Autonomy, Interaction & Learning – A Semantic Web Perspective

Michael Rovatsos

Department of Informatics
Technical University of Munich

15th April, 2004

Claim

Agent and multiagent technology will play a pivotal role in exploiting the full potential of the **Semantic Web** through

- agent **autonomy**
- agent **interaction**
- multiagent **learning**

because these features enable applications that go beyond the "content" view of the Internet.

Claim

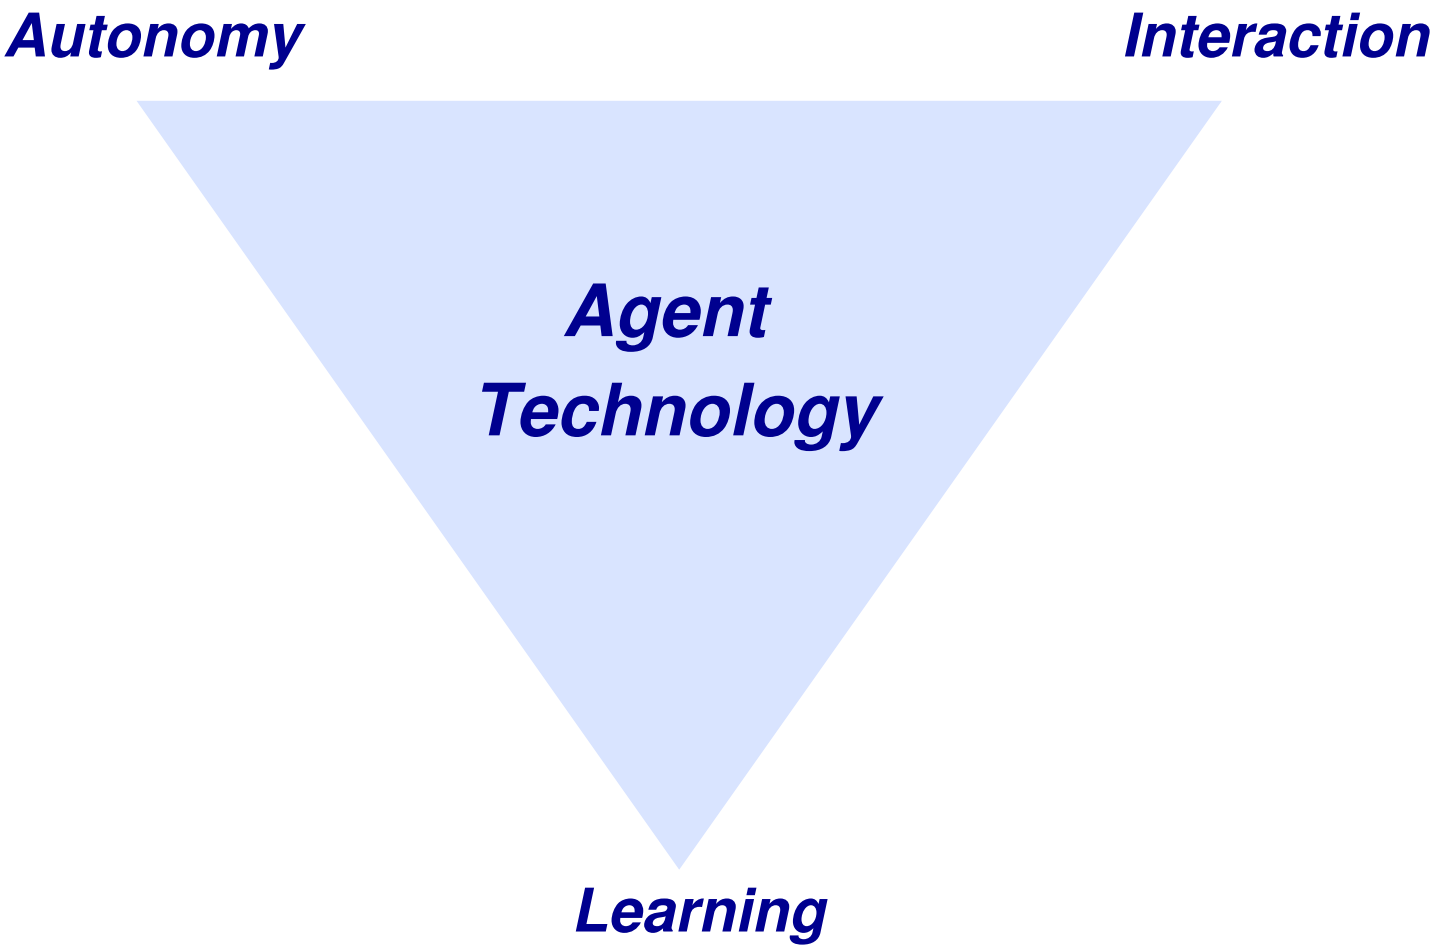

Claim

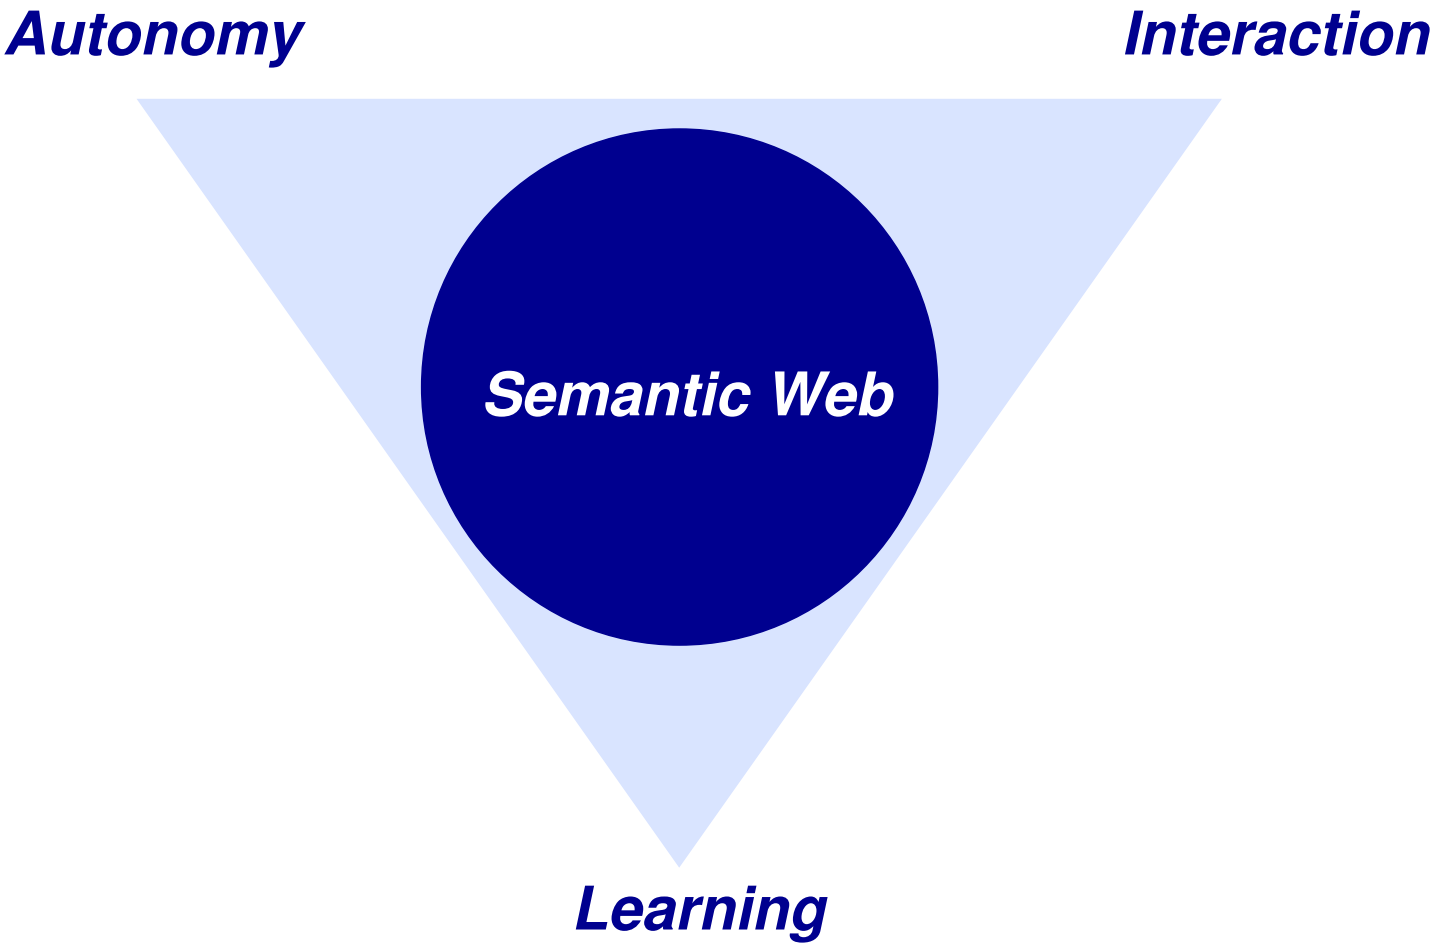

Outline

Outline

Agents

- **Situated** in an environment
- Able to **perceive** and **act upon** an environment (through sensors/effectors)
- **Autonomous**, i.e. able to operate without external guidance

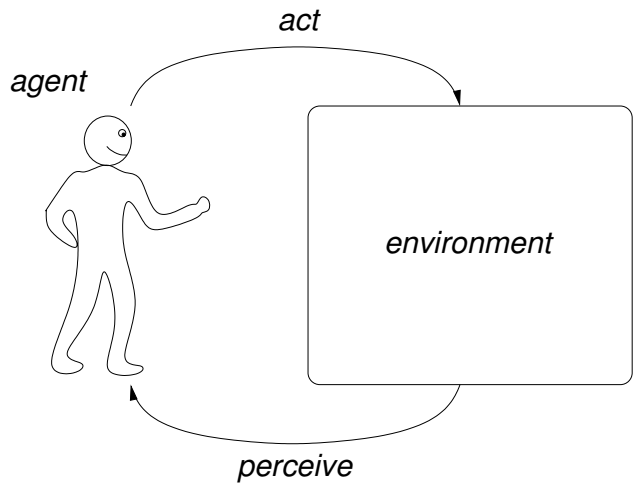

Agents

- **Situated** in an environment
- Able to **perceive** and **act upon** an environment (through sensors/effectors)
- **Autonomous**, i.e. able to operate without external guidance

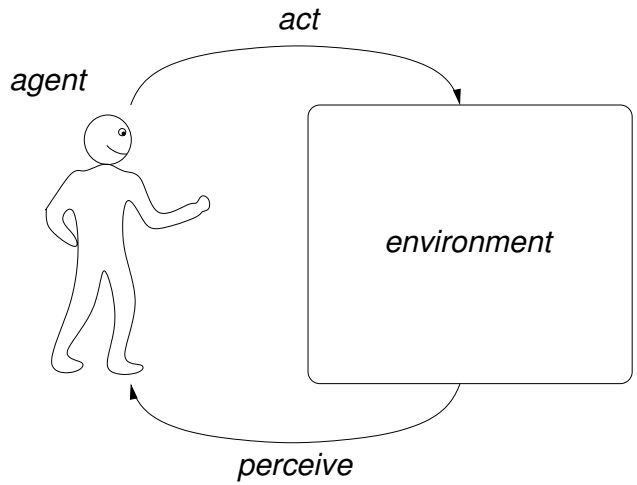

Intelligent Agents

- **Reactive** to changes in the environment
- **Pro-active**, i.e. they take action to achieve their goals
- Able to **interact** with others in a social context

