

CS 681 Fall 2008

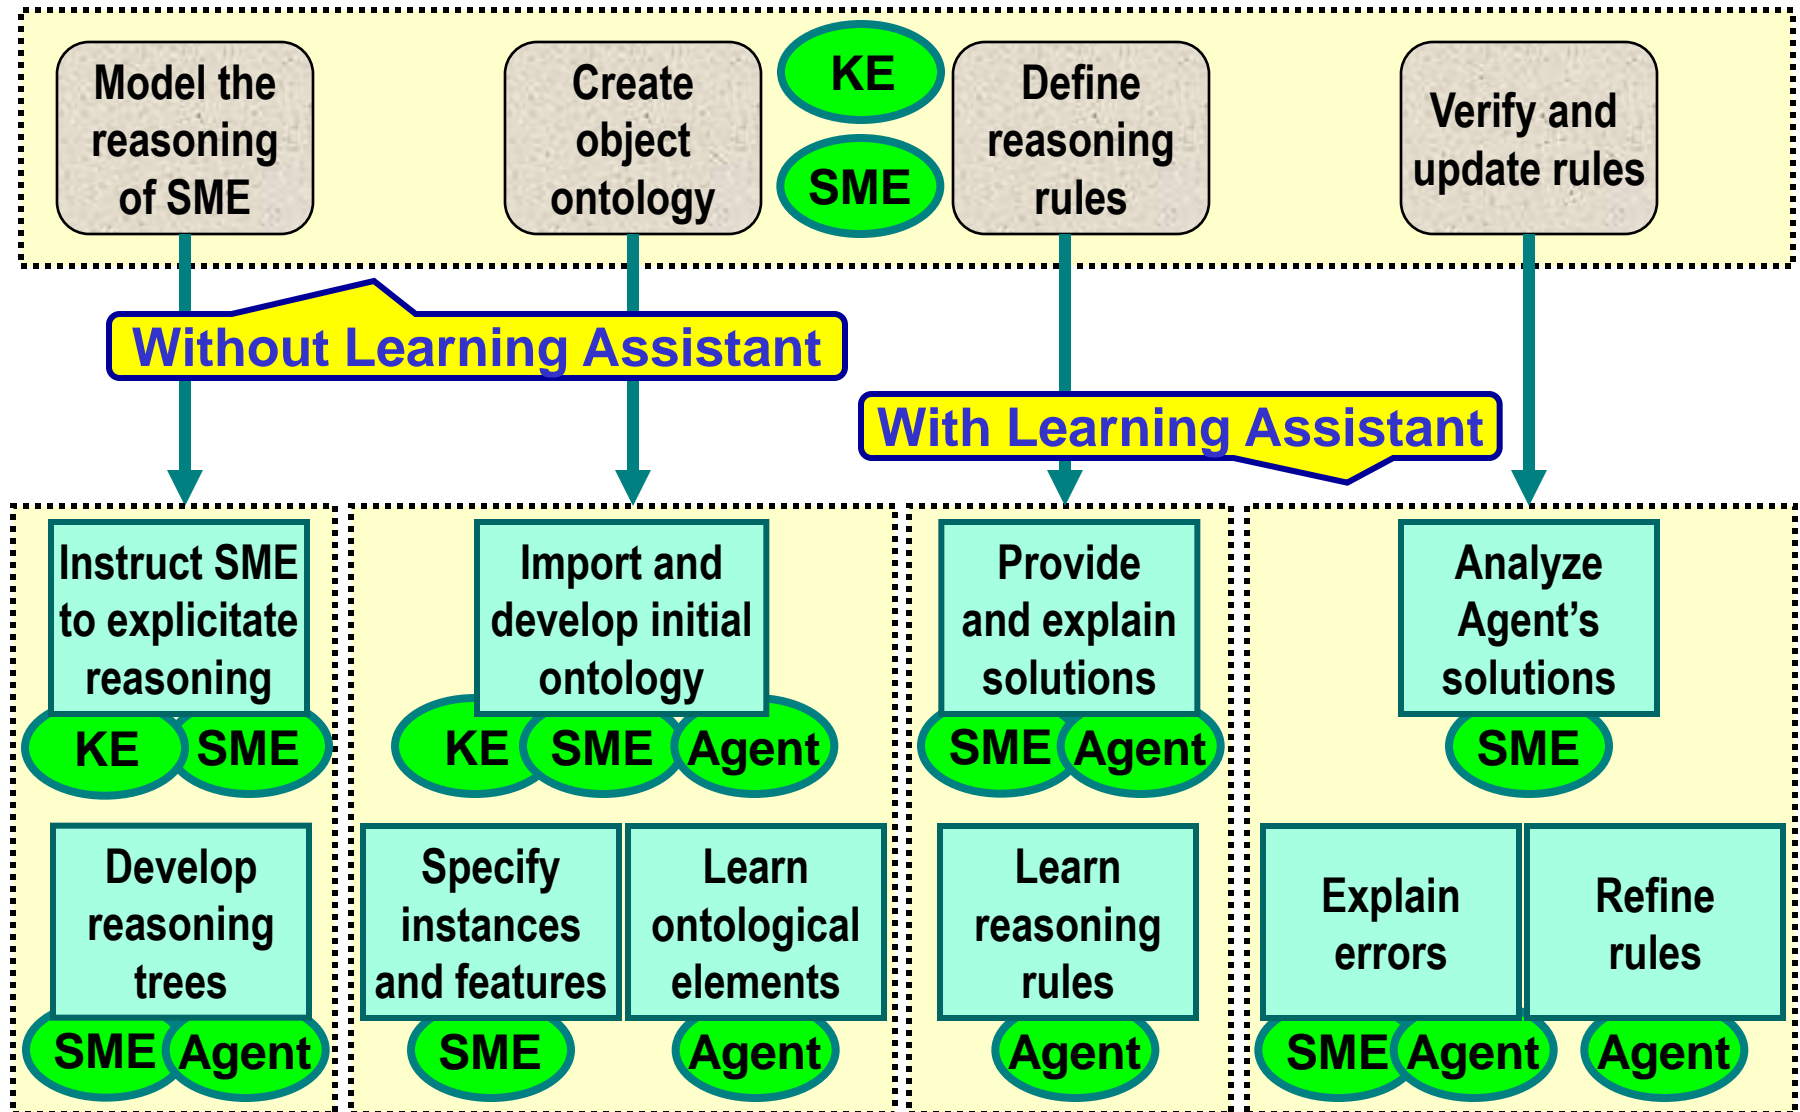
Designing Expert Systems

10. Design Principles for Learning Assistants

Gheorghe Tecuci
tecuci@gmu.edu
<http://lac.gmu.edu/>

**Learning Agents Center
and Computer Science Department
George Mason University**

Agent Development Approaches



Overview: Design Principles



Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning

Multistrategy Learning

Learning with an Evolving Representation Space

Mixed-Initiative Modeling, Learning and Problem Solving

Plausible Reasoning with Partially Learned Knowledge

User Tutoring in Problem Solving

Agent Architecture for Generality-Power Tradeoff

P1. Spiral Development with SMEs and End-Users

Intelligence analysis, Center of gravity determination,
Course of action critiquing, Emergency response
planning, Workaround reasoning, PhD advisor
selection, Teaching higher order thinking skills.

Army War College
Air War College
George Mason University

Development of
systematic approach to
expert problem solving

Working closely
with subject
matter experts to
model their
reasoning



DISCIPLE

Development of the Disciple
theory for agent teaching by
non-computer experts

Development and
application of
Disciple agents

Working closely
with end users to
receive crucial
and timely
feedback



Integration of many areas of Artificial Intelligence

Overview: Design Principles

Spiral Development with SMEs and End-Users



Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning

Multistrategy Learning

Learning with an Evolving Representation Space

Mixed-Initiative Modeling, Learning and Problem Solving

Plausible Reasoning with Partially Learned Knowledge

User Tutoring in Problem Solving

Agent Architecture for Generality-Power Tradeoff

P2. Problem Solving Paradigm for User-Agent Collaboration

Develop a general problem solving paradigm which is:

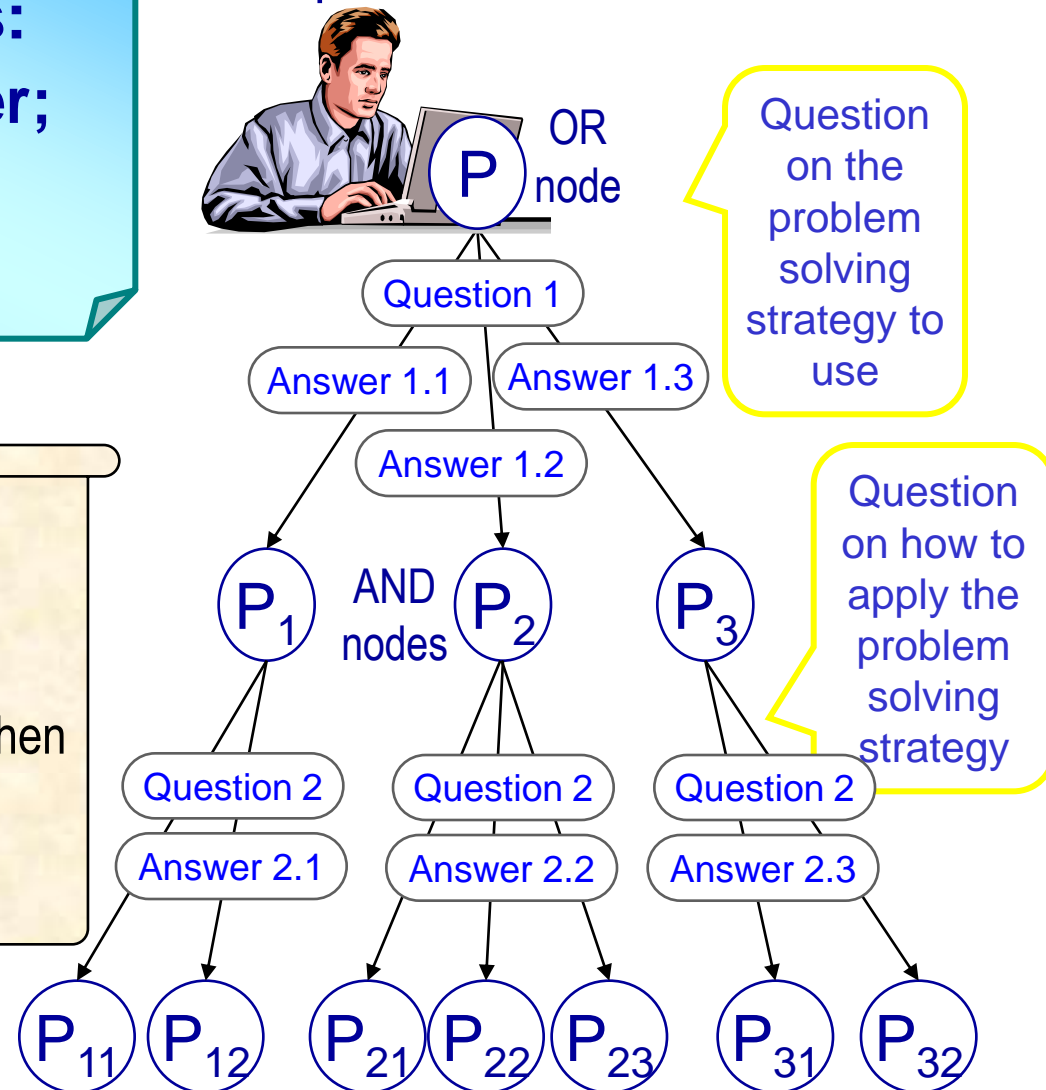
- natural for a human user;
- appropriate for an automated agent.

"I Keep Six Honest..."

I keep six honest serving-men
(They taught me all I knew);
Their names are What and Why and When
And How and Where and Who.

Rudyard Kipling

Problem Reduction Guided by
Introspective Questions and Answers



Solution Synthesis

“One of the most highly developed skills in contemporary Western civilization is dissection; the split-up of problems into their smallest possible components. We are good at it. So good, we often forget to put the pieces back together again.”

Alvin Toffler, *Science and Change*, Forward to Ilya Prigogine and Isabelle Stengers *“Order Out of Chaos: Man’s New Dialogue with Nature”*, Bantam Books, 1984.



Alvin Toffler

Problem-Reduction/Solution Synthesis Paradigm

The reduction representation of a class of problems is a quadruple (P, S, RO, OS) where:

P - class of problems;


S - solutions of problems;

RO - reduction operators, each reducing a problem to sub-problems and/or solutions,

SO - synthesis operators, each synthesizing the solution of a problem from the solutions of its sub-problems.

A problem P_1 is solved by:

- successively reducing it to simpler problems through the application of the reduction operators;
- finding the solutions of the simplest problems;
- successively combining these solutions through the application of synthesis operators until the solution of the initial problem is obtained.



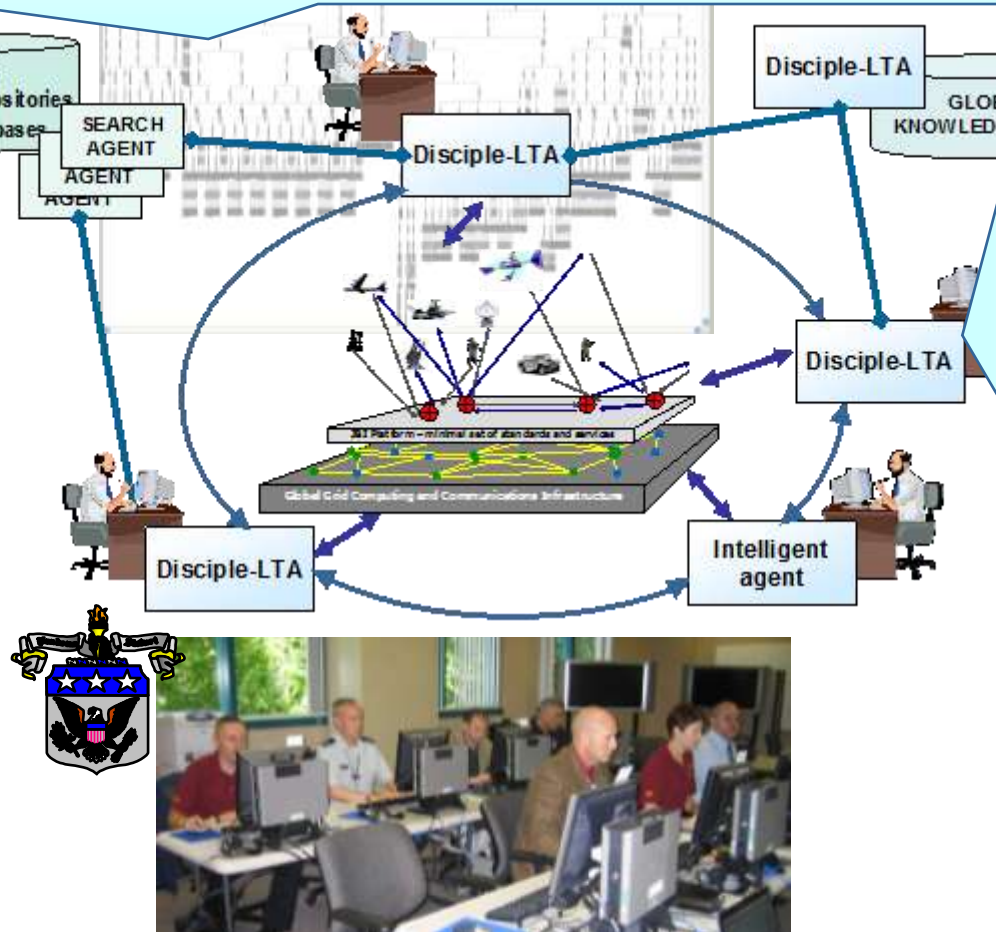
Sample Agent: Analyst's Cognitive Assistant

Analytic Assistance

Empowers the analysts through mixed-initiative reasoning for hypotheses analysis, collaboration with other analysts and experts, and sharing of information.

