

Yet another *Artificial Intelligence 1&2* Summary

Written by Philip K.,* extending “(Yet another)² *Artificial Intelligence 1* Summary”[†] by Lorenz Gorse[‡]

Last updated for the Summer Semester 2021

1 Agents

Def. 1 (Agent). An agent a is an entity that perceives (via sensors) and acts (via actuators). It can be modelled as an *agent function* $f_a: \mathcal{P}^* \mapsto \mathcal{A}$, mapping from percept histories to actions.

Def. 2 (Performance measure). A function that evaluates a sequence of environments.

Def. 3 (Rationality). An **agent** is “rational”, if it chooses the actions that maximizes the expected value of the **performance measure** given the percept history.

Def. 4 (Autonomy). An **agent** is called autonomous, if it does not rely on the prior knowledge of the designer. Autonomy avoids fixed behaviour in changing environments.

Def. 5 (Task environment). The combination of a **performance measure**, environment, actuators and sensors (PEAS) describes a (task) environment e .

Def. 6 (Environment properties). An **environment** is called... **fully observable**, iff a 's sensors give it access to the complete state of e at any point in time, else **partially observable**. **deterministic**, iff the next state of e is completely determined by a 's action and e 's current state, else **stochastic**.

episodic, iff a 's experience is divided into atomic, independent episodes, where it perceives and performs a single action. Non-episodic environments are called **sequential**.

dynamic, iff e can change without an action performed by a , else **static**.

discrete, iff the sets of e 's states and a 's actions are countable, else **continuous**.

single-agent, iff only a acts on e .

Def. 7 (Simple reflex agent). An **agent** that bases its next action only on the most recent percept, $f_a: \mathcal{P} \mapsto \mathcal{A}$.

Def. 8 (Model-based agent). Like **simple reflex agent**, but additionally maintains a (world) model to decide it's next move.

Def. 9 (Goal-based agent). A **model-based agent** that also takes it's goals into consideration when deciding.

Def. 10 (Utility-based agent). An **agent** that combines a world model and a utility function, measuring state-preferences. Its choices attempt to maximize the expected utility, allowing rational decisions where **goals** are insufficient.

Def. 11 (Learning agent). An **agent** that augments the performance element, which chooses actions from percept sequences, with a...

learning element making improvements to the agent's performance element.

critic giving feedback to the *learning element* based on an external performance standard.

problem generator suggesting actions that can lead to new, informative experiences.

Def. 12 (State representation). We call a state representation: **atomic** if it has no internal structure.

factored if each state is characterized by attributes and their values.

structured if the state includes objects and their relations.

*<https://gitlab.cs.fau.de/oj14ozun/ai2-summary>, the source for this document should be accessible as a PDF attachment. The document and the source is distributed under **CC BY-SA 4.0**.

[†]<https://gitlab.cs.fau.de/oj14ozun/ai1-summary>

[‡]<https://gitlab.cs.fau.de/snippets/15>

2 Search

Def. 13 (Search problem). A search problem $\Pi := \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{I}, \mathcal{G} \rangle$ consists of a set \mathcal{S} of states, a set \mathcal{A} of actions, a transition model $\mathcal{T}: \mathcal{A} \times \mathcal{S} \mapsto \mathfrak{P}(\mathcal{S})$ that assigns any action and state to a set of successor states. Certain states in \mathcal{S} are labelled as “goal states” \mathcal{G} and “initial states” \mathcal{I} . A cost function $c: \mathcal{A} \mapsto \mathbb{R}_0^+$ may assigns costs to actions.

Def. 14 (Solution). A sequence of applicable actions that lead from a initial state \mathcal{I} to a goal state $g \in \mathcal{G}$ is a solution.

Def. 15 (Problem types). **Problems** come in many variations: **Single-state problem**: state is always known with certainty (observable, deterministic, static, discrete)

Multiple-state problem: know which states might be in (initial state not/partially observable)

Contingency problem: constructed plans with conditional parts based on sensors (non-deterministic, unknown state space)

Def. 16 (Tree search). An algorithm that explores state spaces, forming a search tree of already-explored states, modelled as nodes. It's fringe are the nodes that have not yet been considered.

Def. 17 (Search strategy). A search strategy picks a **node** from the fringe of a search tree. It's properties are:

Completeness: Does it always find a solution if one exists?

Time complexity: Number of nodes generated/expanded.

Space complexity: Maximum number of nodes held in memory.

Optimality: Does it always find the least-cost solution?

Def. 18 (Uninformed search). **Search strategies** that only employ information from the **problem definition** yield uninformed searches. Examples are breadth-first-search (BFS), uniform-cost-search (UCS, also called “Dijkstra's algorithm”), depth-first-search (DFS), depth-limited search and iterative-deepening-search (IDS).

Def. 19 (Informed search). **Search strategies** that use information about the real world beyond the **problem statement** yield informed searches. The additional information about the world is provided in form of **heuristics**. Examples are **greedy-search** and **A*-search**.

Def. 20 (Heuristic). A heuristic is an evaluation function $h: \mathcal{S} \mapsto \mathbb{R}_0^+ \cup \{\infty\}$ that estimates the cost from a state n to the nearest **goal state**. If $s \in \mathcal{G}$, then $h(s) = 0$. All **nodes** for the same states must have the same h -value.

Def. 21 (Goal distance function). A function $h^*: \mathcal{S} \mapsto \mathbb{R}_0^+ \cup \{\infty\}$ determining the cheapest path from any **nodes** to a **goal state**, or ∞ if no path exists.

Def. 22 (Admissibility and consistency). A **heuristic** h is admissible if $h(s) \leq h^*(s)$ for all states $s \in \mathcal{S}$, i.e. forming a lower bound. h is consistent if $h(s) - h(s') \leq c(a)$ for all $s \in \mathcal{S}$, $a = (s, s') \in \mathcal{A}$ and a cost function c .

Def. 23 (Greedy-search). Greedy-search always expands the **node** that appears to be closest to a goal state, as determined by a **heuristic**.

Def. 24 (A^* -search). Expands the node with the minimum evaluation value $f(s)=g(s)+h(s)$, where $g(s)$ is the path cost. A^* -search is optimal if it uses an **admissible heuristic** h .

Def. 25 (Dominance). Let h_1 and h_2 be two **admissible heuristics** we say that h_1 dominates h_2 if $h_1(s) \geq h_2(s)$ for all $s \in S$. The dominant heuristic is better for search.

Def. 26 (Relaxation). A **search problem** $\Pi := \langle S, \mathcal{A}, \mathcal{T}, \mathcal{I}, \mathcal{G} \rangle$ has a relaxed problem $\Pi^r := \langle S, \mathcal{A}^r, \mathcal{T}^r, \mathcal{I}^r, \mathcal{G}^r \rangle$ iff $\mathcal{A} \subseteq \mathcal{A}^r$, $\mathcal{T} \subseteq \mathcal{T}^r$, $\mathcal{I} \subseteq \mathcal{I}^r$, $\mathcal{G} \subseteq \mathcal{G}^r$. This means that any solution for Π is a solution for Π^r .

Def. 27 (Local search). A search algorithm that only operates on a single space at a time is called a local search. Local search algorithms need constant space, because it doesn't have to remember multiple paths. Examples are hill-climbing, simulated annealing or genetic algorithms.

2.1 Adversarial Search

Def. 28 (Game state space). A 6-tuple $\Theta = \langle S, A, T, I, S^T, u \rangle$ is a game state space, for two players "Max" and "Min" consists of:

- S is the disjoint union of S^{Max} , S^{Min} and S^T (respectively the sets of "Max" 's, "Min" 's and terminal states).
- A is the disjoint union $A^{\text{Max}} \subseteq S^{\text{Max}} \times (S^{\text{Min}} \cup S^T)$ and $A^{\text{Min}} \subseteq S^{\text{Min}} \times (S^{\text{Max}} \cup S^T)$
- I is the initial state.
- $u: S^T \mapsto \mathbb{R}$ is the utility function.

