

(c) Filip Železný, all rights reserved. Any use of this material only with prior approval by the author.

Symbolic Machine Learning

Filip Železný

Contents

1	A General Framework	4
1.1	Percepts and Actions	4
1.2	Nonsequential Cases	6
1.3	Batch Learning	6
1.4	Rewards and Goals	8
1.5	Environment States	9
1.6	Agent States	12
1.7	Nonsequential and Batch Cases with States	13
1.8	Prior Knowledge	15
1.9	Hypothesis Representations	16
1.10	Learning Scenarios	16
2	On-line Concept Learning	17
2.1	Generalizing Agent	22
2.2	The Subsumption Relation	25
2.3	Separating agent	26
2.4	Hypothesis and Concept Classes	28

2.5	Version Space Agent	29
2.6	The Mistake Bound Learning Model	30
3	Batch Concept Learning	32
3.1	Batch Learning with the Generalizing Agent	33
3.2	Batch Learning with General On-line Agents	34
3.3	Consistent Agent	35
3.4	The PAC Learning Model	36
4	Learning First-Order Logic Concepts	38
4.1	Generalizing Agent	41
4.2	Generalization of Clauses	45

1 A General Framework

1.1 Percepts and Actions




Figure 1: The basic situation under study.

- Discrete *time*

$$k = 1, 2, \dots$$

- *Percepts*

$$\forall k : x_k \in X$$

- *Actions*

$$\forall k : y_k \in Y$$

X and Y are finite.

A *history* is a sequence of alternating percepts and actions, i.e.,

$$x_1, y_1, x_2, y_2, \dots, x_k, y_k$$

and is denoted as $xy_{\leq k}$. Similarly, $xy_{< k} = x_1, y_1, x_2, y_2, \dots, x_{k-1}, y_{k-1}$. There is a probability distribution μ on histories

$$\mu(xy_{\leq k}) = \mu(x_1)\mu(y_1|x_1)\mu(x_2|x_1, y_1) \dots \mu(x_k|xy_{< k})\mu(y_k|x_k, xy_{< k}) \quad (1)$$

After the initial ‘kick-off’ x_1 from the environment distributed according to $\mu(x_1)$, any percept x_k generated by the environment at time k depends on the entire preceding history $xy_{< k}$ according to

$$\mu(x_k|xy_{< k}) \quad (2)$$

Actions y_k are determined by agent’s decision *policy* which also depends on the history as well as the current percept and are distributed according to

$\mu(y_k|x_k, xy_{<k})$. We will assume that the policy is *deterministic*. Thus we identify the policy with function $\pi : (X \times Y)^* \times X \rightarrow Y$, so

$$y_k = \pi(xy_{<k}, x_k) \quad (3)$$

This means that $\mu(y_k|xy_{<k}, x_k) = 1$ for $y_k = \pi(xy_{<k}, x_k)$ and 0 otherwise.

The following diagram illustrates the influences between the introduced variables.




Figure 2: Influence diagram for actions y_k and percepts x_k for $1 \leq k \leq 3$ with full lines indicating deterministic influences (via π) and dashed lines showing probabilistic influences (via μ).

While we have yet to define what goals the agent should achieve through interaction with the environment, obviously some histories will be “better” than others in terms of the goal achievement. To maximize the probability (1) of good histories, the agent cannot influence the conditional probability (2), which is inherent to the environment, but it can follow a good policy (3). However, the effect of actions proposed by the policy depends on (2) which is generally not known to the agent. So the agent needs to recognize the environment by experimenting with it. This is formally reflected by (3) where action y_k depends not only on the current percept x_k but also on the history $xy_{<k}$. So the agent will generally make different decisions $y_k \neq y_{k'}$ for $k > k'$ even if $x_k = x_{k'}$ because the experience $xy_{<k}$ at time k is larger than experience $xy_{<k'}$ at time k' . This is our first reflection of *learning*.

How does the agent know how well it is doing? This information comes from the environment through a specially distinguished part of the percepts, called *rewards*. The remaining part of each percept contains *observations*. Formally, $X = O \times R$, $o_k \in O$, $r_k \in R \subset \mathbf{R}$, so

$$x_k = (o_k, r_k) \quad (4)$$

Since X is assumed finite, it follows that rewards have their finite minimum and maximum.

The probability of x_k in (2) can be written in terms of the marginals μ_O and μ_R

$$\begin{aligned}\mu(x_k|xy_{<k}) &= \mu(o_k, r_k|xy_{<k}) = \\ \mu_O(o_k|r_k, xy_{<k})\mu_R(r_k|xy_{<k}) &= \mu_R(r_k|o_k, xy_{<k})\mu_O(o_k|xy_{<k})\end{aligned}$$

which also makes it clear that o_k and r_k are in general not assumed mutually independent, even if conditioned on $xy_{<k}$.

