The Independent Choice Logic: A pragmatic combination of logic and decision theory

David Poole

University of British Columbia

The independent choice logic influences

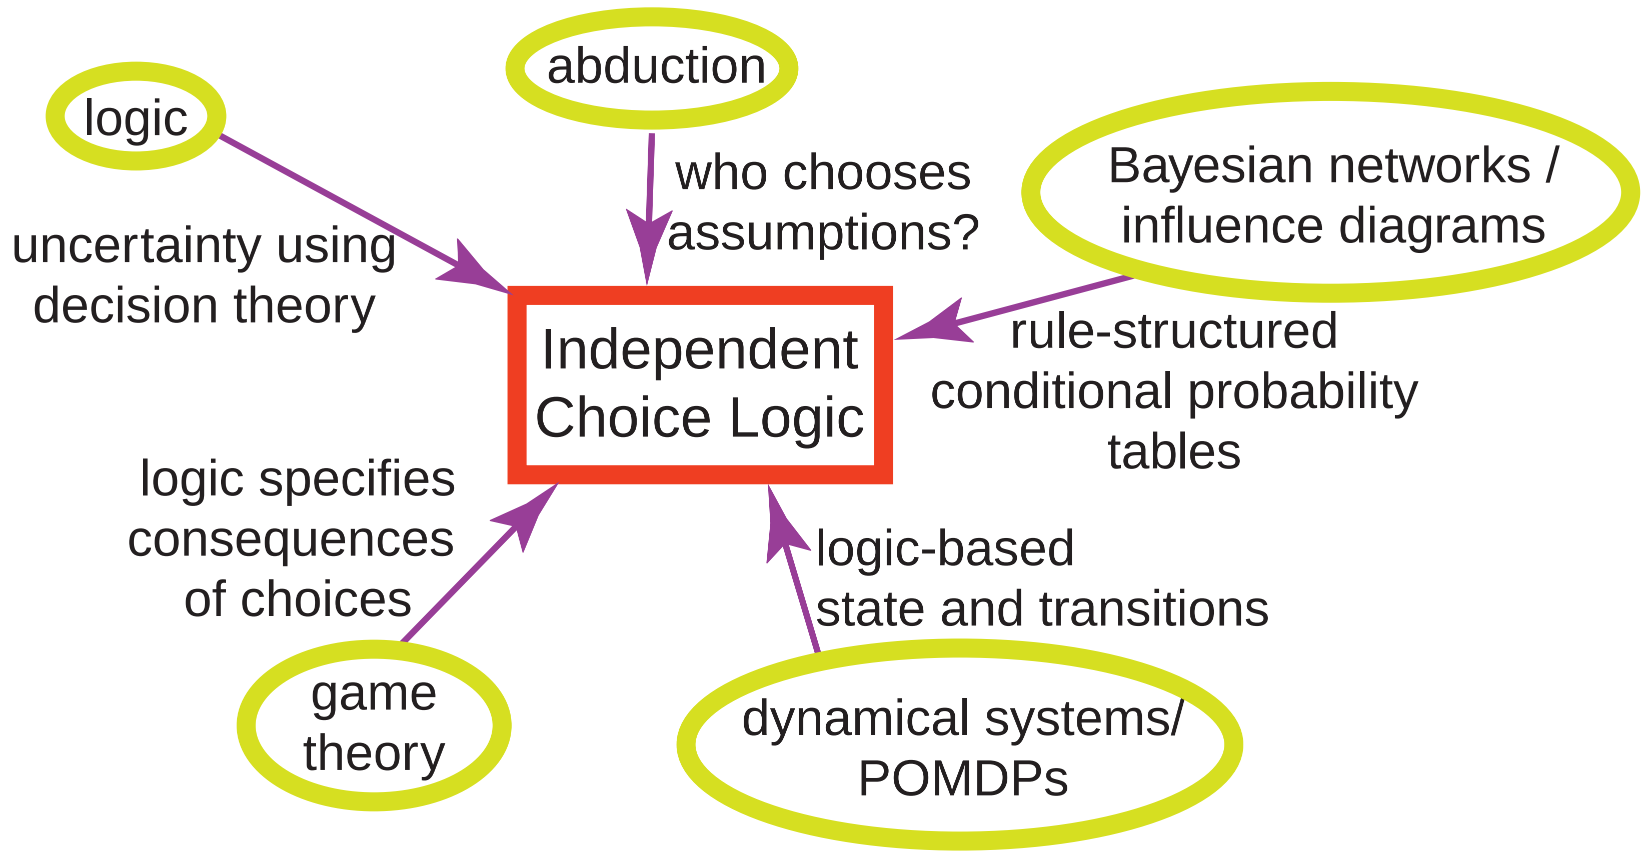

Overview

- Knowledge representation, logic, decision theory.
- Abduction + who chooses the assumptions?
- Logic + handle uncertainty using decision theory.
- Bayesian networks + rule-structured conditional probability tables.
- Dynamical systems and logic.