Def. 29 (Strategy). Let Θ be a **game state space**, and $X \in \{\text{Max}, \text{Min}\}$. A strategy for X is a function $\sigma^X: S^X \mapsto A^X$ so that a is applicable to s whenever $\sigma^X(s) = a$. A strategy is optimal if it yields the best possible utility for X assuming perfect opponent play.

Def. 30 (Minimax Algorithm). The minimax **algorithm** is given by the following function whose input is a state $s \in S^{\text{Max}}$, in which Max is to move. It attempts to find the best move for "Max":

Algo. 1 MinimaxDecision(s) **returns** an action

- 1: $v := \text{MaxValue}(s)$
 - 2: **return** an action yielding value v in the previous function call
-

Algo. 2 MaxValue(s) **returns** a utility value

- 1: **if** TerminalTest(s) **then return** $u(s)$
 - 2: $v := -\infty$
 - 3: **for each** $a \in \text{Actions}(s)$ **do**
 - 4: $v := \max(v, \text{MinValue}(\text{ChildState}(s, a)))$
 - 5: **return** v
-

Algo. 3 MinValue(s) **returns** a utility value

- 1: **if** TerminalTest(s) **then return** $u(s)$
 - 2: $v := +\infty$
 - 3: **for each** $a \in \text{Actions}(s)$ **do**
 - 4: $v := \min(v, \text{MaxValue}(\text{ChildState}(s, a)))$
 - 5: **return** v
-

Def. 31 (Alpha-Beta Search). To avoid evaluating states that are not of interest, *Alpha-Beta Pruning* can be used to accelerate **Minimax search**:

Algo. 4 AlphaBetaSearch(s) **returns** an action

- 1: $v := \text{MaxValue}(s, -\infty, +\infty)$
 - 2: **return** an action yielding value v in the previous function call
-

Algo. 5 MaxValue(s, α, β) **returns** a utility value

- 1: **if** TerminalTest(s) **then return** $u(s)$
 - 2: $v := -\infty$
 - 3: **for each** $a \in \text{Actions}(s)$ **do**
 - 4: $v := \max(v, \text{MinValue}(\text{ChildState}(s, a), \alpha, \beta))$
 - 5: $\alpha := \max(\alpha, v)$
 - 6: **if** $v \geq \beta$ **then return** v $\triangleright v \geq \beta \iff \alpha \geq \beta$
 - 7: **return** v
-

Algo. 6 MinValue(s, α, β) **returns** a utility value

- 1: **if** Terminal-Test(s) **then return** $u(s)$
 - 2: $v := +\infty$
 - 3: **for each** $a \in \text{Actions}(s)$ **do**
 - 4: $v := \min(v, \text{MaxValue}(\text{ChildState}(s, a), \alpha, \beta))$
 - 5: $\beta := \min(\beta, v)$
 - 6: **if** $v \leq \alpha$ **then return** v $\triangleright v \leq \beta \iff \alpha \geq \beta$
 - 7: **return** v
-

Def. 32 (Monte-Carlo Tree Search). If there is no good known evaluation function, *Monte-Carlo Tree Search* decides on an action through sampling average $u(t), t \in S^T$. For Monte-Carlo tree search we maintain a search tree T :

Algo. 7 MonteCarloTreeSearch(s) **returns** an action

- 1: **while** time not up **do**
 - 2: apply actions within T to select a leaf state s'
 - 3: select action a' applicable to s'
 - 4: run random sample from a'
 - 5: add s' to T , update averages etc.
 - 6: **return** an a for s with maximal average $u(t)$
 - 7: When executing a , keep the part of T below a .
-

3 Constraint Satisfaction Problems

Def. 33 (Constraint Satisfaction Problem, CSP). This is a **search problem** where the states are given by a finite set of variables $V := \{X_1, \dots, X_n\}$ over domains $D := \{D_v \mid v \in V\}$ and a goal test, giving legal combinations of values for subsets of variables. The CSP is called...

binary iff all constraint relate at most two variables.

discrete iff all of the variables have countable domains.

continuous iff it is not discrete.

A CSP has a **factored world representation**. Examples include SuDuKo, Map-Colouring, Timetabling and Scheduling.

Def. 34 (Constraint network). A triple $\langle V, D, C \rangle$ is called a constraint network, where V and D as the same as for **CSPs**, and a set of binary constraints

$$C := \{C_{uv} = C_{vu} \subseteq D_u \times D_v \mid u, v \in V \text{ and } u \neq v\}.$$

Any CSP can be represented by a constraint network.

Def. 35 (Constraint Network Graph). For a **constraint network** $\gamma = \langle V, D, C \rangle$, the graph formed by $\langle V, C \rangle$ is called the "constraint graph" of γ .

Def. 36 (Assignment). A partial assignment for a **constraint network** is a partial function $a: V \mapsto \bigcup_{v \in V} D_v$ if $a(v) \in D_v$ for all $v \in V$. If a is total, a , is just called an “assignment”.

Def. 37 (Consistency). A **partial assignment** a is inconsistent, iff there are variables $u, v \in V$ and a constraint $C_{uv} \in C$ and $(a(u), a(v)) \notin C_{uv}$. Otherwise a is called consistent. A consistent, total assignment is a solution.

Def. 38 (Backtracking on CSPs). A straightforward approach to solve a **CSP** is to incrementally try **assigning** variables until a **consistent** solution is found, backtracking if necessary. To improve the efficiency of this approach, the following heuristics can be applied:

Minimum remaining values Assign the variable with the fewest remaining legal values. This is done to reduce the branching factor of the **search tree**.

Degree heuristic Assign the variable with the most constraints on remaining variables. This is done to detect inconsistencies early on.

Least constraining value When assigning a variable, choose the value that rules out the fewest values from the neighbouring domains.

Def. 39 (Equivalent constraint networks). Two **constraint networks** $\gamma = \langle V, D, C \rangle$ and $\gamma' = \langle V, D', C' \rangle$ are equivalent ($\gamma \equiv \gamma'$) iff they have the same **solutions**.

Def. 40 (Tightness). Let $\gamma = \langle V, D, C \rangle$ and $\gamma' = \langle V, D', C' \rangle$ be two **constraint networks**. γ' is “tighter” than γ ($\gamma' \sqsubseteq \gamma$) iff

1. For all $v \in V$, $D'_v \subseteq D_v$
 2. For all $u, v \in V, u \neq v$ and $C'_{uv} \in C'$, $C'_{uv} \notin C$ or $C'_{uv} \subseteq C_{uv}$
- If at least one of these inclusions are strict, γ' is “strictly tighter”.

An equivalent but tighter **constraint network** is preferable, because it has fewer consistent partial assignments.

Def. 41 (Backtracking with Inference). The general algorithm for backtracking with inference, where $\text{Inference}(\gamma)$ is any procedure that delivering a (**tighter**) equivalent **network**.