1.2 Nonsequential Cases

Scenarios where current percepts depend on the history of previous percepts and actions are called *sequential*. The framework described so far is maximally general in that dependence is assumed on the entire history from $k = 1$ on. On the other extreme are *nonsequential* scenarios. Here, observations are independent of the history as well as the current reward, i.e.

$$\mu_O(o_k|r_k, xy_{<k}) = \mu_O(o_k) \tag{5}$$

and thus o_1, o_2, \dots are mutually independent random variables sampled from the same distribution μ_O (they are “i.i.d.”).

Rewards in the nonsequential case are assumed to depend only on the immediately preceding observation and the action taken on it, i.e.

$$\mu_R(r_k|o_k, xy_{<k}) = \mu_R(r_k|o_{k-1}, y_{k-1}) \tag{6}$$

however, since y_{k-1} is functionally determined by the history $xy_{<k-1}$ and percept $x_{k-1} = (o_{k-1}, r_{k-1})$ through (3), we may rewrite (6) as

$$\mu_R(r_k|o_{k-1}, r_{k-1}, xy_{<k-1}) \tag{7}$$

which makes it clear that reward r_k depends on previous rewards, and thus rewards r_1, r_2, \dots are not i.i.d.. This is natural since if they were, it would mean the agent never improves its performance.

1.3 Batch Learning

We will also consider a specific yet important nonsequential case called *batch learning* consisting of two phases switching right after time K




Figure 3: Influence diagram for actions y_k , observations o_k , and rewards r_k for $1 \leq k \leq 3$ with full lines indicating deterministic influences (via π) and dashed lines showing probabilistic influences (via μ) in the nonsequential case.

- the *learning (training, exploration) phase* at $k = 1, 2, \dots, K$
- the *action (testing, exploitation) phase* taking place in $k = K+1, K+2, \dots$

In the action phase, the agent no longer changes its decision making policy, so

$$\text{if } k, k' > K \text{ and } x_k = x_{k'} \text{ then } y_k = y_{k'} \quad (8)$$

and ignores rewards. So the action proposed by the policy depends only on the current observation and the history only up to time K . So for $k > K$, (3) changes here into

$$y_k = \pi(xy_{\leq K}, o_k) \quad (9)$$

and (6, 7) change into

$$\mu_R(r_k | o_{k-1}, y_{k-1}) = \mu_R(r_k | o_{k-1}, xy_{\leq K}) \quad (10)$$

because due to (9), y_{k-1} is determined by o_{k-1} and $xy_{\leq K}$. The observation o_{k-1} does not depend on rewards due to (5). So reward r_k does not depend on previous rewards $r_{k'}, k > k' > K$. Another way to say this is that rewards in the action phase are conditionally independent of each other, given the learning phase history:

$$\mu_R(r_k, r_{k'} | xy_{<K}) = \mu_R(r_k | xy_{<K}) \mu_R(r_{k'} | xy_{<K}) \quad (11)$$

The following figure illustrates the batch-learning situation.




Figure 4: Influence diagram for actions y_k , observations o_k , and rewards r_k in the action phase ($k > K$) of batch learning with full lines indicating deterministic influences (via π) and dashed lines showing probabilistic influences (via μ). The top row indicates the influence of the learning phase on the agent’s decisions in the action phase.

We have observed that rewards $r_k, r_{k'}$ (for any $k, k' > K$) are mutually independent given the learning phase $xy_{<K}$. Note that they are also sampled from the same distribution. This may seem to contradict (10) which stipulates that r_k depends on the observation o_{k-1} whereas $r_{k'}$ depends on $o_{k'-1}$. However, by (5), o_{k-1} and $o_{k'-1}$ are sampled from the same distribution μ_O . So we can express the distribution of r_k ($\forall k > K$) without conditioning on the observations by marginalizing them away from the equation as follows

$$\mu_R(r_k | xy_{\leq K}) = \sum_{o_{k-1} \in \mathcal{O}} \mu_O(o_{k-1}) \mu_R(r_k | o_{k-1}, xy_{\leq K}) \quad (12)$$

So rewards in the action phase indeed are i.i.d. according to the above distribution conditioned only on the history of the learning phase.