Knowledge Representation

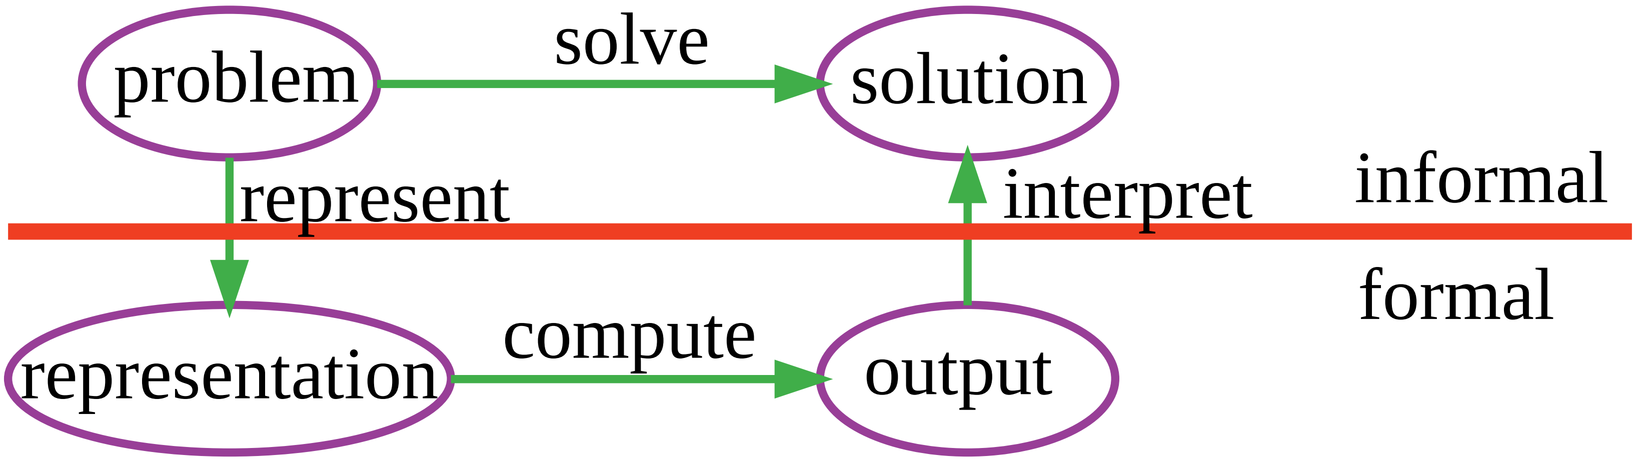

- Find compact / natural representations, exploit features of representation for computational gain.
- Approximate the solution, not the problem!
- Simplicity.

The problem: what should an agent do?

- It depends on its goals / background knowledge / (experience) / observations.
- Two normative traditions:
  - logic semantics (symbols have meaning), proofs
  - decision / game theory tradeoffs under uncertainty

  (use logic at the object-level, not the meta-level)

Assumption-based reasoning

- Given background knowledge / facts F and assumables / possible hypotheses H,
- An explanation of g is a set D of assumables such that

$$F \cup D \not\models \mathit{false}$$

$$F \cup D \models g$$

- abduction is when g is given and you want D
- default reasoning / prediction is when g is unknown

Who chooses the assumptions?

- **sceptical** adversary chooses the assumptions
- **credulous** agent chooses what assumptions it likes
- **probabilistic** nature gets to choose assumptions

	sceptical	credulous	probabilistic
prediction			
abduction			

Independent Choice Logic

- **C**, the choice space is a set of alternatives.
  An alternative is a set of atomic choices.
  An atomic choice is a ground atomic formula.
  An atomic choice can only appear in one alternative.
- **F**, the facts is an acyclic logic program.
  No atomic choice unifies with the head of a rule.