Learning

Rapid acquisition and maintenance of subject matter expertise in intelligence analysis which currently takes years to establish, is lost when experts separate from service, and is costly to replace.

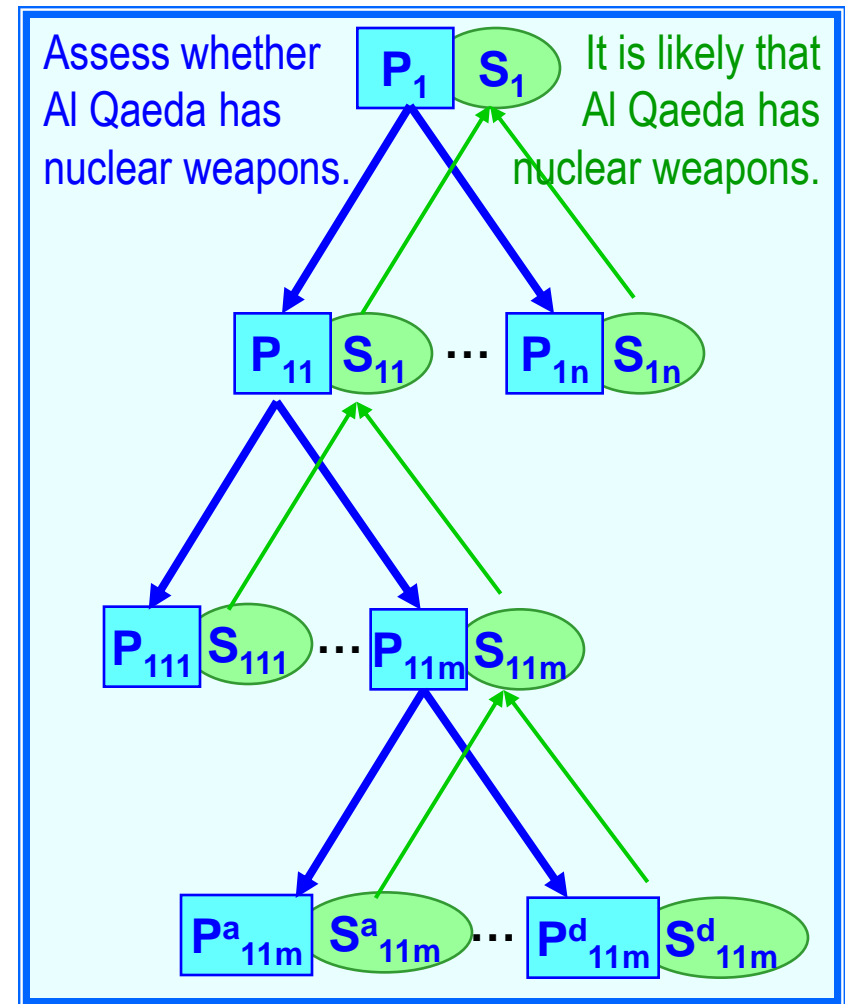


Tutoring

Helps new intelligence analysts learn the reasoning processes involved in making intelligence judgments and solving intelligence analysis problems.

Hypothesis Analysis through Problem Reduction

- 1) A complex hypothesis analysis problem is successively reduced to simpler problems that either have known solutions or can be solved through evidence analysis.
- 2) Potentially relevant pieces of evidence for the unsolved problems are identified.
- 3) The pieces of evidence are analyzed to obtain solutions for the unsolved problems.
- 4) The solutions of the simplest problems are successively combined to obtain the solution of the initial problem.



Remote	Unlikely	Even chance	Probably, Likely	Almost certainly
National Intelligence Council's standard estimative language				

Demo: Analytic Assistance

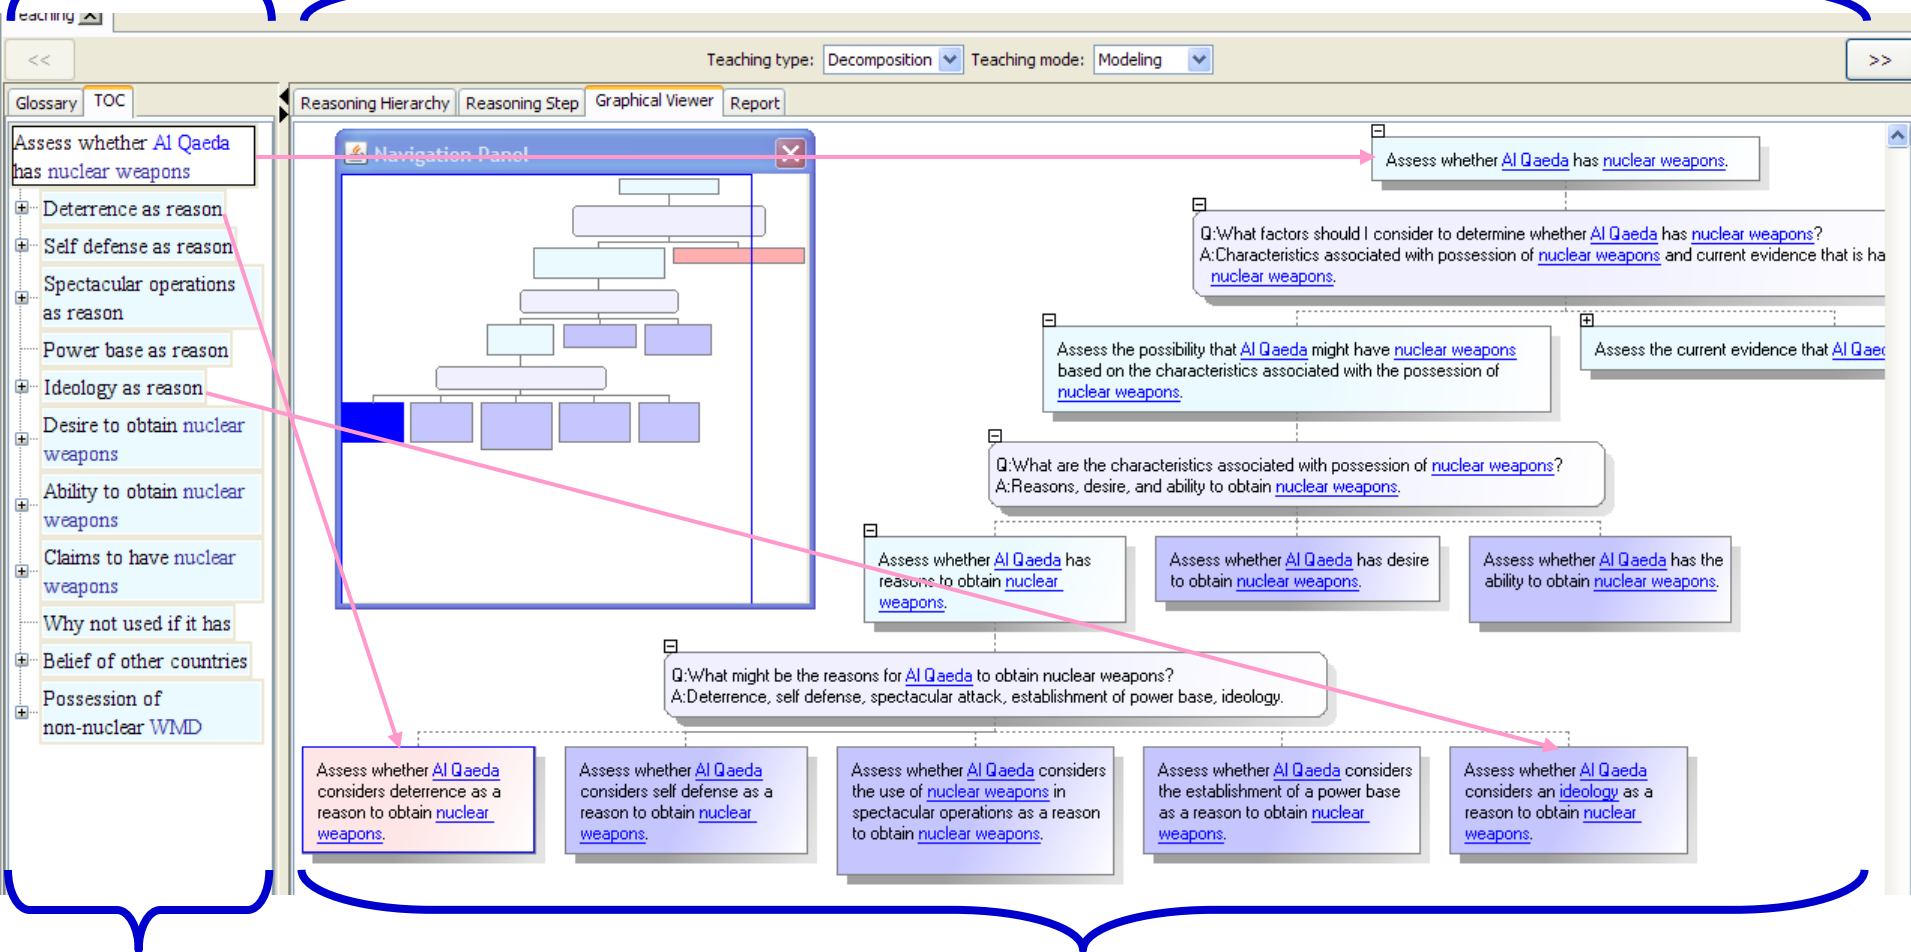


Disciple-LTA Demo: Analytic Assistance



Problem Reduction: Abstract and Detailed Reasoning

Main problems Reduction of a main problem to its main subproblems



Abstract tree

Detailed tree

Problem Reduction and Solution Synthesis

System Ontology Rules Scenario Reasoning Tutoring Help

repository\PNW-cwv-070904-00\AI Qaida

Mixed-Initiative Reasoner X

Reasoning type: Both Reasoning mode: Modeling Plausibility: medium

- Glossary TOC
- Assess whether Al Qaeda has nuclear weapons
- ⊕ Deterrence as a reason
 - ⊖ Self defense as a reason
 - ⊖ Favoring evidence
 - ⊕ EVD-Reuters-01-01c
 - ⊖ EVD-Dawn-Mir01-01c
 - Relevance
 - ⊖ Believability
 - ⊕ Reporter Hamid Mir
 - ⊖ Source Osama bin Laden
 - ⊕ Competence
 - ⊖ Credibility
 - Veracity
 - Objectivity
 - Observational sensitivity
 - ⊖ Disfavoring evidence
 - ⊕ Spectacular operations as a reason
 - ⊖ Establishment of a power base as a reason
 - ⊕ Ideology as a reason
 - ⊕ Desire to obtain nuclear weapons
 - ⊕ Ability to obtain nuclear weapons

Reasoning Hierarchy Reasoning Step Graphical Viewer Report

Assess the credibility of Osama bin Laden as the source of EVD-Dawn-Mir01-01c.

The credibility of Osama bin Laden as the source of EVD-Dawn-Mir01-01c is an even chance.

Q:What facts

A:The veracity, objectivity, and observational sensitivity of Osama bin Laden because EVD-Dawn-Mir01-01c is testimonial evidence based upon direct observation.

Assess the veracity of Osama bin Laden with respect to the information provided in EVD-Dawn-Mir01-01c.

The veracity of Osama bin Laden with respect to the information provided in EVD-Dawn-Mir01-01c is an even chance.

Assess the objectivity of Osama bin Laden with respect to the information provided in EVD-Dawn-Mir01-01c.

The objectivity of Osama bin Laden with respect to the information provided in EVD-Dawn-Mir01-01c is almost certain.

Assess the observational sensitivity of Osama bin Laden with respect to the information provided in EVD-Dawn-Mir01-01c.

The observational sensitivity of Osama bin Laden with respect to the information provided in EVD-Dawn-Mir01-01c is almost certain.

Detailed evidence and source analysis

EVD-Dawn-Mir-01-01c

The following interpretation:

Al Qaeda has nuclear weapons and may use them to defend itself of the following statement by Osama bin Laden:

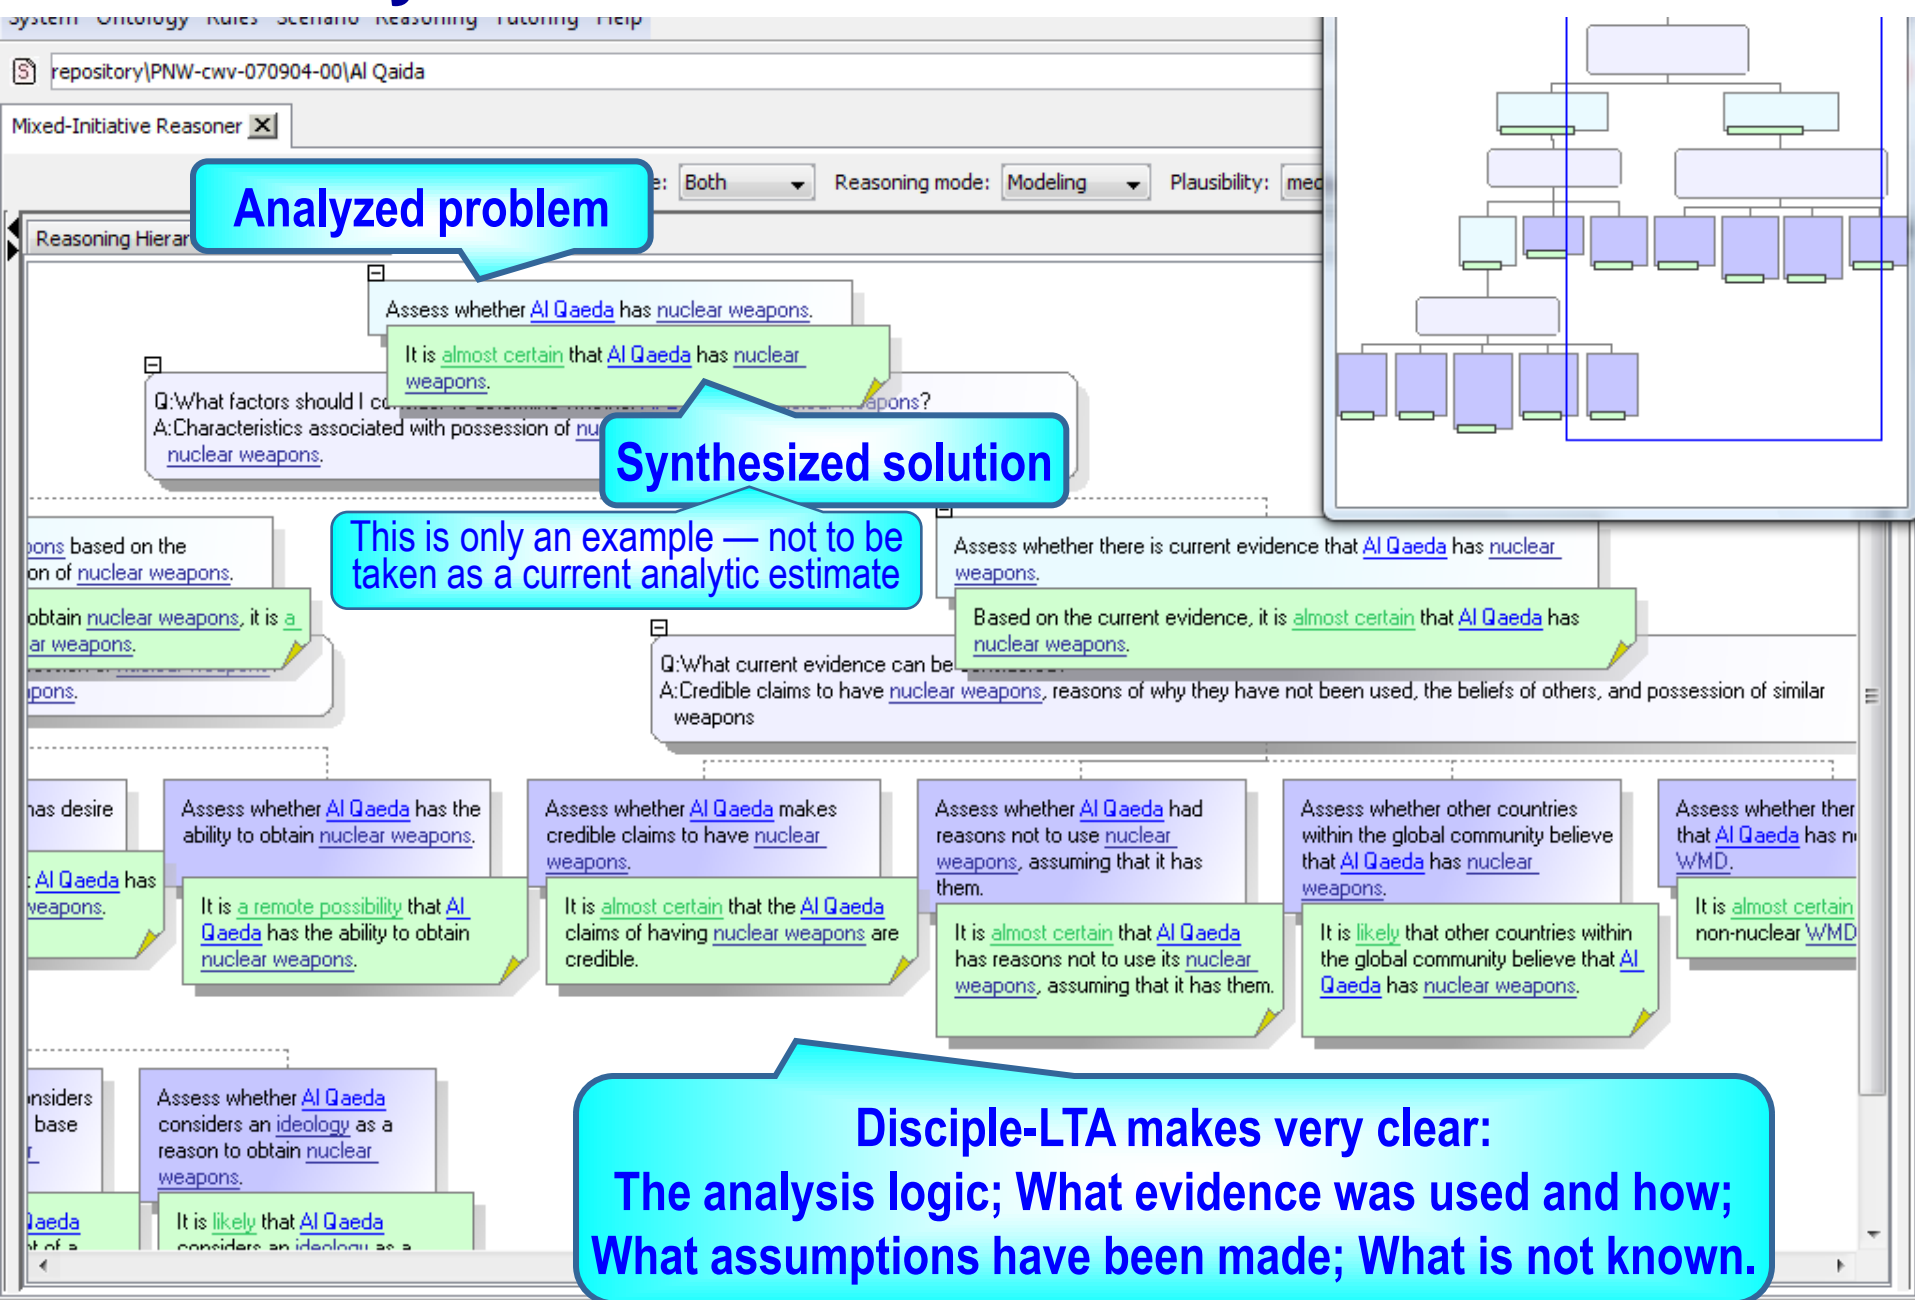
We [Al Qaeda] have chemical and nuclear weapons [...] We ourselves are the target of killings, destruction and atrocities. We are only defending ourselves. This is defensive Jihad. We want to defend our people and our land. That is why I say that if we dont get security, the Americans, too would not get security.

from the document EVD-Dawn-Mir01:

Mir, H. (2001, November 10). Osama claims he has nukes: If US uses N-arms it will get same response. Dawn - Pakistan English Newspaper

<http://www.dawn.com/2001/11/10/top1.htm>

Solution Synthesis: Illustration



Assumptions-based Analysis

System Ontology Rules Scenario Reasoning Tutoring Help

2. Revised, assumption-based, solution

Disciple-LTA allows for: Assumptions checking; Rapid updating of large analysis trees based on new intelligence data and assumptions.

Reasoning type: Both

Graphical Viewer Report

Assess whether Al Qaeda has nuclear weapons.

It is likely that Al Qaeda has nuclear weapons.

Q: What factors should I consider to determine whether Al Qaeda has nuclear weapons?
A: Characteristics associated with possession of nuclear weapons and current evidence that it has nuclear weapons.

Assess whether there is current evidence of Al Qaeda having nuclear weapons.

Based on the current evidence, it is unlikely that Al Qaeda has nuclear weapons.

Q: What current evidence can be used to determine whether Al Qaeda has nuclear weapons?
A: Credible claims to have nuclear weapons, reasons of why they have nuclear weapons.

Assess whether Al Qaeda has the ability to obtain nuclear weapons.

It is a remote possibility that Al Qaeda has the ability to obtain nuclear weapons.

Assess whether Al Qaeda makes credible claims to have nuclear weapons.

It is almost certain that the Al Qaeda claims of having nuclear weapons are credible.

Assess whether Al Qaeda had reasons not to use nuclear weapons, assuming that it has them.

It is unlikely that Al Qaeda has reasons not to use its nuclear weapons, assuming that it has them.

Assess whether Al Qaeda considers an ideology as a reason to obtain nuclear weapons.

It is likely that Al Qaeda considers an ideology as a reason to obtain nuclear weapons.

The following assumption:
It is unlikely that Al Qaeda has reasons not to use its nuclear weapons, assuming that it has them.
is challenged by the system's solution:
It is almost certain that Al Qaeda has reasons not to use its nuclear weapons, assuming that it has them.

Assumptions

☒ Enabled S: It is unlikely that Al Qaeda has reasons not to use its nuclear weapons, assuming that it has them.

☐ X ☐ J: I think that they would use the weapons.

Modify Save Cancel

Over 1700 reasoning nodes

1. Analyst's assumption challenged by Disciple-LTA

New

Overview: Design Principles

Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration



Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning

Multistrategy Learning

Learning with an Evolving Representation Space

Mixed-Initiative Modeling, Learning and Problem Solving

Plausible Reasoning with Partially Learned Knowledge

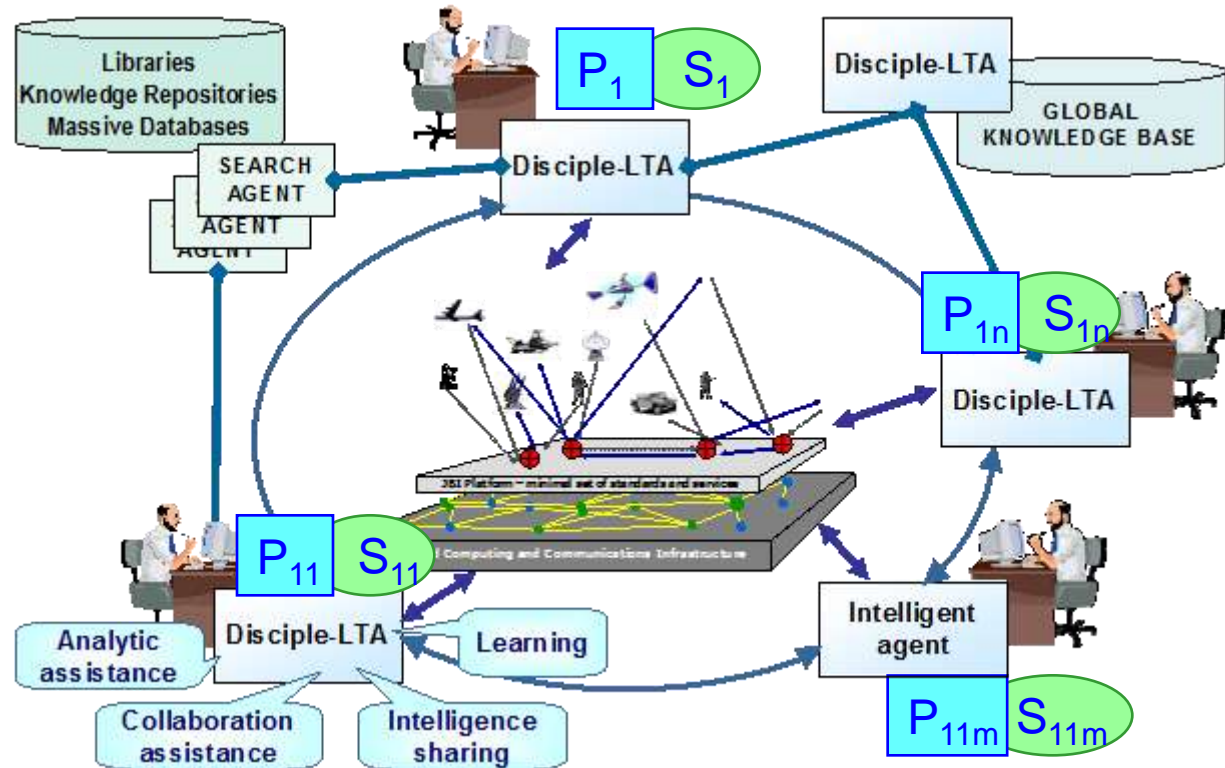

User Tutoring in Problem Solving

Agent Architecture for Generality-Power Tradeoff

P3. Multi-Agent and Multi-Domain Problem Solving

Develop a general problem solving paradigm that facilitates:

- collaboration between users assisted by their agents;
- solving problems requiring multi-domain expertise.



Service-Oriented Disciple Systems

Problem: Assess whether Al Qaeda has nuclear weapons.

Perform analysis

1

Subproblem: Assess whether other countries within the global community believe that Al Qaeda has nuclear weapons.

Subproblem: Assess whether Al Qaeda makes credible claims to have nuclear weapons.

Ask broker for solutions

2

Receive solutions

5

Integrate Solutions

6

Disciple-LTA Client

7

Upload reasoning tree in Catalyst



Hypothesis Analysis Broker

Check UDDI for registered competence

3



Solicit / receive solutions

4

Disciple-LTA Clients



Disciple LTA Servers



Hypothesis Analysis Web Service

Subproblem: Assess whether other countries within the global community believe that Al Qaeda has nuclear weapons.
Solution: It is likely that other countries within the global community believe that Al Qaeda has nuclear weapons.

Subproblem: Assess whether Al Qaeda makes credible claims to have nuclear weapons.
Solution: It is almost certain that the Al Qaeda claims of having nuclear weapons are credible.

Overview: Design Principles

Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving



Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning

Multistrategy Learning

Learning with an Evolving Representation Space

Mixed-Initiative Modeling, Learning and Problem Solving

Plausible Reasoning with Partially Learned Knowledge

User Tutoring in Problem Solving

Agent Architecture for Generality-Power Tradeoff


P4. Knowledge Base Structuring for Knowledge Reuse

Structure the knowledge base into two parts:

- o its more general and reusable components;
- o its more specific components.

Disciple: Knowledge Base Structuring

- The object ontology which may be reused from existing knowledge repositories;
- The problem solving rules which are learned from the subject matter expert.



Knowledge Base = Object Ontology + Rules

The ontology is a hierarchical description of the domain objects.

type-based evidence

authoritative record evidence

testimonial evidence

testimonial evidence obtained at second hand

testimonial evidence as opinion based on analysis

testimonial evidence based upon direct observation

testimonial evidence about tangible evidence

chart evidence

videotape evidence

audiotape evidence

object evidence

image evidence

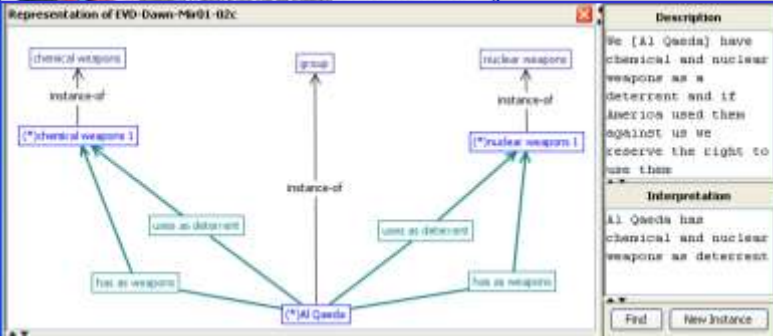
measurement evidence

document evidence

tangible evidence

Interpretation: Al Qaeda has chemical and nuclear weapons as deterrent.


Interpretation: Al Qaeda has chemical and nuclear weapons as deterrent.



Knowledge Base = Ontology + Rules

Rules specify
general
problem reduction
or solution
synthesis steps

Analysis Tree



Partially
learned
rule

IF: Assess whether there are states with nuclear weapons willing to sell nuclear weapons to ?O1.

Q:

Which is a nuclear state which is not an enemy of ?O1 and does not oppose the proliferation of nuclear weapons?

