A DIGITAL ARISTOTLE (PROJECT HALO)

Motivation, basics

Eamples of questions and reasoning

Reasoning techniques

Assessment and lookout

PROJECT HALO – A DIGITAL ARISTOTLE

The vision

- **Put the bulk of the world's scientific and similar knowledge on-line**
- **Answer questions, act as personal tutor, with deep reasoning**

The task

- **College-level science**
  - **Selected as initial domain focus**
  - **Medium wide, medium deep**
  - **Good metrics available: textbook-type exam questions**

Advanced Placement Exam (AP) in Physics, Chemistry, and Biology

- **Selected as initial domain task metric focus**
- **Taken by USA high-school students to get credit for 1st-year college courses**

WHAT IS VULCAN?

- **Vulcan manages Paul Allen's assets**
  - Vulcan creates a variety of world-class endeavors and high impact initiatives that change and improve the way we live, learn, do business, and experience the world.
- **Vulcan Operating Groups**
  - Vulcan Sports and Entertainment
  - Vulcan Productions
  - Vulcan Real Estate
  - Paul G. Allen Family Foundation
  - Allen Institute for Brain Sciences
  - Vulcan Capital
  - Vulcan Technology

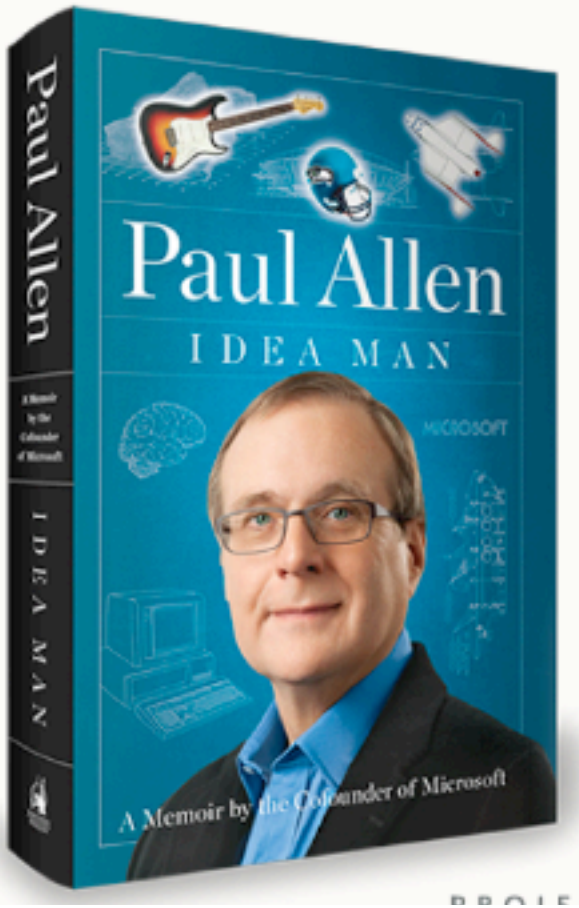

DIGITAL ARISTOTLE

PROJECT HALO®

- **Project Halo** is a staged, long-range research effort by Vulcan Inc. towards the development of a Digital Aristotle
- **Digital Aristotle** – a reasoning system capable of answering novel questions and solving advanced problems in a broad range of scientific disciplines.
- The project envisions two primary classes of application:
  - A **tutor** capable of instructing and assessing students in those subjects
  - A **research assistant** with broad, interdisciplinary skills to help scientists and others in their work

A DIGITAL ARISTOTLE – INGREDIENTS

Major tasks

- **Query understanding**
- **Knowledge acquisition**
- **Reasoning**
- **Explanation/presentation**

Which ones are doable, extesible, bottlenecks, hopeless?

Which ones are indispensible, which ones can be "simplified"?

A DIGITAL ARISTOTLE – INGREDIENTS

Major tasks

• **Query understanding**	***hopeless***	***simplified by "translation"***
• **Knowledge acquisition**	***bottleneck***	***major emphasis later***
• **Reasoning**	***extensible***	***most technical innovations***
• **Explanation/presentation**	***doable***	***quality loss permitted***

REQUIREMENTS ANALYSIS - REASONING

Data material

- **Practice exams - Princeton review guide**
  A scanned copy of each sample exam,
  detailed answers to all questions of each sample exam, and
  a copy of our analysis of each complete example exam are available online
  (http://www.ai.sri.com/halo/public/fois2010/

Domains/structure

- **Physics, chemistry, biology, environmental sciences,**
  microeconomics, US Government & Politics
- **90-95% multiple choice questions, some free response questions**
- **Analysts with KR & R as well as domain expertise**
- **Abstracting away from NLP analysis and open-ended world knowledge**

EXAMPLE 1 - PHYSICS

Typical challenge/knowledge required

- **Determine the suitable model to apply to the problem at hand**
- **Reasoning with structured objects, equations**

Example

How much force is required to lift a 50-newton object with an acceleration of 10 m/s^2?

(A) 10 N
(B) 50 N
(C) 100 N
(D) 150 N
(E) 200 N

Some specific knowledge required

- **Numeric solution of equations, in numbers and units**
- **Default assumption - lifting is vertical, establishing a free body diagram**

EXAMPLE 2 - CHEMISTRY

Typical challenge/knowledge required

- **Similar as with physics, mostly classification reasoning, also logic program rules**

Example

A 100 ml sample of 0.10 molar NaOH solution was added to 100 ml of 0.10 molar H3C6H5O7. After equilibrium was established, which of the ions listed below was present in the greatest concentration:

(A) H2C6H5O7-
(B) HC6H5O72-
(C) C6H5O73-
(D) OH-
(E) H+

Some specific knowledge required

- **Computational knowledge (in equational problem solving)**

EXAMPLE 3 - BIOLOGY

Typical challenge/knowledge required

- **Process representation and reasoning**
  Inspecting process structure and characteristics
- **Discrete, nonnumeric representations and semantic matching**

Example

Which of the following is NOT a characteristic of bacteria?

(A) Circular double-stranded DNA
(B) Membrane-bound cellular organelles
(C) Plasma membrane consisting of lipids and proteins
(D) Ribosomes that synthesize polypeptides

Some specific knowledge required

- **Structured object representation of bacteria**

EXAMPLE 4 - ENVIRONMENTAL SCIENCE

Typical challenge/knowledge required

- **Rich representations of qualitative processes, classifying a specific situation**
- **Comparative analysis of two or more possible behaviors**

Example

Which of the following organisms is the first to be adversely affected by thermal pollution in a stream?

(A) Trees along the bank
(B) Insect larvae in the water
(C) Large fish migrating up stream
(D) Birds drinking the water
(E) Bacteria in the water

Some specific knowledge required

- **Positioning organisms w.r.t. water, and how they use the water**

EXAMPLE 5 - MICROECONOMICS

Typical challenge/knowledge required

- **Reasoning about mathematical functions, and their qualitative properties**
- **Switching back and forth between mathematical and geometrical representations**

Example

What are the effects on supply and demand curves for Frisbees if a new procedure reduces the cost of making a Frisbee®?

Demand	*Supply*
(A) Shifts right	**Shifts right**
(B) No change	**Shifts left**
(C) No change	***Shifts right***
(D) Shifts left	**Shifts right**
(E) Shifts right	**Shifts left**

Some specific knowledge required

- **Effect on demand by default - no effects described that might have impacts**

EXAMPLE 6 - US GOVERNMENT & POLITICS

Typical challenge/knowledge required

- **Heuristic methods for assessing the likelihood of possible answers**
- **Matching domain knowledge available with knowledge referred to**

Example

Which of the following correctly states the relationship between the federal and state judiciaries?

(A) Federal courts are higher courts than state courts and may overturn state decisions on any grounds.

(B) The two are entirely autonomous, and neither ever hears cases that originate in the other.

(C) The two are generally autonomous, although federal courts may rule on the constitutionality of state court decisions.

(D) State courts are trial courts; federal courts are appeal courts.

(E) State courts try all cases except those that involve conflicts between two states, which are tried in federal courts.

Some specific knowledge required

- **Need examples to reason from or applicable axioms**

CATEGORIES OF REASONING NEEDED (MAJOR ONES)

KR&R Category	*Physics*	*Chemistry*	*Biology*	*Envir.*	*Politics*	*Microec.*
Diagrams	**30**	**7**	**14**			**8**
Processes	**3**	**4**	**44**	**72**	**1**	
Vague knowledge						**2**
Qualitative modeling	**3**	**20**	**6**	**72**		
Computational knowledge		**19**	**8**			
Spatial relations and directions	**4**	**1**				
Simultaneous equations	**1**	**6**				
Qualitative reasoning with equations	**18**	**5**	**1**			**6**
Symbolic reasoning with equations	**19**	**18**	**1**			
Numerical solution of equations	**24**	**4**	**2**			
Theorem proving					**9**	
Reasoning by elimination	**8**	**5**	**6**			
Making assumptions	**10**	**2**				
Compare and contrast			**9**	**4**	**3**	**1**
Qualitative reasoning	**2**	**19**		**30**		**15**

AURA - EXPERT SYSTEM DEVELOPED

Technology

- **Novel combination of available techniques from AI**
- **Controlled Natural Language, GUI, Frame-based KR, Problem-Solving**
- **Students as users**
- **Initial version 2004, then refined extensively and tested rigorously**
- **For recent info, see http://www.ai.sri.com/project/aura**

A major challenge

- **How to enable Subject Matter Experts(SMEs), rather than Knowledge Engineers (KEs), to interact to:**
  - **State questions**
  - **Understand answers and explanations**
- **How to reduce cost via KA by SMEs not KEs**
  - **Scaleabilityw.r.t. domains**
  - **Incrementally refine the knowledge base**
  - **Lower the cost of textbook page of knowledge that users enter, edit, and utilize**

