Why Has AI Failed?
And How Can It Succeed?

John F. Sowa

VivoMind Research, LLC

10 May 2015

Extended version of slides for MICAI'14

Problems and Challenges

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 generate 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?

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, technical adviser for the movie *2001*:

“The HAL 9000 is a *conservative estimate* of the level of artificial intelligence in 2001.”

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."

Laura used a larger subset of logic than Montague formalized.

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

*** John Limber, The genesis of complex sentences.**
http://pubpages.unh.edu/~jel/JLimber/Genesis_complex_sentences.pdf

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.

Why Has Progress Been So Slow?

Theorem provers in 1960 were much faster than humans.

Today's computers are a million times bigger and faster.

Why haven't the predictions by Delavenay, Good, and Minsky come true years ago?

Short answer: Animal brains are smarter than silicon brains.

Longer answer:

1. From 1945 to 1970, progress in both hardware and software grew exponentially.
2. The rate of growth in hardware technology is continuing.
3. But the rate of growth in software technology slowed down.

Why can't AI technology improve software development?

Home Computer of 2004

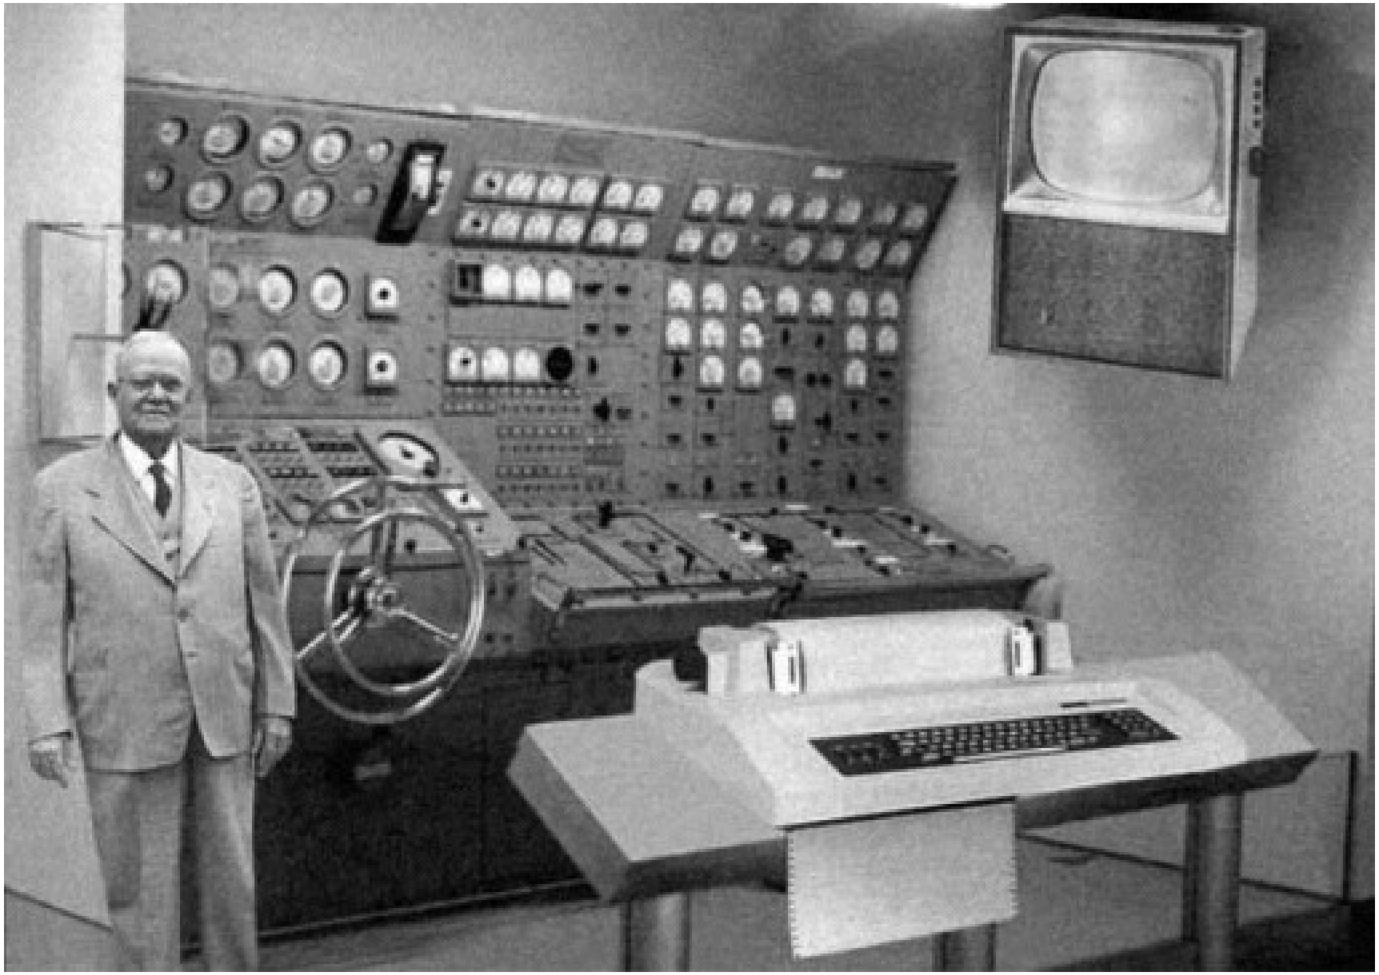

As predicted by the RAND Corporation in 1954: *

"With teletype interface and the FORTRAN language, it will be easy to use."

*** Actually, this claim is a hoax. But BASIC on the IBM PC in 1981 was almost as hard to learn as FORTRAN in 1954, and modern languages such as Java are even harder.**

Historical Perspective

For computer hardware, smaller is better:

- 1940s: Vacuum tubes for the "Giant Brains".
- 1950s: Smaller tubes, denser storage, early transistors.
- 1960s: Integrated circuits with multiple transistors.
- 1970s to present, Moore's law: "The number of transistors per chip doubles every 18 months."

From 1945 to 1980, many centuries of research was digitized:

- Theories of logic and ontology from Aristotle to the present.
- Computing devices from the abacus to the Turing machine.
- Mathematical methods for every branch of science.
- Record keeping in banking, business, schools, and governments converted to a computable form for punched-card machines.

The major programming languages today have more features than FORTRAN in 1954, but they are no easier to use.

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.

Visualization 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.”

Language, Learning, and Reasoning

For everyone from Laura to Einstein, perception, action, and mental imagery are fundamental.

Natural languages are based on those mental images.

The notations for mathematics and logic are abstractions from the symbols and patterns in natural languages.

Computer systems can manipulate those symbols much faster than any human.

- But computers are much less efficient in perception and action.
- How could they support a cognitive cycle of perception, learning, reasoning, and action?
- And the languages people use to talk about every step?

What is Possible Today?

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 most systems are designed around a single paradigm — they cannot take advantage of all the available resources.

What Makes People Intelligent?

Short answer: Flexibility, generality, and adaptability.

- 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 answer: 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.

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 represent the general knowledge of a typical high-school student.
- A high-school education isn't sufficient to win at Jeopardy!
- But IBM devoted a large research team to a single problem.

Longer answer: Single paradigm vs. multiple paradigms.

- Cyc is based on a single paradigm: formal logic, deductive reasoning, a very large ontology, and large volumes of data.
- IBM used a team of researchers with a wide range of expertise.
- They started with many independently developed tools and made them interoperate on different aspects of the problem.

Next question: Is it possible to develop Watson-like systems without requiring three dozen PhD researchers?

Minsky's Challenge

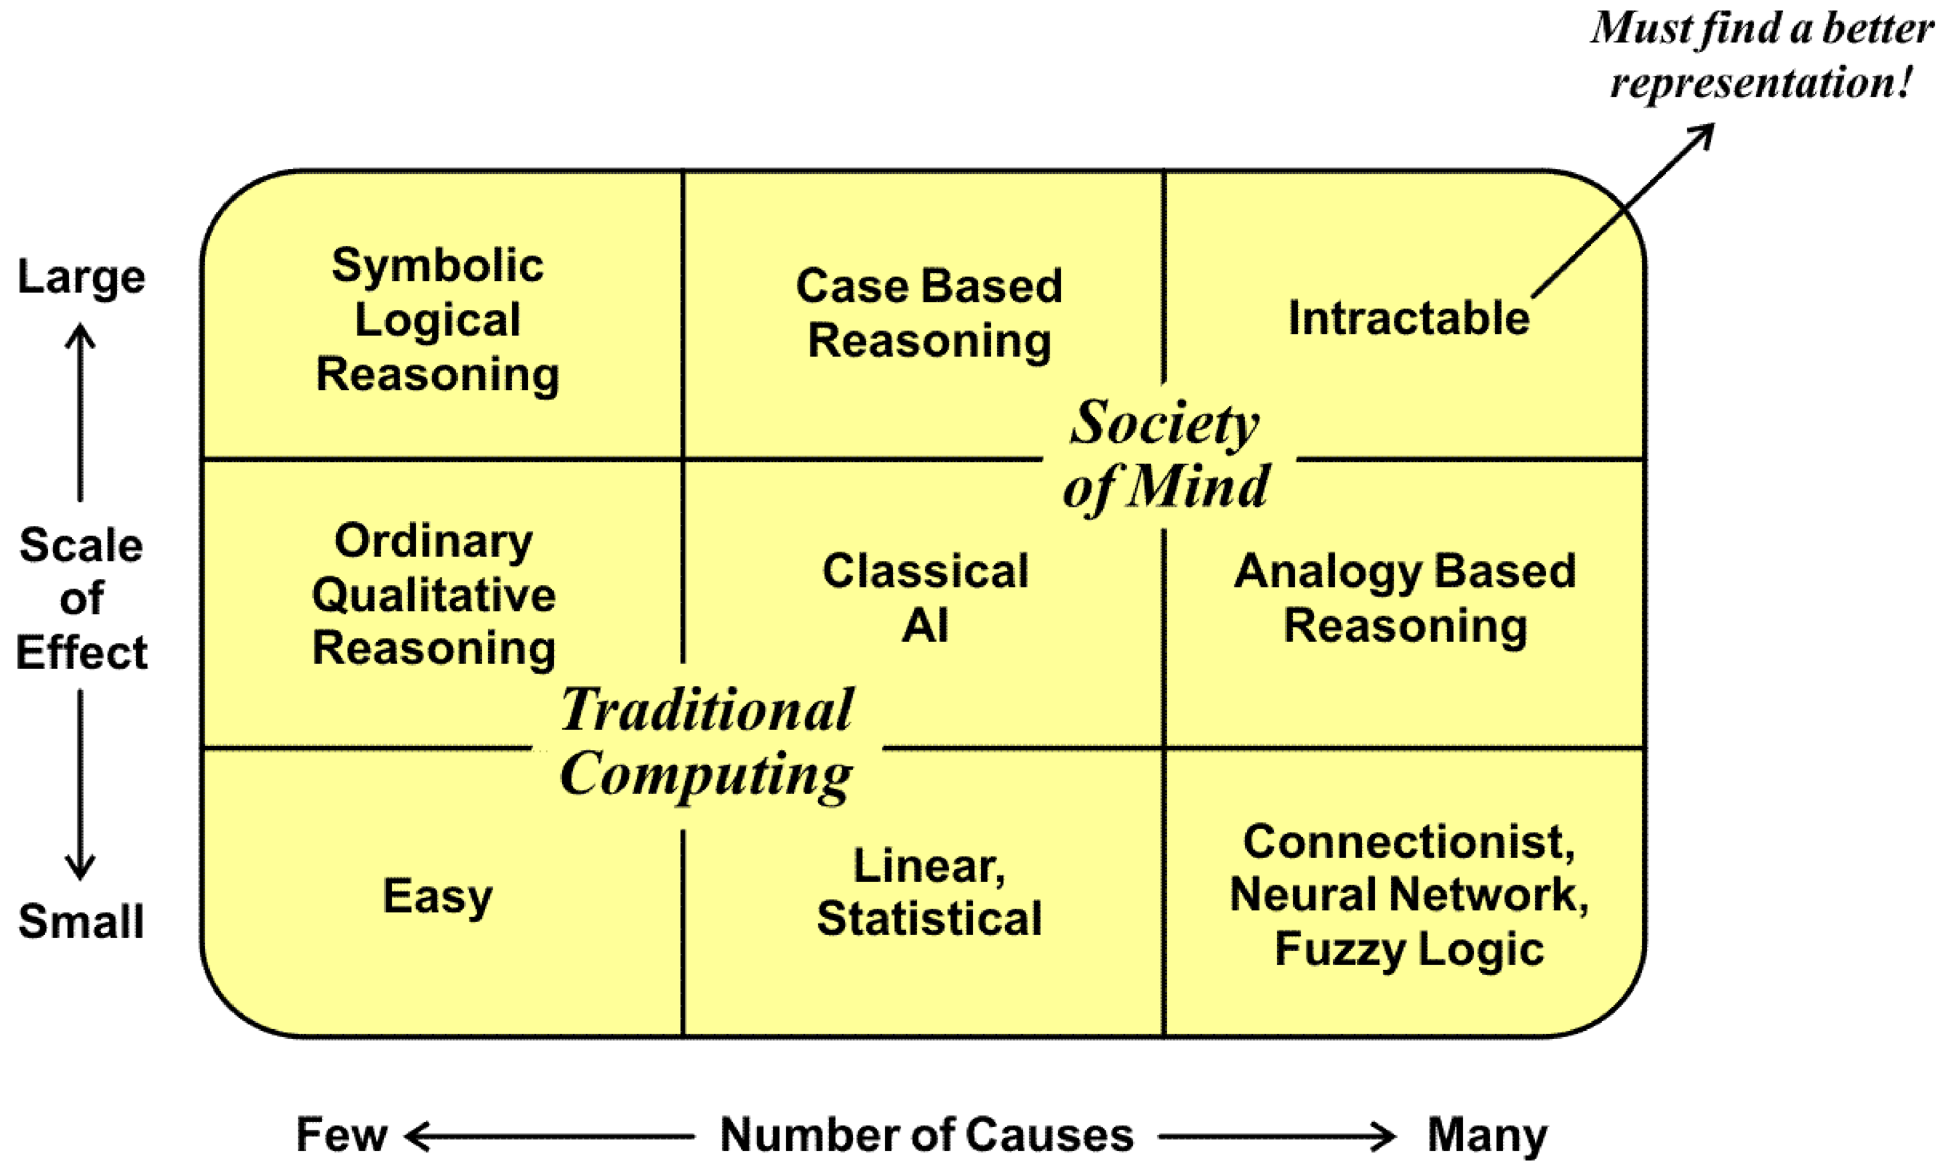

Adapted from a diagram by Minsky, Singh, & Sloman (2004).

