Advanced Knowledge Representation and Reasoning for A-SpaceX and the A-Desk

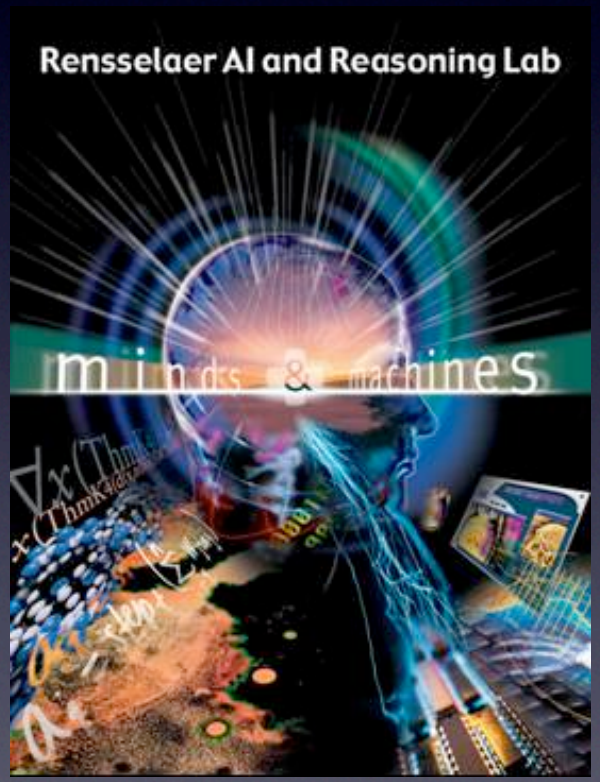

Selmer Bringsjord
Andrew Shilliday, Dan Werner,
Micah Clark, Joshua Taylor
Rensselaer AI & Reasoning (RAIR) Laboratory
also including: Konstantine Arkoudas
Department of Cognitive Science
Department of Computer Science
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 US
@ BWI 9.20.07

Rensselaer | DEPARTMENT OF COGNITIVE SCIENCE

Rensselaer Computer Science

Advanced Knowledge Representation and Reasoning for A-SpaceX and the A-Desk

The Plumbers?

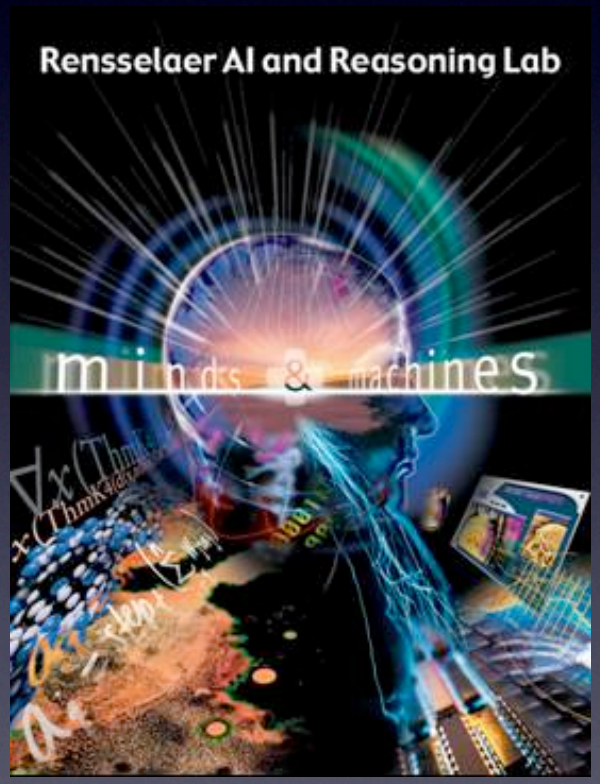

Selmer Bringsjord
Andrew Shilliday, Dan Werner,
Micah Clark, Joshua Taylor
Rensselaer AI & Reasoning (RAIR) Laboratory
also including: Konstantine Arkoudas
Department of Cognitive Science
Department of Computer Science
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 US
@ BWI 9.20.07

Rensselaer DEPARTMENT OF COGNITIVE SCIENCE

Rensselaer Computer Science

It was asked: Why share? Where's the carrot?

It was asked: Why share? Where's the carrot?

The Multi-Mind Effect

Selmer Bringsjord[1], Konstantine Arkoudas[2], Deepa Mukherjee[3]
Andrew Shilliday[4], Joshua Taylor[5], Micah Clark[6], Elizabeth Bringsjord[7]
Department of Cognitive Science[1–6]; Department of Computer Science[1,4,5]
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 USA
State University of New York[7]; SUNY Plaza
Albany NY 12246 USA

Abstract *Courtesy of experiments carried out by such thinkers as Wason, Johnson-Laird, and Kahneman & Tversky, there is overwhelming empirical evidence that the vast majority of logically untrained humans are unable to reason in context-independent, normatively correct fashion. However, the* multi-mind effect*, which is predicted by our earlier success at teaching this kind of reasoning, and also by our general theory of human and machine reasoning, shows that while individual persons (with rare exceptions) are unable to solve problems that demand context-independent reasoning,* groups *of persons can often solve such problems.*

Keywords: multi-mind effect, heterogeneous reasoning, multi-agent reasoning, logic-based computational cognitive modeling

1 Introduction

Experimental study of human reasoning has shown that exceedingly few humans can solve problems demanding normatively correct, context-independent reasoning [1]. Two theories in the field of psychology of reasoning, mental logic (ML) [2] and mental models (MM) [3], both predict such failures, giving the same general explanation: humans generally lack the mental machinery required to solve such problems. Predicted failures include phenomena such as *illusory inferences* [4, 5], in which subjects "see" logically valid inferences that simply aren't there.

In earlier work, we have shown that education of a certain kind in the area of formal logic, specifically education in accordance with our theory of human reasoning, *mental meta-logic* (MML) [6, 7, 8, 9], *contra* the claims of some well-known psychologists (e.g., [10]), can produce humans able to negotiate problems demanding context-independent, normatively correct reasoning [11, 7]. We say that such humans acquire *logical minds*, and hold that our work vindicates Piaget to a considerable degree.[1]

Unfortunately, the training required for an individual to reach this level must take place over an extended period of time, and there is no reason to believe that this individual, without ongoing practice, would retain her hard-won ability.

Nonetheless, the possibility remains that normatively correct reasoning might be possible to achieve in a different manner. But how?

While it's indeed true that the vast majority of individuals are unable to solve problems that require context-independent reasoning (unless suitably trained), *groups* of individuals acting together after being stimulated under the right circumstances *can* often solve such problems, even in the absence of extended training. This result is what we call the *multi-mind effect* (MME).

2 Related Research

Of course, it has long been known that groups can out-perform individuals.[2] However, one must distinguish between groups of individuals that can reason in normatively correct, context-independent fashion to solve so-called "unsolvable" problems (as in

[1] Piaget, as is well known, held that in the course of normal development humans would acquire a capacity to think in accordance with first-order logic [12]. Though Piaget's position has fallen out of favor, with sufficient training in formal logic, humans can in fact *exceed* the level of reasoning Piaget called "formal operations," and reach the level in which they can reason in *many* logical systems, as well as about such systems. In this level, humans can also *create* logical systems. In general, this level of reasoning is reached by professionals in the formal sciences.

[2] E.g., in [13] it's shown that cooperative and collaborative learning in mathematics is very effective, and in [14] it's shown that group performance is generally qualitatively and quantitatively superior to the performance of the average individual.

The Rensselaer AI & Reasoning (RAIR) Lab

RAIR Lab Method

RAIR Lab Method

- Isolate and dissect human ingenuity.

RAIR Lab Method

- Isolate and dissect human ingenuity.
  - With sufficient time, energy, and motivation, humans could communicate and collaborate across groups and agencies w/o any high tech (though some insuperable obstacles would arise).

RAIR Lab Method

- Isolate and dissect human ingenuity.
  - With sufficient time, energy, and motivation, humans could communicate and collaborate across groups and agencies w/o any high tech (though some insuperable obstacles would arise).
- Formalize weak correlate to this ingenuity in advanced logical systems.

RAIR Lab Method

- Isolate and dissect human ingenuity.
  - With sufficient time, energy, and motivation, humans could communicate and collaborate across groups and agencies w/o any high tech (though some insuperable obstacles would arise).
- Formalize weak correlate to this ingenuity in advanced logical systems.
  - We must have such systems, otherwise we will not be able to query that which we've represented, get answers and justifications back.

RAIR Lab Method

- Isolate and dissect human ingenuity.
  - With sufficient time, energy, and motivation, humans could communicate and collaborate across groups and agencies w/o any high tech (though some insuperable obstacles would arise).
- Formalize weak correlate to this ingenuity in advanced logical systems.
  - We must have such systems, otherwise we will not be able to query that which we've represented, get answers and justifications back.
    - We can test the value of a virtual world or 3D representation by asking whether it contains the knowledge to provide answers to: "Are there more than three buildings over 5 stories with flat roofs on *X* street."

RAIR Lab Method

- Isolate and dissect human ingenuity.
  - With sufficient time, energy, and motivation, humans could communicate and collaborate across groups and agencies w/o any high tech (though some insuperable obstacles would arise).
- Formalize weak correlate to this ingenuity in advanced logical systems.
  - We must have such systems, otherwise we will not be able to query that which we've represented, get answers and justifications back.
    - We can test the value of a virtual world or 3D representation by asking whether it contains the knowledge to provide answers to: "Are there more than three buildings over 5 stories with flat roofs on *X* street."
- Implement correlate in working computer programs in accordance with functional/declarative programming paradigm.

RAIR Lab Method

- Isolate and dissect human ingenuity.
  - With sufficient time, energy, and motivation, humans could communicate and collaborate across groups and agencies w/o any high tech (though some insuperable obstacles would arise).
- Formalize weak correlate to this ingenuity in advanced logical systems.
  - We must have such systems, otherwise we will not be able to query that which we've represented, get answers and justifications back.
    - We can test the value of a virtual world or 3D representation by asking whether it contains the knowledge to provide answers to: "Are there more than three buildings over 5 stories with flat roofs on *X* street."
- Implement correlate in working computer programs in accordance with functional/declarative programming paradigm.
- Augment correlate as needed with machine-specific power.

RAIR Lab Method

- Isolate and dissect human ingenuity.
  - With sufficient time, energy, and motivation, humans could communicate and collaborate across groups and agencies w/o any high tech (though some insuperable obstacles would arise).
- Formalize weak correlate to this ingenuity in advanced logical systems.
  - We must have such systems, otherwise we will not be able to query that which we've represented, get answers and justifications back.
    - We can test the value of a virtual world or 3D representation by asking whether it contains the knowledge to provide answers to: "Are there more than three buildings over 5 stories with flat roofs on *X* street."
- Implement correlate in working computer programs in accordance with functional/declarative programming paradigm.
- Augment correlate as needed with machine-specific power.
  - Machines can do it faster; *much* faster.

RAIR Lab Method

- Isolate and dissect human ingenuity.
  - With sufficient time, energy, and motivation, humans could communicate and collaborate across groups and agencies w/o any high tech (though some insuperable obstacles would arise).
- Formalize weak correlate to this ingenuity in advanced logical systems.
  - We must have such systems, otherwise we will not be able to query that which we've represented, get answers and justifications back.
    - We can test the value of a virtual world or 3D representation by asking whether it contains the knowledge to provide answers to: "Are there more than three buildings over 5 stories with flat roofs on *X* street."
- Implement correlate in working computer programs in accordance with functional/declarative programming paradigm.
- Augment correlate as needed with machine-specific power.
  - Machines can do it faster; *much* faster.
  - Security can be retained.

RAIR Lab Method

- Isolate and dissect human ingenuity.
  - With sufficient time, energy, and motivation, humans could communicate and collaborate across groups and agencies w/o any high tech (though some insuperable obstacles would arise).
- Formalize weak correlate to this ingenuity in advanced logical systems.
  - We must have such systems, otherwise we will not be able to query that which we've represented, get answers and justifications back.
    - We can test the value of a virtual world or 3D representation by asking whether it contains the knowledge to provide answers to: "Are there more than three buildings over 5 stories with flat roofs on *X* street."
- Implement correlate in working computer programs in accordance with functional/declarative programming paradigm.
- Augment correlate as needed with machine-specific power.
  - Machines can do it faster; *much* faster.
  - Security can be retained.
- Empower human by handing over system.

Toward a General Logicist Methodology for Engineering Ethically Correct Robots

Selmer Bringsjord, Konstantine Arkoudas, and Paul Bello,
Rensselaer Polytechnic Institute

A deontic logic formalizes a moral code, allowing ethicists to render theories and dilemmas in declarative form for analysis. It offers a way for human overseers to constrain robot behavior in ethically sensitive environments.

As intelligent machines assume an increasingly prominent role in our lives, there seems little doubt they will eventually be called on to make important, ethically charged decisions. For example, we expect hospitals to deploy robots that can administer medications, carry out tests, perform surgery, and so on, supported by software agents, or softbots, that will manage related data. (Our discussion of ethical robots extends to all artificial agents, embodied or not.) Consider also that robots are already finding their way to the battlefield, where many of their potential actions could inflict harm that is ethically impermissible.

How can we ensure that such robots will always behave in an ethically correct manner? How can we know ahead of time, via rationales expressed in clear natural languages, that their behavior will be constrained specifically by the ethical codes affirmed by human overseers? Pessimists have claimed that the answer to these questions is: "We can't!" For example, Sun Microsystems' cofounder and former chief scientist, Bill Joy, published a highly influential argument for this answer.[1] Inevitably, according to the pessimists, AI will produce robots that have tremendous power and behave immorally. These predictions certainly have some traction, particularly among a public that pays good money to see such dark films as Stanley Kubrick's *2001* and his joint venture with Stephen Spielberg, *AI*).

Nonetheless, we're optimists: we think formal logic offers a way to preclude doomsday scenarios of malicious robots taking over the world. Faced with the challenge of engineering ethically correct robots, we propose a logic-based approach (see the related sidebar). We've successfully implemented and demonstrated this approach.[2] We present it here in a general methodology to answer the ethical questions that arise in entrusting robots with more and more of our welfare.

Deontic logics: Formalizing ethical codes

Our answer to the questions of how to ensure ethically correct robot behavior is, in brief, to insist that robots only perform actions that can be proved ethically permissible in a human-selected *deontic logic*. A deontic logic formalizes an ethical code—that is, a collection of ethical rules and principles. Isaac Asimov introduced a simple (but subtle) ethical code in his famous Three Laws of Robotics:[3]

1. A robot may not harm a human being, or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence, as long as such protection does not conflict with the First or Second Law.

Human beings often view ethical theories, principles, and codes informally, but intelligent machines require a greater degree of precision. At present, and for the foreseeable future, machines can't work directly with natural language, so we can't simply feed Asimov's three laws to a robot and instruct it behave in

1541-1672/06/$20.00 © 2006 IEEE
Published by the IEEE Computer Society

Toward a General Logicist Methodology for Engineering Computational Artifacts

Selmer Bringsjord, Konstantine Arkoudas, and Paul Bello,
Rensselaer Polytechnic Institute

A deontic logic formalizes a moral code, allowing ethicists to render theories and dilemmas in declarative form for analysis. It offers a way for human overseers to constrain robot behavior in ethically sensitive environments.

As intelligent machines assume an increasingly prominent role in our lives, there seems little doubt they will eventually be called on to make important, ethically charged decisions. For example, we expect hospitals to deploy robots that can administer medications, carry out tests, perform surgery, and so on, supported by software agents, or softbots, that will manage related data. (Our discussion of ethical robots extends to all artificial agents, embodied or not.) Consider also that robots are already finding their way to the battlefield, where many of their potential actions could inflict harm that is ethically impermissible.

How can we ensure that such robots will always behave in an ethically correct manner? How can we know ahead of time, via rationales expressed in clear natural languages, that their behavior will be constrained specifically by the ethical codes affirmed by human overseers? Pessimists have claimed that the answer to these questions is: "We can't!" For example, Sun Microsystems' cofounder and former chief scientist, Bill Joy, published a highly influential argument for this answer.[1] Inevitably, according to the pessimists, AI will produce robots that have tremendous power and behave immorally. These predictions certainly have some traction, particularly among a public that pays good money to see such dark films as Stanley Kubrick's *2001* and his joint venture with Stephen Spielberg, *AI*).

Nonetheless, we're optimists: we think formal logic offers a way to preclude doomsday scenarios of malicious robots taking over the world. Faced with the challenge of engineering ethically correct robots, we propose a logic-based approach (see the related sidebar). We've successfully implemented and demonstrated this approach.[2] We present it here in a general methodology to answer the ethical questions that arise in entrusting robots with more and more of our welfare.

Deontic logics: Formalizing ethical codes

Our answer to the questions of how to ensure ethically correct robot behavior is, in brief, to insist that robots only perform actions that can be proved ethically permissible in a human-selected *deontic logic*. A deontic logic formalizes an ethical code—that is, a collection of ethical rules and principles. Isaac Asimov introduced a simple (but subtle) ethical code in his famous Three Laws of Robotics:[3]

1. A robot may not harm a human being, or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence, as long as such protection does not conflict with the First or Second Law.

Human beings often view ethical theories, principles, and codes informally, but intelligent machines require a greater degree of precision. At present, and for the foreseeable future, machines can't work directly with natural language, so we can't simply feed Asimov's three laws to a robot and instruct it behave in

1541-1672/06/$20.00 © 2006 IEEE
Published by the IEEE Computer Society

http://www.cogsci.rpi.edu/research/rair

Rensselaer Artificial Intelligence and Reasoning Laboratory

http://www.cogsci.rpi.edu/research/rair/index.php

Google

Rensselaer | DEPARTMENT OF COGNITIVE SCIENCE + Rensselaer Computer Science

Rensselaer > Department of Cognitive Science > Research > RAIR Lab

Rensselaer Artificial Intelligence and Reasoning (RAIR) Laboratory

Home | Projects | People | Lectures | Sponsors | Tour

The **Rensselaer Artificial Intelligence and Reasoning (RAIR) Laboratory** is located in rooms 1112 and 1201 of the Russell Sage Laboratory on the RPI campus.

Research and development in the RAIR Lab ranges across a number of applied projects, as well as across many of the fundamental questions AI raises (e.g., Are we machines ourselves? If so, what sort of machines?). Everything is to a high degree unified by the fact that the formalisms, tools, techniques, systems, etc. that underlie the lab's R&D are invariably based on reasoning.

Because of this, logic plays for us a central role (since, after all, logic is the science of reasoning), but reasoning can be implemented in many ways, and so to reach our goals we happily turn to whatever concretization of reasoning gets the job done.

The sentence below is true
The sentence above is false

f(1) = 0. a_(1,1) a_(1,2) a_(1,3)
f(2) = 0. a_(2,1) a_(2,2) a_(2,3)
f(3) = 0. a_(3,1) a_(3,2) a_(3,3)

```
end_of_list.

list(sos).
x + y = y + x.
(x + y) + z = x + (y + z).
n(n(x + y) + n(x + n(y))) = x.
n(C + D) = n(C).
```

RAIR Lab News

Artificially induced: Teaching computers to read first step in developing consciousness
February 20, 2005

"RPI's work will investigate learning and reasoning, both areas that are key to achieving the vision of cognitive systems," said Jan Walker, with DARPA's external relations department. "In addition, while learning and reasoning are generally important, it is also important to be able to measure when a cognitive system has learned. The RPI project will develop ways to help measure when a system has truly learned something."

Rensselaer Researchers Awarded DARPA Grant to Focus on Learning and Reading

http://www.cogsci.rpi.edu/research/rair/projects.php

RAIR Laboratory – Projects

http://www.cogsci.rpi.edu/research/rair/projects.php

Rensselaer | DEPARTMENT OF COGNITIVE SCIENCE + Rensselaer Computer Science

Rensselaer > Department of Cognitive Science > Research > RAIR Lab > Projects

Rensselaer Artificial Intelligence and Reasoning (RAIR) Laboratory

Home | Projects | People | Lectures | Sponsors | Tour

Projects

Advanced Knowledge Representation and Reasoning for Interactive Visualization (AKRRIV)

The AKRRIV project will develop the necessary tools and frameworks to facilitate interoperability between ARIVA systems at the *visual* level. During AKRRIV three systems will be designed and implemented: Vivid-CL, a logic which handles *visual* information; RASCALS^{IA}, a framework for building models of intelligence analysts, including goals, plans, and beliefs; and Director, a system to manage the interaction between systems. Slate will be the first system enhanced with these new developments.