AURA SYSTEM ARCHITECTURE

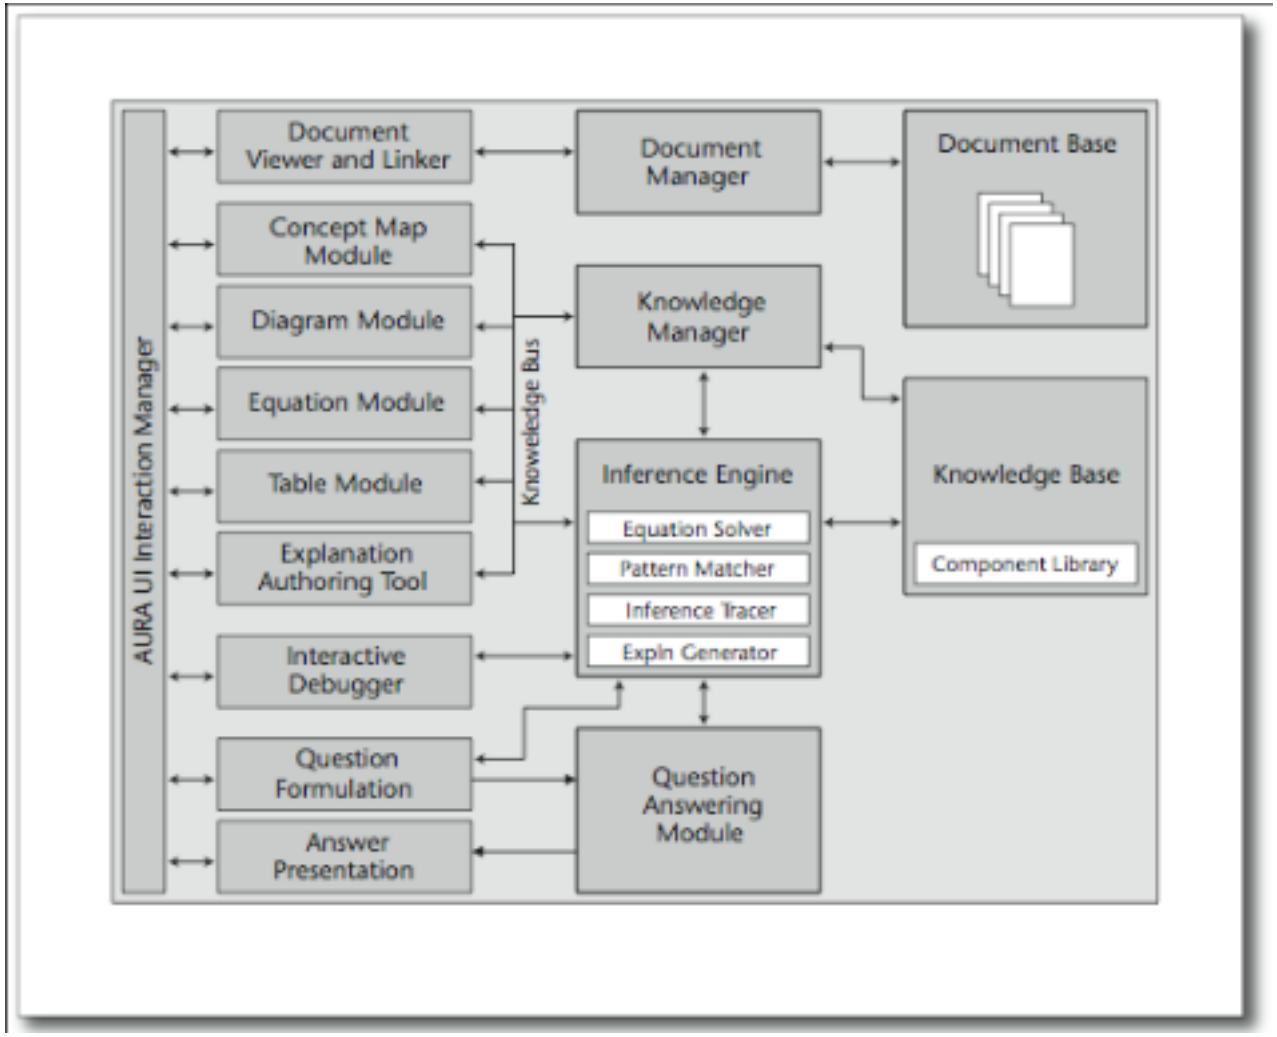

KNOWLEDGE REPRESENTATION/REASONING

Components

- **State-of-the-art frame-based knowledge representation system**
- **Prototypes**
- **Unification mapping**

Query interpretation

- **Query pattern as a logical formula to be proven**
- **Semantic role labeling**
- **Rules for producing underspecified versions (droping unknown modifiers)**

Question answering

- **Query formulated in logic in the vocabulary of the knowledge base**
- **Prototypes**
- **Unification mapping**

EXAMPLE 1 - MITOSIS (BIOLOGY)

EXAMPLE 2 - FREE FALL (PHYSICS)

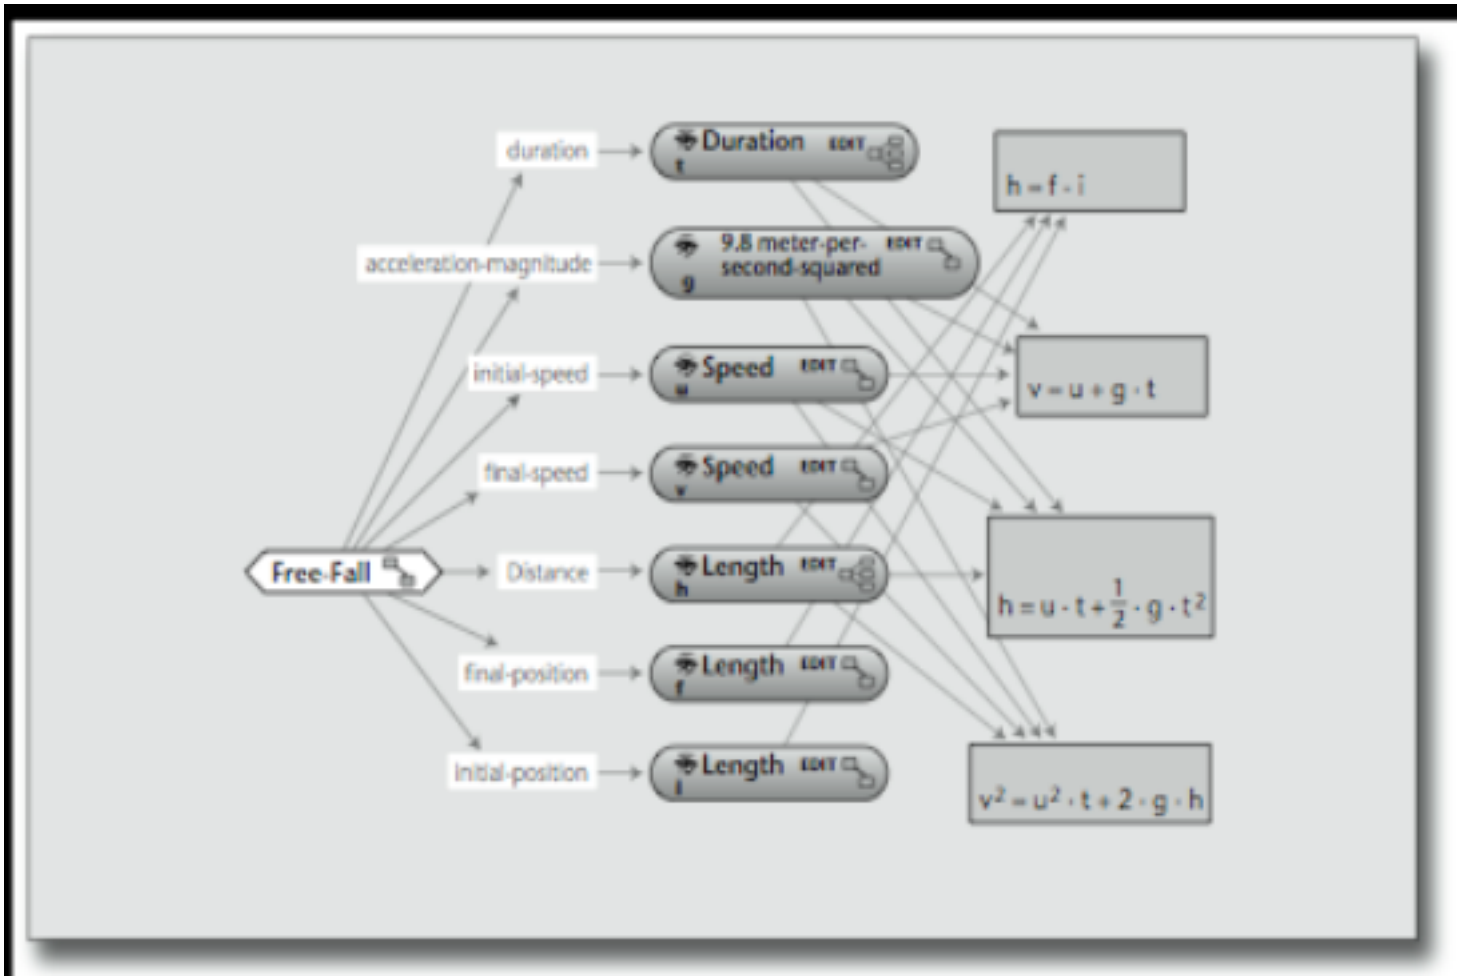

EXAMPLES FOR SIMPLIFIED ENGLISH

Example 1 (Physics)

Original Question:
A car accelerates from 12 m/s to 25 m/s in 6.0 s. How far did it travel in this time?
Reformulation in CPL:
A car is driving.
The initial speed of the car is 12 m/s.
The final speed of the car is 25 m/s.
The duration of the drive is 6.0 s.
What is the distance of the drive?

Example 2 (Chemistry

Original Question:
What two molecules must always be present in the products of a combustion reaction of a hydrocarbon compound?
Reformulation in CPL:
What are the products of a hydrocarbon combustion reaction?

Example 3 (Biology

Original Question:
Crossing over occurs during which of the following phases in meiosis? a. prophase I; b. ...[etc]... ?

Reformulation in CPL:
Does crossing over occur during prophase I?

EXAMPLE 1 - EXPLANATION (PHYSICS)

Answer

s = 111 m

Explanation

motion-with-constant-acceleration: A move of an object such that the acceleration of the object is constant throughout the move.

Given:

- v_1 = 25 m/s [the speed of the final-velocity]
- v_2 = 12 m/s [the speed of the initial-velocity]
- v_3 = 12 m/s [the speed of the initial-velocity]
- v_4 = 25 m/s [the speed of the final-velocity]
- t = 6.0 s [the duration of the motion-with-constant-acceleration]
- u = 12 m/s [the speed of the initial-velocity]
- v_4 = u + a * t [the speed of the final-velocity]

Solving for s ... s = u * t + ((1/2) * a) * t^2
a = (v_4 – u) / t [solving (v_4 = u + a * t) for a]
... a = 2.2 m/s^2 ... s = 111 m
Therefore, the distance of the motion-with-constant-acceleration (s) = 111 m

EXAMPLE 2 - EXPLANATION (CHEMISTRY)