Multiagent Systems (MAS)

- Societies of interacting agents
- Decentralisation of data and control
- Asynchronous computation and communication

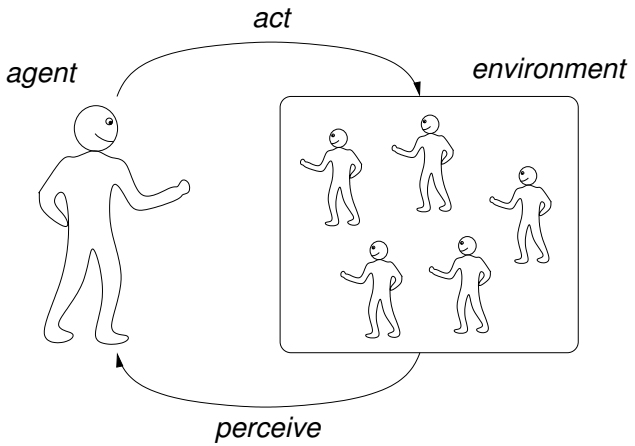

Multiagent Systems (MAS)

- Societies of interacting agents
- Decentralisation of data and control
- Asynchronous computation and communication
- Distributed AI distinguishes between two types:
  - Strictly cooperative, closed MAS (distributed problem-solving, task-oriented)
  - Open MAS (changing populations of self-interested agents, different internal designs)

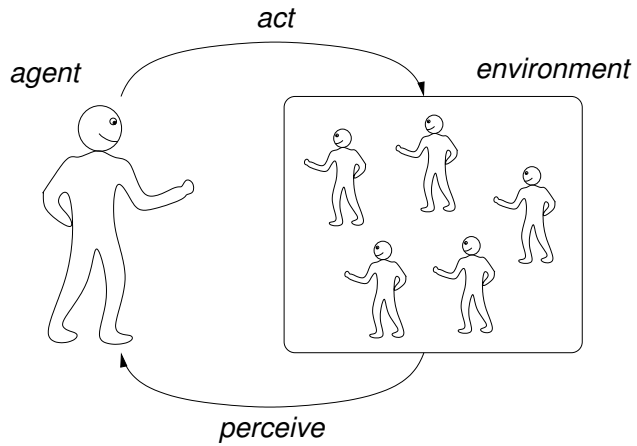

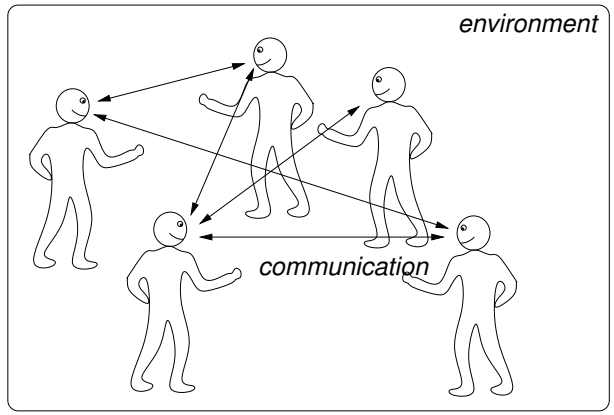

A more specific definition

- Agent concept very general/abstract, debatable

A more specific definition

- Agent concept very general/abstract, debatable
- In this talk, "agent" shall denote a computational entity that

A more specific definition

- Agent concept very general/abstract, debatable
- In this talk, "agent" shall denote a computational entity that
  - is capable of autonomous, goal-oriented action

A more specific definition

- Agent concept very general/abstract, debatable
- In this talk, "agent" shall denote a computational entity that
  - is capable of autonomous, goal-oriented action
  - is able to perform complex tasks specified at a high level

A more specific definition

- Agent concept very general/abstract, debatable
- In this talk, "agent" shall denote a computational entity that
  - is capable of autonomous, goal-oriented action
  - is able to perform complex tasks specified at a high level
  - employs AI techniques to perform its tasks (e.g. KR & R, planning, machine learning, NLP, machine vision, robotic control)

A more specific definition

- Agent concept very general/abstract, debatable
- In this talk, "agent" shall denote a computational entity that
  - is capable of autonomous, goal-oriented action
  - is able to perform complex tasks specified at a high level
  - employs AI techniques to perform its tasks (e.g. KR & R, planning, machine learning, NLP, machine vision, robotic control)
  - interacts with other agents (including humans) through a common physical environment

A more specific definition

- Agent concept very general/abstract, debatable
- In this talk, "agent" shall denote a computational entity that
  - is capable of autonomous, goal-oriented action
  - is able to perform complex tasks specified at a high level
  - employs AI techniques to perform its tasks (e.g. KR & R, planning, machine learning, NLP, machine vision, robotic control)
  - interacts with other agents (including humans) through a common physical environment
  - communicates with them using a high-level symbolic communication language

