

PROPOSITIONAL LOGIC

We say m is a model of a sentence α if α is true in m : $m \models \alpha$.

$$M(\alpha) = \{m \mid m \models \alpha\}$$

$$KB \models \alpha \Leftrightarrow M(KB) \subseteq M(\alpha)$$

$$\alpha \equiv \beta \text{ if } M(\alpha) = M(\beta)$$

α is valid $\Leftrightarrow \alpha$ is true in all models (tautology).

α is satisfiable $\Leftrightarrow M(\alpha) \neq \emptyset$

Resolution rule.

$$\frac{l_1 \vee \dots \vee l_k \quad m_1 \vee \dots \vee m_n}{l_1 \vee \dots \vee l_{i-1} \vee l_{i+1} \vee \dots \vee l_k \vee m_1 \vee \dots \vee m_{j-1} \vee m_{j+1} \vee \dots \vee m_n}$$

where $l_i = \neg m_j$ or $\neg l_i = m_j$ on a syntactic level.

A sentence is in CNF iff it is a conjunction of disjunctions of literals.

Approach 1 (Writing propositional logic in CNF)

1. Replace $S_1 \Leftrightarrow S_2$ by $(S_1 \Rightarrow S_2) \wedge (S_2 \Rightarrow S_1)$.
2. Replace $S_1 \Rightarrow S_2$ by $\neg S_1 \vee S_2$.
3. Push \neg inwards until it hits a literal. Use:

$$\neg(S_1 \wedge S_2) = \neg S_1 \vee \neg S_2$$

$$\neg(S_1 \vee S_2) = \neg S_1 \wedge \neg S_2$$

$$\neg\neg S_1 = S_1$$

4. Distribute \vee over \wedge whenever possible:

$$(S_1 \wedge S_2) \vee S_3 = (S_1 \vee S_3) \wedge (S_2 \vee S_3)$$

Approach 2 (Proof by Resolution for propositional logic)

1. Add α to KB.
2. Write $KB \wedge \neg\alpha$ in CNF.
3. Show that $KB \wedge \alpha \vdash \perp$ by applying the resolution rule until no new clause can be added anymore, or **false** is derived.

PREDICATE LOGIC

Substitution: $\{\text{Old}_1/\text{New}_1, \dots, \text{Old}_n/\text{New}_n\}$

A unifier is a substitution which makes two literals equal. We unify as follows:

1. If T_1 and T_2 are constants, then T_1 and T_2 unify if they are the same constant.
2. If T_1 is a variable and T_2 is any type of term, then T_1 and T_2 unify, and T_1 is instantiated to T_2 (and vice versa).
3. If T_1 and T_2 are complex terms then they unify if:
 - (a) They have the same functor and arity, and,
 - (b) All their corresponding arguments unify, and,
 - (c) The variable instantiations are compatible.

The Most General Unifier (MGU) is a unifier with the least amount of variables replaced by concrete values.

$$\forall_x \neg P \equiv \neg \exists_x P$$

$$\neg \forall_x P \equiv \exists_x \neg P$$

$$\forall_x P \equiv \neg \exists_x \neg P$$

$$\exists_x P \equiv \neg \forall_x \neg P$$

Approach 3 (Removing quantifiers)

First remove existential quantifiers by replacing quantified variables by Skolemisation constants or functions.

$$\forall_{a,b,\dots} \exists_x P(x,y) \equiv \forall_{a,b,\dots} P(S(a,b,\dots),y)$$

Then, drop all universal quantifiers.

Approach 4 (Proof by Resolution for predicate logic)

Same as for propositional logic.

SEARCH

Tree search and graph search can both be used on trees and graphs. The difference is that graph search keeps track of completely visited nodes. Tree search is favorable in acyclic action paths.

A* search is a *best-first* search which uses an evaluation function $f(n) = g(n) + h(n)$ to find the next candidate node n , with $g(n)$ the cost to reach the node n from the start state (known) and $h(n)$ the estimated cost to get from n to the goal.

The tree version of A* is optimal if $h(n)$ is admissible: $h(n)$ never overestimates the cost.

The graph version of A* is optimal if $h(n)$ is consistent: the triangle inequality holds: $h(n) \leq \text{cost}(n, n') + h(n')$.