Algo. 8 BacktrackingWithInference(γ, a) **returns** a solution, or “inconsistent”

- 1: **if** a is inconsistent **then return** “inconsistent”
 - 2: **if** a is a total assignment **then return** a
 - 3: $\gamma' :=$ a copy of γ $\triangleright \gamma' := \langle V, D', C' \rangle$
 - 4: $\gamma' := \text{Inference}(\gamma')$
 - 5: **if** exists v with $D'_v = \emptyset$ **then return** “inconsistent”
 - 6: select some variable v for which a is not defined
 - 7: **for** each $d \in$ copy of D'_v in some order **do**
 - 8: $a' := a \cup \{v = d\}$ \triangleright makes a explicit as a constraint
 - 9: $D'_v := \{d\}$
 - 10: $a'' := \text{BacktrackingWithInference}(\gamma', a')$
 - 11: **if** $a'' \neq$ “inconsistent” **then return** a''
 - 12: **return** “inconsistent”
-

Def. 42 (Forward checking). For a **constraint network** γ and a **partial assignment** a , propagate information about values from the domains of unassigned variables that are in conflict with the values of already assigned variables to obtain a **tighter** network γ' .

Algo. 9 ForwardChecking(γ, a) **returns** modified γ

- 1: **for** each v where $a(v) = d'$ is defined **do**
 - 2: **for** each u where $a(u)$ is undefined and $C_{uv} \in C$ **do**
 - 3: $D_u := \{d \in D_u \mid (d, d') \in C_{uv}\}$
 - 4: **return** γ
-

Def. 43 (Arc consistency). A variable pair $v, u \in V, v \neq u$ is arc consistent, if $C_{uv} \notin C$ or for every value $d \in D_v$ there exists a $d' \in D_u$ such that $(d, d') \in C_{uv}$.

A **constraint network** γ is arc consistent, if every variable pair $v, u \in V, v \neq u$ is arc consistent.

Arc consistency can be “enforced” by reducing domains. $\text{Revise}(\gamma, v, u)$ enforces arc consistency for v relative to u .

Algo. 10 Revise(γ, v, u) **returns** modified γ

- 1: **for** each $d \in D_v$ **do**
 - 2: **if** there is no $d' \in D_u$ with $(d, d') \in C_{vu}$ **then**
 - 3: $D_v := D_v \setminus \{d\}$
 - 4: **return** γ
-

The AC-3 Algorithm ($\mathcal{O}(mk^3)$, for m constraints and maximal domain size k) applies $\text{Revise}(\gamma, u, v)$ up to a fixed point, remembering potentially inconsistent variable pairs:

Algo. 11 AC-3(γ) **returns** modified γ

- 1: $M := \emptyset$
 - 2: **for** each constraint $C_{uv} \in C$ **do**
 - 3: $M := M \cup \{(u, v), (v, u)\}$
 - 4: **while** $M \neq \emptyset$ **do**
 - 5: remove any element (u, v) from M
 - 6: $\text{Revise}(\gamma, u, v)$
 - 7: **if** D_u has changed in the call to revise **then**
 - 8: **for** each constraint $C_{wu} \in C$ where $w \neq v$ **do**
 - 9: $M := M \cup \{(w, u)\}$
 - 10: **return** γ
-

To solve an acyclic constraint network, enforce arc consistency with AC-3(γ) and run backtracking with inference on the arc consistent network. This will find a solution without having to backtrack.

A simpler algorithm, AC-1(γ) has a runtime of $\mathcal{O}(mk^3n)$, where n is the number of variables.

Def. 44 (Acyclic Constraint Graph). Let $\gamma = \langle V, D, C \rangle$ be a **constraint network** with n variables and maximal domain size k , whose **constraint graph** is acyclic. Then we can find a solution for γ , or prove γ to be inconsistent, in time $\mathcal{O}(nk^2)$:

Algo. 12 AcyclicCG(γ) **returns** solution, or “inconsistent”

- 1: Obtain a directed tree from γ 's constraint graph, picking an arbitrary variable v as the root, and directing arcs outwards.
 - 2: Order the variables topologically, i.e., such that each vertex is ordered before its children; denote that order by v_1, \dots, v_n .
 - 3: **for** $i := n, n-1, \dots, 2$ **do**
 - 4: $\text{Revise}(\gamma, v_{\text{parent}(i)}, v_i)$
 - 5: **if** $D_{v_{\text{parent}(i)}} = \emptyset$ **then return** “inconsistent”
 - 6: Run BacktrackingWithInference with forward checking, using the variable order v_1, \dots, v_n .
-

Def. 45 (Cutset conditioning). Let $\gamma = \langle V, D, C \rangle$ be a **constraint network**, and $V_0 \subseteq V$. V_0 is a “cutset” for γ if the sub-graph of γ 's **constraint graph** induced by $V \setminus V_0$ is acyclic. V_0 is optimal if its size is minimal among all cutsets for γ . The cutset conditioning algorithm computes an optimal cutset:

Algo. 13 CutsetConditioning(γ, V_0, a) returns a solution, or “inconsistent”

```

1:  $\gamma' := \text{ForwardChecking}(\text{a copy of } \gamma, a)$ 
2: if ex.  $v$  with  $D'_v = \emptyset$  then return “inconsistent”
3: if ex.  $v \in V_0$  s.t.  $a(v)$  is undefined then
4:   select such  $v$ 
5: else
6:    $a' := \text{AcyclicCG}(\gamma')$ 
7:   if  $a' \neq$  “inconsistent” then return  $a \cup a'$ 
8:   else return “inconsistent”
9: for each  $d \in$  copy of  $D'_v$  in some order do
10:   $a' := a \cup \{v = d\}$ ;  $D'_v := \{d\}$ 
11:   $a'' := \text{CutsetConditioning}(\gamma', V_0, a')$ 
12: if  $a' \neq$  “inconsistent” then return  $a''$ 
13: else return “inconsistent”

```

4 Logic

Def. 46 (Syntax). Rules to decide what are legal statements (formulas).

Def. 47 (Semantics). $\phi \models A$: Rules to decide whether a formula A is true for a given assignment ϕ .

Def. 48 (Model). Consists of a universe and an interpretation (what connectives “do” and assignments).

Def. 49 (Entailment). If for every model ϕ
 $\phi \models A \Rightarrow \phi \models B$
 B is entailed by A , written $A \models B$.

Def. 50 (Calculus). A set of inference rules.

Def. 51 (Deduction). Statements that can be derived from A using a calculus \mathcal{C} (calculus), written $A \vdash_{\mathcal{C}} B$.

Def. 52 (Soundness). A calculus \mathcal{C} is sound if for all formulas A, B it is true that $A \vdash_{\mathcal{C}} B \Rightarrow A \models B$.

Def. 53 (Complete). A calculus \mathcal{C} is complete if for all formulas A, B it is true that $A \models B \Rightarrow A \vdash_{\mathcal{C}} B$.

Def. 54 (Logical System). A logical system is a triple $\langle \mathcal{L}, \mathcal{K}, \models \rangle$, where \mathcal{L} is a formal language, \mathcal{K} is a set and $\models \subseteq \mathcal{K} \times \mathcal{L}$.

For a model $\mathcal{M} \in \mathcal{K}$ and formula $A \in \mathcal{L}$, we call $A \dots$

satisfied by \mathcal{M} , iff $\mathcal{M} \models A$

falsified by \mathcal{M} , iff $\mathcal{M} \not\models A$

satisfiable in \mathcal{K} , iff “ $\exists \mathcal{M} \in \mathcal{K}. \mathcal{M} \models A$ ”

valid in \mathcal{K} (written $\models A$), iff “ $\forall \mathcal{M} \in \mathcal{K}. \mathcal{M} \models A$ ”

falsifiable in \mathcal{K} , iff “ $\exists \mathcal{M} \in \mathcal{K}. \mathcal{M} \not\models A$ ”

unsatisfiable in \mathcal{K} , iff “ $\forall \mathcal{M} \in \mathcal{K}. \mathcal{M} \not\models A$ ”

Def. 55 (Propositional logic, PL^0). $\text{wff}_o(\mathcal{V}_o)$ is the set of “well-formed” (syntactically correct) formulas with variables \mathcal{V}_o . Its model $\langle \mathcal{D}_o, \mathcal{I} \rangle$ consists of a universe $\mathcal{D}_o = \{\text{T}, \text{F}\}$ and an interpretation \mathcal{I} , that assigns connectives values. The value function $\mathcal{I}_\phi: \text{wff}_o(\mathcal{V}_o) \rightarrow \mathcal{D}_o$, assigns values to formulas.

