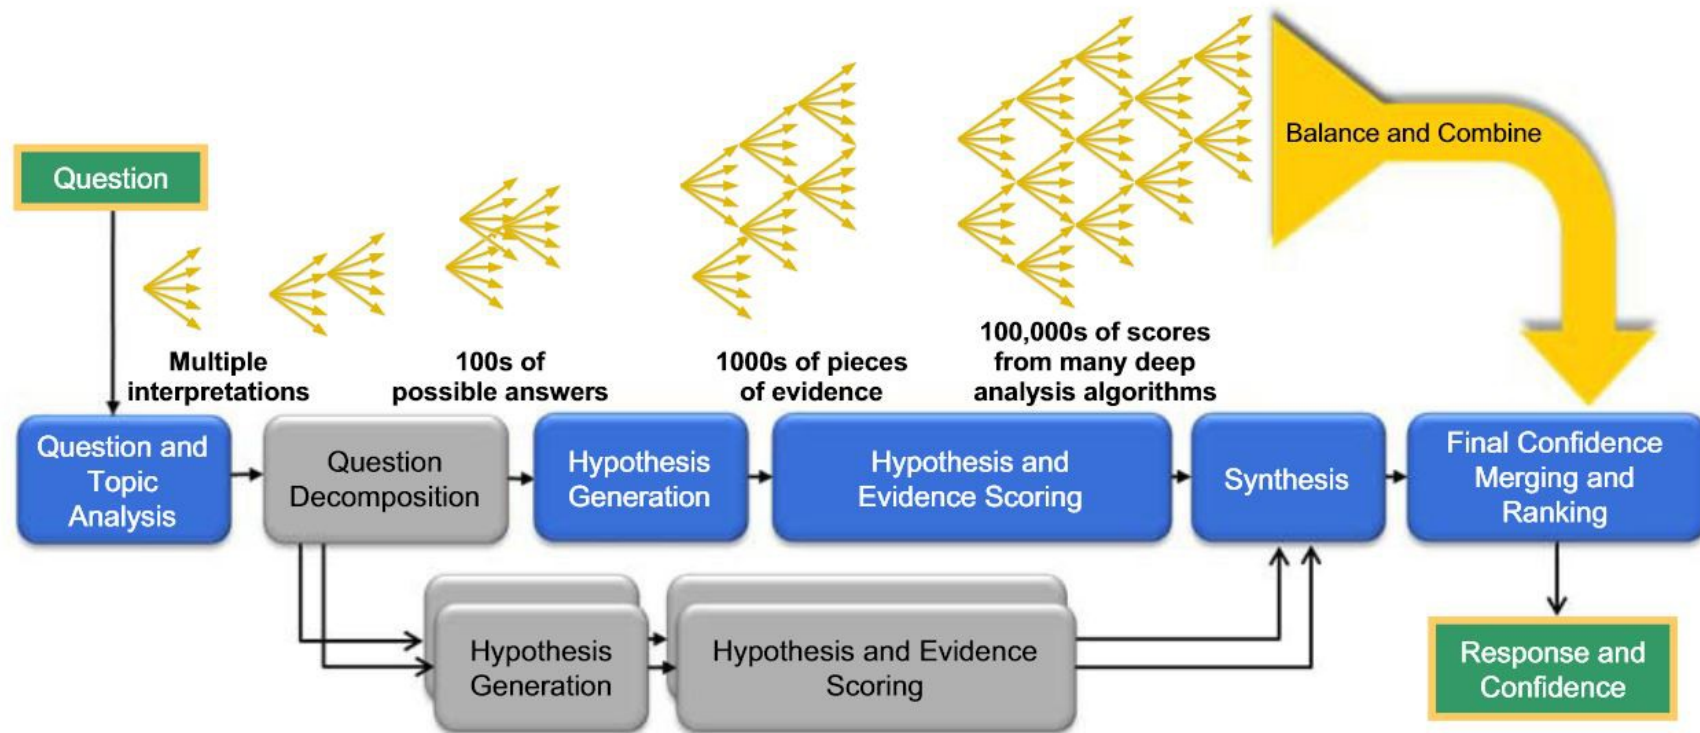
- Cyc was designed to be a general-purpose intelligent assistant.
- But IBM devoted a large research team to a single task.

Longer answer: Watson had more diversity than Cyc.

- As Minsky said, the number of reasoning methods is unlimited.
- The first version of Watson, which performed poorly on Jeopardy!, used 6 different reasoning algorithms.
- The version that won the Jeopardy! challenge used about 100 algorithms optimized for different kinds of questions and data.
- The “dramatically different environment” led the Watson team to design a framework with novel methods of machine learning. *
- And it ran on a supercomputer with 2,880 parallel threads.

* Gondek et al, A framework for merging and ranking answers in DeepQA,
http://researcher.watson.ibm.com/researcher/files/us-heq/W%2816%29%20ANSWERS%20MERGING_RANKING%2006177810.pdf

IBM Watson




Multiple paradigms and a growing number of modules. *

- For Jeopardy!, one API and about 100 reasoning methods.
- Now, a few dozen APIs and growing.
- Versions of all the major paradigms for natural language processing, statistical, symbolic, semantic, pragmatic. **
- But in language learning, nothing can compete with a 3-year-old child.

* Rob High, <http://www.redbooks.ibm.com/redpapers/pdfs/redp4955.pdf>

** Zadrozny, de Paiva, & Moss, <https://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/viewFile/9905/9684>

Finding Associations



The kinds of associations are highly context dependent.

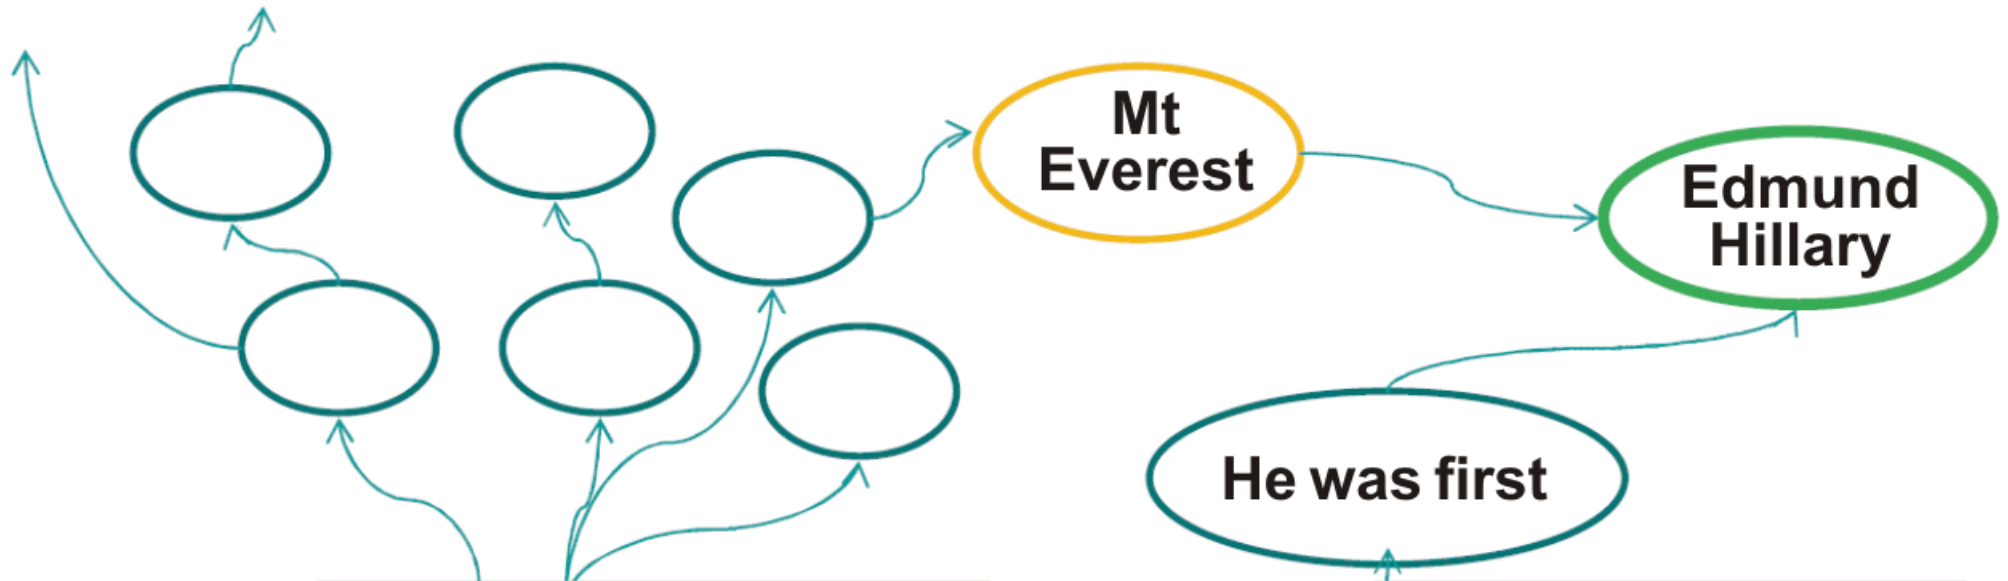
Different algorithms for different kinds.

- Indexes for finding co-occurrences in the data.
- Searching a network to find shortest paths.
- Deduction from definitions and axioms.

But Watson chose climate instead of religion as the context:

- “This kind of meat should not be shipped to Iraq.”
- “What is reindeer?”

Generating a Response for Jeopardy!



On hearing of the discovery of Malory's body, he told reporters he still thought he was first.

Multiple steps that use a variety of algorithms:

1. Parse the question and analyze the relationships among key phrases.
2. Generate hypotheses, find evidence for them, and estimate their quality.
3. Combine the best hypotheses in possible answers.
4. Rank the answers by a confidence measure.
5. Select the best one and respond "Who is Edmund Hillary?"
6. Use feedback about success or failure for "dynamic learning."

Similar methods are used in the **IBM Engagement Adviser**.

IBM Watson for Applications

A hybrid with multiple paradigms.

Scenario-based system for reasoning and Q/A about various applications.