A:

?O2

MAIN CONDITION

Var	Lower Bound	Upper Bound
?O1	(terrorist group)	(actor)
?O2	(nuclear state)	(nuclear state)

EXCEPT WHEN CONDITION 1

Var	Lower Bound	Upper Bound
?O1	(terrorist group)	(actor)
?O2	(superpower, nuclear state)	(actor)

EXCEPT WHEN CONDITION 2

Var	Relationship	Var
?O2	perceives as enemy	?O1

EXCEPT WHEN CONDITION 2

Var	Lower Bound	Upper Bound
?SII	[medium - medium]	[very low - very high]
?O2	(nuclear state)	(actor)

Var	Relationship	Var
?O2	degree of opposition to nuclear weapons proliferation	?SII

THEN: Assess whether ?O2 is willing to sell nuclear weapons to ?O1.

To assess whether there are states that may be willing to sell nuclear weapons to an actor, **one has to consider each nuclear state** and assess whether that state may be willing to sell nuclear weapons to that actor, **except for** the case in which the nuclear state is an enemy of that actor and also **except for** the case when the nuclear state opposes the proliferation of nuclear weapons.

Overview: Design Principles

Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse



Integrated Teaching and Learning

Multistrategy Learning

Learning with an Evolving Representation Space

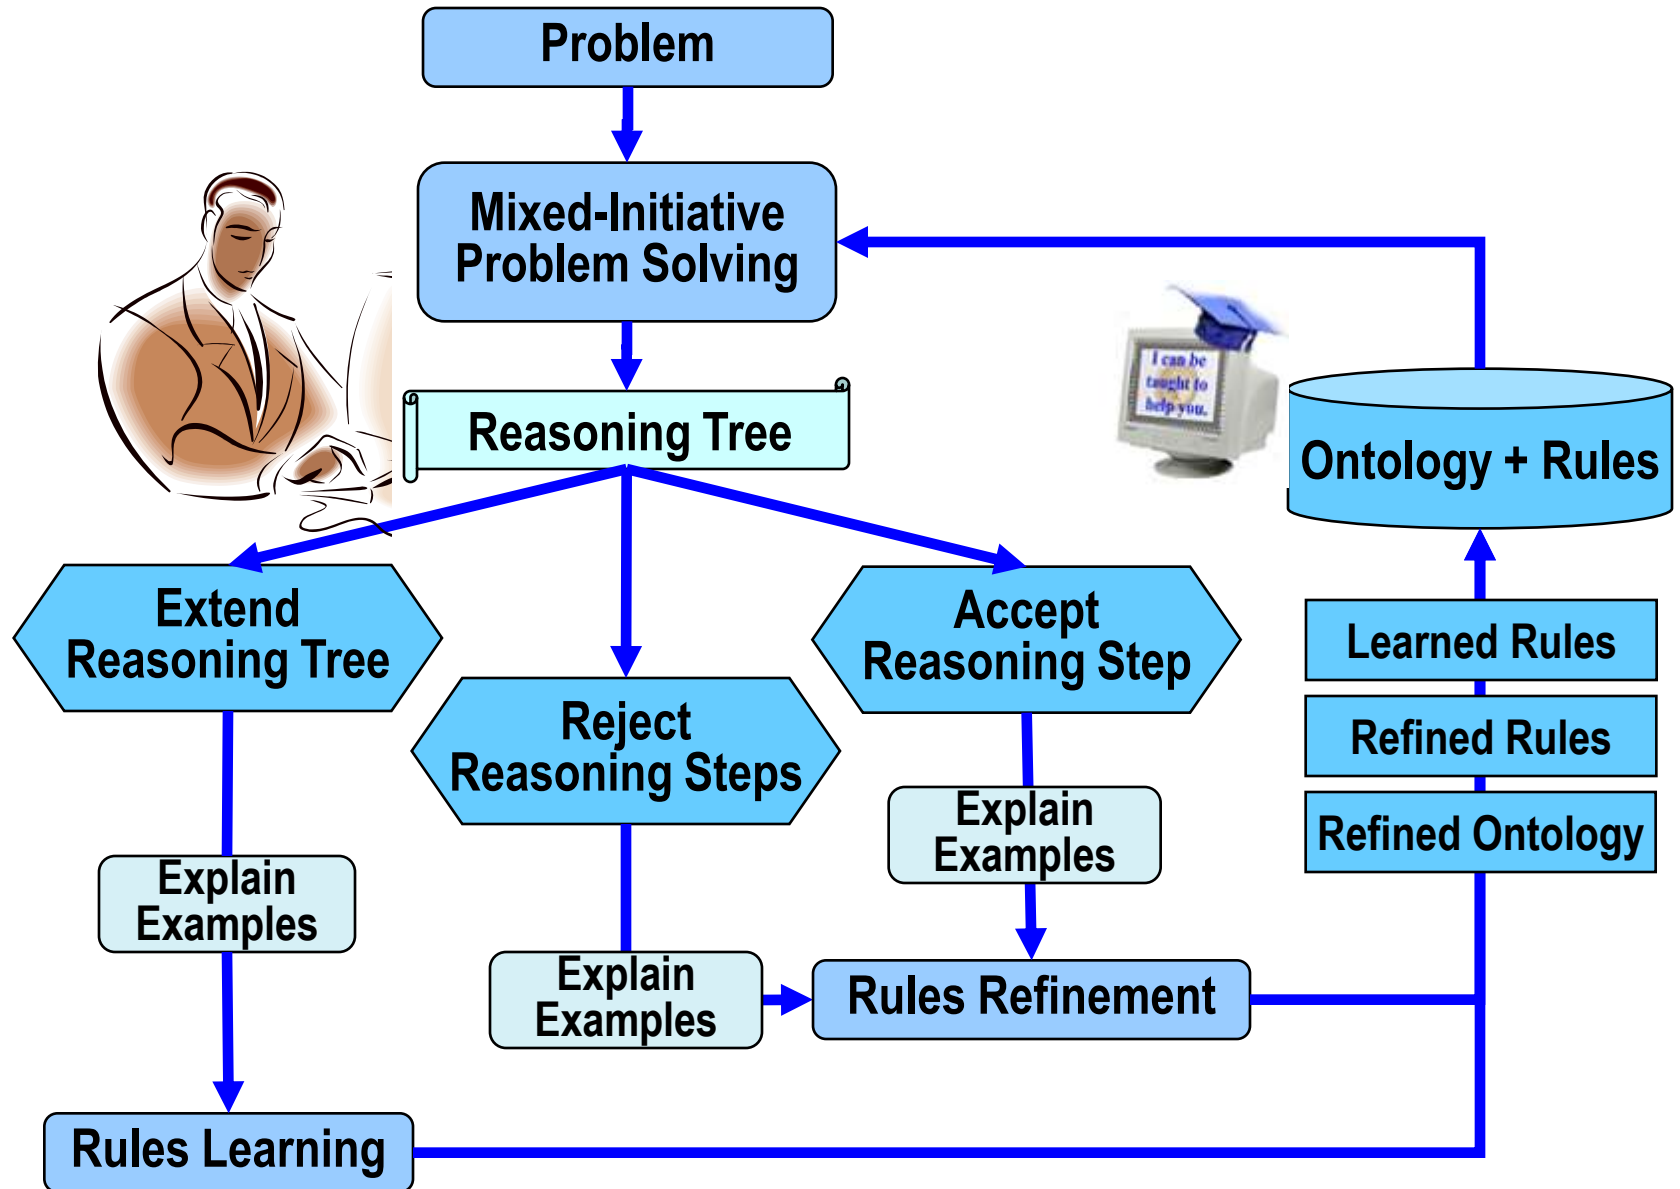
Mixed-Initiative Modeling, Learning and Problem Solving

Plausible Reasoning with Partially Learned Knowledge

User Tutoring in Problem Solving

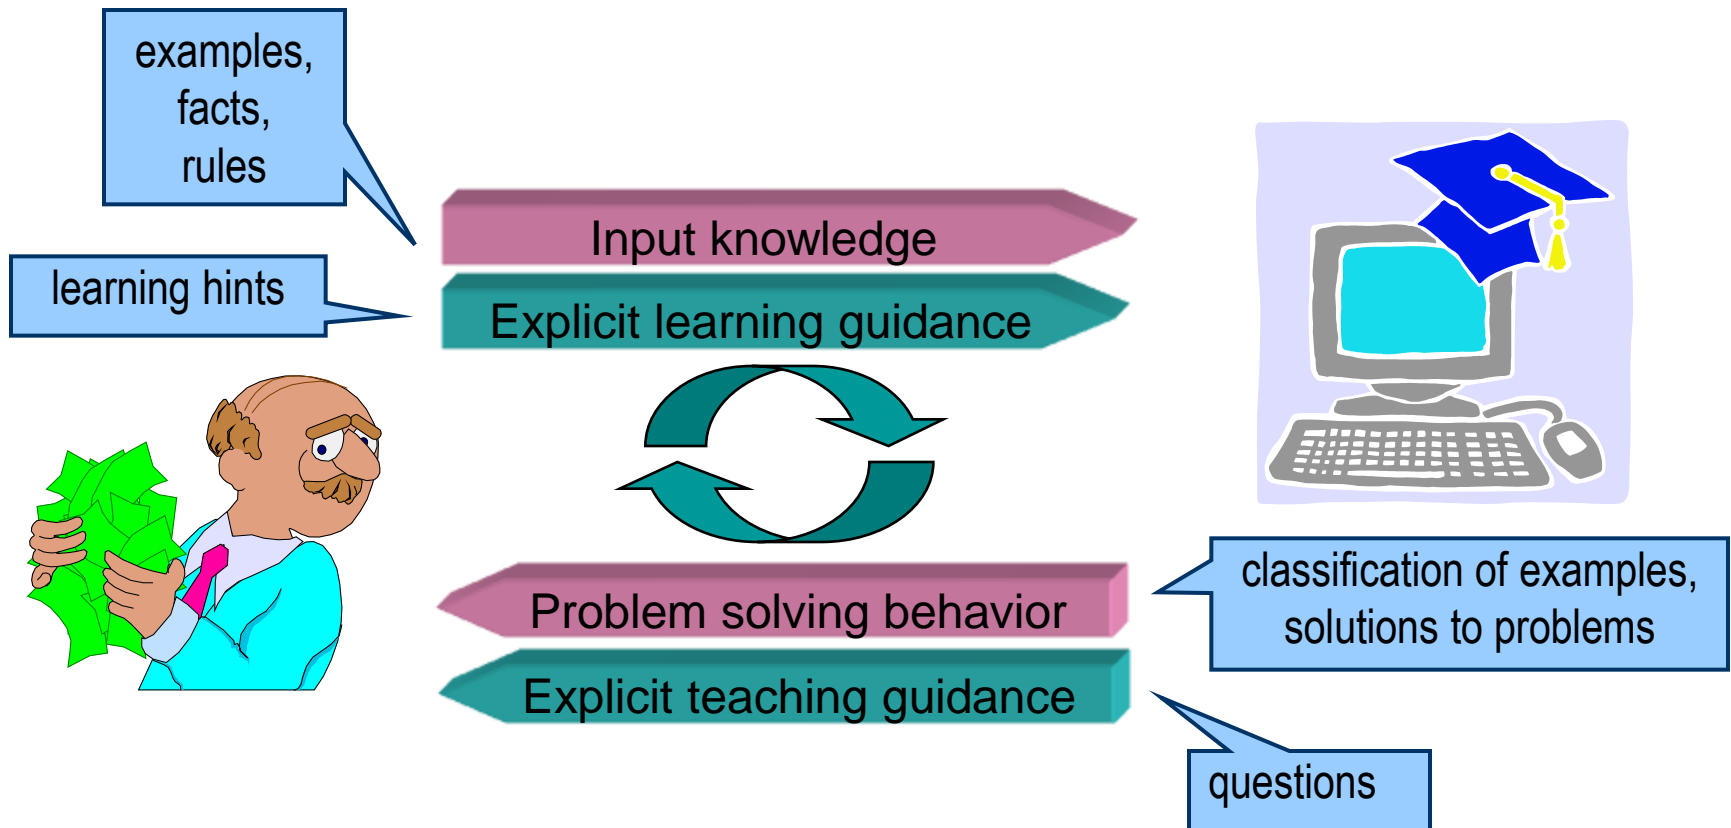
Agent Architecture for Generality-Power Tradeoff