PL^0 is an example for a logical system $\langle \text{wff}_o(\mathcal{V}_o), \mathcal{K}, \models \rangle$, where \mathcal{K} is the set of variable assignments, and $\phi \models A \iff \mathcal{I}_\phi(A) = \text{T}$.

Def. 56 (First order logic, FOL, PL^1). $\text{wff}_l(\Sigma_l)$ is the set of “well-formed” terms over a signature Σ_l (function and skolem constants — individuals). $\text{wff}_o(\Sigma)$ is the set of well-formed propositions over a signature Σ (Σ_l plus connectives and predicate constants — truth values).

Def. 57 (Natural deduction, $\mathcal{N}^{\mathcal{D}1}$). A “natural deduction” calculus for First order Logic:

$$\begin{array}{c}
\frac{A \quad B}{A \wedge B} \wedge I \qquad \frac{}{A = A} = I \qquad \frac{A}{\forall X.A} \forall I \\
\frac{A \wedge B}{A} \wedge E_l \qquad \frac{A = B \quad C[A]_p}{[B/p](C)} = E \\
\frac{A \wedge B}{B} \wedge E_r \qquad \frac{[B/X](A)}{\exists X.A} \exists I \qquad \frac{\forall X.A}{[B/X](A)} \forall E \\
\frac{A}{B} \Rightarrow I \\
\frac{A \Rightarrow B \quad A}{B} \Rightarrow E \qquad \frac{\exists X.A \quad \frac{[c/X](A)}{B}}{B} \exists E \qquad \frac{}{A \vee \neg A} \text{TND}
\end{array}$$

Def. 58 (Analytical tableaux). A tableau calculus for First order Logic: Every formula is labelled as either true (A^T) or false (A^F). To satisfy a formula A^α , it has to be shown that A has a truth value of α . This is done by branching out using the rules below. A branch is closed if it contains F, else open. A tableau is closed (\neq saturated) if all of its branches are closed. A is valid iff there is a closed tableau with A^F at the root.

$$\begin{array}{c}
\frac{A \wedge B^T}{A^T \quad B^T} \mathcal{T}_0 \wedge \qquad \frac{\forall X.A^T \quad C \in \text{cwff}_l(\Sigma_l)}{[C/X](A)^T} \mathcal{T}_1 \forall \\
\frac{A \wedge B^F}{A^F | B^F} \mathcal{T}_0 \vee \qquad \frac{\forall X.A^F \quad c \in (\Sigma_0^{\text{sk}} \setminus \mathcal{H})}{[c/X](A)^F} \mathcal{T}_1 \exists \\
\frac{\neg A^T}{A^F} \mathcal{T}_0 \neg \\
\frac{\neg A^F}{A^T} \mathcal{T}_0 \neg \\
\frac{A \Rightarrow B^T}{A^F | B^T} \mathcal{T}_0 \Rightarrow \\
\frac{A \Rightarrow B^F}{A^T \quad B^F} \mathcal{T}_0 \Rightarrow \\
\frac{A^\alpha \quad B^\beta \quad \alpha \neq \beta, \sigma(A) = \sigma(B)}{F : \sigma} \mathcal{T}_1 F \qquad \frac{A^T \quad A \Rightarrow B^T}{B^T} \mathcal{T}_0 \neg
\end{array}$$

Def. 59 (FOL Unification). For two terms \mathbf{A} and \mathbf{B} , a unification is the problem of finding a substitution σ , s.t. $\sigma(\mathbf{A}) = \sigma(\mathbf{B})$. A substitution σ is “more general” than θ , if there is a substitution φ , s.t. $\theta = \varphi \circ \sigma[W]$. There is no more general unifier than the “most general unifier” (*mgu*).

Def. 60 (Conjunctive Normal Form (CNF)). A formula is in conjunctive normal form if it is a conjunction of disjunction of literals.

For a FOL formula, it can be computed as follows:

1. Rewrite implications $p \Rightarrow q$ into the form $\neg p \vee q$.
2. Move negations inwards, so that only predicates are negated.
3. Rename variables bound by quantifiers making them unique.
4. Replace variables bound by existential quantifiers with new “skolem functions” $f \in \Sigma_k^{\text{sk}}$ over all the free variables X_1, \dots, X_k in the quantified term:

$$\forall X.A \rightarrow [f(X_1, \dots, X^k)/X](A)$$
5. Distribute \vee inwards over \wedge :

$$A \vee (B \wedge C) \rightarrow (A \vee B) \wedge (A \vee C)$$

Def. 61 (FOL Resolution). The resolution calculus for FOL operates on the CNF of a formula.

Like Tableau, it shows shows $\neg T \vdash F$ to prove T . T is transformed into CNF and manipulated using the rules below. If the empty disjunction (“clause set”, \square) is derived, T has been refuted.

$$\frac{P^T \vee A \quad P^F \vee B}{A \vee B}$$

$$\frac{P^T \vee A \quad P^F \vee B \quad \sigma = \text{mgu}(P, Q)}{\sigma(A) \vee \sigma(B)}$$

$$\frac{A^\alpha \vee B^\alpha \vee C \quad \sigma = \text{mgu}(A, B)}{\sigma(A) \vee \sigma(C)}$$

Def. 62 (DPLL). The DPLL procedure is an algorithm to find an interpretation satisfying a clause set.

4.1 Logic Programming

Def. 63 (Fact). A term that is unconditionally true.

Def. 64 (Rule). A term that is true if certain premises are true.

Def. 65 (Clause). **Facts** and **rules** are both clauses.

Def. 66 (Horn clause). A horn clause is a clause with at most one positive literal.

The Prolog rule $H :- B_1, \dots, B_n$ is the implication $B_1 \wedge \dots \wedge B_n \Rightarrow H$ can be written as a horn clause $\neg B_1 \vee \dots \vee \neg B_n \vee H$.

4.2 Knowledge Representation

Def. 67 (Semantic Network). A directed graph representing knowledge. It consist of nodes representing objects/concepts, and edges representing relations between these, also called “links”.

Def. 68 (Isa/Inst). Links may be labelled with “isa” (*is a*) or “inst” (*instance*) to designate concept inclusion or concept membership respectively. They propagate properties encoded by other links.

Def. 69 (TBox). The sub-graph of a **semantic network** between concepts is called terminology, or TBox. It is spanned by **isa** links.

Def. 70 (ABox). The sub-graph of a **semantic network** between objects is called assertions, or ABox. It is spanned by **inst** links and relations between objects.

Def. 71 (Semantic Web). A collaborative movement led by the W3C promoting inclusion of semantic content into web pages. One example is RDF (Resource Description Framework), used for describing resources on the web.

Def. 72 (Ontology). A **logical system** $\langle \mathcal{L}, \mathcal{K}, \models \rangle$ and “concept axioms” about individuals, concept and relations. **Semantic networks** are ontologies.

Def. 73 (Description Logic). A formal system for talking about sets and their relations. A description logic \mathcal{D} has a **D-ontology**, consisting of a **TBox** and **ABox**.