Example: cascaded inverters

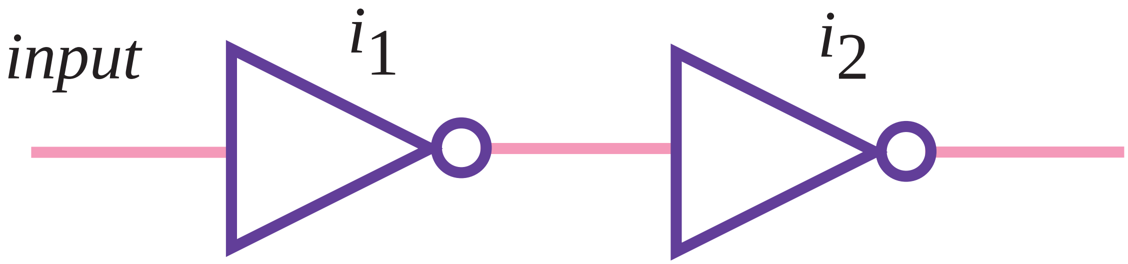

$$\mathbf{C} = \{ \quad \{ok(i_1), shorted(i_1), broken(i_1)\}$$
$$\{ok(i_2), shorted(i_2), broken(i_2)\}$$
$$\{input(on), input(off)\}\}$$

$$\mathbf{F} = \{ \quad out(i_1, on) \leftarrow ok(i_1) \wedge input(off),$$
$$out(i_1, V) \leftarrow shorted(i_1) \wedge input(V),$$
$$out(i_1, off) \leftarrow broken(i_1), \cdots\}$$

Abductive Characterisation of ICL

- The atomic choices are assumable.
- The elements of an alternative are mutually exclusive.
- Each alternative is controlled by an agent. They get to choose the elements of the alternative.

Note that:

- The choices are independent; the facts provide no constraints on choices.
- We can do both abduction and prediction.

Nature choosing assumptions

- Have a probability distribution over alternatives controlled by nature.
- For every alternative $\chi \in \mathbf{C}$ that is controlled by nature, there is a function:

$$P_0 : \chi \rightarrow [0, 1]$$

such that

$$1 = \sum_{\alpha \in \chi} P_0(\alpha)$$

Independent choice logic theory

C is a **choice space**

F, the **facts**, is an acyclic logic program such that no atomic choice unifies with the head of any rule.

A is a finite set of agents. There is a distinguished agent 0 called "nature".

controller is a function from $\mathbf{C} \rightarrow \mathbf{A}$. Let $\mathbf{C}_a = \{\chi \in \mathbf{C} : controller(\chi) = a\}$.

P_0 is a function $\cup\mathbf{C}_0 \rightarrow [0, 1]$ such that $\forall \chi \in \mathbf{C}_0$, $\sum_{\alpha \in \chi} P_0(\alpha) = 1$.

Probabilities of propositions

Suppose the rules are disjoint

$$a \leftarrow b_1$$

$$\cdots \qquad b_i \wedge b_j \text{ for } i \neq j \text{ can't be true}$$

$$a \leftarrow b_k$$

We can define:

$$P(g) = \sum_{E \text{ is a minimal explanation of } g} P(E)$$

$$P(E) = \prod_{h \in E} P_0(h)$$

P satisfies the axioms of probability.

Conditional Probabilities

$$P(g|e) = \frac{P(g \wedge e)}{P(e)} \begin{array}{l} \longleftarrow \text{explain } g \wedge e \\ \longleftarrow \text{explain } e \end{array}$$

- Given evidence e, explain e then try to explain g from these explanations.
- The explanations of $g \wedge e$ are the explanations of e extended to also explain g.
- Probabilistic conditioning is abduction + prediction.

Logic for reasoning

- How can we reconcile the normative arguments for logic and decision theory?
- Logic provides:
  - Symbols have denotation.
  - Way to determine truth of sentences (semantics).
  - Proof procedures.

  . . . so we need at least the first order predicate calculus.

Logic and decisions

- Claim: disjunction is a stupid way to handle uncertainty.
- Idea: lets try to handle all uncertainty using Bayesian decision theory / game theory.
- We want: the strongest logic that includes no uncertainty. Let's use acyclic logic programs (including negation as failure).
- All we have lost is the ability to handle uncertainty using disjunction!

Game theory

The strategic form of a game [von Neumann and Morgenstern, 1953]

- Multiple agents each get to choose a strategy.
- Nature has a probability distribution over strategies.
- A complete game (choice by every agent including nature) has a utility.
- Each player chooses its strategy to maximize its utility.

We use a logic program to specify the consequences of choices.

Semantics of ICL

- A **total choice** is a set containing exactly one element of each alternative in $\mathbf{C}$.
- For each total choice τ there is a **possible world** w_τ.
- Formula f is **true** in w_τ (written $w_\tau \models f$) if f is true in the (unique) stable model of $\mathbf{F} \cup \tau$.