The Semantic Web & Agents

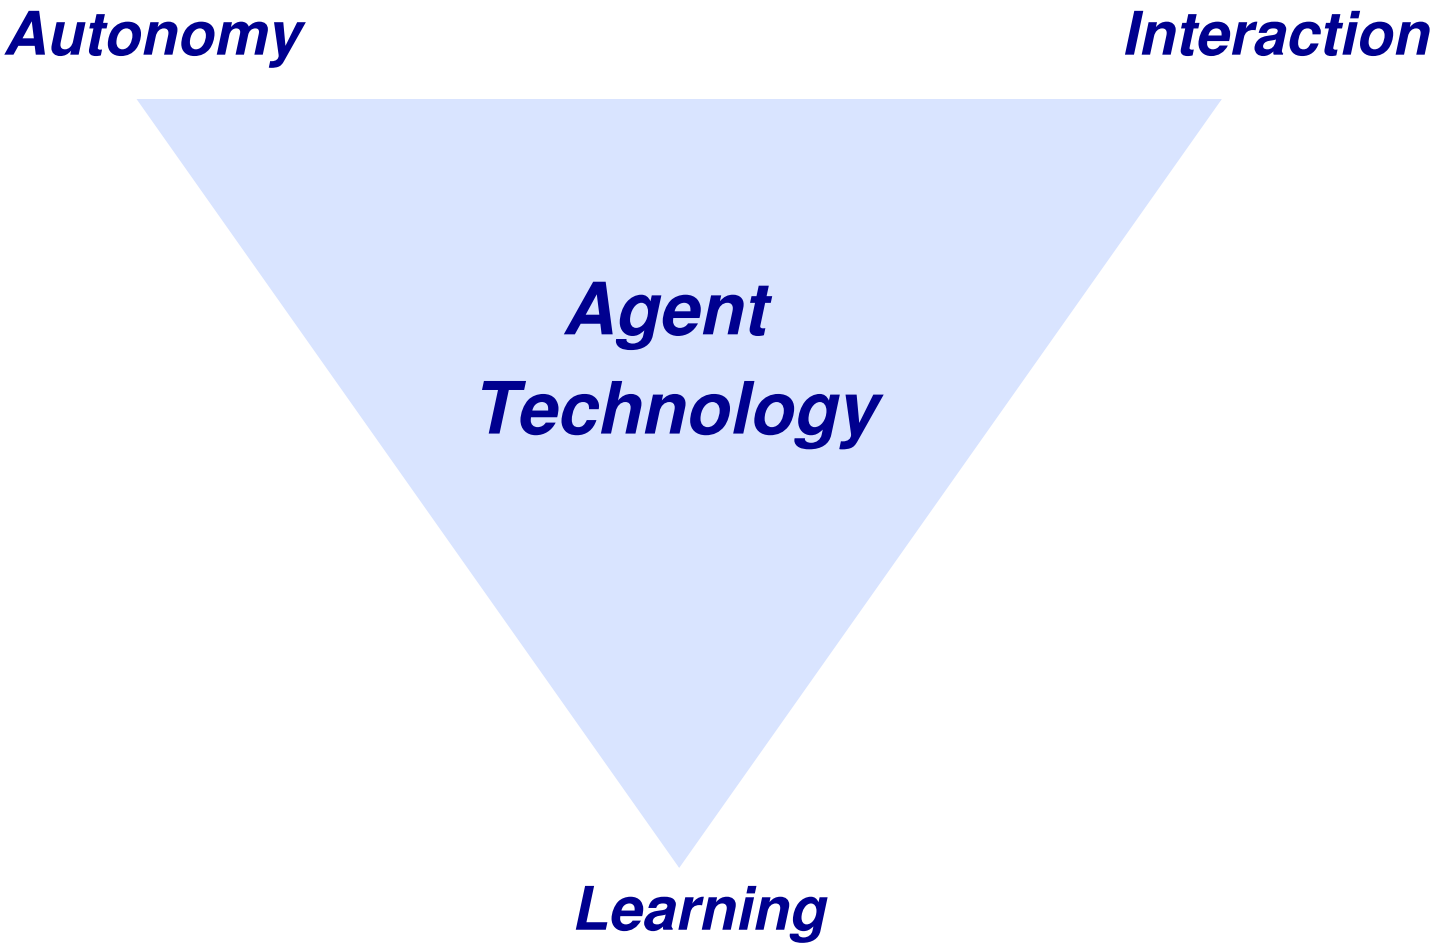

The Semantic Web & Agents

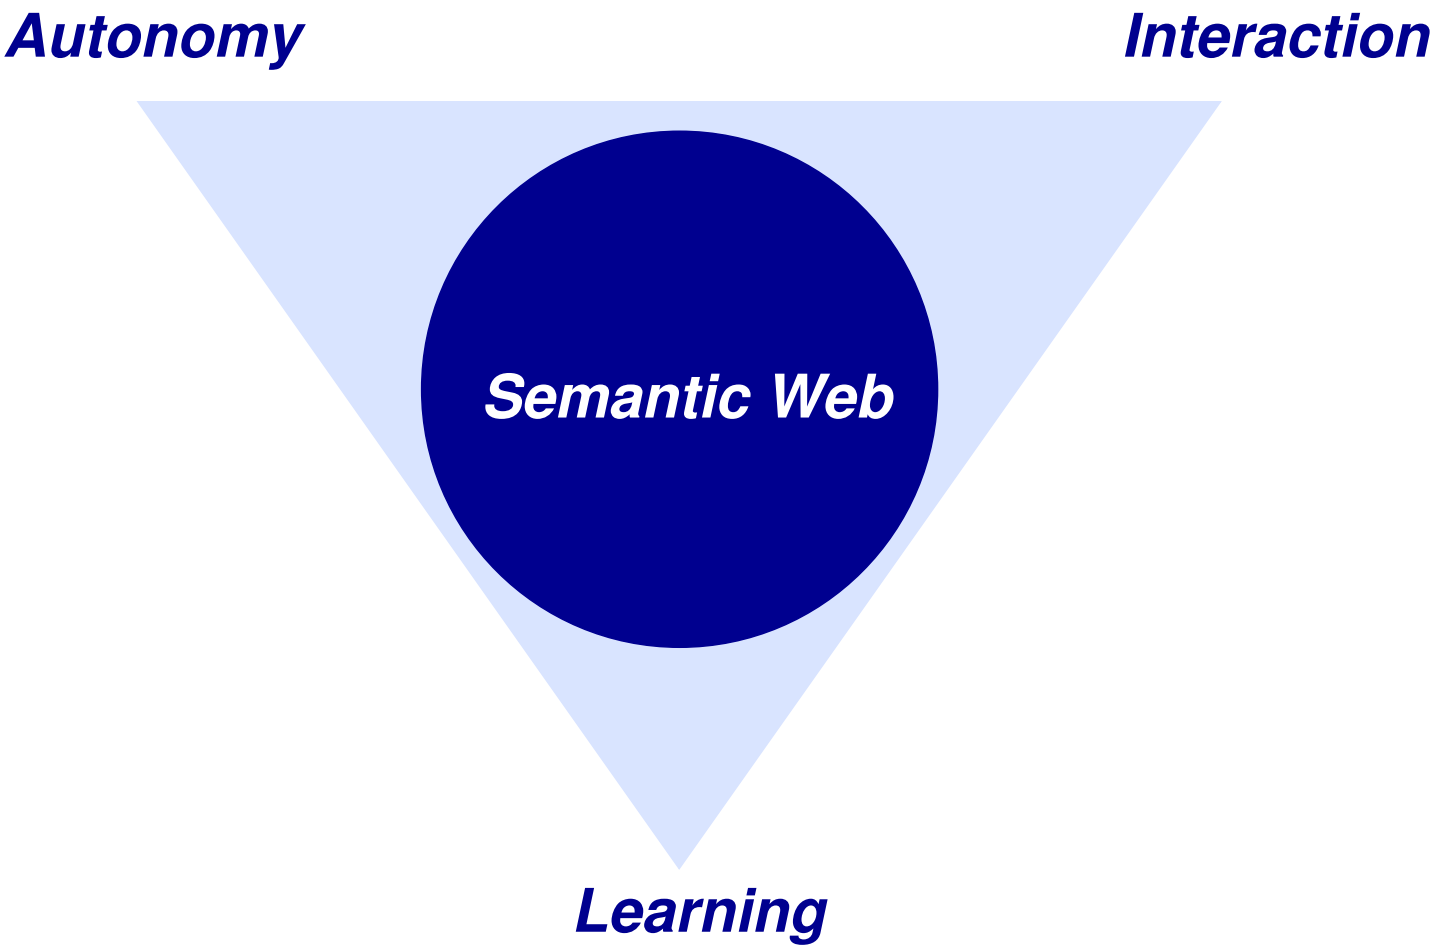

The Semantic Web – A Great Challenge

- Goal: to improve information access via machine-processable meta-data

The Semantic Web – A Great Challenge

- Goal: to improve information access via machine-processable meta-data
- Primarily concerned with
  - Standards for meta-data representation (XML, RDF, UDDI, SOAP, WSDL, etc.),
  - Ontologies (DAML+OIL)

The Semantic Web – A Great Challenge

- Goal: to improve information access via machine-processable meta-data
- Primarily concerned with
  - Standards for meta-data representation (XML, RDF, UDDI, SOAP, WSDL, etc.),
  - Ontologies (DAML+OIL)
- Meta-data facilitates
  - search (for information/services)
  - logical inference/knowledge discovery
  - web mining/clustering/classification

The Semantic Web – A Great Challenge

- Goal: to improve information access via machine-processable meta-data
- Primarily concerned with
  - Standards for meta-data representation (XML, RDF, UDDI, SOAP, WSDL, etc.),
  - Ontologies (DAML+OIL)
- Meta-data facilitates
  - search (for information/services)
  - logical inference/knowledge discovery
  - web mining/clustering/classification
- Problem: Web still regarded as a static source of information in the form of text/multimedia content

The Semantic Web – A Great Challenge

- Goal: to improve information access via machine-processable meta-data
- Primarily concerned with
  - Standards for meta-data representation (XML, RDF, UDDI, SOAP, WSDL, etc.),
  - Ontologies (DAML+OIL)
- Meta-data facilitates
  - search (for information/services)
  - logical inference/knowledge discovery
  - web mining/clustering/classification
- Problem: Web still regarded as a static source of information in the form of text/multimedia content
- Focusing on the **combination** of autonomy, interaction and learning opens entirely new prospects for the SW!

Autonomy

- Essential for delegation of complex tasks to agents

Autonomy

- Essential for delegation of complex tasks to agents
- Still poorly understood, but important
  (e.g. IBMs Autonomic Computing Initiative)

Autonomy

- Essential for delegation of complex tasks to agents
- Still poorly understood, but important
  (e.g. IBMs Autonomic Computing Initiative)
- Our view:

	perspective	
range	internal	external
performative	*capability*	*dependency*
deliberative	*motivation*	*control*
normative	*commitment*	*expectation*

Autonomy

- Essential for delegation of complex tasks to agents
- Still poorly understood, but important
  (e.g. IBMs Autonomic Computing Initiative)
- Our view:

	perspective	
range	internal	external
performative	*capability*	*dependency*
deliberative	*motivation*	*control*
normative	*commitment*	*expectation*

- Issues:
  - How much autonomy is desirable?
  - At what level should it be specified?
  - How do we respond to other agents' autonomy?

Interaction

- The Web obtains its "semantics" only through the ways it is used
  - textual/visual content: interaction of reader/viewer with author
  - services: search, matchmaking, brokering, access (e.g. Web services)
  - markets: negotiation, contracting, financial transactions
  - forums: discussions, blackboards, collaborative authoring
  - organisations: knowledge management, process management

Interaction

- The Web obtains its "semantics" only through the ways it is used
  - textual/visual content: interaction of reader/viewer with author
  - services: search, matchmaking, brokering, access (e.g. Web services)
  - markets: negotiation, contracting, financial transactions
  - forums: discussions, blackboards, collaborative authoring
  - organisations: knowledge management, process management
- "Meaning" = the significance of communicative actions in the context of human/agent activity

Interaction

- The Web obtains its "semantics" only through the ways it is used
  - textual/visual content: interaction of reader/viewer with author
  - services: search, matchmaking, brokering, access (e.g. Web services)
  - markets: negotiation, contracting, financial transactions
  - forums: discussions, blackboards, collaborative authoring
  - organisations: knowledge management, process management
- "Meaning" = the significance of communicative actions in the context of human/agent activity
- Ongoing process of construction of meaning rather than pre-defined semantics

Learning

- Agent autonomy and complex environments induce unpredictability

Learning

- Agent autonomy and complex environments induce unpredictability
- Impossible to pre-specify system behaviour

Learning

- Agent autonomy and complex environments induce unpredictability
- Impossible to pre-specify system behaviour
- Acquisition of knowledge from observation may
  - serve as a minimal-assumption approach
  - ensure that system dynamics is taken into account
  - provide information for long-term optimisation (either from human designer or agent perspective)

Learning

- Agent autonomy and complex environments induce unpredictability
- Impossible to pre-specify system behaviour
- Acquisition of knowledge from observation may
  - serve as a minimal-assumption approach
  - ensure that system dynamics is taken into account
  - provide information for long-term optimisation (either from human designer or agent perspective)
- Focus: **communication learning**
  - modelling and learning of interaction structures
  - evolutionary semantics of agent communication
  - autonomy-respecting intervention through strategic communication

Example Applications

- Electronic Travel Assistant on a PDA
  - Books tickets, provides directions, displays sightseeing information retrieved from the Web
  - How much **autonomy** should it be endowed with?

Example Applications

- Electronic Travel Assistant on a PDA
  - Books tickets, provides directions, displays sightseeing information retrieved from the Web
  - How much **autonomy** should it be endowed with?
- Biological eScience Forum
  - Platform for negotiating access to biological data (industry/academia mix)
  - How to **interact** strategically given predominant communicative conventions?

Example Applications

- Electronic Travel Assistant on a PDA
  - Books tickets, provides directions, displays sightseeing information retrieved from the Web
  - How much **autonomy** should it be endowed with?
- Biological eScience Forum
  - Platform for negotiating access to biological data (industry/academia mix)
  - How to **interact** strategically given predominant communicative conventions?
- Electronic auctions
  - "Buyers who won't pay" problem, what if distrust reaches dangerous level?
  - **Learning** how to identify critical situations, how to intervene

