Computers vs. Common Sense:
The Bottleneck to Artificial Intelligence Really Helping in the Developing World

Dr. Douglas B. Lenat

Email: **Lenat@cyc.com**
Phone: 512-342-4001

"Real AI" Applications which will Redefine the Developing World

- **Education:** Individualized learning; Teaching a population to how to be effective tutors/teachers
  - False hope: Just-in-time "learning"; Passive videos
- **Healthcare:** Education; Individualized treatment; Lightening the elder/invalid/(child) care burden
  - False hope: Fully automated M.D.; Cheaper pharma.
- **Economy**: Education; Global knowledge economy; Superhuman natural language/speech understanding
  - Real hope (but not due to Real AI): Industrial robots; and decreasing importance of "uncorrected" physical location
  - False hope: Home robots; Self-driving cars; 3D printers

cYcorp

Intelligence

Have knowledge about ***x***

Able to introspect on it

Use that knowledge when it's appropriate

Red herring: Turing Test

ELIZA (DOCTOR)

1965 Joe Weizenbaum, MIT

- Carl Rogers advocated ***reflection***

  Patient: "I swear a lot."

  "How do you feel about the fact that you swear a lot?"

- Counted a lot on anthropomorphism
- Count on patient-to-doctor respect/awe/…

  Patient: "My dog's mother died recently."

  "Tell me more about your mother."

Veneer of intelligence, not intelligence

cYcorp

ELIZA (DOCTOR)

"blah blah blah **mother** blah blah blah"

Patient: "My dog's mother died recently."

"Tell me more about your mother."

Veneer of intelligence, not intelligence

cYcorp

The Knowledge Principle

A system exhibits intelligent understanding and action at a high level of competence primarily because of the ***knowledge*** it can bring to bear: the concepts, facts, representations, models, methods, metaphors, and heuristics about its domain of endeavor, and the meta-knowledge (wisdom) about when/how to apply all that.

Veneer of intelligence, not intelligence

The Breadth Hypothesis

To behave intelligently in **unexpected** situations, the problem solver often must

…fall back on increasingly general knowledge
…analogize to specific far-flung knowledge

1984: Software was Brittle

- Medical diagnosis system
  - ***Are there spots on the body? More on the trunk than elsewhere? What color are they?***

2014: Software is still Brittle

"Who was when Barrack Obama was born?"

2-Stage Research Program for AI

1. **Slowly hand-code a huge, broad ontology + KB**
2. **When enough knowledge is present, the system should actively help with the continuing expansion of its own knowledge base.**
   - **a**. Semi-automatically learn by reading online texts, DBs
   - **b**. Semi-automatically learn by gathering and examining data, formulating hypotheses, doing experiments

“Cyc”

cYcorp

Cyc's "upper model" contains:

32,000 Predicates
172,000 Collections
500,000 Individuals
13,000,000 Assertions

Represented in:

- Higher Order Logic
- Context Logic

Inference engine:

- General theorem prover
- 1050 special reasoners

Interfaces:

- Interactive English NLU/NLG dialogue
- Structured KA tools
- Automated KA
- Sem. DB Integration