Meaningless Example

$\mathbf{C} = \{\{c_1, c_2, c_3\}, \{b_1, b_2\}\}$

$$
\begin{aligned}
\mathbf{F} = \{\; & f \leftarrow c_1 \wedge b_1, && f \leftarrow c_3 \wedge b_2, \\
& d \leftarrow c_1, && d \leftarrow \sim c_2 \wedge b_1, \\
& e \leftarrow f, && e \leftarrow \sim d, \\
& u(a_1, 5) \leftarrow \sim e, && u(a_1, 0) \leftarrow e \wedge f, \\
& u(a_1, 9) \leftarrow e \wedge \sim f, && \\
& u(a_2, 7) \leftarrow d, && u(a_2, 2) \leftarrow \sim d\}
\end{aligned}
$$

There are 6 possible worlds:

$$
\begin{array}{lllllllll}
w_1 & \models & c_1 & b_1 & f & d & e & u(a_1, 0) & u(a_2, 7) \\
w_2 & \models & c_2 & b_1 & \sim f & \sim d & e & u(a_1, 9) & u(a_2, 2) \\
w_3 & \models & c_3 & b_1 & \sim f & d & \sim e & u(a_1, 5) & u(a_2, 7) \\
w_4 & \models & c_1 & b_2 & \sim f & d & \sim e & u(a_1, 5) & u(a_2, 7) \\
w_5 & \models & c_2 & b_2 & \sim f & \sim d & e & u(a_1, 9) & u(a_2, 2) \\
w_6 & \models & c_3 & b_2 & f & \sim d & e & u(a_1, 0) & u(a_2, 2)
\end{array}
$$

Abductive and semantic view

- The explanations of g form a concise (DNF) description of the worlds where g is true.
- The abductive characterisation is sound and complete with respect to the semantics.
- The possible worlds view shows how we can handle negation as failure and non-disjoint rules.
- The abductive characterisation can be extended to include negation as failure and non-disjoint rules.

Probabilities of Propositions

- When all choices are made by nature (& finite **C**):

$$P(w_\tau) = \prod_{a \in \tau} P_0(a)$$

$$P(f) = \sum_{\tau : w_\tau \models f} P(w_\tau)$$

- Theorem: the probabilities from the semantic view correspond to the probabilities in the abductive view.

Power of ICL

Surely independent hypotheses aren't powerful enough for real applications.

- No! Independent hypotheses can represent any probability distribution.
- The ICL can represent any probability represented in a Bayesian network.
- The ICL is more compact than a Bayesian network.

Factorization of probability distribution

- Bayesian networks provide a decomposition of a joint probability.

- Totally order the variables, $x_1, \ldots, x_n$.

$$
\begin{aligned}
P(x_1, \ldots, x_n) &= \prod_{i=1}^{n} P(x_i | x_{i-1} \ldots x_1) \\
&= \prod_{i=1}^{n} P(x_i | \pi_{x_i})
\end{aligned}
$$

- π_{x_i} are parents of x_i: set of variables such that the predecessors are independent of x_i given its parents.

(Bayesian) Belief Networks

- Graphical representation of dependence.
- DAGs with nodes representing random variables.
- Arcs from parents of a node into the node.
- If $b_1, \cdots, b_k$ are the parents of a, we have an associated conditional probability table