Meeting the Challenge

As Minsky's diagram shows, AI methods cannot process large numbers of causes and complex effects as efficiently as humans.

Statistical methods and neural networks can relate many causes (input variables), but only small-scale effects (simple outputs).

Logic can reason about complex effects (multiple interrelated phenomena), but only with simplified causes (few axioms).

In his *Society of Mind* and *Emotion Engine*, Minsky proposed systems of heterogeneous, interacting modules or agents:

- **Can those agents improve computational efficiency?**
- **Can psycholinguistics and neuroscience guide the design of agents?**
- **What kind of logic, reasoning, and semantics would they support?**
- **Would they use symbolic, statistical, or image-like representations?**
- **Or would they use an open-ended variety of representations?**

Making and Using Tools

The ability to make tools is a critical sign of intelligence:

- All animals, including humans, are born with a set of built-in tools.
- Birds can't wash dishes for the same reason that humans can't wash dishes with a sewing machine: they have the wrong tools.
- Humans make and use the most elaborate tools, but biologists keep discovering more species that make and use tools.

The role of instinct:

- Birds have an instinct to build nests, beavers have an instinct to build dams, and humans have an instinct to speak a language.
- But the details of the nests, dams, and languages depend on the animals' built-in tools, the environment, learning from parents, and creativity.

For every physical tool or activity, there is a kind of mental tool:

- Minsky's mental agents correspond to the diversity of human activities.
- The diversity of Wittgenstein's language games results from the many ways of using the human vocal tools for talking about the activities.

VivoMind Technology

Tools for describing, analyzing, and reasoning about anything.

Conceptual graphs for mapping to and from languages and notations of any kind — natural or artificial, formal or informal.

Intellitex parser for natural and artificial languages.

- Translate any notation to or from conceptual graphs.

Cognitive Memory™ for associative storage of graphs.

- For any query graph, find all approximate matches in (log N) time.

Flexible Modular Framework (FMF) for integrating components:

- Influenced by Marvin Minsky's Society of Mind, John McCarthy's Elephant 2000, and David Gelernter's Linda.

Automated tools for learning by reading documents.

Role of Analogy

Basis for perception, reasoning, and language understanding.

Logic is a disciplined special case of analogical reasoning:

- Essential for precise reasoning in mathematics and science.
- Important for precision in any field.
- But even in science and engineering, analogy is necessary for knowledge discovery and innovation.

Conceptual graphs support logical and analogical methods:

- They are defined by the ISO/IEC standard 24707 for Common Logic.
- But they also support semantic-distance measures for analogy.
- They provide a bridge between informal language and formal logic.

CGs derived from natural languages can be used for analogies.

But CGs used for formal logic should be derived from formal languages or be corrected by comparison to formal CGs.

Computational Complexity

Research by Falkenhainer, Forbus, & Gentner: *

- Pioneers in finding analogies with the Structure Mapping Engine.
- Showed that SME algorithms take time proportional to N-cubed, where N is the number of frames or graphs in the knowledge base.
- MAC/FAC: Reduce the number of candidates before using SME.
- But how do you determine which candidates are best?

VivoMind approach:

- Encode graph structure and ontology in a Cognitive Signature™.
- For any graph, find closely matching signatures in log(N) time.
- Only graphs with similar signatures are likely candidates.
- The log-time algorithms scale to the size of the WWW.

* See http://www.qrg.northwestern.edu/papers/papers.html

Efficient Algorithms

Graphs of organic molecules are similar to conceptual graphs:

- **Atoms ⇒ concept nodes labeled by the name of the element.**
- **Chemical bonds ⇒ relation nodes labeled by the name of the bond type.**
- **But conceptual graphs have many more types of concepts and relations.**

Chemical graphs inspired Peirce to design existential graphs as a notation for "the atoms and molecules of logic."

Chemists found highly efficent methods for processing graphs, including log-time search algorithms:

Mining patents using molecular similarity search, by Rhodes et al., http://psb.stanford.edu/psb-online/proceedings/psb07/rhodes.pdf

VivoMind algorithms also use geometric algebras and Rvachev * functions: US Patent 8566321, Relativistic concept measuring system.

*** For Rvachev functions, see http://docs.lib.purdue.edu/dissertations/AAI3263546/**

Cognitive Memory™

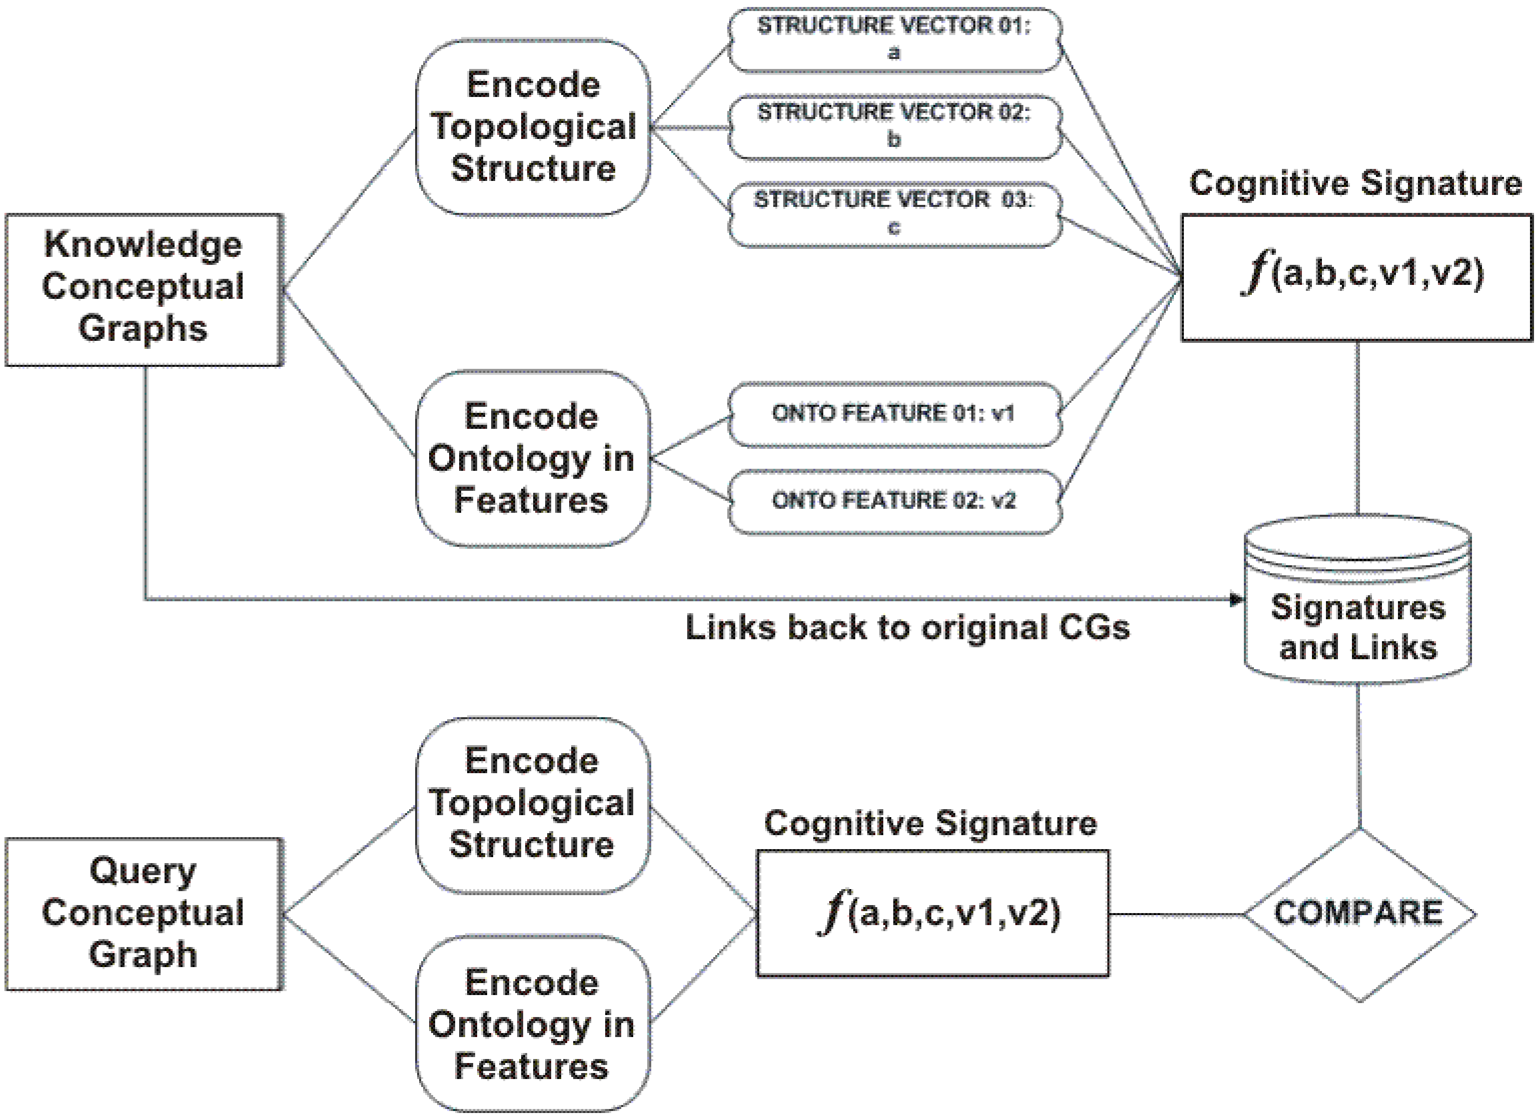

Basis for the VivoMind Analogy Engine (VAE)

Applications of VivoMind Software

General approach:

- Intellitex and VAE for analyzing English and other languages.
- Basic ontology and lexical resources in conceptual graphs.
- Add the client’s preferred terminology and ontology (if any).
- Derive further information from any NL texts the client uses.
- Integrate information from all sources, structured or unstructured.
- Translate the results to any notation or format the client prefers.

Some sample applications:

1. Evaluate student answers in free-form English sentences.
2. Information extraction from research reports.
3. Derive CGs from textbooks and research reports about oil and gas fields, and answer English queries by a geologist.

Evaluating Student Answers

Multiple-choice questions are easy to evaluate by computer.

Long essays are often evaluated by statistical methods.

But short answers about mathematics are very hard to evaluate.

Sample question:

> ***The following numbers are 1 more than a square: 10, 37, 65, 82.***
> ***If you are given an integer N that is less than 200,***
> ***how would you determine whether N is 1 more than a square?***
> ***Explain your method in three or four sentences.***

An example of a correct answer:

> ***To show that N is 1 more than a square, show that N−1 is a square.***
> ***Find some integer x whose square is slightly less than N−1.***
> ***Compare N−1 to the squares of x, x+1, x+2, x+3, ...,***
> ***and stop when some square is equal to or greater than N−1.***
> ***If the last square is N−1, then N is one more than a square.***

Even experienced teachers must spend a lot of time checking and correcting such answers.

Publisher's Current Procedure

To evaluate new exam questions, the publisher normally gives the exam to a large number of students.

For each problem, they would get about 50 different answers:

- Some are completely correct
  — but stated in different ways.
- Some are partially correct
  — and the teacher says what is missing.
- Others are wrong
  — in many different ways.

Result: 50 pairs of student answer and teacher's response.

Each answer-response pair is a case for case-based reasoning.

Case-Based Reasoning

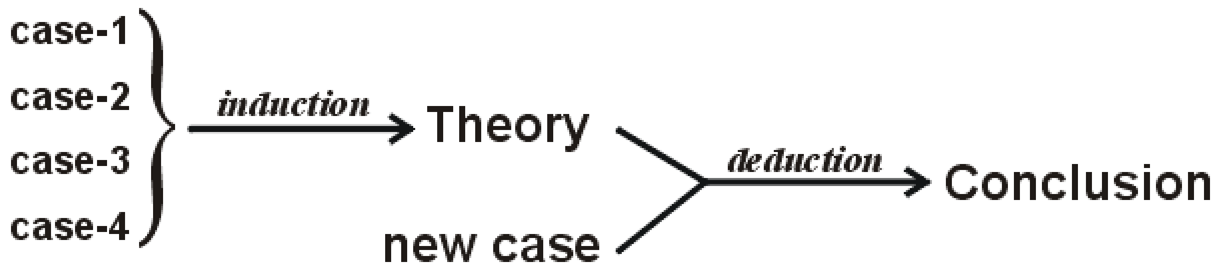

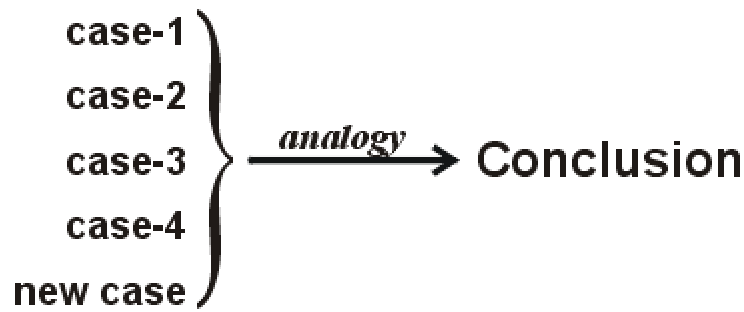

Given the same cases, analogy takes one step to derive an answer, but induction and deduction take multiple steps.