Current Research and Future Challenges

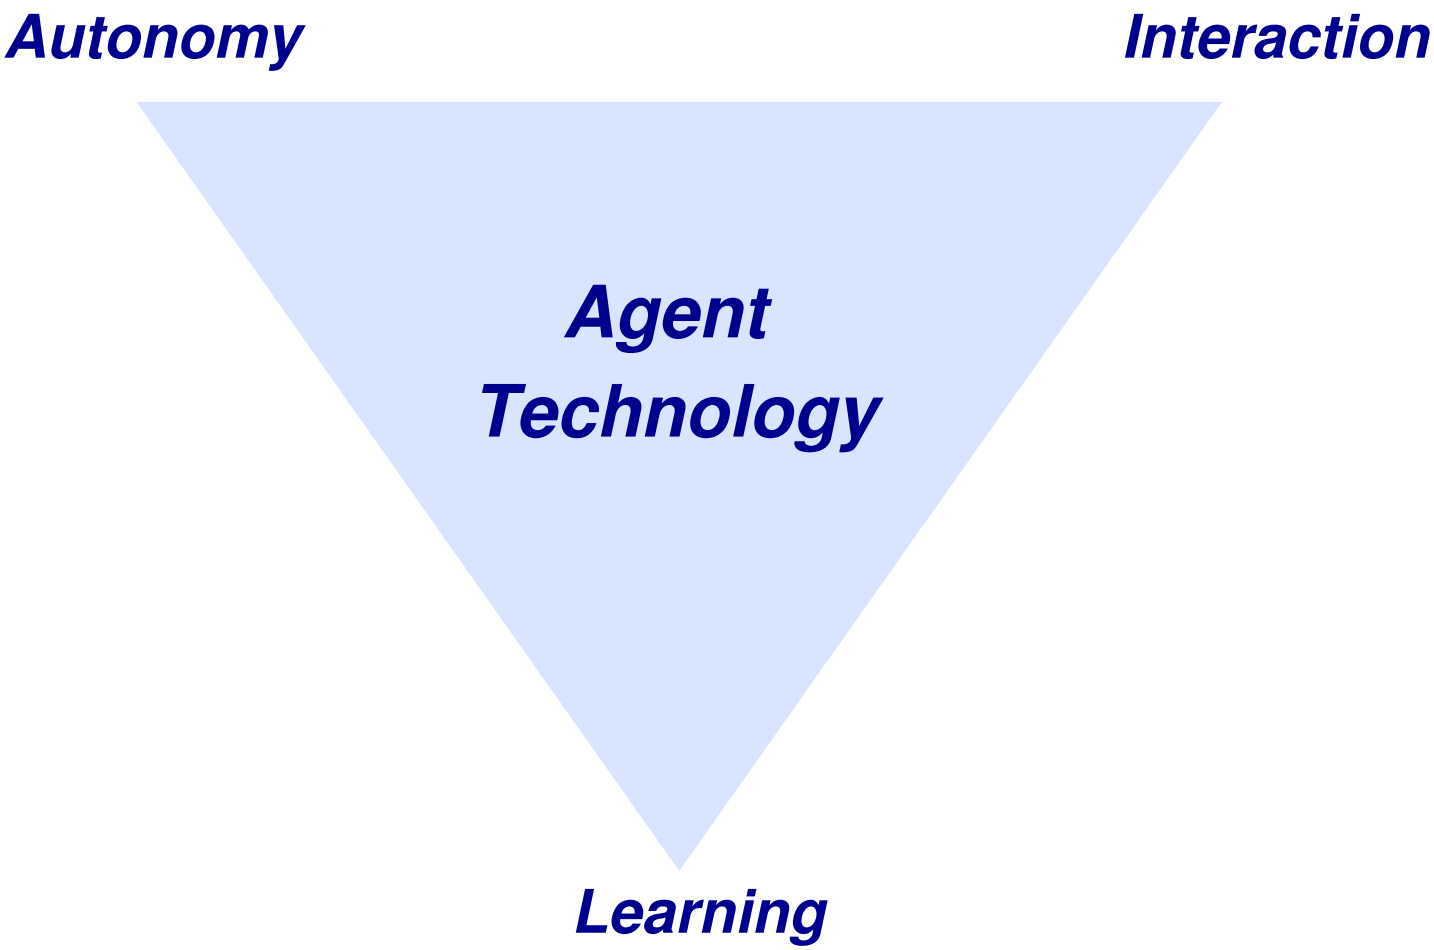

Current Research and Future Challenges

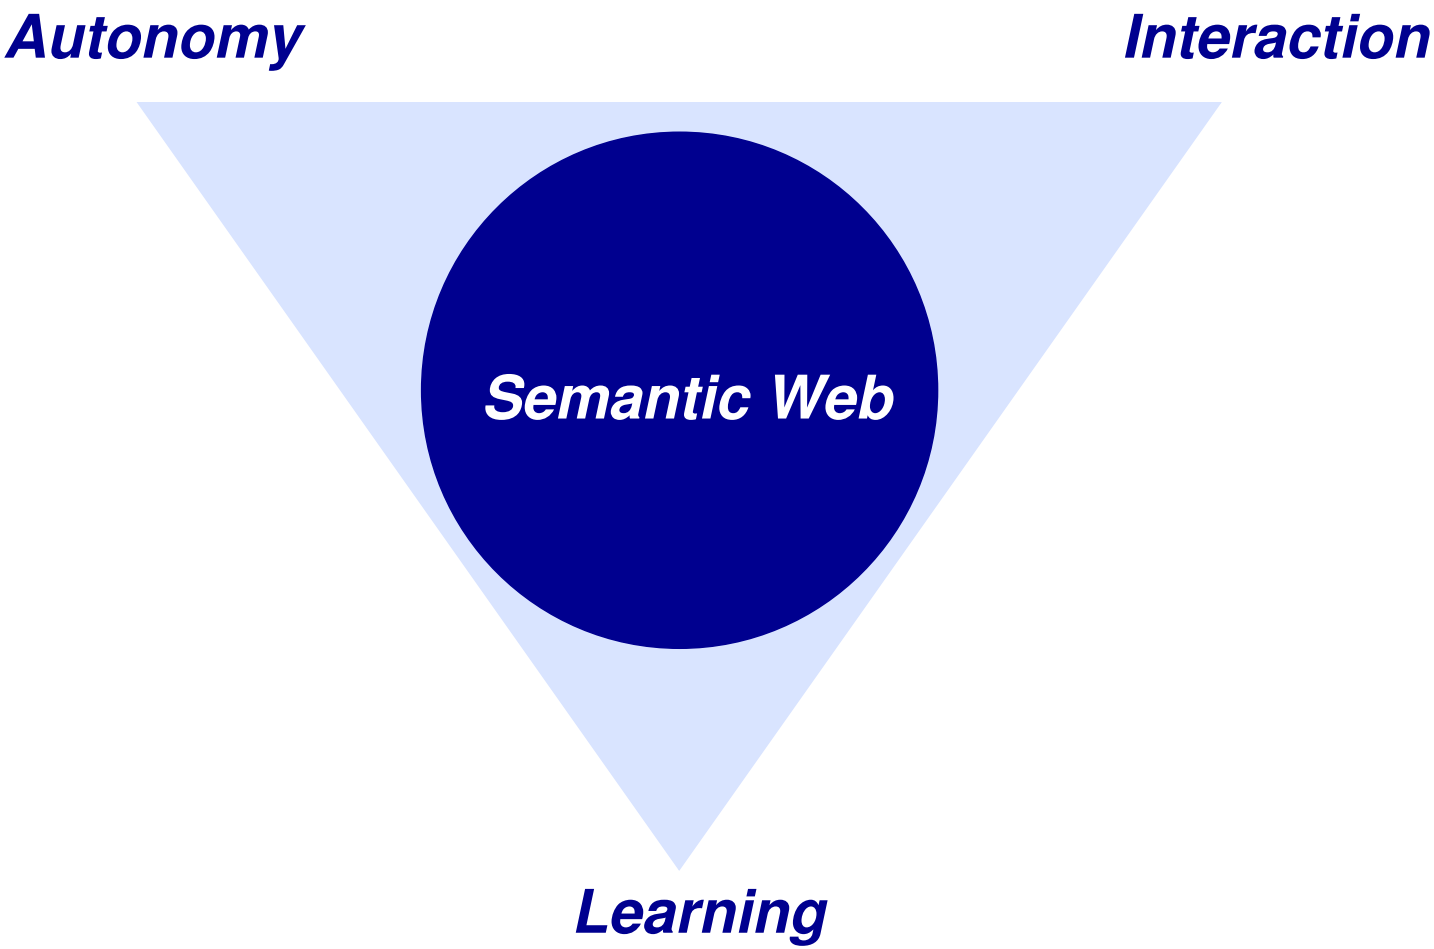

Autonomy: RNS – Roles, Norms and Sanctions

- Example: basic activity

ACT deliver (*material*,*quantity*)
{ STATUS RANGE
<IND> : NORM <**P**> <NO> + SANC <NO> <NO>
<DEP EACH> : NORM <**O**> <*quantity* $\leq$ 100> + SANC <**PU**> <withdraw_role>
<DEP AssemblyMg> : NORM <**I**> <*material* = steel> + SANC <**PU**> <pay_fine>
}

- Example: request activity

ACT REQUEST **(** EACH **;** USsupplier, EUROsupplier **;** NOT deliver (*material*, *quantity*) **)**
{ STATUS RANGE
<IND> : NORM <**P**> < (*material* = steel) AND (rating(*material*) = poor)> +
SANC <NO> <NO>
NORMATIVE IMPACT
NORM <**I**> <*material* = steel>
}

Interaction: The Communication Systems Approach

Formalism for describing structure and evolution of interaction practices using **expectation networks**

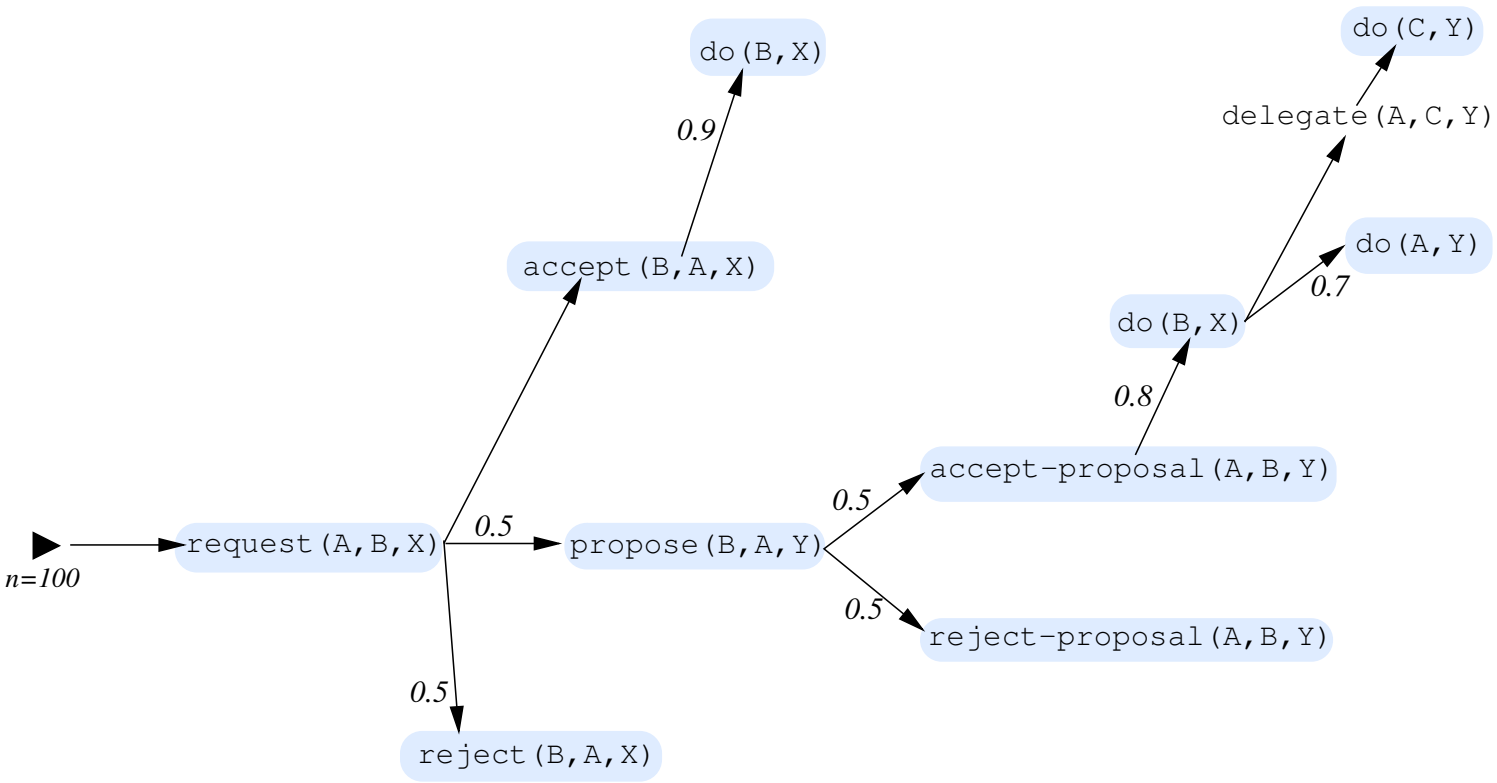

Interaction: The Communication Systems Approach

Formalism for describing structure and evolution of interaction practices using **expectation networks**