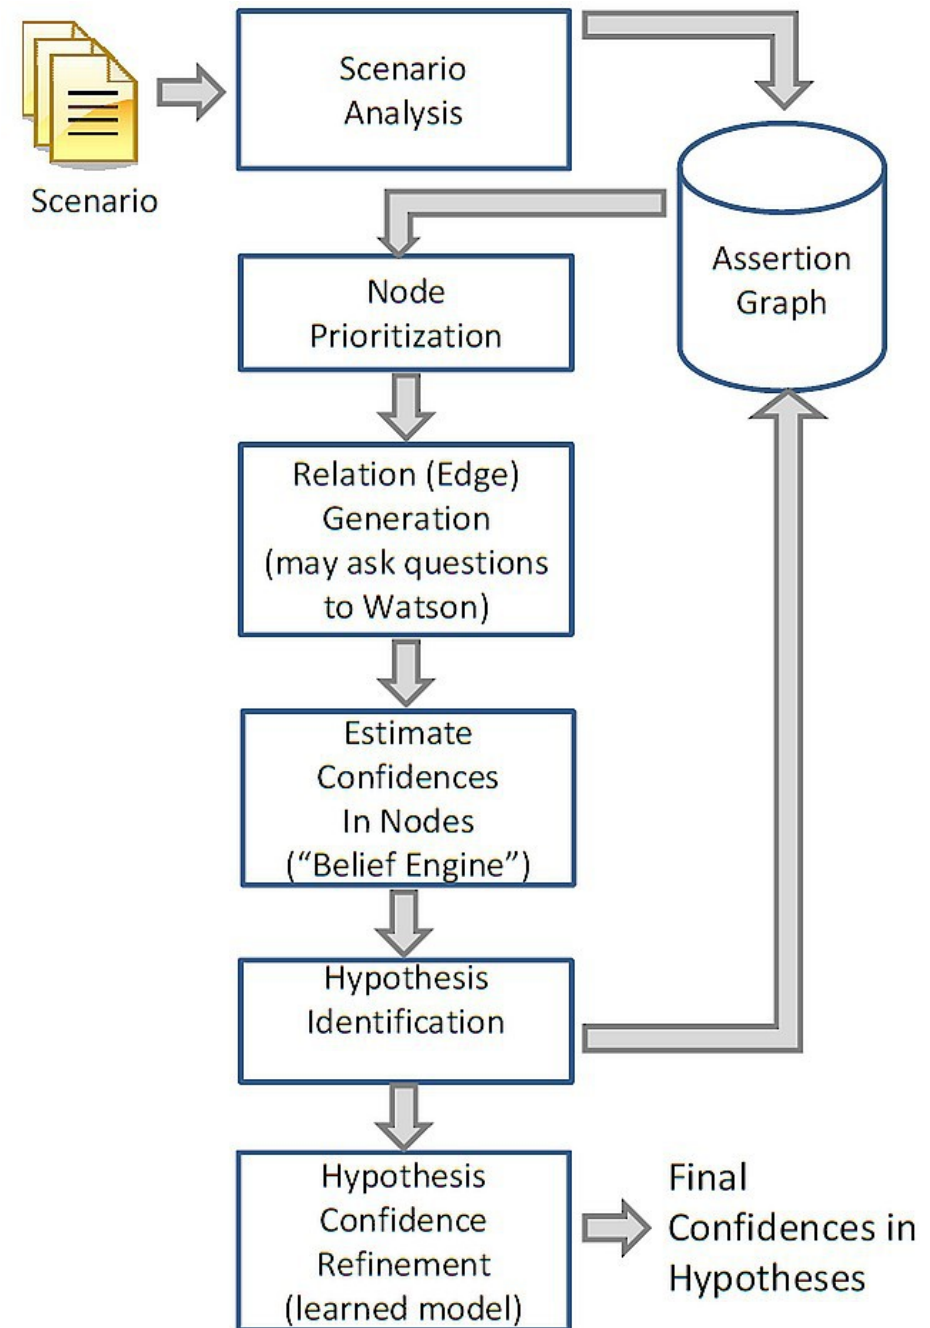
The input scenario describes some situation, e.g. a patient's symptoms.

The first step translates the input to an assertion graph.

Instead of answering one question, as in Jeopardy!, Watson Paths does extended reasoning to generate a more complex assertion graph.

The reasoning may continue until the assertion graph satisfies some task-dependent criteria.

See Lally et al. (2014) [Watson Paths](#).



3. Cognitive Computing

Support learning at a deeper, more human level than DNNs.

Cognitive Memory™ can use knowledge from any source:

- **Associative retrieval from background knowledge in $\log(N)$ time.**
- **Approximate pattern matching for analogies and metaphors.**
- **Precise pattern matching for logic and mathematics.**
- **Hybrid systems that can use any or all methods, including DNNs.**

Analogies support informal, case-based reasoning:

- **Long-term memory stores large numbers of previous experiences.**
- **Any new case is matched to similar cases in long-term memory.**
- **Close matches are ranked by a measure of semantic distance.**

The cycle of pragmatism can explain each step of reasoning:

- **Induction: Observe similarities to derive generalizations.**
- **Abduction: Use insights from any source to form hypotheses.**
- **Deduction: Derive a conclusion by formal or informal methods.**
- **Testing: Act upon the conclusion, evaluate the results, and repeat.**

Designing Hybrid Systems

In his *Society of Mind* and *Emotion Engine*, Minsky proposed a society of heterogeneous, interacting modules or agents.

- Could that society take advantage of recent developments in AI?
- How would the agents communicate and cooperate among themselves?
- What methods of learning and reasoning would they support?
- How could they relate symbols, images, and feelings to neural networks?
- How could a society of agents produce a unified personality?
- How could they agree on common goals, strategic plans, and tactical moves that support the strategy?
- Could they explain their results in a way that people could understand?

Requirements for supporting a society of agents:

- A system of communication and coordination. *
- Methods for sharing and using information in long-term memory. **
- A cycle of perception, learning, reasoning, acting, and evaluating.

* J. F. Sowa (2002) Architectures for intelligent systems. <http://www.jfsowa.com/pubs/arch.pdf>

** A. K. Majumdar & J. F. Sowa (2009) Two paradigms are better than one, and multiple paradigms are even better. <http://www.jfsowa.com/pubs/paradigm.pdf>

Kyndi Cognitive Architecture

A framework for perception, learning, reasoning, and acting.

Theory of signs (semiotic) by Charles Sanders Peirce.

- Natural language semantics represented in conceptual graphs.
- The full range of precision from ISO Common Logic to fuzzy sensations, vague tweets, and tentative guesses.
- Cycle of observation, induction, abduction, deduction, and action.

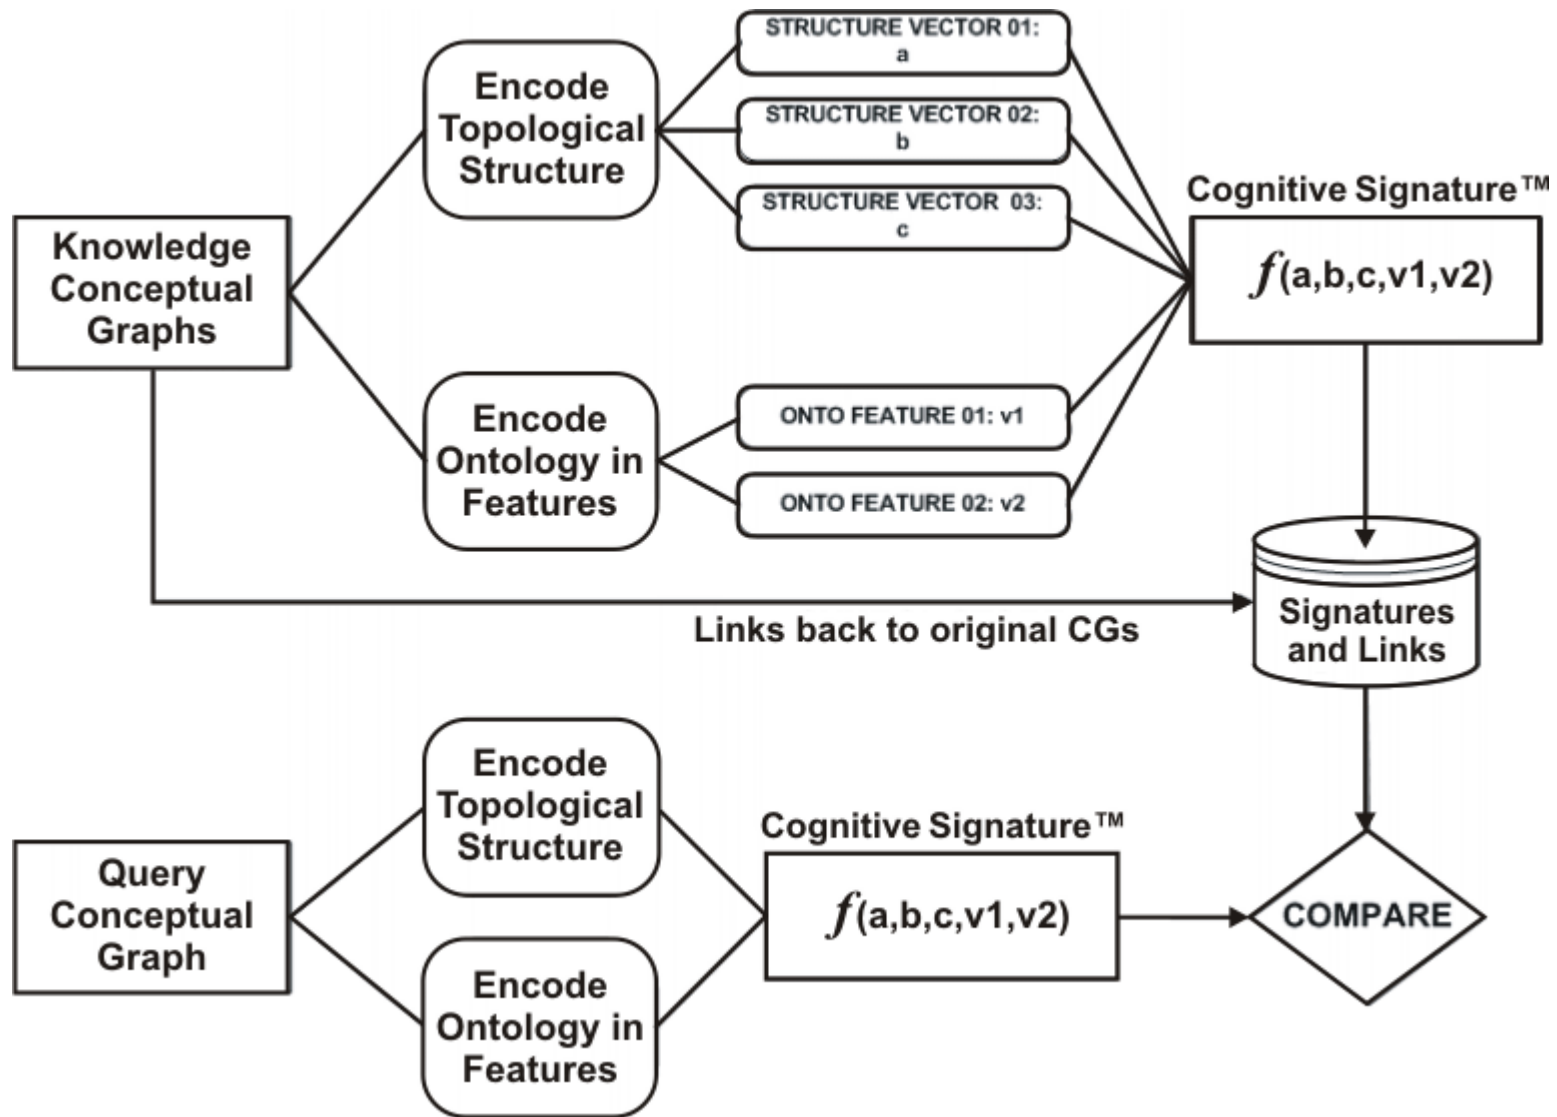
Cognitive Memory™ for associative retrieval of patterns.

- Networks of anything – signs, symbols, images, words, or texts.
- Finding exact or approximate matches in logarithmic time.

Market-Driven Learning™ (MDL) for dynamic evolution.

- Agents are organized in a managerial hierarchy.
- Patterns of activity can dynamically reorganize the hierarchy.
- The theoretical foundation is based on the π -calculus, Minsky's Society of Mind, McCarthy's Elephant 2000, and Gelernter's Linda.

Cognitive Memory™



Role of CM in Cognitive Computing

Cognitive Memory™ is essential for human-like cognition:

- Used in every stage of perception, language, reasoning, and action.
- Perception by matching new sensations to prior percepts.
- Simultaneous syntactic and semantic analysis of language.
- Approximate pattern matching for analogies and metaphors.
- Precise pattern matching for logic and mathematics.

Analogies support informal, case-based reasoning:

- Long-term memory (LTM) stores all previous experiences.
- Any new case is matched to similar cases in LTM.
- Close matches are ranked by a measure of semantic distance.

Formal reasoning is based on a disciplined use of analogy:

- Induction: Generalize multiple cases to create new rules or axioms.
- Deduction: Match (unify) a new case with part of some rule or axiom.
- Abduction: Form a hypothesis based on aspects of similar cases.

Any or all Kyndi modules can call upon CM at any step.

Applications

General approach:

- The following examples used earlier versions of Kyndi technology.
- For each one, clients specified requirements and paid for the results.
- The Kyndi modular structure enables rapid design and development by combining new modules with a library of base modules.

Three projects:

1. Extract information from research reports and map to a relational DB.
2. Legacy re-engineering: Analyze 40 years of legacy software and relate it to the documentation – manuals, reports, memos, and comments.
3. Oil and gas exploration: Extract information from research reports and answer English queries by a geologist.