Control of Modeling, Learning and Problem Solving

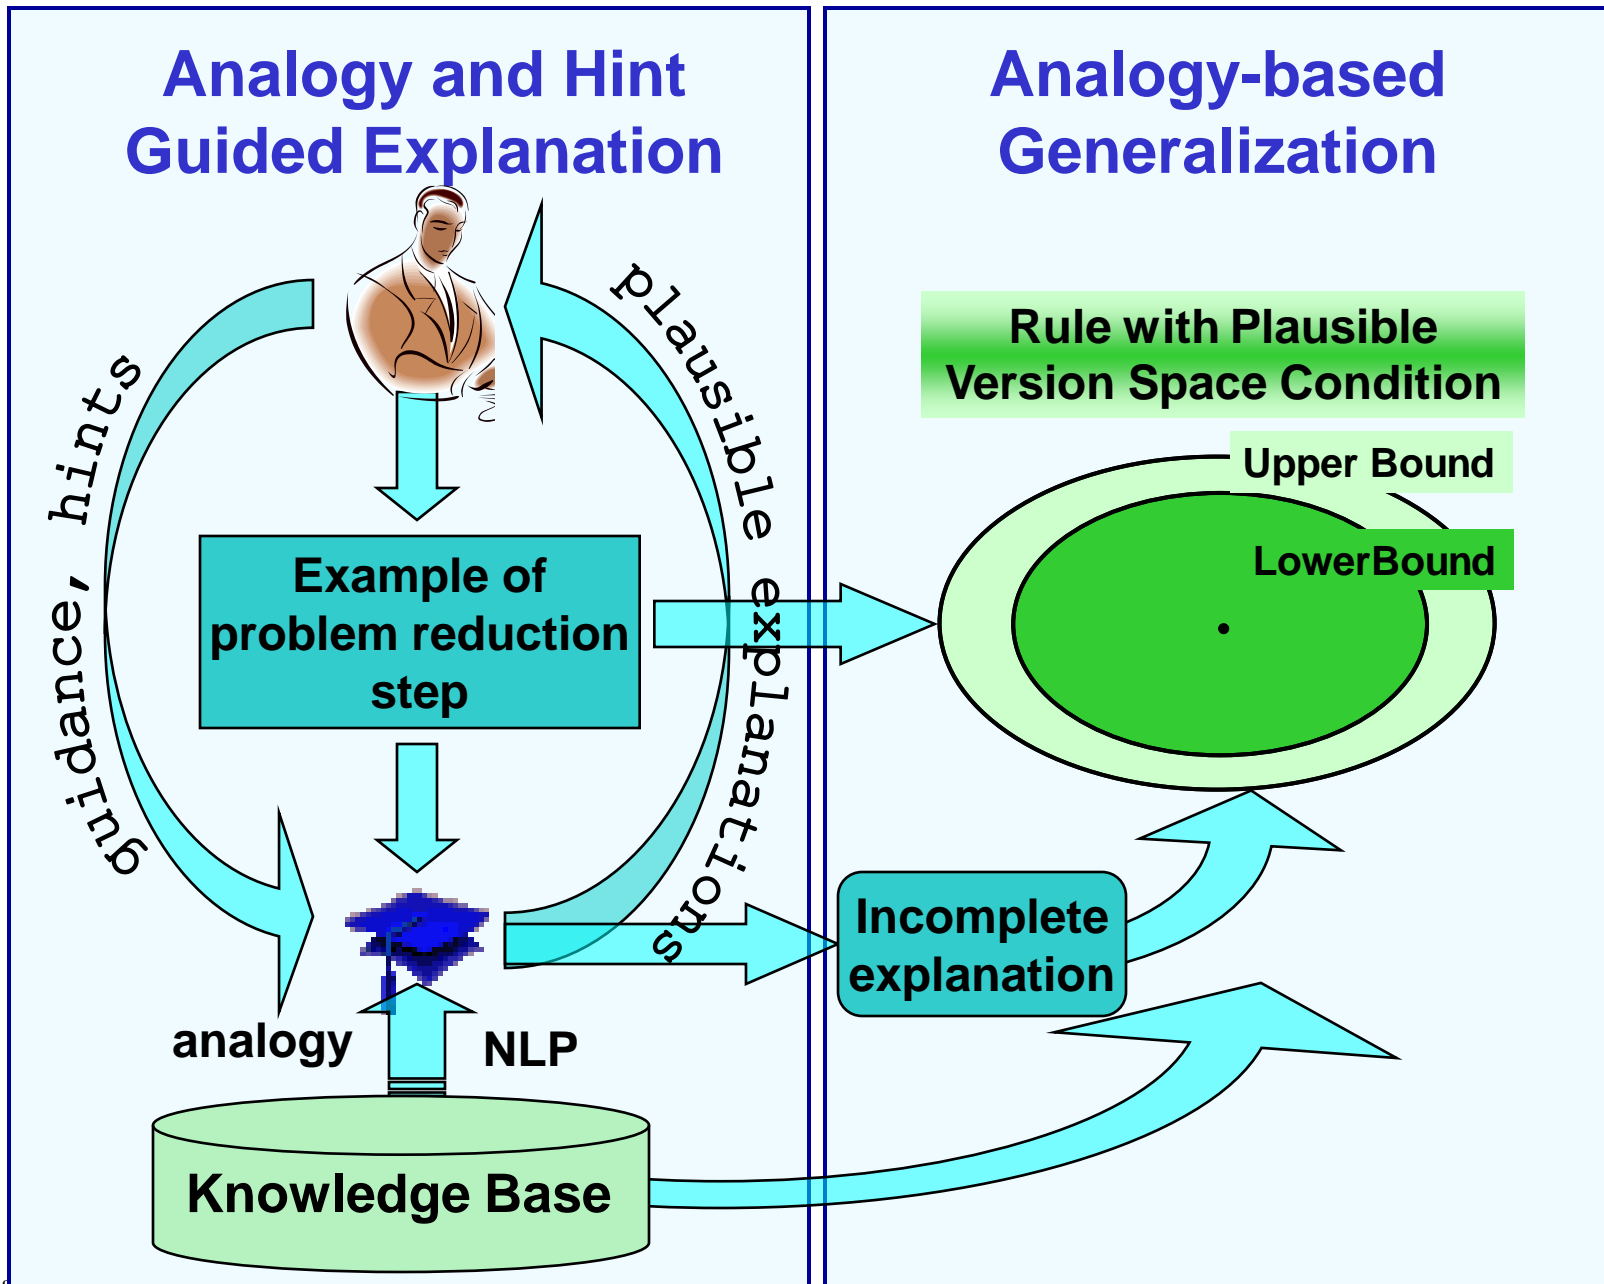


P5. Integrated Teaching and Learning

Develop agent teaching and learning methods where the subject matter expert helps the agent to learn (e.g. by giving examples, hints and explanations), and the agent helps the expert to teach it (e.g. by asking relevant questions).



Integrated Teaching and Learning in Disciple



Reasoning Rules Learned from Analyst's Solution



1. The analyst extends the analysis logic

Assess whether there are states with nuclear weapons that may be willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?
North Korea

Assess whether North Korea may be willing to sell nuclear weapons to Al Qaeda.

What might be a possible reason for North Korea to sell nuclear weapons to Al Qaeda?
United States is perceived as a common enemy of North Korea and Al Qaeda.

Assess to what extent the perception that United States is a common enemy of North Korea and Al Qaeda might be a good reason for North Korea to sell nuclear weapons to Al Qaeda.

2. Disciple learns reasoning rules



Learned Rule

DECOMPOSITION RULE DDR.00104 FORMAL DESCRIPTION

IF: Assess whether there are states with nuclear weapons that may be willing to sell nuclear weapons to ?O1.

Q:	Which is a nuclear state?
A:	?O2

MAIN CONDITION

Var	Lower Bound	Upper Bound
?O1	(terrorist group)	(actor)
?O2	(nuclear state)	(nuclear state)

THEN: Assess whether ?O2 may be willing to sell nuclear weapons to ?O1.

To assess whether there are states that may be willing to sell nuclear weapons to an actor, **one has to consider each nuclear state** and assess whether that state may be willing to sell nuclear weapons to that actor.

Learned Rule

DECOMPOSITION RULE DDR.00105 FORMAL DESCRIPTION

IF: Assess whether ?O1 may be willing to sell nuclear weapons to ?O2?

Q:	What might be a possible reason for ?O1 to sell nuclear weapons to ?O2?
A:	?O3 is perceived as a common enemy of ?O1 and ?O2

MAIN CONDITION

Var	Lower Bound	Upper Bound
?O1	(terrorist group)	(actor)
?O2	(terrorist group)	(actor)
?O3	(superpower)	(actor)

Var	Relationship	Var
?O1	perceives as enemy	?O3
?O2	perceives as enemy	?O3

THEN: Assess to what extent the perception that ?O3 is a common enemy of ?O1 and ?O2 might be a good reason for ?O1 to sell nuclear weapons to ?O2.

Overview: Design Principles

Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning



Multistrategy Learning

Learning with an Evolving Representation Space

Mixed-Initiative Modeling, Learning and Problem Solving

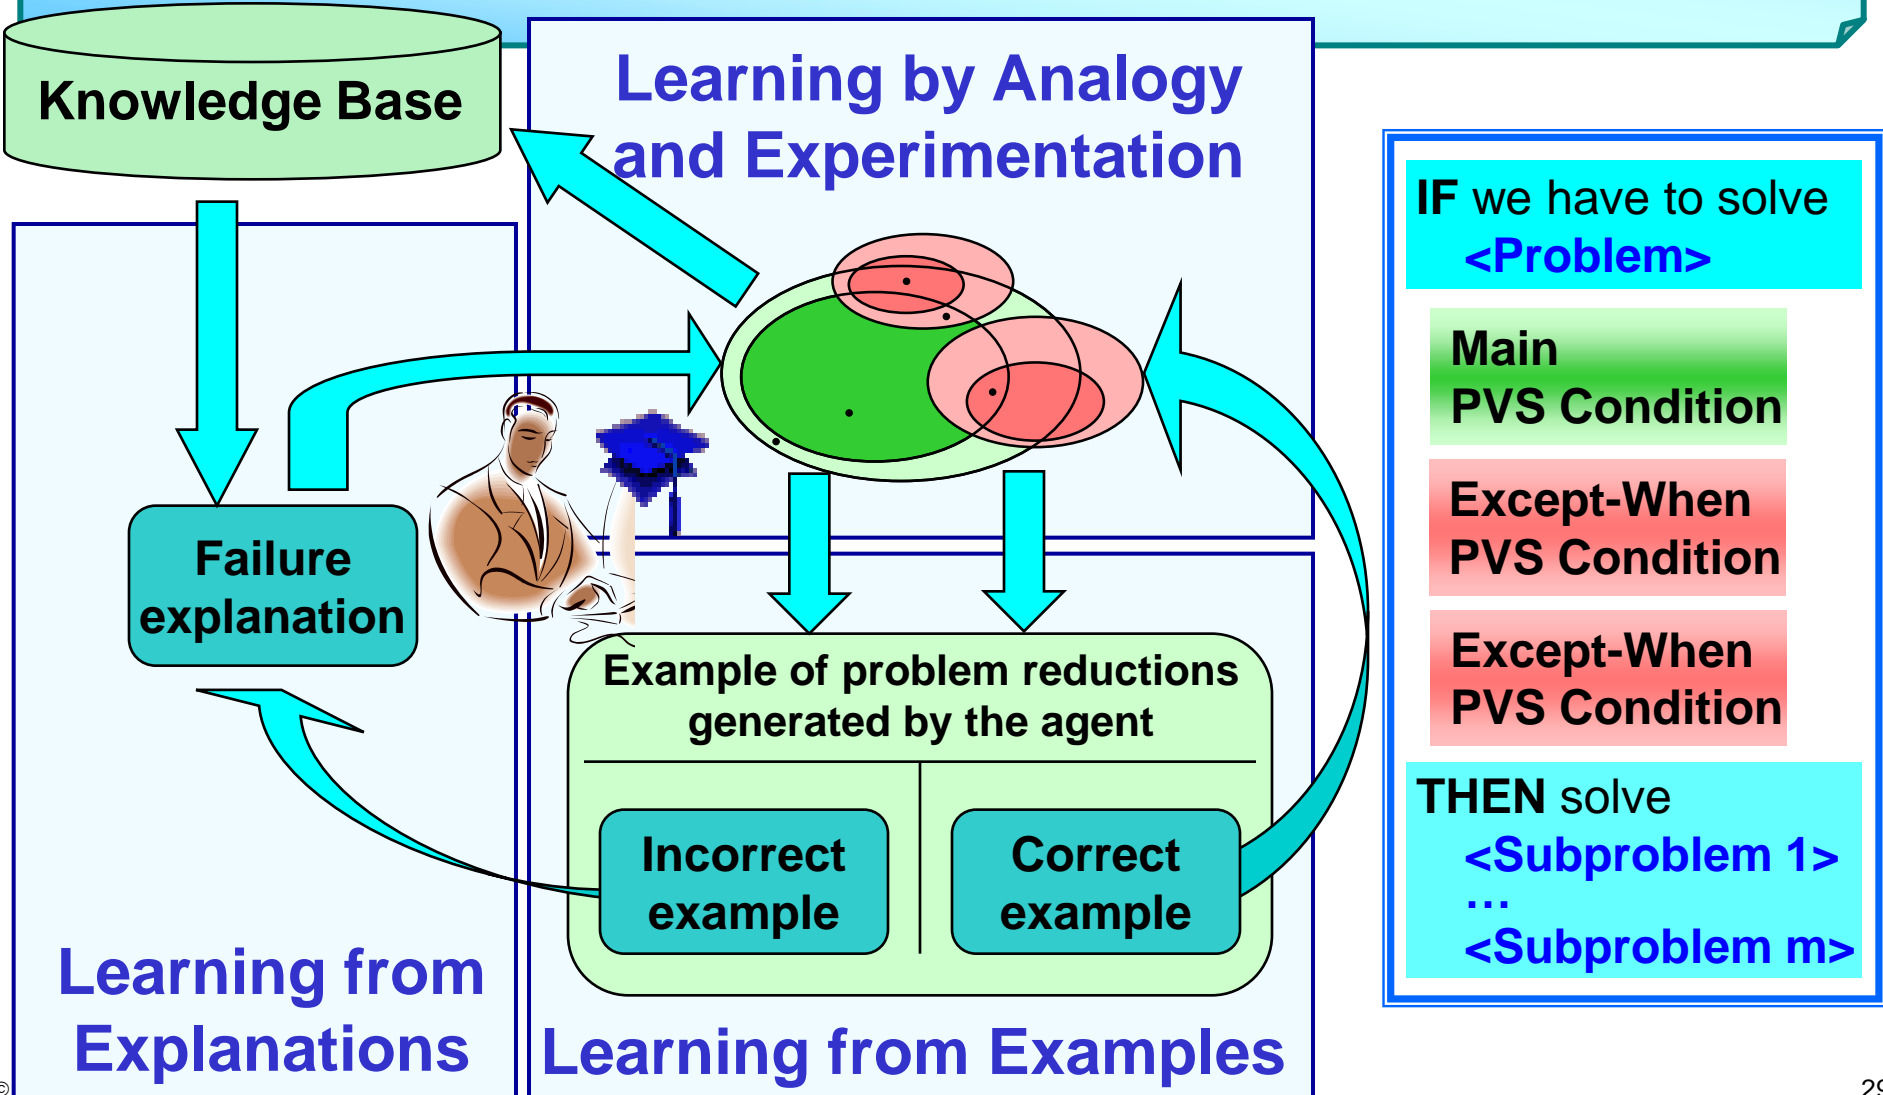
Plausible Reasoning with Partially Learned Knowledge

User Tutoring in Problem Solving

Agent Architecture for Generality-Power Tradeoff

P6. Multistrategy Learning

Develop multistrategy learning methods that integrate complementary learning strategies to take advantage of their strengths to compensate for each other's weaknesses



Rules Refined based on Analyst's Critique

2. The analyst critiques the reasoning

1. Disciple applies the learned rule

3. Disciple-LTA refines the rule with an except-when condition

Disciple Agent KB



This is wrong!

France will not sell nuclear weapons to Al Qaeda because it perceives it as an enemy.

Assess whether there are states with nuclear weapons willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?
North Korea

Assess whether North Korea is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?
France

Assess whether France is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?
India

Assess whether India is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?
Iran

Assess whether Iran is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?
Israel

Assess whether Israel is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?
Pakistan

Assess whether Pakistan is willing to sell nuclear weapons to Al Qaeda.

To assess whether there are states that may be willing to sell nuclear weapons to an actor, **one has to consider each nuclear state** and assess whether that state may be willing to sell nuclear weapons to that actor, **except for** the case in which the nuclear state is an enemy of that actor.

Rule Viewer

Refined Rule LE DDR.00205 FORMAL

DESCRIPTION

IF: Assess whether there are states with nuclear weapons willing to sell nuclear weapons to ?O1.

Q: Which is a nuclear state?

A: ?O2

MAIN CONDITION

Var	Lower Bound	Upper Bound
?O1	(terrorist group)	(actor)
?O2	(nuclear state)	(nuclear state)

EXCEPT WHEN CONDITION 1

Var	Lower Bound	Upper Bound
?O2	(nuclear state)	(actor)
?O1	(terrorist group)	(actor)

Var	Relationship	Var
?O2	perceives as enemy	?O1

THEN: Assess whether ?O2 is willing to sell nuclear weapons to ?O1.

Rules Refined based on Analyst's Critique

2. The analyst critiques the reasoning



This is wrong!

Russia will not sell nuclear weapons to Al Qaeda because it opposes the proliferation of nuclear weapons.

1. Disciple applies the refined rule

Assess whether there are states with nuclear weapons willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state which is not an enemy of Al Qaeda?
North Korea

Assess whether North Korea is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state which is not an enemy of Al Qaeda?
Iran

Assess whether Iran is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state which is not an enemy of Al Qaeda?
Pakistan

Assess whether Pakistan is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state which is not an enemy of Al Qaeda?
Russia

Assess whether Russia is willing to sell nuclear weapons to Al Qaeda.