Solomon

While current Q&A systems are competent and useful with respect to the information they process, they are very limited when compared to a conversation an analyst could have with a human who has read the same information. Solomon, a radically new Q&A system that will transcend the limitations of existing systems by approaching real conversation with real humans.

RAIR Lab News

Artificially induced: Teaching computers to read first step in developing consciousness
February 20, 2005

"RPI's work will investigate learning and reasoning, both areas that are key to achieving the vision of cognitive systems," said Jan Walker, with DARPA's external relations department. "In addition, while learning and reasoning are generally important, it is also important to be able to measure when a cognitive system has learned. The RPI project will develop ways to help measure when a system has truly learned something."

Rensselaer

http://www.cogsci.rpi.edu/research/rair/projects.php

RAIR Laboratory – Projects

http://www.cogsci.rpi.edu/research/rair/projects.php

Rensselaer | DEPARTMENT OF COGNITIVE SCIENCE + Rensselaer Computer Science

Rensselaer > Department of Cognitive Science > Research > RAIR Lab > Projects

Rensselaer Artificial Intelligence and Reasoning (RAIR) Laboratory

Home | Projects | People | Lectures | Sponsors | Tour

Projects

Advanced Knowledge Representation and Reasoning for Interactive Visualization (AKRRIV)

The AKRRIV project will develop the necessary tools and frameworks to facilitate interoperability between ARIVA systems at the *visual* level. During AKRRIV three systems will be designed and implemented: Vivid-CL, a logic which handles *visual* information; RASCALS^{IA}, a framework for building models of intelligence analysts, including goals, plans, and beliefs; and Director, a system to manage the interaction between systems. Slate will be the first system enhanced with these new developments.

Solomon

While current Q&A systems are competent and useful with respect to the information they process, they are very limited when compared to a conversation an analyst could have with a human who has read the same information. Solomon, a radically new Q&A system that will transcend the limitations of existing systems by approaching real conversation with real humans.

RAIR Lab News

Artificially induced: Teaching computers to read first step in developing consciousness
February 20, 2005

"RPI's work will investigate learning and reasoning, both areas that are key to achieving the vision of cognitive systems," said Jan Walker, with DARPA's external relations department. "In addition, while learning and reasoning are generally important, it is also important to be able to measure when a cognitive system has learned. The RPI project will develop ways to help measure when a system has truly learned something."

Rensselaer

Enhancing/Leveraging ...

Slate

www.cogsci.rpi.edu/slate

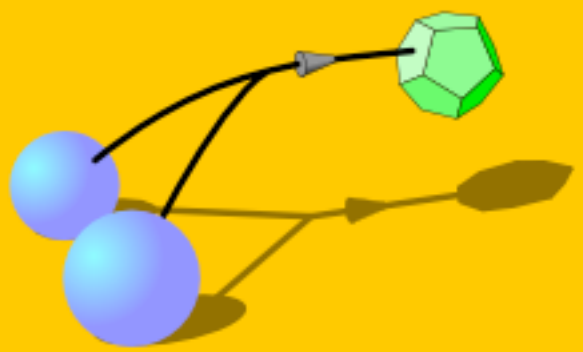

Slate was designed and developed by:
Selmer Bringsjord
Andrew Shilliday
Joshua Taylor

With valuable suggestions from:
Marc Destefano, Wayne Gray,
Michael Schoelles, Jason Wodicka,
and Micah Clark.

Slate is the property of Rensselaer Polytechnic Institute (RPI) and the Rensselaer Artificial Intelligence and Reasoning (RAIR) Lab. When officially released, sponsors and general contractors enjoy an unrestricted license to the system.

The Problem; The Solution ...

Four Overarching Goals*

*These goals are applicable to current RAIR Lab projects in A-SpaceX. No claim is made that such goals apply to other projects, other prospective projects, and so on.

Four Overarching Goals*

1. The IC functions as a seamlessly integrated organization through A-SpaceX, on the shoulders of SOAs & provability-based semantic interoperability.
2. Machines augment the unique intelligence of humans to reach a level of reasoning and decision-making power beyond what any adversary of the US can muster.
3. All recommendations, hypotheses, ... from IAs are mechanically certified relative to input, and are articulated in accordance with normatively correct, bias-free reasoning.
4. IAs train using state-of-the-art intelligent software on case studies and in simulations, and cultivate their capacity for normatively correct, bias-free reasoning.

*These goals are applicable to current RAIR Lab projects in A-SpaceX. No claim is made that such goals apply to other projects, other prospective projects, and so on.

Four Overarching Goals*

1. The IC functions as a seamlessly integrated organization through A-SpaceX, on the shoulders of SOAs & provability-based semantic interoperability.
2. Machines augment the unique intelligence of humans to reach a level of reasoning and decision-making power beyond what any adversary of the US can muster.
3. All recommendations, hypotheses, ... from IAs are mechanically certified relative to input, and are articulated in accordance with normatively correct, bias-free reasoning.
4. IAs train using state-of-the-art intelligent software on case studies and in simulations, and cultivate their capacity for normatively correct, bias-free reasoning.

*These goals are applicable to current RAIR Lab projects in A-SpaceX. No claim is made that such goals apply to other projects, other prospective projects, and so on.

Stories are Needed, but it's Reasoning in the End, and Prescription/Education Can't be Avoided

UNCLASSIFIED // FOUO

A Primer on the Use of Arguments in Slate

Analyzing Pearl Harbor in Late 1941

Micah Clark, Daniel Werner,
Andrew Shilliday, Selmer Bringsjord

Rensselaer AI & Reasoning Laboratory (RAIR)
Department of Cognitive Science
Department of Computer Science
Rensselaer Polytechnic Institute (RPI)
110 8th St., Troy NY 12180

y — Pearl Harbor in Late Nov. 1941

Pearl Harbor case to illustrate how to construct and use a valid **argument**.
gnificant intelligence question that remains familiar; the goal is not to
complete study of Pearl Harbor.[5] Argumentative analysis is not only a
ation, but can be applied just as effectively to many intelligence questions

nsider the available evidence and gaps and try to determine intuitively
on, or, in this case, decision. As one common analogy has it, the analyst
he puzzle — recognizing that many pieces are missing and some that are
r, to use a metaphor common after September 11, 2001, the analyst tries

pothesis and Identifying Primary Reasons

nber 1941; you are an intelligence analyst weighing potential Japanese
rst, identify which hypothesis to investigate using brainstorming and other
r's Analysis of Competing Hypotheses to narrow the range of plausible
pothesis — in this case, that Japan will attack Pearl Harbor — to the Slate

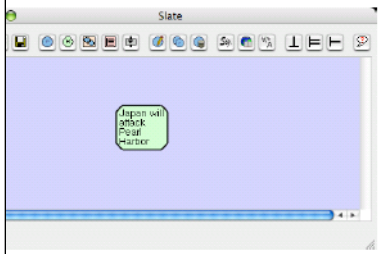

e hypothesis, our next task is to identify the primary reason: the most
reason for accepting the hypothesis. This type of backward-chaining
om conclusions to justifying premises — is often called "goal-directed"

arbor is vast, and this paper uses only a small portion of the available body of evidence
ment. For clarity and illustrative purposes, much of the supporting argument has been
was taken by Hal Ford, *The Purposes and Problems of National Intelligence Estimating*
, 1989), chapter 1, on a hypothetical December 4, 1941, Special National Intelligence

er, Jr., *Psychology of Intelligence Analysis* (Government Printing Office, 1999), chapter
potheses, pp. 95–110.

licators

l for grouping reasons into categories. For example, the reasons for
ch a war against the US can be usefully divided into two types:

war with the US positively.

at Japan will attack US forces soon.

ntive) is a general supposition (or sometimes, a proposition) that out-
e (or disadvantage) in pursuing some course of action.[9] An indicator
f information that suggests that the adversary is planning something.
sitions, i.e. they are typically known to be true.

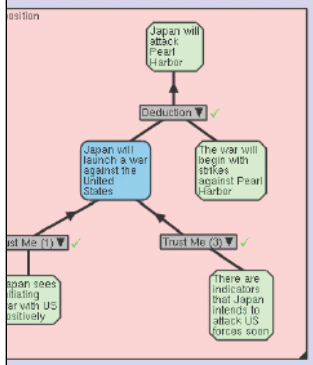

incentive/indicator distinction is found in the 1996 terrorist attack on United States
ers in Saudi Arabia. Terrorists had abundant **incentives** or general reasons for
ding undermining US will to maintain a military presence on Saudi territory. In
rs, including the seizure of explosives being smuggled into Saudi Arabia, threats
detection of surveillance of Khobar Towers. Even more recently, analysts trying
ass destruction (WMD) efforts could identify clear incentives for Iraqi pursuit of
eedom studies detail how analysts misinterpreted and overstated indicators of Iraqi

Why A-SpaceX?

(elevator)

Why A-SpaceX?

(elevator)

"If we all easily and quickly came to know what's important in what others know, and to be able to do that's valuable to us in what others do, we'd each be a *lot* smarter."

“The Problem”

You speak only English, and work in your office with data expressed in English. You want to ask questions of, and get answers back from, thousands of others outside your office, and outside your institution — but unfortunately they speak *other* languages (Spanish, Arabic, Chinese, Norwegian, and so on). Furthermore, all these other people want to keep a good chunk of their data private; you feel the same way about some of your own data. In addition, no has the time to meet face-to-face and go through the painstaking process of learning the other’s language, and then sharing, by hand, piece by piece of information in ways that (in the desired cases) preserve privacy.

The Problem

You have your own way of representing data, and work in your office within *X*/16 with data expressed in this manner. For example, you use particular schemas to represent relational data about financial transactions, suspects, potential sleepers, political figures, etc. You want to "broadcast" information to, ask questions of, and get answers back from, thousands of other analysts and systems outside your office, and outside *X* — but unfortunately they represent their data using different schema, ontologies, formats, and relational databases, work in different domains, and work in different ways. Furthermore, all these other people want to keep a large portion of their data private; you feel the same way about some of your own data. In addition, no one has the time to meet face-to-face and go through the painstaking process of learning the other's schemas and formats, and then sharing, by hand, piece by piece of information in ways that (in certain cases) preserve privacy.

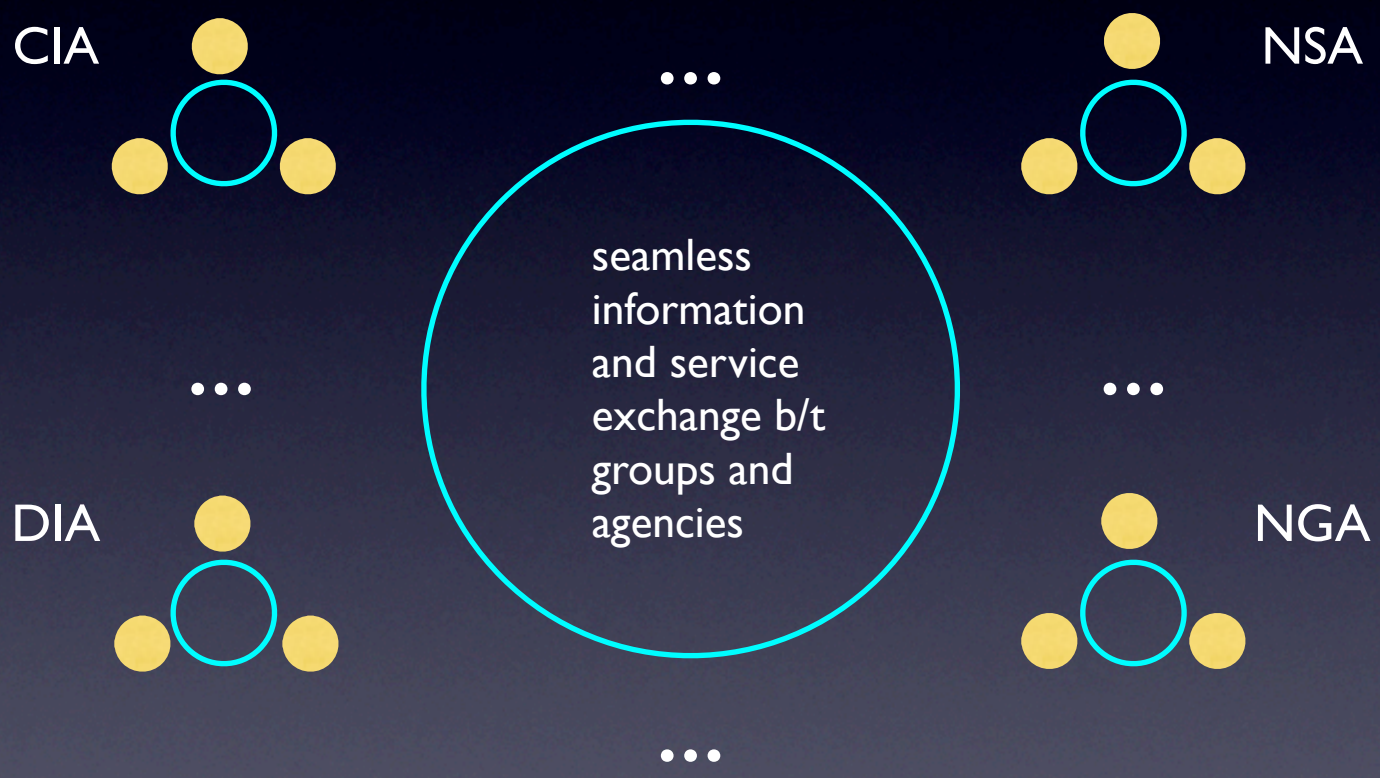
CIA
NSA
DIA
NGA
seamless information and service exchange b/t groups and agencies

The Problem is Solved When Analysts Enter A-SpaceX from the A-Desk

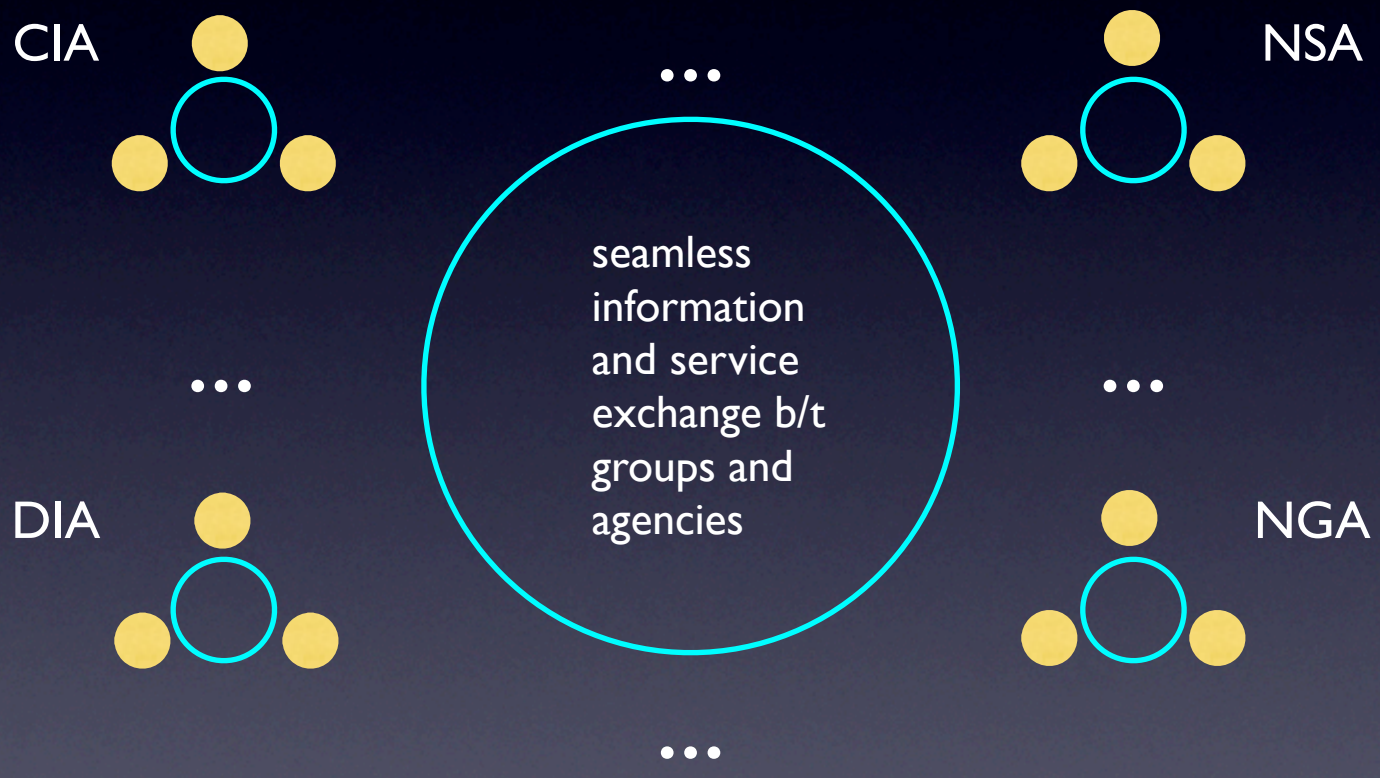

Clarifying 'Interoperability' Tech Presupposed by A-SpaceX/A-Desk

- Service-Oriented Architecture
  - E.g., Expedia uses Sabre. Google's Froogle.
  - For significant SOA-based interaction between systems, PBSI will be required.
- Provability-Based Semantic Interoperability (PBSI)
  - "Hva gjorde Ibsen skriver?" ⇒ "Among other things, Norway's great dramatist wrote *The Master Builder.*"
  - The key to real interoperation.
- Sneaker-Based Node-to-Node Translation/Interoperation
  - Could match PBSI, but out of the question.
- (Microsoft Office) Inter-App Interoperability
  - You imbed an Excel spreadsheet into a Word document, and double-clicking on the table in Word allows you to work with the data etc.
  - By comparison, trivial.
- Mere syntactic translation from one formal language to another.
  - Mother-Of(Rhonda,Josh) ⇒ Mom(Rhonda,Josh)
  - Needed as *part* of PBSI.
- Reducibility
  - Can be exceedingly nuanced, but all syntactic in the end.

Provability-Based Semantic Interoperability ...