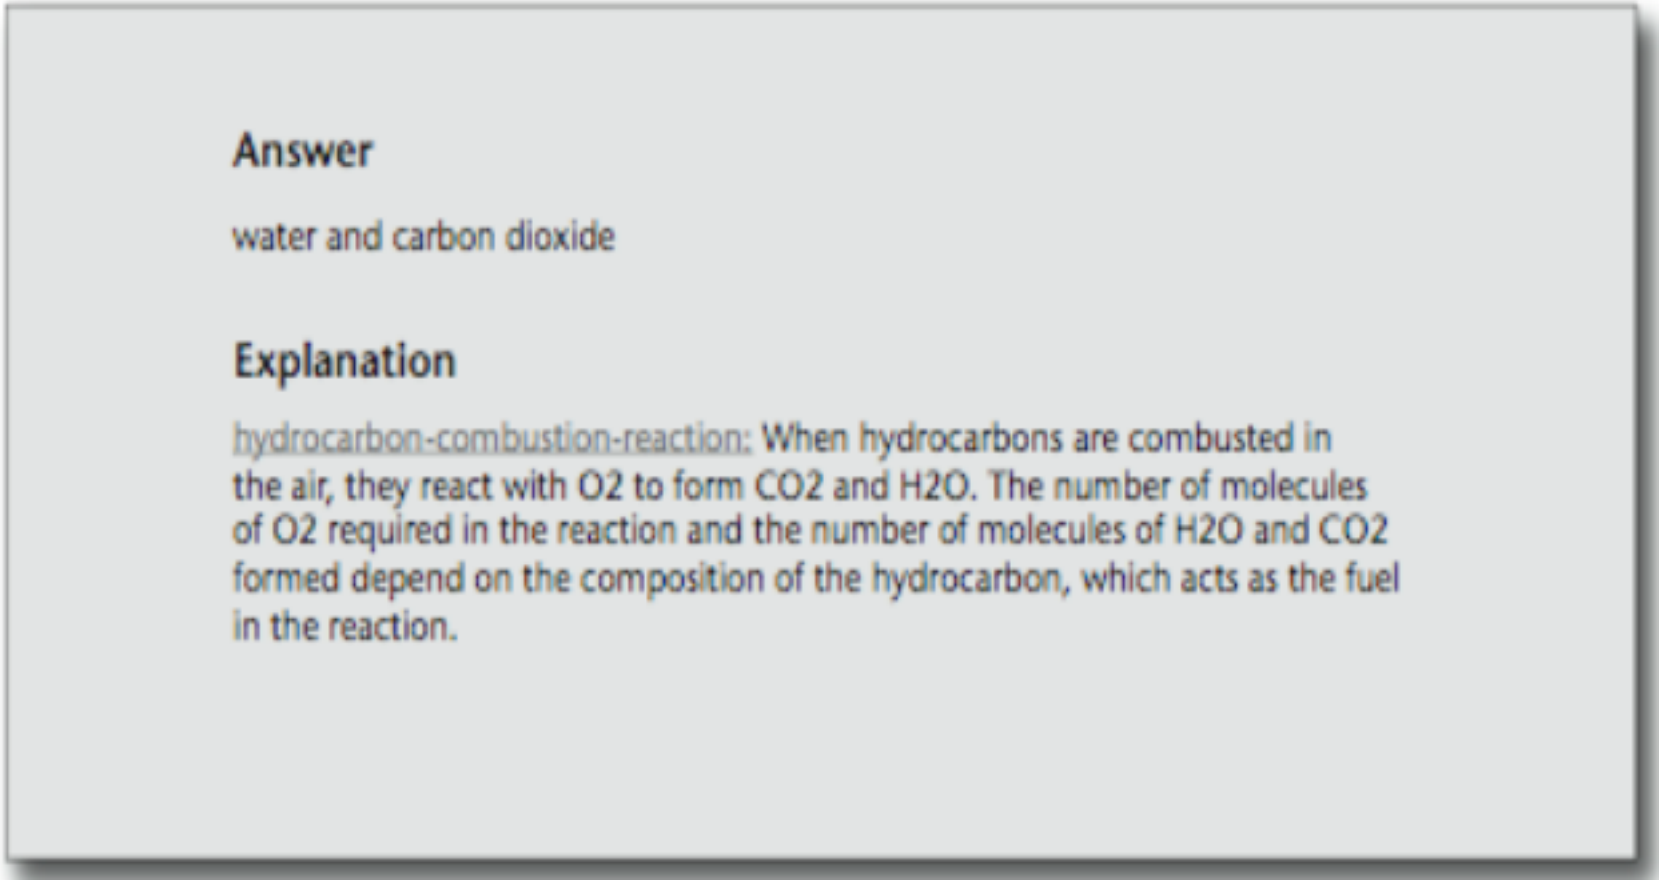

KNOWLEDGE FORMULATION

Textbook example (biology - necessary properties of a cell)

All cells have several basic features in common:
they are all bounded by a membrane, called a plasma membrane.
Within the membrane is a semifluid substance, cytosol,
in which organelles are found.
All cells contain chromosomes,
carrying genes in the form of DNA.
And all cells have ribosomes, tiny organelles
that make proteins according to instructions from the genes.

Extra text - sufficient properties of a cell

In contrast, the eukaryotic cell (Greek eu, true, and karyon) has a true nucleus,
by a membranous nuclear envelope.

KNOWLEDGE FORMULATION - GRAPH EDITED

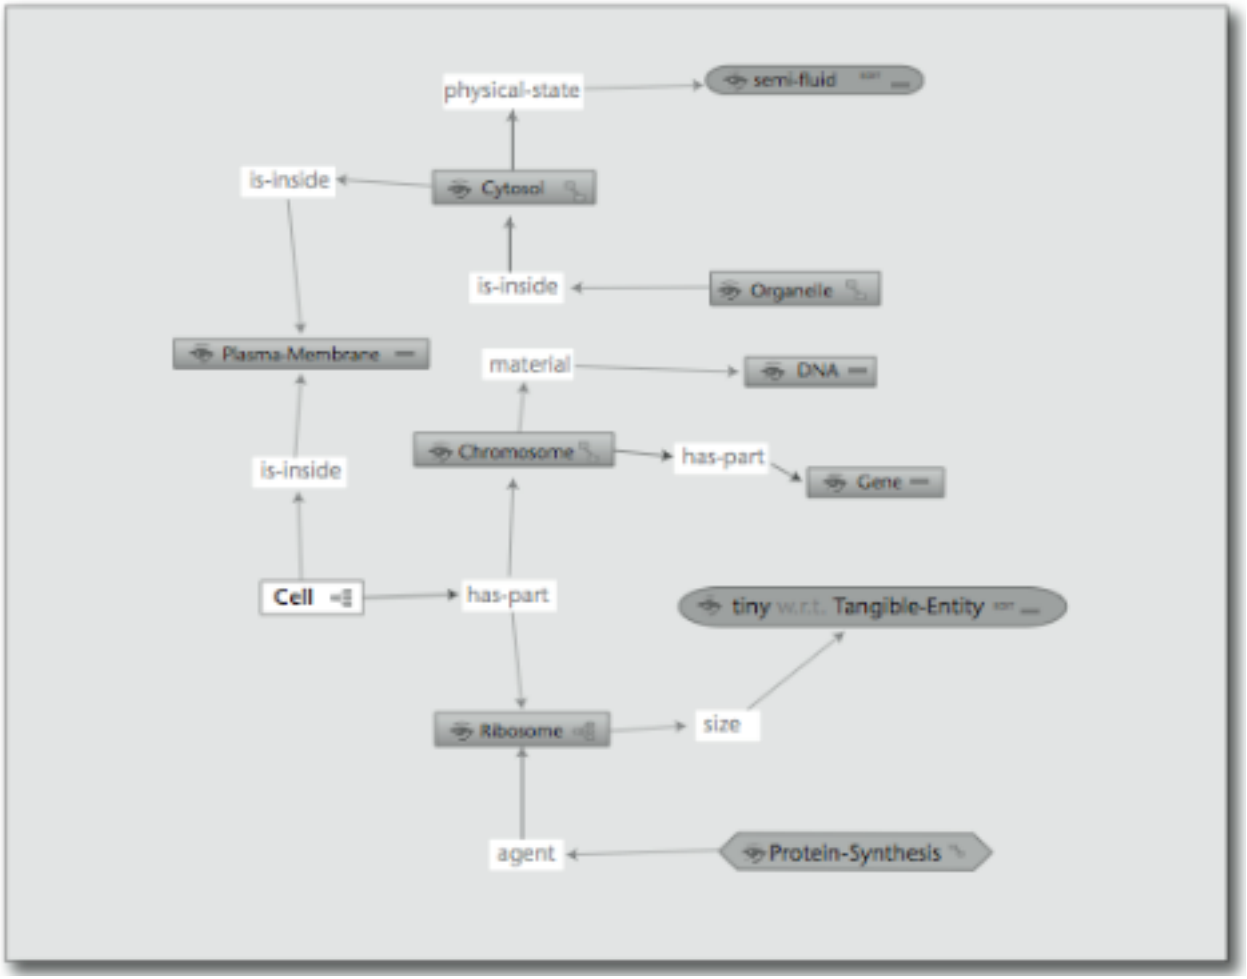

KNOWLEDGE FORMULATION - RULE REPRESENTED

```
(forall ?c
  (=> (instance-of ?c Cell)
    (exists ?r ?ch ?g ?d ?p ?pm ?cy ?o
    (and
      (instance-of ?r Ribosome)
      (instance-of ?ch Chromosome)
      (instance-of ?g Gene)
      (instance-of ?d DNA)
      (instance-of ?p Protein-Synthesis)
      (instance-of ?pm Plasma-Membrane)
      (instance-of ?cy Cytosol)
      (instance-of ?o Organelle)
      (has-part ?c ?ch) (has-part ?ch ?g)
      (has-part ?c ?r) (material?ch ?d)
      (size ?r (scalar-value tinyTangible-Entity))
      (agent ?p ?r)
      (is-inside ?c ?pm) (is-inside ?cy ?pm)
      (is-inside ?o ?cy) (physical-state ?cy semi-fluid))
```