$$P(a|b_1, \cdots, b_k)$$

Bayesian Network for Overhead Projector

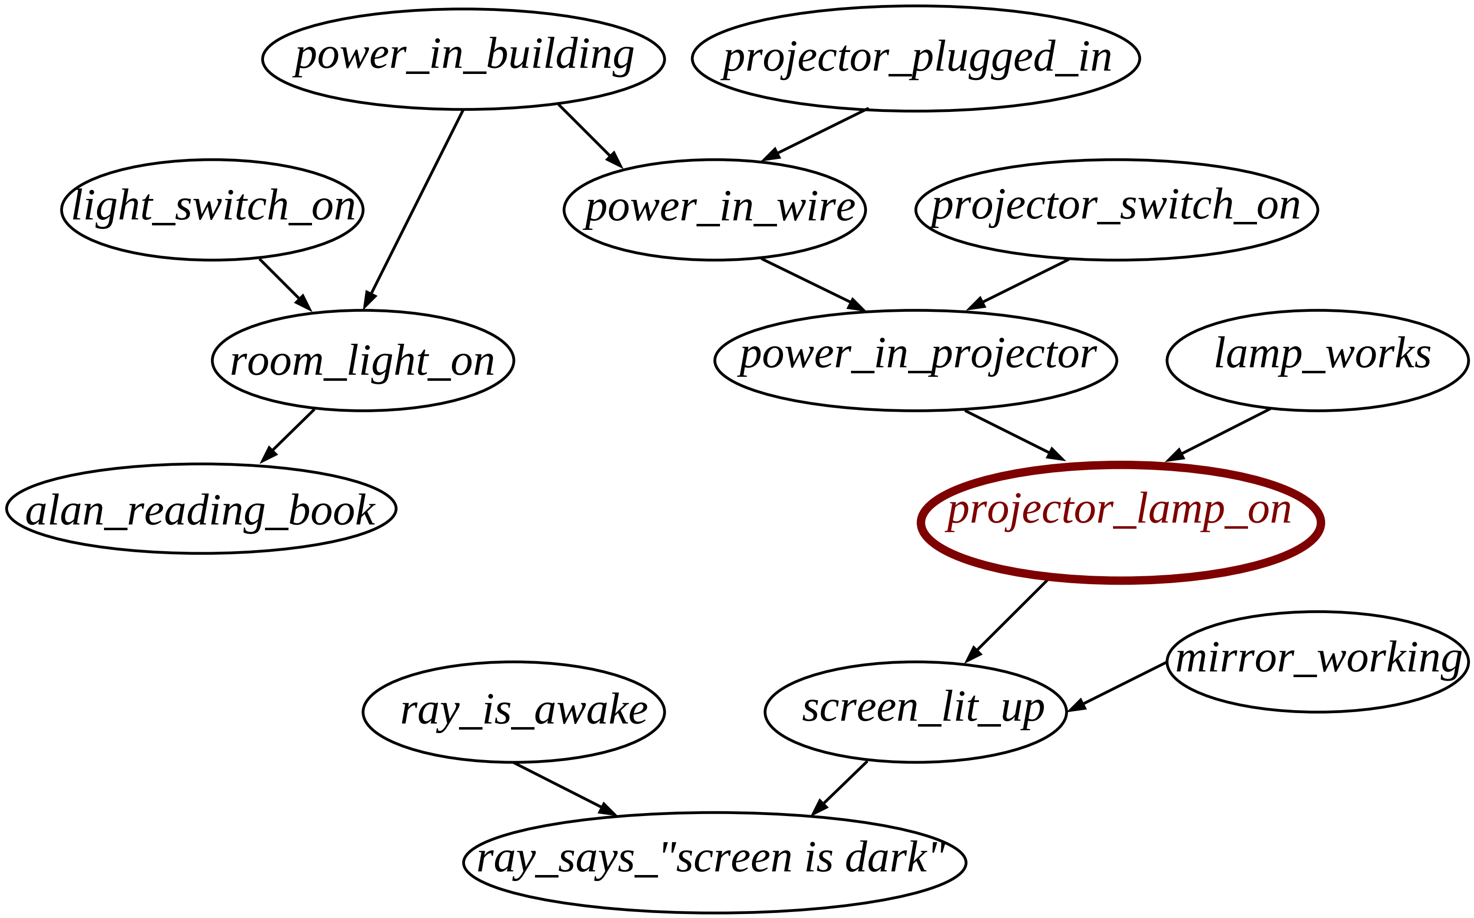

Bayesian networks as logic programs

$$
\begin{aligned}
&\mathit{projector\_lamp\_on} \leftarrow \\
&\qquad \mathit{power\_in\_projector} \wedge \\
&\qquad \mathit{lamp\_works} \wedge \\
&\qquad \mathit{projector\_working\_ok}. \longleftarrow \text{atomic choice} \\
&\qquad\qquad\qquad\qquad\qquad\qquad\quad\;\; \text{chosen by nature}
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{projector\_lamp\_on} \leftarrow \\
&\qquad \mathit{power\_in\_projector} \wedge \\
&\qquad \sim\mathit{lamp\_works} \wedge \\
&\qquad \mathit{working\_with\_faulty\_lamp}.
\end{aligned}
$$

Probabilities of hypotheses

$$
\begin{aligned}
& P_0(\mathit{projector\_working\_ok}) \\
& \quad = \quad P(\mathit{projector\_lamp\_on} \mid \\
& \qquad\qquad\quad \mathit{power\_in\_projector} \wedge \mathit{lamp\_works})
\end{aligned}
$$