Def. 74 (\mathcal{ALC}). A **description logic** more expressive than PL^0 , but less complex than **FOL**. It relates “Concepts” (classes of objects, C) with “Roles” (binary relations, R). Its **Syntax** is as follows:

$$F_{\mathcal{ALC}} := C \mid \top \mid \perp \mid \overline{F_{\mathcal{ALC}}} \mid F_{\mathcal{ALC}} \sqcap F_{\mathcal{ALC}} \mid F_{\mathcal{ALC}} \sqcup F_{\mathcal{ALC}} \mid \\ \exists R.F_{\mathcal{ALC}} \mid \forall R.F_{\mathcal{ALC}}$$

where \top and \perp are the special concepts designating “all” and “none” respectively.

Def. 75 (\mathcal{ALC} Tableau Calculus). The **Tableau calculus** for \mathcal{ALC} :

$$\frac{x:c \quad x:\bar{c}}{\perp} \mathcal{T}_\perp \quad \frac{x:\phi \sqcup \psi}{x:\phi \mid x:\psi} \mathcal{T}_\sqcup \quad \frac{x:\exists R.\phi}{xRy \quad y:\phi} \mathcal{T}_\exists \\ \frac{x:\phi \sqcap \psi}{x:\phi \quad x:\psi} \mathcal{T}_\sqcap \quad \frac{x:\forall R.\phi \quad xRy}{y:\phi} \mathcal{T}_\forall$$

5 Planning

Def. 76 (Planning language/task). A logical description of the components of a **search problem**:

- a set of possible states
- an initial state I
- a goal condition G
- a set of actions A in terms of preconditions and effects, constituting a planning task. This approach allows a solver to gain insight into the problem structure, resulting in a **structured** world representation.

Def. 77 (Satisficing planning). A procedure that takes as input a **planning problem** and outputs a plan or “unsolvable”, if no such plan exists.

Def. 78 (Optimal planning). A procedure that takes as input a **planning problem** and outputs an optimal plan or “unsolvable”, if no such plan exists.

Def. 79 (STRIPS planning task). This is an encoding of a **planning problem** using a quadruple $\Pi = \langle P, A, I, G \rangle$ where

- P is a finite set of facts
- A is a finite set of actions, each given as a triple of “preconditions”, an “add list” and a “delete list”.
- $I \subseteq P$ is the initial state
- $G \subseteq P$ is the goal.

Satisficing planning for STRIPS is called “PlanEx”, and **optimal planning** is called “PlanLen”, that tries to find the shortest plan.

A **heuristic** for Π with states S is function $h: S \mapsto \mathbb{N} \cup \infty$ so that $h(g) = 0$ for a goal state g . The perfect heuristic h^* assigns every $s \in S$ the length of the shortest path to g or ∞ if non-existent.

Def. 80 (Partial Order Planning). A partially ordered plan is a collection of causal links $S \xrightarrow{p} T$ and temporal ordering $S \prec T$ where p is an effect of S and precondition of T . If the causal links and temporal ordering induce a partial ordering, it is called “consistent”. If every precondition is achieved, it is called “complete”.

Partial order planning is the process of computing a complete and consistent partially order plan.

Def. 81 (Delete relaxation). This is a **relaxation** Π^+ of a given **STRIPS task** Π all actions have empty delete lists.

Def. 82 (Relaxed plan). For a **STRIPS task** $\Pi = \langle P, A, I, G \rangle$ and state s , then $\langle P, A, \{s\}, G \rangle^+$ is a **relaxed plan** for I/Π .

PlanEx for relaxed problems is called PlanEx⁺.

Def. 83 (h^+ -heuristic). For a **planning task** $\Pi = \langle P, A, I, G \rangle$, the optimal heuristic calculates the length of the optimal **relaxed plan** for s or ∞ if no plan exists. h^+ is admissible. The heuristic h^{FF} approximates h^+ , since calculating h^+ is in NP.

Def. 84 (Real World Planning). When planning in real-world situations, the **agent** the **task environment** is **partially observable** and **non-deterministic**, which invalidates the previous assumptions. Variations on **planning** try to overcome this:

Conditional Extend the possible action in plans by **conditional steps** that execute sub-plans conditionally.

Conformant Tries to find a plan without sensing, instead relying on the its (fully observable) belief states.

Contingent Generate a plan with conditional branching based on percepts.

Def. 85 (Online Search). Interleaving of search and actions, basing action on incoming perceptions. A planner P can be turned into an online problem server by adding an action $\text{Replan}(g)$, that re-starts P in the current state with goal g .

6 Probability Theory

Def. 86 (Probability Model $\langle \Omega, P \rangle$). consists of a countable sample space Ω and a probability function $P: \Omega \rightarrow [0;1]$, s.t. $\sum_{\omega \in \Omega} P(\omega) = 1$

Def. 87 (Event). When a random variable X takes on a value x .

Def. 88 (Conditional/Posterior Probability). The probability

$$P(a|b) = \frac{P(a \wedge b)}{P(b)},$$

i.e. the chance that event “a” takes place, given the event “b”.

Def. 89 (Conditional Independence). Two events a and b are conditionally independent, if $P(a \wedge b | c) = P(a | c)P(b | c)$.

Def. 90 (Probability Distribution). A vector for $\mathbf{P}(X)$ relating each element of the **sample space** to a probability:

$$\langle P(\omega_1), \dots, P(\omega_n) \rangle.$$

Related concepts:

Joint PD Given $Z \subseteq \{X_1, \dots, X_n\}$, results in a array the probabilities of all **events**.

Full joint PD Joint PD for all random variables.

Conditional PD Given X and Y , results in a table for every probability $P(X|Y)$.

Def. 91 (Product Rule). $P(a \wedge b) = P(a|b)P(b)$

Def. 92 (Chain Rule). Extension of the **product rule**,
 $P(X_1, \dots, X_n) = P(X_n | X_{n-1}, \dots, X_1) \dots P(X_2 | X_1) P(X_1)$

Def. 93 (Marginalisation). $\mathbf{P}(\mathbf{X}) = \sum_{y \in \mathbf{Y}} \mathbf{P}(\mathbf{X}, y)$

Def. 94 (Normalisation). Given $\mathbf{P}(X|e)$, and a normalization constant

$$\alpha = \frac{1}{P(x_1|e) + \dots + P(x_n|e)},$$

normalization scales each element of the probability distribution s.t. $\sum \alpha \mathbf{P}(X|e) = 1$

Def. 95 (Bayes’ Rule). Given two propositions a and b ,

$$P(a|b) = \frac{P(b|a)P(a)}{P(b)},$$

where $P(a) \neq 0$ and $P(b) \neq 0$.

Def. 96 (Naive Bayes’ Model). In this model, the **full joint probability distribution** is

$$\mathbf{P}(c | e_1, \dots, e_n) = \mathbf{P}(c) \prod_i \mathbf{P}(e_i | c),$$

i.e. a *single* cause c influences a number of **cond. independent** effects e_i .

6.1 Bayesian Networks

Def. 97 (Bayesian Network). A directed, acyclic graph, where each node corresponds to a random variable, connected by links designating “parent” variables.

A “diagnostic” link points from cause to effect, a “causal” from effect to cause.

Def. 98 (Conditional Probability Table). A table specifying the probability for each node of a **Bayesian network** given the values of the parent variables.

Def. 99 (Constructing Bayesian Network). Given any fixed order of variables X_1, \dots, X_n a **Bayesian network** can be constructed by iteratively finding a minimal set $\text{Parent}(X_i) \subseteq \{X_1, \dots, X_i\}$ s.t. $\mathbf{P}(X_i | X_{i-1}, \dots, X_1) = \mathbf{P}(X_i | \text{Parent}(X_i))$, linking X_i with every $X_j \in \text{Parent}(X_i)$ and associating a **conditional probability table** with X_i corresponding to $\mathbf{P}(X_i | \text{Parent}(X_i))$.