1.4 Rewards and Goals

It has been obvious that the agent’s goal is to maximize rewards. Here we formalize this goal. Since rewards come at each point of the history, we want the agent to maximize their sum up to a finite time *horizon* $k' \in N$

$$r_1 + r_2 + \dots + r_{k'}$$

or, more generally, maximize the *discounted* sum

$$\sum_{k=1}^{\infty} r_k \delta_k$$

where $\forall k : \delta_k \geq 0$ and $\sum_{i=1}^{\infty} \delta_i < \infty$. The latter condition guarantees that the sum above converges, which is because the r_k 's are bounded by a constant (c.f. Section 1.1).

But since rewards are probabilistic, the agent should choose a sequence $y_{\leq k'}$ of actions leading to a high *expected* cumulative reward

$$\sum_{r_{\leq k'}} \mu_R(r_{\leq k'} | y_{\leq k'}) (r_1 + r_2 + \dots + r_{k'})$$

or, in the discounted case

$$\lim_{k' \rightarrow \infty} \sum_{r_{\leq k'}} \mu_R(r_{\leq k'} | y_{\leq k'}) \sum_{k=1}^{k'} r_k \delta_k$$

where the first sum in both cases goes over all possible reward sequences $r_{\leq k'}$ (since R and k' are finite, there is a finite number of them).

However, for the specific case of *batch learning*, we establish a more appropriate learning goal. First, we do not care about maximizing rewards in the *learning phase* as the purpose of this phase is to probe the environment even at the price of possibly poor rewards. Second, in the *action phase* after time K , the rewards r_k , $k > K$ are sampled independently from the same distribution (12) so we can simply maximize their expectation with respect to this distribution

$$\sum_{r_k \in R} \mu_R(r_k | xy_{\leq K}) r_k \tag{13}$$

It is again obvious from the formula that the expected reward only depends on the learning phase history $xy_{\leq K}$, after which the agent no longer changes its action policy. Note also that the batch learning scenario allowed us to define an objective (13) without the need to choose the parameters k' or δ_k ($k = 1, 2, \dots$) needed in the sequential scenario.

1.5 Environment States

With the exception of the non-sequential scenario, our framework has been very general in that percepts x_k generally depend on entire histories $xy_{<k}$. In the real world, many histories may be equivalent, i.e. leading to the same probabilities

of x_k conditioned on action y_{k-1} . This can be formalized through the notion of *environment state*. Intuitively, the state acts as the environment’s ‘memory’ carrying all the information from the history, which is important for generating percepts. To formalize this, we will assume there is a set E of all possible states, and instead of 2, we will prescribe that percepts at time k only depend on the state $e_k \in E$ at that time, i.e. they are generated according to

$$\mu(x_k|e_k) \tag{14}$$

But how are the states determined? We will first explore a principle, which—in a sense—will turn out to be maximally general. In particular, assume that the initial state e_1 is fixed to an ‘empty’ (or ‘dummy’) value $e_1 = s^*$ such that

$$\mu(x_1|s^*) = \mu(x_1) \tag{15}$$

so the first percept is generated just as in (1). Afterwards, any state e_k ($k > 1$) is established probabilistically by the preceding state, the last percept, and the last agent’s action through the following state *update* distribution

$$\mathcal{E}(e_k|e_{k-1}, x_{k-1}, y_{k-1}) \tag{16}$$

We said earlier that this principle was ‘maximally general.’ To say this more precisely, for any environment generating percepts by (2), we can set up E and \mathcal{E} such that

$$\mu(x_k|e_k) = \mu(x_k|xy_{<k}) \tag{17}$$

Indeed, if we allow E to be infinite, then there could simply exist a distinct state for each possible history (there is an infinite number of possible histories for unbounded k). Then we can instantiate the distribution (16) to the functional dependence

$$e_k = e_{k-1} \parallel (x_{k-1}, y_{k-1}) \tag{18}$$

where \parallel denotes concatenation. As a result, e_k will simply collect the entire history and in (17), e_k would be just a different name for $xy_{<k}$.

However, we will make an important assumption, which will significantly simplify the framework, that the number of possible states is bounded by a finite constant $E_{\max} \in \mathbb{N}$ which does not depend on k

$$|E| < E_{\max} \tag{19}$$

In practical tasks, there will be far fewer states than possible histories.

Moreover, we can afford an additional simplifying assumption which will further lessen the generality of the framework, while keeping it able to encompass the learning scenarios we are going to elaborate. In particular, we will assume quite naturally that the influence between environment states and the emitted

percepts are single-directional in the sense that the percepts depend on states by (14) but not vice versa, so we remove x_{k-1} from (16)

$$\mathcal{E}(e_k|e_{k-1}, x_{k-1}, y_{k-1}) = \mathcal{E}(e_k|e_{k-1}, y_{k-1}) \quad (20)$$

As we have discussed already, the finiteness of E means that the environment has a ‘finite memory.’ Thus, from a current state e , one cannot in general identify the entire preceding history of states and actions. However, the framework as defined can still model environments that remember a *bounded number* of previous states and actions. We will demonstrate this through an example with a one-step memory.