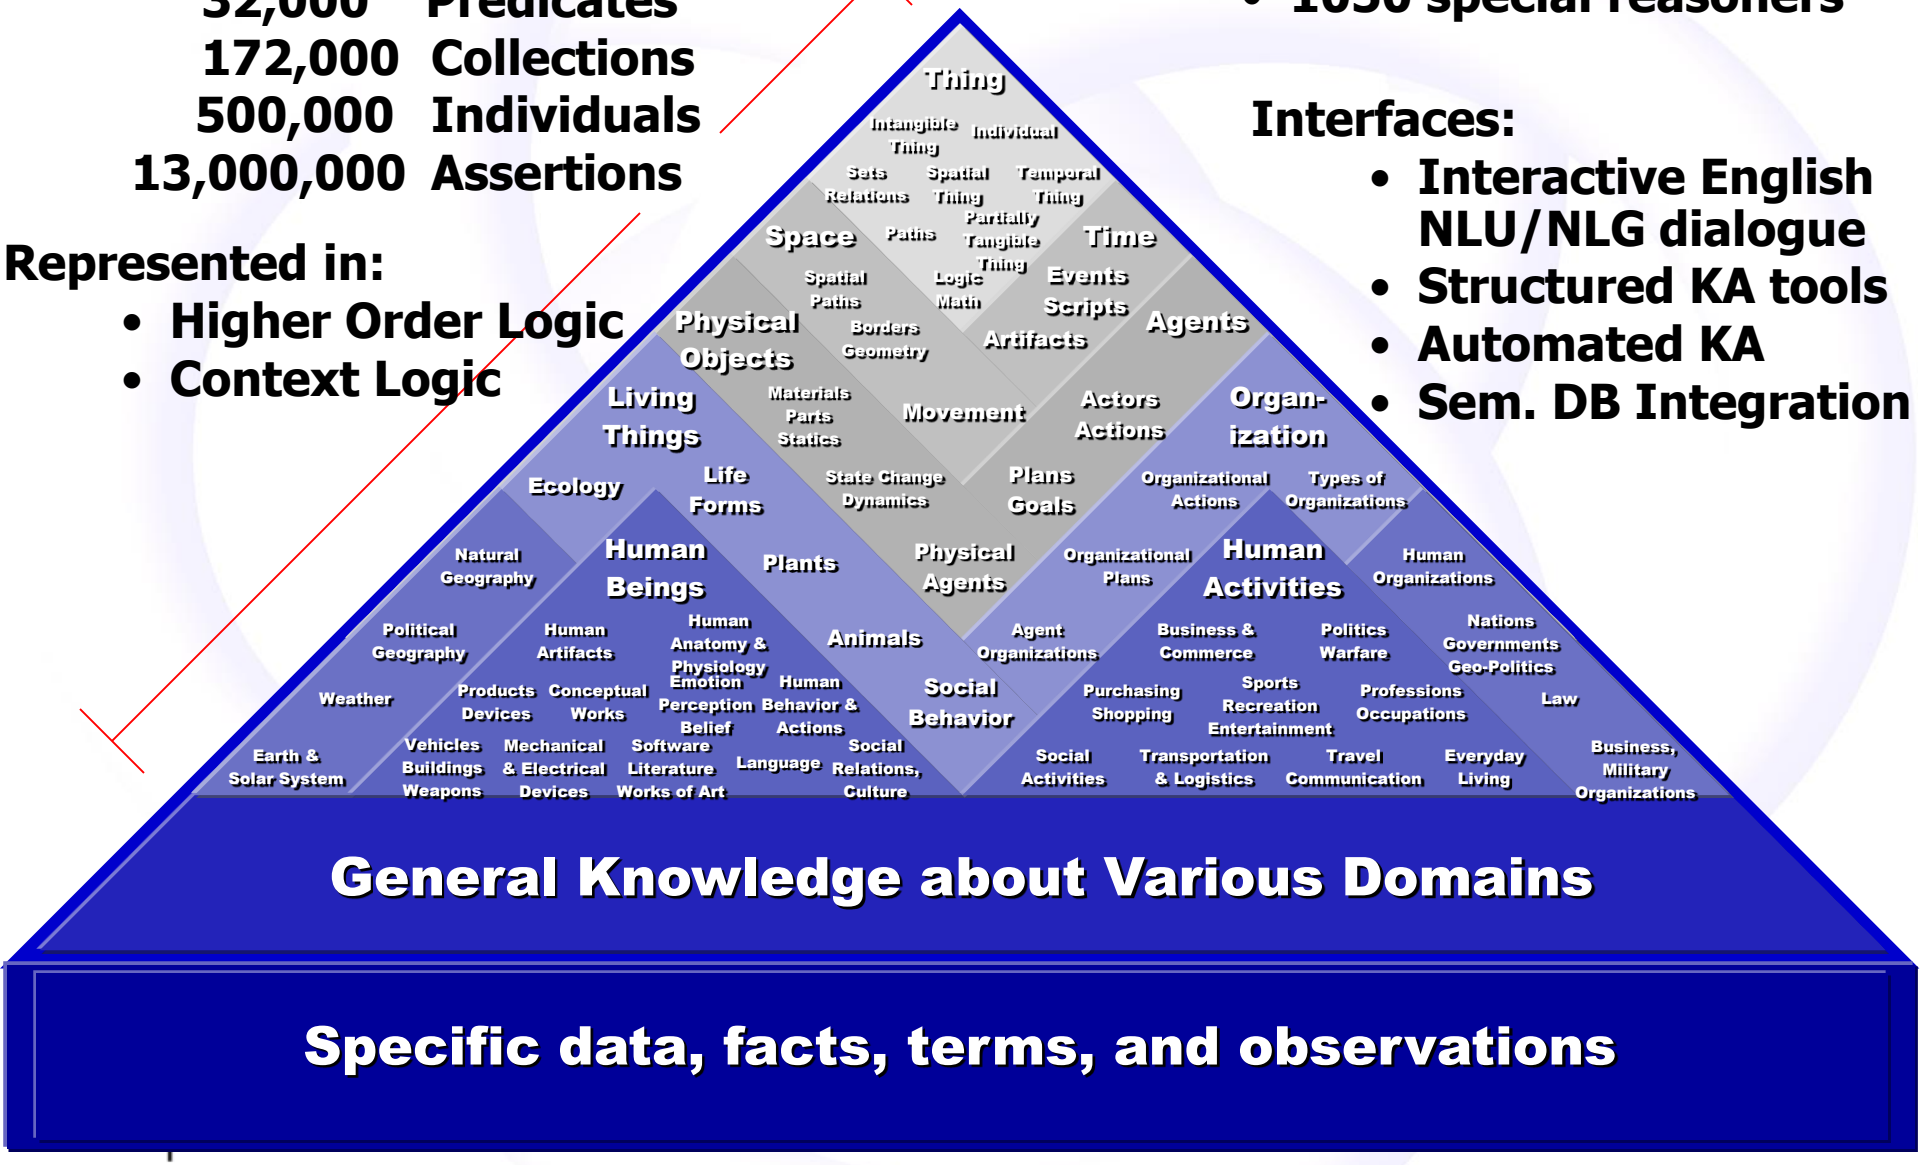

Representations of Knowledge

Each representation makes certain operations fast, at the expense of other operations being slow

- **E.g., natural language sentences**
- **E.g., node & link diagrams**
- **E.g., *n*th–order logic formulae**
- **E.g., high-resolution imagery**
- **E.g., database tuples**
- **E.g., algebraic equations**

=> NL is sometimes an inefficient representation

cYcorp

Five friends get together to play 5 doubles matches, with a different group of 4 players each time. The sums of the ages of the players for the different matches are 124, 128, 130, 136 and 142 years.
What is the age of the youngest player ?

$$v+w+x+y = 124$$
$$v+w+x+z = 128$$
$$v+w+y+z = 130$$
$$v+x+y+z = 136$$
$$w+x+y+z = 142$$

Okay, so let's tell the computer the same sorts of things that human beings know about cars, and colors, heights, movies, time, driving to a place, etc. ← all the other stuff that everybody knows.

The basic idea:

Get the computer to ***understand***, not just store, information. Then it can ***reason*** to answer your queries.

Example: Searching Text Documents

How could the results be more relevant if the search engine had some understanding?

The basic idea:

Get the computer to ***understand***, not just store, information. Then it can ***reason*** to answer your queries.

How ***formalized knowledge*** helps search

```
(ForAll ?P (ForAll ?C
      (implies
            (and
               (isa ?P Person)
               (children ?P ?C))
            (loves ?P ?C))))
```

When you become happy, you smile.

You become happy when someone you love accomplishes a milestone.

Taking one's first step is a milestone.

Parents love their children.

