Reasoning and Decision-Making under Uncertainty

3. Session: „Probabilistic turn“ -- Uncertainty, Degrees of Belief & Probability Calculus

Prof. Dr.-Ing. Stefan Kopp
Center of Excellence „Cognitive Interaction Technology“
AG Sociable Agents

CITEC Sociable Agents

Intelligent agent

Reasoning, inference

Decision-making, action selection

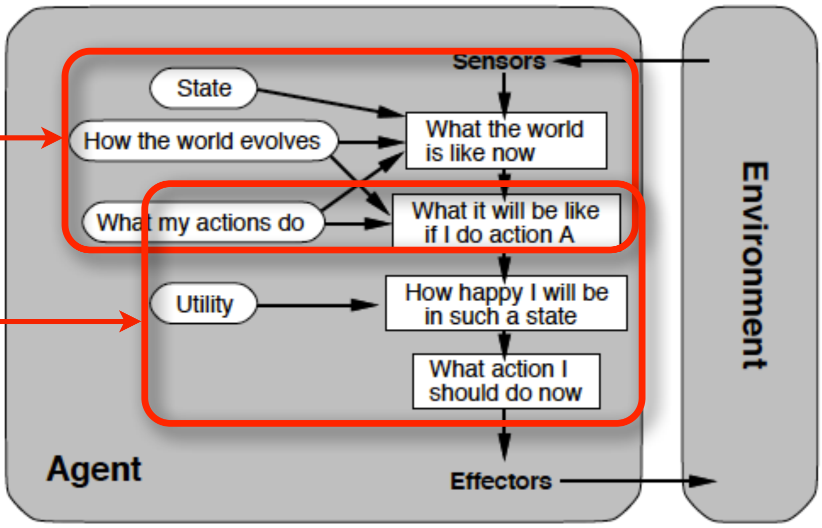

The need to cope with uncertainty

Many causes for uncertainty in reasoning and decision-making

- incomplete knowledge, „invisible facts“
  - environment not fully observable and non-deterministic
  - state of the world might have changed already
  - actions might not have desired effects
  - reliance on default assumptions
- non-adequate formalism (calculus), trading adequacy for tractability
  - frame problem, qualification problem, ramification problem
  - modularity (locality and detachment) of logical inference
  - limited horizon in planning, bounded rationality
  - cumulating uncertainties when drawing inferences
  - combinatorial explosion when accounting for contingencies and indeterminacy

How to deal with uncertainty?

Two principled approaches:

- extensional (rule-based): assign certainties locally to formulae and update during inference
  - computationally convenient, semantically sloppy

> rule $A \xrightarrow{x} B$ reads as follows: "If you see the certainty of A undergoing a change δ_A, then regardless of what other things the knowledge base contains and regardless of how δ_A was triggered, you are given an unqualified license to modify the current certainty of B by some amount δ_B, which may depend on x, on δ_A, and on the current certainty of B."†

- intensional (model-based): assign certainties globally to *possible worlds* which can be more or less exactly specified
  - computationally clumsy, semantically clear

Judea Pearl, Probabilistic reasoning in intelligent systems, Morgan Kaufmann, 1988.

How to deal with uncertainty?

Example: SARS diagnosis

$Fever(X) \wedge Cough(X) \Rightarrow SARS(X)$

Problem: Too many (unknown) pre-conditions

$SARS(X) \Rightarrow Fever(X) \wedge Cough(X)$

Problem: Lack in theory to determine certain consequences

Idea: Instead of enumerating all exact antecedents/conclusions and explicit rules, and assigning them local uncertainty values, summarize their co-occurrences implicitly by numbers and interpret as probabilities

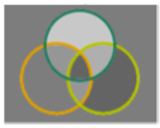

$Pr(Fever, Cough|SARS),\ Pr(SARS|Fever, Cough)$

Summarize in probability numbers

How to deal with uncertainty?