Consider an environment with states E and update distribution \mathcal{E} , and assume that the current state e_k does not allow to infer the previous state e_{k-1} or action y_{k-1} . Then we can always define an extended (yet still finite) set of states as

$$E_{\text{ext}} = E \times E \times Y \quad (21)$$

(with the latter two factors acting as memory) and an extended update distribution

$$\mathcal{E}_{\text{ext}} \{(e_k, \text{olde}_k, \text{old}y_k) | (e_{k-1}, \text{olde}_{k-1}, \text{old}y_{k-1}), y_{k-1}\} \quad (22)$$

such that

- e_k is distributed according to $\mathcal{E}(e_k|e_{k-1}, y_{k-1})$
- $\text{olde}_k = e_{k-1}$ and $\text{old}y_k = \text{old}y_{k-1}$, both with probability 1.

So the three components of the extended state correspond, respectively, to the original state at current time k , the same at the previous time $k - 1$ and the agent’s actions at time $k - 1$.

We will often model environments with a natural (interpretable) set of states E , producing percepts that depend not only on the current state but also on the state and agent’s action one-step back in history. We have just seen that such a situation can still be modeled with the simple assumption (14) by extending E towards a state set with a memory as in the above example. However, this leads to some cumbersome notation as in (22). We will avoid these complications by keeping the original state set E rather than extending it with memory for e_{k-1} and y_{k-1} . Instead, we will add the latter two explicitly to the conditional part of (14), thus obtaining

$$\mu(x_k|e_k, e_{k-1}, y_{k-1}) \quad (23)$$

Regarding the two components of the percepts x_k , the observations will depend only on the current state (and not on the previous one) and the last agent’s action

$$\mu(o_k|e_k, e_{k-1}, y_{k-1}) = \mu_o(o_k|e_k, y_{k-1}) \quad (24)$$

and the reward will depend on the previous state (and not on the current one) and the action taken immediately on it

$$\mu(r_k|e_k, e_{k-1}, y_{k-1}) = \mu_r(r_k|e_{k-1}, y_{k-1}) \quad (25)$$

The formulation (23)–(25) is convenient for interpretability and to avoid the complex notation such as in (22). One should however keep in mind that with a suitable definition of the state variable, it is possible to avoid the variables e_{k-1}, y_{k-1} in the conditional part and adhere to the simple prescription (14).

1.6 Agent States

A reasoning similar to the previous section applies to the agent, whose actions generally depend on the entire history as in (3). Again, many histories can lead to the same mapping from percepts to actions, for example because the agent has built the same hypothesis about the environment throughout the different histories. So analogically to the environmental states, we introduce the notion of agent’s *state* $a \in A$ and postulate that

$$|A| < A_{\max} \quad (26)$$

for some constant $A_{\max} \in N$.

We adopt a counterpart of (14), meaning that the action will be determined by the agent’s state, again through a functional prescription

$$y_k = \pi(a_k) \quad (27)$$

The current action thus does not depend explicitly on the current percept x_k . This is because the latter can be simply stored as a part of the state equipped with memory as can be shown through a reasoning very similar to that in the previous section. However, the dependence on the current percept, and specifically on its observation part, will be so typical that we will make it explicit. Again, this will save us notational complications entailed by the need to memorize percepts within states. We will thus use the formula

$$y_k = \pi(a_k, o_k) \quad (28)$$

Regarding the update rule, analogically to (20) we will assume that the state is updated (deterministically) given the previous one and the current percept.

$$a_k = \mathcal{A}(a_{k-1}, x_k) \quad (29)$$

This seems not fully analogical to (20) as the current percept, rather than the previous one is taken into account. This is due to our setting of the agent-environment communication, in which y_{k-1} is the last action taken before the

environment updates its state, whereas x_k is the last percept received by the agent for its state update.

The formalization using environment states and agent hypotheses results in the agent and environment structures depicted in Fig. 5. The diagram of variable influences is shown in Fig. 6.




Figure 5: The state-based scheme of agent-environment interaction according to (14), (20) and (27), and (29). Full and dashed lines denote functional and probabilistic influences, respectively. The $k - 1$ nodes denote a one-step time lag. Note that we have introduced further dependencies through (23) and (28), which are not shown in the picture. These dependencies can be avoided through a reformulation of the state variable and the state update function.