Each node an analyst, or close-knit team of analysts. Each link direct, face-to-face translation/interoperation.

$$\frac{n!}{(n-2)!2!}$$

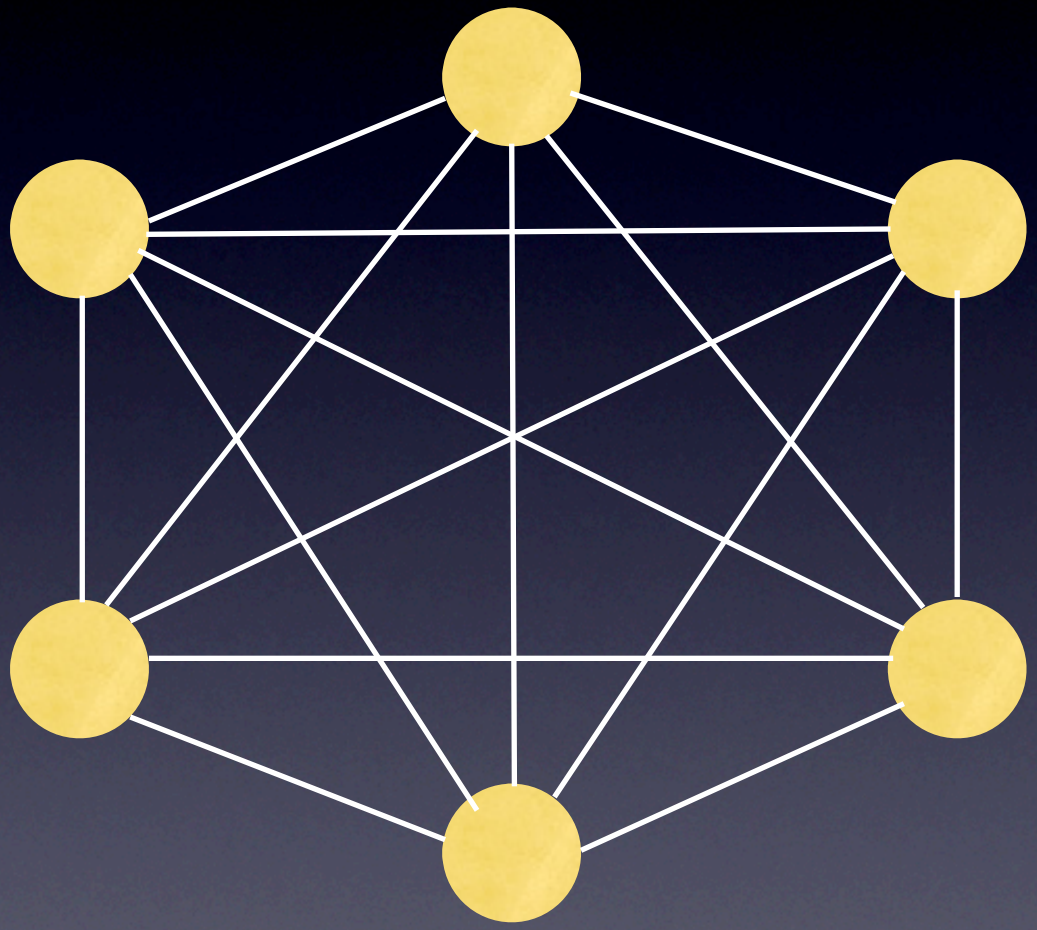

Not a Smart Way to Go

Each node an analyst, or close-knit team of analysts. Each link direct, face-to-face translation/interoperation.

$$\frac{n!}{(n-2)!2!}$$

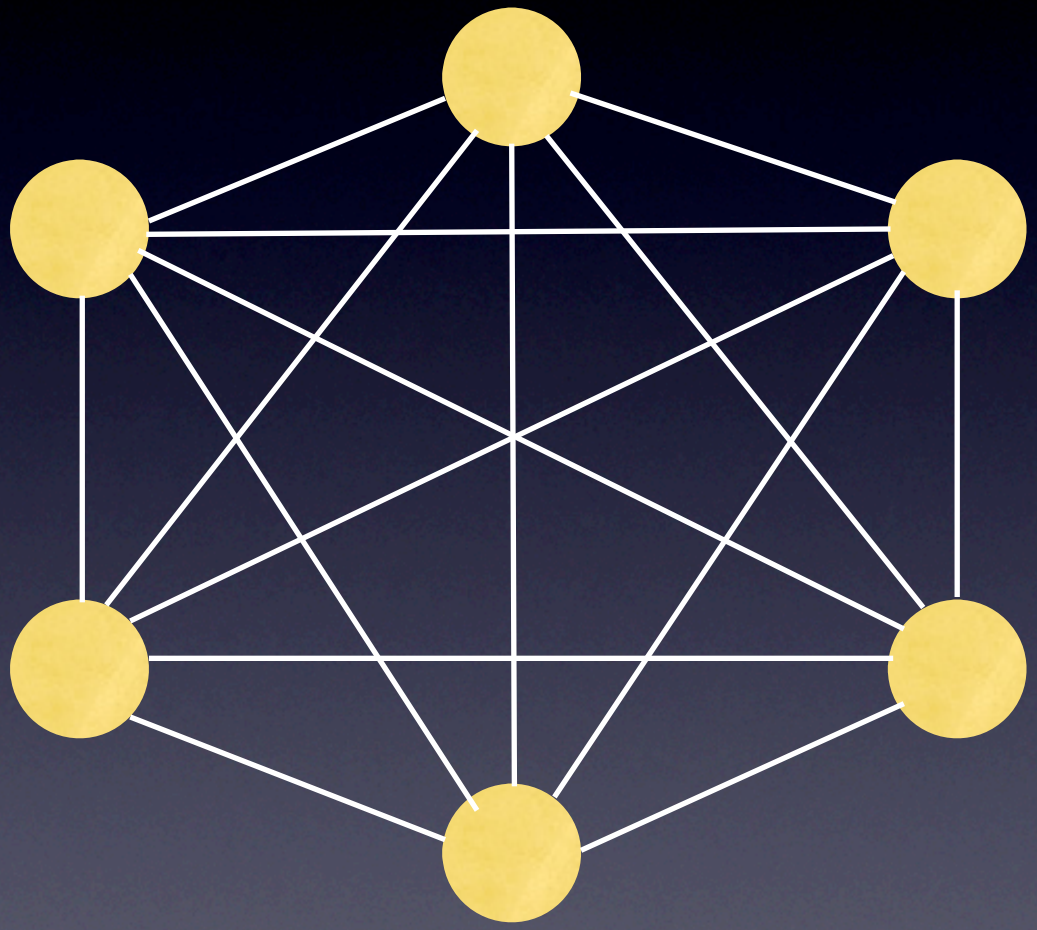

n

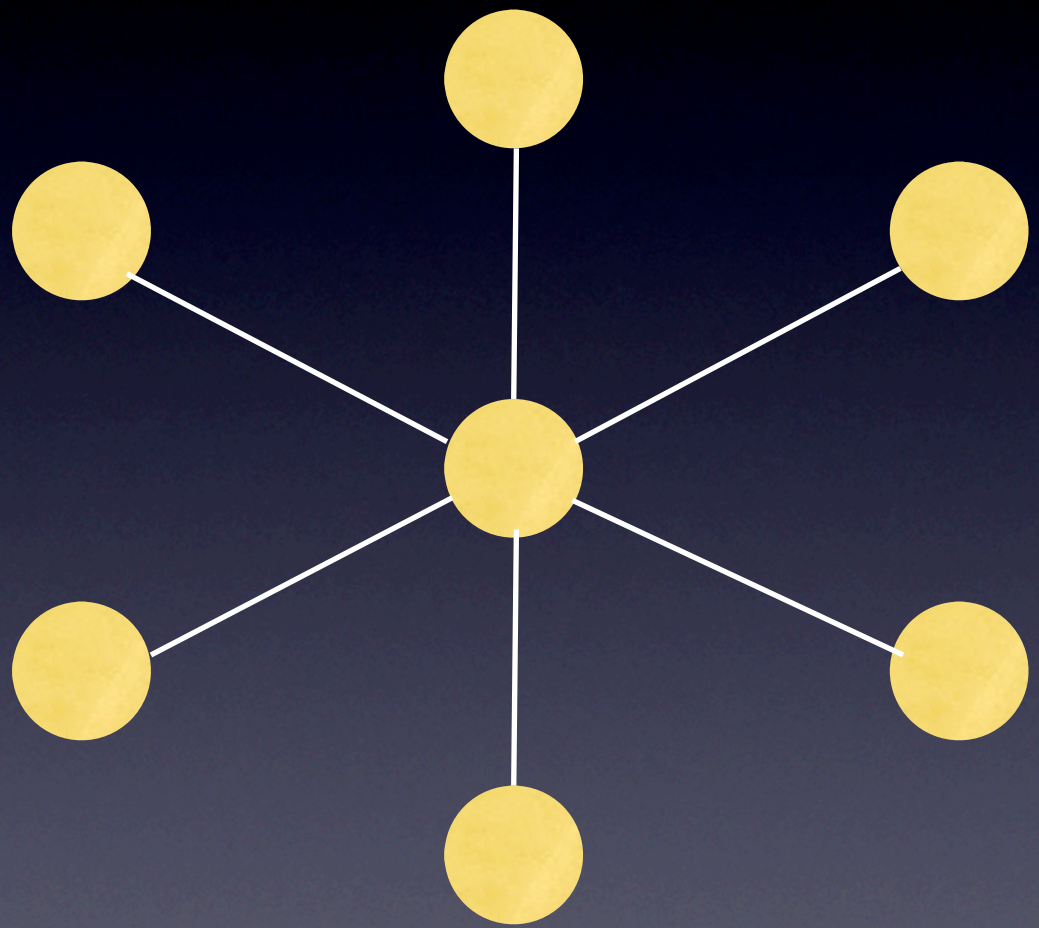

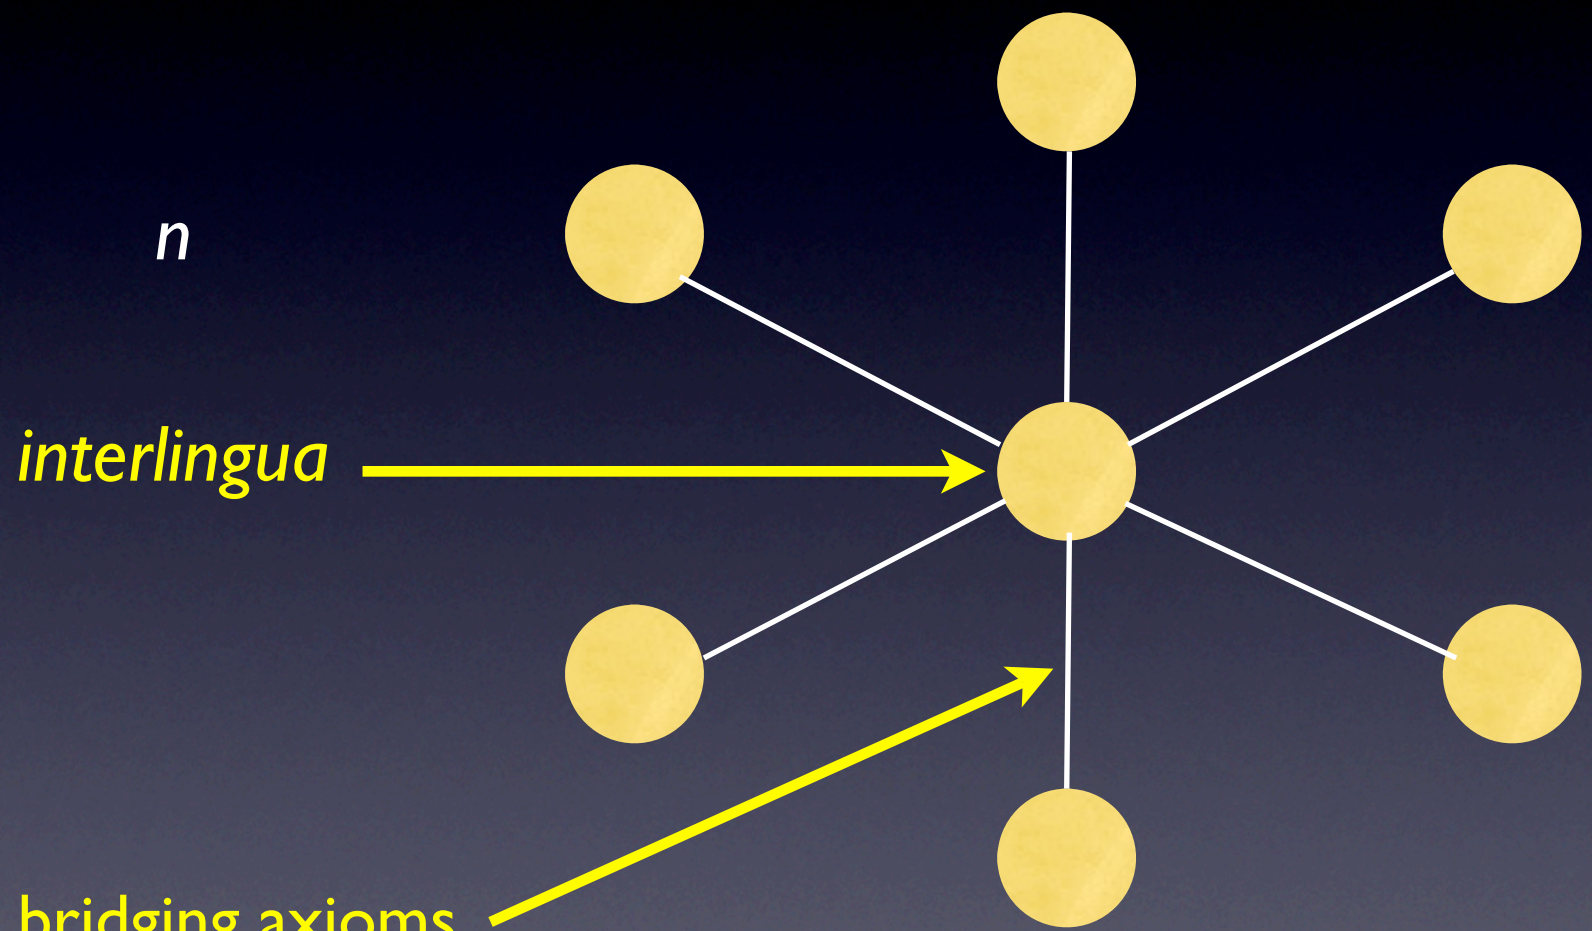
n
interlingua
bridging axioms

This (= Provability-Based Semantic Interoperability) is the Way to Go!