For more detail, see the presentation on **Cognitive Memory™**:

- Using conceptual graphs (CGs) for finding analogies in NL texts.
- Application: Evaluating student answers in free-form English.
- Method: Translate the answers to CGs and use analogies.

Information Extraction Project

The next slide shows a table derived from research reports.

To extend the semantics, an ontology for chemistry was added to the basic Kyndi ontology.

Then for each report,

- Map each sentence to a conceptual graph (CG). *
- Analyze anaphoric references to link pronouns to named entities.
- The result is a large CG that represents every sentence in the document.
- Store that graph (including subgraphs) in Cognitive Memory.
- Query Cognitive Memory for the data in each row of the table.
- Store the answers in the table.

In a competition among twelve NLP systems,

- The Kyndi system got 96% of the entries correct.
- The second best score was 73%. Most scores were below 50%.

* For an overview of CG methods, see <http://www.jfsowa.com/pubs/template.pdf>

Information Extracted from Documents

DOE BREMS PROJECT.xlsx

Search in Sheet

Home Layout Tables Charts SmartArt Formulas Data Review

Edit Font Alignment Number Format Cells Themes

Calibri (Body) 24 abc Wrap Text General

Conditional Formatting Styles Insert Delete Format Themes

| | A | B | C | D |
|----|--|------------------------------|---|---|
| 1 | COMPOUND | CURIE TEMP. | SOURCE | |
| 2 | Mn ₃ [Cr(CN) ₆] ₂ · 16H ₂ O | 89 K | A solid-state hybrid density functional theory study | |
| 3 | Sr ₃ Ir ₂ O ₇ in Sr ₃ Ir ₂ O ₇ single-cr | ~ 280 K | Canted antiferromagnetic ground state in Sr ₃ Ir ₂ O ₇ | |
| 4 | PrPt ₂ B ₂ C | 6 K | Coexistence of superconductivity and magnetic ord | |
| 5 | La _{0.3} Nd _{0.7} Pt ₂ :1- | 2.8 K | Coexistence of superconductivity and magnetic ord | |
| 6 | NdPt ₂ :1B ₂ :4C ₁ :2 | 3 K | Coexistence of superconductivity and magnetic ord | |
| 7 | NdPt _{1.5} Au _{0.6} B ₂ C | 3 K | Coexistence of superconductivity and magnetic ord | |
| 8 | SmNiC ₂ | = 17.7 K | Commensurate charge-density wave with frustrate | |
| 9 | Co _{0.22} Zn _{0.88} Fe ₂ O ₄ . in CdxCo1- | ~ 780 K | Does Ti+4 ratio improve the physical properties of C | |
| 10 | Zn _{0.88} Co _{0.12} O in ZnO | ~ 540 K | Effect of Co doping on the structural; optical and m | |
| 11 | La in Sr _{2-x} LaxFeMoO ₆ | 425 K | Effect of La doping on the properties of Sr _{2-x} LaxFe | |
| 12 | Fe in Sr _{2-x} LaxFeMoO ₆ | ~ 1040 K | Effect of La doping on the properties of Sr _{2-x} LaxFe | |
| 13 | FeSe | ~ 305 K | Electronic and magnetic properties of FeSe thin film | |
| 14 | Ni-Mn-Ga | = 376 K | Electronic and structural properties of ferromagnet | |
| 15 | LaFexSi _{1-x} in La _{1-z} Prz(Fe) | ~ 190 K | Enhancement of magnetocaloric effects in La _{1-z} Prz | |
| 16 | LaFe _{0.88} Si _{0.12} in La _{1-z} Prz(| = 195 K | Enhancement of magnetocaloric effects in La _{1-z} Prz | |
| 17 | Co ₂ MnGa in Co ₂ MnGa | 600 K | Ferromagnetic resonance in Co ₂ MnGa films with va | |
| 18 | HoCrO ₄ in HoCrO ₄ | 17.6 K | Ferromagnetism vs. antiferromagnetism of the dim | |
| 19 | Mn ₃ (HCOO) ₆ in Mn ₃ (HCOO) ₆ | 8.0 K | Guest-induced chirality in the ferrimagnetic nanop | |
| 20 | NaZn ₁₃ - in La _{0.5} Pr _{0.5} (Fe _{0.88} | range from 195 K to 185 K | Large isothermal magnetic entropy change after hy | |
| 21 | La _{2/3} Ba _{1/3} MnO ₃ in La ₂₋₃ Ba ₁ | range from 300 K to 250 K | Magnetic and neutron diffraction study of La ₂₋₃ Ba ₁ | |
| 22 | CuMnSb in Co _{1-x} CuxMnSb | range from 476 K to 300 K | Magnetic properties of half-metallic semi Heusler C | |
| 23 | Nd ₂ in Nd _{2-y} DyyFe _{17-x} Six | range from 61.46 °C to 236 ° | Magnetic properties of iron-rich Nd _{2-y} DyyFe _{17-x} Six | |
| 24 | Tb ₂ Fe ₁₇ in Nd _{2-y} DyyFe _{17-x} S | ~ 80 °C | Magnetic properties of iron-rich Nd _{2-y} DyyFe _{17-x} Six | |

Sheet1

Application to Legacy Re-engineering

Analyze the software and documentation of a corporation.

Programs in daily use, some of which were up to 40 years old.

- **1.5 million lines of COBOL programs.**
- **100 megabytes of English documentation – reports, manuals, e-mails, Lotus Notes, HTML, and program comments.**

Goal:

- **Analyze the COBOL programs.**
- **Analyze the English documentation.**
- **Compare the two to determine:**
 - Data dictionary of all data used by all programs.**
 - English glossary of all terms with index to the software.**
 - Evolution of terminology over the years.**
 - Structure diagrams of the programs, files, and data.**
 - Discrepancies between programs and documentation.**

An Important Simplification

An extremely difficult and still unsolved problem:

- **Translate English specifications to executable programs.**

Much easier task:

- **Translate the COBOL programs to conceptual graphs (CGs).**
- **Those CGs provide the ontology and background knowledge.**
- **The CGs derived from English may have ambiguous options.**
- **In parsing English, use CGs from COBOL to resolve ambiguities.**
- **The COBOL CGs show the most likely options.**
- **They can also provide missing information or detect errors.**

The CGs derived from COBOL provide a formal semantics for the informal English texts.

Excerpt from the Documentation

The input file that is used to create this piece of the Billing Interface for the General Ledger is an extract from the 61 byte file that is created by the COBOL program BILLCRUA in the Billing History production run. This file is used instead of the history file for time efficiency. This file contains the billing transaction codes (types of records) that are to be interfaced to General Ledger for the given month.

For this process the following transaction codes are used: 32 — loss on unbilled, 72 — gain on uncollected, and 85 — loss on uncollected. Any of these records that are actually taxes are bypassed. Only client types 01 — Mar, 05 — Internal Non/Billable, 06 — Internal Billable, and 08 — BAS are selected. This is determined by a GETBDATA call to the client file.

Note that none of the files or COBOL variables are named.

By matching graphs derived from English to graphs derived from COBOL, all names of files and COBOL variables were determined.

Interpreting Novel Patterns

Many documents contain unusual or ungrammatical patterns.

They may be elliptical forms that could be stored in tables.

But some authors wrote them as phrases:

- *32 — loss on unbilled*
- *72 — gain on uncollected*
- *85 — loss on uncollected*

The dashes were represented by a default relation (Link):

[Number: 32]→(Link)→[Punctuation: “-”]→(Link)→[Loss]→(On)→[Unbilled]

This CG, which was derived from an English document, matched CGs derived from COBOL programs:

- The value 32 was stored as a constant in a COBOL program.
- The phrase “loss on unbilled” was in a comment that followed the value 32 in that program.

Results

Job finished in 8 weeks by Arun Majumdar and André LeClerc.


- **Four weeks for customization:**
Design, ontology, and additional programming for I/O formats.
- **Three weeks to adapt the software that used Cognitive Memory:**
Matches with strong evidence (close semantic distance) were correct.
Weak matches were confirmed or corrected by Majumdar and LeClerc.
- **One week to produce a CD-ROM with the desired results:**
Glossary, data dictionary, data flow diagrams, process architecture diagrams, system context diagrams, and list of errors detected.

A major consulting firm estimated that the job would take 40 people two years to analyze the documentation and find all cross references.

With Cognitive Memory, the task was completed in 15 person weeks.

Discrepancy Detected

A diagram of relationships among data types in the database:



Question: Which location determines the market?

- According to the documentation: Business unit.
- According to the COBOL programs: Client HQ.

For many years, management had been making decisions based on incorrect assumptions.

Contradiction Detected

From the ontology used for interpreting English:

- Every employee is a human being.
- No human being is a computer.

From analyzing COBOL programs:

- Some employees are computers.

What is the reason for this contradiction?

Quick Patch in 1979

A COBOL programmer made a quick patch:

- **Two computers were used to assist human consultants.**
- **But there was no provision to bill for computer time.**
- **Therefore, the programmer named the computers Bob and Sally, and assigned them employee ids.**

For more than 20 years:

- **Bob and Sally were issued payroll checks.**
- **But they never cashed them.**

The software discovered two computer “employees.”

Relating Formal and Informal CGs

The legacy-reengineering task required two kinds of processing.

Precise reasoning:

- **Analyzing the COBOL programs and translating them to CGs.**
- **Detecting discrepancies between different programs.**
- **Detecting discrepancies between programs and documentation.**

Indexing and cross references:

- **Creating an index of English terms and names of programs.**
- **Mapping English documents to the files and programs they mention.**

Conceptual graphs derived from COBOL are precise.

- **But CGs derived from English are informal and unreliable.**
- **Informal CGs are adequate for cross-references between the English documents and the COBOL programs.**
- **All precise reasoning was performed on CGs from COBOL or on CGs from English that were *corrected* by CGs from COBOL.**

Application to Oil and Gas Exploration

Source material:

- 79 documents, ranging in length from 1 page to 50 pages.
- Some are reports about oil or gas fields, and others are chapters from a textbook on geology used as background information.
- English, as written for human readers (no semantic annotations).
- Additional data from relational DBs and other structured sources.
- Lexical resources derived from WordNet, CoreLex, IBM-CSLI Verb Ontology, Roget's Thesaurus, and other sources.
- An ontology for the oil and gas domain written in controlled English by geologists from the University of Utah.

Queries:

- A paragraph that describes a potential oil or gas field.
- Analogies compare the query to the documents.