KNOWLEDGE ENGINE

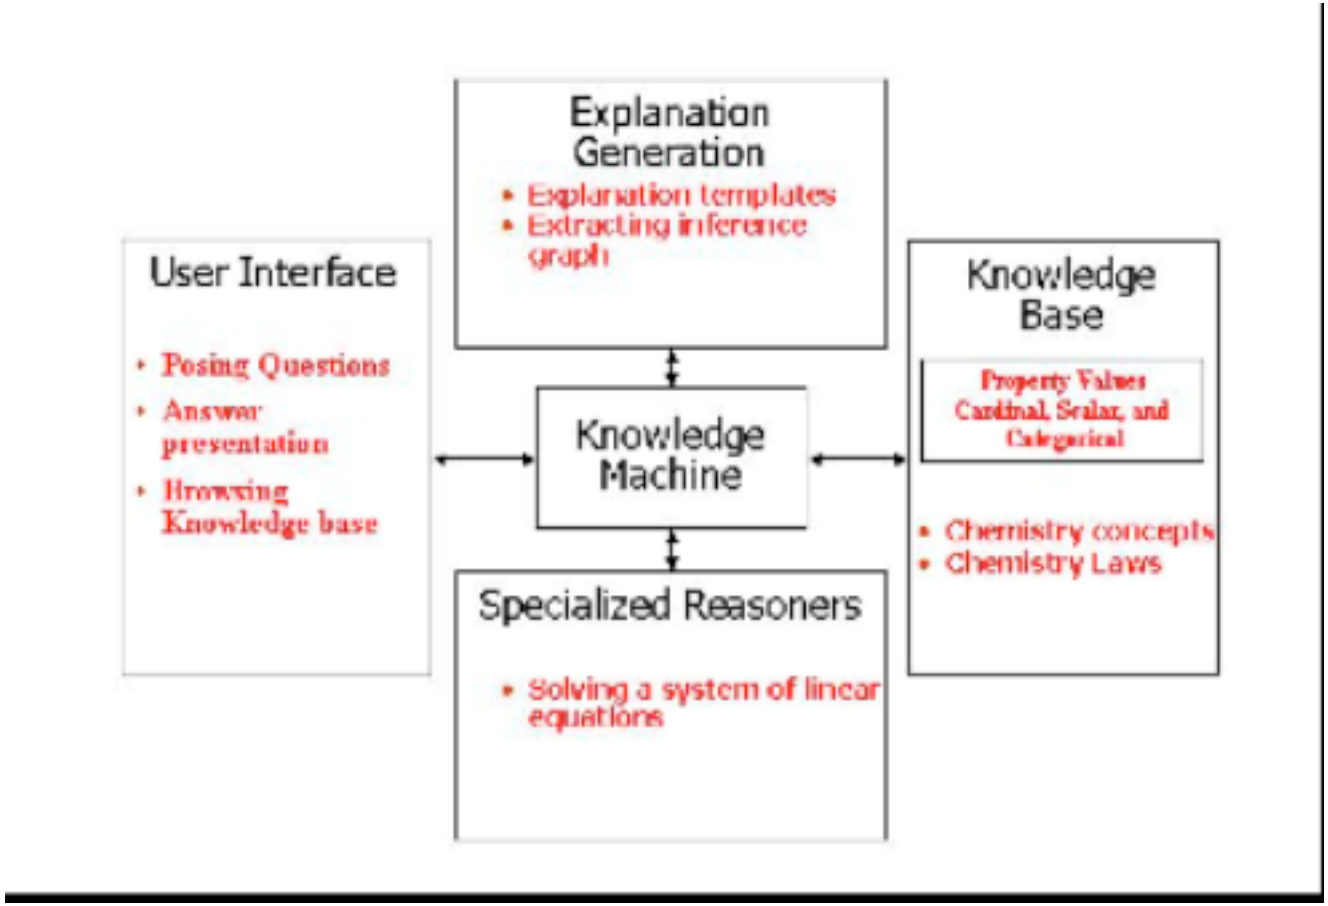

REASONING TECHNIQUES (1)

Defaults/Exceptions/Defeasible

(incl. nonmonotonicreasoning, theory revision, argumentation, truth maintenance)

- **A kinematics problem situation has standard earth gravity, and no air resistance. [physics AP]**
- **A given organism has the anatomy/behavior that is typical/normal for its species, e.g., a bat has 2 wings and flies. [bio AP]**
- **Price info for an airplane ticket on Alaska Air's website is accurate and up to date. [e-shopping]**

Practical reasoning almost always involves a potential for exceptions – Hypotheticals

- **If Apollo astronaut Joe golfed a ball on the moon, then standardearth gravity would not apply. [negative hypothetical] [*conflict* between defaults, resolved by *priority* among them]**
- **If I had swerved my car 5 seconds later than I did, I would have hit the debris in the left lane with mytire. [*counterfactual*]**

REASONING TECHNIQUES (2)

Actions and Causality

- **If a doorkeyis incompletely inserted into the keyhole, turning the key will fail. [*precondition*]**
- **During the mitotic stage of prometaphase, a cell's nuclear envelope fragments**

  [biology AP]
- **After a customer submits an order on the website, Amazon will email a confirmation and ship the item. [Event-Condition-Action (*ECA*) rule] [policy]**

Processes

- **Mitosis has five stages; its successful completion results in two cells.**

  [compose] [partial description]
- **If Amazon learns that it will take an unexpectedly long time to stock an ordered item, then it emails the customer and offers to cancel the order without penalty.**

  [exception handling]
- **A Stillcosensor-based negative feedback thermal regulator is adequate**

  to ensure the overnight vat fermentation of the apple mash will proceed

  within desired bounds of the alcohol concentration parameter.

  [science-based business process]

SYSTEM FEATURES

KR System with multiple software components

- **Logical Language, incl. Syntax and Semantics**
- **Reasoning, incl. Backward and Forward Inferencing**
- **Web Knowledge Interchange, incl. Translators**
- **KA/UI Support, incl. for Editing and Explanation**

Evolutionary Approach

- **Start from known coreKR**
- **Add more features in principledfashion**
- **Requirements, use cases, benchmarking, KB building;**
  system design (incl. theory, usability), implementation, testing (incl. task)

REASONING METHODS

- **It integrates several major LP extensions never previously combined:**
- **Higher-orderandFramesand Skolemization, cf. F-Logic**
- **+Defaults, cf. Courteous LP (and DefeasibleLogic)**
- **Newly generalized and modified approach**
- **+Weakened Full Classical Logic, cf. Hypermonotonicmapping**
  - **Greatly generalizes the approach of Description LP and OWL 2 RL**
  - **Unrestricted clauses, plus skolemization**
  - **Leverages Courteous feature**
  - **Give up disjunction / reasoning by cases, so is weakened, but**
  - **Behaves robustly in face of knowledge quality errors and conflictful merging**

CAUSAL PROCESS REASONING – EXAMPLE

- **/* Toxic discharge into a river causes fish die-off. */**
- **/* Initial facts */**
  - **occupies(trout,Squamish).**
  - **fishCount(s0,Squamish,trout,400).**
- **/* Action/event description that specifies causal effect on nextstate */**
  - **{tdf1} fishCount(?s+1,?r,?f,0) :-occurs(?s,toxicDischarge,?r) and occupies(?f,?r).**
- **/* Persistence (“frame”) axiom */**
  - **{pe1} fishCount(?s+1,?r,?f,?p) :-fishCount(?s,?r,?f,?p).**
- **/* Action effect axiom has higher priority than persistence axiom */**
  - **{pr1} overrides(tdf1,pe1).**
- **/* An action instance occurs */**
  - **{UhOh} occurs(s0+1,toxicDischarge,Squamish).**
  - ***As desired:* |= fishCount(s0+1,Squamish,trout,400) and fishCount(s0+2,Squamish,trout,0).**

SUMMARY – A DIGITAL ARISTOTLE

The task

- **Question-answering in technical domains (e.g., chemistry)**
- **Multiple domains and multiple reasoning techniques**
- **Explanations presented in a concise form, user-adaquate form**

Insights

- **Knowledge acquisition is the bottleneck – Investigations in automation**
- **Query formulation in "simplified English" – may be tricky**
- **Explanations patterns: (rule, nested dependents, conclusion) essential**
- **Assessment scores close to average human performance**

Failures

- **Asking for explicit descriptions of strategies**
- **Asking about reasons for an implausible result of a standard procedure**
- **Asking about similarities and differences (e.g., chemistry language)**

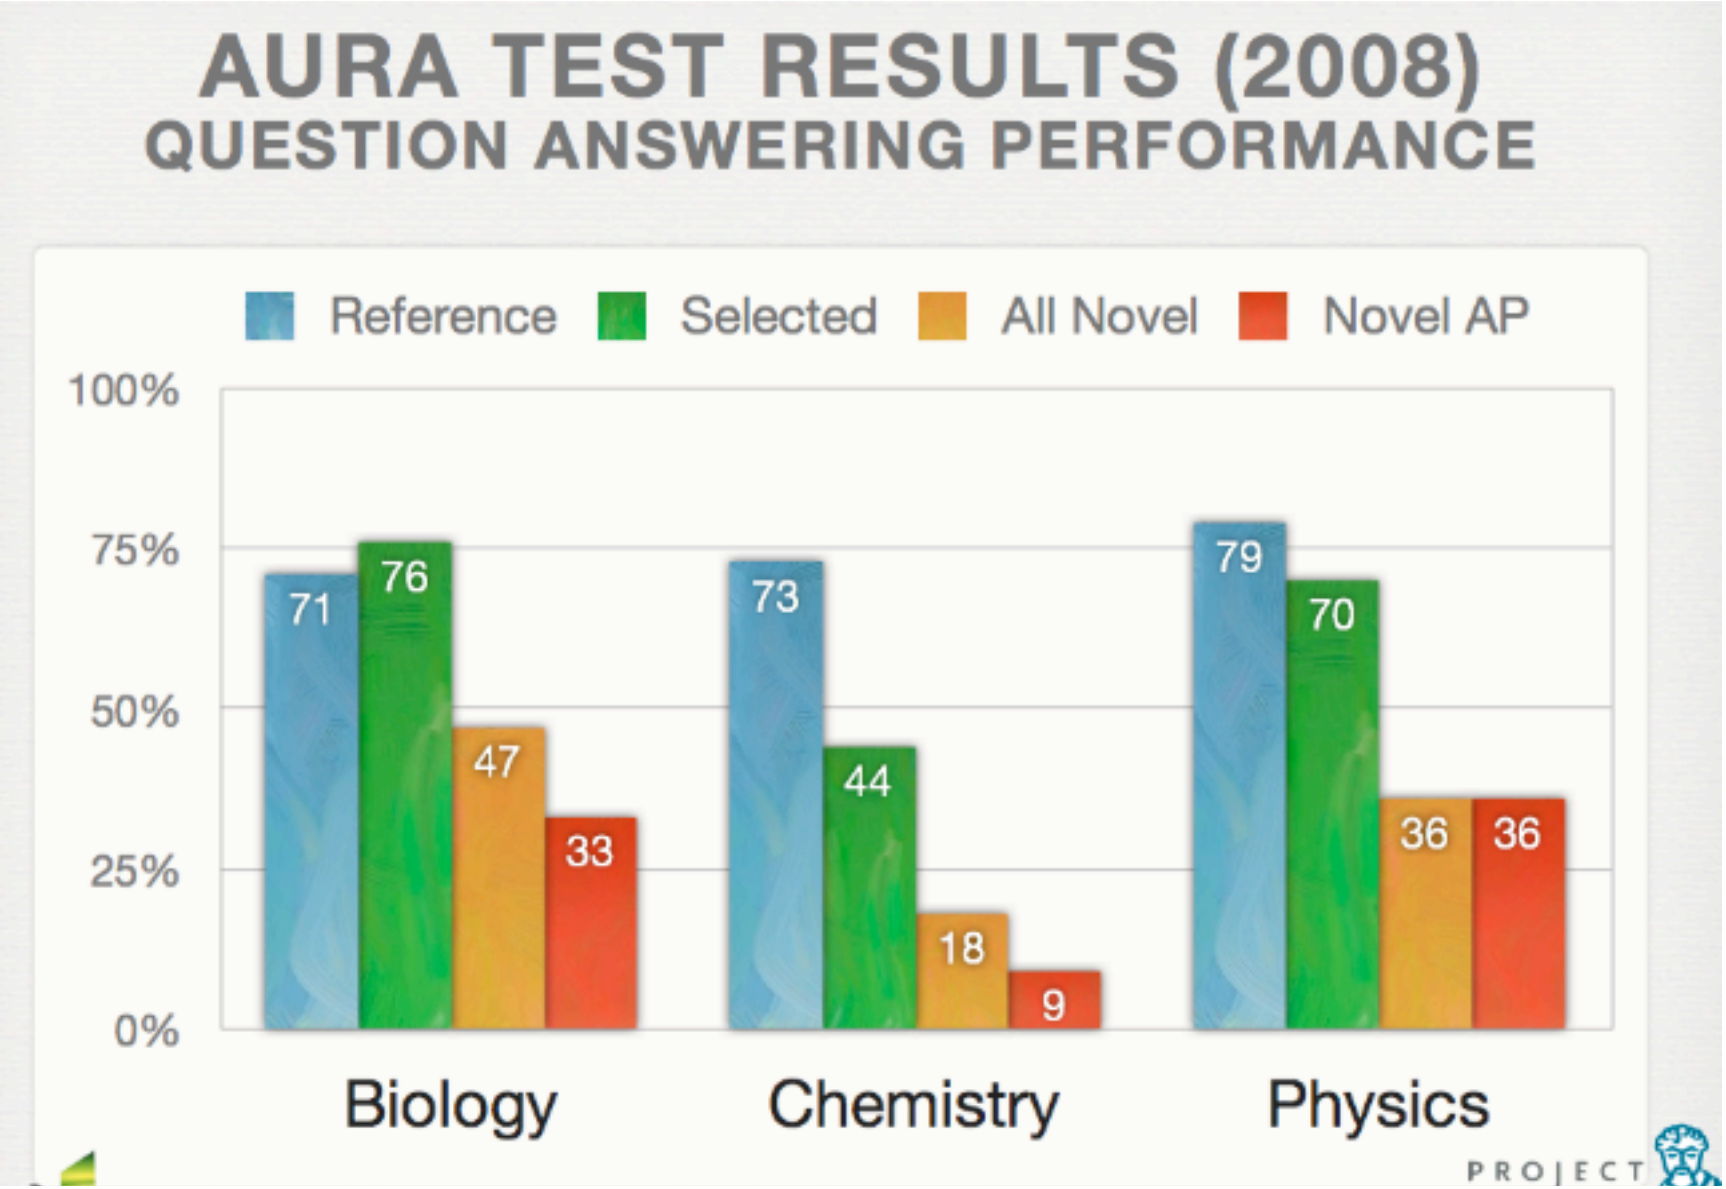
AURA TEST RESULTS (2008)
QUESTION ANSWERING PERFORMANCE
Reference
Selected
All Novel
Novel AP
100%
75%
50%
25%
0%
71
76
47
33
73
44
18
9
79
70
36
36
Biology
Chemistry
Physics
PROJECT