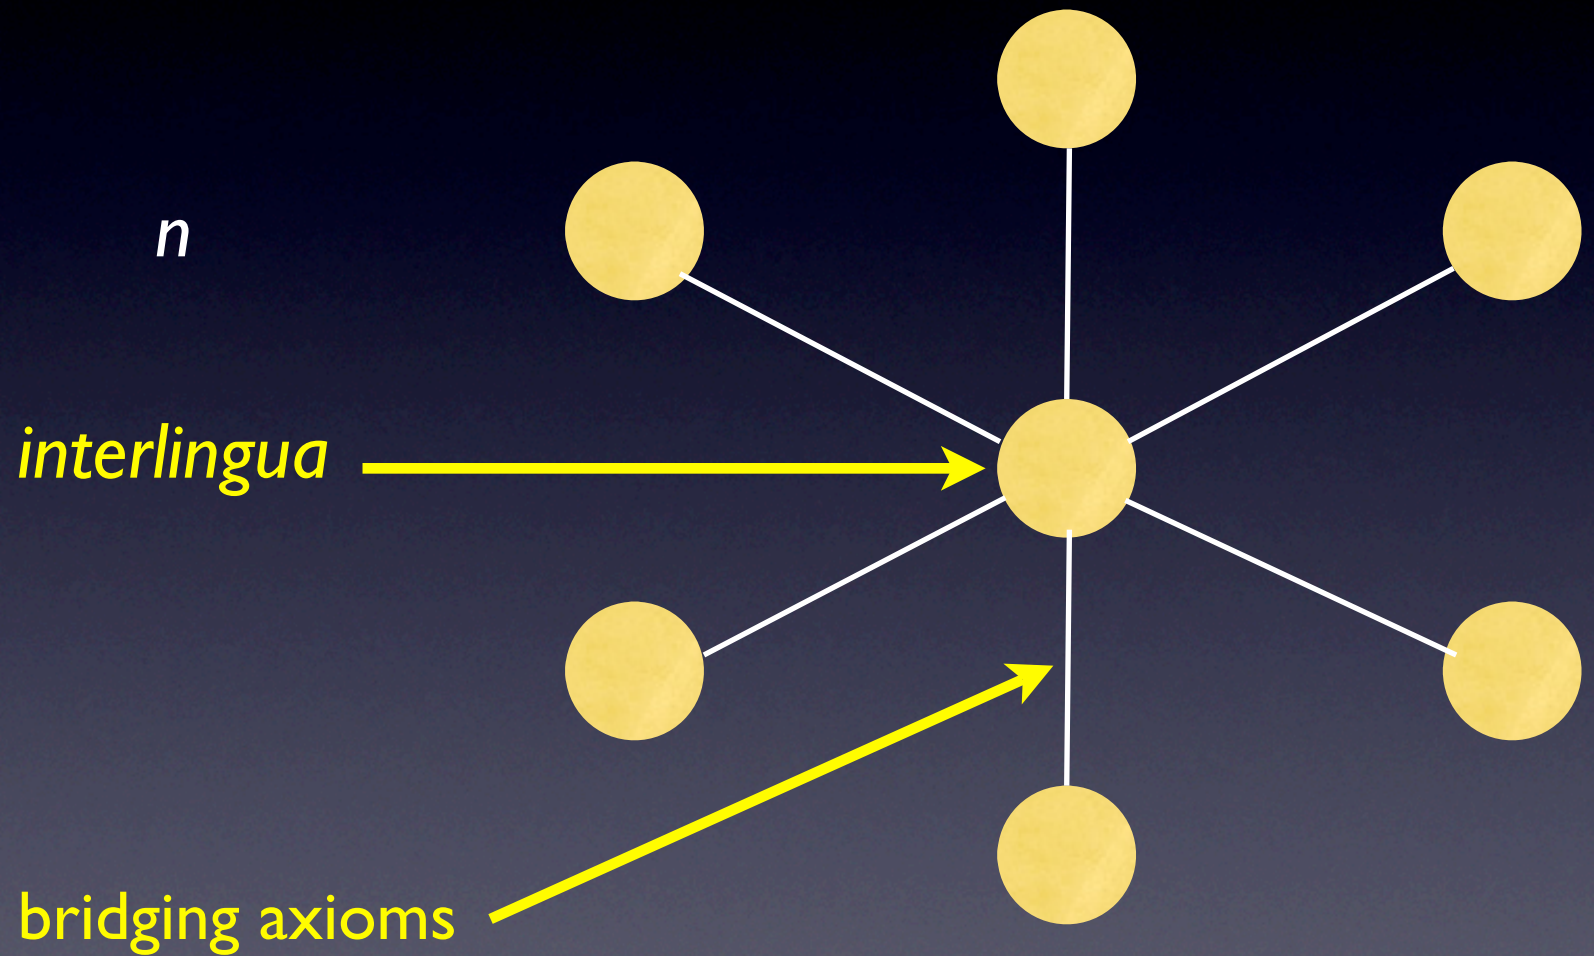

Bridging Axioms for PBSI Derived from Translation Graphs

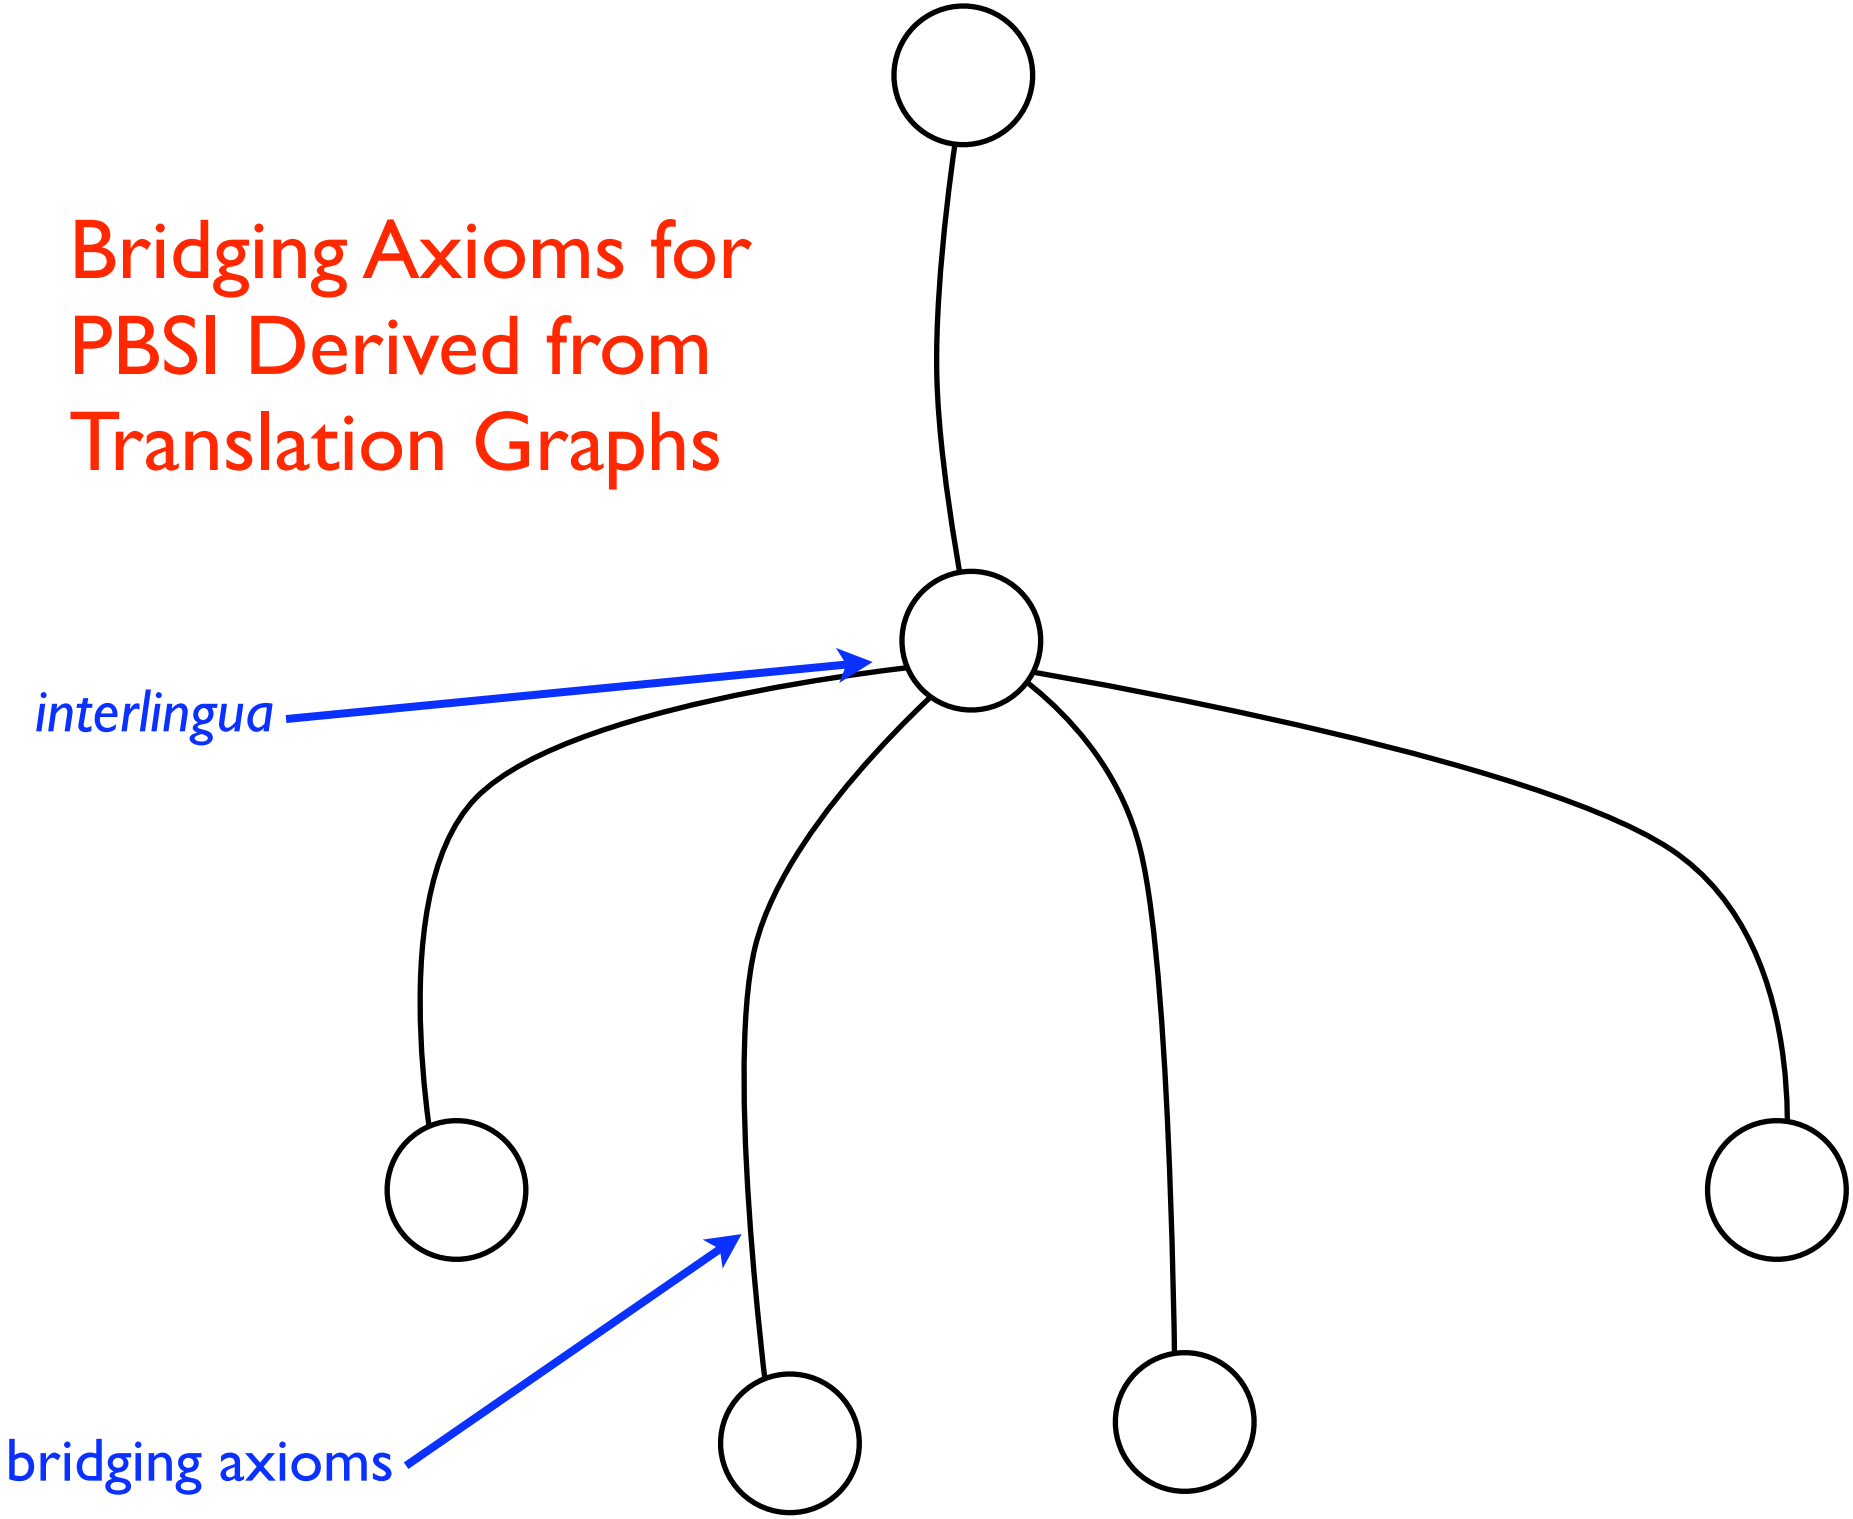

Bridging Axioms for PBSI Derived from Translation Graphs

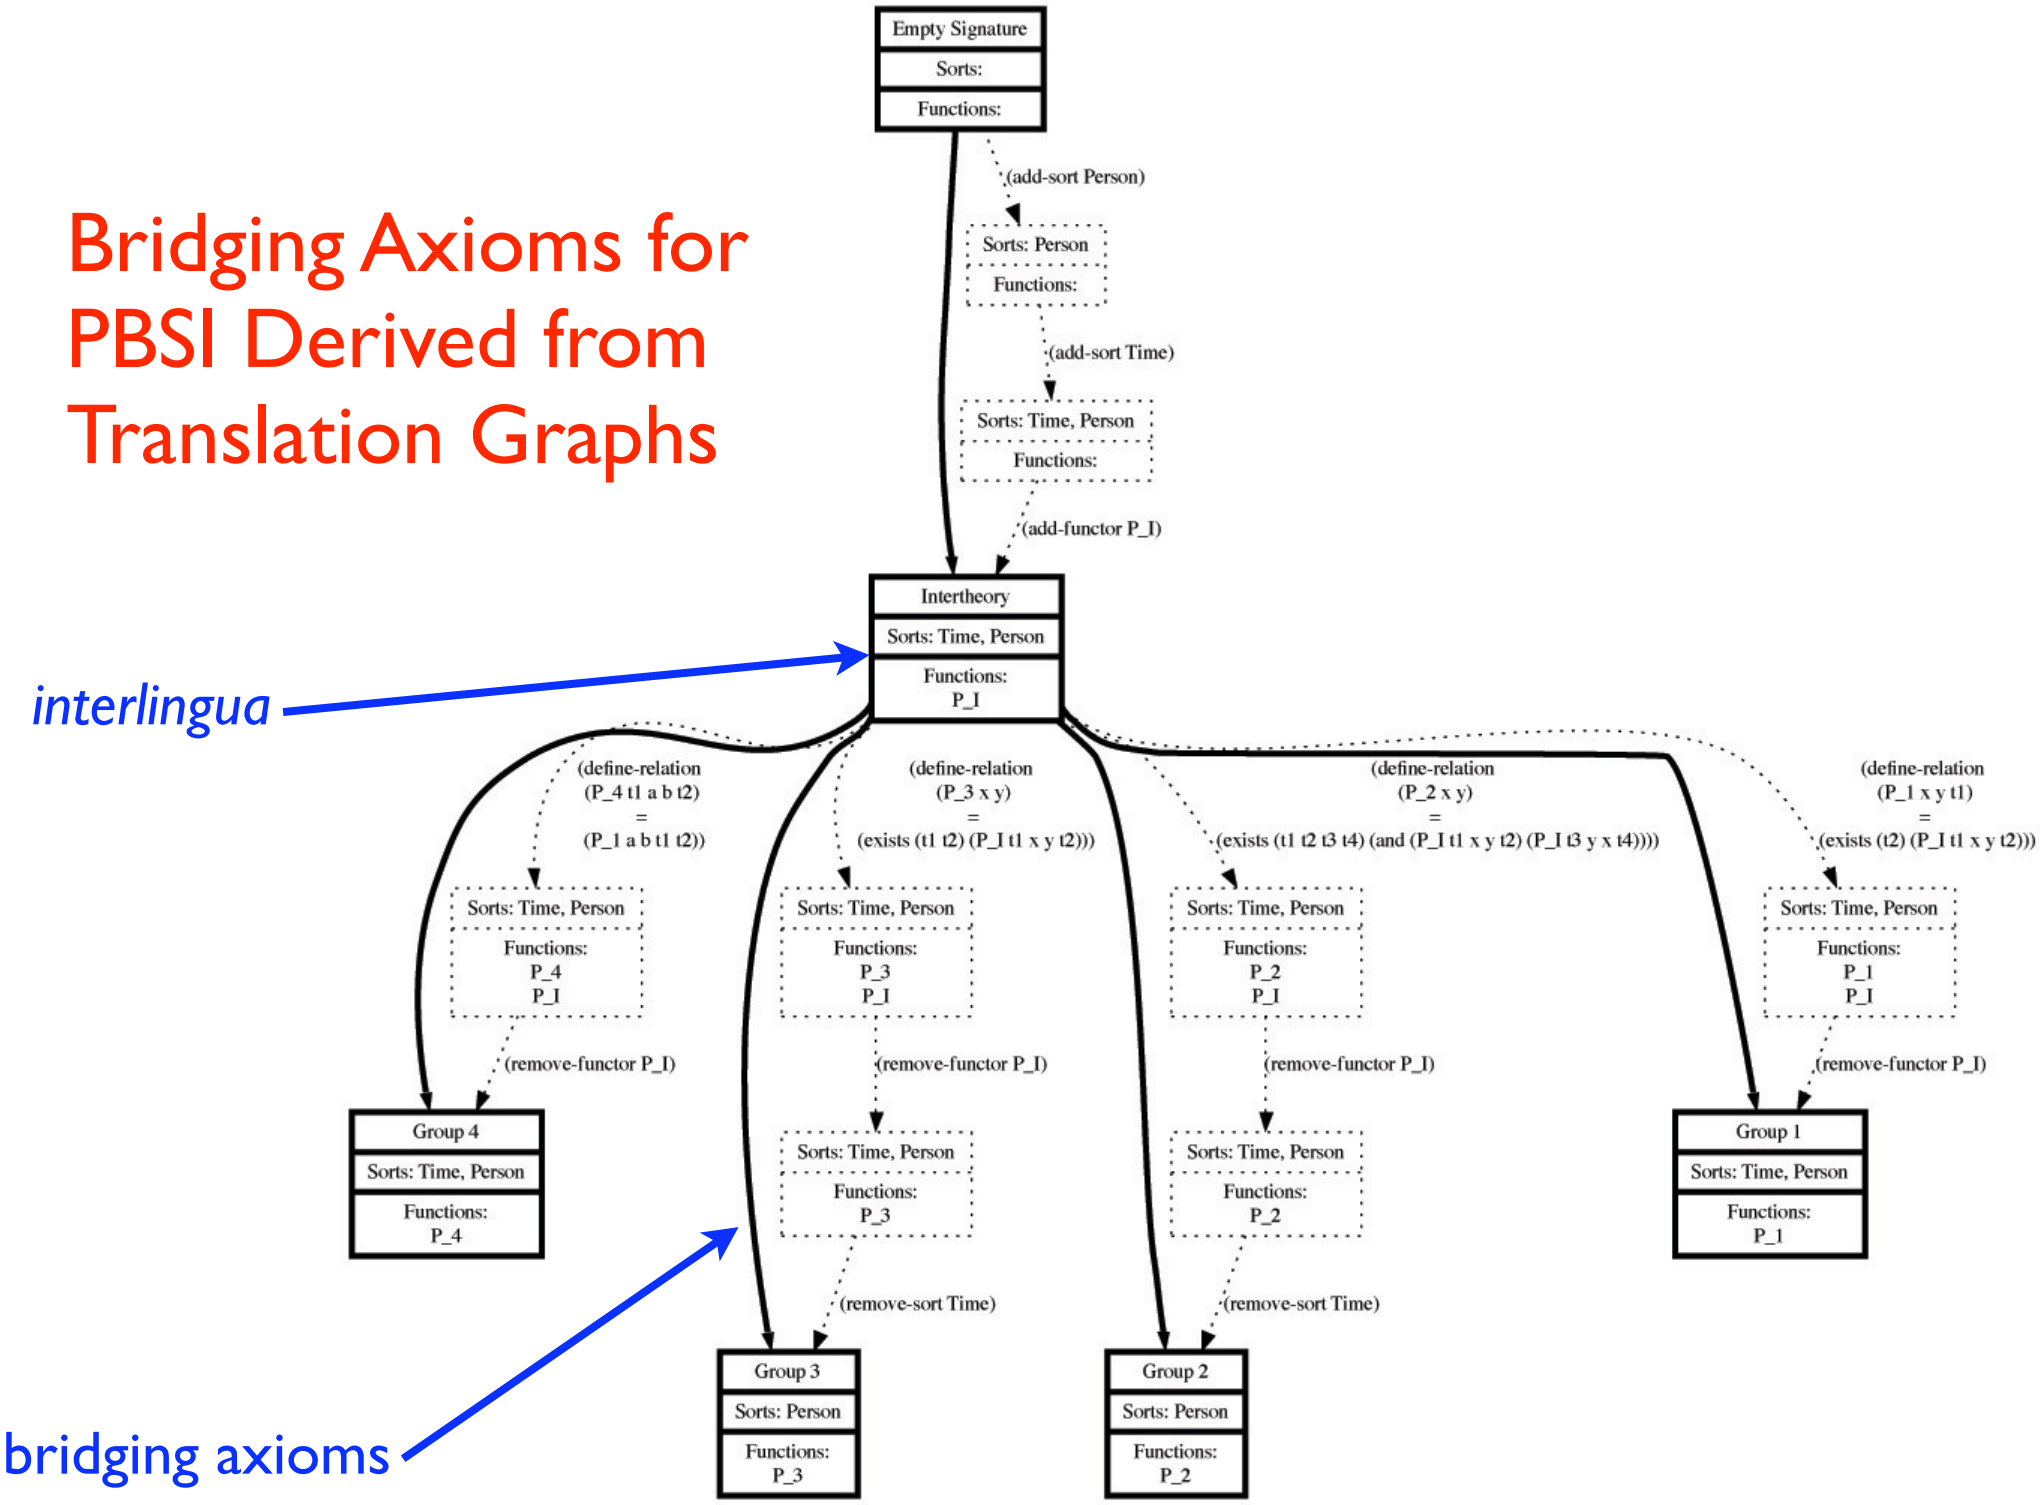

IKL (and its antecedents over the past two decades) is rather powerful. E.g., it's more expressive than atomic formulas (relational databases) and the simple formats e.g. underlying the Semantic Web (RDF, Owl, etc.).

5.3 Integrating New Ontologies into Translation Graphs

Let's introduce a fifth company, χ, who tracks only those phone calls made by its employees to its customers. These
shown, an excerpt of which is shown i
in Figure 4. Suppose that χ has dec
data, and to omit the associations b
to reason, then, that the consumers
about the specifics of each phone c
placed at particular times and with

Table 1. Ex

Customer ID	
43	03/
234	01/
173	01/

Group	Sorts	
χ	$s_{\chi_0}, s_{\chi_1}, s_{\chi_2}$	called $\mapsto f_{\chi}$
		x $\mapsto f_{1_1}$
		t $\mapsto f_{1_2}$
		d $\mapsto f_{1_3}$

Fig. 4. A possible signatur

Provability-Based Semantic Interoperability via Translation Graphs

Joshua Taylor and Andrew Shilliday and Selmer Bringsjord

tayloj,shilla,selmer@rpi.edu
Rensselaer AI & Reasoning (RAIR) Lab
Departments of Cognitive and Computer Science
Rensselaer Polytechnic Institute (RPI), Troy NY 12180, USA

Abstract. *Provability-based semantic interoperability* is a kind of interoperability that transcends mere syntactic translation to allow for robust, meaningful information exchange across systems employing otherwise unresolvable ontologies, and which can be evaluated by provability-based (PB) queries. We introduce a system of *translation graphs* to formalize the relationships between diverse ontologies and knowledge representation and reasoning systems, and to automatically generate the translation axioms governing PB information exchange and inter-system reasoning. We demonstrate the use of translation graphs on a small number of systems to achieve interoperability.

"Provability-Based Semantic Interoperability via Translation Graphs" for ONISW2007

NCOR

http://ncor.us/oic2007

Google

SUNY System...ss – Logon | Ultra Hardw...re Products | Wolfe's Bald...rized Dealer | Screen Door ... Von Morris | Screen Door ... Von Morris | Apple | Amazon | eBay | Yahoo! | News ▾ | .Mac

THE NATIONAL CENTER FOR ONTOLOGICAL RESEARCH - BUFFALO PRESENTS:

OIC-2007

ONTOLOGY FOR THE INTELLIGENCE COMMUNITY:

Towards Effective Exploitation and Integration of Intelligence Resources

November 28 - 29, 2007

Hilton Columbia Hotel • Columbia, Maryland

CALL FOR PAPERS *Closed July 22*	REGISTER NOW!	VENUE & DIRECTIONS	POSTER *Coming Soon*	PROGRAM

The increasing volume, variety and velocity of intelligence analysis requires new approaches that will enable greater flexibility, precision, timeliness and automation of analysis to maximize valuable human resources in responding to fast-evolving threats. Ontology-based technology as applied in areas such as bioinformatics has demonstrated the possibility of gains along all of these dimensions. The time is ripe to extend these gains also to other spheres.

This conference will bring together experts on ontology-based technology with particular experience in the problems facing the intelligence community. It will feature invited talks from prominent ontologists and intelligence community leaders, as well as submitted papers focusing especially on the creation of public-domain ontology resources to support the work of intelligence analysts.

No Reason Why This Approach Can't Be Used to Enable Unified Queries in Today's *Database* Situation in the IC/DoD

Agency 1/Group 1

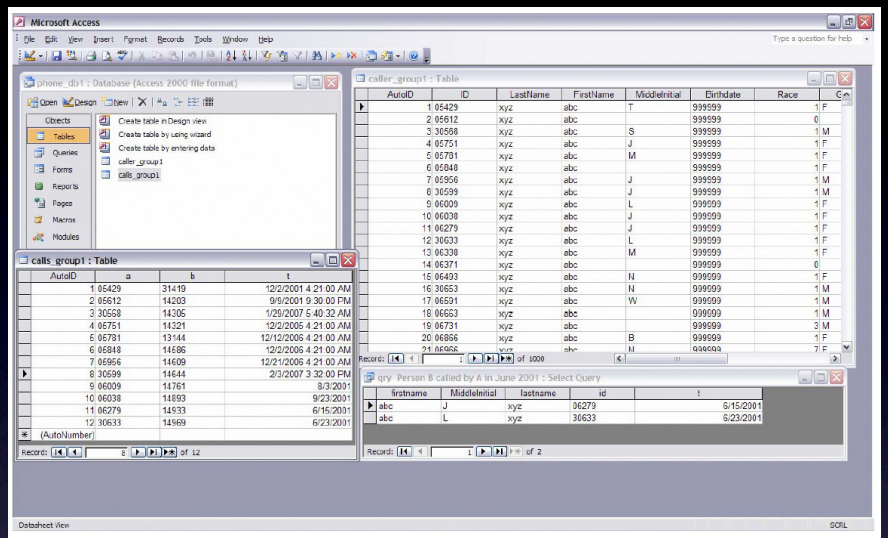

Figure 1. Group 1 – Person b was phoned by a at time t (MS Access).

