Dimensions

- Research proceeds by making simplifying assumptions, and gradually reducing them.
- Each simplifying assumption gives a dimension of complexity
  - Can be multiple values in a dimension: values go from simple to complex
  - Simplifying assumptions can be relaxed in various combinations
- Much of the history of AI can be seen as starting from the simple and adding in complexity in some of these dimensions.

Dimensions of Complexity

- Flat or hierarchical
- Explicit states or features or objects and relations
- Static or finite stage or indefinite stage or infinite stage
- Fully observable or partially observable
- Deterministic or stochastic dynamics
- Goals or complex preferences
- Single-agent or multiple agents
- Knowledge is given or knowledge is learned from experience
- Perfect rationality or bounded rationality

Modularity

- You can model the system at one level of abstraction: flat
- You can model the system at multiple levels of abstraction: hierarchical
- Example: Planning a trip from here to a resort in Cancun, Mexico
- Flat representations are adequate for simple systems, but complex biological systems, computer systems, organizations are all hierarchical
- A flat description is either continuous or discrete. Hierarchical reasoning is often a hybrid of continuous and discrete.

Succinctness and Expressiveness

Much of modern AI is about finding compact representations and exploiting that compactness for computational gains.
A agent can reason in terms of:

- explicit states — a state is one way the world could be

Succinctness and Expressiveness

Much of modern AI is about finding compact representations and exploiting that compactness for computational gains.
A agent can reason in terms of:

- explicit states — a state is one way the world could be
- features or propositions.
  - It's often more natural to describe states in terms of features.
  - 30 binary features can represent $2^{30} = 1,073,741,824$ states.

Succinctness and Expressiveness

Much of modern AI is about finding compact representations and exploiting that compactness for computational gains.
A agent can reason in terms of:

- explicit states — a state is one way the world could be
- features or propositions.
  - It's often more natural to describe states in terms of features.
  - 30 binary features can represent $2^{30} = 1,073,741,824$ states.
- individuals and relations
  - There is a feature for each relationship on each tuple of individuals.
  - Often we can reason without knowing the individuals or when there are infinitely many individuals.

Planning horizon

...how far the agent looks into the future when deciding what to do.

- Static: world does not change
- Finite stage: agent reasons about a fixed finite number of time steps
- Indefinite stage: agent is reasoning about finite, but not predetermined, number of time steps
- Infinite stage: the agent plans for going on forever (process oriented)

Uncertainty

There are two dimensions for uncertainty. In each dimension we can have

- no uncertainty: the agent knows which world is true
- disjunctive uncertainty: there is a set of worlds that are possible
- probabilistic uncertainty: a probability distribution over the worlds.

Sensing Uncertainty

What the agent can determine the state from the observations:

- Fully-observable: the agent knows the state of the world from the observations.
- Partially-observable: there can be many states that are possible given an observation.

Uncertain dynamics

If the agent knew the initial state and the action, could it predict the resulting state?
The dynamics can be:

- Deterministic: the state resulting from carrying out an action in state is determined from the action and the state
- Stochastic: there is uncertainty over the states resulting from executing a given action in a given state.

Goals or complex preferences

- achievement goal is a goal to achieve. This can be a complex logical formula.
- complex preferences that may involve tradeoffs between various desiderata, perhaps at different times.
  Either ordinal or cardinal (e.g., utility)
- Examples: coffee delivery robot, medical doctor

Single agent or multiple agents

- Single agent reasoning is where an agent assumes that any other agents are part of the environment.
- Multiple agent reasoning is when an agent needs to reason strategically about the reasoning of other agents.

Agents can have their own goals: cooperative, competitive, or goals can be independent of each other

Learning from experience

Whether the model is fully specified a priori:

- knowledge is given
- knowledge is learned from data or past experience

Perfect rationality or bounded rationality

- Perfect rationality: the agent can determine the best course of action, without taking into account its limited computational resources.
- Bounded rationality: the agent mast make good decisions based on its perceptual, computational and memory limitations.

Dimensions of Complexity

- Flat or hierarchical
- Explicit states or features or objects and relations
- Static or finite stage or indefinite stage or infinite stage
- Fully observable or partially observable
- Deterministic or stochastic dynamics
- Goals or complex preferences
- Single-agent or multiple agents
- Knowledge is given or knowledge is learned from experience
- Perfect rationality or bounded rationality

State-space Search

- flat or hierarchical
- explicit states or features or objects and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic actions
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or learned
- perfect rationality or bounded rationality

Classical Planning

- flat or hierarchical
- explicit states or features or objects and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic actions
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or learned
- perfect rationality or bounded rationality

Influence Diagrams

- flat or hierarchical
- explicit states or features or objects and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic actions
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or learned
- perfect rationality or bounded rationality

Decision-theoretic Planning

- flat or hierarchical
- explicit states or features or objects and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic actions
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or learned
- perfect rationality or bounded rationality

Reinforcement Learning

- flat or hierarchical
- explicit states or features or objects and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic actions
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or learned
- perfect rationality or bounded rationality

Comparison of Some Representations

	CP	MDPs	IDs	RL	POMDPs	GT
hierarchical	✔					
properties	✔		✔	✔		
relational	✔					
indefinite stage	✔	✔		✔	✔	
stochastic dynamics		✔	✔	✔	✔	✔
partially observable			✔		✔	✔
values		✔	✔	✔	✔	✔
dynamics not given				✔		
multiple agents						✔
bounded rationality						

The Dimensions Interact in Complex Ways

- Partial observability makes multi-agent and indefinite horizon reasoning more complex
- Modularity interacts with uncertainty and succinctness: some levels may be fully observable, some may be partially observable
- Three values of dimensions promise to make reasoning simpler:
  - Hierarchical reasoning
  - Objects and relations
  - Bounded rationality

Base Formalisms

- Classical AI is built on the foundation of logic
- Reinforcement learning and decision-theoretic planning is built on the foundations of decision/game theory (e.g., Markov decision processes).

These each provide some of the tools to tackle some of the dimensions.