Def. 100 (Probabilistic Inference Task). Given a **Bayesian network** \mathcal{B} , calculating $\mathbf{P}(\mathbf{X} | \mathbf{e})$, where \mathbf{X} is a set of “query variables” and \mathbf{E} is a set of “evidence variables” assigned by an event \mathbf{e} . The remaining variables \mathbf{Y} are referred to as “hidden”.

This problem can be solved using “Inference by Enumeration”:

1. **Normalize** and **marginalize**:

$$\mathbf{P}(X|\mathbf{e}) = \alpha \begin{cases} \mathbf{P}(X, \mathbf{e}) & \text{if } \mathbf{Y} = \emptyset \\ \sum_{\mathbf{y} \in \mathbf{Y}} \mathbf{P}(X, \mathbf{e}, \mathbf{y}) & \text{otherwise} \end{cases}$$

2. **Chain rule**, by ordering X_1, \dots, X_n to be consistent with \mathbf{B} :

$$\mathbf{P}(X_1, \dots, X_n) = \mathbf{P}(X_n | X_{n-1}, \dots, X_1) \dots \mathbf{P}(X_2 | X_1) \mathbf{P}(X_1)$$

3. Exploit **conditional independence**:

$$P(X_i | X_{i-1}, \dots, X_1) = P(X_i | \text{Parents}(X_i))$$

Alternative methods such as “Variable Elimination” avoid redundant computations by using dynamic programming.

7 Making Simple Decisions Rationally

Def. 101 (Expected Utility). Given a **state-evaluation function** $U: S \mapsto \mathbb{R}^+$, the expected utility of an action a given evidence \mathbf{e} is

$$EU(a|\mathbf{e}) = \sum_{s' \in S} P(R(a) = s' | a, \mathbf{e}) U(s'),$$

that rational **agents** attempt to maximize, where $R(a)$ is the result of the action a .

Def. 102 (Preference). Given two states, one can say that an **agent** prefers A over B ($A \succ B$), is indifferent ($A \sim B$) or does not prefer B over A ($A \succeq B$).

A preference is rational if orderability, transitivity, continuity, substitutability, monotonicity and decomposability all hold for the relation. If given, there must be a **utility function** U with

$$U(A) \geq U(B) \iff A \succeq B$$


and

$$U([p_1, S_1; \dots; p_n, S_n]) = \sum_i p_i U(S_i),$$

for states S_i and probabilities p_i (Ramsey’s Theorem).

Def. 103 (Micromort). A unit of utility, describing a $1/10^{-6}$ chance of death.

Def. 104 (Decision Network). A **Bayesian network** with “action” and “utility” nodes, to aid with rational decisions. An example (here “umbrella” is an action and “utility” is an utility node):



To decide, maximize the **expected utility** for the utility nodes by comparing every value for action nodes.

Def. 105 (Value of perfect Information). For a random variable F over D , the VPI given evidence E is

$$\text{VPI}_E(F) := \left(\sum_{f \in D} P(F = f | E) EU(\alpha_f | E, F = f) \right) - EU(\alpha | E),$$

where

$$EU(\alpha | E) = \max_{a \in A} \left(\sum_{s \in S_a} U(s) P(s | E, a) \right),$$

$$EU(\alpha_f | E, F = f) = \max_{a \in A} \left(\sum_{s \in S_a} U(s) P(s | E, a, F = f) \right),$$

$$\alpha_f = \operatorname{argmax}_{a \in A} EU(a | E, F = f).$$

8 Temporal Models

Def. 106 (Markov Property). The variable \mathbf{X}_t only depends on a subset $\mathbf{X}_1, \dots, \mathbf{X}_{t-1}$ ($\mathbf{X}_{0:t-1}$). The n th-order *Markov Property* is given when $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{0:t-1}) = \mathbf{P}(\mathbf{X}_t | \mathbf{X}_{t-n:t-1})$.

Def. 107 (Markov Process). A sequence of random variables with the **Markov Property**. The variables can be divided into “state variables” X_t and “evidence variables” E_t .

Given the initial prior probability $\mathbf{P}(\mathbf{X}_0)$, the **full joint probability distribution** can be computed as

$$P(X_{0:t}, E_{1:t}) = P(X_0) \prod_{i=0}^t P(X_i | X_{i-1}) P(E_i | X_i).$$

Def. 108 (Transition Model). A transition model of a **Markov Process** is given by $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{t-1})$. If $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{t-1})$ is the same for all t , the process is said to be “stationary”, making the model finite in size.

Def. 109 (Sensor Model). A sensor model of a **Markov Process** predicts the influence of percepts on the belief state. It the “sensor Markov property” iff $\mathbf{P}(\mathbf{E}_t | \mathbf{X}_{0:t}, E_{1:t-1}) = \mathbf{P}(\mathbf{E}_t | \mathbf{X}_t)$

Def. 110 (Filtering). Computing the belief state from prior experience by dividing up the evidence (1), using **Bayes’ rule** (2) and applying the **sensor Markov property** (3),

$$\begin{aligned} \mathbf{P}(\mathbf{X}_t | \mathbf{e}_{1:t}) &= \mathbf{P}(\mathbf{X}_t | \mathbf{e}_{1:t-1}, \mathbf{e}_t) & (1) \\ &= \alpha \mathbf{P}(\mathbf{e}_t | \mathbf{X}_t, \mathbf{e}_{1:t-1}) \mathbf{P}(\mathbf{X}_t | \mathbf{e}_{1:t-1}) & (2) \\ &= \alpha \underbrace{\mathbf{P}(\mathbf{e}_t | \mathbf{X}_t)}_{\text{transit. model}} \underbrace{\mathbf{P}(\mathbf{X}_t | \mathbf{e}_{1:t-1})}_{\text{recursion}} & (3) \end{aligned}$$

Def. 111 (Prediction). Computing a future state distribution, i.e. **filtering** without new evidence:

$$\mathbf{P}(\mathbf{X}_{t+k+1} | \mathbf{e}_{1:t}) = \sum_{\mathbf{x}_{t+k}} \mathbf{P}(\mathbf{X}_{t+k+1} | \mathbf{x}_{t+k}) \underbrace{\mathbf{P}(\mathbf{x}_{t+k} | \mathbf{e}_{1:t})}_{\text{recursion}}$$

where $0 < k$.

Def. 112 (Smoothing). Improving a past belief state, using **Bayes’ rule** (1) and **cond. independence** (2),

$$\begin{aligned} \mathbf{P}(\mathbf{X}_k | \mathbf{e}_{1:t}) &= \mathbf{P}(\mathbf{X}_k | \mathbf{e}_{1:k}, \mathbf{e}_{k+1:t}) \\ &= \alpha \mathbf{P}(\mathbf{X}_k | \mathbf{e}_{1:k}) \mathbf{P}(\mathbf{X}_{k+1} | X_k, \mathbf{e}_{1:k}) & (1) \\ &= \alpha \mathbf{P}(\mathbf{X}_k | \mathbf{e}_{1:k}) \underbrace{\mathbf{P}(\mathbf{X}_{k+1} | X_k)}_{\text{recursion}} & (2) \end{aligned}$$

where $0 \leq k < t$.