Answering Questions

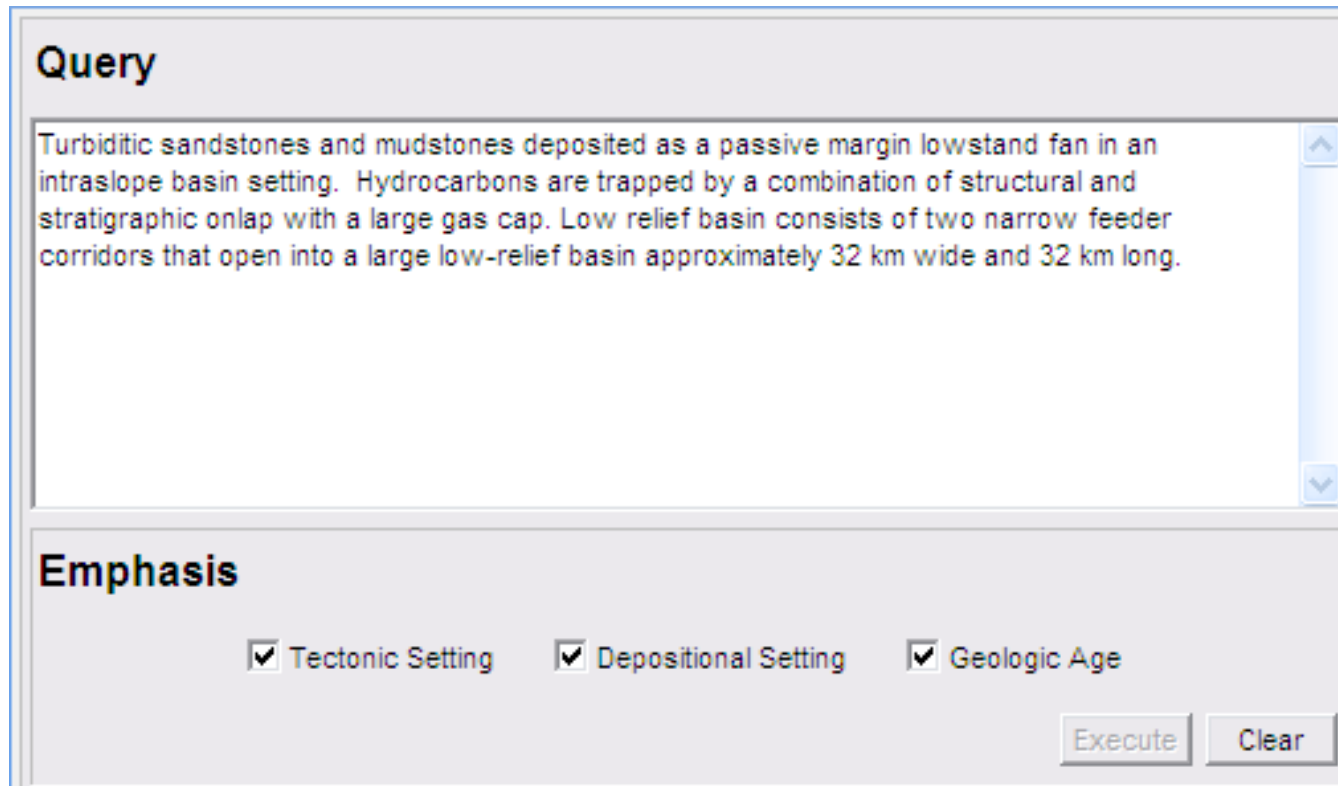
For the sources, either NL documents or structured data:

- Translate the text or data to conceptual graphs.
- Translate all CGs to Cognitive Signatures™ in time proportional to $(N \log N)$, where N is the total number of CGs.
- Store each Cognitive Signature in Cognitive Memory™ with a pointer to the original CG and the source from which that CG was derived.
- Use previously translated CGs to help interpret new sentences.

For a query stated as an English sentence or paragraph,

- Translate the query to conceptual graphs.
- Find matching patterns in the source data and rank them in order of semantic distance. The time is proportional to $(\log N)$.
- For each match within a given threshold, use structure mapping to verify which parts of the query CG match the source CG.
- As answer, return the English sentences or paragraphs in the source document that had the closest match to the query.

A Query Written by a Geologist



The image shows a software interface for writing a query. It has a title bar 'Query' and a text input area containing a geological description. Below the text area is an 'Emphasis' section with three checked checkboxes: 'Tectonic Setting', 'Depositional Setting', and 'Geologic Age'. At the bottom right are 'Execute' and 'Clear' buttons.

Query

Turbiditic sandstones and mudstones deposited as a passive margin lowstand fan in an intraslope basin setting. Hydrocarbons are trapped by a combination of structural and stratigraphic onlap with a large gas cap. Low relief basin consists of two narrow feeder corridors that open into a large low-relief basin approximately 32 km wide and 32 km long.


Emphasis

Tectonic Setting Depositional Setting Geologic Age

Execute Clear

Turbiditic sandstones and mudstones deposited as a passive margin lowstand fan in an intraslope basin setting. Hydrocarbons are trapped by a combination of structural and stratigraphic onlap with a large gas cap. Low relief basin consists of two narrow feeder corridors that open into a large low-relief basin approximately 32 km wide and 32 km long.

Statistics



Show:

Query Results

- 10) 17%- Vautreuil
- 23) 16%- Hogsnyta Type II Shelf Margin
- 25) 15%- Tanqua Karoo Subbasin
- 36) 15%- des
- 3) 14%- Espy Ranch, Spine 1, and Rattlesnake Ridge
- 8) 14%- Songpan-Ganzi Complex
- 19) 14%- Pukearuhe Beach
- 31) 11%- Waikiekie South Beach and Inland
- 2) 10%- Brushy Canyon Outcrop Belt
- 16) 10%- Atlapexco Road Cut
- 35) 10%- depocenter
- 22) 9%- Storvola Type 1 Shelf Margin
- 21) 8%- Punta Marone

Sort By:

Query Results: Analog Summary

[Source Visualization](#)

NAME : Vautreuil
 COUNTRY : France
 FORMATION : Gres d Annot Formation (Annot Sandstones)
 AGE : Eocene-Oligocene

Query Results: Evidence

vautreuil chapter 44 lomas, et. al. onlapping
 sheet sandstones in the gres d_annot, vautreuil, france cliffs
 forming the east side of the vautreuil de laverq (44?18-n;
 valley, west of the foret domaniale 6?29-e) region:
 provence-alpes-cote d_azur, departement: alpes-de-haute-provence
 france overview montage: 2700 m (8850 ft), detailed panel: 800 m

Preferences

Emphasis:

- On: Tectonic Setting
- On: Depositional Setting
- On: Geologic Age

Weights:

- On: Provenance
- On: Profile

Sources:

Corporate

- On: Exploration
- On: Production
- On: Financial

Vendor

- On: AAPG
- On: Wood

External

Details of the closest matching hydrocarbon fields

Vautreuil.ren

File Edit Interface Selection Arrange Display Clustering DetailControl Window Help

Left-click or drag to select; <Shift> to mod sel; drag to move; middle-click for info.; middle-click-<Shift> for contents info.

Mouse Mode

NEW!

Hold space bar and drag to pan, use mouse wheel to zoom.

Use right-click to get menu of contextual operations.

Memory: 65%

start 16 W... GA 4 Wi... 59. Ir... Windo... Fundi... UESStu... 4 No... Scree... 3 Ja... Windo... 12:53 PM

Turbiditic sandstones and mudstones deposited as a passive margin lowstand fan in an intraslope basin setting. Hydrocarbons are trapped by a combination of structural and stratigraphic onlap with a large gas cap. Low relief basin consists of two narrow feeder corridors that open into a large low-relief basin approximately 32 km wide and 32 km long.

NAME : Vautreuil
 COUNTRY : France
 FORMATION : Gres d'Annot
 Formation (Annot Sandstones)
 AGE : Eocene-Oligocene

00004: The Annot Sandstone (Gres d'Annot) of southeast France and its correlative deposits (e.g., the Champsaur Sandstone) form a widespread unit of lower Tertiary turbidites deposited in the Alpine foreland basin. This is an ideal system in which to characterize sandstone geometries developed against confining slopes, because the basin floor was bathymetrically complex, being divided into a series of discrete subbasins. This division is related to the development of a piggyback basin, and the Tertiary subbasins are interpreted as the surface expression of a thrust system within the underlying Mesozoic section. In the Maritime Alps, mild post depositional deformation and good exposure aid the characterization of pinch-out geometries at the margins of these subbasins. The outcrop studies detailed here focus on confining slopes preserved at the margins of the Annot and Peira Cava subbasins. Our analysis is divided into two sections: characterization of sandstone geometries developed against the confining slope and characterization of facies changes observed approaching the slope.

00006: The basin margin bounded the subbasin preserved around the village of Annot; intrabasin highs related to ramps in the underlying thrust system separated it from other subbasins. This subbasin contains at least two temporally distinct turbidite systems, of which the older Oligocene Braux system is included in this article. The Braux system constitutes a moderately sandy sheet complex, point-sourced in the east, that has a sand/shale ratio of about 2:1 overall. The section described in this article was deposited after earlier sandstones had buried the initial basin-floor topography, so the turbidity currents were able to expand across a relatively flat basin floor until confined by an east-northeast-dipping slope on the southwest side of the subbasin. This basin-margin slope provides an example of oblique frontal confinement. Its gradient before compaction has been estimated at about 12°.

Chapter 44.bt
 Vautreuil
 @QUERY:0
 CompositeEvidence
 Chapter 45.bt
 McCaffrey and Kneller_2001.bt
 evidence#6 : 0.98798