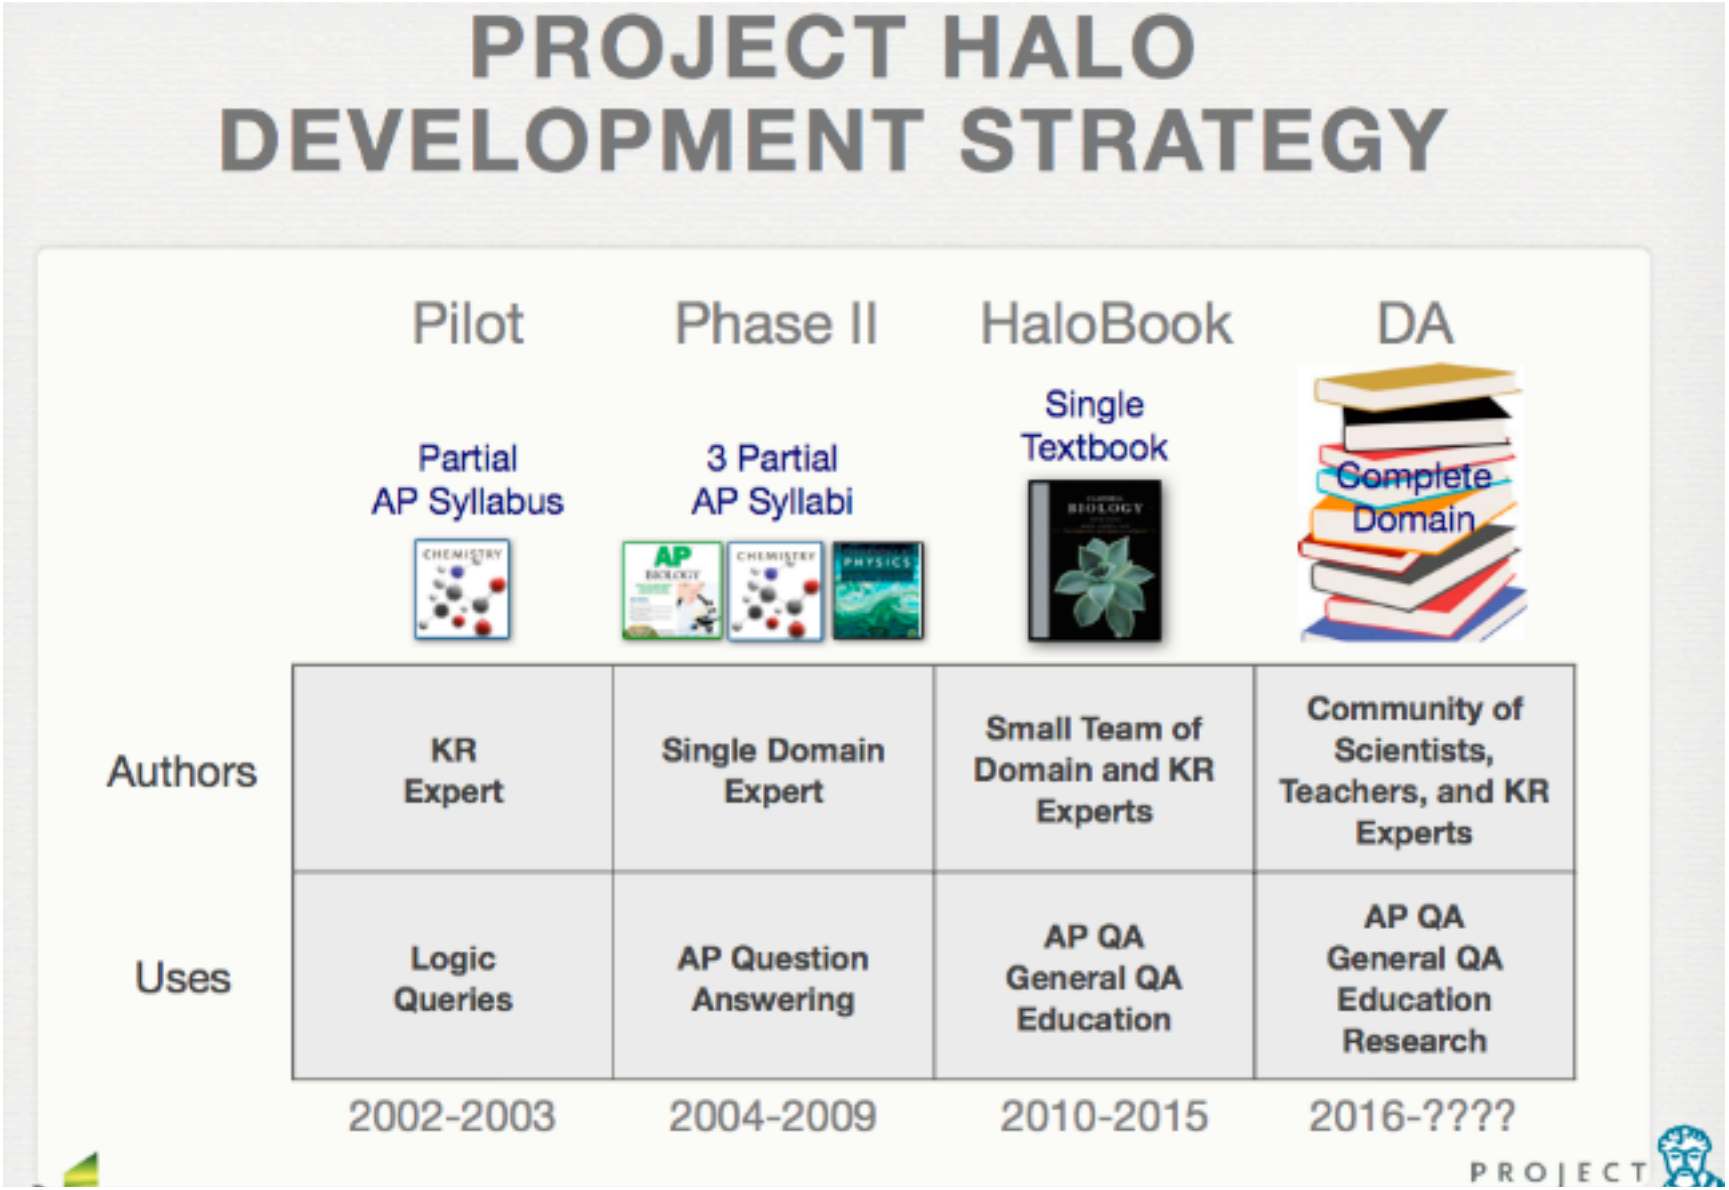
PROJECT HALO
DEVELOPMENT STRATEGY
Pilot
Phase II
HaloBook
DA
Partial
AP Syllabus
3 Partial
AP Syllabi
Single
Textbook
Complete
Domain
Authors
KR
Expert
Single Domain
Expert
Small Team of
Domain and KR
Experts
Community of
Scientists,
Teachers, and KR
Experts
Uses
Logic
Queries
AP Question
Answering
AP QA
General QA
Education
AP QA
General QA
Education
Research
2002-2003
2004-2009
2010-2015
2016-????
PROJECT