Analogy is fast and flexible, but a theory is important if it can be used and reused in many applications.

Using Intellitex and VAE

Translate all answers to conceptual graphs (CGs):

- **Intellitex uses a link grammar and a variety of semantic knowledge.**
- **Whenever a new CG is derived, it is stored in Cognitive Memory.**
- **With high-speed search, any CG can be used to analyze any sentence.**
- **The next slide shows a formally defined CG for Multiply.**

VAE compares each new answer to the 50 cases:

- **CGs for the answers of all 50 cases are stored in Cognitive Memory.**
- **Compare the CG for each new answer to the CGs in Cognitive Memory.**
- **If there is a good match, print out the teacher's previous response.**
- **Otherwise, send the new student answer to some teacher to evaluate.**
- **Add the new answer-response pair to the collection of cases.**

VAE supports both formal and informal reasoning.

Conceptual Graph for Multiply

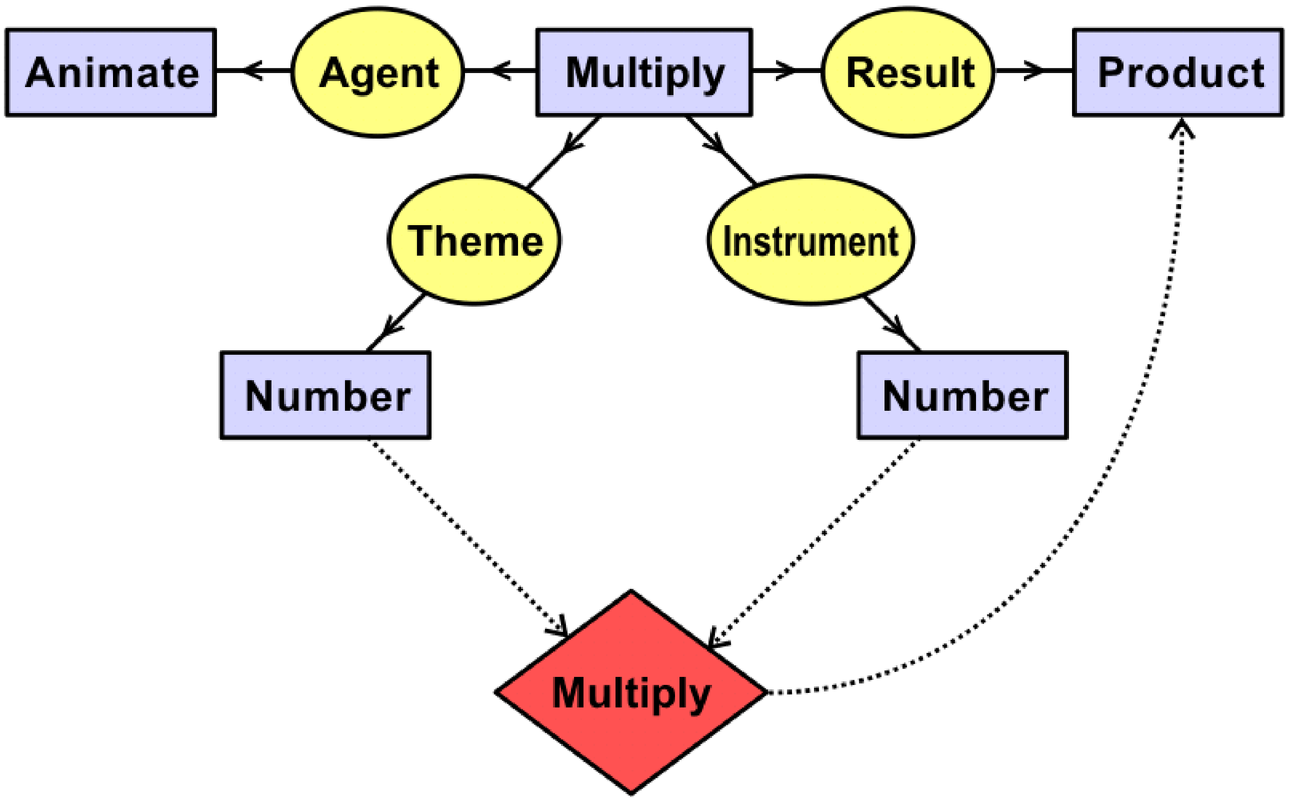

This CG represents a pattern or schema for the concept Multiply:

[Someone] multiplies a number by a number to get a product.

The diamond node, called an actor, represents a function that multiplies two numbers to compute their product.

Using Syntax and Semantics

Intellitex uses Cognitive Memory to find conceptual graphs that approximately match the input phrases and sentences.

Any CG can be used to supplement the syntax:

- **Intellitex uses a link grammar to generate dependency graphs.**
- **Those graphs can be mapped to CGs, but the mapping is rarely one-to-one, and Intellitex usually requires more information.**
- **CGs from any source can be used to check constraints, provide default information, resolve ambiguities, or suggest alternatives.**

Many kinds of CGs can be used during the analysis:

- **IBM-CSLI verb ontology* for case roles and type constraints.**
- **CGs for an ontology of arithmetic (see the previous slide).**
- **CGs automatically derived from a textbook for the course.**
- **CGs automatically derived from answers written by the teachers.**

*** See http://lingo.stanford.edu/vso/**

Results

VAE found a good match for most student answers.

The answers were often poorly written and not precise enough for a formal parser and theorem prover.

But the stored CGs made the parsing more robust —

- By resolving ambiguities and showing expected combinations,
- By correcting and compensating for errors in syntax,
- By providing defaults for missing arguments, and
- By relating multiple fragments to form complete sentences.

Approximate matches within a reasonable semantic distance were adequate for evaluating student answers.

For matches outside the limits, student answers could be sent to a teacher, who would write a new evaluation.

Information Extraction Project

The next slide shows a table derived from research reports.

To analyze the reports, an ontology of CGs for chemistry and superconductivity was added to Cognitive Memory.

Then for each report,

- Map each sentence to a conceptual graph.
- Analyze all anaphoric references between the CGs.
- The result is a single, tightly connected CG that links all the CGs for every sentence in the report.
- Store that CG 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,

- VivoMind got 96% of the entries correct.
- The second best had 73% correct. Two others were slightly above 50%. All the others were below 50%.

A Spreadsheet Derived from Published Reports

DOE BREMS PROJECT.xlsx

	COMPOUND	CURIE TEMP.	SOURCE
2	Mn3[Cr(CN)6]2 ·16H2O	89 K	A solid-state hybrid density functional theory study
3	Sr3Ir2O7 in Sr3Ir2O7 single-cr	~ 280 K	Canted antiferromagnetic ground state in Sr3Ir2O7
4	PrPt2B2C	6 K	Coexistence of superconductivity and magnetic ord
5	La0:3Nd0:7Pt2:1-	2.8 K	Coexistence of superconductivity and magnetic ord
6	NdPt2:1B2:4C1:2	3 K	Coexistence of superconductivity and magnetic ord
7	NdPt1:5Au0:6B2C	3 K	Coexistence of superconductivity and magnetic ord
8	SmNiC2	= 17.7 K	Commensurate charge-density wave with frustrate
9	Co0.2Zn0.8Fe2O4. in CdxCo1-	~ 780 K	Does Ti+4 ratio improve the physical properties of
10	Zn0.88Co0.12O in ZnO	~ 540 K	Effect of Co doping on the structural; optical and m
11	La in Sr2-xLaxFeMoO6	425 K	Effect of La doping on the properties of Sr2-xLaxFe
12	Fe in Sr2-xLaxFeMoO6	~ 1040 K	Effect of La doping on the properties of Sr2-xLaxFe
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	LaFexSi1 - x13 in La1-zPrz(Fe	~ 190 K	Enhancement of magnetocaloric effects in La1-zPrz
16	LaFe0.88Si0.1213 in La1-zPrz(	= 195 K	Enhancement of magnetocaloric effects in La1-zPrz
17	Co2MnGa in Co2MnGa	600 K	Ferromagnetic resonance in Co2MnGa films with va
18	HoCrO4 in HoCrO4	17.6 K	Ferromagnetism vs. antiferromagnetism of the dim
19	Mn3(HCOO)6 in Mn3(HCOO)(	8.0 K	Guest-induced chirality in the ferrimagnetic nanopo
20	NaZn13- in La0.5Pr0.5(Fe0.88	range from 195 K to 185 K	Large isothermal magnetic entropy change after hy
21	La2/3Ba1/3MnO3 in La2-3Ba1	range from 300 K to 250 K	Magnetic and neutron diffraction study of La2-3Ba1
22	CuMnSb in Co1-xCuxMnSb	range from 476 K to 300 K	Magnetic properties of half-metallic semi Heusler C
23	Nd2 in Nd2-yDyyFe17-xSix	range from 61.46 °C to 236 °	Magnetic properties of iron-rich Nd2-yDyyFe17-xSix
24	Tb2Fe17 in Nd2-yDyyFe17-xSi	~ 80 °C	Magnetic properties of iron-rich Nd2-yDyyFe17-xSix

Sheet1

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 tagging).
- 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 with VAE

For the sources, either NL documents or structured data:

- Use Intellitex to 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 a conceptual graph.
- Find matching patterns in the source data and rank them in order of semantic distance.
- 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

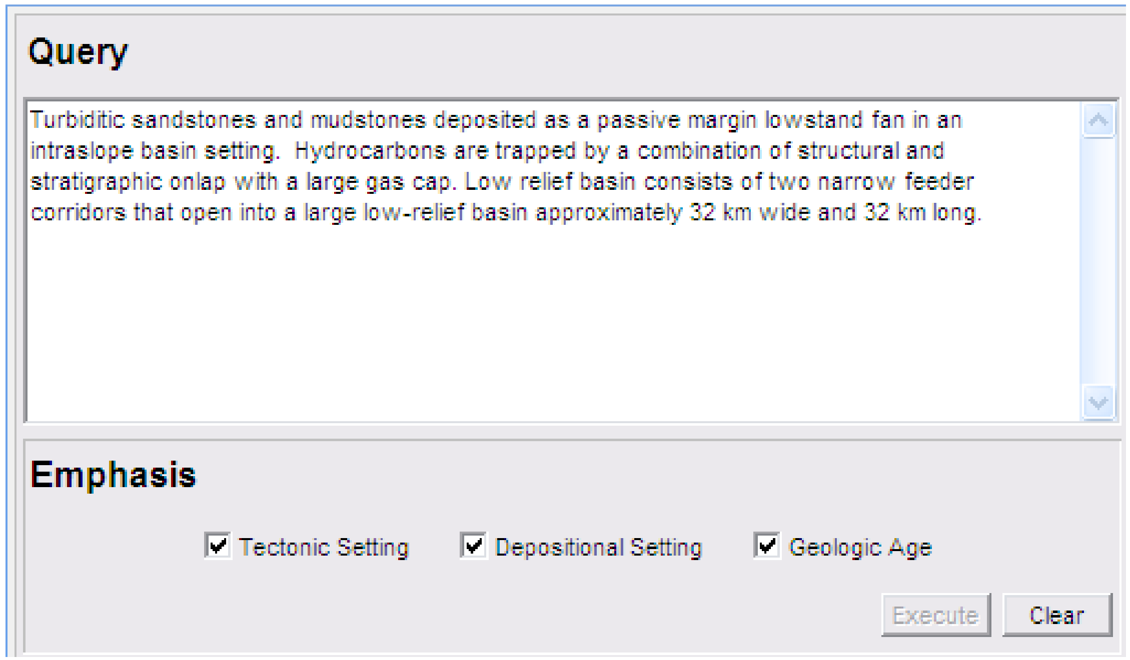

Turbiditic sandstones and mudstones deposited as a passive margin lowstand fan in an intraslope basin setting. Hydrocarbons are trapped by a combination of structural and stratigraphic onlap with a large gas cap. Low relief basin consists of two narrow feeder corridors that open into a large low-relief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