Linking the query to the paragraphs that contain the answer

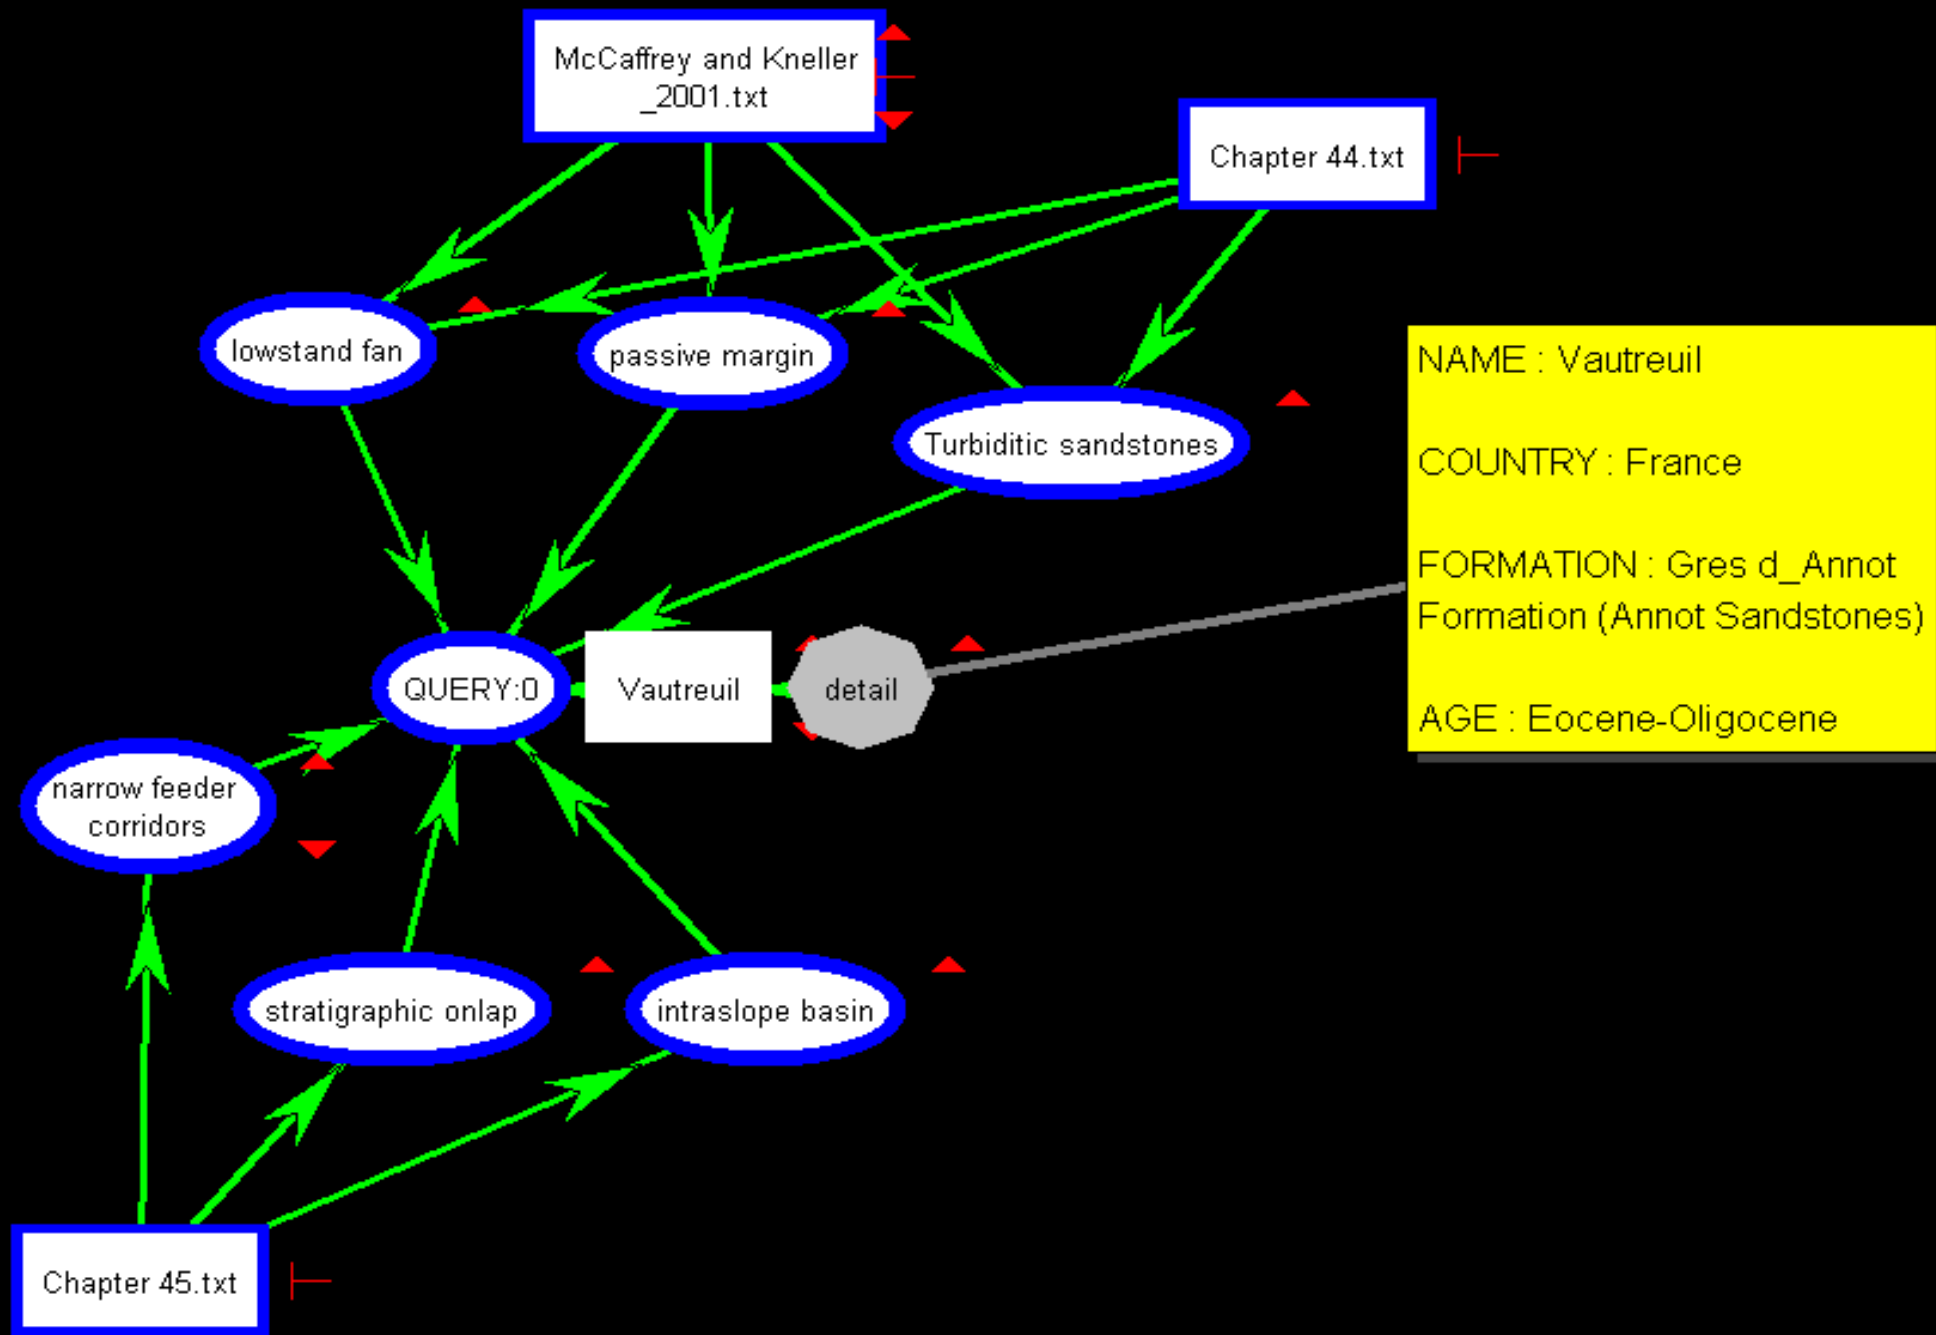
What the Screen Shots Show

Information shown in the previous screen shot:

- The query in the green box describes some oil or gas field.
- The data in the small yellow box describes the Vautreuil field.
- The large yellow box shows the paragraphs in a report by McCarthy and Kneller from which that data was extracted.

The next screen shot shows how the answer was found:

- Many terms in the query were not defined in the ontology: *lowstand fan, passive margin, turbiditic sandstones, narrow feeder cables, stratigraphic onlap, intraslope basin*.
- Generate tentative CGs for these phrases and look in Cognitive Memory to find similar CGs derived from other sources.
- Chapters 44 and 45 of the textbook on geology contained those CGs as subgraphs of larger graphs that had related information.
- Patterns found in the larger graphs helped relate the CGs derived from the query to CGs derived from the report that had the answer.



Using background knowledge from a textbook to find the answer

Emergent Knowledge

When reading the 79 documents,

- **Translate the sentences and paragraphs to CGs.**
- **But do not do any further analysis of the documents.**

When a geologist asks a question,

- **Look for related phrases in Cognitive Memory.**
- **To connect those phrases, further searches may be needed.**
- **Some sources may be textbooks with background knowledge that may help interpret the research reports.**
- **The result consists of CGs that relate the query to paragraphs in research reports that contain the answer.**
- **The new CGs can be added to Cognitive Memory for future use.**

By a “Socratic” dialog, a geologist can lead the system to explore novel paths and discover unexpected patterns.

Knowledge Discovery

Observation by Immanuel Kant:

Socrates said he was the midwife to his listeners, i.e., he made them reflect better concerning that which they already knew, and become better conscious of it. If we always knew what we know, namely, in the use of certain words and concepts that are so subtle in application, we would be astonished at the treasures contained in our knowledge...

Platonic or Socratic questions drag out of the other person's cognitions what lay within them, in that one brings the other to consciousness of what he actually thought.

From his *Vienna Logic*

We need tools that can play the role of Socrates.

They should help us discover the implicit knowledge and use it to process the huge volumes of digital data.

4. Cycles of Learning and Reasoning


Children learn language by starting with words and patterns of words that are grounded in perception and purposive action.

By trial and error, children and adults revise, extend, and adjust their beliefs to make better predictions about the world:

- **Observations generate low-level facts.**
- **Induction derives general axioms from multiple facts.**
- **A mixture of facts and axioms is an unstructured *knowledge soup*.**
- **Abduction selects facts and axioms to form a hypothesis (theory).**
- **Analogies may relabel a theory of one topic and apply it to another.**
- **Deductions from a theory generate predictions about the world.**
- **Actions test the predictions against reality.**
- **The effects of the actions lead to new observations.**

Cycles within cycles may be traversed at any speed – from seconds to minutes to research projects that take years.

Observing, Learning, Reasoning, Acting



The cognitive cycle, as described by Charles Sanders Peirce. Similar cycles occur in every aspect of life, including science.

Knowledge Soup

A heterogeneous, loosely linked mixture:

- **Fluid, lumpy, and dynamically changing.**
- **Many lumps are or can be structured in a computable form.**
- **But they may be inconsistent or incompatible with one another.**

In anybody's head, knowledge soup is

- **The totality of everything in memory.**

In the Internet, knowledge soup is

- **The totality of everything people downloaded from their heads, recorded automatically, or derived by any computable method.**

Linked Open Data is good for finding and classifying anything in the soup – whether loose items or structured lumps.

But understanding the contents of the LOD poses the same challenge as understanding natural language.

Human Learning Requires Language

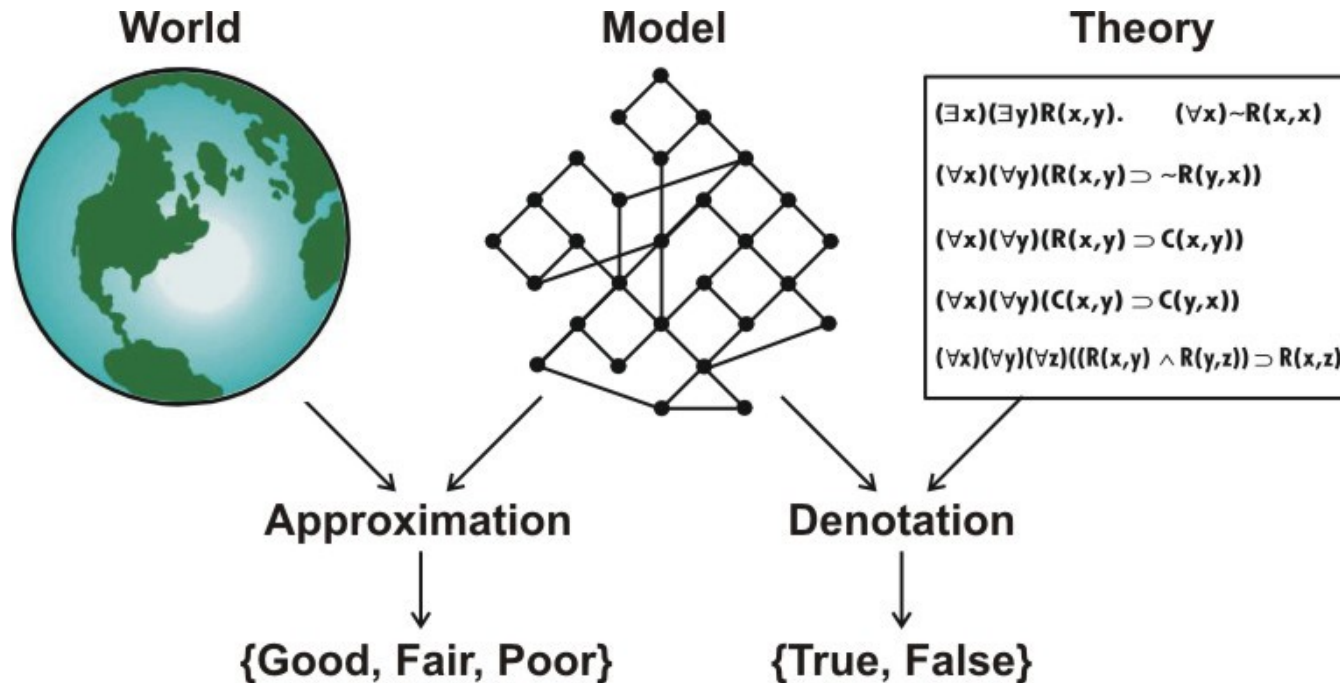
People use language to express every aspect of life.

The cognitive cycle integrates all aspects, including language:

- **New data (experiences) accumulate from observations in life.**
- **Statistical methods are useful for finding generalizations.**
- **But those generalizations must be integrated with previous knowledge.**
- **Routine abduction may use statistics to select patterns from the soup.**
- **But creative abduction is necessary to invent new patterns.**
- **Belief revision integrates various patterns into larger, better structured patterns called hypotheses or theories.**
- **Deduction generates predictions from the theories.**
- **Actions in and on the world test the predictions.**
- **New observations provide supervision (rewards and punishments).**

Language is essential for expressing novel patterns and for learning the novel patterns discovered by other people.

Relating Language to the World



Model-theoretic semantics relates symbols to symbols:

- It determines truth values by relating logic to a Tarski-style model.
- But the symbols of a model are, at best, approximations to the world.
- As engineers say, “All models are wrong, but some are useful.”