PROBABILISTIC REASONING

$$P(A \wedge B) = P(A \mid B)P(B)$$

$$P(A \vee B) = P(A) + P(B) - P(A, B)$$

$$P(X_1, \dots, X_n) = \prod_{i=1}^n P(X_i \mid X_1, \dots, X_{i-1})$$

$$\mathbf{P}(X) = \langle P(X = x_1), \dots, P(X = x_n) \rangle$$

$$\mathbf{P}(Y) = \sum_z \mathbf{P}(Y, Z = z)$$

$$\mathbf{P}(Y) = \sum_z \mathbf{P}(Y \mid z)P(z)$$

Independence (equivalent statements):

$$P(a \mid b) = P(a)$$

$$\mathbf{P}(X \mid Y) = \mathbf{P}(X)$$

$$P(b \mid a) = P(b)$$

$$\mathbf{P}(Y \mid X) = \mathbf{P}(Y)$$

$$P(a \wedge b) = P(a)P(b)$$

$$\mathbf{P}(X, Y) = \mathbf{P}(X)\mathbf{P}(Y)$$

Bayes rule: $P(b \mid a) = \frac{P(a|b)P(b)}{P(a)}$

$$P(A|B) = \left[\begin{array}{l} P(a, a), P(\neg a, b) / \alpha_b \\ P(a, \neg b), P(\neg a, \neg b) / \alpha_{\neg b} \end{array} \right]$$

$$\alpha_b = P(a, a) + P(\neg a, b)$$

$$\alpha_{\neg b} = P(a, \neg b) + P(\neg a, \neg b)$$

BAYESIAN NETWORKS

A Bayesian network is a directed acyclic graph $G(V, E)$ with random variables as nodes and direct influence between random variables as edges. We define $\text{Parents}(X_i) = \{p \mid (p, X_i) \in E\}$ the parents of node X_i .

$$P(X_1, X_2, \dots, X_n) = \prod_{i=1}^n P(X_i \mid \text{Parents}(X_i))$$

For the set of nodes A, B and C , $A \perp B \mid C$ if all paths from A to B are blocked.

$$A \overset{\pm}{\rightarrow} W \overset{\pm}{\rightarrow} B \Rightarrow \begin{cases} A \text{ does not affect } B & \text{if } W \text{ observed} \\ A \text{ affects } B & \text{otherwise} \end{cases}$$

$$A \overset{\pm}{\leftarrow} W \overset{\pm}{\rightarrow} B \Rightarrow \begin{cases} A \text{ does not affect } B & \text{if } W \text{ observed} \\ A \text{ affects } B & \text{otherwise} \end{cases}$$

$$A \overset{\pm}{\rightarrow} W \overset{\pm}{\leftarrow} B \Rightarrow \begin{cases} A \text{ does not affect } B & \text{if } W \text{ not observed} \\ A \text{ affects } B & \text{otherwise} \end{cases}$$

MACHINE LEARNING

- **Unsupervised learning:** we only supply data and give no labels.
- **Semi-supervised learning:** the data is only partly labeled.
- **Supervised learning:** the data is fully labeled.

$$D = \left\{ \left\{ \xi^\mu, S^\mu \right\} \right\}_\mu = 1^P$$

$$D = \cup_{i=1}^n D^{(i)} \quad \text{all data}$$

$$D_{\text{train}}^{(i)} = D \setminus D^{(i)} \quad \text{training data } (i)$$

$$D_{\text{test}}^{(i)} = D^{(i)} \quad \text{test data } (i)$$

We wish to aim two competing goals:

- **Low bias:** the system only slightly deviates from the "true solution".
- **Low variance:** the system weakly depends on the training set.

		Actual	
		P	N
Pred.	P	True Positive (TP)	False Positive (FP)
	N	False Negative (FN)	True Negative (TN)

REINFORCEMENT LEARNING

$$U_h [s_0, s_1, s_2, \dots] = R(s_0), \gamma R(s_1), \gamma^2 R(s_2), \dots$$

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

$$= \text{Sensitivity} = \text{True Positive Rate}$$

$$\text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}}$$

$$= \text{True Negative Rate}$$

$$\text{F-score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

$$E = \frac{1}{2} \sum_{n=1}^N (y(X_n) - t_n)^2$$

Two internal measures for clustering: cohesion, separation.

$$\text{WSS} = \sum_i \sum_{x \in C_i} (x - m_i)^2$$

$$\text{BSS} = \sum_i |C_i| (m - m_i)^2$$

DECISION TREES

$$\text{Entropy}(S) = \sum_{i=1}^c -P_i \log_2(P_i)$$

$$\text{Gain}(S, A) = \text{Entropy}(S) - \sum_{v \in A} \frac{|S_v|}{|S|} \text{Entropy}(S_v)$$

Approach 5 (Iterative Dichotomiser)

Input. A data set with attributes.