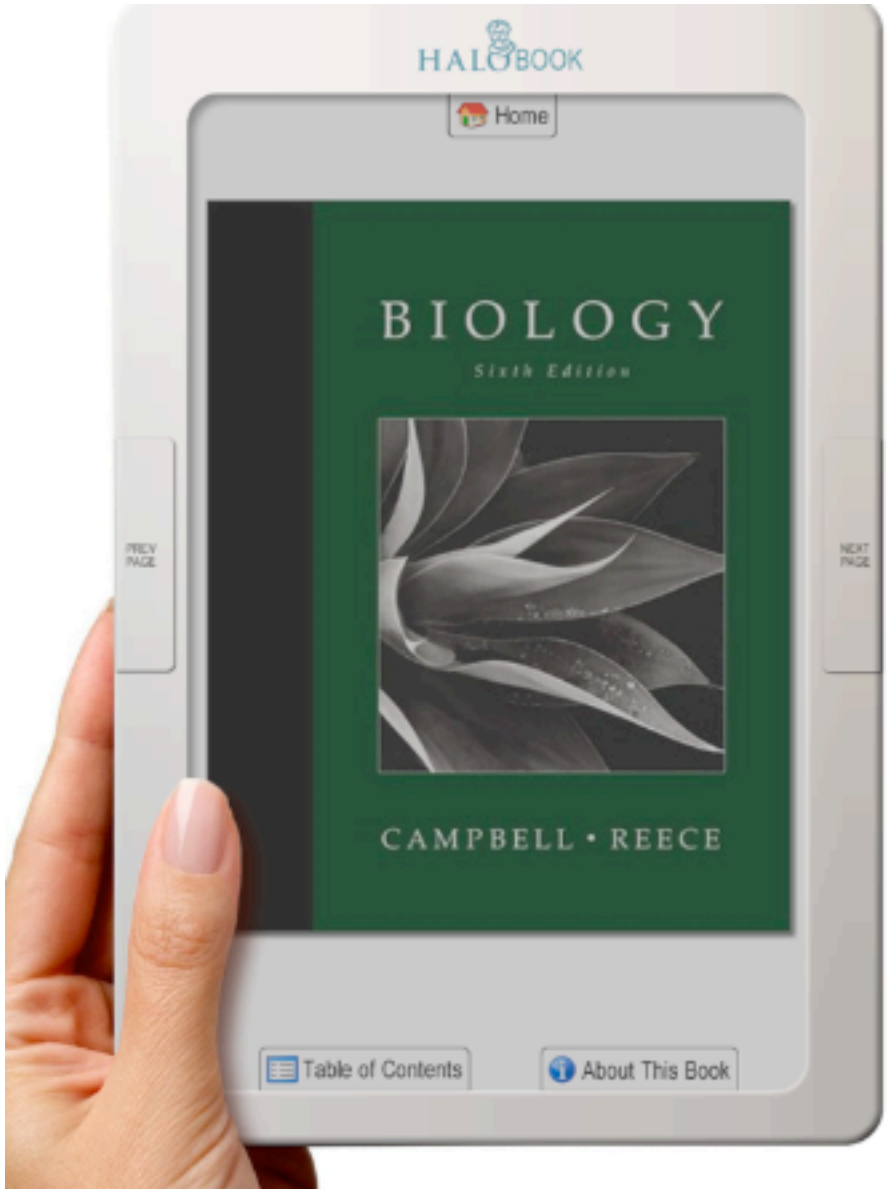

HaloBook 2015

- HaloBook is a concept for a new kind of electronic textbook
  - that contains an underlying knowledge-base of the book's contents
  - that answers the reader's questions and provides tailored instruction
- This briefing presents of sketch of how a HaloBook might look in 2015
  - assuming continued Halo development and modest success in solving research challenges
  - An initial HaloBook prototype will be developed in 2010

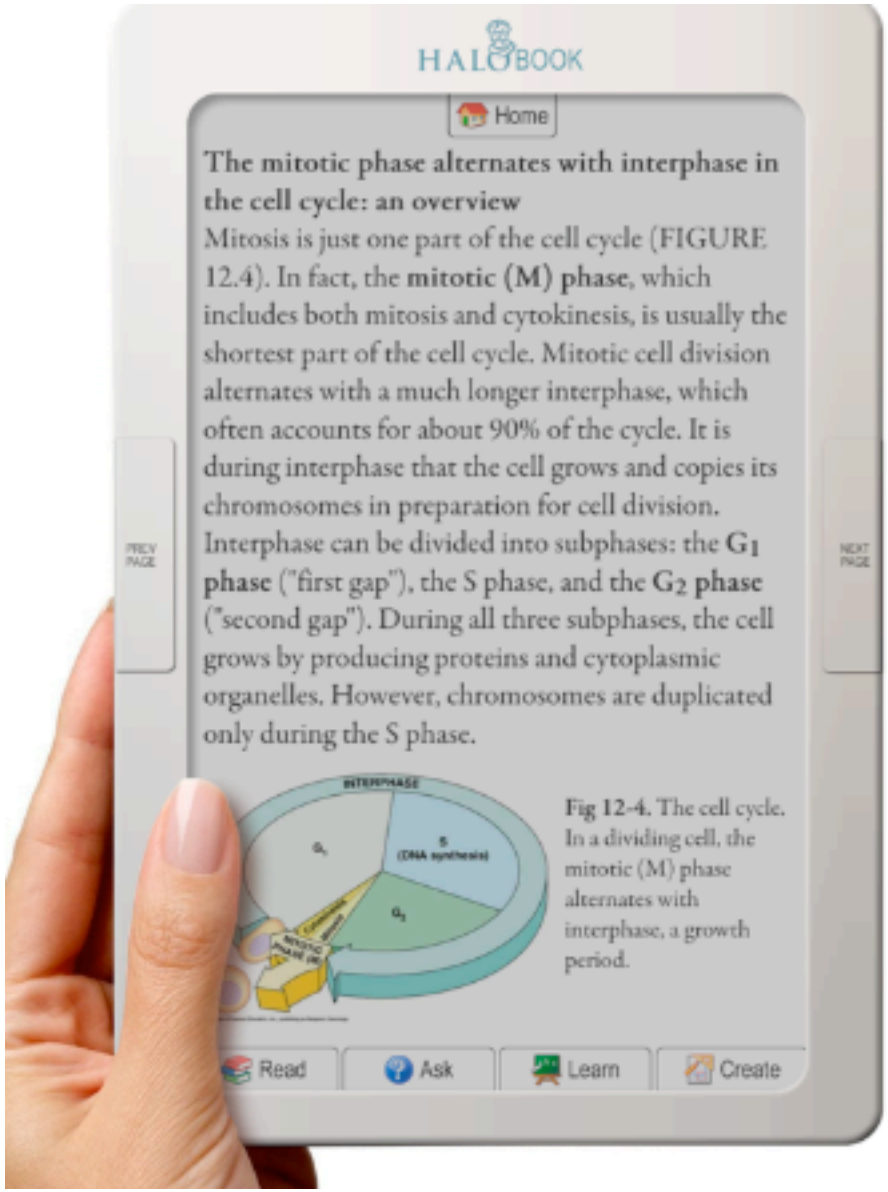

HaloBook Functions

A HaloBook user can:

- **READ** the contents dynamically, interactively
- **ASK** questions and get explanations on any subject in the book
- **LEARN** to master the subject through individualized tutoring
- **CREATE** and explore their own conceptualizations of the material

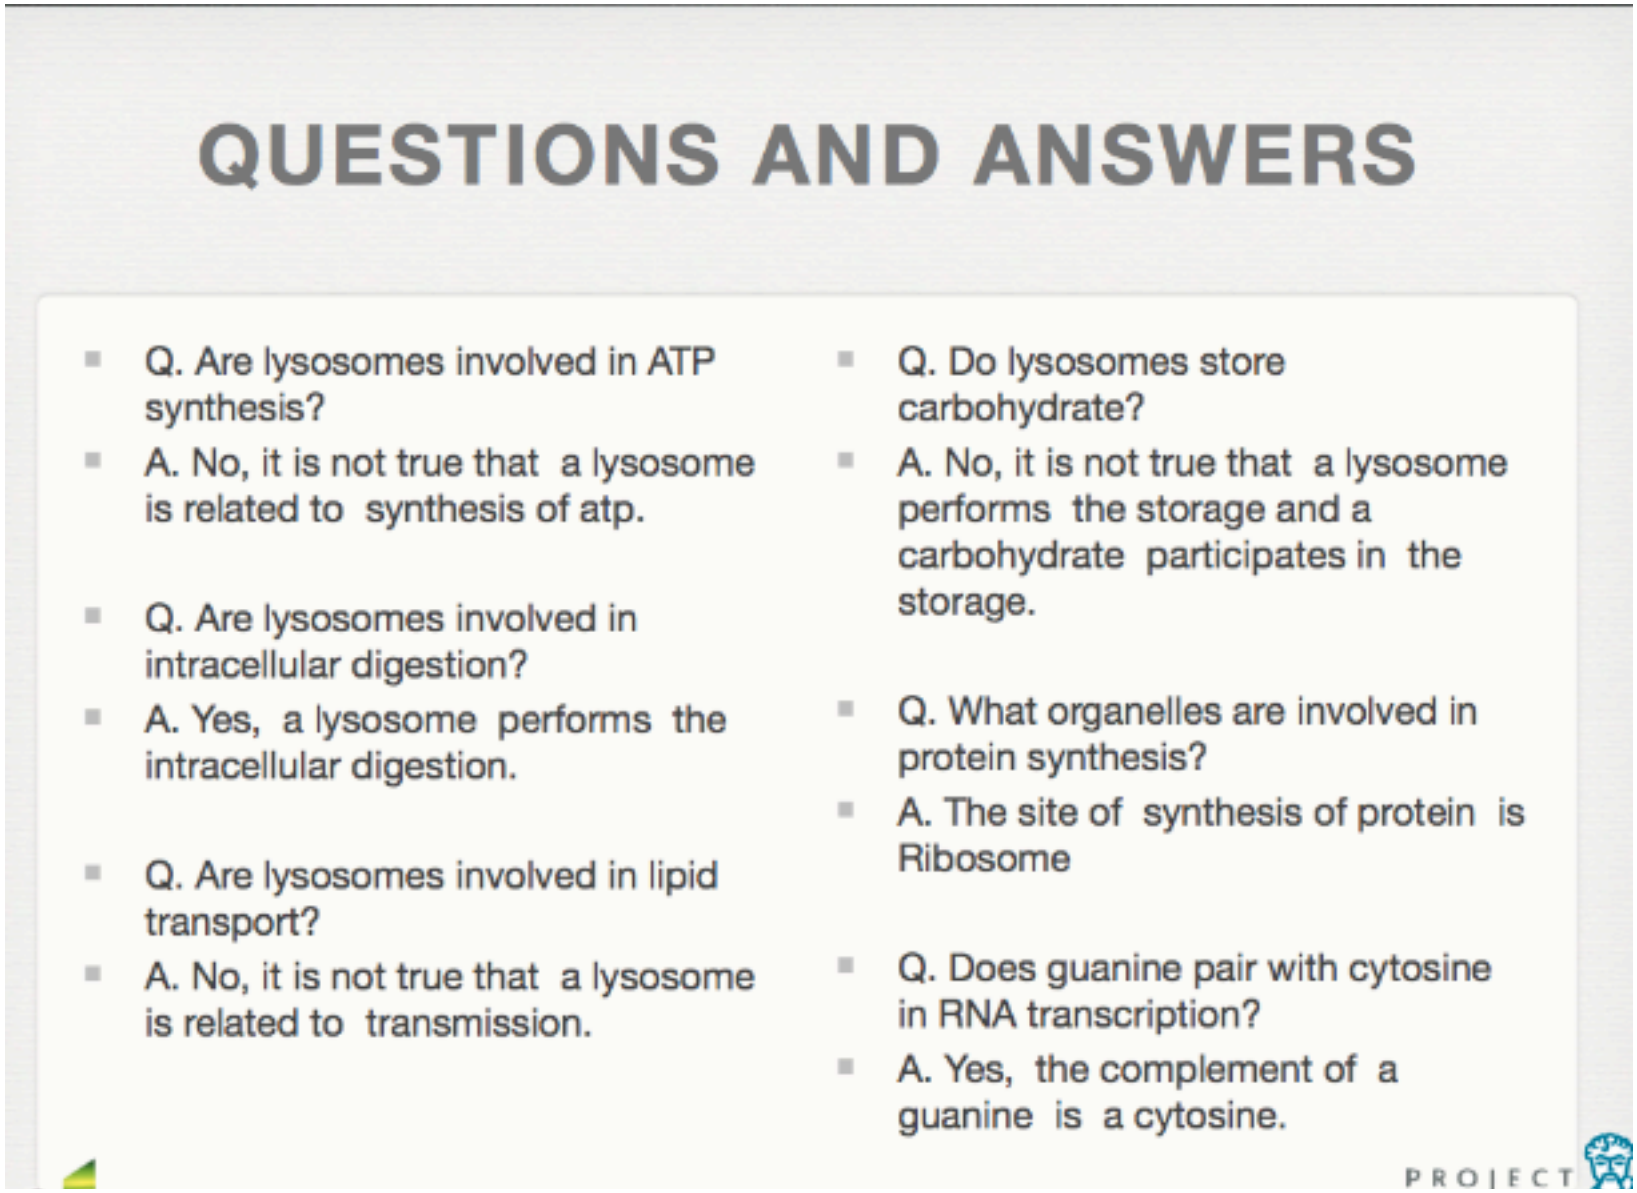
QUESTIONS AND ANSWERS
Q. Are lysosomes involved in ATP synthesis?
A. No, it is not true that a lysosome is related to synthesis of atp.
Q. Are lysosomes involved in intracellular digestion?
A. Yes, a lysosome performs the intracellular digestion.
Q. Are lysosomes involved in lipid transport?
A. No, it is not true that a lysosome is related to transmission.
Q. Do lysosomes store carbohydrate?
A. No, it is not true that a lysosome performs the storage and a carbohydrate participates in the storage.
Q. What organelles are involved in protein synthesis?
A. The site of synthesis of protein is Ribosome
Q. Does guanine pair with cytosine in RNA transcription?
A. Yes, the complement of a guanine is a cytosine.
PROJECT