1.7 Nonsequential and Batch Cases with States

Just like in the framework using entire histories, also with the formulation based on states and hypotheses the situation simplifies a lot in the *nonsequential case*. Here, the environment has no memory at all so the conditioning factors in (20) and states are updated by i.i.d. sampling from the marginal distribution

$$\mathcal{E}(e_k) \tag{30}$$

Furthermore, observations o_k no longer depend on agent's last action as in (24) so given e_k , an observation is sampled from

$$\mu_O(o_k|e_k) \tag{31}$$

Since e_k 's are i.i.d., the o_k 's are also i.i.d. from

$$\mu_O(o_k) = \sum_{e_k \in E} \mu_O(o_k|e_k)\mathcal{E}(e_k) \tag{32}$$




Figure 6: Influence diagram for states a , actions y_k and percepts x_k for $1 \leq k \leq 3$ with full lines indicating deterministic influences (via π and \mathcal{A}) and dashed lines showing probabilistic influences (via μ and \mathcal{E}). See caption to Fig. 5 for further relevant remarks.

Rewards, given by (25), are however still generally non-i.i.d. as they depend on the agent's actions, which in turn depend on the evolving agent's hypothesis. Fig. 7 shows the complete set of influences in the nonsequential case.

A further simplification comes in the special *batch-learning* scenario of the non-sequential case. While in the learning phase of the latter, the agent uses the update rule (29), in the action phase it no longer updates its state, so

$$a_k = a_K, \forall k \geq K \quad (33)$$

This is illustrated in Fig. 8. Special attention is needed regarding the variables at time K . Reward r_K (part of percept x_K) is the last training reward, according to which the last update is conducted towards the final h_K . Observation o_K (another part of percept x_K) is, however, the first *testing* observation.

Since y_{k-1} is determined by o_{k-1} and a_{k-1} through (28), we can rewrite (25) as

$$\mu_r(r_k | e_{k-1}, o_{k-1}, a_{k-1}) \quad (34)$$

Because environment states e_k and observations o_k in the non-sequential setting are i.i.d. (sampled from the respective distributions (30) and (32) independent of k), we can marginalize the above as

$$\mu_r(r_k | a_{k-1}) = \sum_{e_{k-1} \in E} \sum_{o_{k-1} \in O} \mu_r(r_k | a_K, e_{k-1}, o_{k-1}) \mu_O(o_{k-1} | e_{k-1}) \mathcal{E}(e_{k-1}) \quad (35)$$




Figure 7: Influence diagram for hypothesis h_k , actions y_k , observations o_k , and rewards r_k for $1 \leq k \leq 3$ with full lines corresponding to deterministic influences (via π and \mathcal{H}) and dashed lines showing probabilistic influences (via μ and \mathcal{E}) in the nonsequential case.

In the testing phase ($k > K$) where a_k is fixed to $a_k = a_K$, rewards r_k are thus i.i.d. according to the distribution $\mu_r(r_k|a_K)$ depending only on the final state a_K of training. This is analogical to the state-free formulation (12). Similarly to (13), an agent operating in the batch-learning scenario with states will be assessed by the expected reward in the testing phase

$$\sum_{r_k \in R} \mu_R(r_k|a_K)r_k \quad (36)$$

so in the training phase, it should reach a state a_K maximizing this quantity.

1.8 Prior Knowledge

- *Implicit*: the setting of A (“hard bias”) and \mathcal{A} (“soft bias”)
- *Explicit*: the setting of a_1 (“background knowledge”)




Figure 8: Influence diagram for actions y_k , observations o_k , states e_k , and rewards r_k in the action phase ($k > K$) of batch learning with full lines indicating deterministic influences (via π) and dashed lines showing probabilistic influences (via μ). The top row indicates the influence of the agent’s last state in the training phase on the action phase. The dependence of r_{K+1} on e_K and y_K is not shown.

1.9 Hypothesis Representations

See Fig. 9.

1.10 Learning Scenarios

1. on-line concept learning
2. batch concept learning
3. query-based and active learning (not covered here)
4. reinforcement learning
5. universal learning (not covered here)