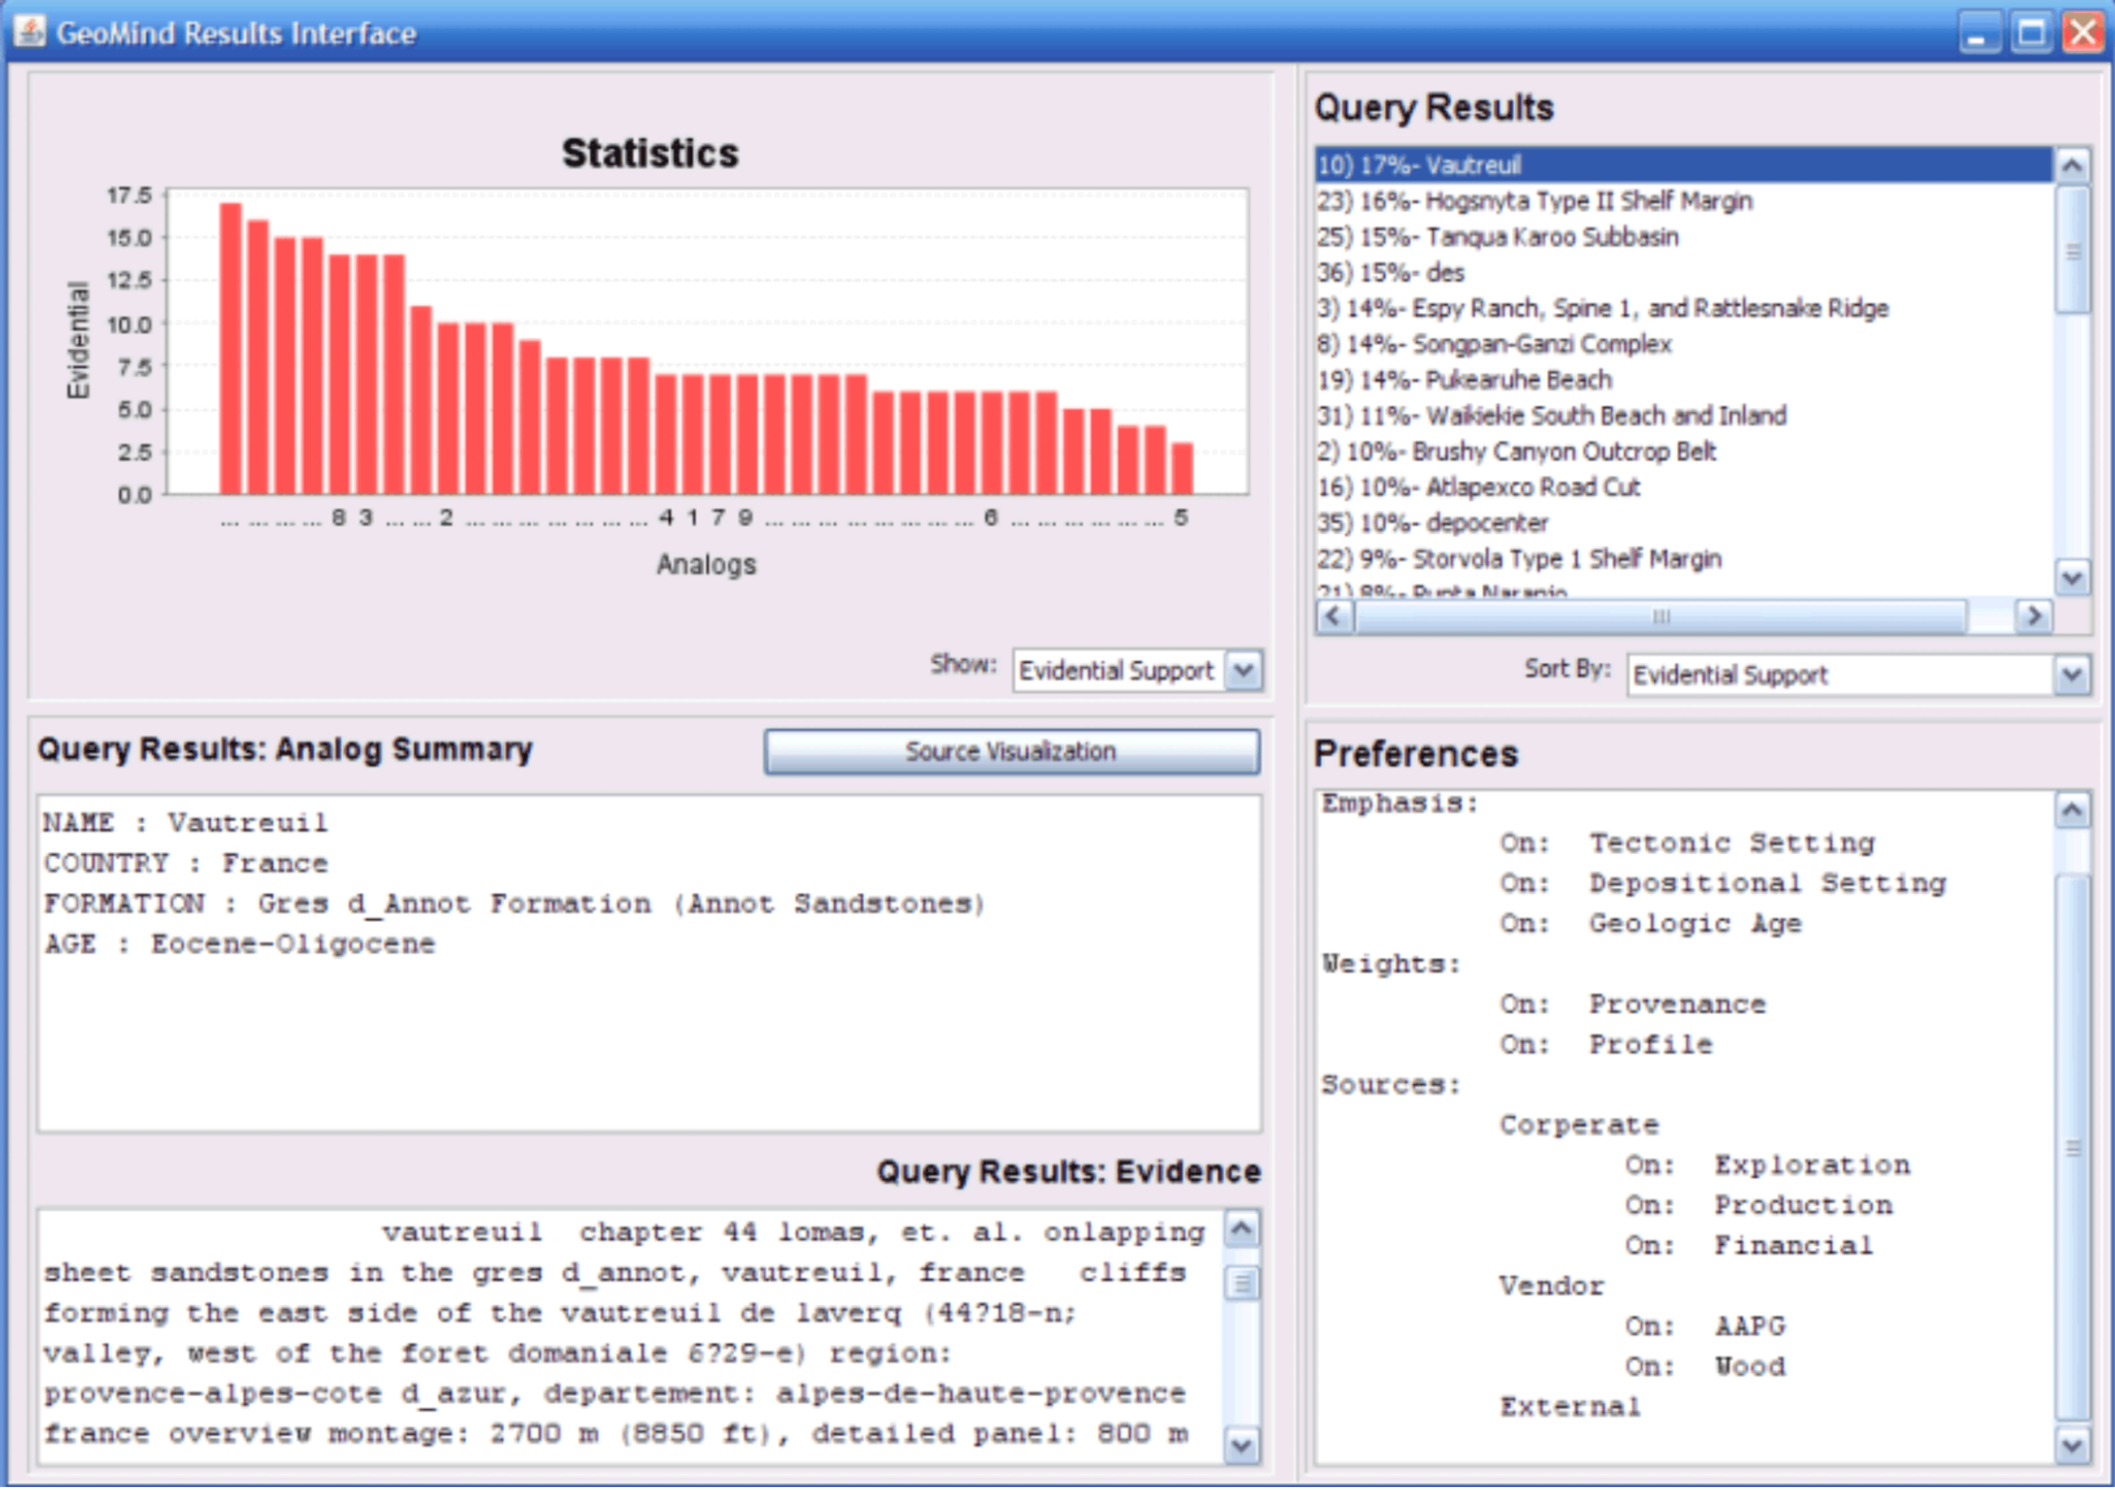

Details of the closest matching hydrocarbon fields

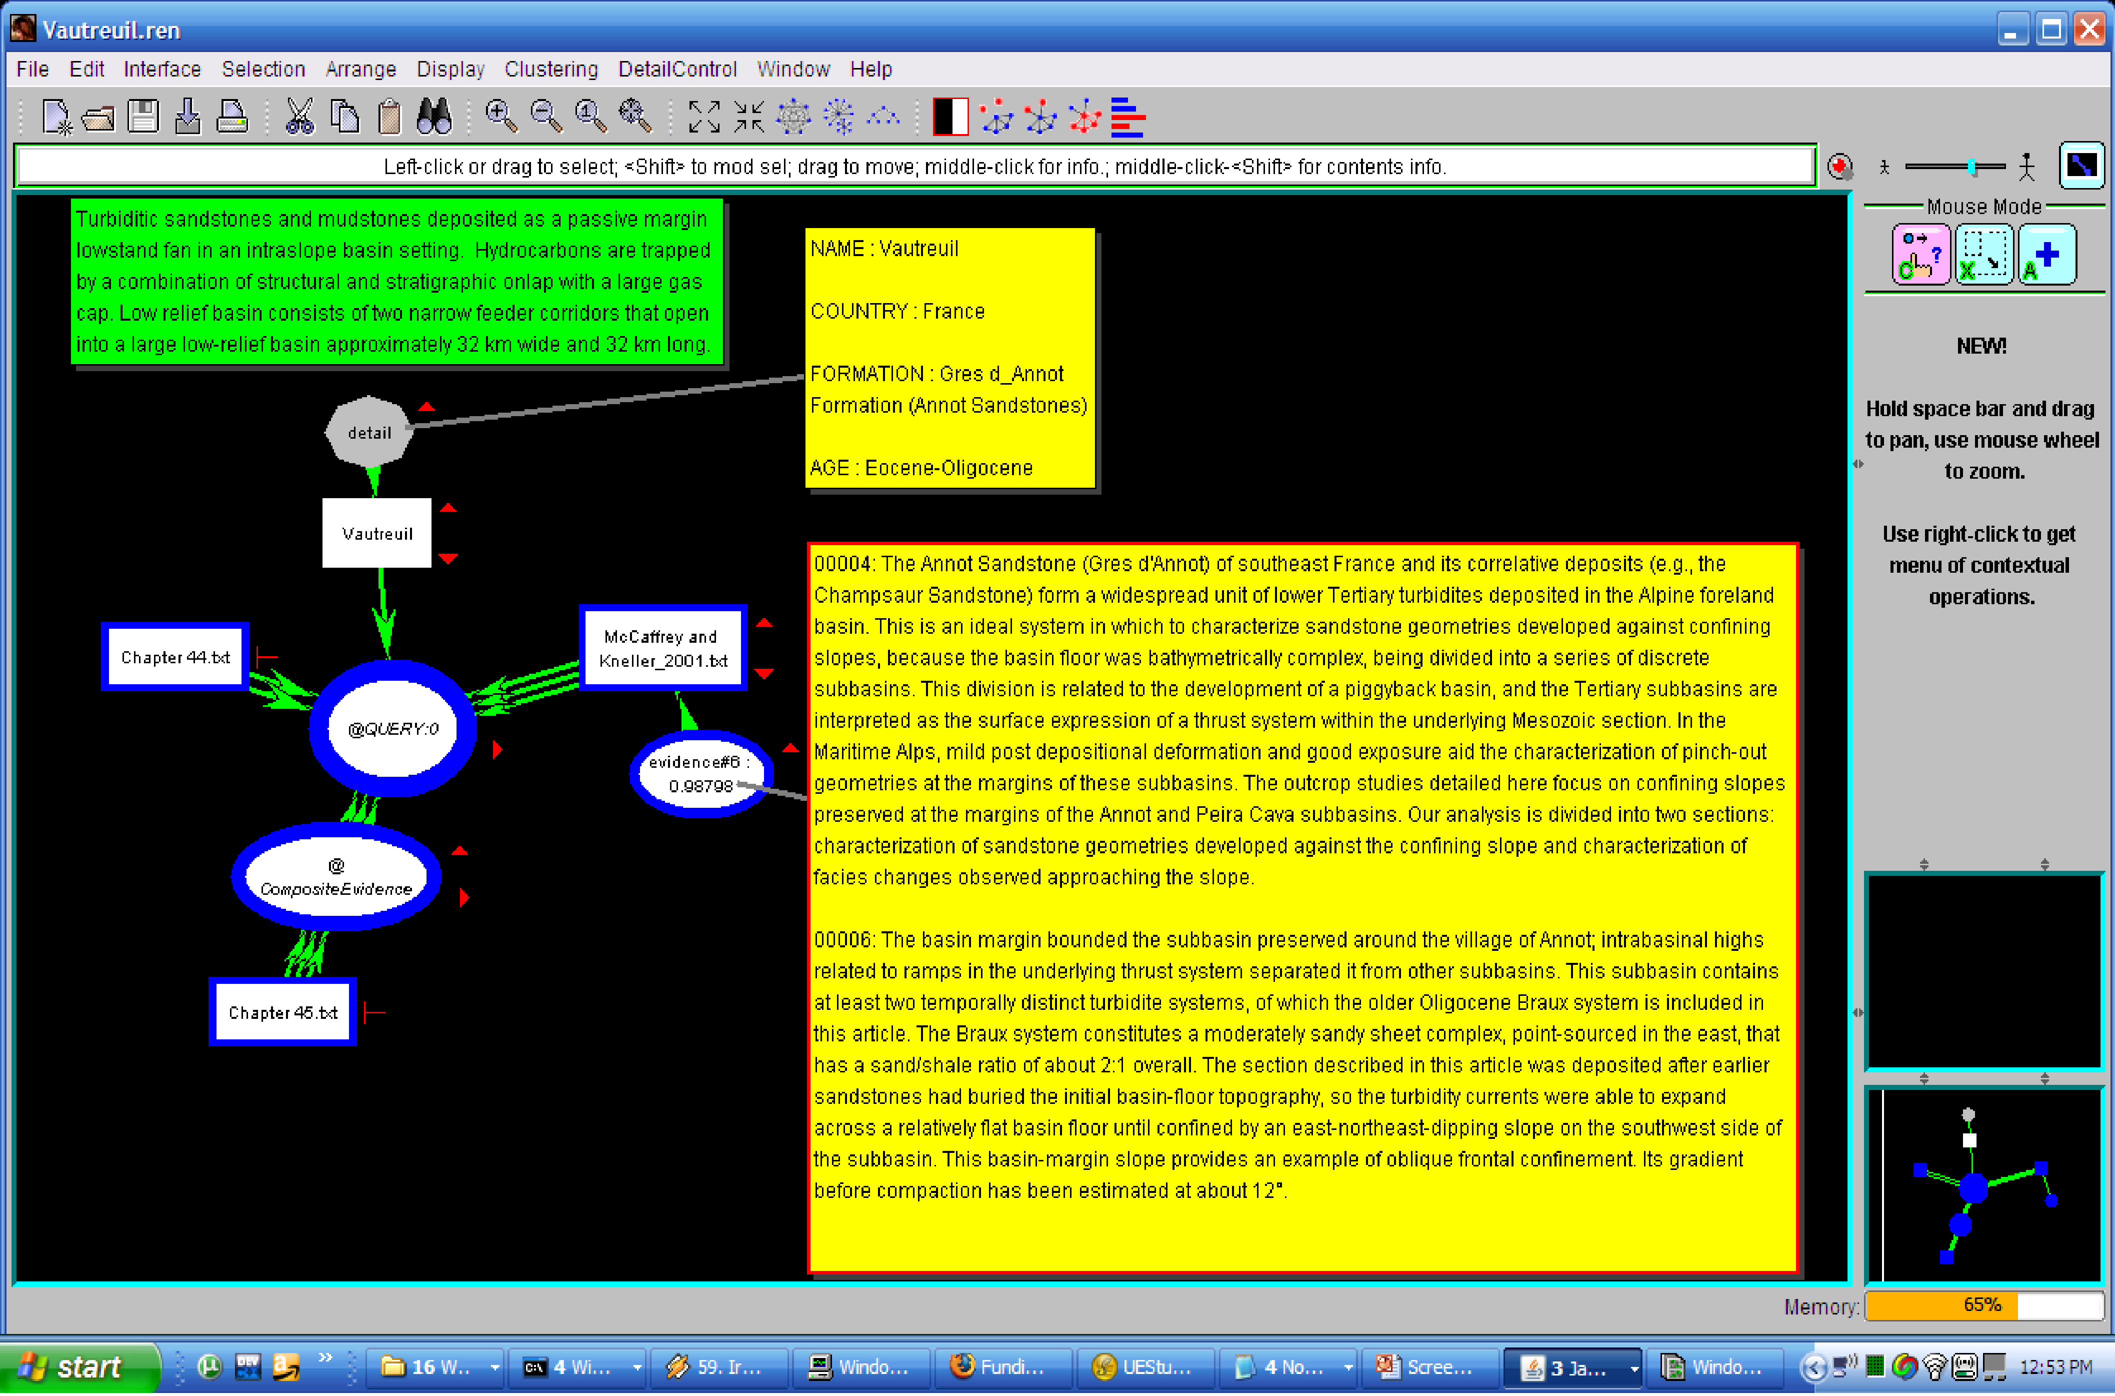

Linking the query to the paragraphs that contain the answer

What the Diagrams Show

Information shown in the previous diagram:

- 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 diagram 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*.
- Intellitex generated tentative CGs for these phrases and looked 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.

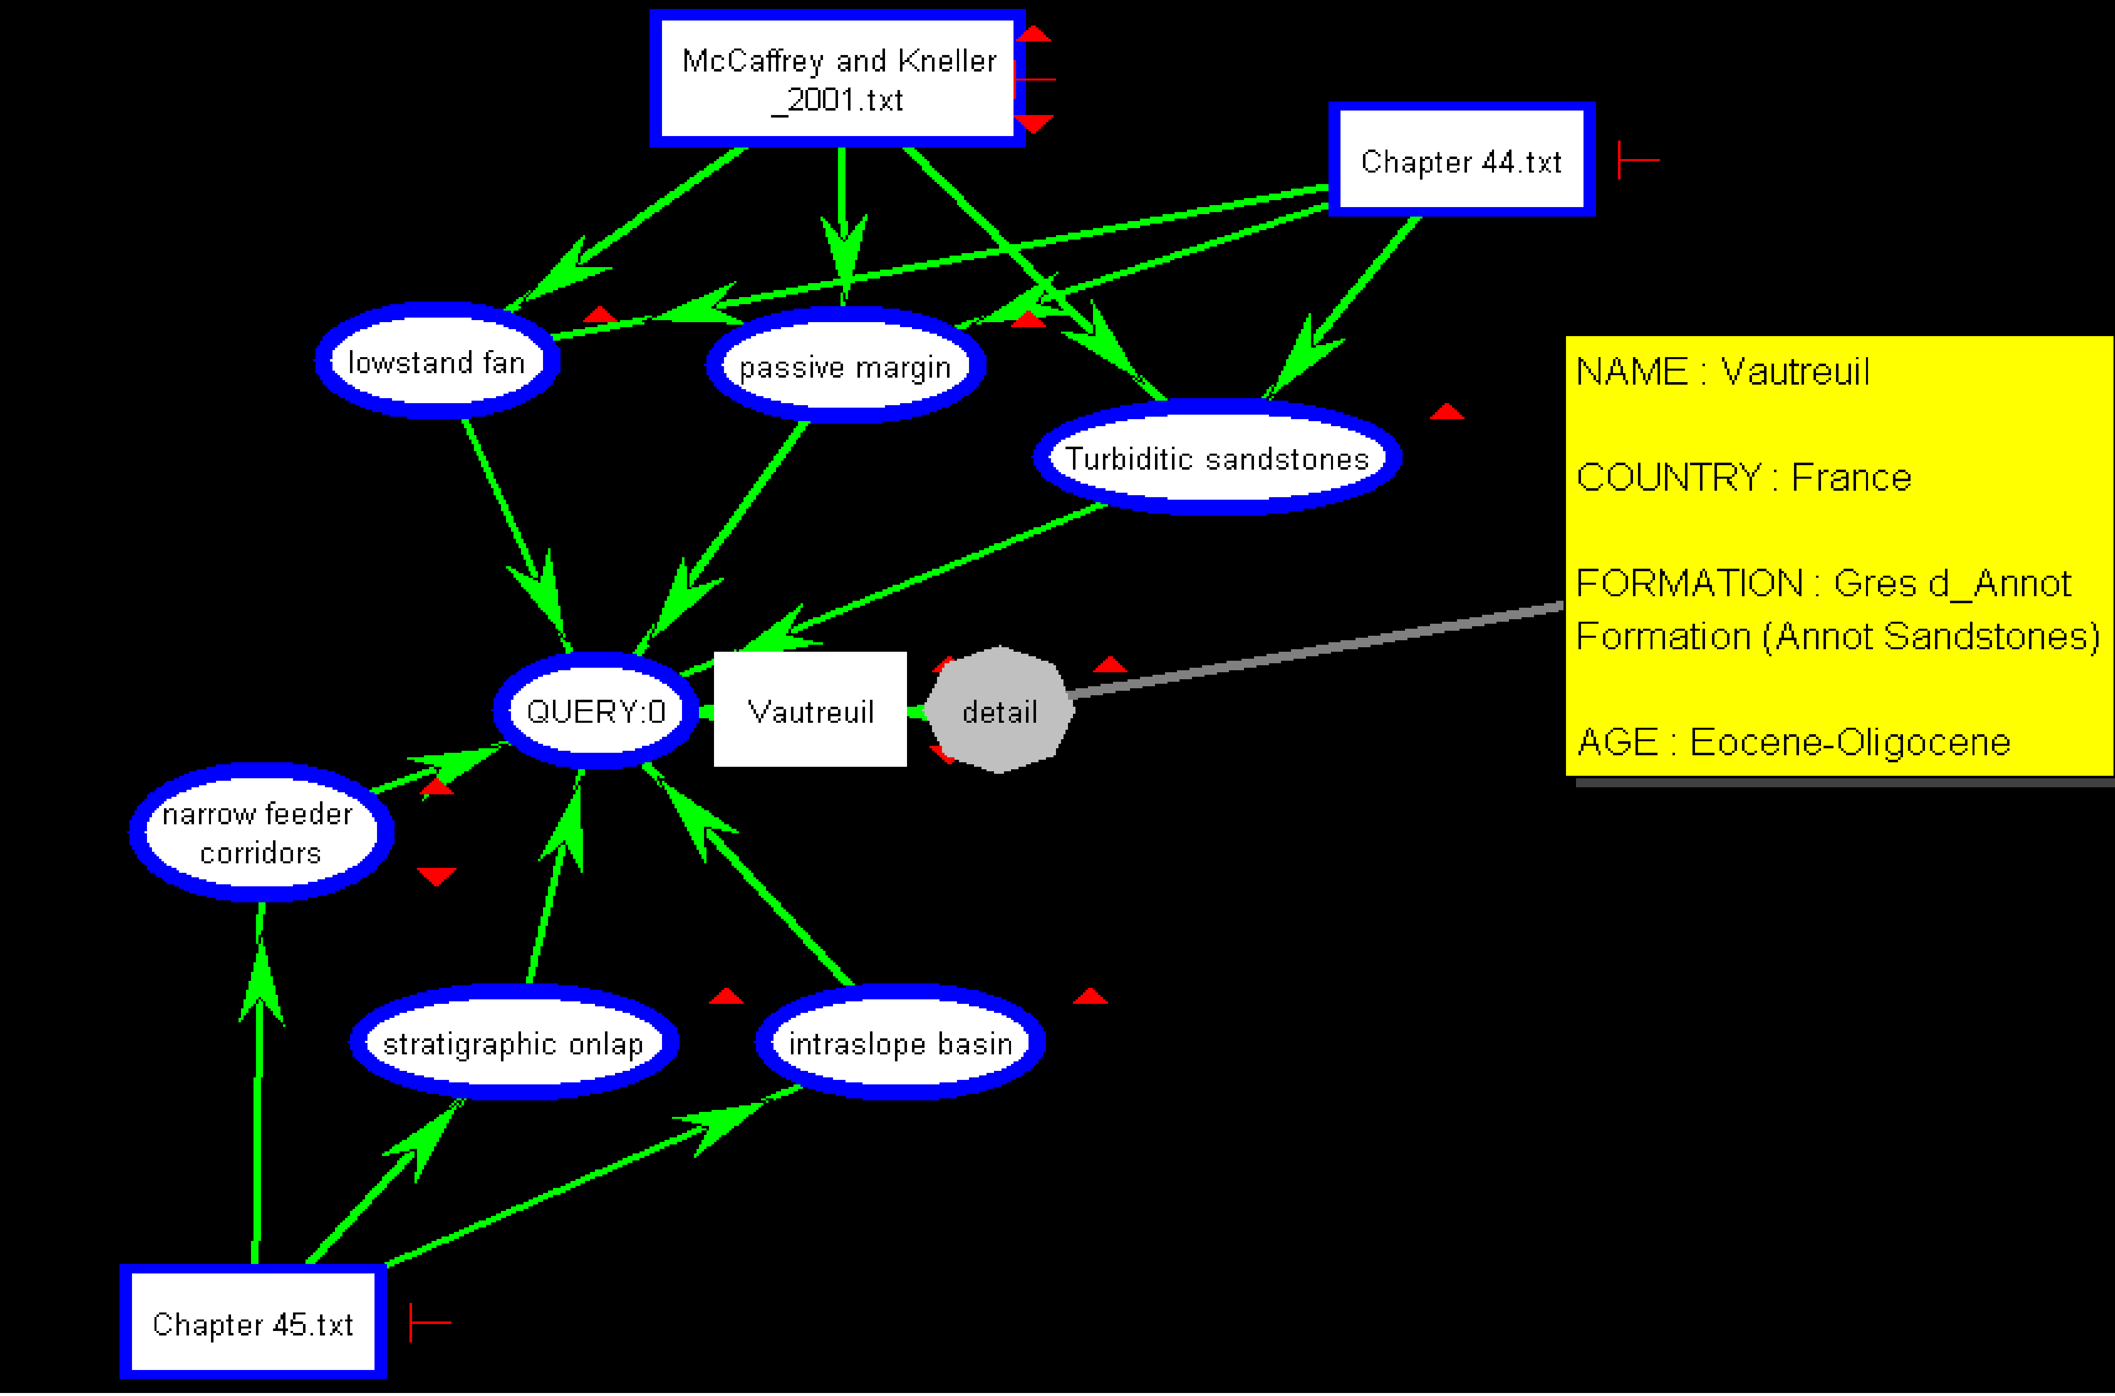

Using background knowledge from a textbook to find the answer

Emergent Knowledge

When reading the 79 documents,

- Intellitex translates the sentences and paragraphs to CGs.
- But it does not do any further analysis of the documents.

When a geologist asks a question,

- The VivoMind system may find related phrases in many sources.
- To connect those phrases, it may need to do further searches.
- Some sources can be textbooks with background knowledge that helps Intellitex interpret the research reports.
- The result consists of conceptual graphs 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, the geologist can lead the system to explore novel paths and discover unexpected patterns.