Agency 2/Group 2

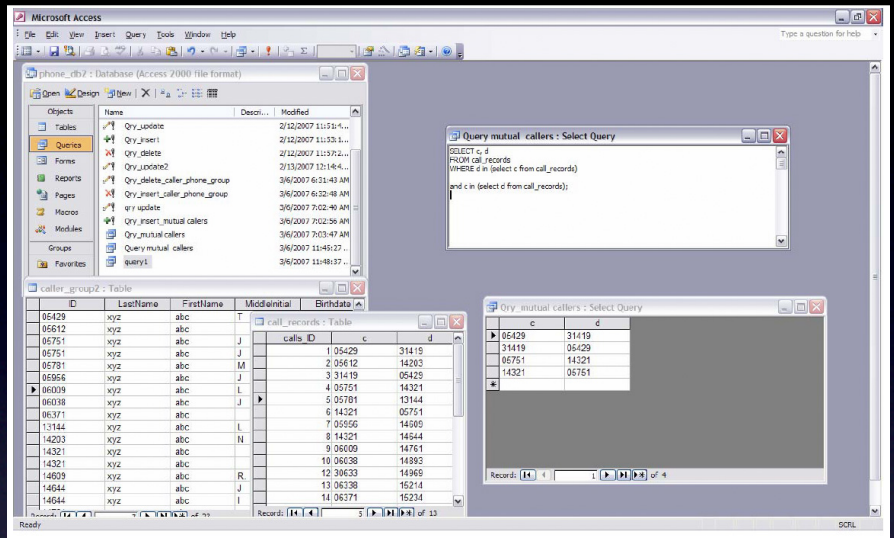

Figure 2. Group 2 – Person c and d phoned each other (MS Access).

Agency 3/Group 3

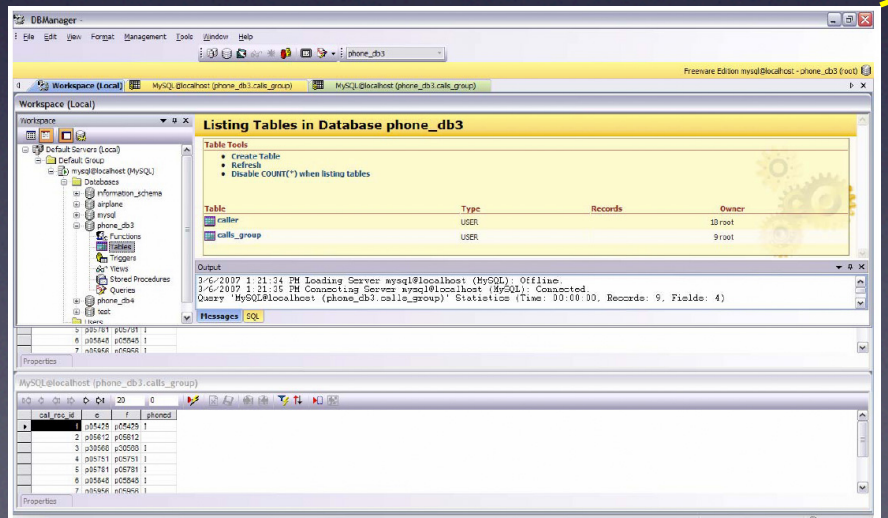

Figure 3. Group 3 – Person e phoned person f (MySql)

Agency 4/Group 4

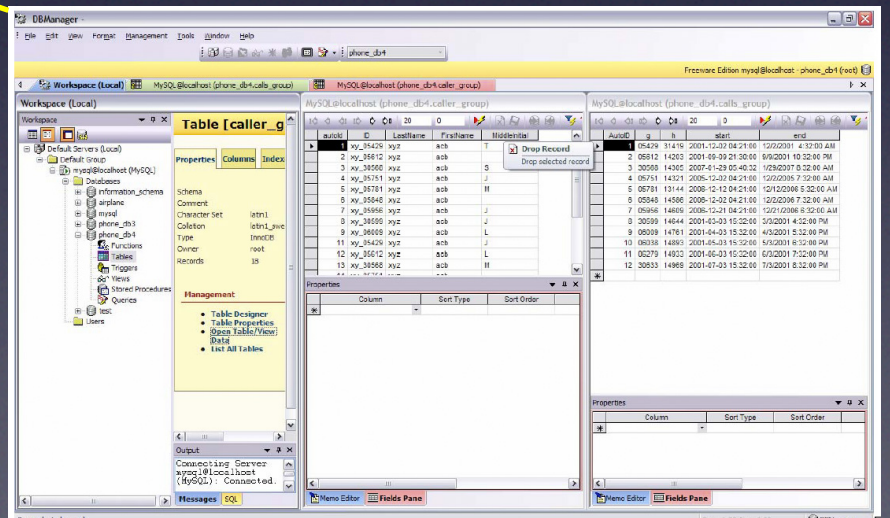

Figure 4. Group 4 – Person g phoned person h at t_start and talked until t_end (MySql)

Semantic Interoperability Demos ...

Slate and GeoTime

A walkthrough of Chart E from
Case Study 4: Sign of the Crescent

06.03.07

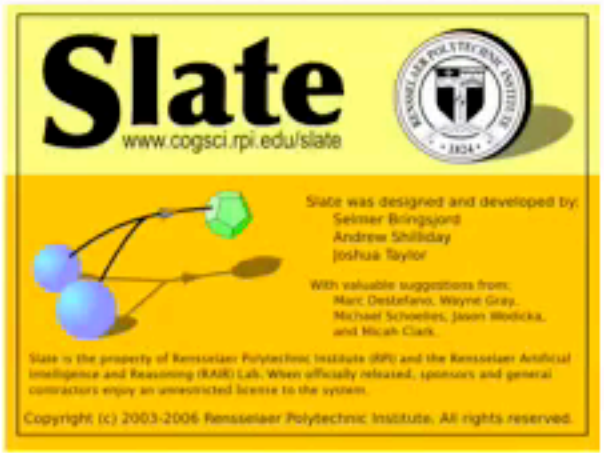

Selmer Bringsjord
Andrew Shilliday
Joshua Taylor
Dan Werner

Rensselaer | DEPARTMENT OF COGNITIVE SCIENCE + Rensselaer Computer Science

Yesterday:
Three-Way Demo:
“Unmasking Palar”

Unmasking Palar

The Cultural Analyst reconstructs and analyzes a report using data received from Oculus' GeoTime and PNNL's AKEA

Selmer Bringsjord
Andrew Shilliday
Micah Clark
Dan Werner

07.09.19

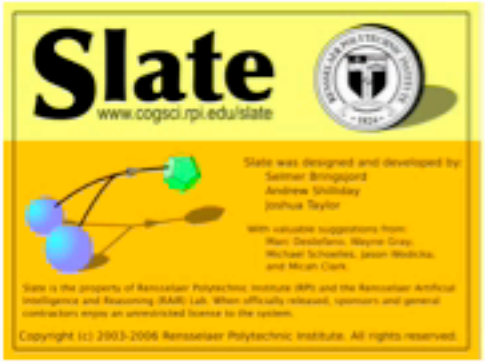

Rensselaer | DEPARTMENT OF COGNITIVE SCIENCE + Rensselaer Computer Science

AKEA

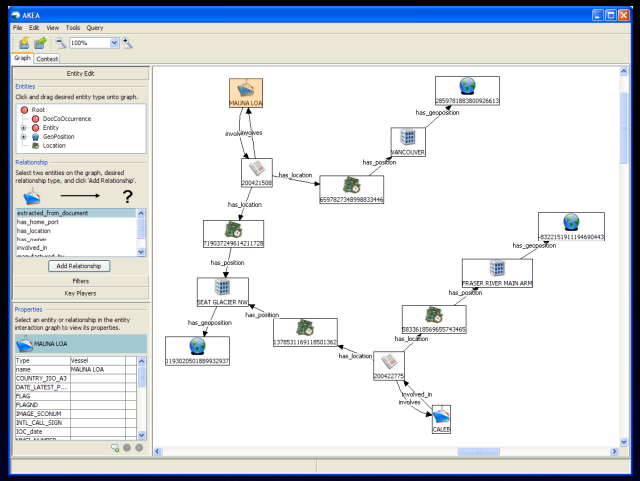

GeoTime

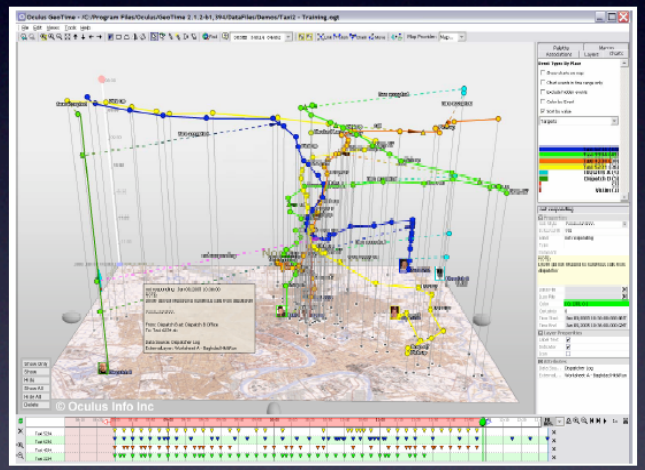

Slate

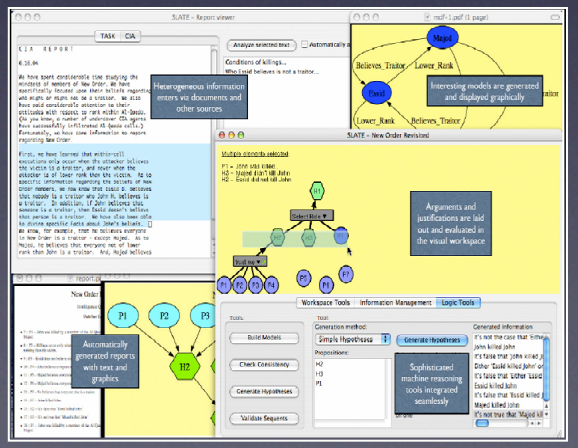

AKEA

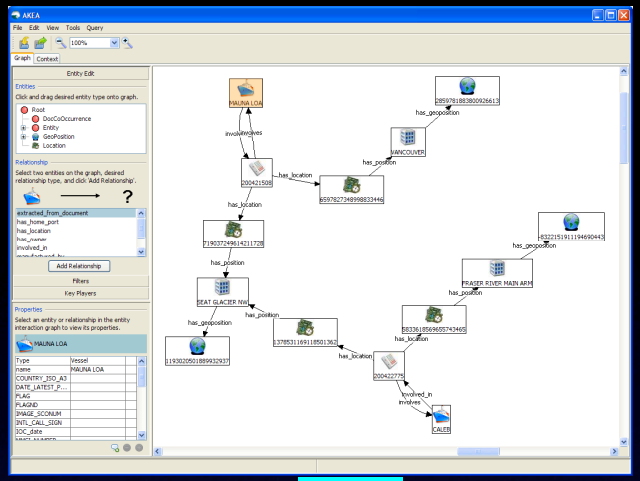

GeoTime

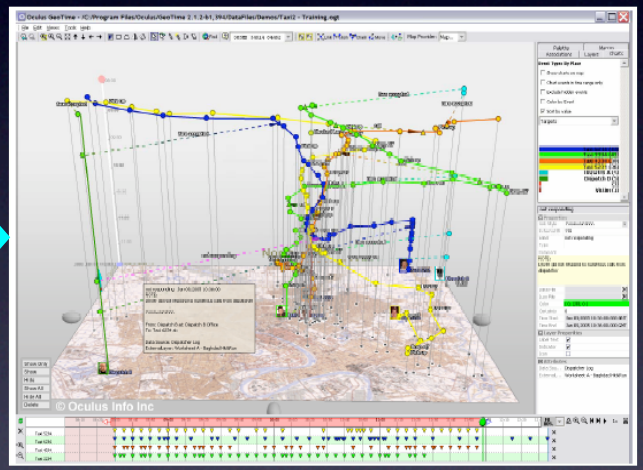

Slate

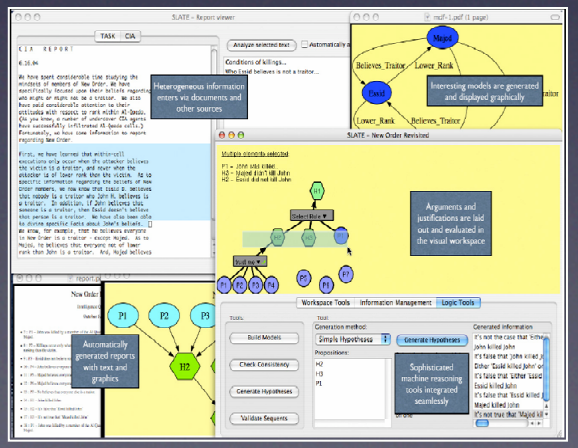

AKEA

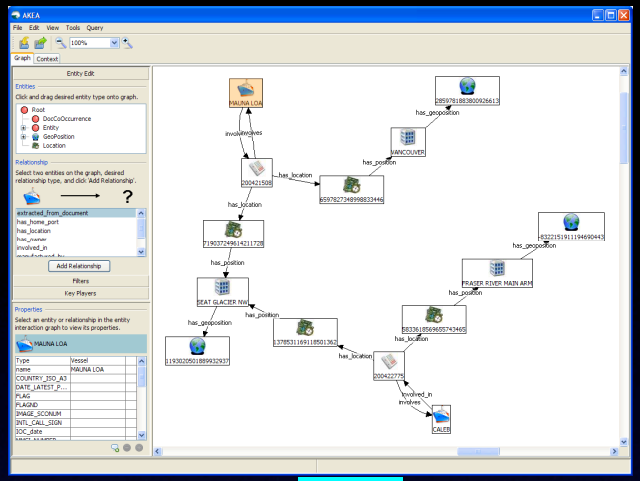

GeoTime

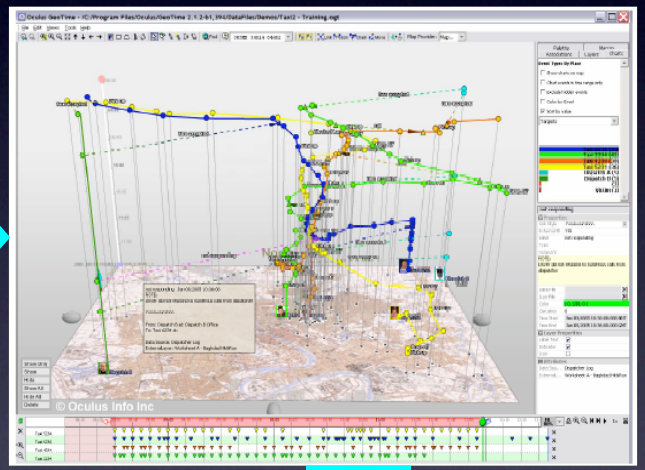

Slate

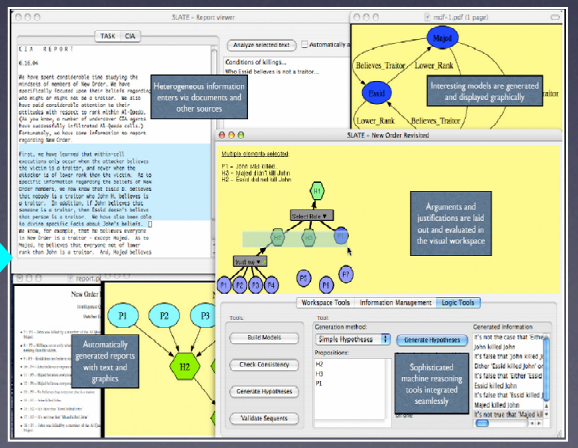

AKEA

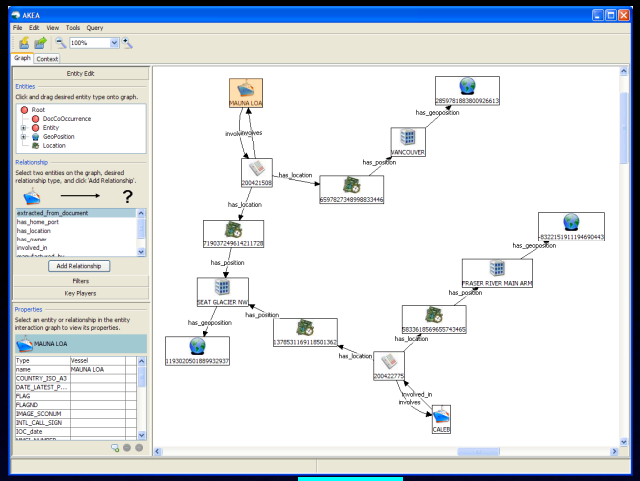

GeoTime

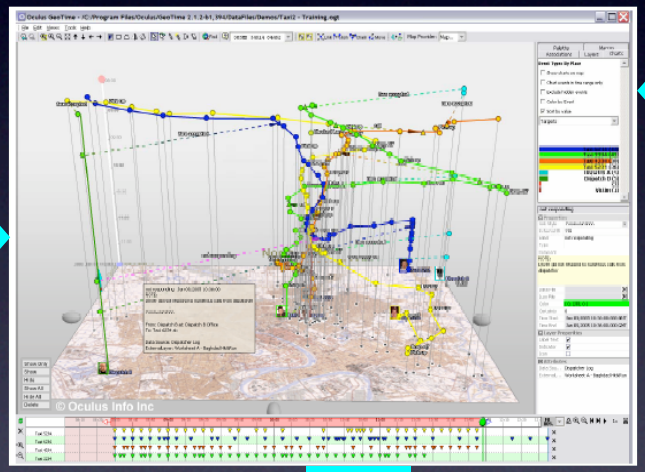

Slate

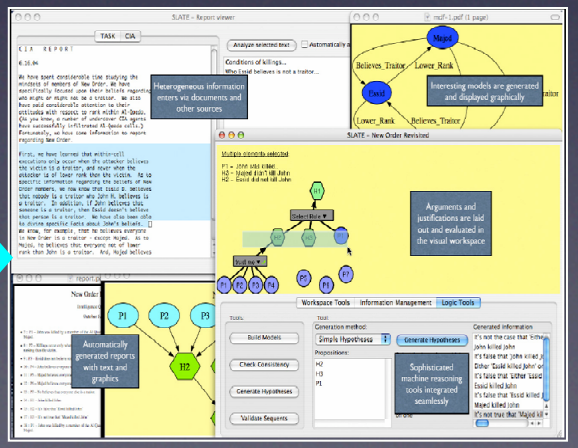

AKEA

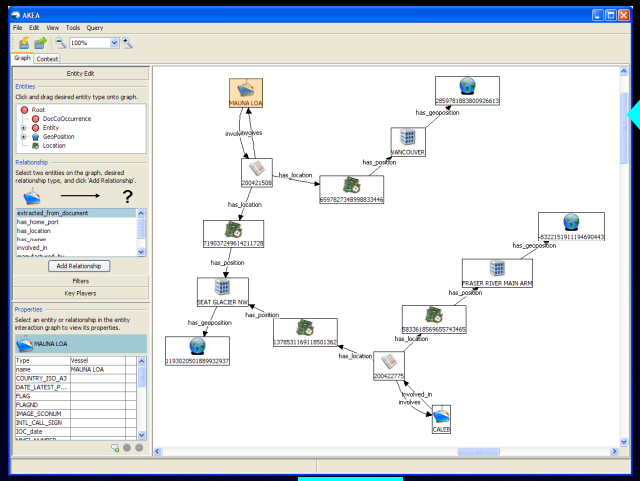

GeoTime

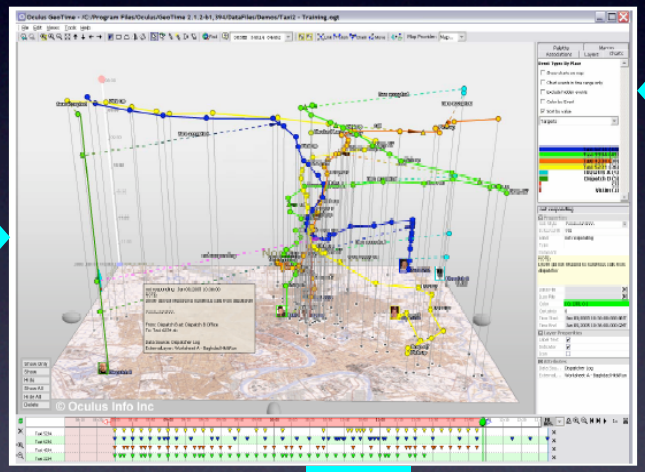

Slate

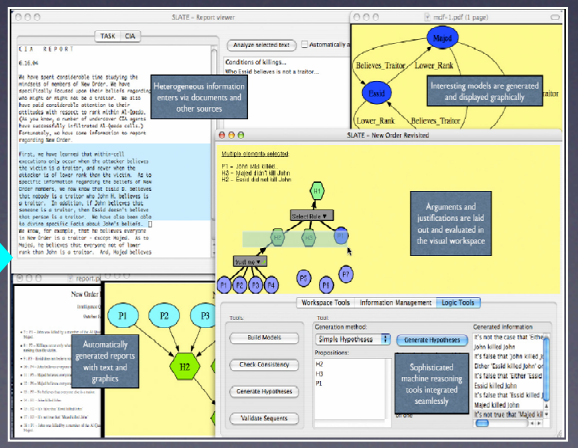

X

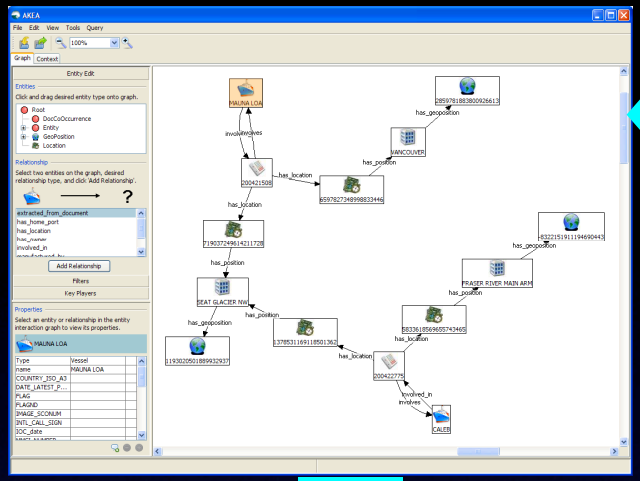
AKLA
File Edit View Tools Query
Entity Edit
Entities
Click and drag desired entity type onto graph.
Root
DocCoOccurrence
Entity
GeoPosition
Location
Relationship
Select two entities on the graph, desired relationship type, and click 'Add Relationship'.
Add Relationship
Filters
Key Players
Properties
Select an entity or relationship in the entity interaction graph to view its properties.
MAUNA LOA
Type Vessel
name MAUNA LOA
COUNTRY_ISO_A3
DATE_LATEST_P...
FLAG
FLAGID
IMAGE_SCONUM
INTL_CALL_SIGN
DOC_date
MAUNA LOA
VANCOUVER
FRASER RIVER MAIN ARM
SEAT GLACIER NW
CALEB
has_geoposition
has_position
has_location
involves
involved_in