Two principled approaches

extensional

- A (= true/false)
- A → B (= true/false)

- B (= true/false)
  deductive inference

intensional

- P(A) = prop. of A
- P(B|A) = prop. of B given A

- P(B) => P(B|A)
  update, conditioning

Vagueness vs. Uncertainty
Probabilities ~ uncertainty if a prop. is true (→ degree of belief)
Not vagueness, i.e. gradual truth of propositions (→ fuzzy-metrics)

Degree of belief and probability

<u>Two alternative interpretations of probability values:</u>

- frequentist
  - probabilities = frequencies of (co-)occurrence
- Bayesian (subjective)
  - probabilities = „degrees of beliefs"
  - conditional probabilities represent organization of human knowledge, where an event *A* is always measured against a context (evidence) *B*

We use the subjective (Bayesian) interpretation:

- Probabilities quantify the degree of belief in a proposition, i.e., relate facts to the agent's own state of knowledge („*how sure am I that....?*")
- Using probability theory to model the form & dynamics of degrees of belief, and the influencing between degrees of different beliefs

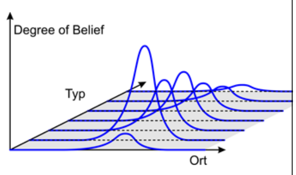

Degree of belief and probability calculus

<u>Example</u>:

- *A = „Ted Kennedy will seek nomination for president in 2012"*
- *Pr(A|K)* ~ agent's subjective (degree of) belief in the event described by *A*, given an available body of knowledge *K*
- may change entirely when new evidence *K'* arrives!
- how to update the degree of belief in A, conditioned upon knowledge K and K' ?

Degrees of belief obey the laws of probability theory, i.e. can use probability calculus for modeling and updating degrees of belief

Propositional logics

World = state of affairs in which each propositional variable is known

- variable assignment with values

ω

Models = worlds that satisfy a sentence

- every sentence represents a *set of worlds* = (atomic) event

$$Mods(\alpha) = \{\omega : \omega \models \alpha\}$$

World	**Earthquake**	**Burglary**	**Alarm**
w1	*true*	*true*	*true*
w2	*true*	*true*	*false*
w3	*true*	*false*	*true*
w4	*true*	*false*	*false*
w5	*false*	*true*	*true*
w6	*false*	*true*	*false*
w7	*false*	*false*	*true*
w8	*false*	*false*	*false*

$$Mods(\alpha \wedge \beta) = Mods(\alpha) \cap Mods(\beta)$$
$$Mods(\alpha \vee \beta) = Mods(\alpha) \cup Mods(\beta)$$
$$Mods(\neg\alpha) = \overline{Mods(\alpha)}$$

Propositional logics

Important properties of sentences

- consistent / satisfiable $Mods(\alpha) \neq \{\}$
- valid $Mods(\alpha) \neq \Omega$ $\models \alpha$

Important relationships of sentences

- equivalent $Mods(\alpha) = Mods(\beta)$
- mutually exclusive $Mods(\alpha) \cap Mods(\beta) = \{\}$
- exhaustive $Mods(\alpha) \cup Mods(\beta) = \Omega$
- implies $\alpha \models \beta$ $Mods(\alpha) \subseteq Mods(\beta)$

Monotonicity of logical reasoning

World	Earthquake	Burglary	Alarm
w1	true	true	true
w2	true	true	false
w3	true	false	true
w4	true	false	false
w5	false	true	true
w6	false	true	false
w7	false	false	true
w8	false	false	false

$$\alpha : (Earthquake \vee Buglary) \Rightarrow Alarm$$
$$Mods(\alpha) = \{\omega_1, \omega_3, \omega_5, \omega_7, \omega_8\}$$

+

$$\beta : Earthquake \Rightarrow Burglary$$

$$Mods(\alpha \wedge \beta) = Mods(\alpha) \cap Mods(\beta) = \{\omega_1, \omega_5, \omega_7, \omega_8\}$$

Monotonicity

learning new information can *only rule out* worlds!