Levels of Cognition

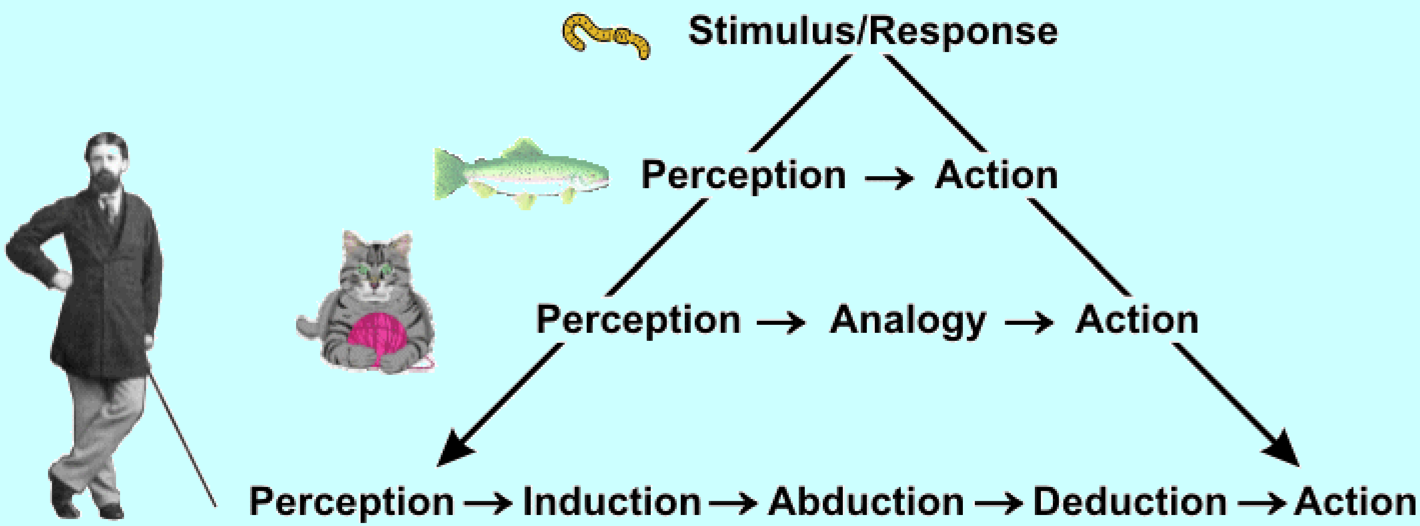

For all animals, cognition relates perception to action.

- **All forms of cognition are based on signs (semiotics).**
- **Bigger brains can process more complex patterns of signs.**

C. S. Peirce analyzed the semiotic processes at each level.

Goal of AI: Simulate the semiotic processes efficiently.

Observing, Learning, Reasoning, Acting

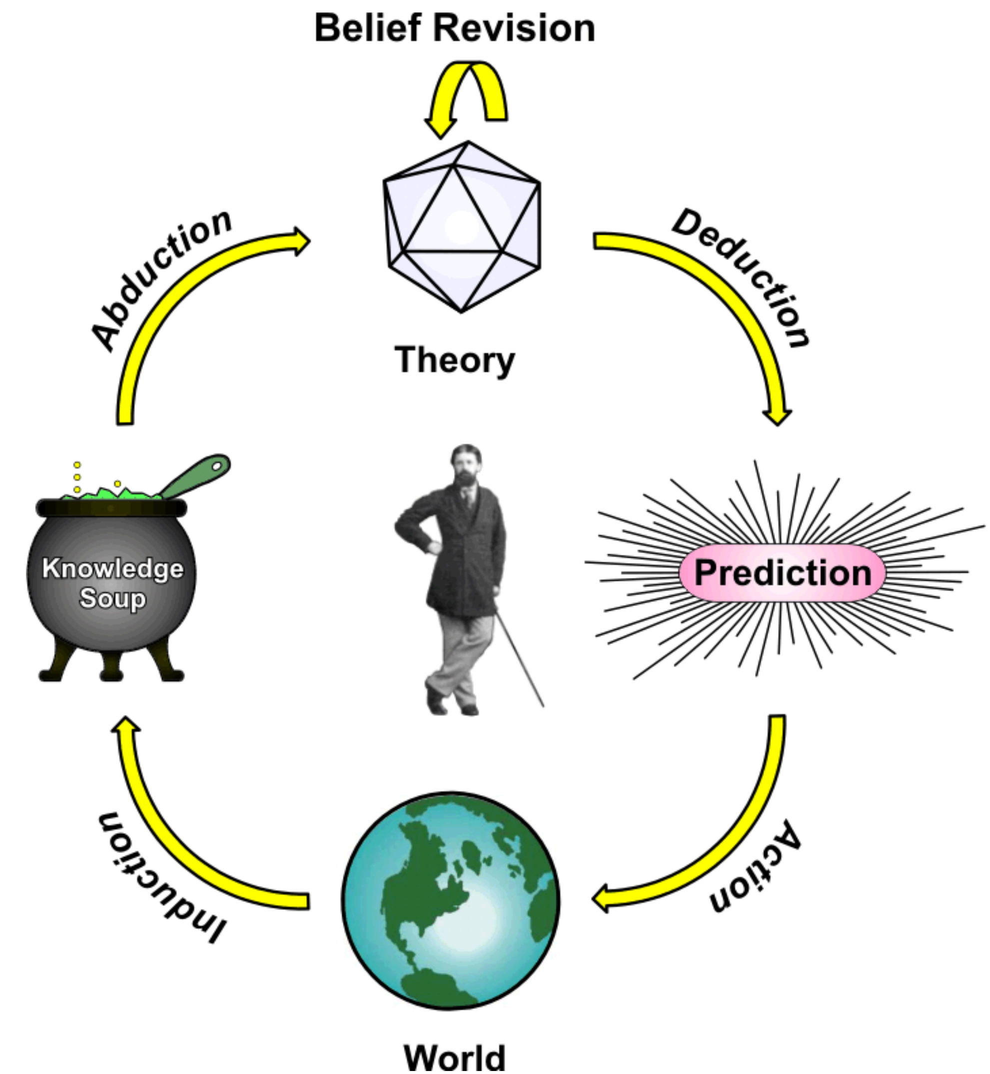

The human cycle, as described by Charles Sanders Peirce.
Similar cycles occur in every aspect of life, including science.

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.

Boyd’s OODA Loop

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

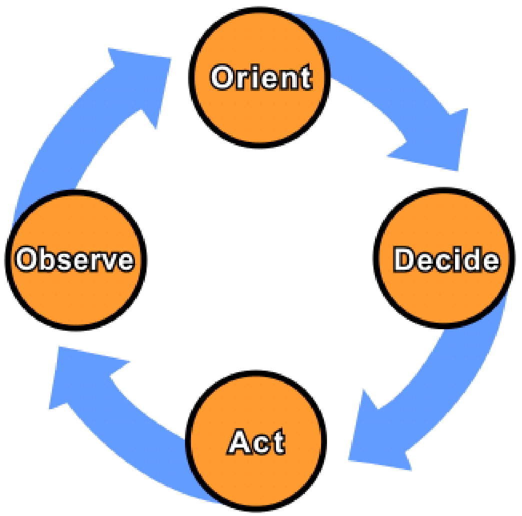

The first two steps – Observe and Orient – involve the occipital, parietal, and temporal lobes.

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 cortex relates vision and touch in “cognitive maps.”
3. Decide: Reasoning is under the control of the frontal lobes, but other areas are also involved.
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

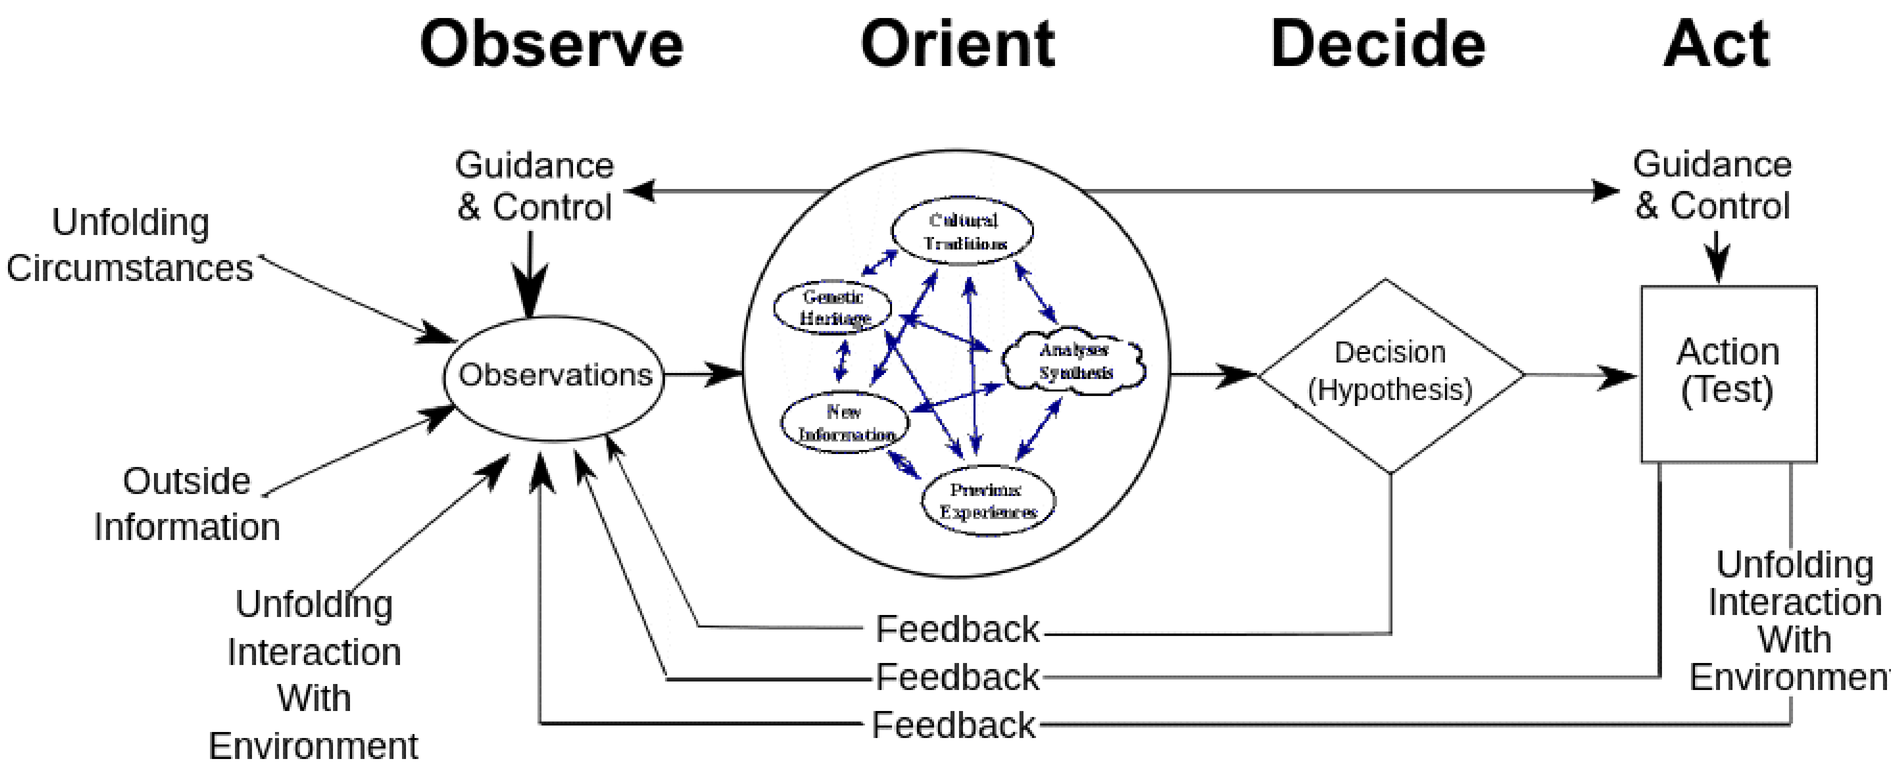

Over the years, Boyd added more detail to the OODA Loop.

He applied it to decision-making processes of any kind.

Both versions are consistent with Peirce's cycle.