Y

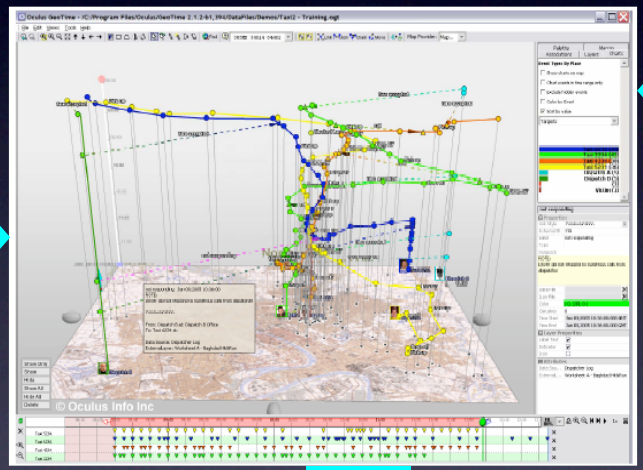
© Oculus Info Inc

Z

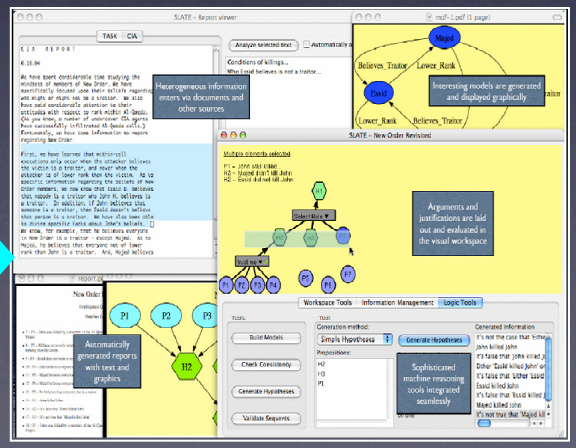
SLATE - Report viewer
TASK CIA
Heterogeneous information enters via documents and other sources
Interesting models are generated and displayed graphically
SLATE - New Order Revisited
Arguments and justifications are laid out and evaluated in the visual workspace
Workspace Tools
Information Management
Logic Tools
Automatically generated reports with text and graphics
Sophisticated machine reasoning tools integrated seamlessly
Build Models
Check Consistency
Generate Hypotheses
Validate Sequents