- if *a* implies *c*, then (*a and b)* will imply *c* as well
- especially problematic in light of qualification problem

Modeling degrees of belief as probabilities

Degree of belief or probability of a *world*

- in fuzzy logic, interpreted as possibility/ vagueness (*not the view adopted here*)

$$Pr(\omega)$$

Degree of belief or probability of a *sentence*

$$Pr(\alpha) := \sum_{\omega \models \alpha} Pr(\omega)$$

State of belief or joint probability distribution

World	Earthquake	Burglary	Alarm	Pr(.)
w1	true	true	true	.0190
w2	true	true	false	.0010
w3	true	false	true	.0560
w4	true	false	false	.0240
w5	false	true	true	.1620
w6	false	true	false	.0180
w7	false	false	true	.0072
w8	false	false	false	.7128

$$\sum_{\omega_i} Pr(\omega_i) = 1$$

$$Pr(Earthquake) = .1$$
$$Pr(Burglary) = .2$$
$$Pr(Alarm) = .2442$$

Properties of beliefs

Properties of (degrees of) beliefs

- bound — $0 \leq Pr(\alpha) \leq 1 \quad \forall\alpha$
- baseline for inconsistent sentences — $Pr(\alpha) = 0 \; \forall\alpha \; inconsistent$
- baseline for valid sentences — $Pr(\alpha) = 1 \quad \forall\alpha \; valid$

Junctions of beliefs

- disjunction — $Pr(\alpha \vee \beta) = Pr(\alpha) + Pr(\beta) - Pr(\alpha \wedge \beta)$
- conjunction — $Pr(\alpha \wedge \beta) = \sum_{\omega \models \alpha, \omega \models \beta} Pr(\omega)$

$$Pr(\alpha \wedge \beta) = 0 \; if \; \alpha, \beta \; mutually \; exclusive$$

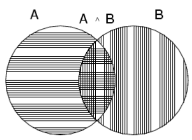

$$Pr(Earthquake \wedge Burglary) = Pr(\omega_1) + Pr(\omega_2) = .02$$

$$Pr(Earthquake \vee Burglary) = .1 + .2 - .02 = .28$$

Degree of belief and probability calculus

Given an agent‘s belief state (= degree of beliefs about all possible values of a variable *X*) its uncertainty about the outcome of the event described by *X* can be quantified using the (Shannon) entropy:

$$ENT(X) := -\sum_x Pr(x) log_2 Pr(x) \qquad (0 \; log 0 := 0)$$

<u>Minimum entropy</u>

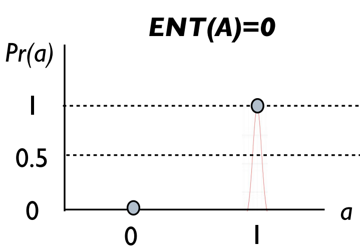

<u>Maximum entropy</u>

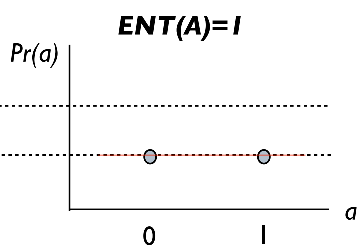

Information theory (Shannon): measure of the amount of information that is missing before reception, expected amount of information in a message

Uncertainty and entropy

Example:

$$ENT(X) := -\sum_x Pr(x) log_2 Pr(x)$$

$$(0\ log 0 := 0)$$

World	**Earthquake**	**Burglary**	**Alarm**	*Pr(.)*
w1	*true*	*true*	*true*	*.0190*
w2	*true*	*true*	*false*	*.0010*
w3	*true*	*false*	*true*	*.0560*
w4	*true*	*false*	*false*	*.0240*
w5	*false*	*true*	*true*	*.1620*
w6	*false*	*true*	*false*	*.0180*
w7	*false*	*false*	*true*	*.0072*
w8	*false*	*false*	*false*	*.7128*

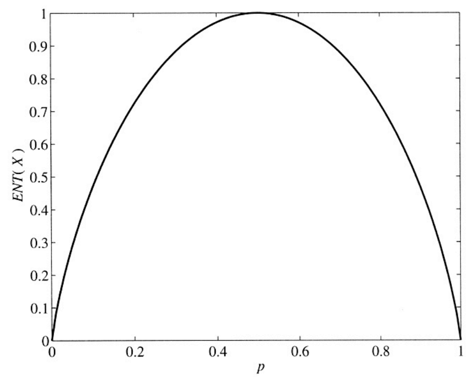

	Earthquake	**Burglary**	**Alarm**
true	*.1*	*.2*	*.2442*
false	*.9*	*.8*	*.7558*
ENT(.)	*.469*	*.722*	*.802*