Diagram adapted from http://en.wikipedia.org/wiki/OODA_loop

Ohlsson's Deep Learning Cycle

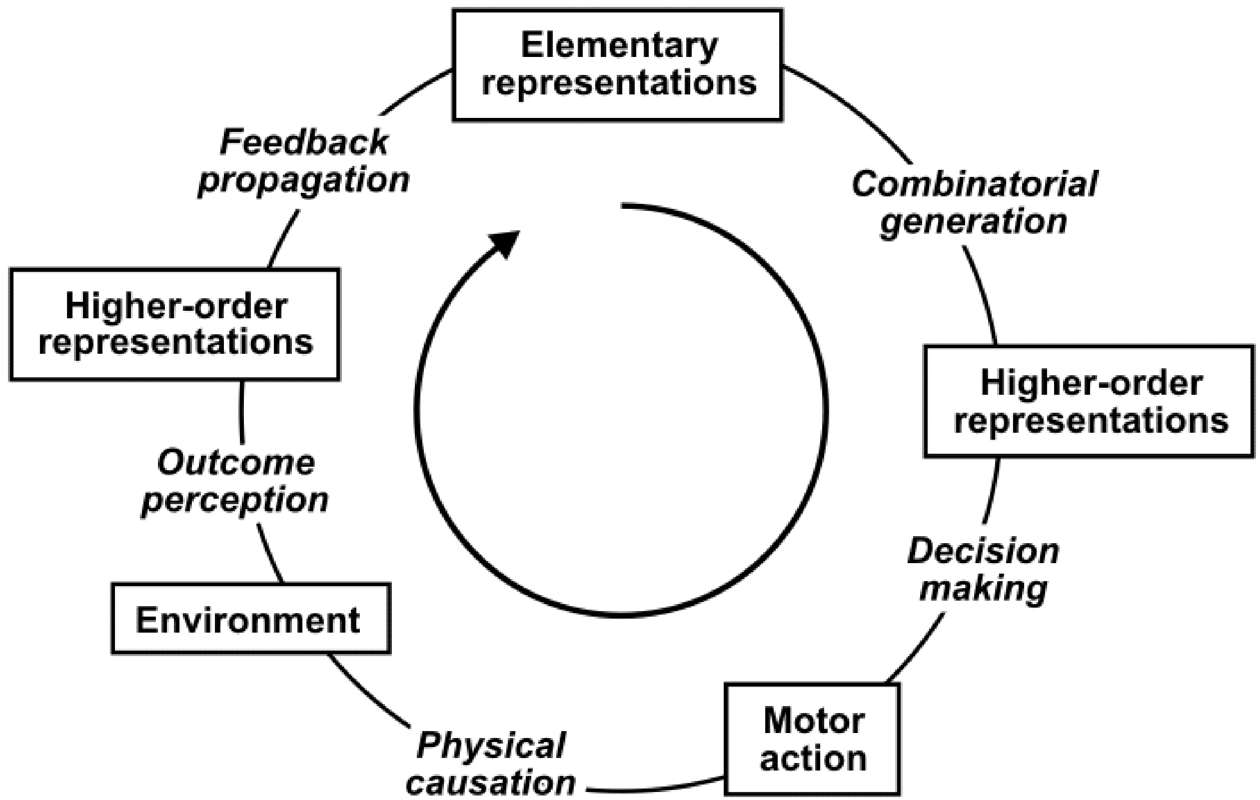

Deep learning is non-monotonic cognitive change: *

- **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

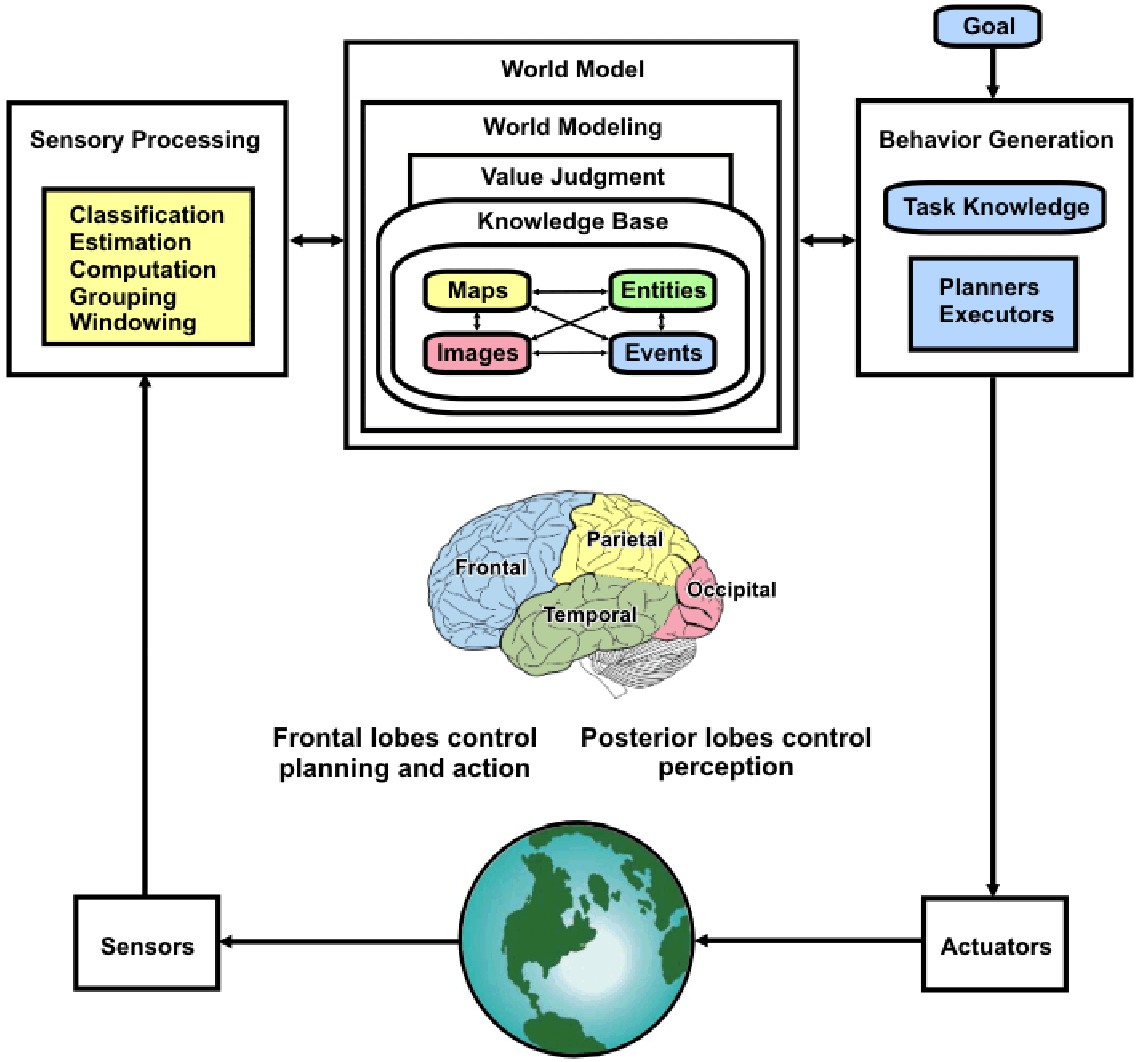

A diagram that resembles the cycles by Peirce, Boyd, and Ohlsson.

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

Albus Cognitive Architecture

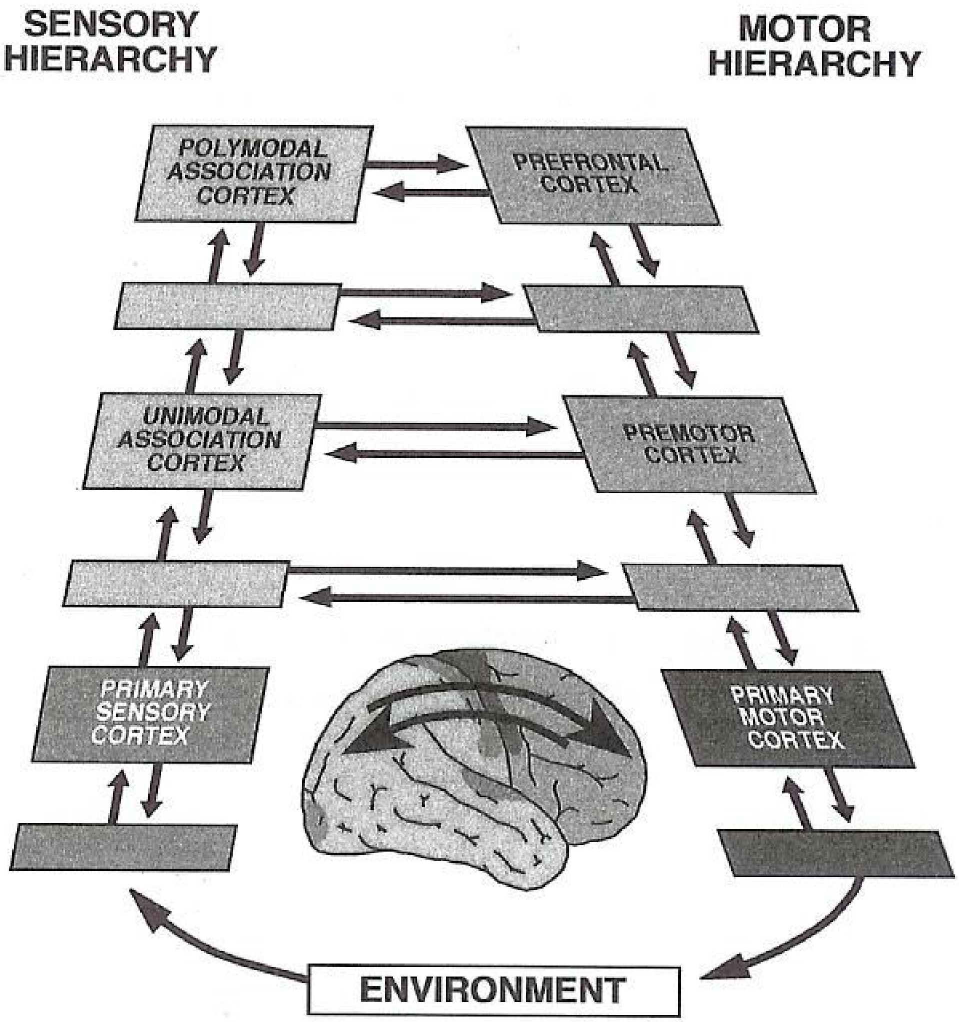

Two tightly integrated hierarchies support the cognitive cycle.

Albus adapted neuroscience to engineering — and vice versa.

Diagram from Figure 6 of http://www.james-albus.org/docs/ModelofComputation.pdf

Real-Time Control System (RCS)

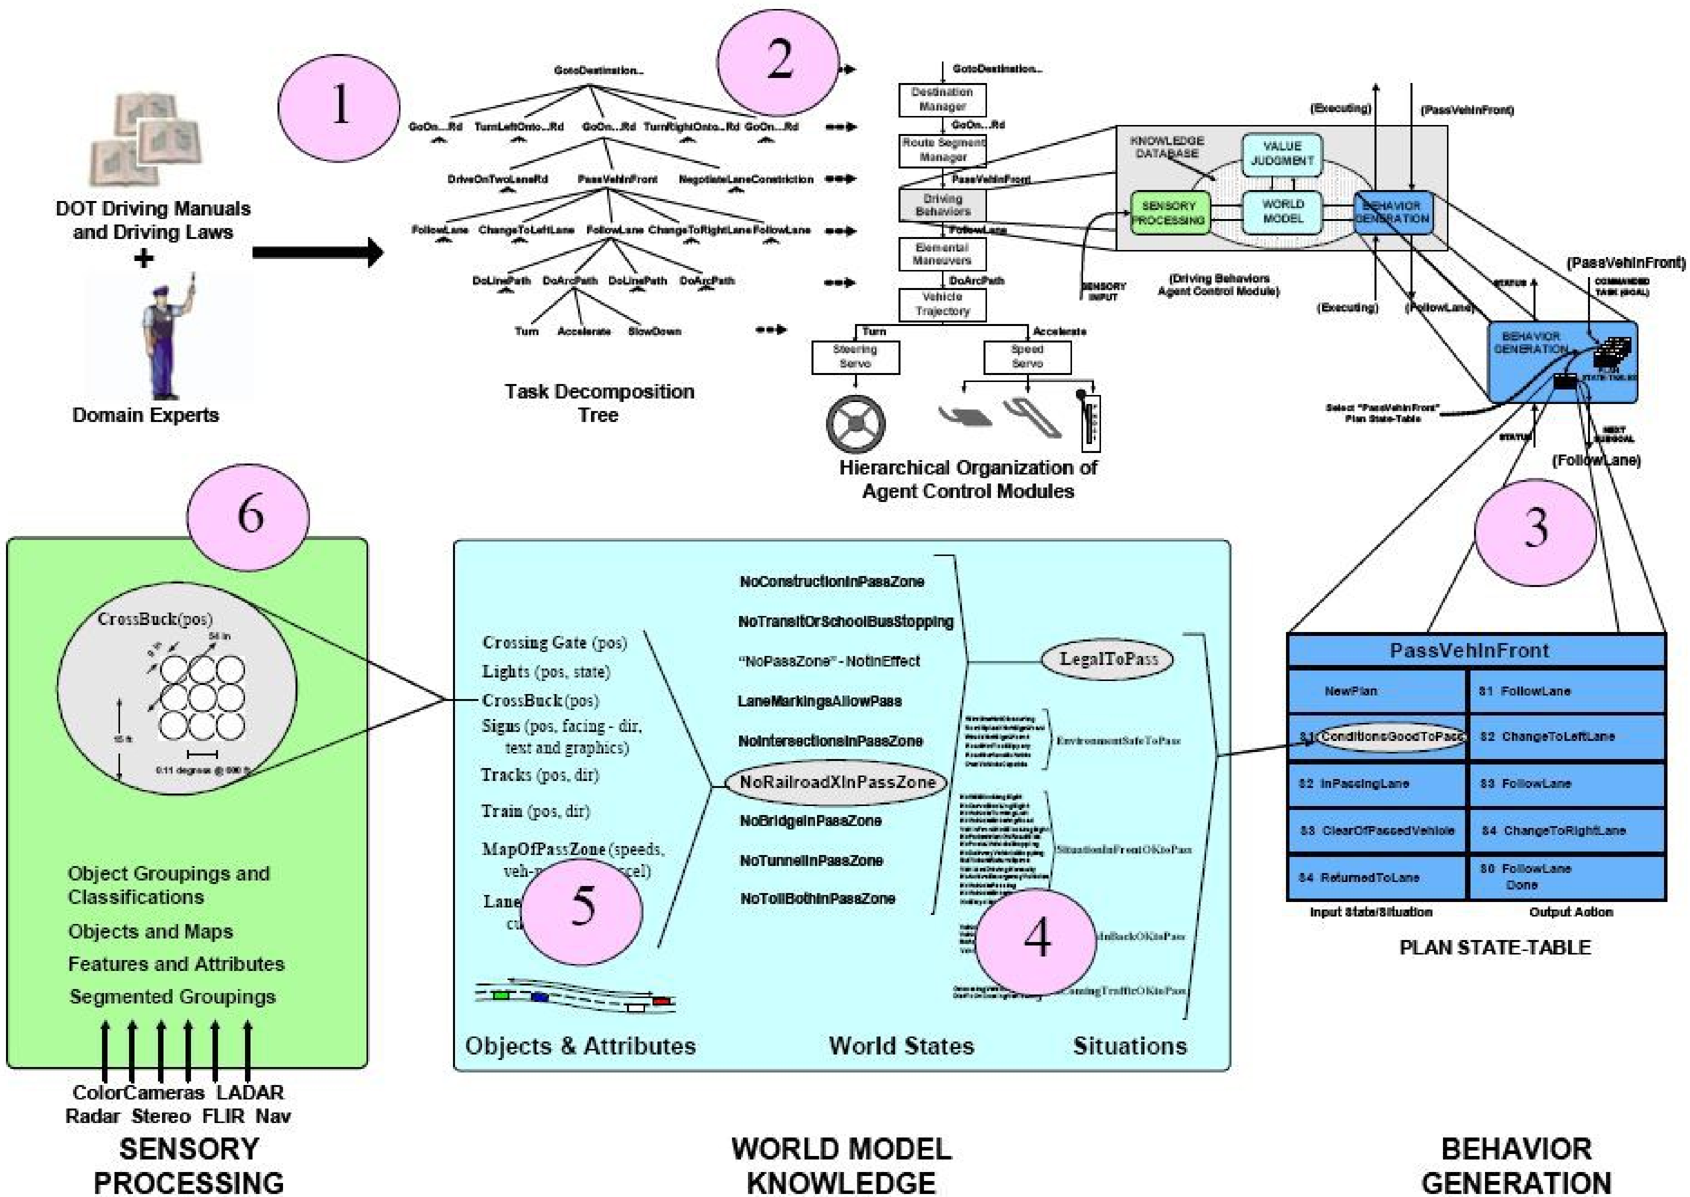

Designed by Albus and colleagues: http://en.wikipedia.org/wiki/Real-time_Control_System