- Query:
  "Someone smiling"

find information by inference (+KB)

- Caption: "A man helping his daughter take her first step"

How ***formalized knowledge***
+ symbolic rule based ***inference*** helps search

find information by inference (+KB)

Analyst's Query:

"*Government buildings damaged in terrorist events in Beirut in the decade before 9/11*"

cycorp

How ***formalized knowledge***
+ symbolic rule based ***inference*** helps search

2 reports about the Khobar Towers bombing:

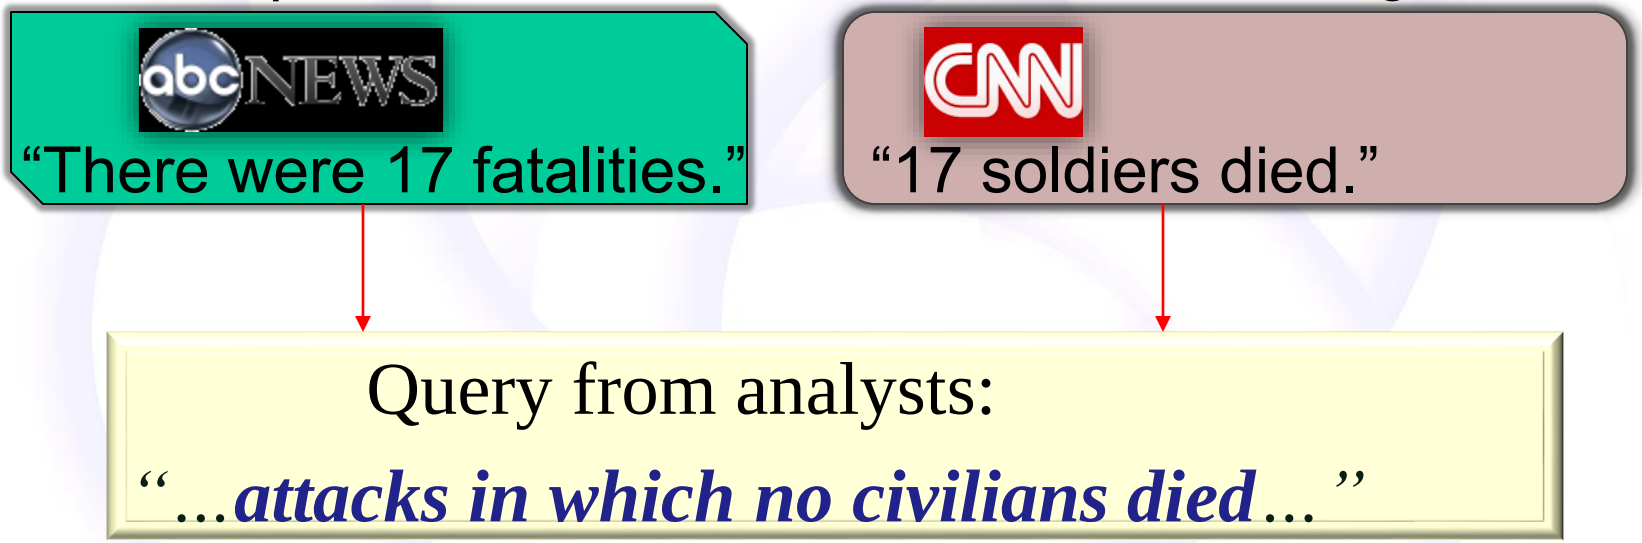

cYcorp

Disambiguation draws on knowledge

1. Tom was mad at Joe, because he stole his lunch.

   Tom was mad at Joe, so he stole his lunch.

cYcorp

Disambiguation draws on knowledge

2. John saw his brother skiing on TV.

The fool didn't have a coat on!

The fool didn't recognize him!

cYcorp

Disambiguation draws on knowledge

3. The pen is in the box.

The box is in the pen.

cYcorp

Disambiguation draws on knowledge

4. Nicole and Stephanie are sisters.

cYcorp

Disambiguation draws on knowledge

4. Nicole and Stephanie are sisters.

Nicole and Stephanie are mothers.

cYcorp

Disambiguation draws on knowledge

5. Every Fedex deliveryman has a mother.

For each Fedex deliveryman x, there exists a y, such that y is the mother of x.