To assess whether there are states that may be willing to sell nuclear weapons to an actor, **one has to consider each nuclear state** and assess whether that state may be willing to sell nuclear weapons to that actor, **except for** the case in which the nuclear state is an enemy of that actor and also **except for** the case when the nuclear state opposes the proliferation of nuclear weapons.

3. Disciple refines the rule with a new except-when condition

Disciple Agent KB

Refined Rule

DECOMPOSITION RULE DDR.00205 FORMAL DESCRIPTION

IF: Assess whether there are states with nuclear weapons willing to sell nuclear weapons to ?O1.

Q: Which is a nuclear state which is not an enemy of ?O1?

A: ?O2

MAIN CONDITION

Var	Lower Bound	Upper Bound
?O1	(terrorist group)	(actor)
?O2	(nuclear state)	(nuclear state)

EXCEPT WHEN CONDITION 1

Var	Lower Bound	Upper Bound
?O2	(nuclear state)	(actor)
?O1	(terrorist group)	(actor)

Var	Relationship	Var
?O2	perceives as enemy	?O1

EXCEPT WHEN CONDITION 2

Var	Lower Bound	Upper Bound
?O2	(nuclear state)	(actor)
?SI1	[medium - medium]	[very low - very high]

Var	Relationship	Var
?O2	degree of opposition to nuclear weapons proliferation	?SI1

Assess whether ?O2 is willing to sell nuclear weapons to ?O1.



Disciple-LTA Demo: Solving, Modeling, and Learning

Discussion: Rule Refinement with Negative Example




- Learn Except When Condition (C, Ex)
- Keep as Negative Exception (C, Ex)

Rule Condition C



- Keep as Negative Example (C, Ex)



- Specialize Upper Bound of Main Cond (C, Ex)
- Learn Except When Condition (C, Ex)
- Keep as Negative Exception (C, Ex)



- Generalize Lower Bound of Except When Condition (C, Ex)



- Generalize Lower Bound of Except When Cond (C, Ex)
- Learn Except When Condition (C, Ex)
- Keep as Negative Exception (C, Ex)



- Generalize Lower Bound of Except When Cond (C, Ex)
- Specialize Upper Bound of Main Cond (C, Ex)
- Learn Except When Condition (C, Ex)
- Keep as Negative Exception (C, Ex)

Overview: Design Principles

Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning

Multistrategy Learning



Learning with an Evolving Representation Space

Mixed-Initiative Modeling, Learning and Problem Solving

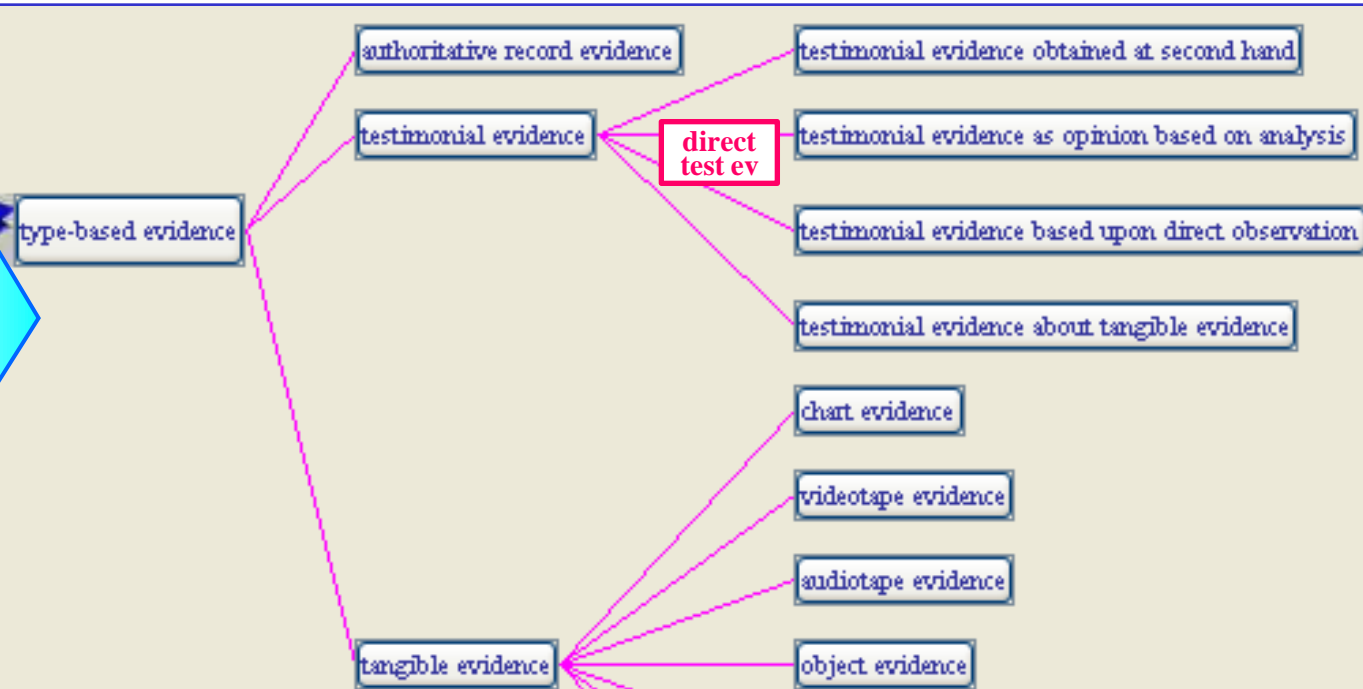
Plausible Reasoning with Partially Learned Knowledge

User Tutoring in Problem Solving

Agent Architecture for Generality-Power Tradeoff

P7. Learning with an Evolving Representation Space

Develop methods that allow continuous adaptation of the previously learned rules to the evolution of the ontology.



IF <Problem>


PVS Condition

Except-When
PVS Condition

THEN <Subproblem 1>
...
<Subproblem m>

PVS Condition

Except-When
PVS Condition



Characterization of the Disciple Learning Methods

Uses the explanation of the first positive example to generate a much smaller version space than the classical version space method.

Conducts an efficient heuristic search of the version space, guided by explanations, and by the maintenance of a single upper bound condition and a single lower bound condition.

Will always learn a rule, even in the presence of exceptions.

Learns from a few examples and an incomplete knowledge base.

Uses a form of multistrategy learning that synergistically integrates learning from examples, learning from explanations, and learning by analogy, to compensate for the incomplete knowledge.

Uses mixed-initiative reasoning to involve the expert in the learning process.

Is applicable to complex real-world domains, being able to learn within a complex and evolving representation language.

Overview: Design Principles

Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning

Multistrategy Learning

Learning with an Evolving Representation Space



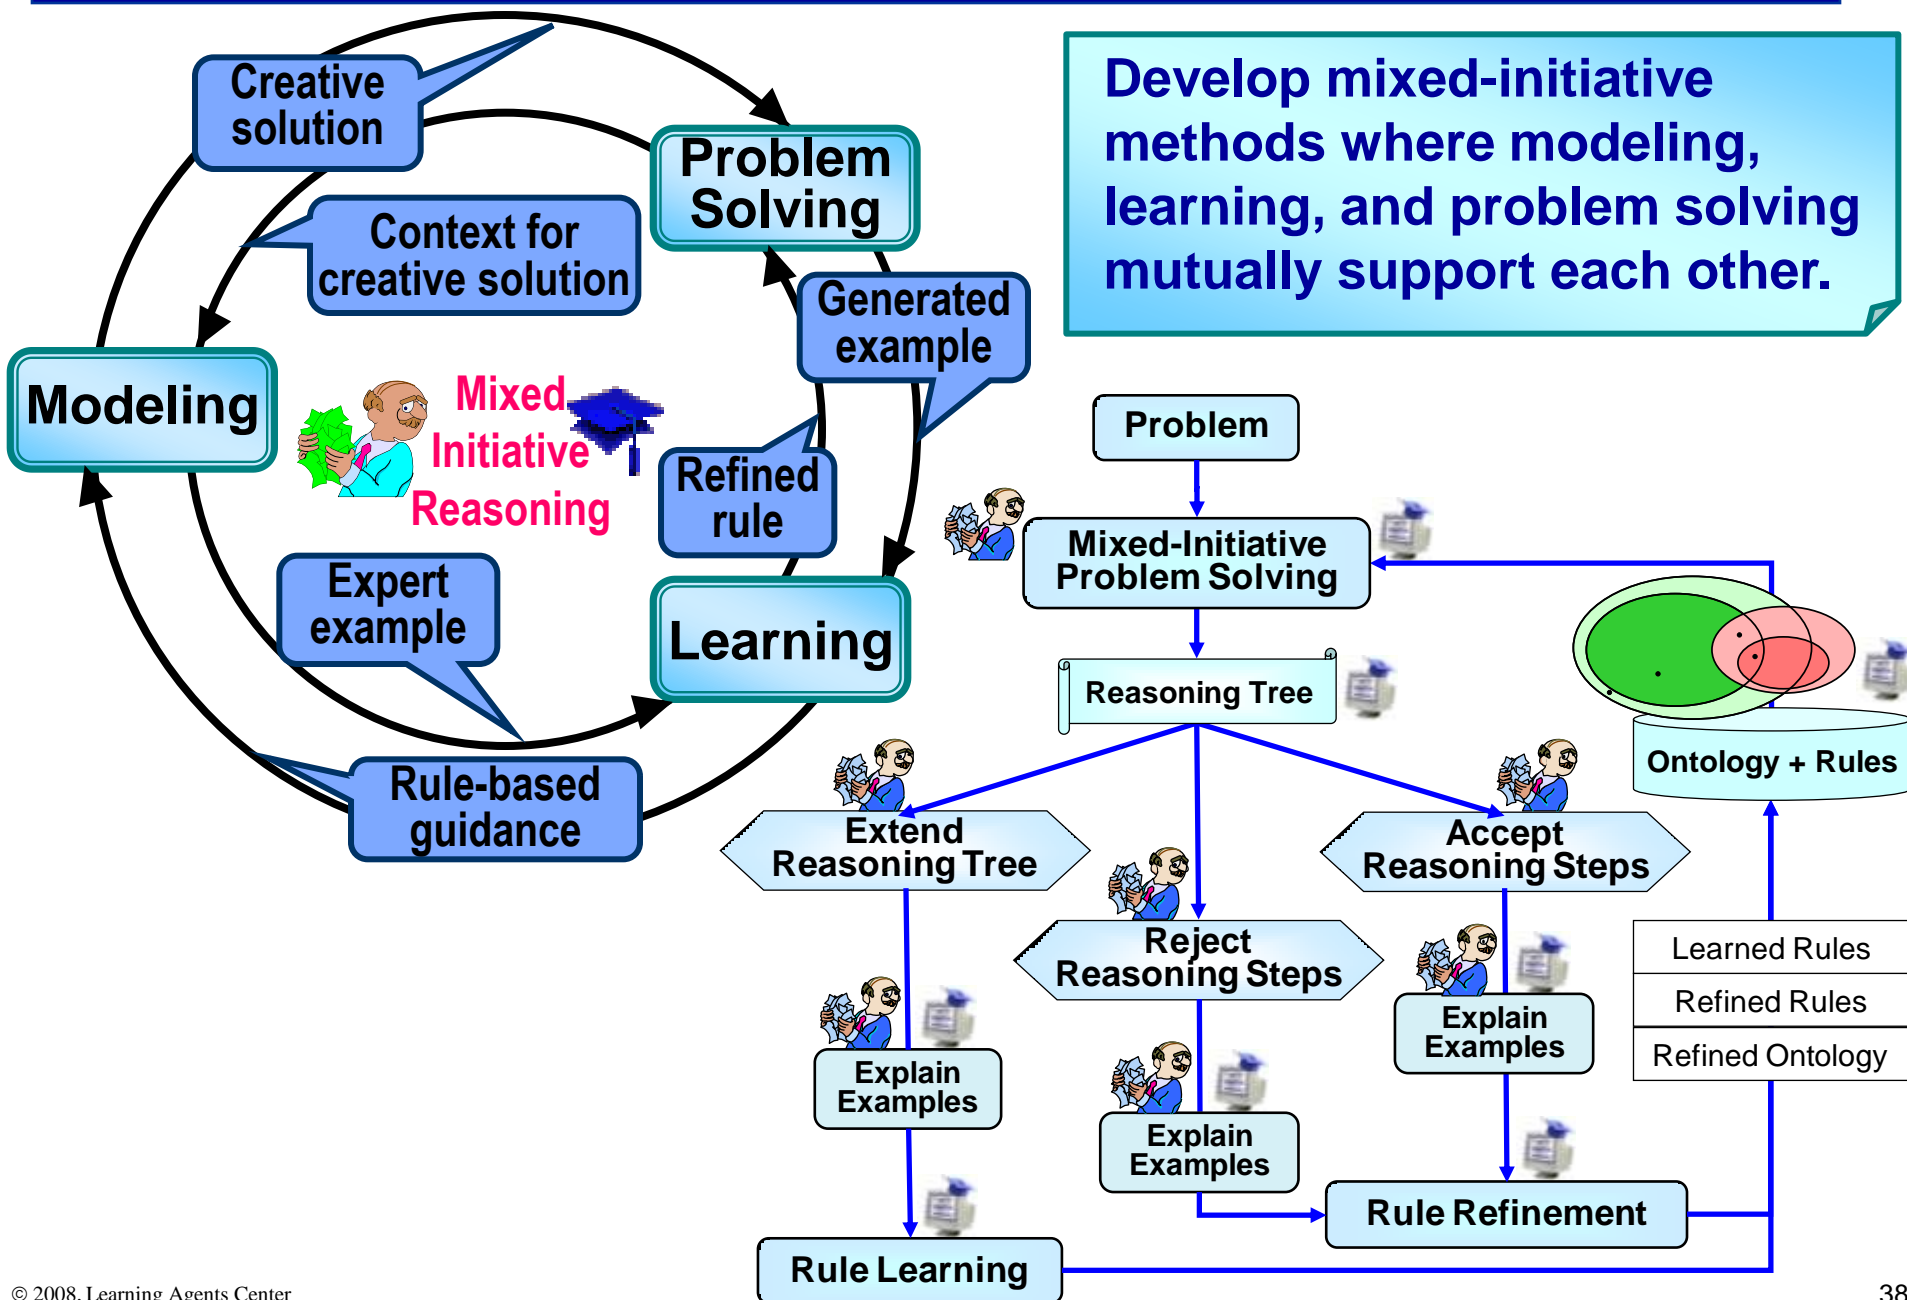
Mixed-Initiative Modeling, Learning and Problem Solving

Plausible Reasoning with Partially Learned Knowledge

User Tutoring in Problem Solving

Agent Architecture for Generality-Power Tradeoff

P8. Mixed-Initiative Modeling, Learning and Pb. Solv.



Overview: Design Principles

Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning

Multistrategy Learning

Learning with an Evolving Representation Space

Mixed-Initiative Modeling, Learning and Problem Solving

➡ Plausible Reasoning with Partially Learned Knowledge

User Tutoring in Problem Solving


Agent Architecture for Generality-Power Tradeoff

P9. Plausible Reasoning with Partially Learned Knowledge

Develop reasoning methods based on confidence levels that allow efficient use of partially learned rules for modeling expert's reasoning, learning and problem solving.