Output. A decision tree.

1. Find the attribute that provides the highest information gain within this data set.
2. Create a node using this attribute.
3. Sort out the data set to the corresponding split.
4. For each split data set: if all data points correspond to the same class, create a leaf node.
5. Otherwise, apply this algorithm to the split data set and insert the resulting tree here.
 - If all the records in the current data set have the same class, then stop recursing.
 - If all records have the same set of input attributes, then stop recursing.

Approach 6 (Pruning decision trees)

1. Consider a node with p positive and n negative examples. Let p_1, n_1 and p_2, n_2 be the positive and negative examples for the two alternatives of the node.
2. Define $\hat{p}_1 = p \cdot \frac{p_1 + n_1}{p + n}$ and \hat{n}_1, \hat{p}_2 , and \hat{n}_2 .
3. Calculate $\Delta = \frac{(p_1 - \hat{p}_1)^2}{\hat{p}_1} + \frac{(n_1 - \hat{n}_1)^2}{\hat{n}_1} + \frac{(p_2 - \hat{p}_2)^2}{\hat{p}_2} + \frac{(n_2 - \hat{n}_2)^2}{\hat{n}_2}$
 $\Delta \sim \chi_{d-1}^2$, with d the number of splits. Small values for Δ imply no rejection of the null hypothesis, and hence justify the pruning of the tree.

REINFORCEMENT LEARNING

$$U_h [s_0, s_1, s_2, \dots] = R(s_0), \gamma R(s_1), \gamma^2 R(s_2), \dots$$

$$U^\pi(s) = \mathbb{E} \sum_{t=0}^{\infty} \gamma^t R(S_t)$$

$$\pi^*(s) = \operatorname{argmax}_{a \in A(s)} \sum_{s'} P(s' | s, a) U(s')$$

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

Value Iteration

$$U_{i+1}(s) \leftarrow R(s) + \gamma \max_{a \in A(s)} \sum_{s'} P(s' | s, a) U_i(s')$$

Policy Iteration

Start with random $\pi = \pi^l$. Until there is no change anymore, for each $s \in S$, consider a such that

$$\left[R(s) + \gamma \sum_{s' \in S} P(s' | s, a) U(\pi, s') \right] > U(\pi, s)$$

Direct Utility Estimation

$$U^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(S_t, \pi(S_t), S_{T+1}) \right]$$

Adaptive Dynamic Programming

$$U^\pi(s) = R(s) + \gamma \sum_{s'} P(s' | s, \pi(s)) U^\pi(s')$$

$$P(s' | s, a) = \frac{\text{times you visited } s' \text{ from } s \text{ with action } a}{\text{times you performed action } a \text{ at state } s}$$

Active learning

Estimate $P(s' | s, a)$ for more than one action. The action the agent should choose is

$$\pi(s) = \operatorname{argmax}_{a \in A(s)} \sum_{s' \in S} P(s' | s, a) U(s')$$

To balance exploitation and exploration, use an exploration function $f(u, n)$:

$$U^+(s) \leftarrow \max_a f \left(\sum_{s'} P(s' | s, a) [R(s, a, s') + \gamma U^+(s')], N(s, a) \right)$$

$$f(u, n) = \begin{cases} R^+ & \text{if } n < N_e \\ u & \text{otherwise} \end{cases}$$

with $N(s, a)$ is the number of times the action a has been executed in state s .

MACHINE LEARNING FOR SECURITY

$$\bar{\phi}_i(d) = \frac{\phi_i(d) - \min_i}{\max_i - \min_i} \in [0, 1]$$

$$\mu_+ = \operatorname{argmin}_{\mu \in \mathbb{R}^N} \max_{1 \leq i \leq n} \|x_i - \mu\|^2$$

$$\min_{R, \mu, \xi} R^2 + \frac{1}{vn} \sum_{i=1}^n \xi_i$$

subject to $\|x_i - \mu\|^2 \leq R^2$ for all $i = 1, \dots, n$.