Every Fedex deliveryman has an employer.

There exists a y, such that for each Fedex deliverman x, y is x's employer.

cYcorp

Disambiguation draws on knowledge

6. “and” sometimes means “or”:

   “Patients who had appendectomies in 2008 **and** 2009.”

cYcorp

Hardest lessons to learn:

- Object/attribute/value "triples" are not expressive enough → logic
- "***is related to***", "***reminds me of***",… is often dangerously shallow
-

Everyone who loves all animals is loved by someone. Anyone who kills an animal is loved by no one. Jack loves all animals. Either Jack or Curiosity killed the cat, who is named Tuna. Did Curiosity kill the cat?

$$\forall x.(\exists y : Animal\ .\ Kills(x, y)) \Rightarrow \forall z.\neg\ Loves(z, x).$$
$$\forall x : Animal\ .\ Loves(Jack, x).$$
$$Kills(Jack, Tuna) \vee Kills(Curiosity, Tuna).$$
$$Cat(Tuna).$$
$$\forall x : Cat\ .\ Animal(x).$$

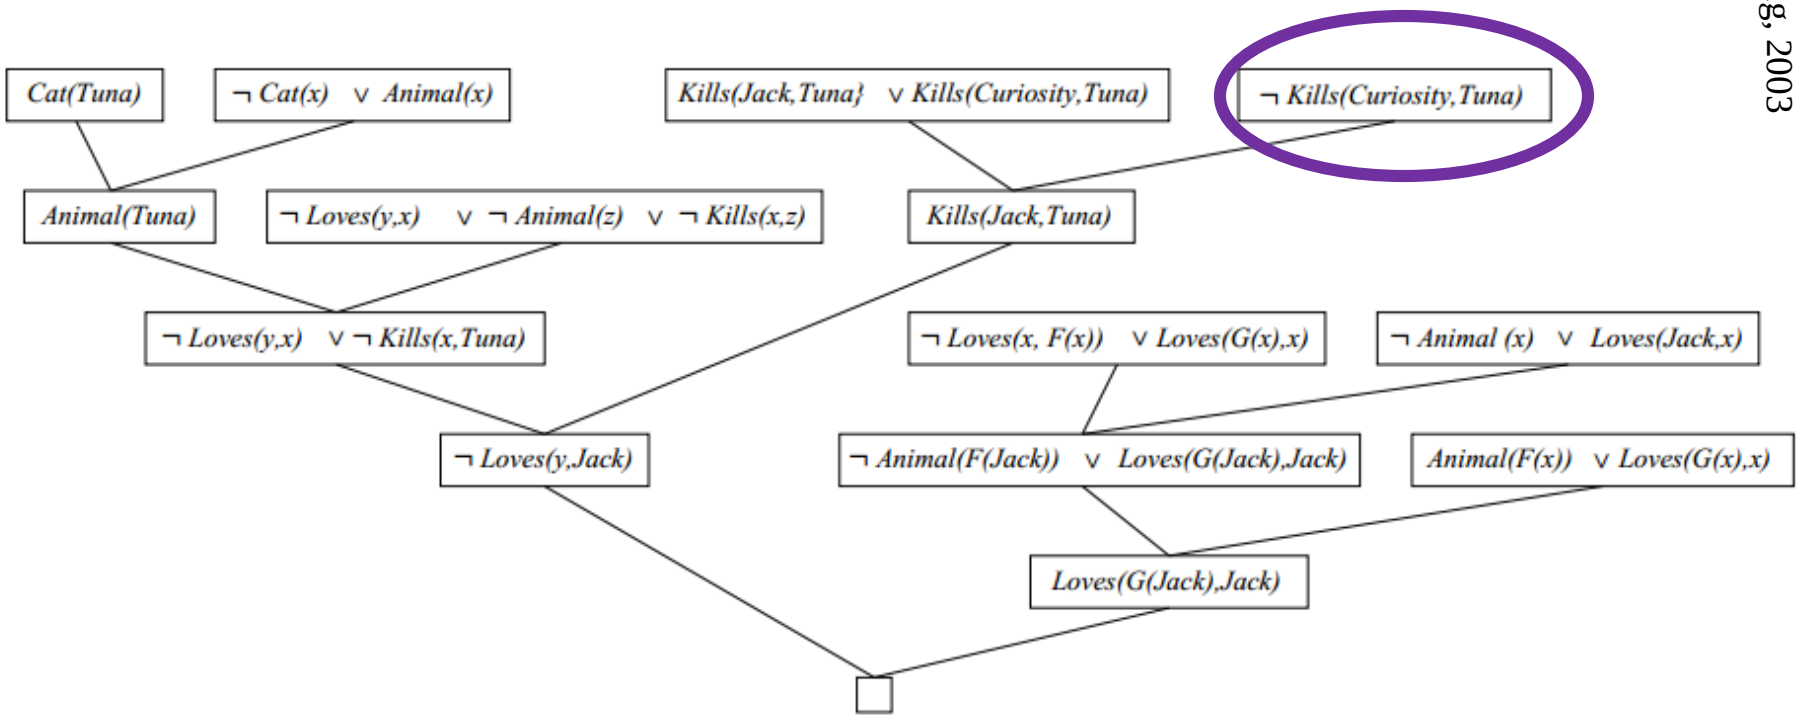

Russell and Norvig, 2003

Logical Inference

"Ariel Sharon was in Jerusalem during 2005 "

"Condoleezza Rice made a ten-day trip to Jerusalem in February of 2005"

temporal reasoning

Both of them were in Jerusalem during February 2005

diplomacy reasoning

They had a meeting, then (even if it didn't make the news.)