Peirce’s theory of symbol grounding:

- The signs of perception and purposive action are primary.
- The symbols of language are interpretations of the primary signs.
- Logic is an abstraction from language, not a foundation for language.


Boyd's OODA Loop

John Boyd drew a four-step diagram for training fighter pilots to observe and respond rapidly.

The first two steps – Observe and Orient – involve the occipital, parietal, and temporal lobes.

The next two steps – Decide and Act – involve the frontal lobes for reasoning and motor control.

The four steps and the associated brain areas:

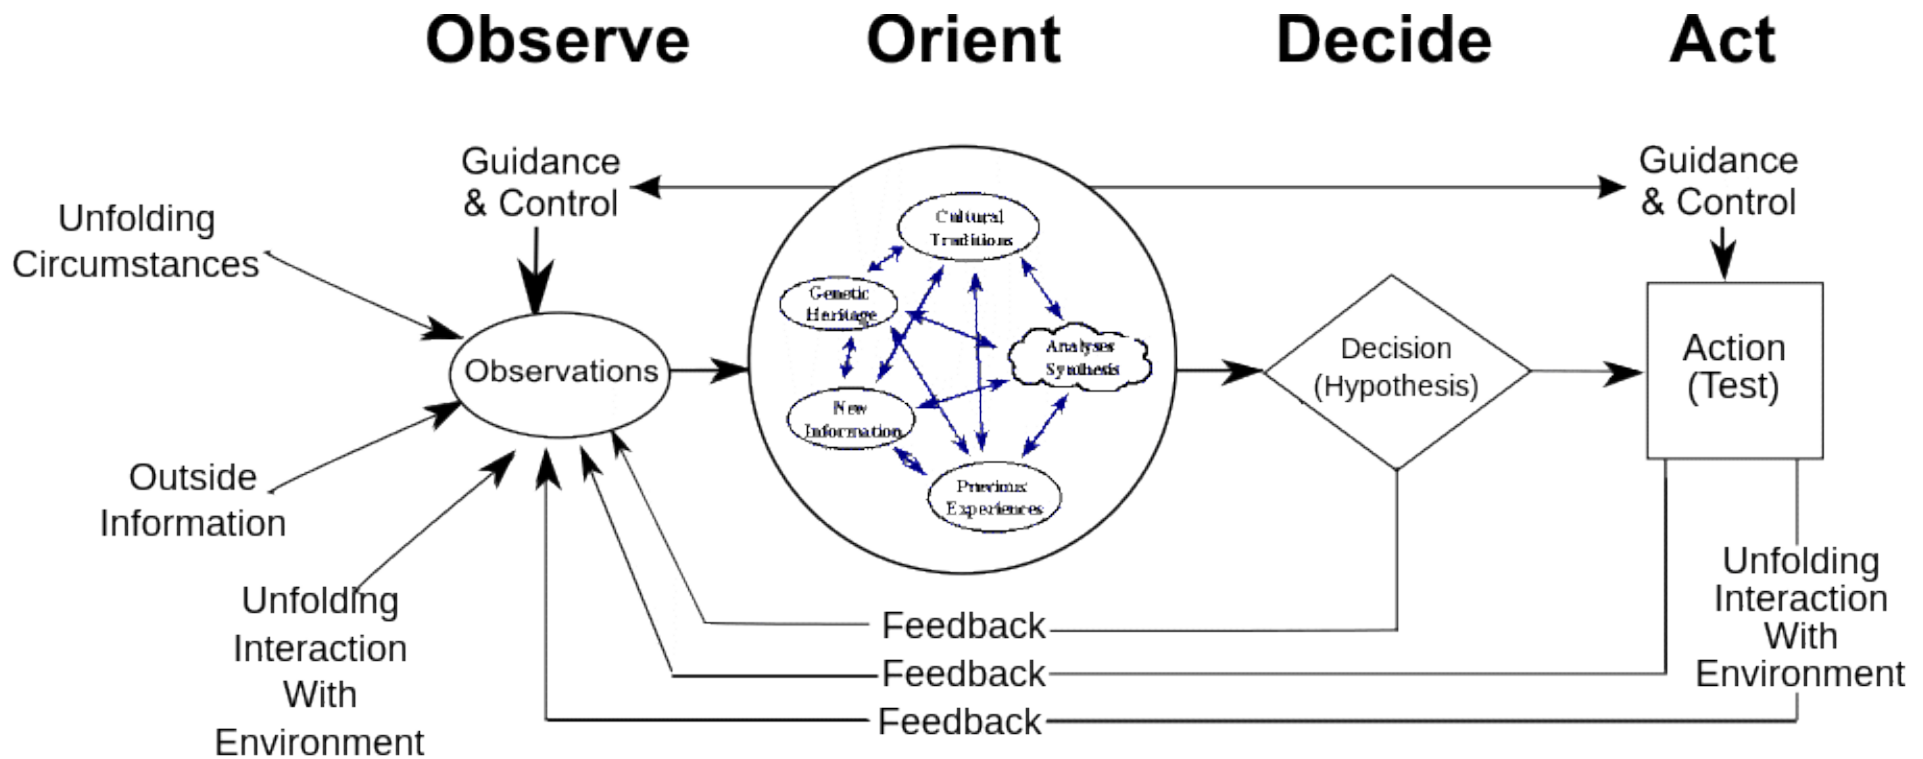


1. **Observe:** Visual input goes to the primary visual cortex (occipital lobes), but object recognition and naming involve the temporal lobes.
2. **Orient:** Parietal lobes relate vision, touch, and sound in “cognitive maps.”
3. **Decide:** Reasoning is under the control of the frontal lobes, but other areas store the “knowledge soup” and the “mental models.”
4. **Act:** “Action schemata” are patterns in the premotor cortex of the frontal lobes. Signals from the motor cortex go to the muscles.

Each step must be traversed in milliseconds for rapid response.

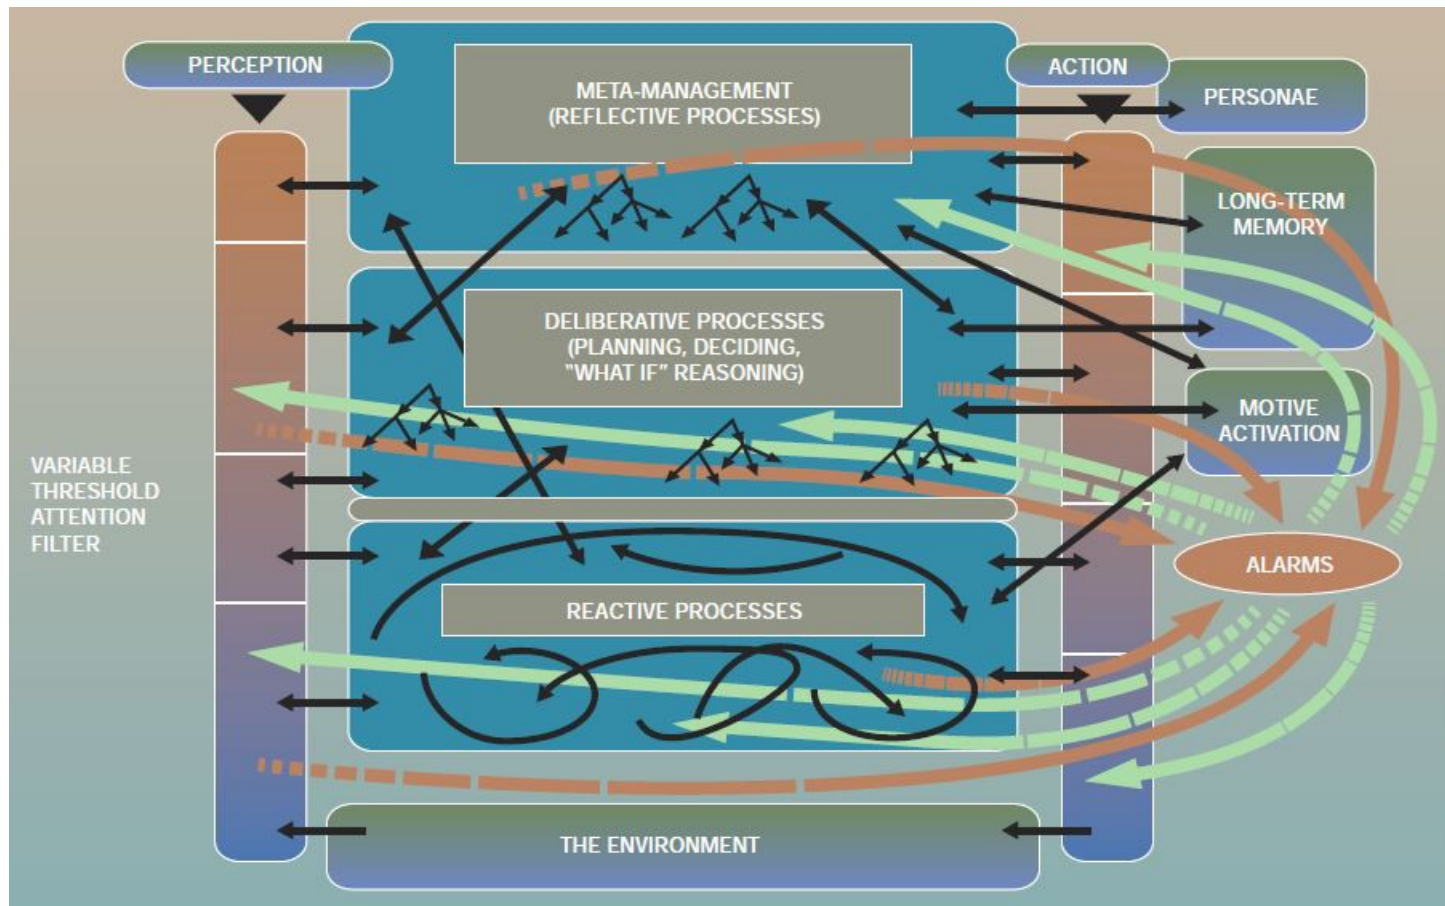
The time constraints require high-speed matching of overlearned patterns.

Extended OODA Loop



Over the years, Boyd added more detail to the OODA Loop. He applied it to decision-making processes of any kind. Both versions are consistent with Peirce's cycle.