Report Analyzed Into Slate

Report Analyzed Into Slate

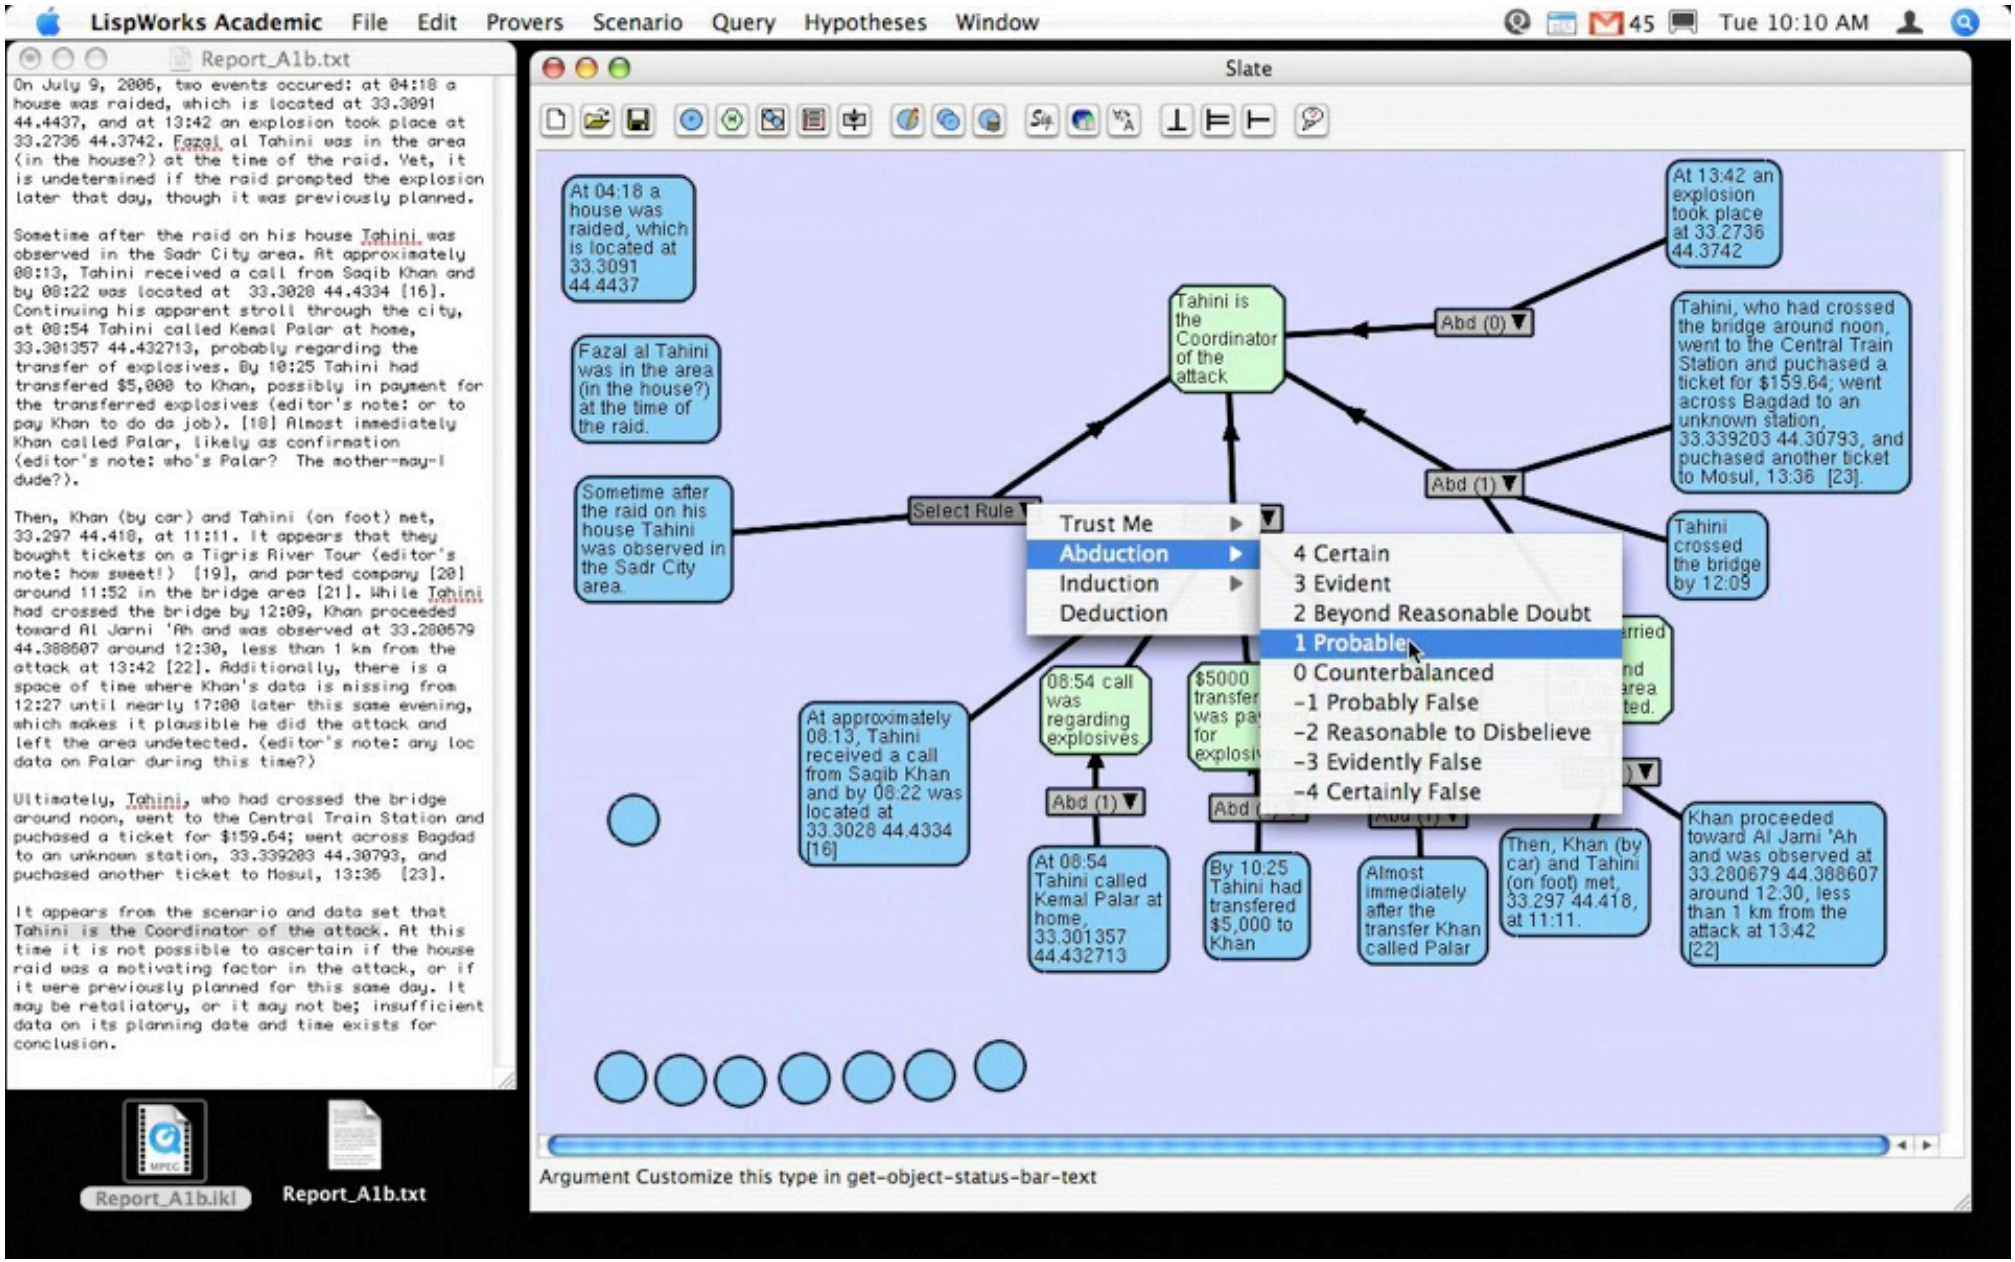

The New Hypothesis Confirmed; Sent

The New Hypothesis Confirmed; Sent

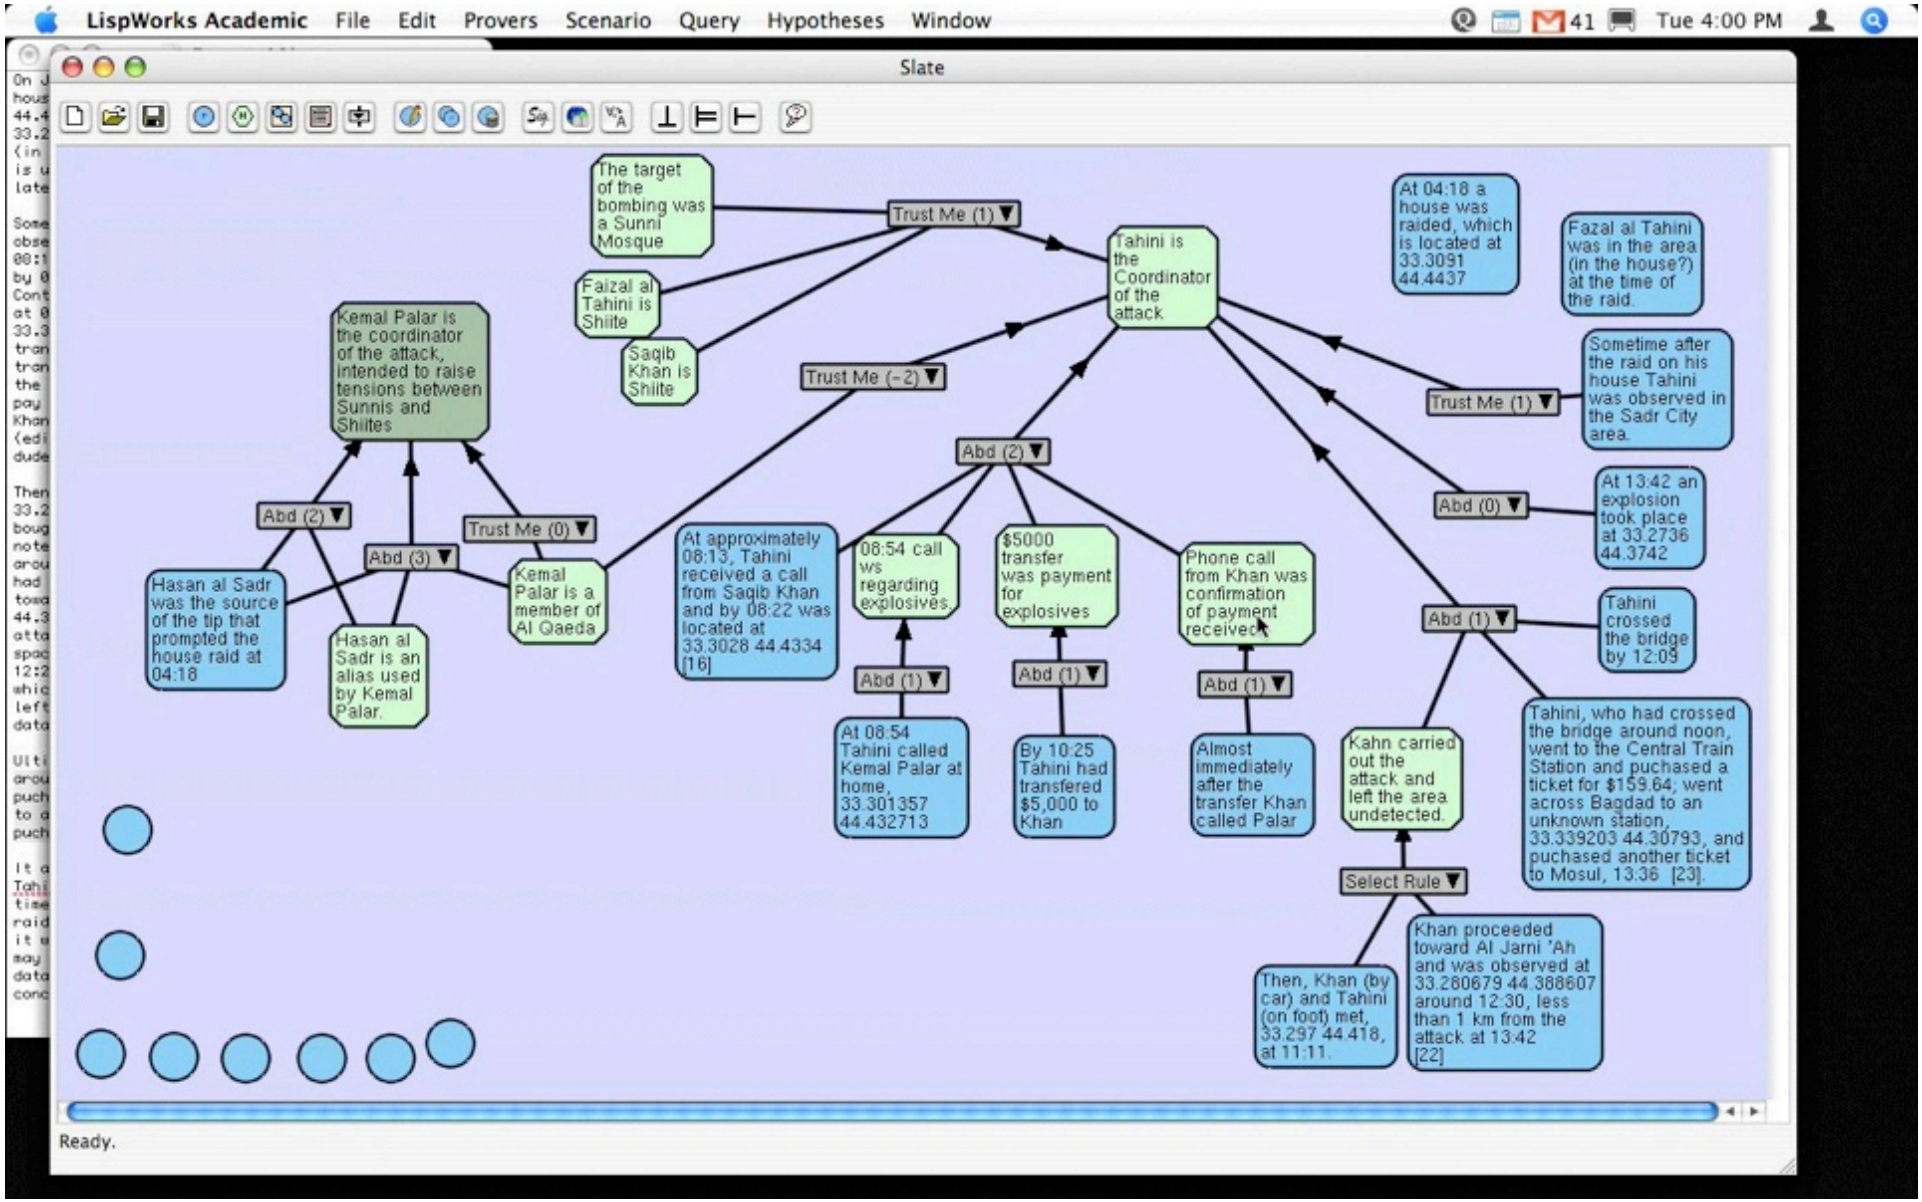

The New Hypothesis Confirmed; Sent

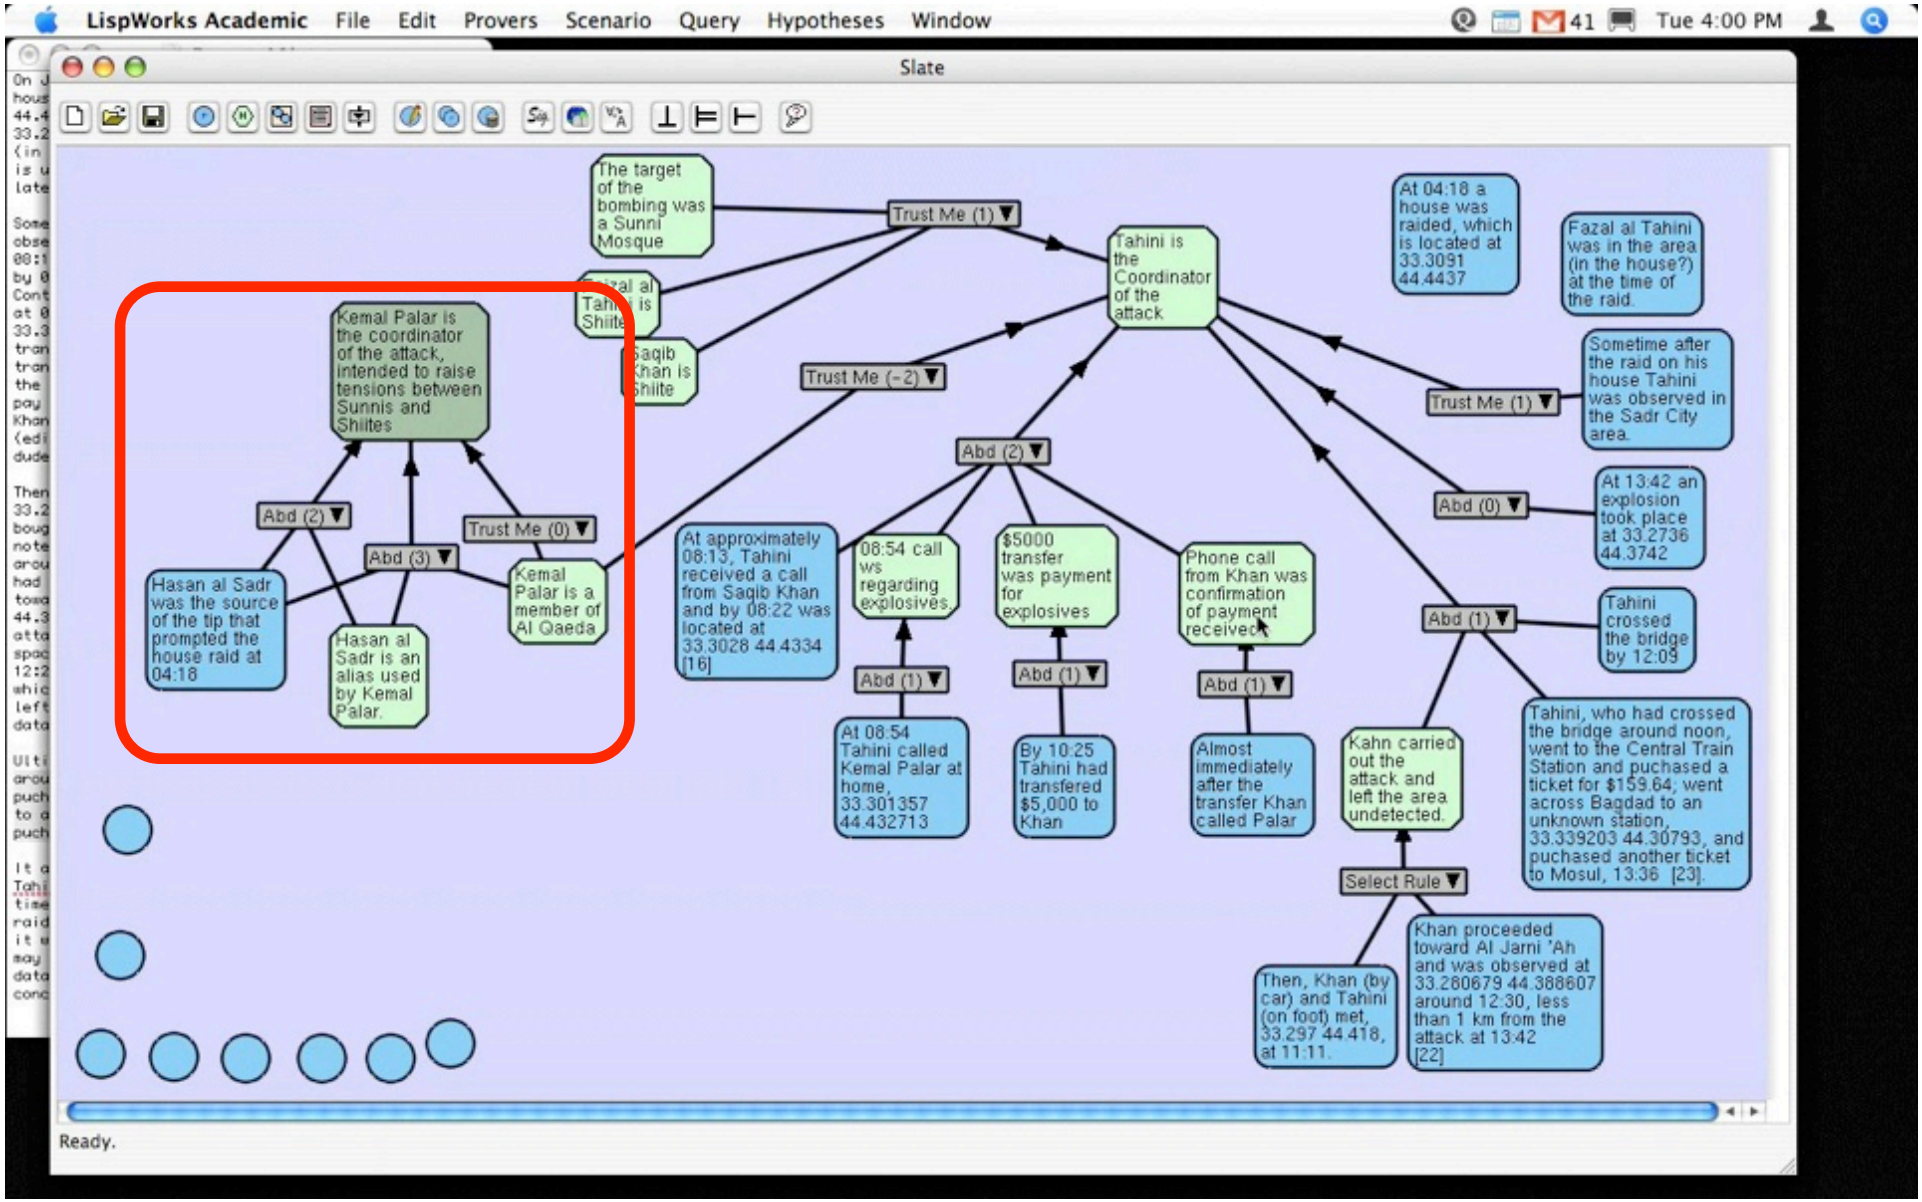

And ...

Cognitively Robust Virtual Characters for Virtual Worlds and the Time Machine ...

Not Cognitively Robust — Yet

Not Cognitively Robust — Yet

Not Cognitively Robust — Yet

Not Cognitively Robust — Yet

Not Cognitively Robust — Yet

Not Cognitively Robust — Yet

Galatea.z8

You come around a corner, away from the noise of the opening.

There is only one exhibit. She stands in the spotlight, with her back to you: a sweep of pale hair on paler skin, a column of emerald silk that ends in a pool at her feet. She might be the model in a perfume ad; the trophy wife at a formal gathering; one of the guests at this very opening, standing on an empty pedestal in some ironic act of artistic deconstruction --

You hesitate, about to turn away. Her hand balls into a fist.

"They told me you were coming."

Not Cognitively Robust — Yet

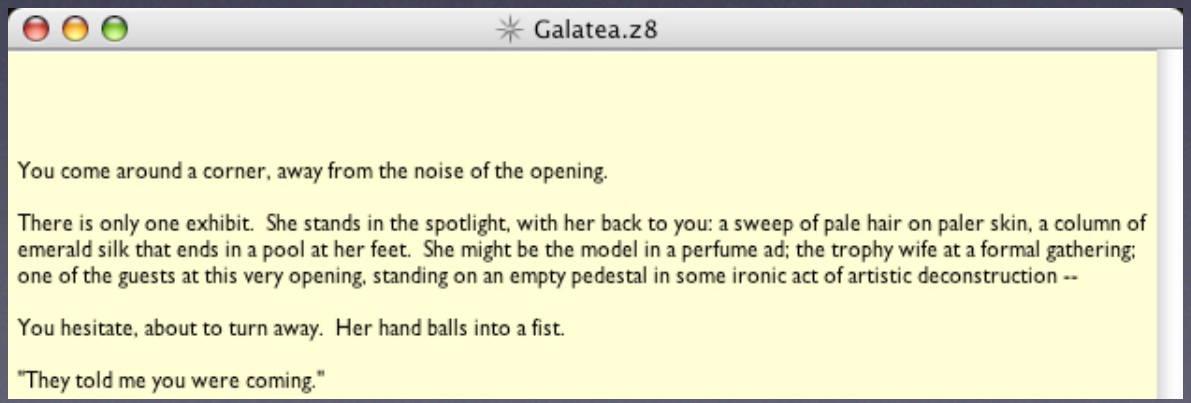

Same Thing Here:
Definitely *Not* a Cognitively Robust SC!

Same Thing Here:
Definitely *Not* a Cognitively Robust SC!

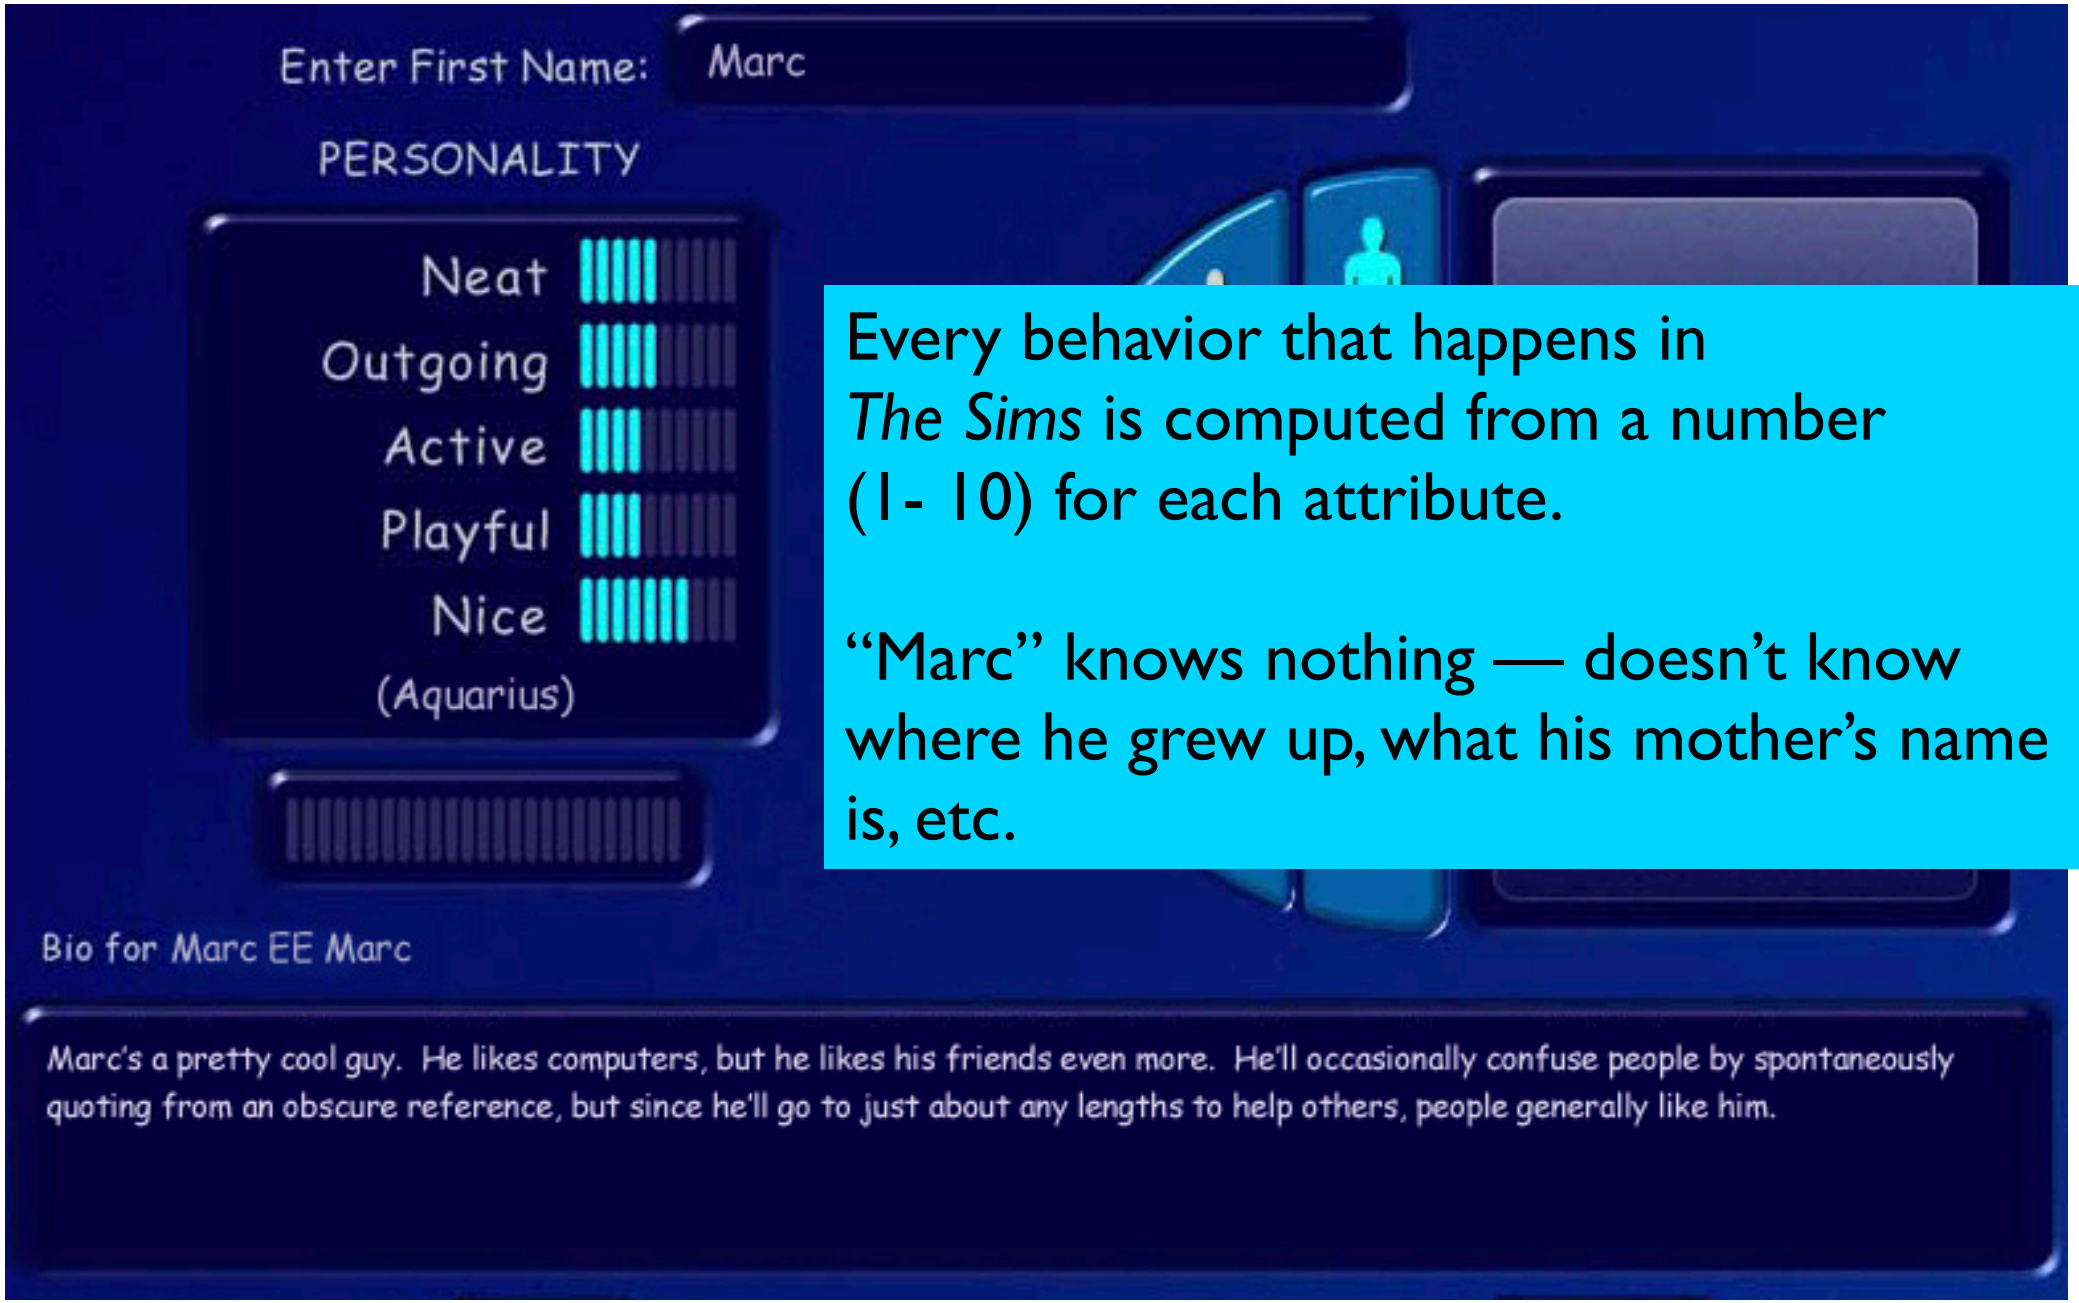

Every behavior that happens in *The Sims* is computed from a number (1- 10) for each attribute.

"Marc" knows nothing — doesn't know where he grew up, what his mother's name is, etc.

Same Thing Here:
Definitely *Not* a Cognitively Robust SC!

Same Thing Here:
Definitely *Not* a Cognitively Robust SC!