Def. 113 (Most likely Explanation). Used to explain the what is the most probable sequence of **events**, that caused the perceived evidence $\max_{\mathbf{x}_1, \dots, \mathbf{x}_n} \mathbf{P}(\mathbf{x}_1, \dots, \mathbf{x}_n, \mathbf{X}_{t+1} | \mathbf{e}_{1:t+1})$, by calculating:

$$\mathbf{P}(\mathbf{e}_{t+1} | \mathbf{X}_{t+1}) \max_{\mathbf{x}_t} \left(\mathbf{P}(\mathbf{X}_{t+1} | \mathbf{x}_t) \max_{\mathbf{x}_1, \dots, \mathbf{x}_n} \mathbf{P}(\mathbf{x}_1, \dots, \mathbf{x}_{t-1}, \mathbf{x}_t | \mathbf{e}_{1:t}) \right).$$

Def. 114 (Hidden Markov Model). A **Markov model** with a single state variable $X_t \in \{1, \dots, S\}$ and a single evidence variable.

Using a transition matrix $T_{ij} = P(X_t = j | X_{t-1} = i)$ and $O_{tii} = P(e_t | X_t = i)$ one can reinterpret Markov inference:

$$\begin{aligned} \text{HMM filtering equation} & \mathbf{f}_{1:t+1} = \alpha (\mathbf{O}_{t+1} \mathbf{T}^\top \mathbf{f}_{1:t}) \\ \text{HMM smoothing equation} & \mathbf{b}_{k+1:t} = \mathbf{TO}_{k+1} \mathbf{b}_{k+2:t} \end{aligned}$$

Def. 115 (Dynamic Bayesian Network). A **Bayesian network** with random variables indexed by a time structure, and can therefore be seen as a network with an infinite number of variables.

9 Making Complex Decisions Rationally

Def. 116 (Sequential Decision Problem). The **agent’s** utility depends on a sequence of decisions, incorporating utilities, uncertainty and sensing.

Def. 117 (Markov Decision Problem). A **sequential decision problem** consisting of a set S of states, A of actions, a transition model $P(s' | s, a)$ and a reward function $R : S \mapsto \mathbb{R}$, in a fully observable, stochastic environment. The goal is to find an optimal policy $\pi : S \mapsto A$, mapping every state to the best action.

Def. 118 (Stationary Preferences). **Preferences** are called “stationary” iff

$$[r, r_0, r_1, \dots] \succ [r', r'_0, r'_1, \dots] \iff [r_0, r_1, \dots] \succ [r'_0, r'_1, \dots]$$

The only ways to combine these over time is

additive $U([s_0, s_1, \dots]) = \sum_i R(s_i)$

discounted $U([s_0, s_1, \dots]) = \sum_i \gamma^i R(s_i)$, where γ is called “discount factor”.

Def. 119 (Utility of States). The optimal policy $\pi^* = \pi_s^* = \max_{\pi} U^\pi(s)$ where for a given policy π and s_t being the state an agent reaches at time t

$$U^\pi(s) := E \left[\sum_{t=0}^{\infty} \gamma^t R(s_t) \right],$$

the utility $U(s)$ of a state s is $U^{\pi^*}(s)$.

Def. 120 (Bellman Equation). The **utilities of states** is given by the solution to the equation

$$U(s) = R(s) + \gamma \max_{a \in A(s)} \left(\sum_{s'} U(s') P(s' | s, a) \right)$$

Def. 121 (Value Iteration Algorithm). A method to find the fix-point of the **Bellman equation**:

Algo. 14 ValueIteration(mdp, ϵ) returns a utility fn.

Input mdp a **MDP** ($S, A(s), P(s' | s, a), R(s), \gamma$), the maximum permitted error ϵ

```

1: repeat
2:    $U := U', \delta := 0$ 
3:   for each state  $s$  in  $S$  do
4:      $U'[s] := R(s) + \gamma \max_{a \in A(s)} \sum_{s'} U[s'] P(s' | s, a)$ 
5:      $\delta := \max\{|U'[s] - U[s]|, \delta\}$ 
6: until  $\delta < \epsilon(1 - \gamma) / \gamma$ 

```

Def. 122 (Policy Iteration Algorithm). Algorithm for iteratively evaluating and improving policies until no changes are made:

Algo. 15 PolicyIteration(mdp) returns a policy

Input mdp a **MDP** ($S, A(s), P(s' | s, a), R(s), \gamma$)

```

1: repeat
2:    $U := \text{PolicyEvaluation}(\pi, U, mdp)$ 
3:   unchanged := true
4:   for each state  $s$  in  $S$  do
5:      $a^* := \text{argmax}_{a \in A(s)} (\sum_{s'} P(s' | s, a) U(s'))$ 
6:     best :=  $\sum_{s'} P(s' | s, a^*) U(s')$ 
7:     if best >  $\sum_{s'} P(s' | s, \pi[s']) U(s')$  then
8:        $\pi[s] := a^*, \text{unchanged} := false$ 
9: until unchanged  $\triangleright U$  satisfies Bellman equation

```

Def. 123 (Partially Observable MDP). A **MDP** with a **sensor model** O that is stationary ($O(s, e) = P(e | s)$), i.e. the agent

does not know in what state it is. To update the belief state the agent calculates

$$b'(s') = \alpha P(e|s') \sum_s P(s'|s,a) b(s)$$

where a is the action taken. An agent searches through the belief state by executing the best assumed action, and updating the belief state based on the percept e .

Def. 124 (Dynamic Decision Network). The extension of a **DBN** by **action and utility** nodes.

10 Machine Learning

Def. 125 (Inductive Learning). Learning by examples of the form (x,y) , where x is an “input sample and y a “classification”. The set of examples S is called consistent if a function.

An inductive learning problem $\mathcal{P} = \langle \mathcal{H}, f \rangle$ attempts to find a hypothesis $h \in \mathcal{H}$ for a consistent training set f ($f \simeq h|_{\text{dom}(f)}$).

Def. 126 (Decision Tree). A tree that given examples described by attribute (“attribute-based representation”), labels non-leaf nodes with attribute-choices and leafs with classifications.

Having more “significant” attributes closer to the root helps generate a compact tree. The act of trying to find a smaller tree is called “decision tree learning”.

Def. 127 (Entropy of the Prior). The information of an answer to the prior probabilities $\langle p_1, \dots, p_n \rangle$ is

$$I(\langle p_1, \dots, p_n \rangle) = - \sum_{i=1}^n p_i \log_2(p_i).$$

Def. 128 (Information Gain). Said for testing an attribute A ,

$$\text{Gain}(A) = I(\mathbf{P}(C)) - \sum_a P(A=a) I(\mathbf{P}(C|A=a)),$$

and given previous results $B_1 = b_1, \dots, B_n = b_n$ is

$$\text{Gain}(A|b) = I(\mathbf{P}(C|b)) - \sum_a P(A=a|b) I(\mathbf{P}(C|a,b)),$$

where C is a classification with an estimate of the probability distribution, e.g.

$$\mathbf{P}(C) = \left\langle \frac{p}{p+n}, \frac{n}{p+n} \right\rangle,$$

for p positive and n negative examples.

Def. 129 (Learning Curve). Percentage of correct test results as a function of the training set size. May encounter difficulties, if the to be approximated function is not in the hypothesis space.

Def. 130 (Overfitting). When a hypothesis h assumes an error is significant part of the underlying data. Conversely, “underfitting” occurs when h cannot capture the underlying trend of the data.

Def. 131 (Decision Tree Pruning). For learned **decision trees**: Repetitively finding test nodes and replacing it with leaf nodes if it has a low **information gain**, as determined by a statistical significance test.