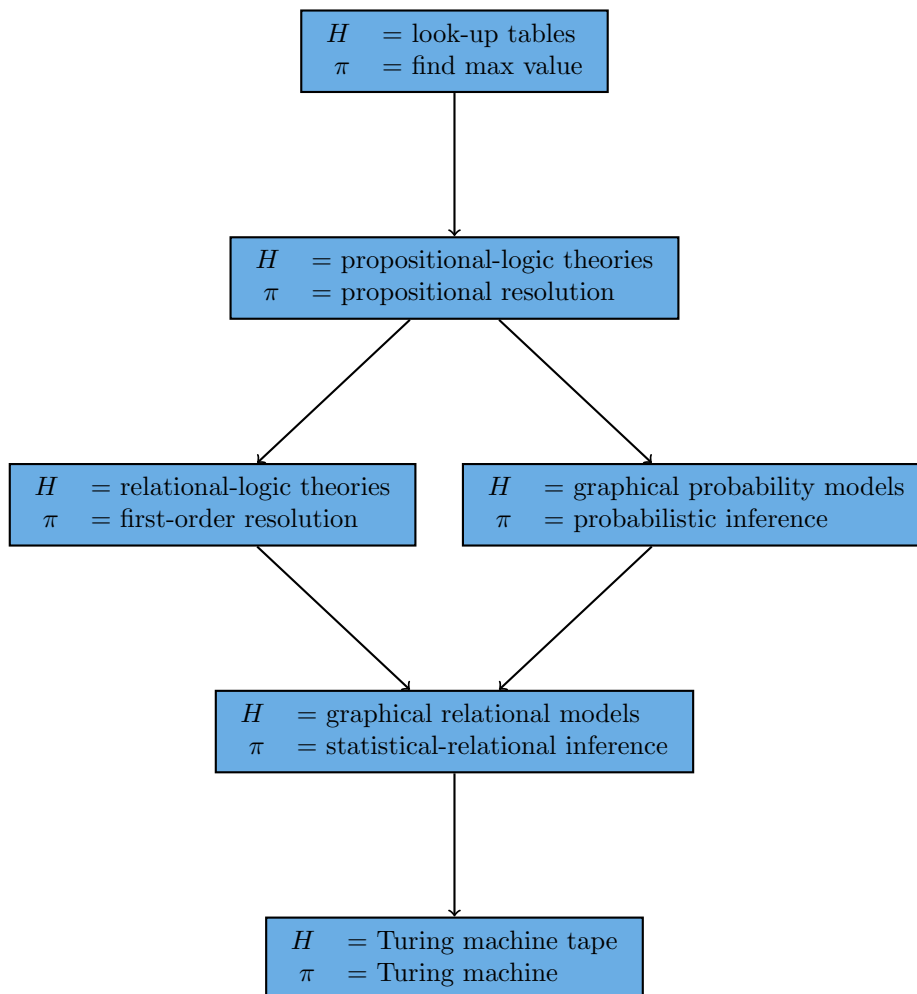


Figure 9: Hypothesis representations and their corresponding policy classes (interpreters) considered in this course. Arrow directions indicate increasing expressiveness.

2 On-line Concept Learning

We first motivate the on-line concept learning scenario with an example, in which the agent is an artificial scientist. The agent conducts repeated experiments with a living cell, which represents the environment. In each experiment, it observes two proteins of interest in the cell. More specifically, the agent detects whether the proteins are present in the cell at all, and it also determines

whether they are in an active state (a special spatial conformation of a protein). The agent suspects that these proteins (both or only one of them) initiate apoptosis (cellular suicide). After each observation of the proteins, it tries to predict whether the cell will die or not. If the prediction is incorrect, the agent receives a negative reward. This can be for example a cut-down on the agent’s salary by the boss of the lab who is not happy with wrong biological predictions, in which case the boss would be a part of the environment. However, we will simply model such a reward with the number -1 for wrong predictions and with 0 for correct predictions.

experiment number	apoptosis initiated	protein 1 present	protein 1 active	protein 2 present	protein 2 active	prediction	reward
k	e_k	o_k^1	o_k^2	o_k^3	o_k^4	y_k	r_k
1	0	0	0	0	0	0	0
2	0	0	0	1	0	1	0
3	0	1	0	0	0	1	-1
4	1	1	1	0	0	0	-1
5	0	1	0	1	1	0	-1
6	1	1	1	1	0	1	0
7	1	1	1	0	0	1	0
8	0	1	0	1	1	0	0
(etc.)							

Table 1: An on-line concept learning experiment.

Table 2 illustrates a history of such agent-environment interaction, in which the agent eventually learns that apoptosis is induced if and only if protein 1 is present and it is in the active form. From time $k = 5$ on, the agent makes correct predictions and is no longer punished with negative reward.

In the sequel, we will see how to model the illustrated scenario in our frameworks and we will see examples of agents able to learn as the agent-scientist has in the story above.

As Table 2 already indicated, the unknown variable guessed by the agent corresponds to the unknown state of the environment e . The variable is binary so we set

$$E = \{0, 1\} \tag{37}$$

We will accommodate (37) as a general assumption in the forthcoming text, unless we specify otherwise. This is because the binary setting is the simplest non-trivial one, to which richer state sets can usually be reduced.