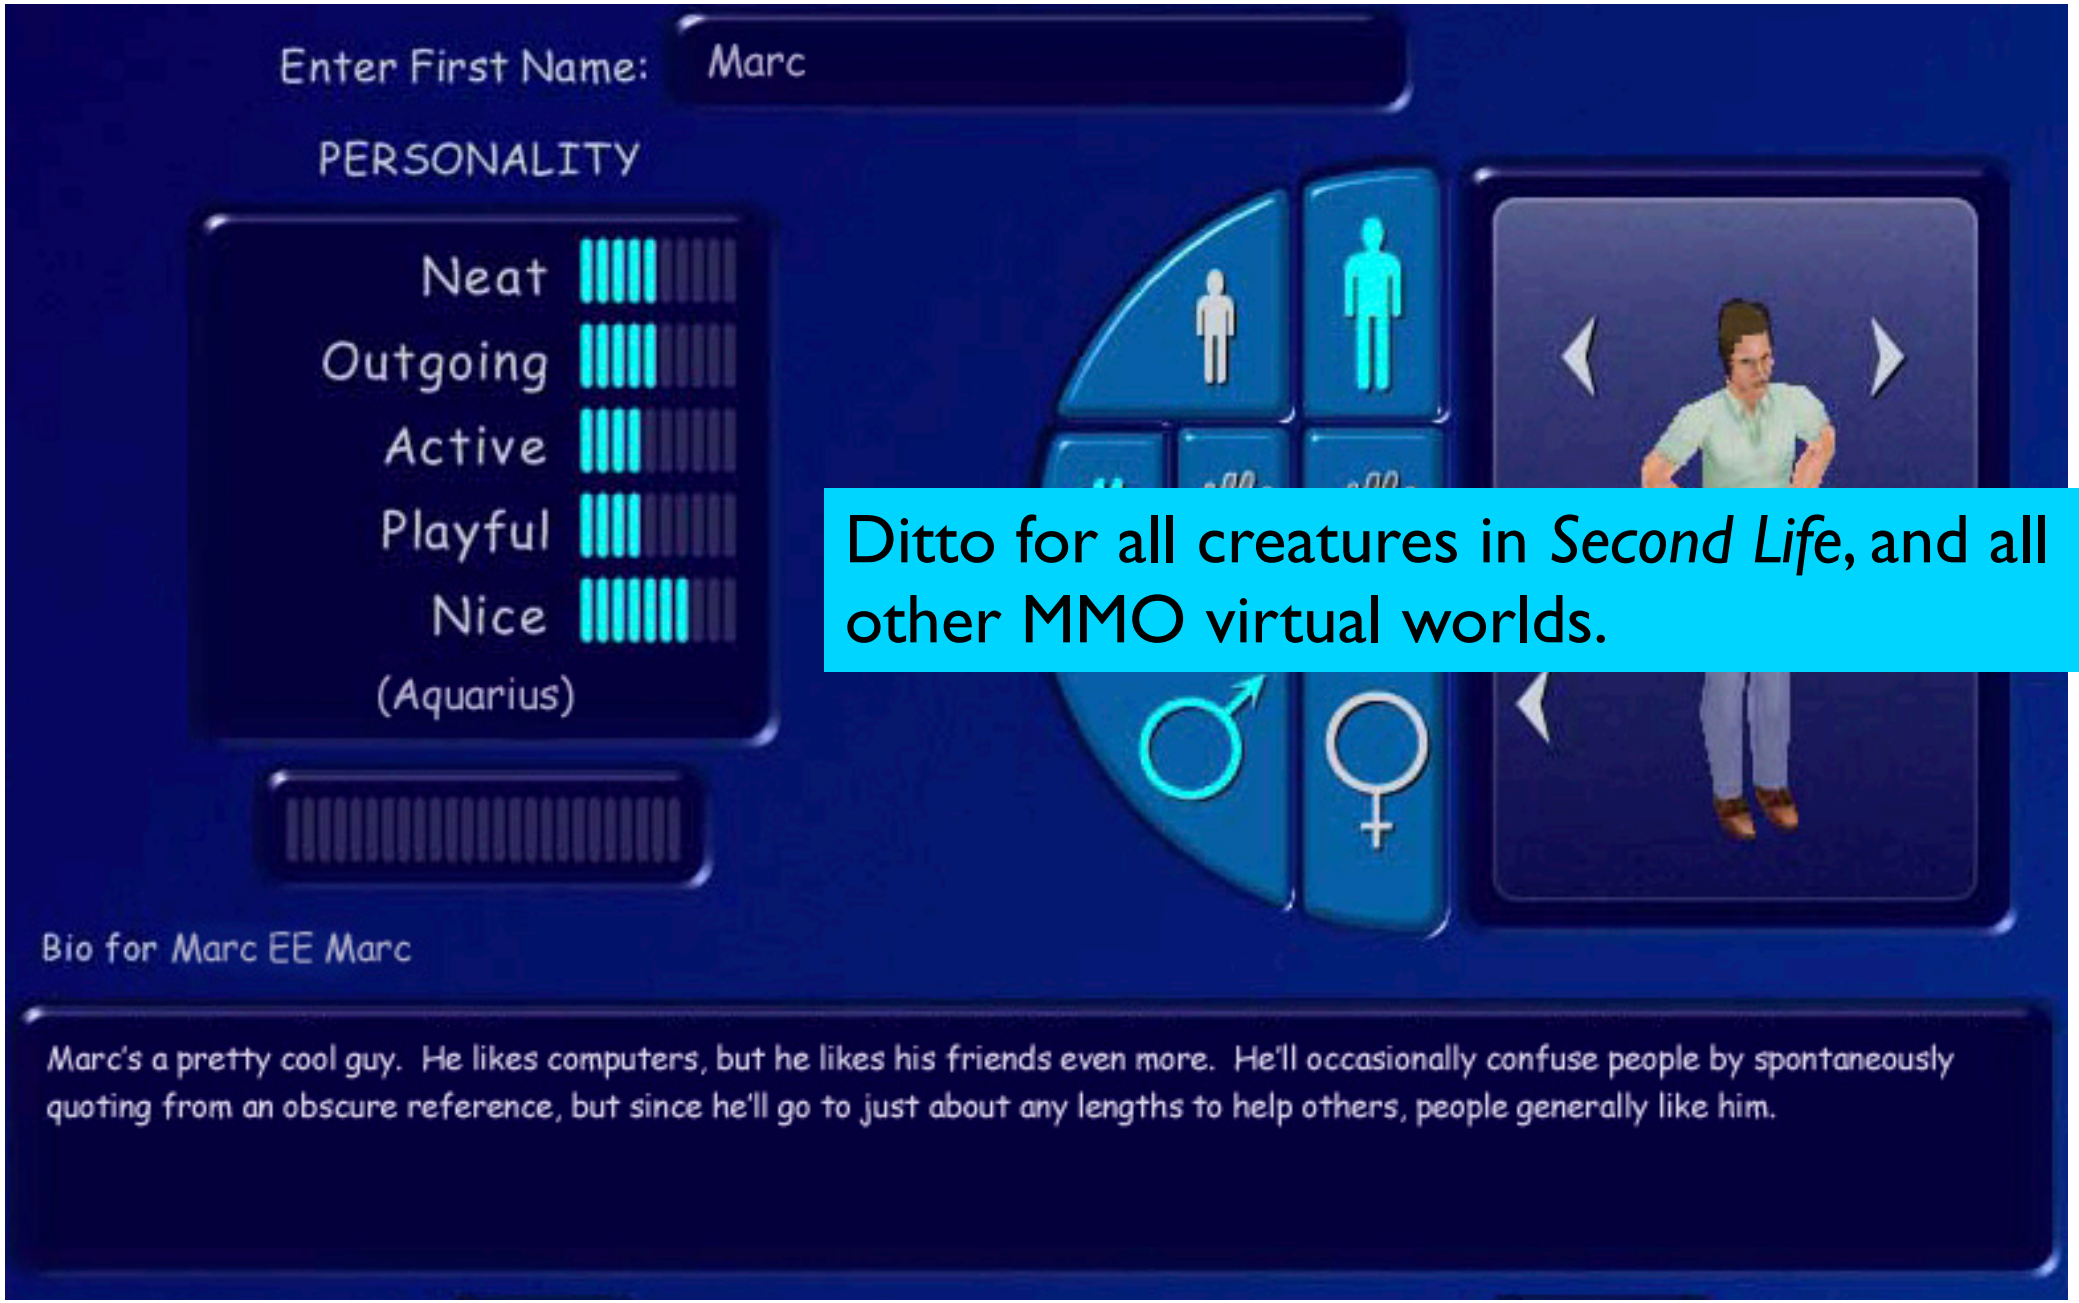

Ditto for all creatures in *Second Life*, and all other MMO virtual worlds.

Mere FSA

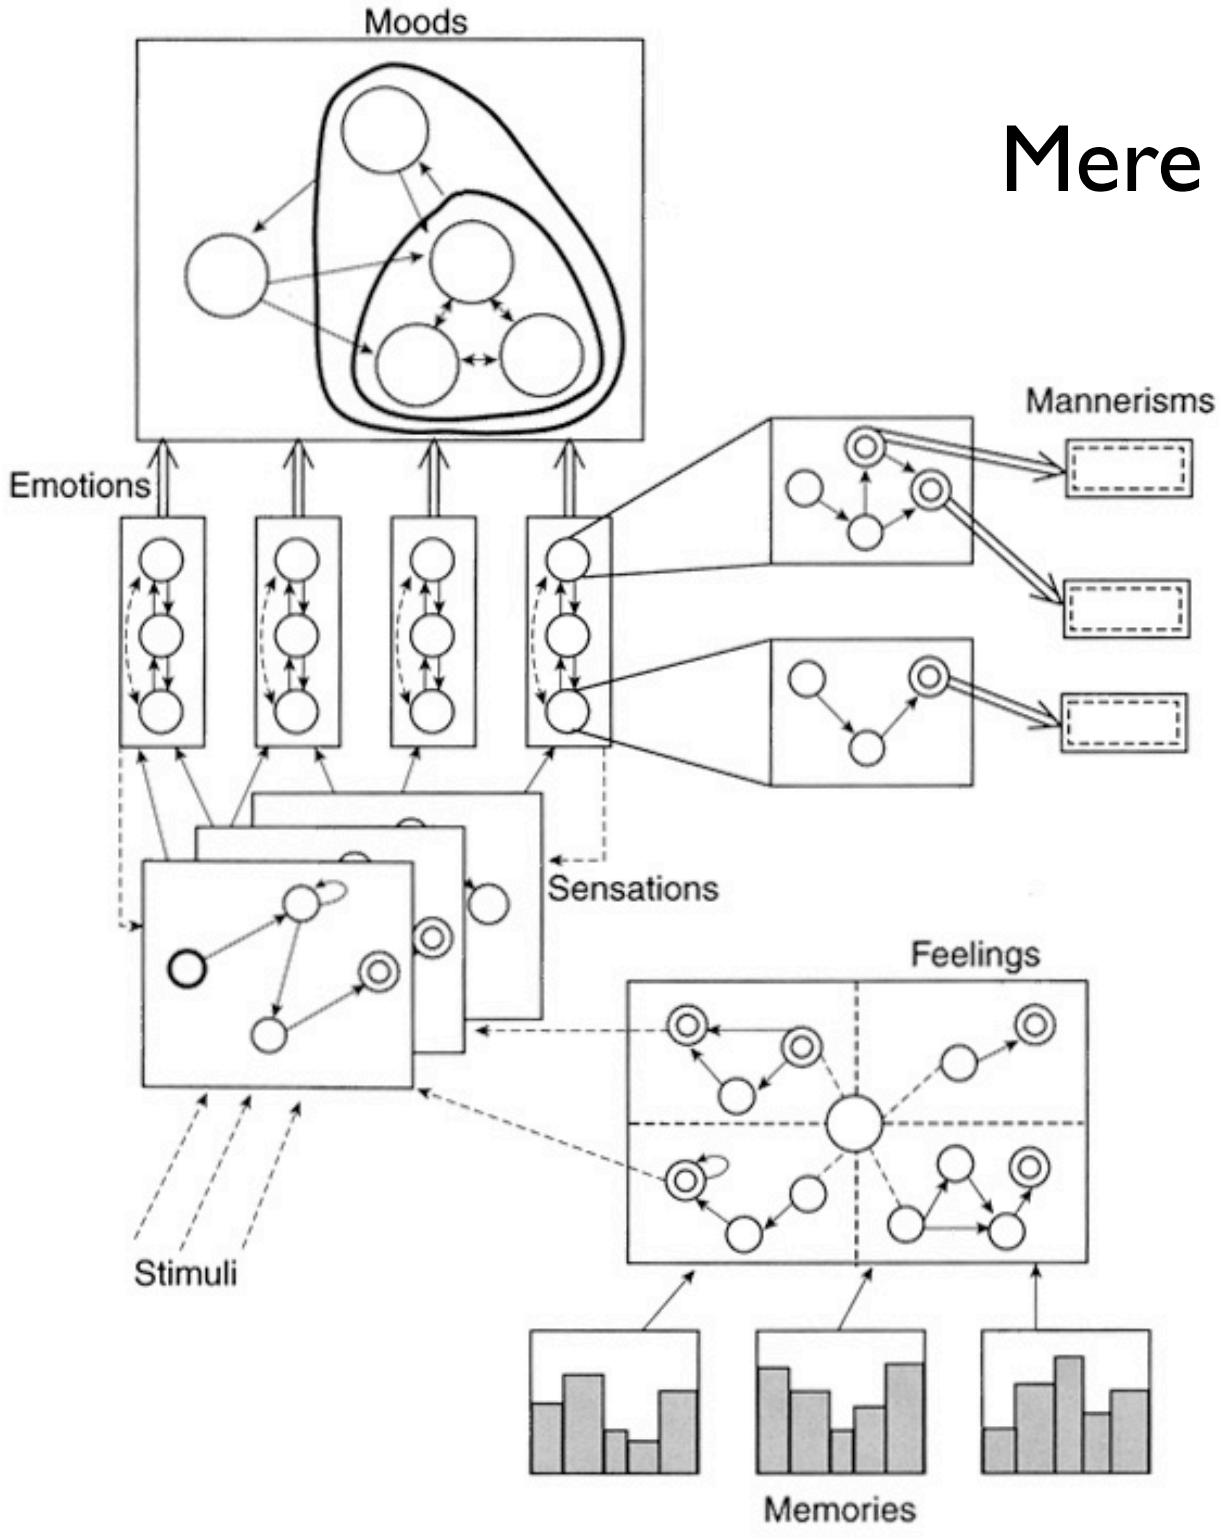
Moods
Emotions
Mannerisms
Sensations
Feelings
Stimuli
Memories

Cognitively Robust Synthetic Characters

```
SNARK-USER 14 >
(in-immature-scenario
  (prove '(t-retrieve subject
                      teddybear
                      ?c)
         :answer '(looks-in ?c)))
```

```
(Refutation
(Row 1
   (or (not (person ?x)) (not (object ?
y)) (not (container ?z)) (not (in ?y ?
z)) (bel-in ?x ?y ?z))
   assertion)
(Row 2
   (or (not (person ?x))
       (not (container ?y))
       (not (object ?z))
       (not (w-retrieve ?x ?z))
       (not (bel-in ?x ?z ?y))
       (t-retrieve ?x ?z ?y))
   assertion)
(Row 4
   (person subject)
   assertion)
(Row 6
   (container c2)
   assertion)
(Row 7
   (object teddybear)
   assertion)
(Row 8
   (in teddybear c2)
   assertion)
(Row 9
   (w-retrieve subject teddybear)
   assertion)
(Row 10
   (not (t-retrieve subject teddybear ?
x))
   negated_conjecture
   Answer (looks-in ?x))
(Row 11
   (or (not (person ?x)) (bel-in ?x
teddybear c2))
   (rewrite (resolve 1 8) 6 7))
(Row 25
   (bel-in subject teddybear c2)
   (resolve 11 4))
(Row 28
   (t-retrieve subject teddybear c2)
   (rewrite (resolve 2 25) 9 7 6 4))
(Row 30
   false
   (resolve 10 28)
   Answer (looks-in c2)))

:PROOF-FOUND
```

```
SNARK-USER 15 > (answer t)
(LOOKS-IN C2)
```

```
SNARK-USER 12 >
(in-mature-scenario
  (prove '(t-retrieve subject
                      teddybear
                      ?c)
         :answer '(looks-in ?c)))
```

```
(Refutation
(Row 1
   (or (not (person ?x))
       (not (container ?y))
       (not (object ?z))
       (not (w-retrieve ?x ?z))
       (not (bel-in ?x ?z ?y))
       (t-retrieve ?x ?z ?y))
   assertion)
(Row 2
   (or (not (person ?x)) (not (object ?
y)) (not (container ?z)) (not (p-in ?
x ?y ?z)) (bel-in ?x ?y ?z))
   assertion)
(Row 4
   (person subject)
   assertion)
(Row 5
   (container c1)
   assertion)
(Row 7
   (object teddybear)
   assertion)
(Row 8
   (p-in subject teddybear c1)
   assertion)
(Row 9
   (w-retrieve subject teddybear)
   assertion)
(Row 10
   (not (t-retrieve subject teddybear ?
x))
   negated_conjecture
   Answer (looks-in ?x))
(Row 11
   (bel-in subject teddybear c1)
   (rewrite (resolve 2 8) 5 7 4))
(Row 25
   (t-retrieve subject teddybear c1)
   (rewrite (resolve 1 11) 9 7 5 4))
(Row 26
   false
   (resolve 10 25)
   Answer (looks-in c1))
)

:PROOF-FOUND
```

```
SNARK-USER 13 > (answer t)
(LOOKS-IN C1)
```

Please note:

Even dirt-simple beliefs like:

"I believe that Jones believes that the object is in the first of the two boxes."

exceed what can be represented in things like RDF.

(Things are the same on the visual side: standard schemes for representing and generating physical objects and environments do not include the declarative information that can be reasoned over in order to answer obvious queries.)

Standalone Cognitively Robust Virtual Characters in *Second Life*

Logically Controlled English

- Logically controlled English is a subset of everyday English that has limits on what can be stated and rules on how to state it.
- This form of controlled English is then translated to a logical representation.

Controlled English to Logic Translation (CELT)

- CELT is a parser for controlled English that translates the English to logic.
- It can be much more powerful than Attempto Controlled English (ACE) because it leverages the Suggested Upper Merged Ontology (SUMO), which incorporates useful knowledge such as geometric axioms.

Second Life and CELT

- Example controlled English:

  "Micah puts the teddy bear in the box."

- CELT's translation to first order logic:

```
(exists
  (?box ?event ?teddy_bear)
  (and
    (instance Micah Human)
    (attribute ?box Box)
    (agent ?event Micah)
    (instance ?event Putting)
    (instance ?teddy_bear Artifact)
    (destination ?event ?box)
    (patient ?event ?teddy_bear)))
```

Analysis of Intelligence Reports using Slate ...

Iraq DIA/CIA BW Report Analyzed into Slate

Fermentor. Exhaust gas compressor.

Interior view of fermentor, media tank, water supply tanks, and gas cylinders connected by pipes.

such equipmenta fermentor, refrigeration, and a gas capture systemand agree with the experts that BW agent production is the only consistent, logical purpose for these vehicles.

- The capability of the system to capture and compress exhaust gases produced during fermentation is not required for legitimate biological processes and strongly indicates attempts to conceal production activity.
- The presence of caustic in the fermentor combined with the recent painting of the plant may indicate an attempt to decontaminate and conceal the plants purpose.
- Finally, the data plate on the fermentor indicates that this system was manufactured in 2002 and yet it was not declared to the United Nations, as required by Security Council Resolutions.

Some coalition analysts assess that the trailer found in late April could be used for bioproduction but believe it may be a newer prototype because the layout is not entirely identical to what the source described.

A *New York Times* article on 13 May 2003 reported that an agricultural expert suggests the trailers might have been intended to produce biopesticides near agricultural areas in order to avoid degradation problems. The same article also reported that a former weapons inspector suggests that the trailers may be chemical-processing units intended to refurbish Iraqs antiaircraft missiles.

- Biopesticide production requires the same equipment and technology used for BW agent production; however, the off-gas collection system and the size of

Manufacturer's data plate on the fermentor.

the equipment are unnecessary for biopesticide production. There is no need to produce biopesticides near the point of use because biopesticides do not degrade as quickly as most BW agents and would be more economically produced at a large fixed facility. In addition, the color of the trailer found in mid-April is indicative of military rather than civilian use.

- Our missile experts have no explanation for how such a trailer could function to refurbish antiaircraft missiles and judge that such a use is unlikely based on the scale, configuration, and assessed function of the equipment.
- The experts cited in the editorial are not on the scene and probably do not have complete access to information about the trailers.

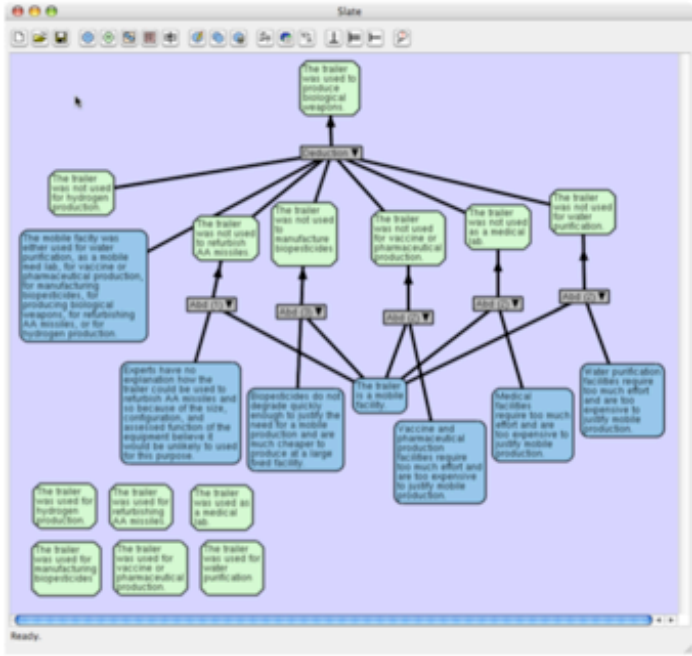

Hydrogen Production Cover Story

Senior Iraqi officials of the al-Kindi Research, Testing, Development, and Engineering facility in Mosul were shown pictures of the mobile production trailers, and they claimed that the trailers were used to chemically produce hydrogen for artillery weather balloons. Hydrogen production would be a plausible cover story for the mobile production units.

- The Iraqis have used sophisticated denial and deception methods that include the use of cover stories that are designed to work. Some of the features of the trailera gas collection system and the presence of causticare consistent with both bioproduction and hydrogen production.

The plants design possibly could be used to produce hydrogen using a chemical reaction, but it would be inefficient. The capacity of this trailer is larger than typical units for hydrogen production for weather balloons. Compact, transportable hydrogen generation systems are commercially available, safe, and reliable.

Sample Collection and Analysis

We continue to examine the trailer found in mid-April and are using advanced sample analysis techniques to determine whether BW agent is present, although we do not

The End