Def. 132 (Generalization). Given a set of examples \mathcal{E} and a prior probability $\mathbf{P}(X,Y)$ the *Generalized Loss* is

$$\text{GenLoss}_L(h) := \sum_{(x,y) \in \mathcal{E}} L(y, h(x)) P(x,y),$$

where L is a loss function, quantifying the lost utility by a hypothesis h such as

absolute value $L_1(y, \hat{y}) = |y - \hat{y}|$

squared error $L_2(y, \hat{y}) = (y - \hat{y})^2$

0/1 $L_{0/1}(y, \hat{y}) = 0$ if $y = \hat{y}$ otherwise, 1

Def. 133 (PAC learning). Any algorithm that returns a probably approximatly correct hypothesis, ie. that after a sufficiently large training set, it is unlikely to be seriously wrong. The “error rate” function

$$\text{error}(h) := \text{GenLoss}_{L_{0/1}}(h)$$

describes the probability that h will misclassifying a new example.

Def. 134 (Decision List). A sequence of **literal conjunctions** each specifying the value to be returned if satisfied, or continuing on to the next test otherwise.

Def. 135 (Classification and Regression). An **inductive learning problem** $\langle \mathcal{H}, f \rangle$ is a *classification* problem, iff $\text{codom}(f)$ is discrete and *regression* otherwise.

Def. 136 (Linear Regression). Given a weight vector $\mathbf{w} = (w_0, w_1)$ and $h_w(x) = w_1 x + w_0$, the task of finding the best \mathbf{w} for a set of examples is called *linear regression*. The space of weight combinations is called the *weight space*.

Def. 137 (Gradient Descent). An algorithm for finding the minimum of a continuous function f by **hill climbing** in the direction of steepest descent. For each vector component the calculation

$$w_i \leftarrow w_i - \alpha \frac{\partial}{\partial w_i} f(\mathbf{w})$$

until \mathbf{w} converges, where α is called the “learning rate”.

Def. 138 (Perceptron Learning Rule). A learning rule given an example (\mathbf{x}, y) updating

$$w_i \leftarrow w_i + \alpha(y - h_{\mathbf{w}}(\mathbf{x})) x_i$$

Def. 139 (Logistic Regression). Using a “softer” learning rule, **linear regression** can be replaced by using a logistic function

$$h_{\mathbf{w}}(\mathbf{x}) = \frac{1}{1 + e^{-\mathbf{w}\mathbf{x}}}.$$

10.1 Artificial Neuronal Networks

Def. 140 (Neuronal Network). A directed graph, propagating *activations* a_i from *unit* i to *unit* j via *links* with *weights* $w_{i,j}$.

If the network has cycles, it is called “recurrent”, otherwise “feed-forward” and it is said to have layers $\{L_0, \dots, L_n\}$

Def. 141 (McCulloch-Pitts unit). A unit model where each activation is computed by

$$a_i = g \left(\sum_j w_{j,i} a_j \right),$$

given a activation function g . If g is a threshold function (e.g. **logistic function**) the unit is called a “perceptron unit”.

Def. 142 (Perceptron Network). A feed-forward network of **perceptron units**. Units not part of the input or output layer are called “hidden”.

Def. 143 (Backpropagation). **Learn** in a **perceptron network** is implemented by updating weights

$$w_{k,j} \leftarrow w_{k,j} + \alpha a_k \Delta_j$$

where

$$\Delta_j \leftarrow g' \left(\sum_j w_{j,i} a_j \right) \sum_i w_{j,i} \Delta_i.$$

Def. 144 (Backpropagation Algorithm). Given a network and a set of examples, the algorithm **propagates the example input** through the network, then **propagates deltas backwards** towards the input layer and then updates every weight using the calculated delta values.

10.2 Statistical Learning

Def. 145 (Bayesian Learning). Calculating the probabilities of each hypothesis, and acting upon these predictions

$$P(\mathbf{d} | h_i) = \alpha P(\mathbf{d} | h_i) P(h_i),$$

where $P(\mathbf{d} | h_i)$ is the likelihood to observe data $\mathbf{d} \in \mathbf{D}$ given a hypothesis, and $P(h_i)$ is the “hypothesis prior”.

To predict a unknown quantity X , one uses

$$P(X | \mathbf{d}) = \sum_i P(X | h_i) P(h_i | \mathbf{d}).$$

Def. 146 (Naive Bayes’ Models for Learning). Using a **naive Bayes’ model**, a single “class” variable C is predicted given “attribute” variables X_i :

$$P(C | X_1, \dots, X_n) = \alpha P(C) \prod_i P(x_i | C),$$

whereafter the most likely class is chosen.

11 Natural Language Processing

Def. 147 (Natural Language Processing). The intersection between computer science, artificial intelligence and linguistics, attempting to understand and generate natural language.

Def. 148 (Linguistically Realized). When a piece of information i can be traced back to a fragment of an utterance U .

Def. 149 (Language Model). A probability distribution of a sequence of characters or words.

Def. 150 (Text Corpus). A large, structured set of texts, used for statistical analysis.

Def. 151 (n -gram model). A $n - 1$ ’th order **Markov chain**, that generates a character/word sequence (n -gram) of the length n . The probability of a character sequence $\mathbf{c}_{1:N}$ is

$$P(\mathbf{c}_{1:N}) = \prod_{i=1}^N P(c_i | \mathbf{c}_{1:i-1}).$$

This can be used for genre classification, named entity recognition, language generation, among other things.

For language identification, the most probable language ℓ^* of a **text corpus** can be approximated by applying **Bayes’ rule** (1) and the **Markov property** (2):

$$\begin{aligned} \ell^* &= \operatorname{argmax}_{\ell} (P(\ell | \mathbf{c}_{1:N})) \\ &= \operatorname{argmax}_{\ell} (P(\ell) P(\mathbf{c}_{1:N} | \ell)) \end{aligned} \quad (1)$$

$$= \operatorname{argmax}_{\ell} \left(P(\ell) \prod_{i=1}^N P(c_i | \mathbf{c}_{1-n:i-1}) \right), \quad (2)$$

given the, if necessary estimated, prior probabilities for $P(\ell)$.

Def. 152 (Term Frequency). The number of times a word t occurs in a document d ($\operatorname{tf}(t, d)$).

Def. 153 (Inverse document frequency). Calculated for a document collection $D = \{d_1, \dots, d_n\}$:

$$\operatorname{idf}(t, D) = \log_{10} \left(\frac{n}{|\{d \in D | t \in d\}|} \right)$$

Def. 154 (Term Frequency/Inverse Document Frequency). Combining **term frequency** and **inverse document frequency** into

$$\operatorname{tfidf}(t, d, D) = \operatorname{tf}(t, d) \operatorname{idf}(t, D).$$

Def. 155 (Word Embeddings). A mapping of words into a \mathbb{R}^n vector space. For **if-idf** it is given by

$$e: t \mapsto \langle \operatorname{tfidf}(t, d_1, D), \dots, \operatorname{tfidf}(t, d_{\#(D)}, D) \rangle.$$

Def. 156 (Cosine Similarity). Calculated for two vectors A and B , and the angle θ between them,

$$\cos \theta = \frac{A \cdot B}{\|A\|_2 \|B\|_2}$$

holds. Used by **word embeddings** for information retrieval.

Def. 157 (Grammar). A tuple $\langle N, \Sigma, P, S \rangle$ where N is a finite set of non-terminal symbols, Σ is a finite set of terminal symbols, P is a set of production rules, S is a distinguished “start symbol”

These can be categorized as “context-sensitive”, “context-free”, “regular”, etc.

Def. 158 ((Formal) Language). A set of sentences that can be generated by a **grammar**, written $L(G)$.

Def. 159 (Ambiguity). Real languages pose issues for **natural language processing**, such as *ambiguity*, *anaphora*, *indexicality*, *vagueness*, *discourse structure*, *metonymy*, *metaphor* and *noncompositionality*.