cYcorp

Logical Inference

2 years ago you bought a lot of diapers

Last month you bought a bicycle

You might be interested in this new product:

Toddler-carrier for bike-riders

Hardest lessons to learn:

- Object/attribute/value "triples" are not expressive enough → logic
- "***is related to***", "***reminds me of***",… is often dangerously shallow
- 99% of the knowledge is language-independent
- 99% of the knowledge is true by default, not absolutely true
- …or worse: only default-true in some context. → argumentation

So: we had to give up Global Consistency (!)

We replaced that with: Local Consistency of the KB

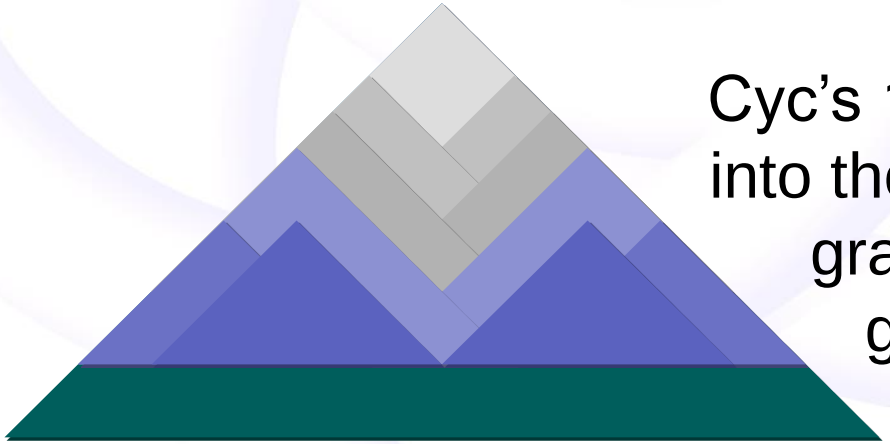

Cyc's 13M axioms are divided into thousands of *contexts* by: granularity, topic, culture, geospatial place, time,...

cYcorp

What factors argue <for/against> the conclusion that <ETA> <performed> <the March 2004 Madrid attacks>?

For:

- ETA often executes attacks near national election
- ETA has performed multi-target coordinated attacks
- Over the past 30 years, ETA performed 75% of all terrorist attacks in Spain
- Over the past 30 years, 98% of all terrorist attacks in Spain were performed by Spain-based groups, and ETA is a Spain-based group.