Updating beliefs

Evidence = a piece of information known to hold β

➔ requires to update state of belief such that $Pr(.) \to Pr(.|\beta)$

- worlds that contradict evidence get zero prob $Pr(\beta|\beta) = 1$ $Pr(\omega|\beta) = 0 \quad for\ all\ \omega \models \neg\beta$
- normalized $\sum_{\omega \models \beta} Pr(\omega|\beta) = 1$
- retain impossible worlds $Pr(\omega) = 0 \to Pr(\omega|\beta) = 0$
- worlds consistent with evidence and positive prob. retain relative beliefs in possible worlds $\frac{Pr(\omega)}{Pr(\omega')} = \frac{Pr(\omega|\beta)}{Pr(\omega'|\beta)}$ $\forall \omega, \omega' \models \beta, Pr(\omega) > 0, Pr(\omega') > 0$

Updating beliefs

➔ update old state of beliefs through **conditioning** on evidence β

$$Pr(\omega|\beta) := \begin{cases} 0 & \omega \models \neg\beta \\ \frac{Pr(\omega)}{Pr(\beta)} & \omega \models \beta \end{cases}$$

new belief state = old belief state, normalized with old belief in new evidence

Earthquake	**Burglary**	**Alarm**	*Pr(.)*
true	*true*	*true*	*.0190*
true	*true*	*false*	*.0010*
true	*false*	*true*	*.0560*
true	*false*	*false*	*.0240*
false	*true*	*true*	*.1620*
false	*true*	*false*	*.0180*
false	*false*	*true*	*.0072*
false	*false*	*false*	*.7128*

Alarm=true →

Earthquake	**Burglary**	**Alarm**	*Pr(.|Alarm)*
true	*true*	*true*	*.0190/.2442*
true	*true*	*false*	*0*
true	*false*	*true*	*.0560 /.2442*
true	*false*	*false*	*0*
false	*true*	*true*	*.1620 /.2442*
false	*true*	*false*	*0*
false	*false*	*true*	*.0072 /.2442*
false	*false*	*false*	*0*

$$Pr(Burglary) = .2 \rightarrow Pr(Burglary|Alarm) = .741$$

Updating beliefs

More efficient: direct update of a *local* sentence from new evidence through **Bayesian conditioning**

$$Pr(\alpha|\beta) = \frac{Pr(\alpha \wedge \beta}{Pr(\beta)}$$

follows from the following commitments:

- worlds that contradict evidence have zero prob
- worlds that have zero prob continue to have zero prob
- worlds that are consistent with evidence and have positive prob will maintain their relative beliefs

Note: Bayesian conditioning is nothing else than application of the basic product rule

$$Pr(\alpha \wedge \beta) = Pr(\alpha|\beta) \cdot Pr(\beta)$$

Updating beliefs

Example: State of belief from above

	Pr(Earthquake)	Pr(Burglary)	Pr(Alarm)
true	*.1*	*.2*	*.2442*

Conditioning on first evidence: *Alarm=true*

	Pr(E|Alarm)	Pr(B|Alarm)	Pr(A|Alarm)
true	*.307*	*.741*	*1*

Conditioning on second evidence: *Earthquake=true*

	Pr(E|A∧E)	Pr(B|A∧E)	Pr(A|A∧E)
true	*1*	*.253*	*1*

➔ *Note*: belief dynamics under incoming evidence is a **consequence of the initial state of beliefs** one had!

Updating beliefs

Updating the belief state is possible, but computationally costly as soon as worlds become complex, i.e. many variables with large domains

- need to sum over all worlds consistent with new evidence, eventually need joint distribution to have prob for any combination of evidence
- no.'s of worlds exponential in no.'s of variables, so is the joint distribution
  - O(dn) with n random variables and domain size d

Idea: Exploit independencies in the world, i.e. assumptions that certain variables have nothing to do with each other, and learning about one doesn't change (degree of) belief in the other

- *„our most basic, robust, and commonly available knowledge about uncertain environments"* -- especially a specific kind of independence...