PVS Condition

Except-When PVS Condition



Overview: Design Principles

Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning

Multistrategy Learning

Learning with an Evolving Representation Space

Mixed-Initiative Modeling, Learning and Problem Solving

Plausible Reasoning with Partially Learned Knowledge



User Tutoring in Problem Solving


Agent Architecture for Generality-Power Tradeoff

P10. User Tutoring in Problem Solving

Develop approaches to user tutoring that allow the agent to easily and rapidly teach the user its problem solving paradigm, to facilitate their collaboration.



Disciple-LTA Demo: Tutoring



Lesson Fragment: Hypothesis support from a piece of evidence

Lesson: Hypothesis support from [piece of evidence](#).

Assess to what extent the [piece of evidence](#) favors the hypothesis.

The information provided by the [piece of evidence](#) and the extent to which it is believable.

Assess to what extent the [piece of evidence](#) favors the hypothesis, assuming that we believe the information provided by the [piece of evidence](#).

Assess the extent to which the information provided by the [piece of evidence](#) is believable.

Abstract reduction strategy

The [piece of evidence](#) is [testimonial evidence obtained at second hand](#).

The [piece of evidence](#) is [testimonial evidence about tangible evidence](#).

The [piece of evidence](#) is [direct testimonial evidence](#).

Assess the [believability](#) the reporter of the [piece of evidence](#).

Assess the [believability](#) the source of the [piece of evidence](#).

Assess the [believability](#) the reporter of the [piece of evidence](#).

Assess the [credibility](#) of the [tangible evidence](#).

Assess the [believability](#) the source of the [piece of evidence](#).

Lesson on Evidence

Assess to what extent the piece of evidence [EVD-Dawn-Mir01-02c](#) favors the hypothesis that [Al Qaeda](#) considers deterrence as a reason to obtain [nuclear weapons](#).

Q: What factors determine how a piece of evidence favors a hypothesis?
A: The information provided by the piece of evidence and the extent to which it is believable.

Assess to what extent [EVD-Dawn-Mir01-02c](#) favors the hypothesis that [Al Qaeda](#) considers deterrence as a reason to obtain [nuclear weapons](#) given the information provided by

Assess the extent to which the information provided by [EVD-Dawn-Mir01-02c](#) is believable.

Automatically generated illustration of the abstract strategy

Q: How was [EVD-Dawn-Mir01-02c](#) obtained?

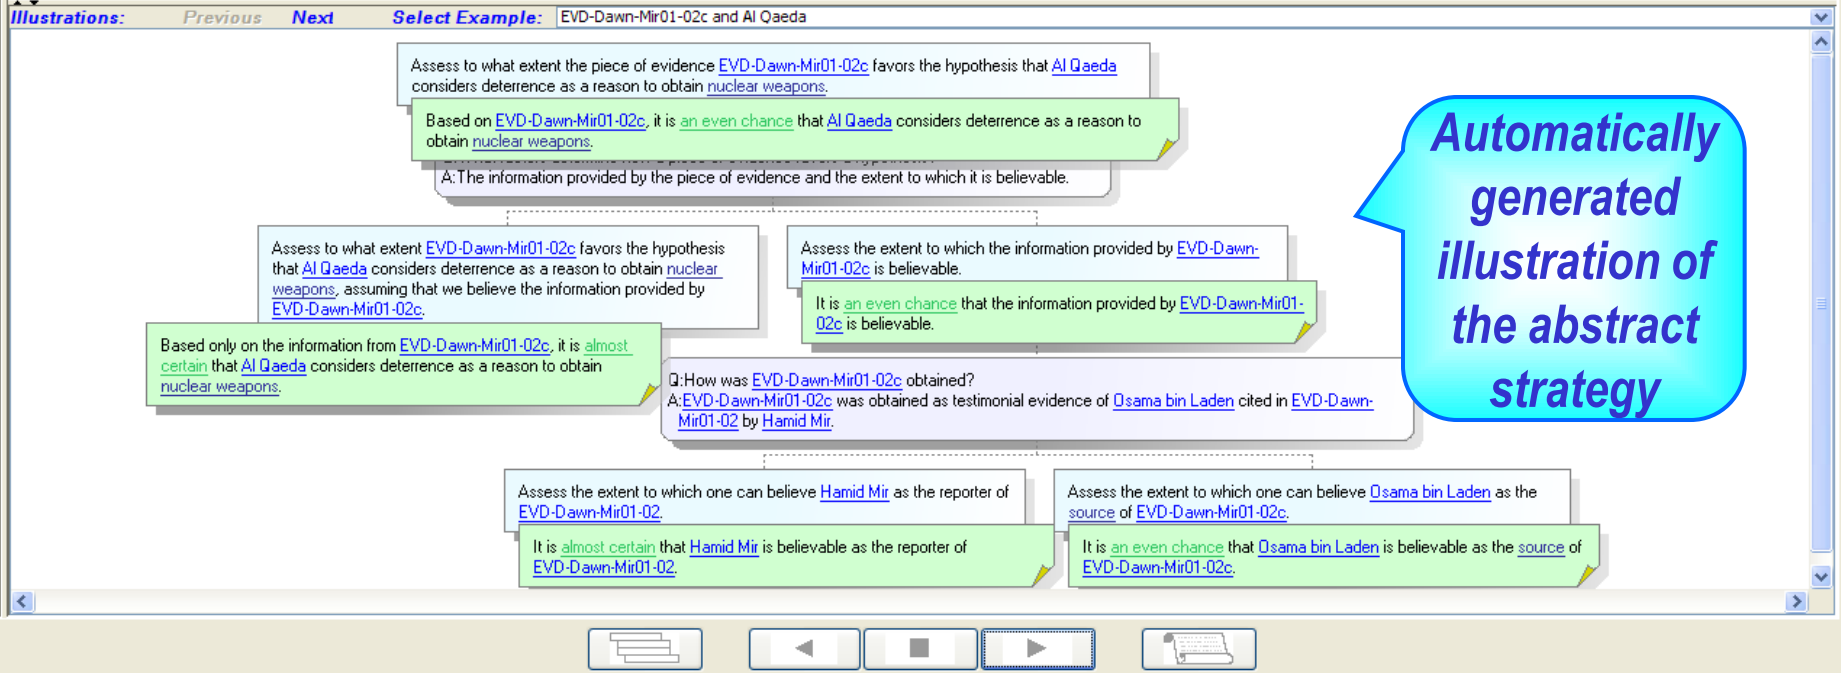
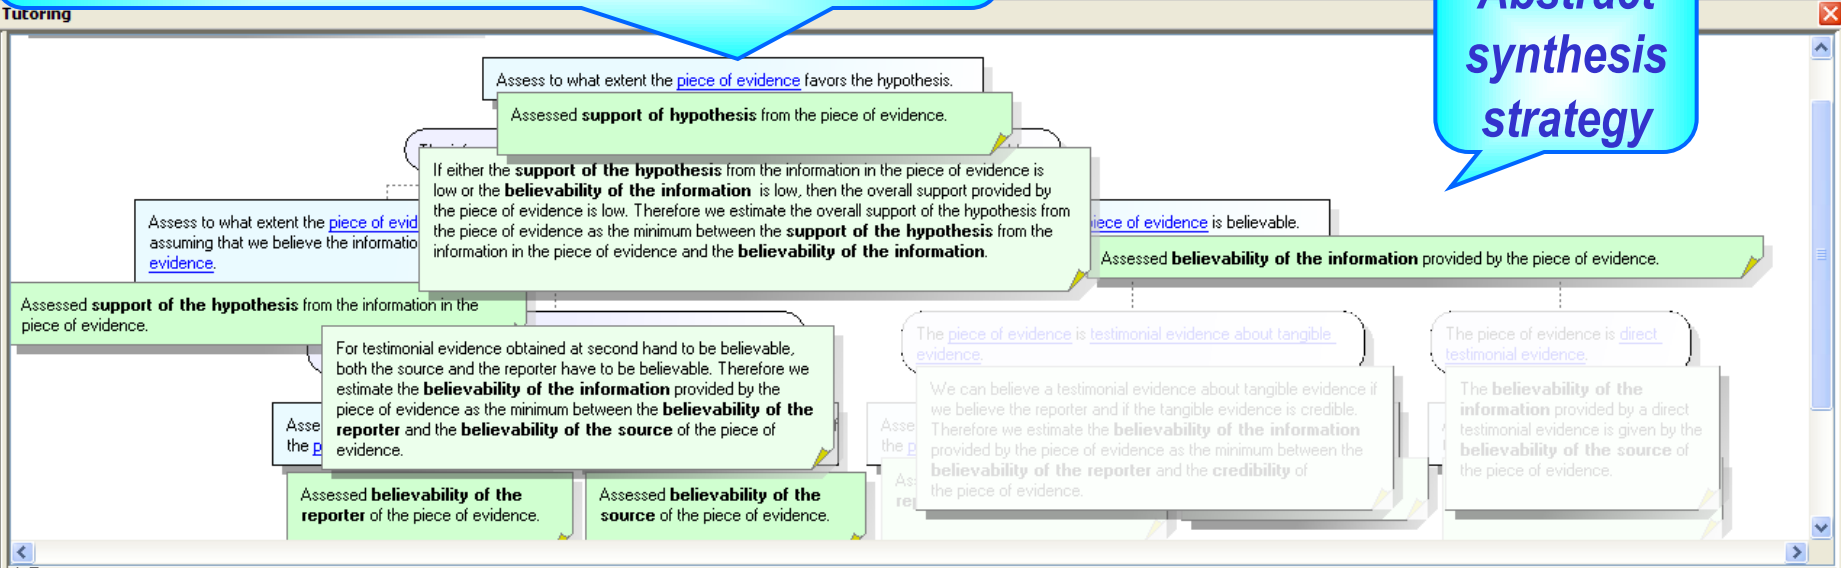
A: [EVD-Dawn-Mir01-02c](#) was obtained as testimonial evidence of [Osama bin Laden](#) cited in [EVD-Dawn-Mir01-02](#) by [Hamid Mir](#).

Assess the extent to which one can believe [Hamid Mir](#) as the reporter of [EVD-Dawn-Mir01-02](#).

Assess the extent to which one can believe [Osama bin Laden](#) as the source of [EVD-Dawn-Mir01-02c](#).

Lesson Fragment: Hypothesis support from a piece of evidence

Abstract synthesis strategy



Automatically generated illustration of the abstract strategy

Overview: Design Principles

Spiral Development with SMEs and End-Users

Problem Solving Paradigm for User-Agent Collaboration

Multi-Agent and Multi-Domain Problem Solving

Knowledge Base Structuring for Knowledge Reuse

Integrated Teaching and Learning

Multistrategy Learning

Learning with an Evolving Representation Space

Mixed-Initiative Modeling, Learning and Problem Solving

Plausible Reasoning with Partially Learned Knowledge

User Tutoring in Problem Solving

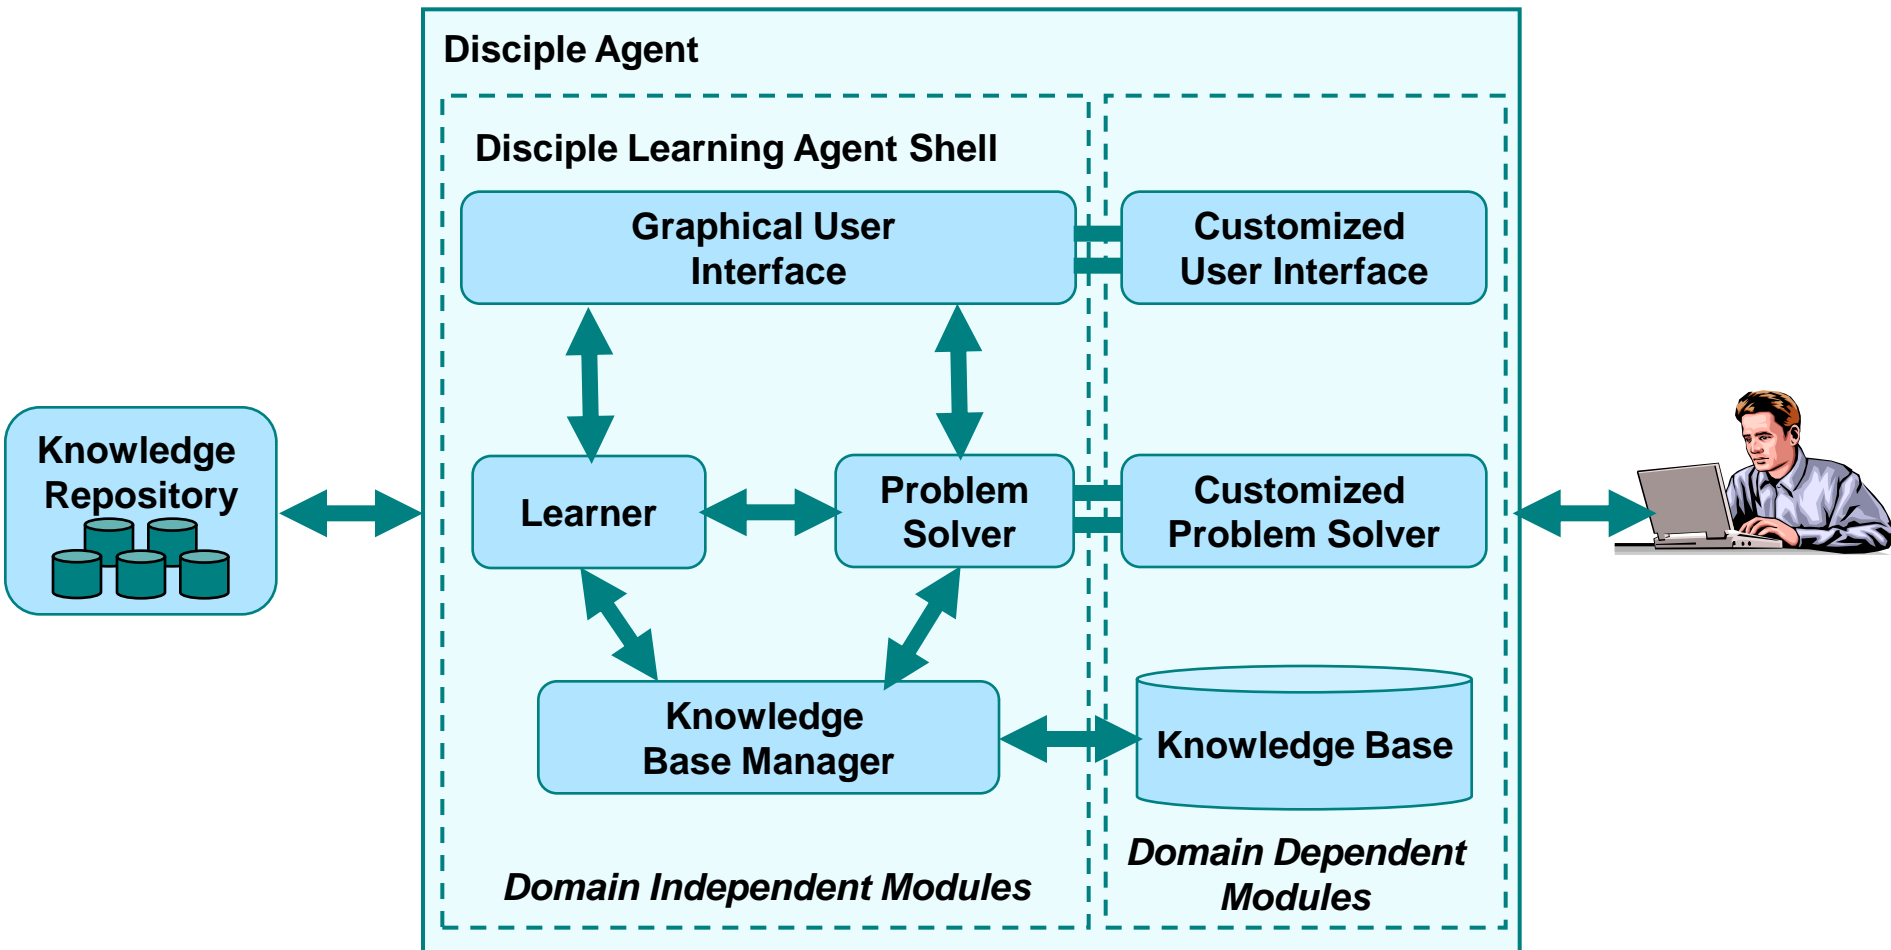
Agent Architecture for Generality-Power Tradeoff