Against:

- ETA warns (a few minutes ahead of time) of attacks that would result in a high number civilian casualties, to prevent them. There was no such warning prior to this attack.
- ETA generally takes responsibility for its attacks, and it did not do so this time.
- ETA has never been known to falsely deny responsibility for an attack, and it did deny responsibility for this attack.

Applying Cyc

- Cyc is a power source, not a single application.

  Like oil, electricity, telephony, computers,…
  Cyc can spawn and sustain a ***knowledge utility industry***.

- It can cost-effectively underlie almost all apps.

  (Provide a common-sense layer to reduce brittleness when faced with unexpected inputs/situations)

- To apply Cyc, we **extend** its ontology, its KB, and possibly its suite of specialized reasoning modules

Semantic Searching

Cohort Finding

Typed-in Query

Partially parse

Understood Fragments

Apply domain & commonsense constraints

Combined into a Logical Query

Diagram temporal relns.

Paraphrase into English

LOUD

Consensus KB-Building

- Augment
- Manage
- Replace

5 yrs

2009

2008

Other enterprises

Other Healthcare

Other Institutes

Heart Institute Data Sources

CVIR

SQL

SPARQL

SDB

SPARQL

Other (Future) Sem. Web DBs

Upper Ontology

Middle Ontology

Domain Models

Domain Models

Web Services

Data Query

Data and Knowledge Management

Data Entry

Process Mgmt

Text, Forms, Instruments

DEMO's

- **"Who can I ask about…?"** Field your question to someone in your large organization likely to be able to answer it.
- **BELLA:** Help 6th graders understand math more deeply, by giving them a simulated "learning by teaching" game experience
- **Helpful Textbook:** Pose relevant questions to (a) help me learn (b) test my understanding

Technology Takeaway

- Symbolic Logic is a different source of power from statistics
- But the two can often cooperate synergistically

Cyc-NIH Pathway Discovery

A → B →… → Y → Z

Heuristic Reasoning over Symbolic, Causal Models

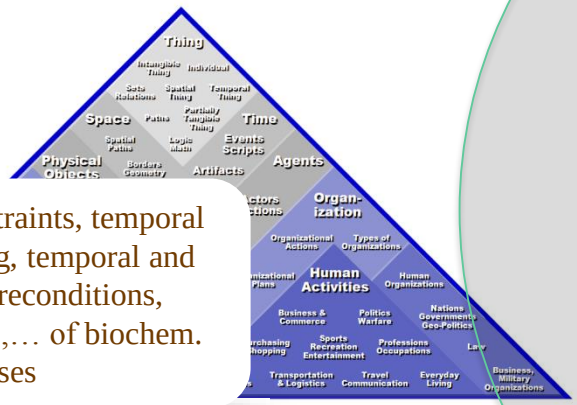

Knowledge of constraints, temporal and spatial ordering, temporal and physical parts, preconditions, consequences, rates,… of biochem. processes