The central assumption of concept learning is that the state e_k is fully *determined* by the observation $o_k = (o_k^1, o_k^2, \dots)$. In other words, an observation o generated by state 0 cannot be generated by state 1 and vice versa. In terms of

the probability notation, this means that

$$\mu_O(o|0)\mu_O(o|1) = 0 \quad (38)$$

i.e., at most one of the probabilities is non-zero. Note that although the state is determined by the observation, the agent does not know *how* until it *learns* such knowledge.

The set of observations which can be generated from state 1 is called the *concept* C (of the environment)

$$C = \{o \in O \mid \mu_O(o|1) > 0\} \quad (39)$$

and the observations coming from this concept are *positive* observations (or, ‘examples’) of it, while the remaining observations, i.e. those in $O \setminus C$ are termed *negative*. Since the environment states are partitioned into two *classes* (positive and negative), the agent’s guessing is an act of *classification* and the concept learning task is a special case of what is commonly termed *classification learning*.

As we have indicated already, the agent guesses the current state through its decision variable $y \in Y$. It is thus natural to set

$$Y = E = \{0, 1\} \quad (40)$$

Whenever the agent makes an incorrect guess $y_k \neq e_k$, it will receive a unit negative reward -1 at the next time instant, so we instantiate (25), i.e., $\mu_R(r_{k+1}|e_k, y_k)$ to assign probability 1 to r_{k+1} such that

$$r_{k+1} = \begin{cases} 0 & \text{if } e_k = y_k \\ -1 & \text{otherwise} \end{cases} \quad (41)$$

Note that now the rewards are determined functionally rather than probabilistically, except for the first reward r_1 , which is immaterial and is still sampled from the marginal $\mu_R(r_1)$.

In a more general setting, the unit punishment -1 could be replaced by a value (called *loss*) which could be different for the two different cases of $e_k \neq y_k$ possible in the present binary setting. Of course, in richer than binary settings, more different mistake kinds exist. We will not bother here with such generalizations.

The agent’s guesses are given by the policy (28) and at any time k they induce a subset of observations analogical to the unknown concept C

$$C(a_k) = \{o \in O \mid \pi(a_k, o) = 1\} \quad (42)$$

$C(a_k)$ is the *agent’s concept* or the *hypothesized concept*. Note that we distinguish the concept C inherent to the environment for the agent’s concept $C(a_k)$




Figure 10: A target concept C and a hypothesized concept $C(a_k)$.

at time k only by indicating the latter to be a function of the agent’s state a_k . To maximize rewards, the agent should evolve its state a_k so that it hypothesized concept co-incides with the target concept C , i.e. from some k on,

$$C(a_k) = C \tag{43}$$

If this is the case, the agent has *identified* the target concept. Figure 10 illustrates the setting.

A crucial question is *how* the agent’s state should be updated so that the guessing accuracy improves. A simple idea would be to let the agent remember a finite number of past observations and their true classes. When a new observation comes, the agent would look up the most similar observation in this memory and guess the class associated with it. Here, similarity could be for example determined by the Hamming distance on the binary observation tuples.

We thus let the state act as a memory for a finite number m of observations and their true classes.

$$a_k = [(o_{k-m+1}, e_{k-m+1}), \dots, (o_{k-1}, e_{k-1}), (o_k, y_k)] \tag{44}$$

The true class e_k of observation o_k can be determined at time $k + 1$ in the obvious way, given the guess y_k made for o_k and the reward r_{k+1} received for that guess. At time k the agent does not know r_{k+1} so for the newest observation it just stores its own guess y_k made according to the class of the most similar observation among $o_{k-m+1}, \dots, o_{k-1}$. In the state update step, the previous guess is replaced by the actual true state and the less recent item is discarded from a_k .

The similarity based approach just explained assumes that similar observations tend to have the same classes. Whether or not such an assumption is justified, the approach hardly merits to be called learning as it rests in plain memorization of observations. We would prefer an agent capable of *generalizing* the

observation towards a *theory*, or *hypothesis* prescribing how observations determine environment states. Such a hypothesis can, for example, take the form of a set of logical rules, an equation, or a program in a programming language.