— provided as part of Bayesian network

Mapping Bayesian networks into ICL

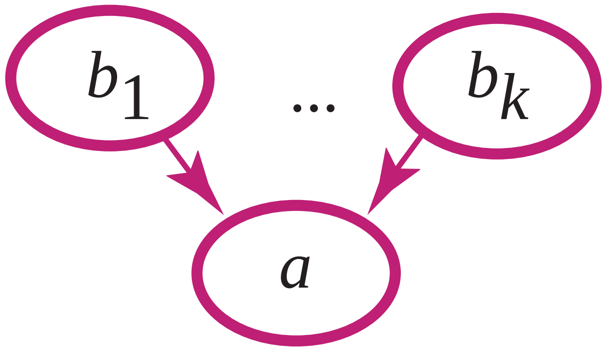

- Translated into the rules

$$a(V) \leftarrow b_1(V_1) \wedge \cdots \wedge b_k(V_k) \wedge h(V, V_1, \ldots, V_k).$$

- and the alternatives

$$\forall v_1 \cdots \forall v_k \{h(v, v_1, \ldots, v_k) | v \in domain(a)\} \in \mathbf{C}$$

Bayesian networks and the ICL

- The probabilities for the Bayesian network and the ICL translation are identical.
- In the translation, the ICL requires the same number of probabilities as the Bayesian network.
- Often the ICL theory is more compact than the corresponding conditional probability table.
- The probabilistic part of the ICL can be seen as a representation for the independence of Bayesian networks.

What can we learn from the mapping?

ICL adds

- rule-structured conditional probability tables
- logical variables and negation as failure in rules
- arbitrary computation in the network
- choices by other agents
- algorithms

Bayesian networks add

- theory of causation
- algorithms
- ties to MDPs, Neural networks, …

Where to now?

Algorithms:

- Extend model-based diagnosis algorithms to Bayes nets
- Extend Bayes net algorithms to exploit rule-structure

Decisions:

- Choices make by various agents
- Utility
- Actions contingent on observations (conditional plans)

Dynamical systems:

- Represent change using the situation calculus
- Logic-based partially observable Markov decision processes

Learning:

- represent the task of learning in the ICL
- combining inductive logic programming and Bayesian network (and neural network) learning

Decision Theory

- An agent makes choices to maximize its expected utility.
- What an agent should do now depends on what it will do in the future.
- What an agent will do in the future depends on what it will observe.
- An agent adopts a policy (strategy), a function from observations (and past actions) into actions.