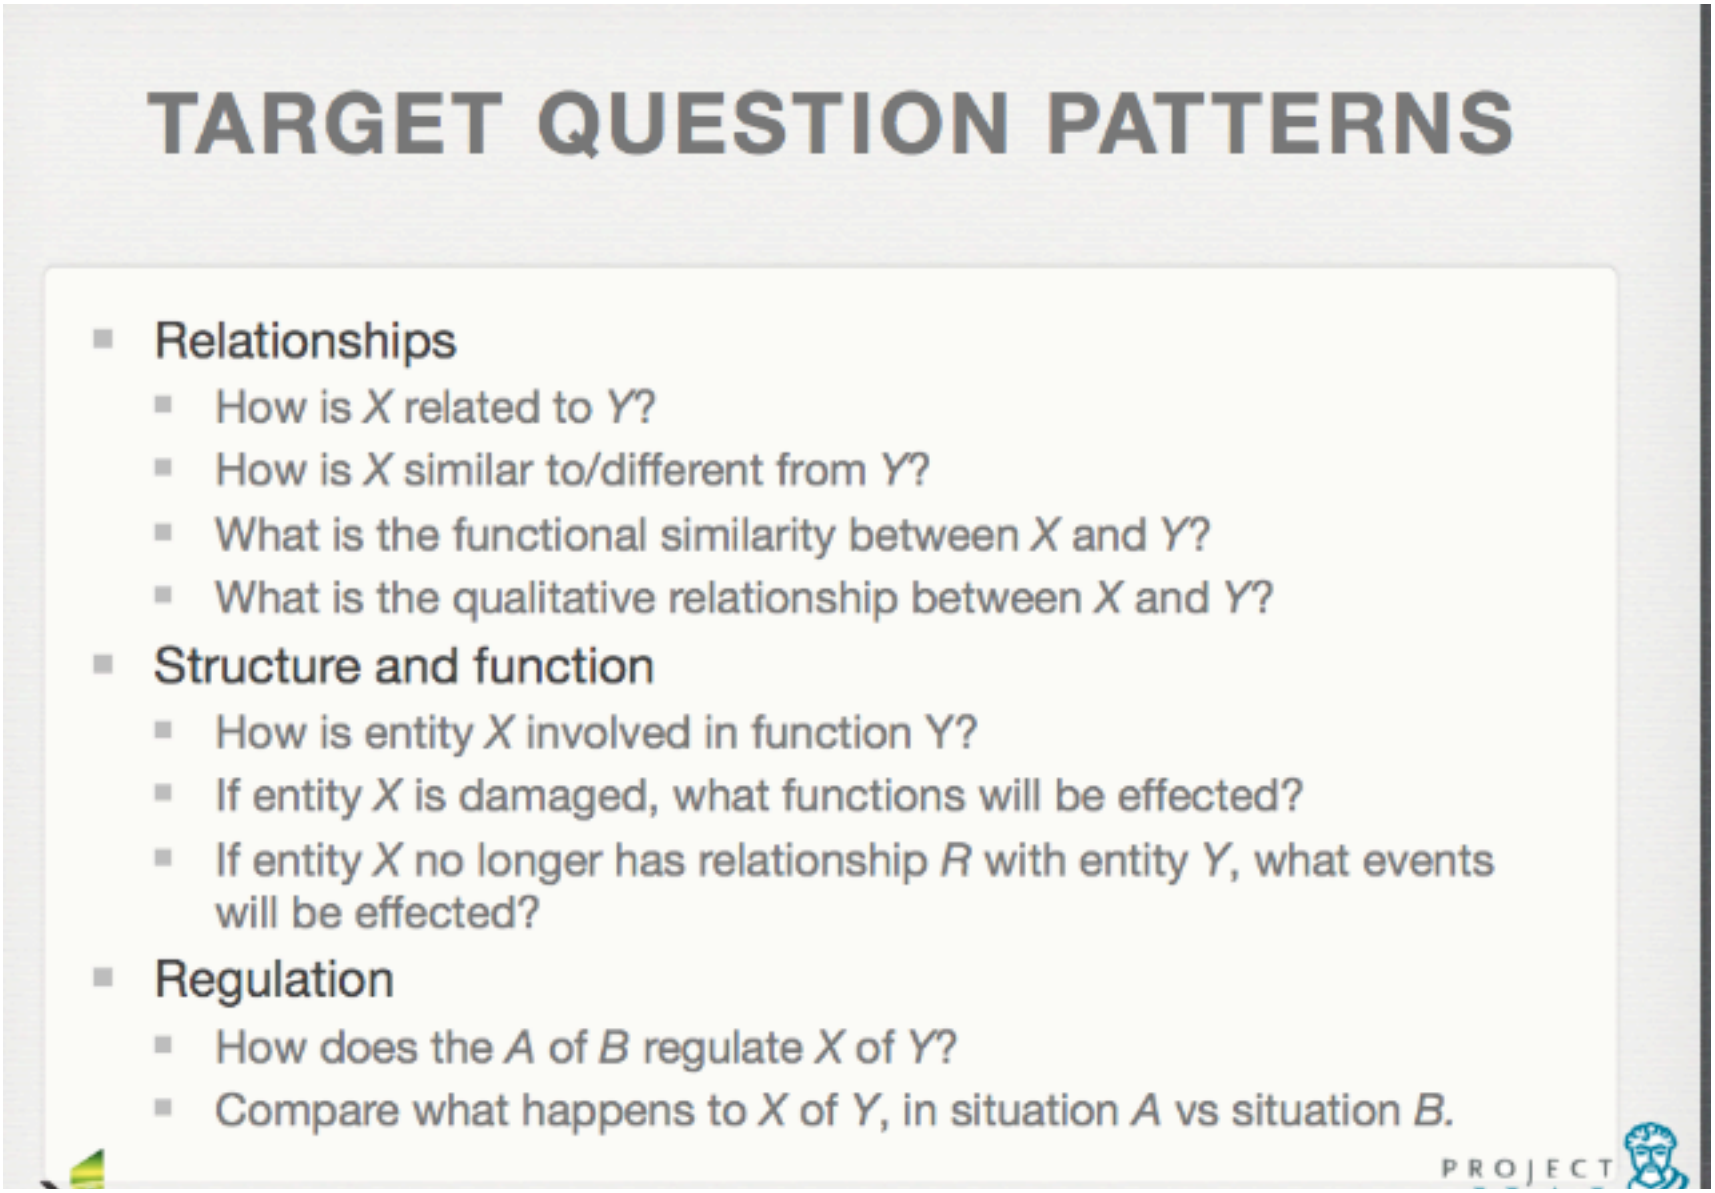
TARGET QUESTION PATTERNS
Relationships
How is X related to Y?
How is X similar to/different from Y?
What is the functional similarity between X and Y?
What is the qualitative relationship between X and Y?
Structure and function
How is entity X involved in function Y?
If entity X is damaged, what functions will be effected?
If entity X no longer has relationship R with entity Y, what events will be effected?
Regulation
How does the A of B regulate X of Y?
Compare what happens to X of Y, in situation A vs situation B.
PROJECT