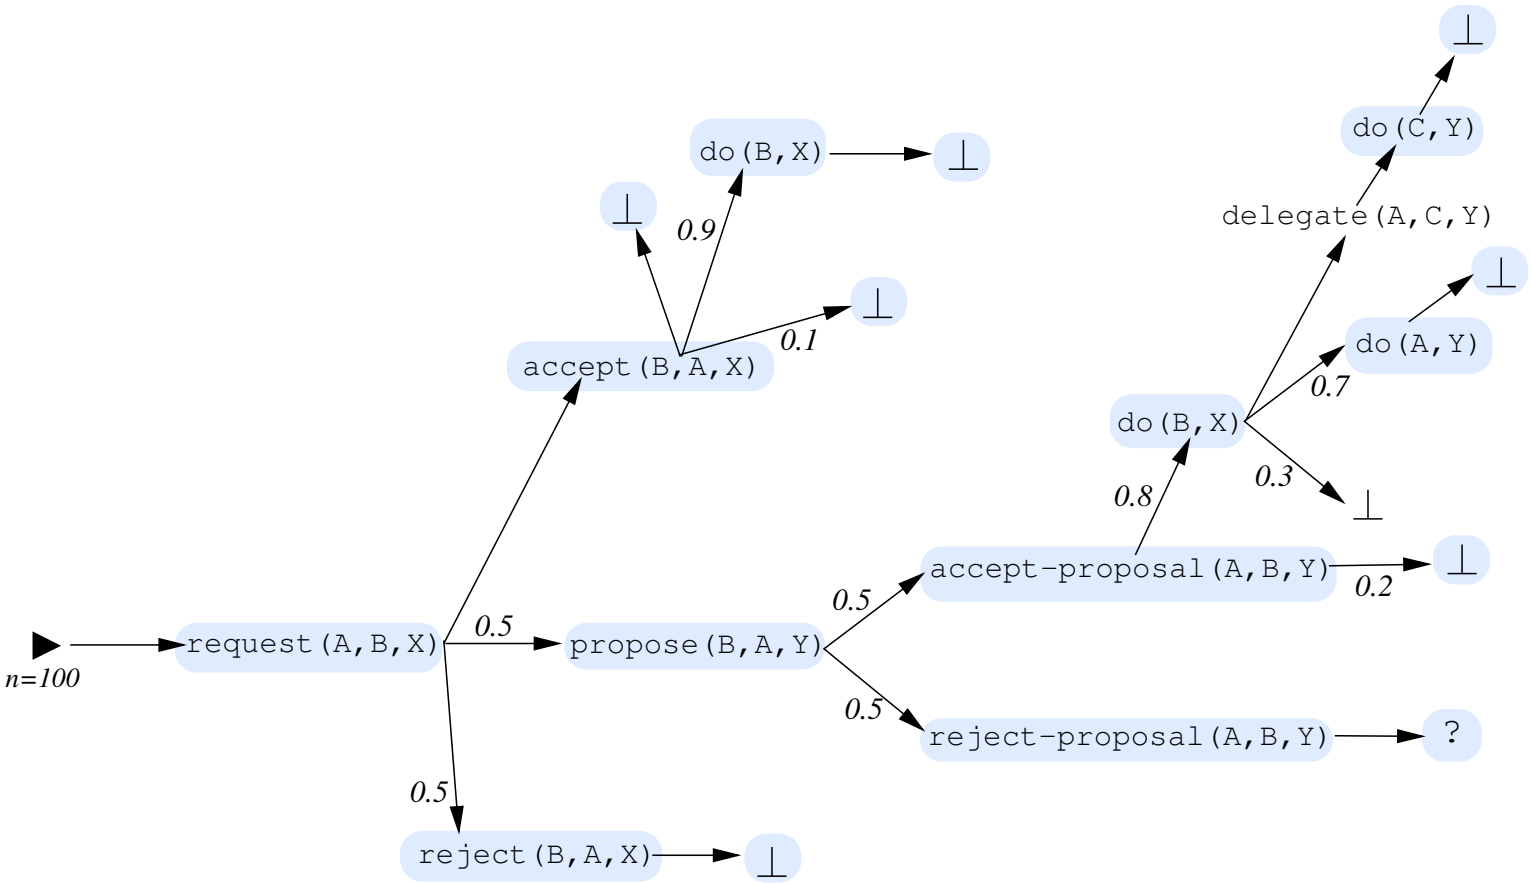

Interaction: The Communication Systems Approach

Formalism for describing structure and evolution of interaction practices using **expectation networks**

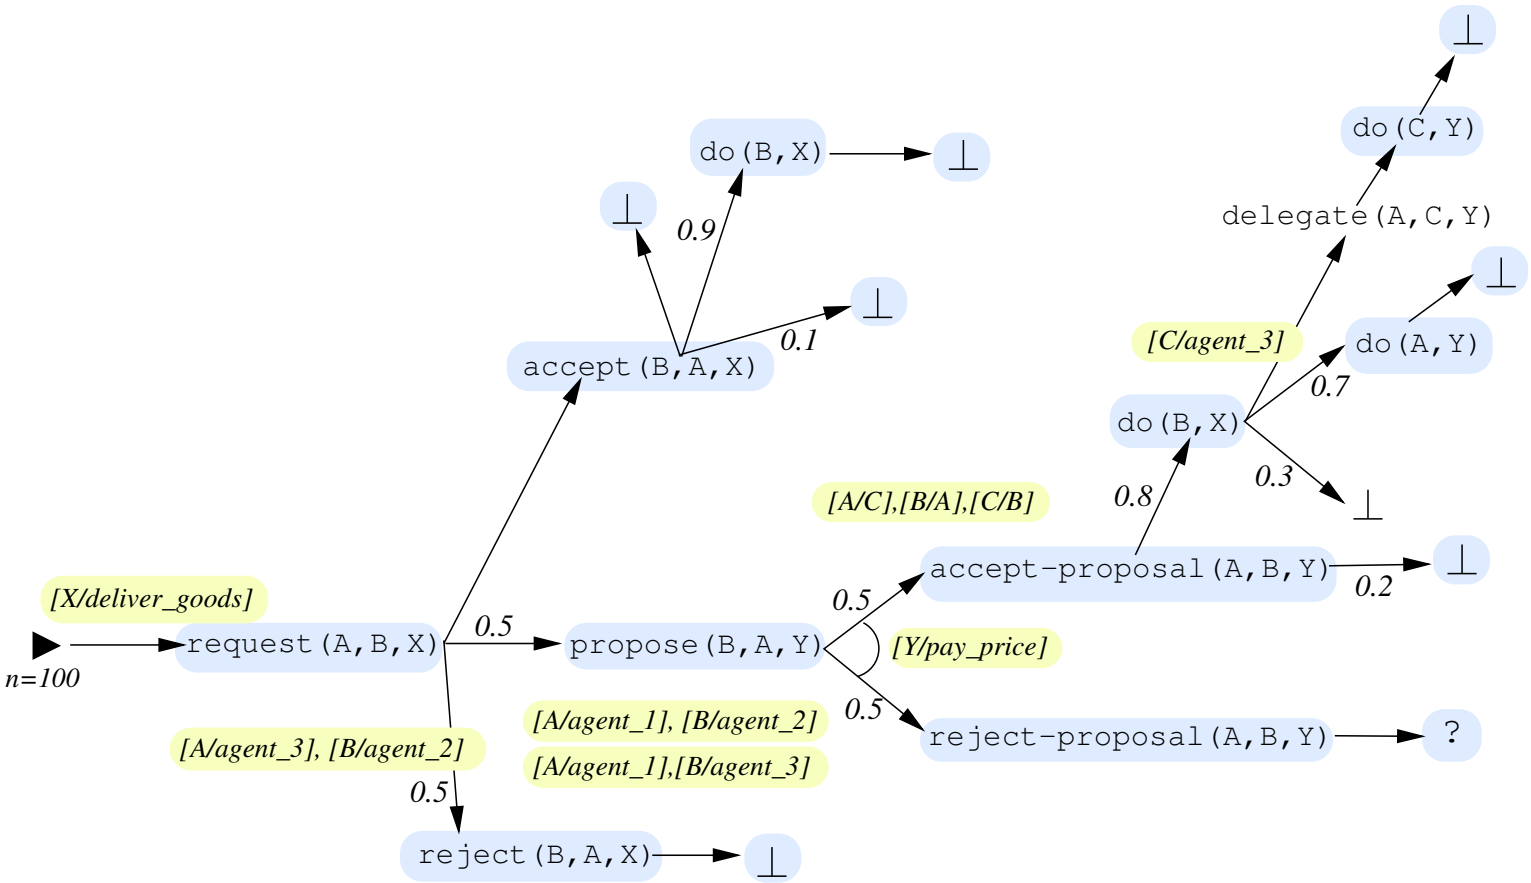

Interaction: The Communication Systems Approach

Formalism for describing structure and evolution of interaction practices using **expectation networks**

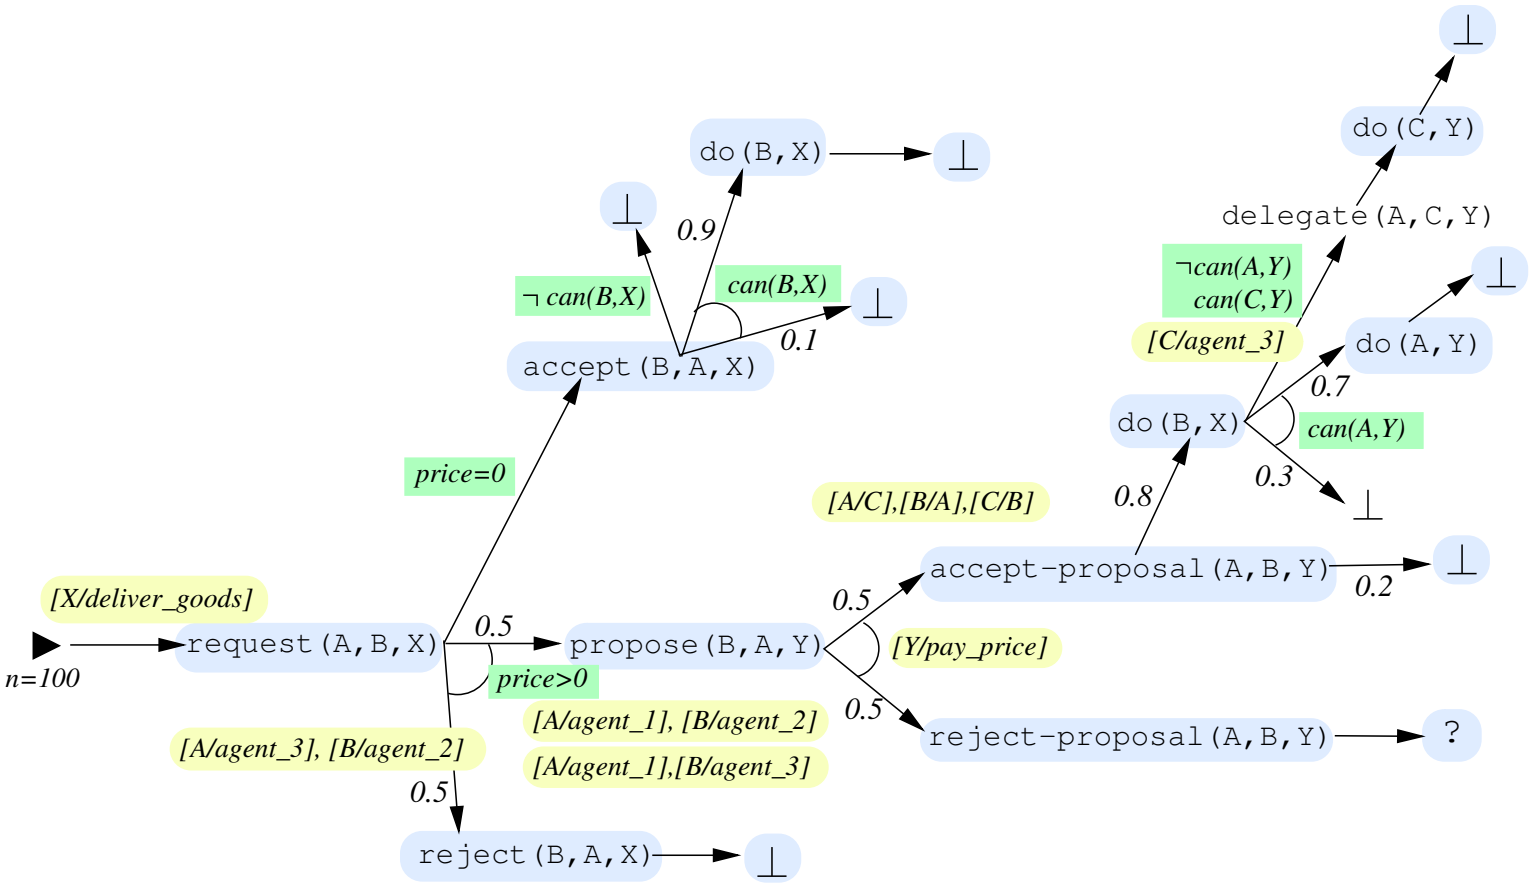

Interaction: The Communication Systems Approach

Advantages over other models of communication semantics:

Interaction: The Communication Systems Approach

Advantages over other models of communication semantics:

- No mentalistic assumptions

Interaction: The Communication Systems Approach

Advantages over other models of communication semantics:

- No mentalistic assumptions
- Can be used by agents/system observers

Interaction: The Communication Systems Approach

Advantages over other models of communication semantics:

- No mentalistic assumptions
- Can be used by agents/system observers
- Allows for context-sensitivity and uncertainty

Interaction: The Communication Systems Approach

Advantages over other models of communication semantics:

- No mentalistic assumptions
- Can be used by agents/system observers
- Allows for context-sensitivity and uncertainty
- Captures the dynamics of evolving meaning

Learning: Interaction Frames/InFFrA

- **In**teraction **F**rames and **Fr**aming **A**rchitecture

Learning: Interaction Frames/InFFrA

- **In**teraction **F**rames and **Fr**aming **A**rchitecture
- **Frame** = model of a class of interactions
  - courses of interactions
  - roles of participants
  - context conditions
  - agent beliefs

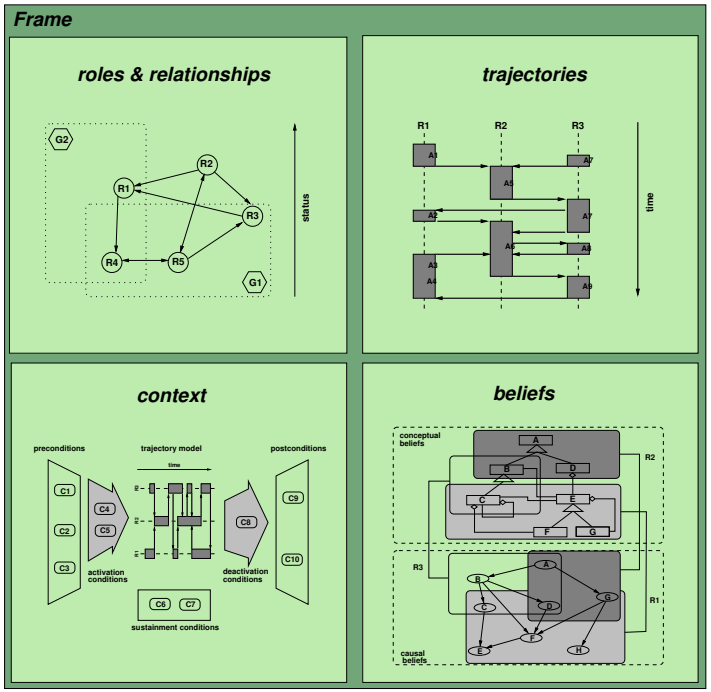

Learning: Interaction Frames/InFFrA

- **In**teraction **F**rames and **Fr**aming **A**rchitecture
- **Frame** = model of a class of interactions
  - courses of interactions
  - roles of participants
  - context conditions
  - agent beliefs
- **Framing** = process of constructing and applying frames

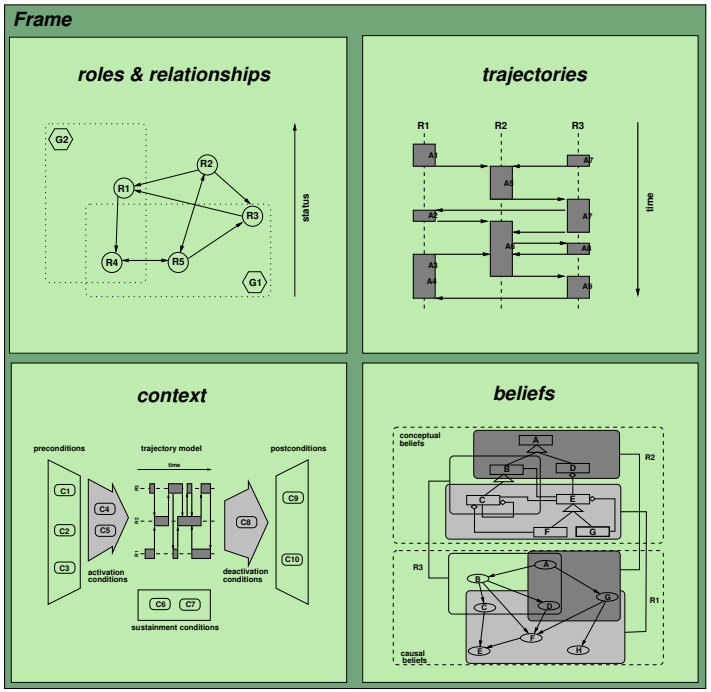

Learning: Interaction Frames/InFFrA

- **In**teraction **F**rames and **Fr**aming **A**rchitecture
- **Frame** = model of a class of interactions
  - courses of interactions
  - roles of participants
  - context conditions
  - agent beliefs
- **Framing** = process of constructing and applying frames
- Specific architecture for reasoning about communication systems at micro-conversation level