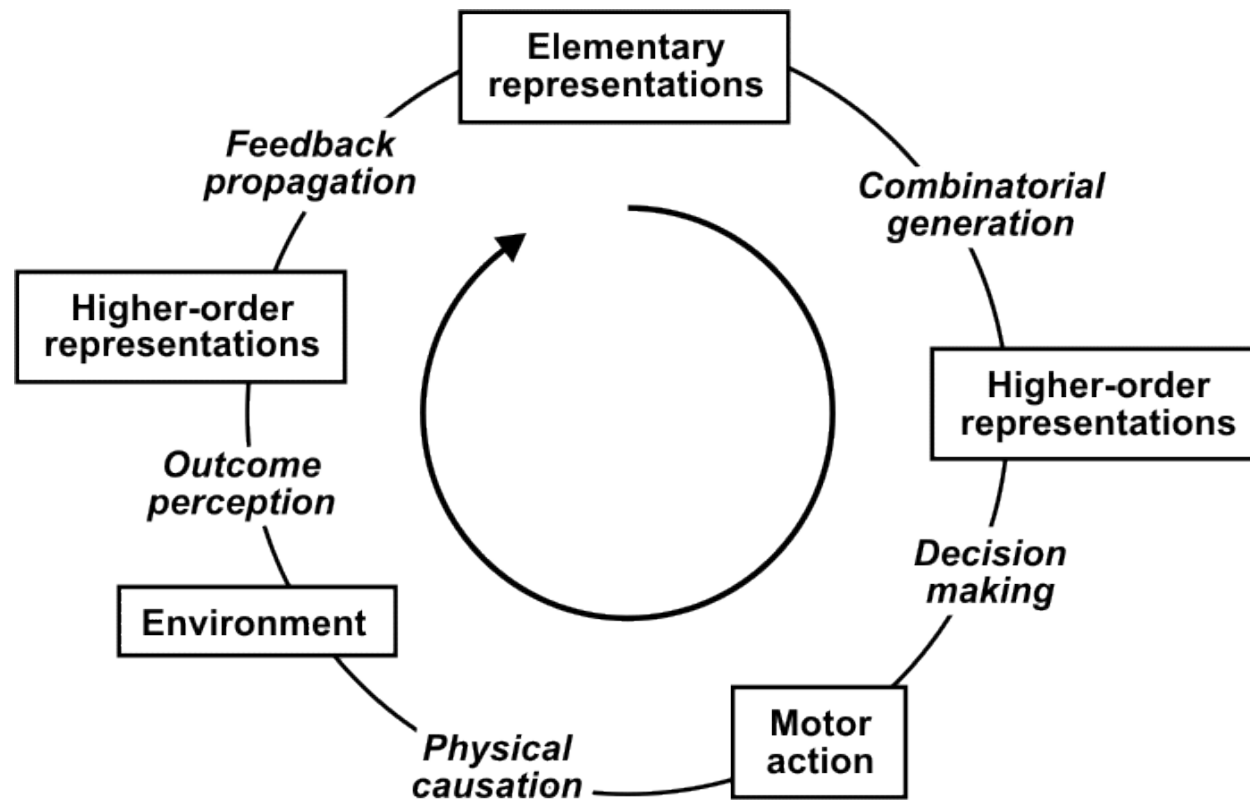
The H-Cogaff Architecture



The Hierarchical Cognition Affect architecture by Aaron Sloman includes a cycle similar to Peirce's or Boyd's. *

* From [An architecture of diversity for Commonsense Reasoning](#) by McCarthy, Minsky, Sloman, et al. (2002)

Ohlsson's Deep Learning Cycle

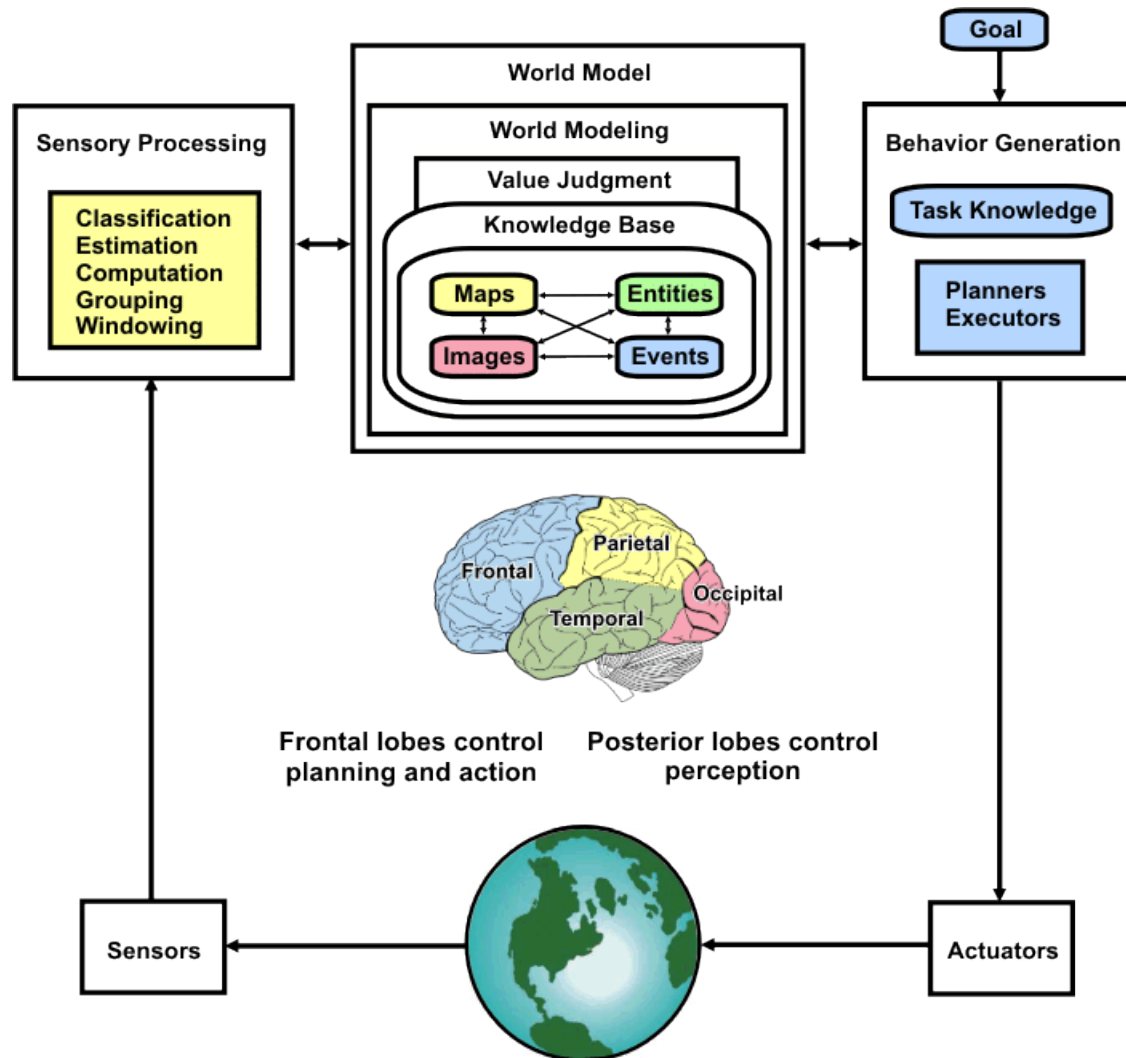


Deep learning is non-monotonic cognitive change: *

- Create novel structures that are incompatible with previous versions.
- Adapt cognitive skills to changing circumstances.
- Test those skills by action upon the environment.

* Stellan Ohlsson (2011) *Deep Learning: How the Mind Overrides Experience*, Cambridge: University Press.

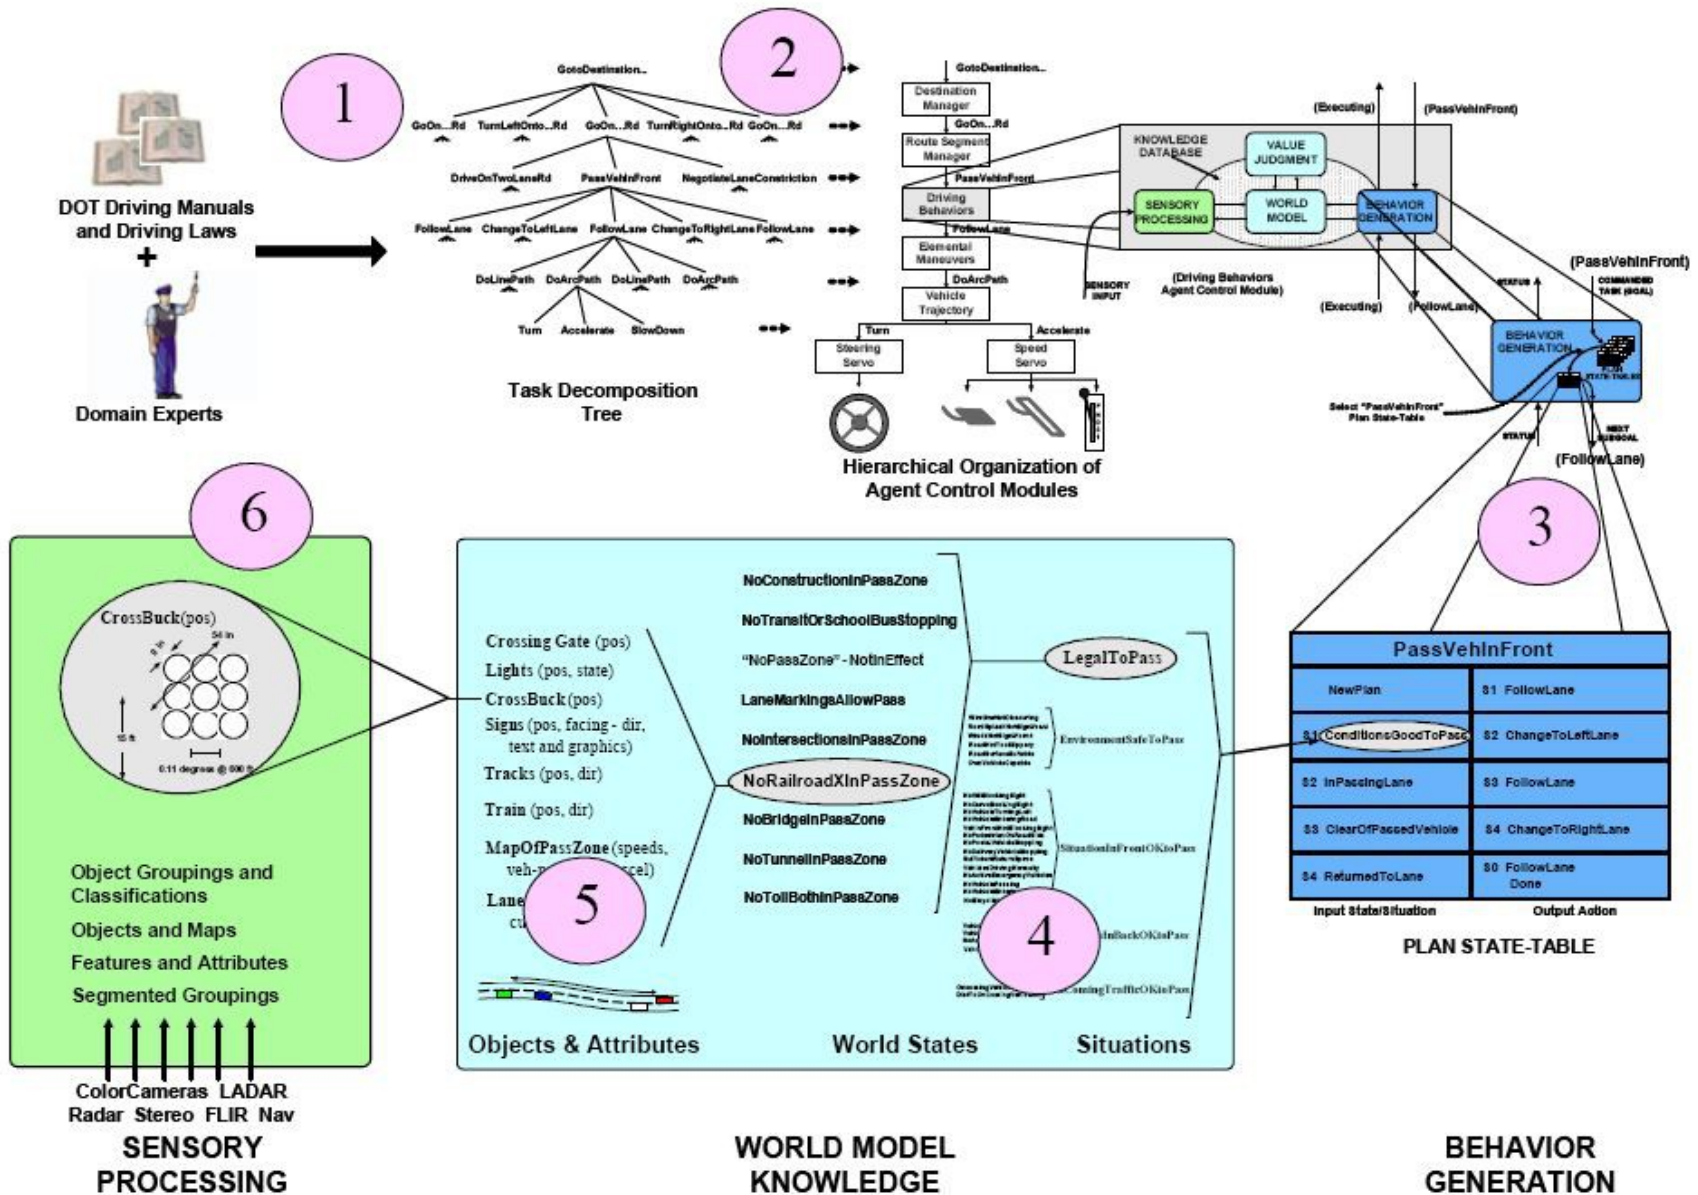
Albus Cognitive Architecture



A cycle that resembles those by Peirce, Boyd, Sloman, and Ohlsson.