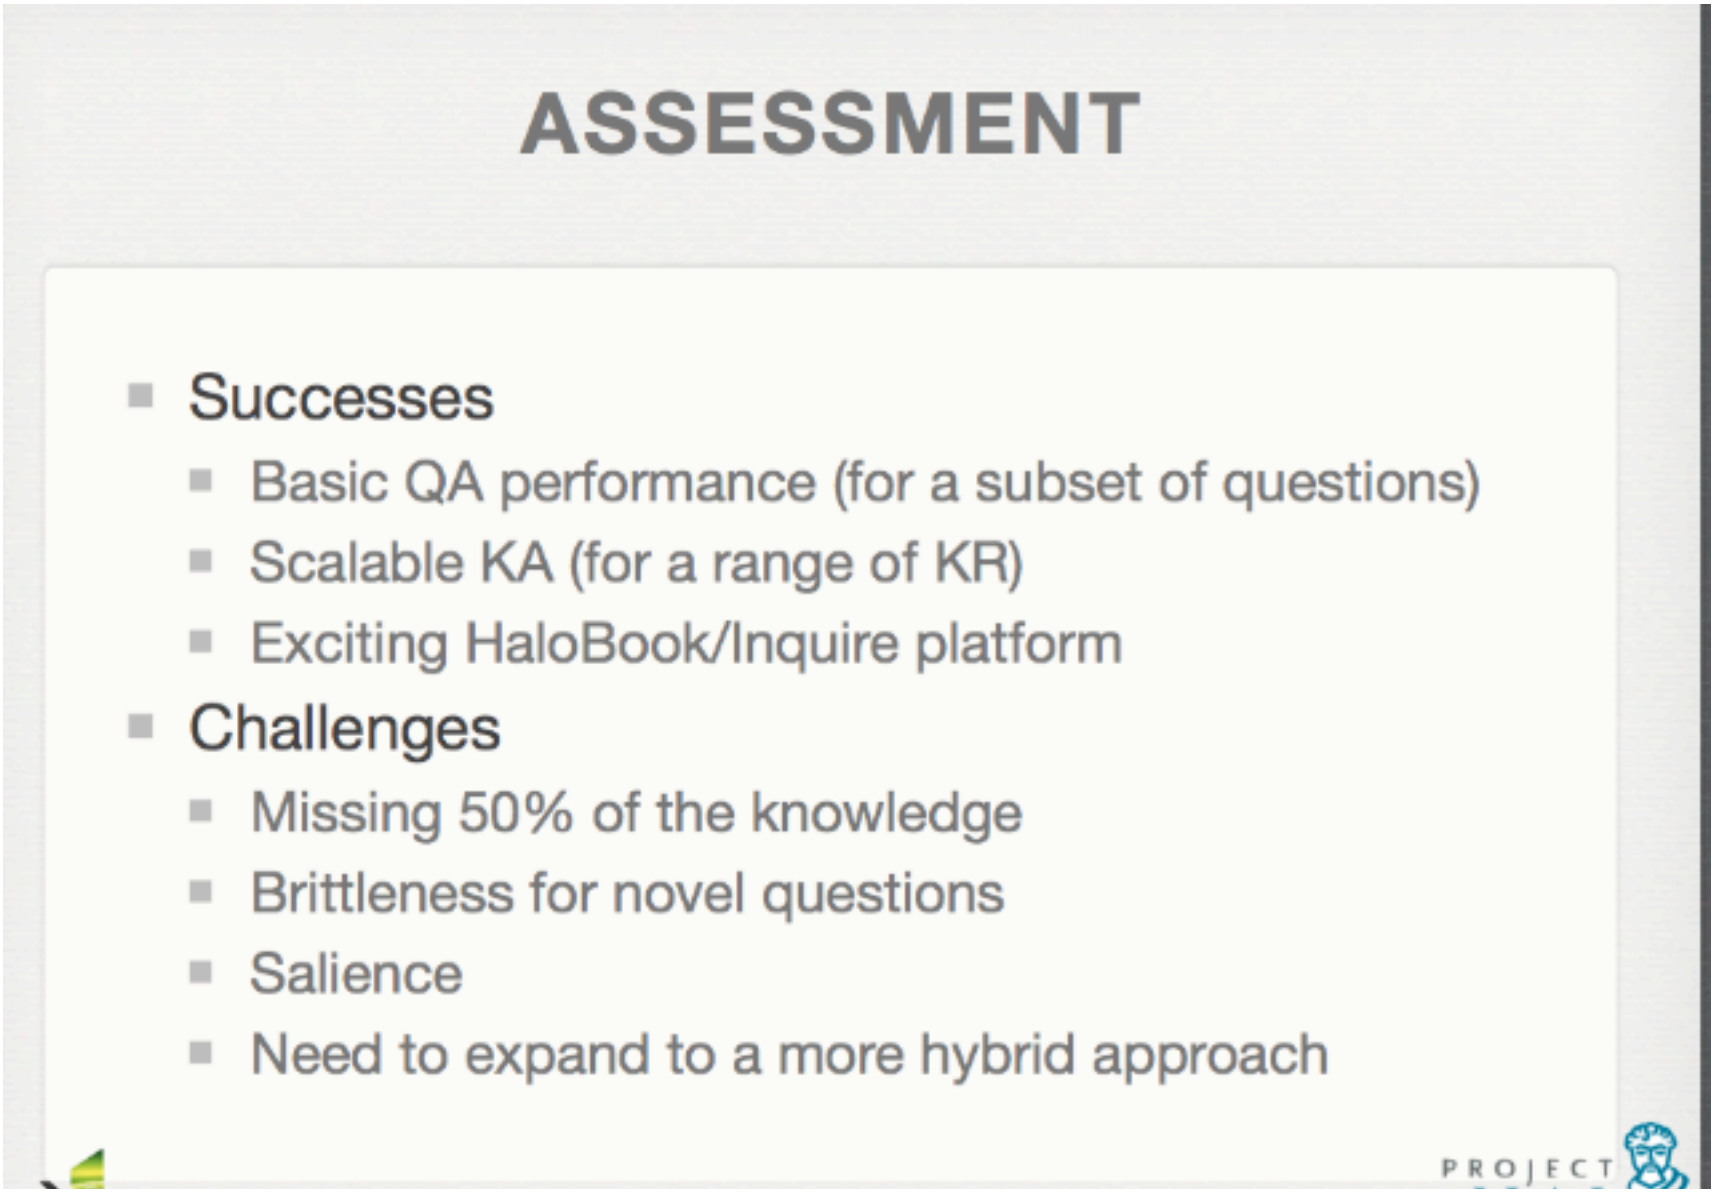
ASSESSMENT
Successes
Basic QA performance (for a subset of questions)
Scalable KA (for a range of KR)
Exciting HaloBook/Inquire platform
Challenges
Missing 50% of the knowledge
Brittleness for novel questions
Salience
Need to expand to a more hybrid approach
PROJECT

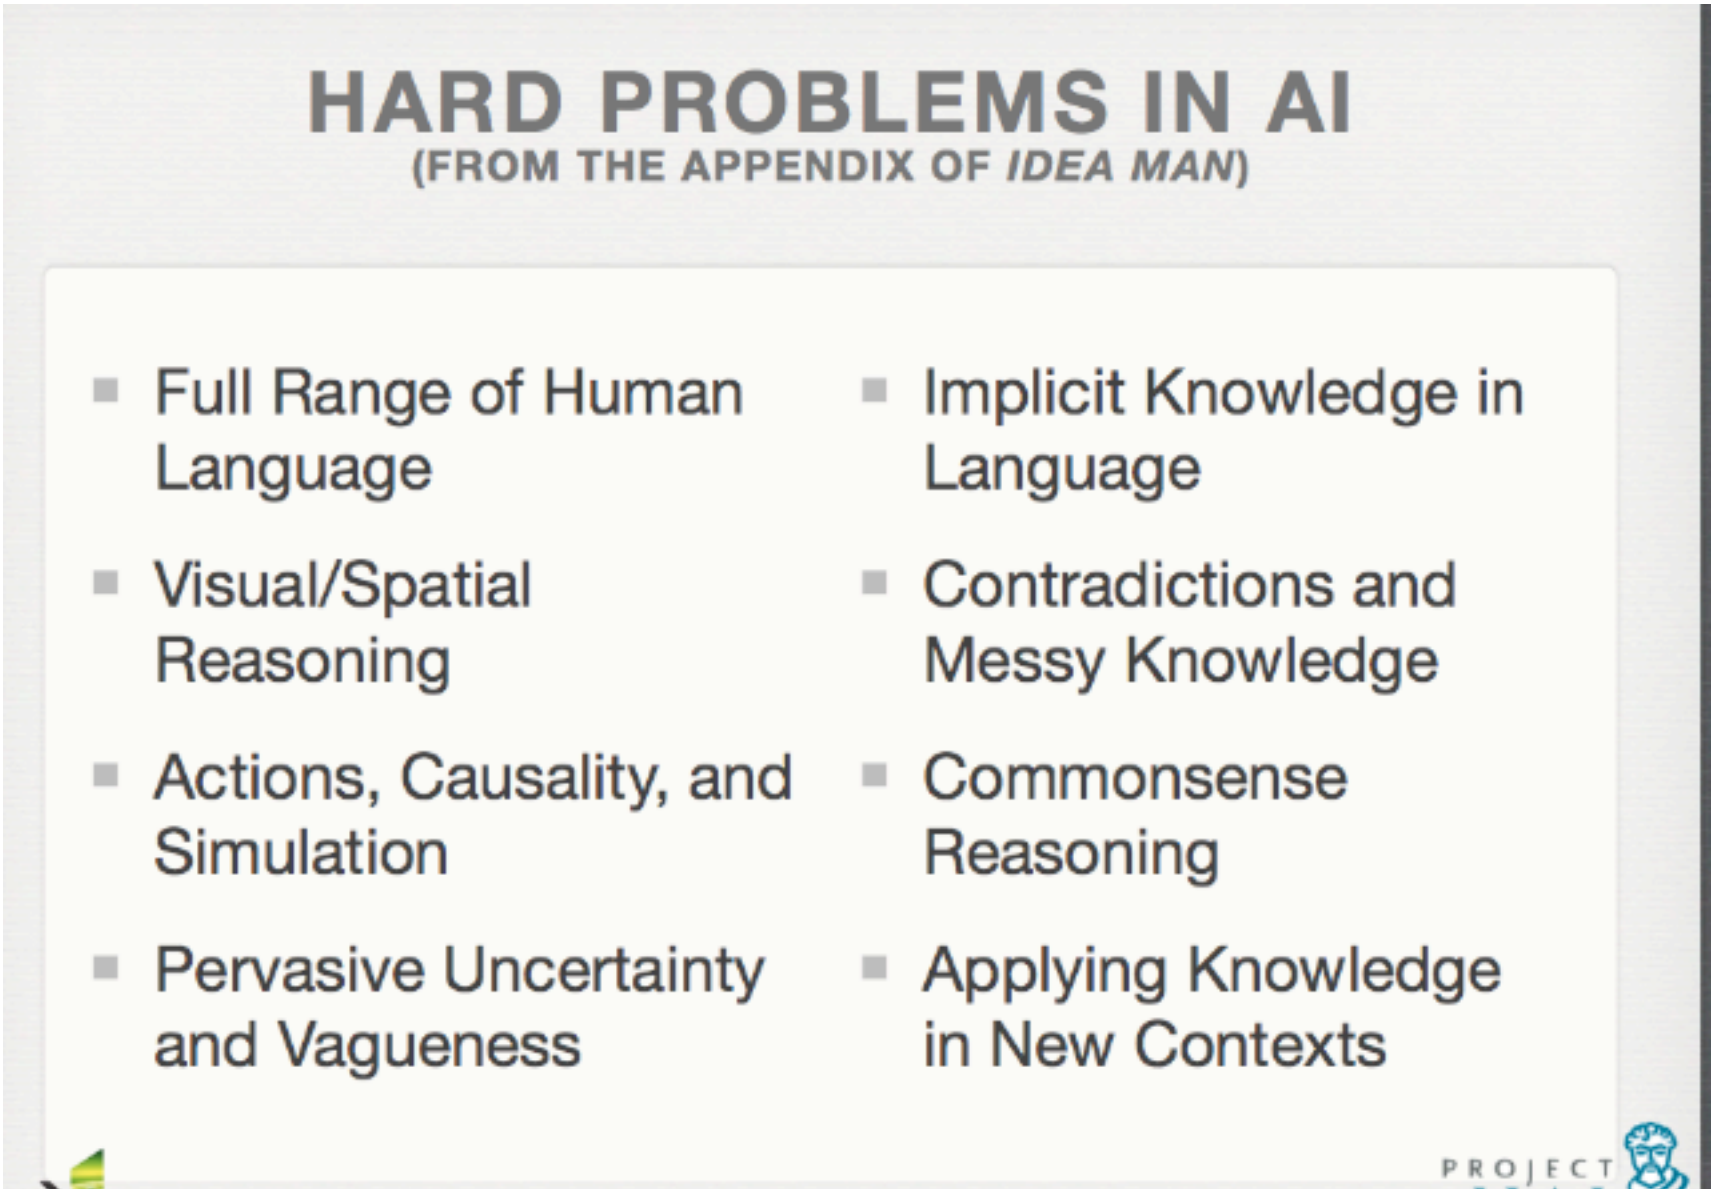
HARD PROBLEMS IN AI
(FROM THE APPENDIX OF IDEA MAN)
Full Range of Human Language
Visual/Spatial Reasoning
Actions, Causality, and Simulation
Pervasive Uncertainty and Vagueness
Implicit Knowledge in Language
Contradictions and Messy Knowledge
Commonsense Reasoning
Applying Knowledge in New Contexts
PROJECT