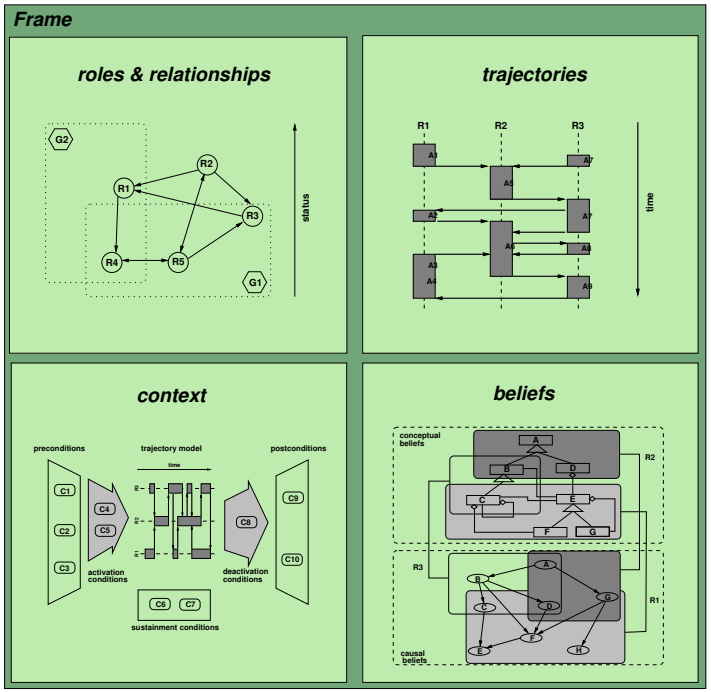

InFFrA Reasoning Process

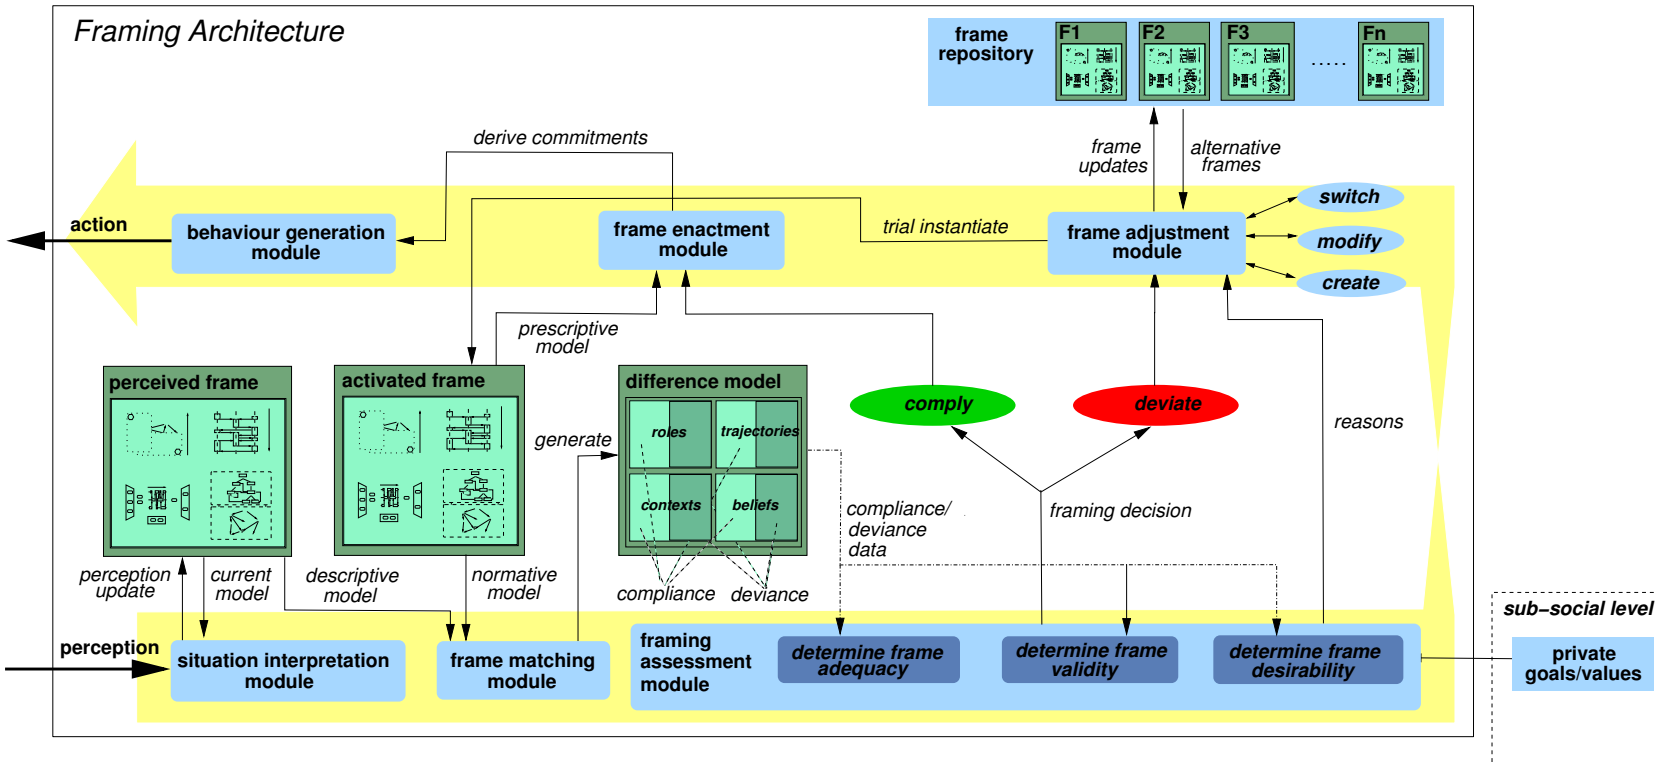

InFFrA Reasoning Process

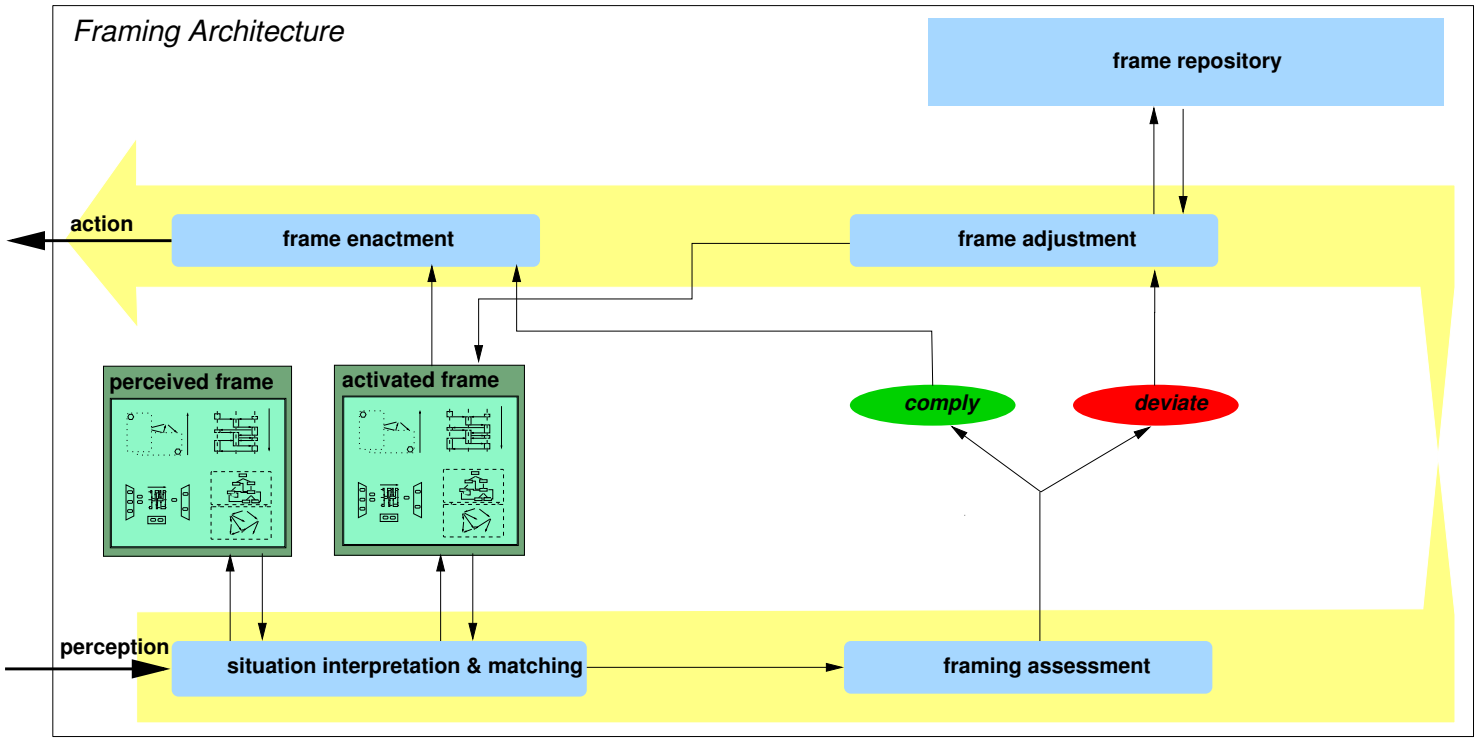

Hierarchical Reinforcement Learning View

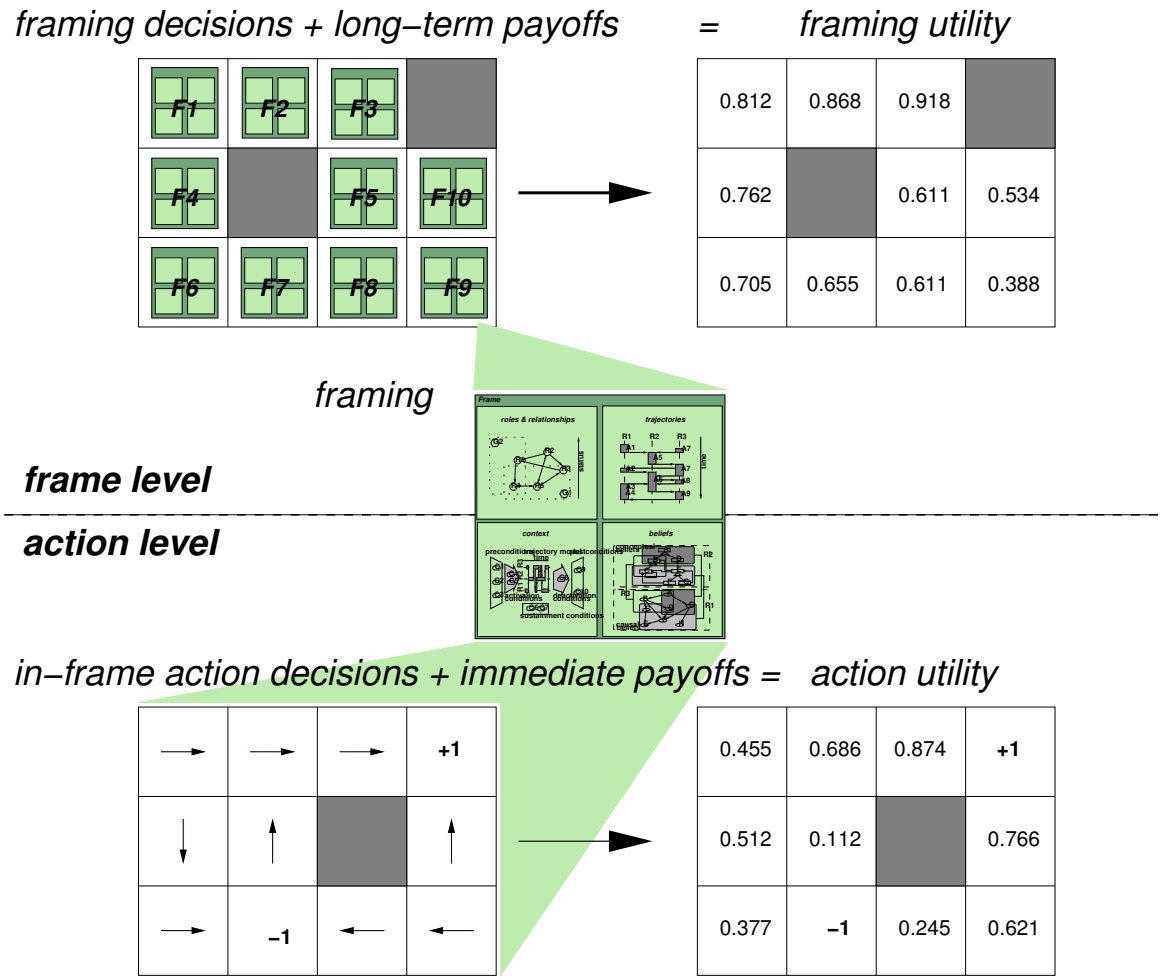

Challenges for Future Research: Autonomy

- Methods for specifying and dealing with autonomy
  - norms and conventions
  - trust and reputation
  - authentication and security
  - autonomy-respecting intervention

Challenges for Future Research: Autonomy

- Methods for specifying and dealing with autonomy
  - norms and conventions
  - trust and reputation
  - authentication and security
  - autonomy-respecting intervention
- Autonomy and mobility
  - lack of theory in context-aware and ubiquitous computing
  - particular relationship between user and digital assistant autonomy
  - virtual co-presence, awareness, **social** context

Challenges for Future Research: Interaction

- Evolving and negotiated ontologies
  - emergent semantics & pragmatics
  - integration with a priori semantics
  - conversation mining and communication process modelling
  - relationship between NLP and agent communication (bringing Web content and agent interaction together)

Challenges for Future Research: Interaction

- Evolving and negotiated ontologies
  - emergent semantics & pragmatics
  - integration with a priori semantics
  - conversation mining and communication process modelling
  - relationship between NLP and agent communication (bringing Web content and agent interaction together)
- Integration of different autonomous components in agent and system architectures (agent-oriented software engineering)

Challenges for Future Research: Learning

- Learning how to communicate
  - Building models of social communication systems
  - Active learning of communication strategies

Challenges for Future Research: Learning

- Learning how to communicate
  - Building models of social communication systems
  - Active learning of communication strategies
- Issues:
  - How to derive appropriate (manageable) state and policy abstractions
  - Combination of MDPs, interaction protocols and knowledge-based inference
  - Merging global and local views of communication systems