Sequential decision problem

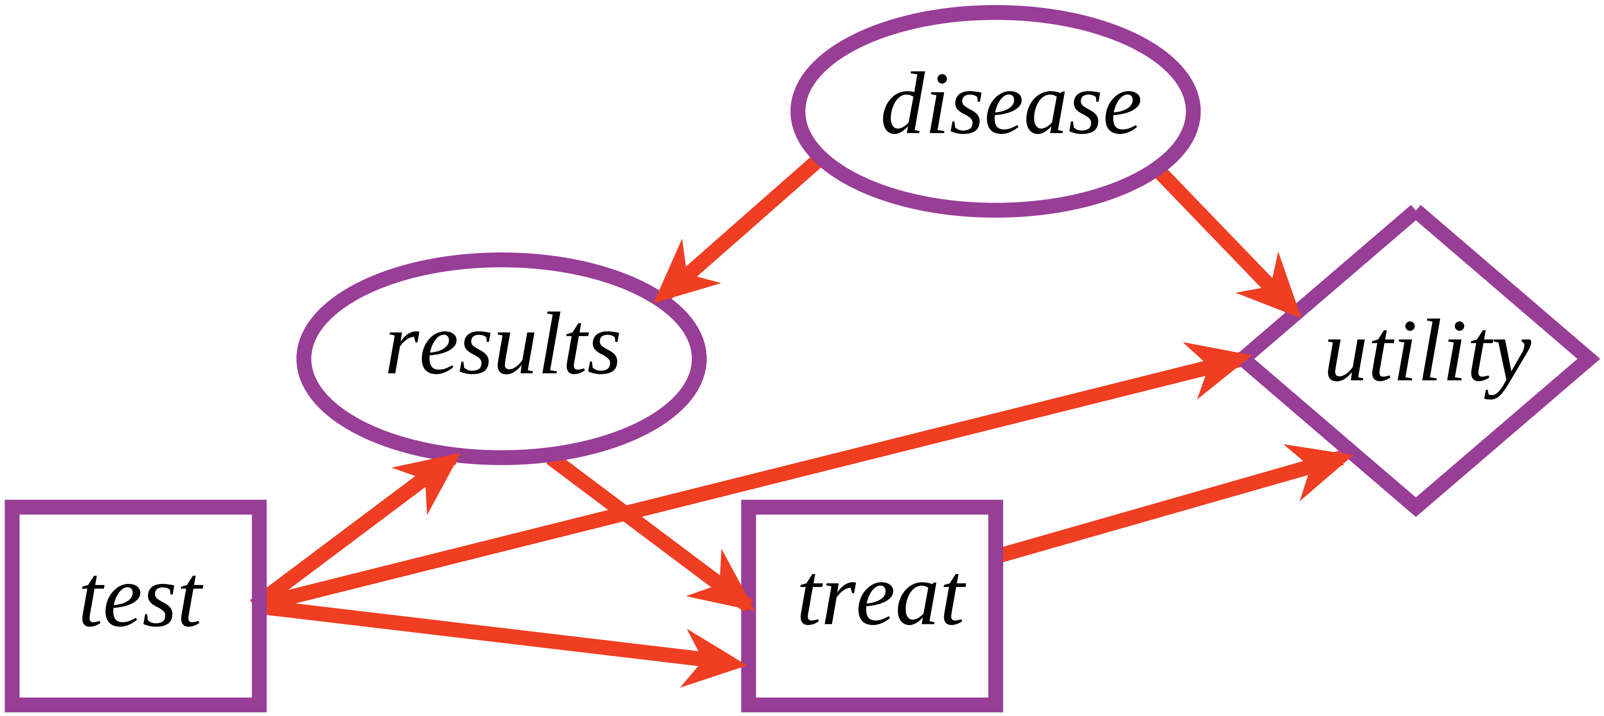

Representing the decision problem

You represent the problem with rules such as:

$$result(none) \leftarrow \sim test$$

$$result(positive) \leftarrow test \wedge disease \wedge \sim false\_neg.$$

$$result(positive) \leftarrow test \wedge \sim disease \wedge false\_pos.$$

$$utility(20) \leftarrow test \wedge disease \wedge treat.$$

A policy is something like:

$$test.$$

$$treat \leftarrow result(positive).$$

All of these rules imply an expected utility.