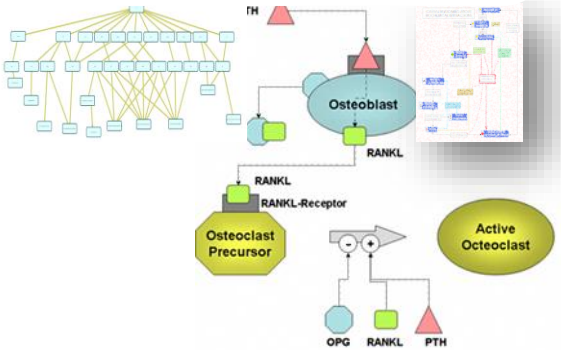

A → Z

A particular sNP correlates with a particular phenotypic condition

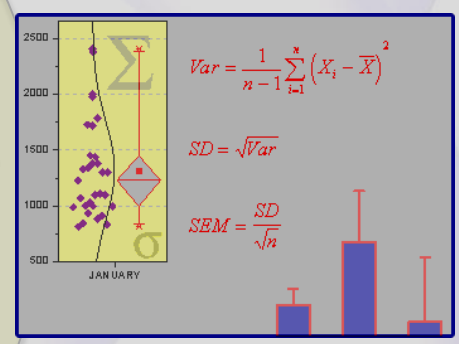

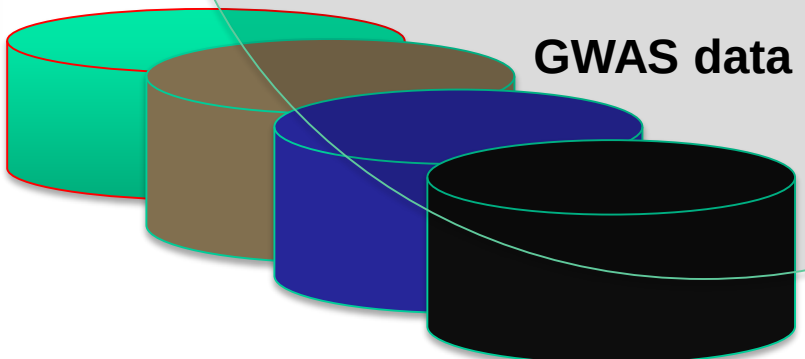

GWAS data

Cyc-NIH Pathway Discovery

A → B →… → Y → Z

Framingham SHARe

- CD40 ligand from plasma
- TNF alpha, pg/ml
- Osteoprotegerin from plas
- Vitamin D, ng/l
- Cystatin C, mg/L
- P-selectin from plasma, n
- . . . . hundreds more

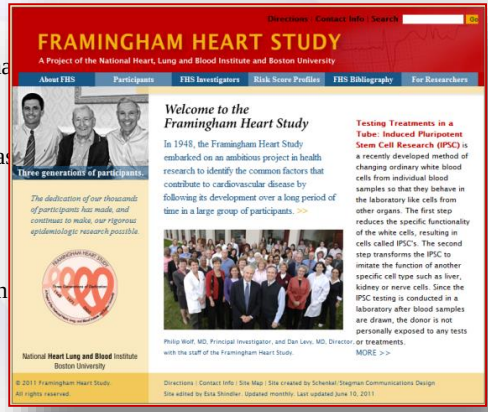

Binds to

Complexes with

Binds to chromatin

Participates in

PARD3 — SMAD3 — E2F4 — CYP27B1 — Bone Mineralization

Enzyme catalyzes

Bioactive Vitamin D

cYcorp

"Real AI" Applications which will Redefine the Developing World

- **Education:** Individualized learning; Teaching a population to how to be effective tutors/teachers
  - False hope: Just-in-time "learning"; Passive videos
- **Healthcare:** Education; Individualized treatment; Lightening the elder/invalid/(child) care burden
  - False hope: Fully automated M.D.; Cheaper pharma.
- **Economy**: Education; Global knowledge economy; Superhuman natural language/speech understanding
  - Real hope (but not due to Real AI): Industrial robots; and decreasing importance of "uncorrected" physical location
  - False hope: Home robots; Self-driving cars; 3D printers

cYcorp

What's the holdup?

- **Education:** Teaching squarely aimed at maximizing standardized test scores; Innumeracy (calculators); Misled by the asymmetry of regression to the mean
  - Google (Just-in-time "learning"); Khan (passive video)
- **Healthcare:** 99.9% conditions are cost-ineffective "orphan diseases"; STEM Ignorance ("…magic")
  - IBM (hyping Dr. Watson, M.D.); Big Pharma
- **Economy**: Knowledge Economy is a "virtuous circle" which requires resources to "spin up"; Faith that statistics (shallow AI) can and will conquer any task

"Real AI" Applications which will Redefine the Developing World

- **Education:** Individualized learning; Teaching a population to how to be effective tutors/teachers

  Zero noticeable impact for a decade, then complete game-changer

- **Healthcare:** Education; Individualized treatment; Lightening the elder/invalid/(child) care burden

  Simple apps out in the coming year, then rapid (<5 yrs) improvement

- **Economy**: Education; Global knowledge economy; Superhuman natural language/speech understanding

  "Stealthed" improvement, timing ~100% driven by competition

cYcorp