Conclusions

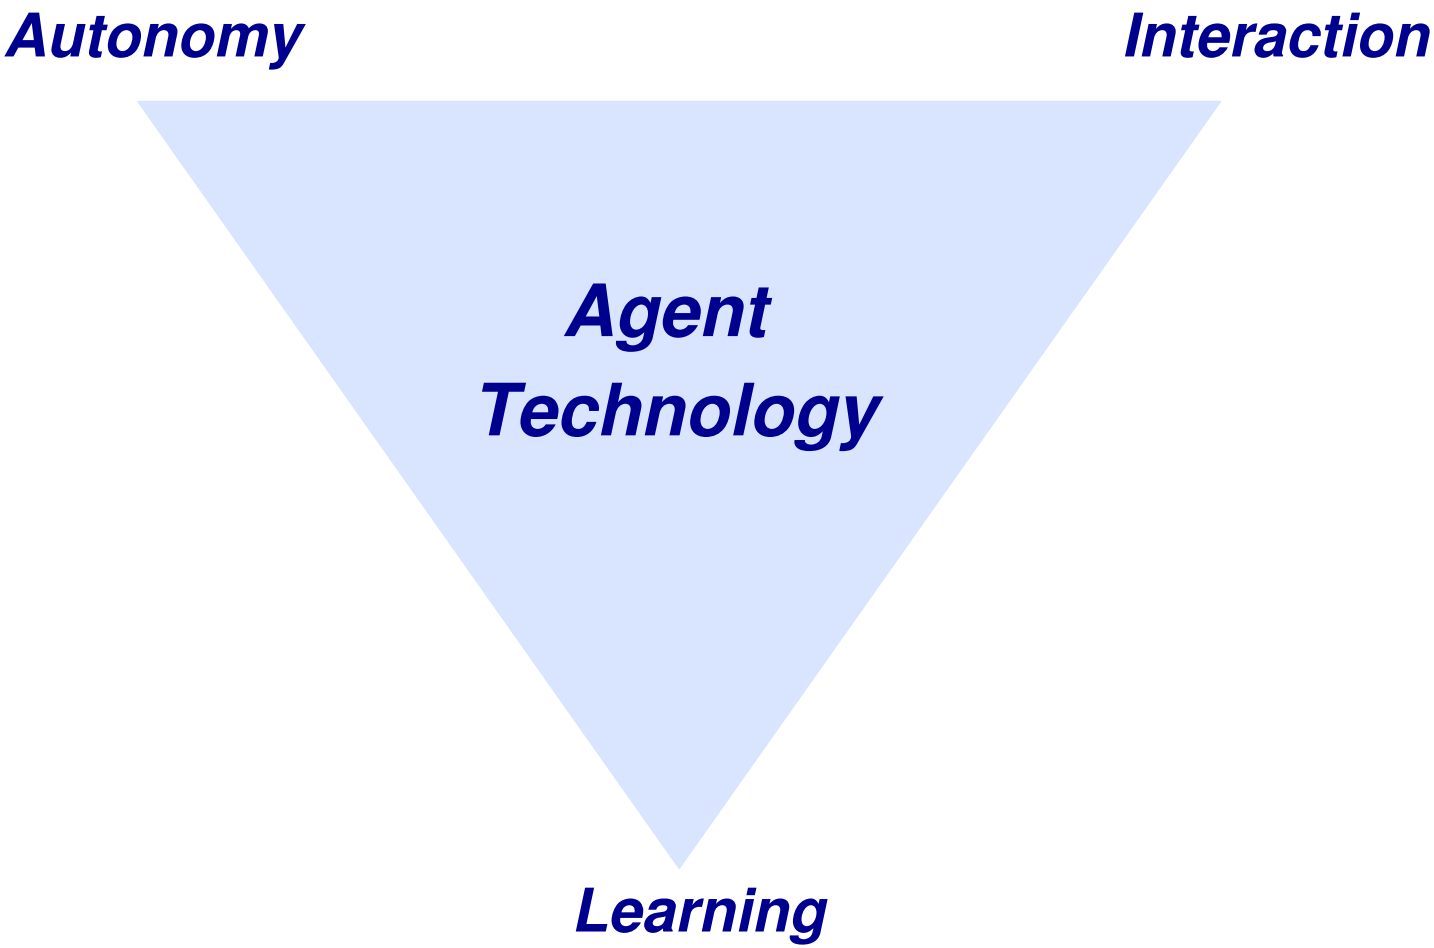

Conclusions

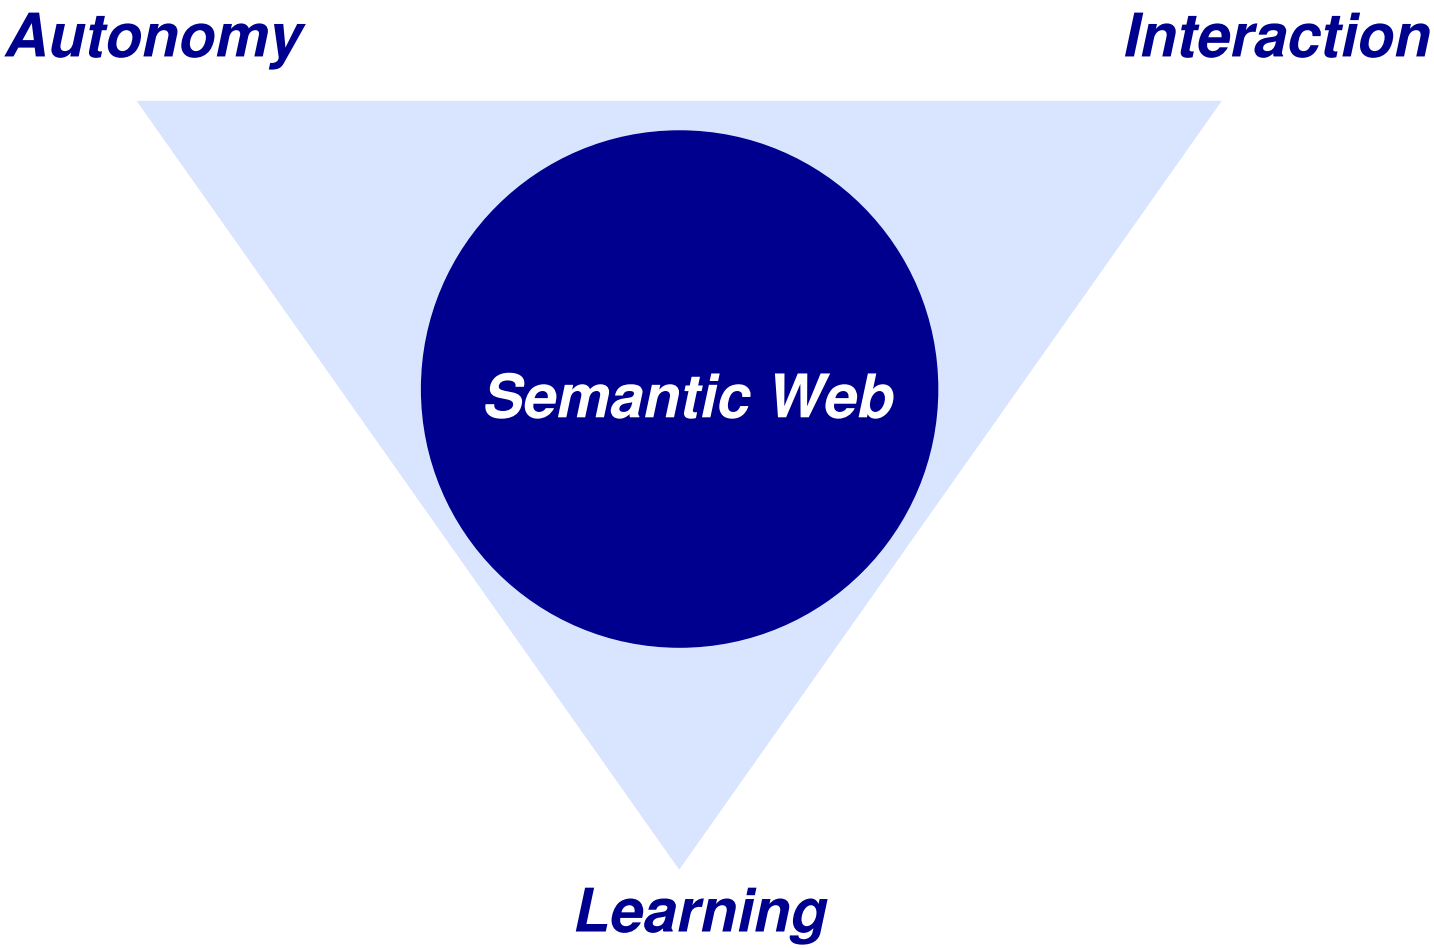

Conclusions

- Made the case for exploiting agent **autonomy**, **interaction** and **learning** to advance Semantic Web technologies

Conclusions

- Made the case for exploiting agent **autonomy**, **interaction** and **learning** to advance Semantic Web technologies
- **Anthropocentric** view: the Web is made of how people use it (and not of HTML data)

Conclusions

- Made the case for exploiting agent **autonomy**, **interaction** and **learning** to advance Semantic Web technologies
- **Anthropocentric** view: the Web is made of how people use it (and not of HTML data)
- Ongoing agent research not sufficiently adapted to the needs of the Semantic Web and vice versa

Conclusions

- Made the case for exploiting agent **autonomy**, **interaction** and **learning** to advance Semantic Web technologies
- **Anthropocentric** view: the Web is made of how people use it (and not of HTML data)
- Ongoing agent research not sufficiently adapted to the needs of the Semantic Web and vice versa
- Research proposed here may help avoid:
  - The Web becoming a huge "information wasteland"
  - Agent technology degenerating to a "nice idea" for lack of killer apps

"The World is not the Web"

Variety of other relevant application areas:

"The World is not the Web"

Variety of other relevant application areas:

- Agents for Context-Aware Computing (lack of theory)

"The World is not the Web"

Variety of other relevant application areas:

- Agents for Context-Aware Computing (lack of theory)
- Autonomic Computing (lack of open systems view)

"The World is not the Web"

Variety of other relevant application areas:

- Agents for Context-Aware Computing (lack of theory)
- Autonomic Computing (lack of open systems view)
- Agent-Supported Peer-to-Peer Computing

"The World is not the Web"

Variety of other relevant application areas:

- Agents for Context-Aware Computing (lack of theory)
- Autonomic Computing (lack of open systems view)
- Agent-Supported Peer-to-Peer Computing
- Agents in Bioinformatics (algorithm level!)

"The World is not the Web"

Variety of other relevant application areas:

- Agents for Context-Aware Computing (lack of theory)
- Autonomic Computing (lack of open systems view)
- Agent-Supported Peer-to-Peer Computing
- Agents in Bioinformatics (algorithm level!)
- Knowledge Management/Business Process Management (Agent-Supported CSCW, Enterprise Application Collaboration, eLearning, eScience)

"The World is not the Web"

Variety of other relevant application areas:

- Agents for Context-Aware Computing (lack of theory)
- Autonomic Computing (lack of open systems view)
- Agent-Supported Peer-to-Peer Computing
- Agents in Bioinformatics (algorithm level!)
- Knowledge Management/Business Process Management (Agent-Supported CSCW, Enterprise Application Collaboration, eLearning, eScience)
- Agent based Anti-Spam/Anti-Fraud Systems

Thank you for your attention!

Criticism/Misconceptions

Agent concept: overtly general, highly abstract, vague

Criticism/Misconceptions

Agent concept: overtly general, highly abstract, vague

- Widely misused as a buzzword
  - eBay bidding “agents”, mail transport “agent”, etc.

Criticism/Misconceptions

Agent concept: overtly general, highly abstract, vague

- Widely misused as a buzzword
  - eBay bidding "agents", mail transport "agent", etc.
- Ongoing theoretical dispute
  - Is a thermostat/word processor/Google an agent?
  - How about ALife/ant algorithms/GAs?
  - How about robots?

Criticism/Misconceptions

Agent concept: overtly general, highly abstract, vague

- Widely misused as a buzzword
  - eBay bidding "agents", mail transport "agent", etc.
- Ongoing theoretical dispute
  - Is a thermostat/word processor/Google an agent?
  - How about ALife/ant algorithms/GAs?
  - How about robots?
- Necessity to distinguish between
  - design metaphor
  - set of methods, technologies and tools

Criticism/Misconceptions

Agent concept: overtly general, highly abstract, vague

- Widely misused as a buzzword
  - eBay bidding "agents", mail transport "agent", etc.
- Ongoing theoretical dispute
  - Is a thermostat/word processor/Google an agent?
  - How about ALife/ant algorithms/GAs?
  - How about robots?
- Necessity to distinguish between
  - design metaphor
  - set of methods, technologies and tools
- Chasm between concept of agent autonomy and classical engineering stance

Example: Intelligent Travel Assistant

- An agent on a PDA that uses Web information to assist a human traveller by ...
  - making sightseeing suggestions
  - booking tickets for attractions
  - displaying travel guide material
  - providing directions
  - planning itineraries

Example: Intelligent Travel Assistant

- An agent on a PDA that uses Web information to assist a human traveller by ...
  - making sightseeing suggestions
  - booking tickets for attractions
  - displaying travel guide material
  - providing directions
  - planning itineraries
- Should it be allowed to ...
  - purchase tickets without user approval?
  - suggest alternative routes according to his internal travel agenda or seek constant user feedback?
  - pro-actively negotiate with other agents (e.g. travel agents) online?

Example: Biological eScience Forum

- Platform for negotiating access to biological data and algorithms
  - Mixture of industry and academic labs
  - Reciprocal exchange and commercial trading

Example: Biological eScience Forum

- Platform for negotiating access to biological data and algorithms
  - Mixture of industry and academic labs
  - Reciprocal exchange and commercial trading
- Different protocols for automated
  - matchmaking (service directories, middle agents)
  - negotiation (auctions, argumentation)
  - reputation management (recommender agents)
  - contracting (e.g. trusted third parties)

Example: Biological eScience Forum

- Platform for negotiating access to biological data and algorithms
  - Mixture of industry and academic labs
  - Reciprocal exchange and commercial trading
- Different protocols for automated
  - matchmaking (service directories, middle agents)
  - negotiation (auctions, argumentation)
  - reputation management (recommender agents)
  - contracting (e.g. trusted third parties)
- Evolving culture of communicative conventions, e.g.
  - reputation of certain labs requires specific strategies
  - chasm or gradual convergence of academia/industry cultures
  - multidisciplinary cooperation (biologists, doctors, computing people)

Example: Electronic Auctions

- "Buyers who won't pay" problem

Example: Electronic Auctions

- "Buyers who won't pay" problem
- eBay offers statistics of customer satisfaction with buyers/sellers and displays comment logs

Example: Electronic Auctions

- "Buyers who won't pay" problem
- eBay offers statistics of customer satisfaction with buyers/sellers and displays comment logs
- What if global level of distrust reaches dangerous level?

Example: Electronic Auctions

- "Buyers who won't pay" problem
- eBay offers statistics of customer satisfaction with buyers/sellers and displays comment logs
- What if global level of distrust reaches dangerous level?
- Necessity of
  - prompt identification of such deviance from normative rules of behaviour (constant monitoring and data analysis)
  - "soft" methods of intervention (e.g. selective filtering of comments, providing incentives for cooperative behaviour, appealing to other institutions)

Application Scenario: Web Linkage Negotiations