RCS Computational Node

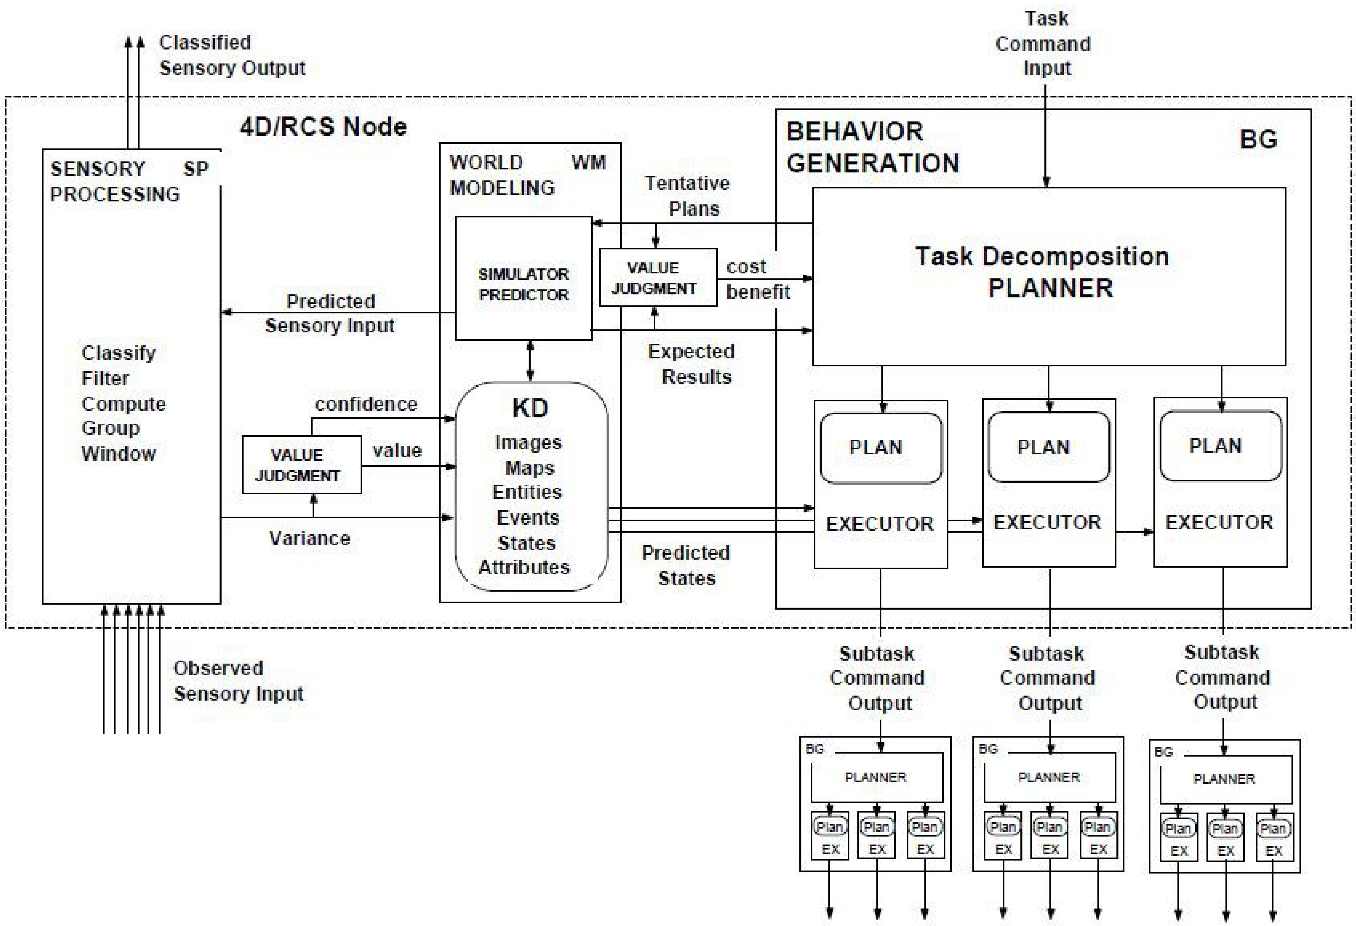

A hierarchy of computational nodes do the planning for tasks and subtasks.

Behavior generation (BG) uses a world model (WM) constructed from sensory processing (SP) and a knowledge database (KD). Value judgment (VJ) uses a cost-benefit analysis to evaluate plans in terms of expected results.

For details, see **http://www.robotictechnologyinc.com/images/upload/file/4DRCS.pdf**

Applied AI and Neuroscience

James Albus began his career as an engineer at NASA: *

- 1957 - 1958: Antenna design for the Vanguard I satellite.
- 1958 - 1969: Developed sensors and image compression hardware for satellites. Began studying AI, brain modeling, and neurophysiology.
- 1969 - 1973: Head, Cybernetics and Subsystems at NASA. Developed the Marr-Albus theory of cerebellar function for his PhD dissertation.
- 1973 - 2008: From a project manager to a Senior Fellow at NIST. Continued R & D in AI, neuroscience, and applications to robotics.
- 2008 - 2011: Senior Fellow, Krasnow Institute at George Mason University. Publishing and lecturing on his cognitive architecture.

Reverse-engineering the human brain.

- Nobody really knows how the brain works or how to simulate it.
- With his focus on applications, Albus showed how to adapt, apply, and integrate insights from AI, neuroscience, and engineering.
- He confirmed Minsky's claims about the need for multiple paradigms.

* From his CV: http://www.james-albus.org/cv.html

Implementing the Cycles

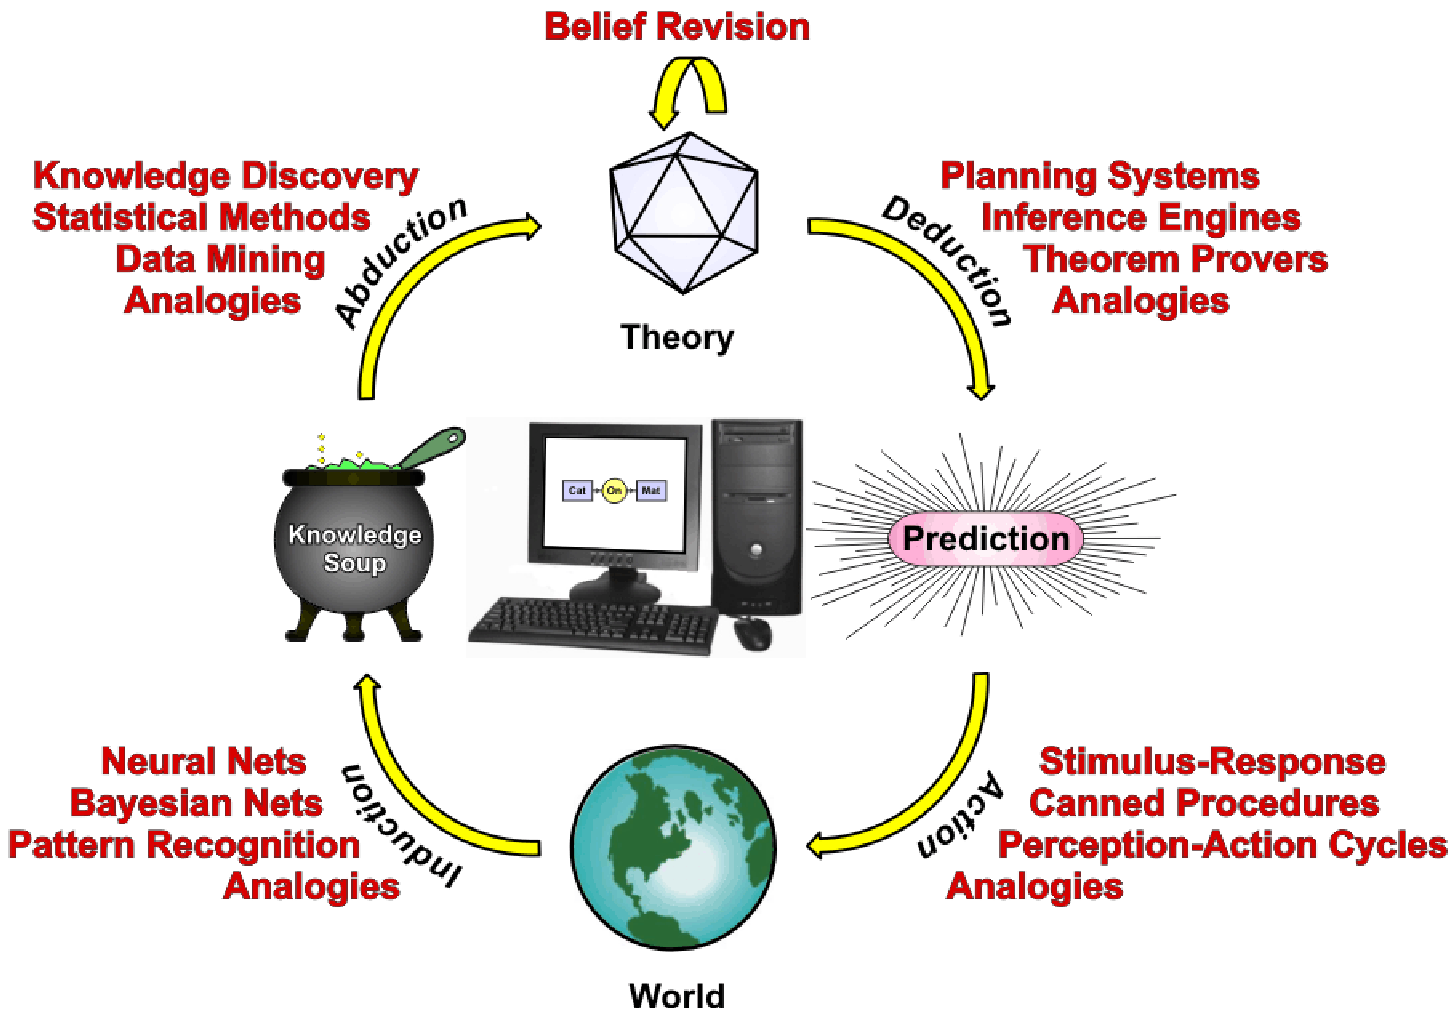

Computational methods for learning, reasoning, and acting.

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 rarely feel the need to create new patterns.
- But they can learn new patterns created by other people.

Learning: Deep, Active, and Cognitive

Geoffrey Hinton, a leader in the development of artificial neural nets, showed that multiple levels of nets are better. *

- **Nets at the early levels learn low-level features.**
- **Later levels learn combinations of the lower-level features.**
- **The multilevel version that he called *deep learning* can learn complex patterns more quickly and accurately than a single, larger net.**

Other researchers claim that *active learning* is better. **

- **Different aspects of a pattern are signficant for different purposes.**
- **Animals constantly shift their attention from one aspect to another.**
- **Active learning should use feedback from other cognitive processes.**

Learning must be integrated with the full cognitive cycle.

- **Pattern recognition is important at every step.**
- **But as Ohlsson observed, the mind can override experience.**

* **http://www.cs.toronto.edu/~hinton/**

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

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.

Future Directions for AI

Language is as general and flexible as human thought.

It requires an interpreter — human or robot — to relate a text to the current task, context, and goals.

- **That process changes the interpreter's background knowledge.**
- **But the kind of change depends critically on the task and the interpreter's goals and background knowledge.**
- **No two interpreters will understand a text in exactly the same way.**
- **With different contexts, goals, or knowledge, an interpreter may understand the same text in different ways at different times.**

For intelligent systems, the cognitive cycle is more fundamental than any particular notation or algorithm.

By integrating perception, learning, reasoning, and action, the cycle can reinvigorate AI research and development.

Related Readings

Some slides in this talk were copied or adapted from slides in the following collection: http://www.jfsowa.com/talks/goal.pdf

John F. Sowa (2014) The paper with the same title as this talk, http://www.jfsowa.com/pubs/micai.pdf

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

John F. Sowa (2006) The challenge of knowledge soup, http://www.jfsowa.com/pubs/challenge.pdf

John F. Sowa (2010) The role of logic and ontology in language and reasoning, http://www.jfsowa.com/pubs/rolelog.pdf

John F. Sowa (2011) Future directions for semantic systems, http://www.jfsowa.com/pubs/futures.pdf

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

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

References

For more information about VivoMind software:

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

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

Sowa, J. F. (2003) Laws, facts, and contexts, http://www.jfsowa.com/pubs/laws.htm

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

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

Sowa, J. F. (2009) Conceptual Graphs for Conceptual Structures, http://www.jfsowa.com/pubs/cg4cs.pdf

VivoMind also uses earlier methods for processing CGs:

Sowa, J. F. (2000) *Knowledge Representation: Logical, Philosophical, and Computational Foundations*, Brooks/Cole Publishing Co., Pacific Grove, CA

Sowa, J. F. (1999) Relating templates to language and logic, http://www.jfsowa.com/pubs/template.htm

Sowa, J. F. (1995) Syntax, semantics, and pragmatics of contexts, http://www.jfsowa.com/pubs/fs95.pdf

Sowa, J. F. (1992) Logical structures in the lexicon, http://www.jfsowa.com/pubs/loglex.pdf

For other references, see the combined bibliography for this site:

http://www.jfsowa.com/bib.htm