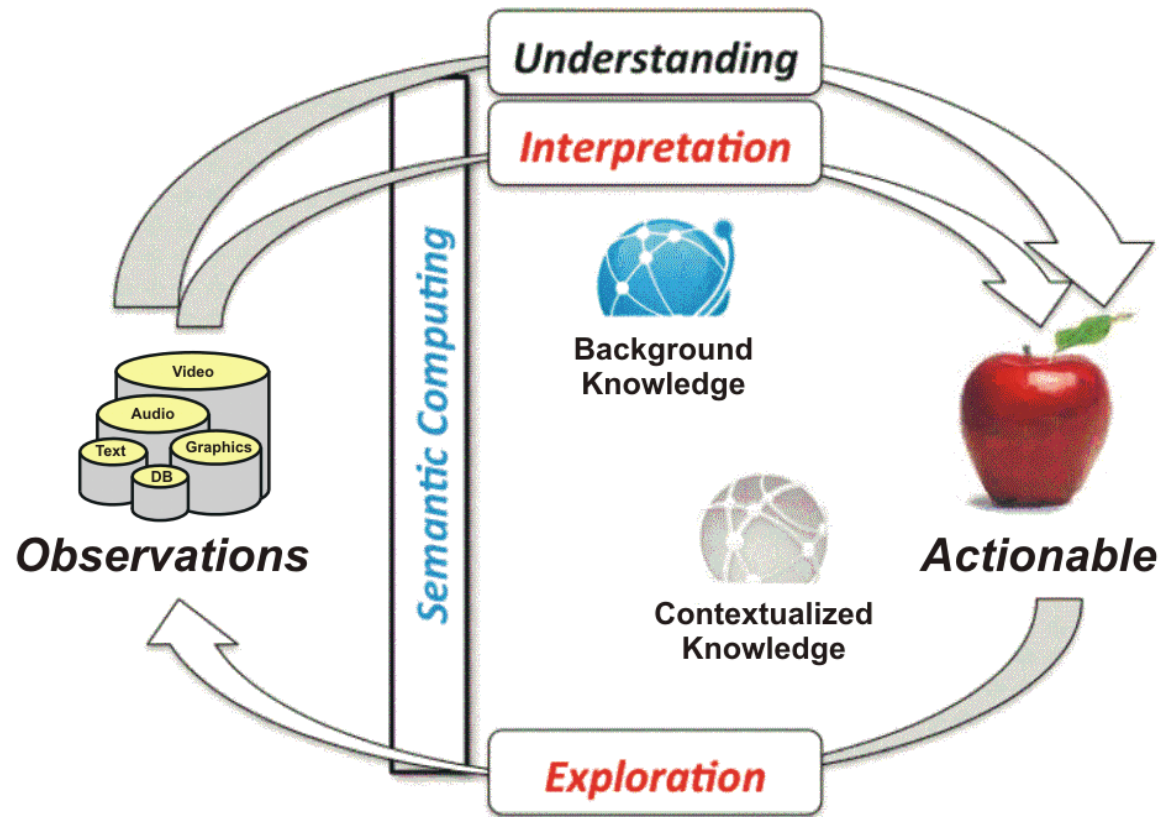
See Albus (2010), <http://www.james-albus.org/docs/ModelofComputation.pdf>

Real-Time Control System (RCS)



Designed by Albus and colleagues: http://en.wikipedia.org/wiki/Real-time_Control_System

Levels of AI Computation

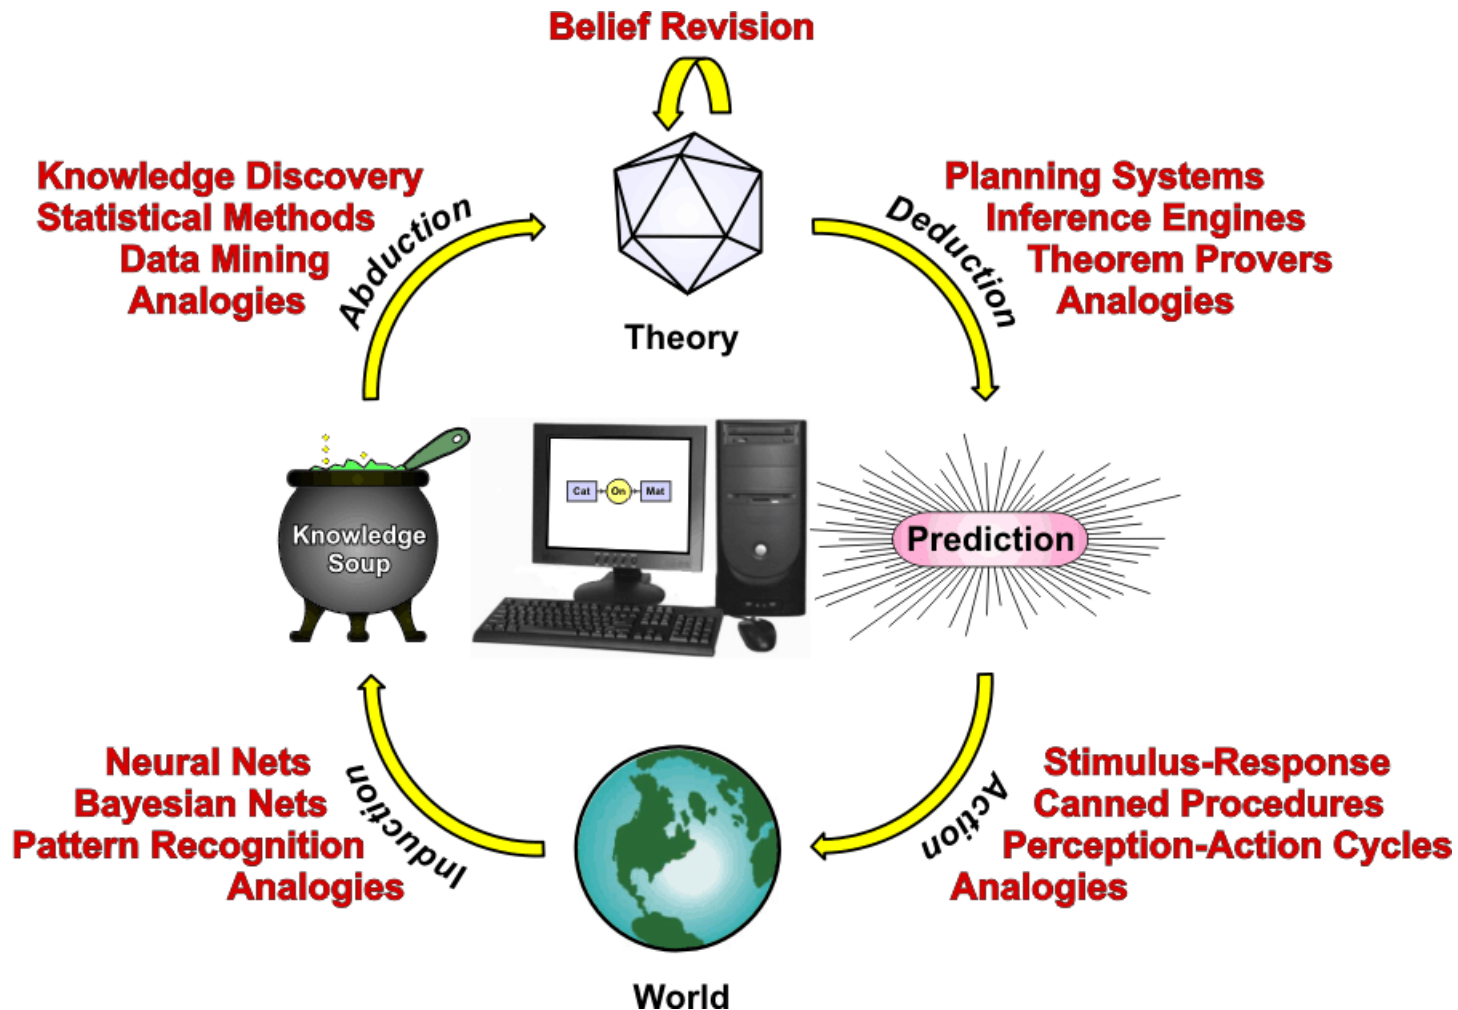


Sheth, Anantharam, and Henson distinguished three levels: *

- **Semantic:** A semantic network for representing knowledge.
- **Perception:** Using background knowledge to interpret sensory data.
- **Cognitive:** Understanding the knowledge in context and acting upon it.
- **But more research is necessary on all three levels.**

* Diagram adapted from <http://arxiv.org/ftp/arxiv/papers/1510/1510.05963.pdf>

Implementing the Cycles



An open-ended variety of methods for learning and reasoning.

Creative Abduction

Creativity, by definition, introduces something totally new.

Observation and abduction are the sources of novelty:

- **Observation is the ultimate source of all information.**
- **Routine observations classify new information in familiar patterns.**
- **Induction generalizes multiple observations by simplifying patterns.**
- **Routine abduction makes selections from familiar patterns.**
- **Belief revision modifies a theory by adding and deleting patterns.**
- **Deduction uses systematic rules for combining and relating patterns.**
- **But creative abduction (invention) introduces novel patterns.**

For young children, almost everything is unfamiliar.

- **They are the most creative people on earth.**

For most adults, most things are familiar.

- **They seldom feel the need to create radically new patterns.**
- **But they can learn new patterns created by other people.**

References

Research that established the foundations for Kyndi technology:

Majumdar, Arun K. (2013) Relativistic concept measuring system for data clustering, US Patent 8,526,321 B2.

Majumdar, Arun K. (2015) Cognitive memory encoding networks for fast semantic indexing, storage, and retrieval, US Patent 9,158,847 B1.

Majumdar, Arun K., John F. Sowa, & John Stewart (2008) Pursuing the goal of language understanding, <http://www.jfsowa.com/pubs/pursuing.pdf>

Majumdar, Arun K., & John F. Sowa (2009) Two paradigms are better than one and multiple paradigms are even better, <http://www.jfsowa.com/pubs/paradigm.pdf>

Majumdar, Arun K., & John F. Sowa (2014) Quantum cognition, <http://www.jfsowa.com/pubs/qcog.pdf>

Sowa, John F. (2002) Architectures for intelligent systems, <http://www.jfsowa.com/pubs/arch.pdf>

Sowa, John F., & Arun K. Majumdar (2003) Analogical reasoning, <http://www.jfsowa.com/pubs/analog.htm>

Sowa, John F. (2005) The challenge of knowledge soup, <http://www.jfsowa.com/pubs/challenge.pdf>

Sowa, John F. (2006) Worlds, models, and descriptions, <http://www.jfsowa.com/pubs/worlds.pdf>

Sowa, John F. (2008) Conceptual graphs, http://www.jfsowa.com/cg/cg_hbook.pdf

Sowa, John F. (2010) Role of Logic and Ontology in Language and Reasoning, <http://www.jfsowa.com/pubs/rolelog.pdf>

Sowa, John F. (2011) Cognitive architectures for conceptual structures, <http://www.jfsowa.com/pubs/ca4cs.pdf>

Sowa, John F. (2013) From existential graphs to conceptual graphs, <http://www.jfsowa.com/pubs/eg2cg.pdf>

ISO/IEC standard 24707 for Common Logic, [http://standards.iso.org/ittf/PubliclyAvailableStandards/c039175_ISO_IEC_24707_2007\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/c039175_ISO_IEC_24707_2007(E).zip)