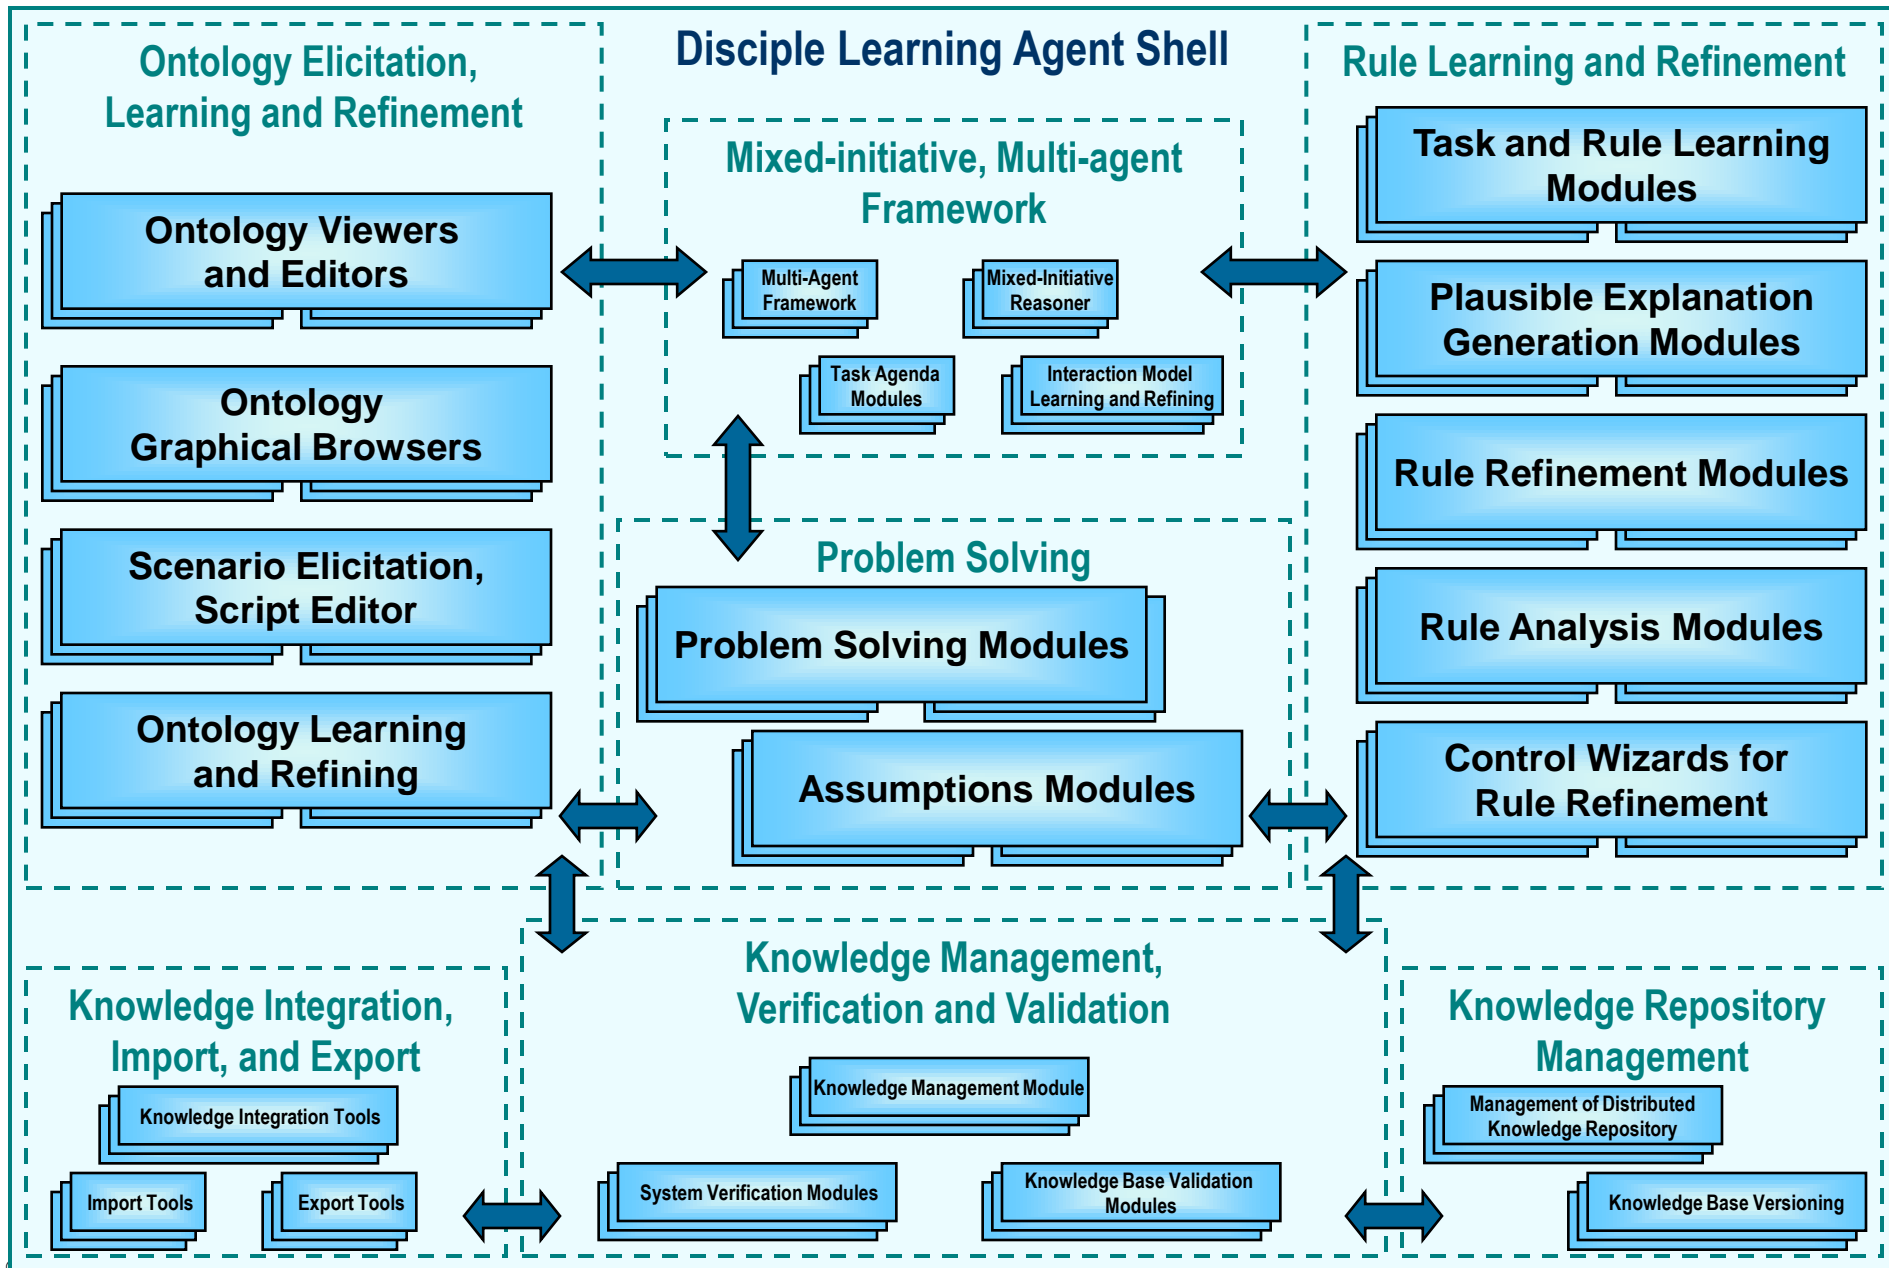
P11: Architecture for Generality-Power Tradeoff

Structure the architecture of the agent into two parts:

- o a reusable domain-independent learning agent shell;
- o domain specific modules.



Overall Architecture of the Disciple Shell



Customization of the Disciple Shell

Customization
for each type
of user

Subject Matter Expert



Field Application

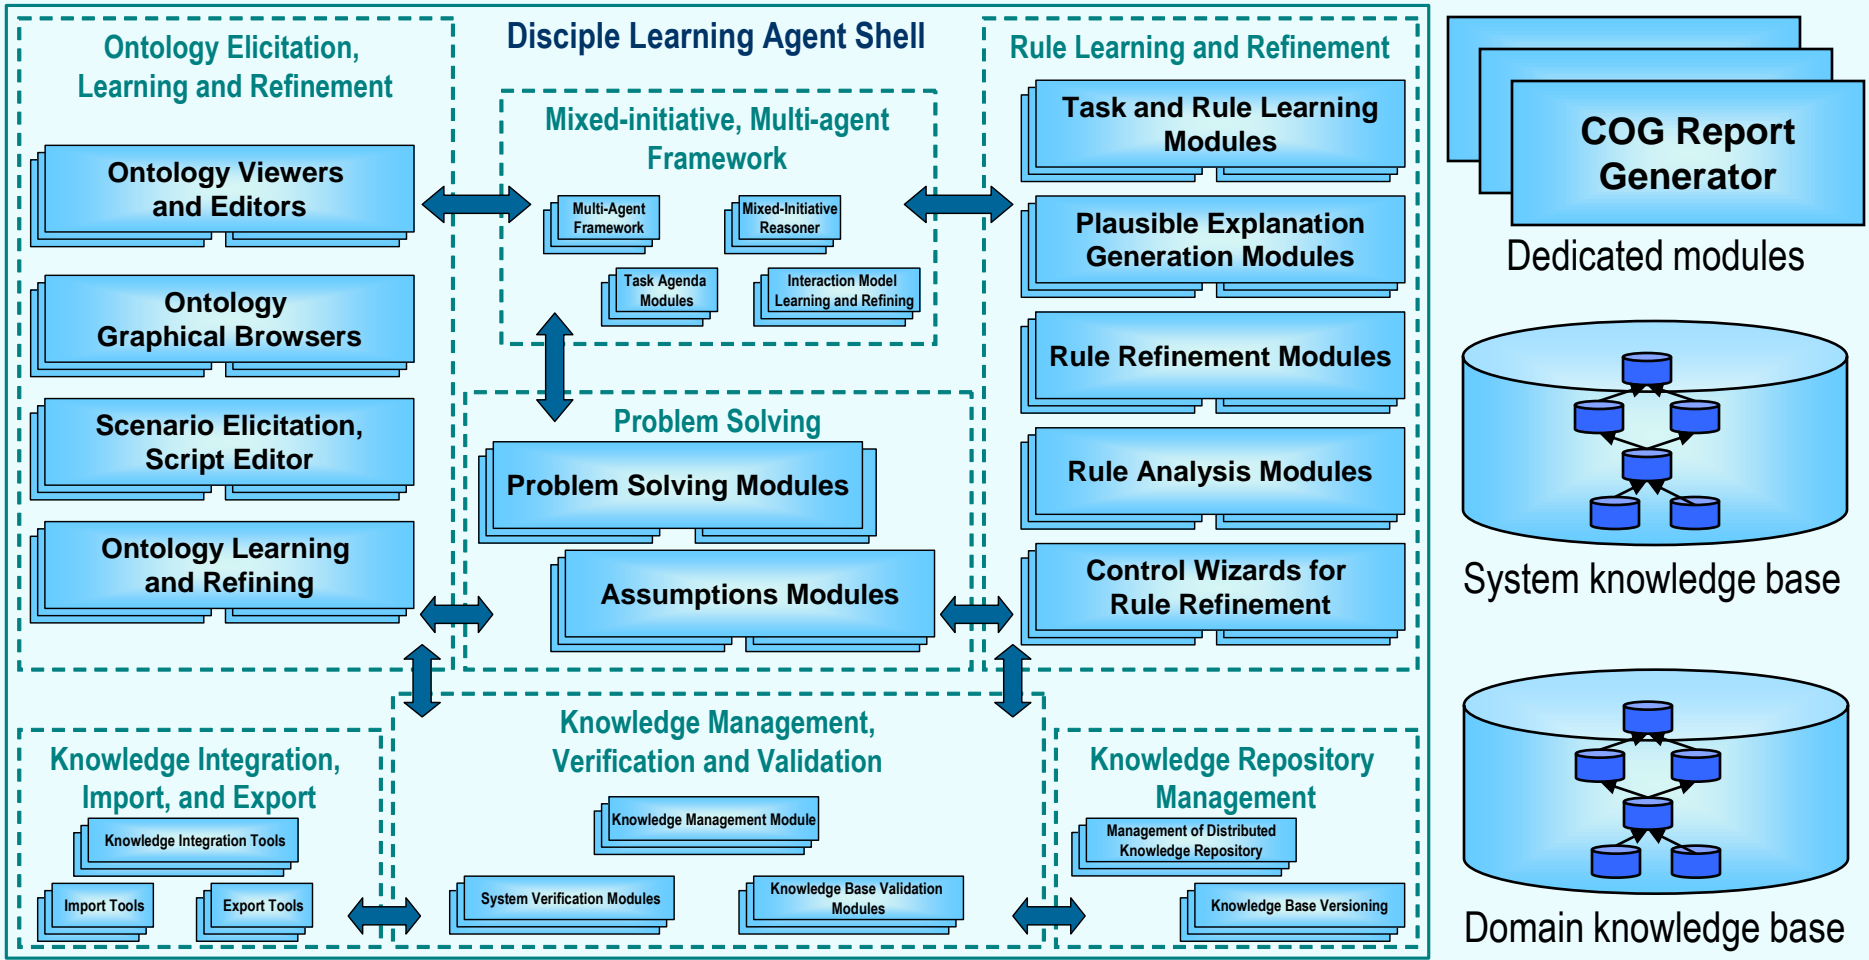


Education and training



Disciple COG

Customization for each domain



Conclusion: Research Vision for Learning Assistants