Example: a simple robot domain

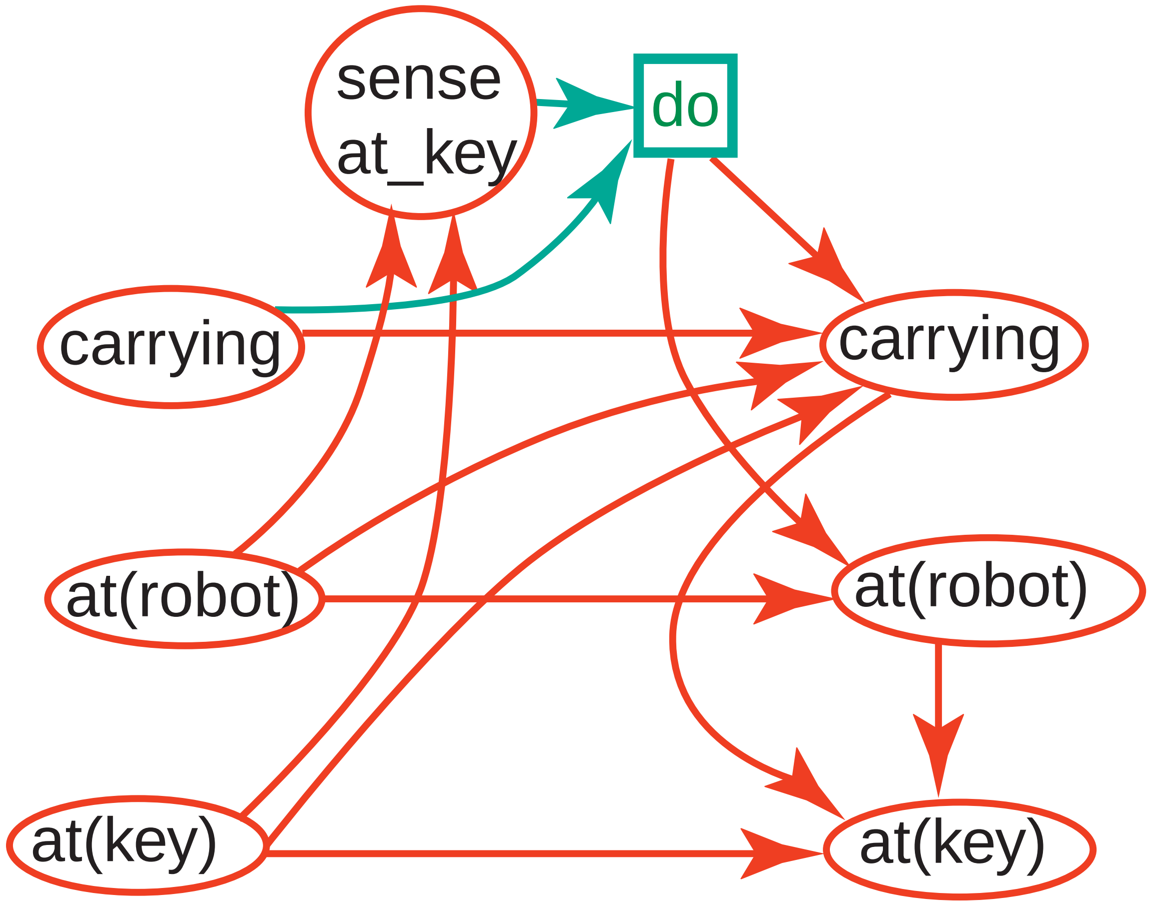

Axiomatising the simple robot domain

$$carrying(key, T+1) \leftarrow$$
$$do(pickup(key), T) \wedge$$
$$at(robot, Pos, T) \wedge at(key, Pos, T) \wedge$$
$$pickup\_succeeds(T).$$

$$carrying(key, T+1) \leftarrow$$
$$do(A, T) \wedge$$
$$carrying(key, T) \wedge$$
$$A \neq putdown(key) \wedge A \neq pickup(key) \wedge$$
$$keeps\_carrying(key, T).$$

Alternatives

$$\forall T\ \{pickup\_succeeds(T), pickup\_fails(T)\} \in \mathbf{C}_0$$

$P_0(pickup\_succeeds(T))$ is the probability the robot is carrying the key after the *pickup*(*key*) action when it was at the same position as the key, and wasn't carrying the key.

$$\forall S\ \{keeps\_carrying(key, T), drops(key, T)\} \in \mathbf{C}_0$$

Imperfect Sensors

A sensor is symptomatic of what is true in the world.

$$
\begin{aligned}
&sense(at\_key, T) \leftarrow \\
&\qquad at(robot, P, T) \wedge \\
&\qquad at(key, P, T) \wedge \\
&\qquad sensor\_true\_pos(T).
\end{aligned}
$$

$$
\begin{aligned}
&sense(at\_key, T) \leftarrow \\
&\qquad at(robot, P_1, T) \wedge \\
&\qquad at(key, P_2, T) \wedge \\
&\qquad P_1 \neq P_2 \wedge \\
&\qquad sensor\_false\_pos(T).
\end{aligned}
$$

Utility Axioms

Utility complete if $\forall w_\tau \forall T$, there exists unique U such that $w_\tau \models \mathit{utility}(U, T)$

$$
\begin{aligned}
&\mathit{utility}(R + P, T) \leftarrow \\
&\qquad \mathit{prize}(P, T) \wedge \\
&\qquad \mathit{resources}(R, T).
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{prize}(-1000, T) \leftarrow \mathit{crashed}(T). \\
&\mathit{prize}(1000, T) \leftarrow \mathit{in\_lab}(T) \wedge \sim\!\mathit{crashed}(T). \\
&\mathit{prize}(0, T) \leftarrow \sim\!\mathit{in\_lab}(T) \wedge \sim\!\mathit{crashed}(T).
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{resources}(200, s_0).\\
&\mathit{resources}(R - \mathit{Cost}, T + 1) \leftarrow\\
&\qquad \mathit{do}(\mathit{goto}(\mathit{To}, \mathit{Route}), T) \land\\
&\qquad \mathit{at}(\mathit{robot}, \mathit{From}, T) \land\\
&\qquad \mathit{pathcost}(\mathit{From}, \mathit{To}, \mathit{Route}, \mathit{Cost}) \land\\
&\qquad \mathit{resources}(R, T).\\
&\mathit{resources}(R - 10, T + 1) \leftarrow\\
&\qquad \mathit{do}(A, T) \land\\
&\qquad \sim\mathit{gotoaction}(A) \land\\
&\qquad \mathit{resources}(R, T).\\
&\mathit{gotoaction}(\mathit{goto}(\mathit{Pos}, T)).
\end{aligned}
$$

Example Policy

$$do(pickup(key), T) \leftarrow$$
$$sense(at\_key, T) \wedge$$
$$\sim carrying(key, T).$$
$$do(goto(door1, direct), T) \leftarrow$$
$$carrying(key, T).$$
$$do(goto(key\_cupboard, direct), T) \leftarrow$$
$$\sim sense(at\_key, T) \wedge$$
$$\sim carrying(key, T).$$

Conclusions

- ICL is a representation that combines logic and Bayesian decision theory / game theory.
- Generalises acyclic logic programs, Bayesian networks, the strategic form of a game, …
- All rules can be interpreted logically. All numbers can be interpreted as probabilities.
- It's (reasonably) simple.
- Applications: diagnosis, robot control, multimedia presentation, user modelling, …