First consider that the agent’s state is fully defined by the agent’s current hypothesis, which we denote h_k , so

$$a_k = h_k \tag{45}$$

Then the decision policy (28) becomes

$$y_k = \pi(h_k, o_k) \tag{46}$$

and it is then natural to view π as an interpreter (a logical inference mechanism, an equation solver, a program interpreter, etc.) of h_k , while o_k plays the role of input data for h_k , according to which the decision is made. Similarly, (42) can be rewritten into

$$C(h_k) = \{ o \in O \mid \pi(h_k, o) = 1 \} \tag{47}$$

and the equality (43) is rephrased as

$$C(h_k) = C \tag{48}$$

As can be expected, the state update function (29) should change the last hypothesis h_{k-1} towards the current one h_k whenever a wrong guess y_{k-1} was made (recall that due to (41) such a wrong guess is indicated to the agent by $r_k = -1$). The change should lead to a correction of the hypothesis so that the same mistake does not happen again. To conduct such an update step at time k , we will need to know what the previous, wrongly classified observation o_{k-1} was. Since o_{k-1} is not an argument in (29), we need to memorize it within the agent state. This means we cannot get rid completely of a memory component of the agent state. So instead of (45) we rather consider the state to be a tuple consisting of the memorized previous observation and the current hypothesis

$$a_k = (o_k, h_k) \tag{49}$$

The update rule (29) then takes the more specific form

$$\mathcal{A}(a_{k-1}, x_k) = \mathcal{A}((o_{k-1}, h_{k-1}), (o_k, r_k)) = (o_k, h_k) \tag{50}$$

where h_k is determined from h_{k-1} , o_{k-1} , and r_k in a way depending on the particular learning strategy. We will visit some strategies in the coming sections.

While we needed the memorized observation o_k stored as part of the agent’s state, this was just for the purpose of updating the hypothesis. To produce the decision y_k , the memorized observation is not needed so in the remainder of

this chapter, we will still use the notation (46)-(48) including h_k rather than $a_k = (o_k, h_k)$ as an argument.

For simplicity, we also will assume the observations to have binary components (as in the running example in Table 2), so they will be n -tuples ($n \in \mathbb{N}$) from

$$O = \{0, 1\}^n \quad (51)$$

2.1 Generalizing Agent

Recall the example from Table 2 and observe that the components of each observation o_k (columns 2-4) correspond to logical propositions such as “Protein 1 is present” and “Protein 1 is active”, which we can denote p_1, \dots, p_4 (one for each of the four observation columns), as customary in propositional logic. Given (51), the values o_k^1, \dots, o_k^4 carry the logical (truth) values assigned to these respective propositions at time k .

An idea then suggests itself, that the agent’s hypothesis h_k would be a propositional logic formula built with truth-valued variables p_1, \dots, p_4 . Its truth value for the assignment o_k^1, \dots, o_k^4 to the variables would determine the decision y_k . For example, the hypothesis that apoptosis initiates if and only if *protein 1 is present and it is active* would be written as

$$h_k = p_1 \wedge p_2 \quad (52)$$

If h_k is a logical conjunction, then the decision policy (46) takes the more specific form

$$y_k = \pi(h_k, o_k) = \begin{cases} 1 & \text{if } o_k \models h_k \\ 0 & \text{otherwise} \end{cases} \quad (53)$$

where $o_k \models h_k$ means h_k is true given the truth-value assignments o_i to variables p_i , $1 \leq i \leq n$. More precisely, we say that positive (negative, respectively) literal p_i ($\neg p_j$) is *consistent* with observation o_k if $o_k^i = 1$ ($o_k^i = 0$). Then, $o_k \models h_k$ holds if and only if all literals of conjunction h_k are consistent with o_k .

Let us design an agent that learns an unknown conjunction such as the above. The plan is to start with the most specific hypothesis (a conjunction of all literals, i.e. all propositional variables as well as their negations) and then successively delete all literals inconsistent with the received observations. So the initial hypothesis is gradually *generalized* towards the correct one.

In the present example, the agent has the initial hypothesis

$$h_1 = p_1 \wedge \neg p_1 \wedge p_2 \wedge \neg p_2 \wedge p_3 \wedge \neg p_3 \wedge p_4 \wedge \neg p_4 \quad (54)$$

This is the most specific hypothesis as it conjoins all possible conditions (literals). At the same time, this conjunction can of course never be true as it is self-contradictory. However, the agent's strategy is to successively remove from it all the literals that are inconsistent with the coming observations.

After the first percept has been received, i.e. for $k > 1$, the update rule (50) determines h_k according to

$$h_k = \begin{cases} h_{k-1} & \text{if } r_k = 0 \\ \text{delete}(h_{k-1}, o_{k-1}) & \text{otherwise} \end{cases} \quad (55)$$

where

$$\text{delete} \left(\bigwedge_{i \in I} p_i \bigwedge_{j \in J} \neg p_j, (o^1, o^2, \dots, o^n) \right) = \bigwedge_{\substack{i \in I \\ o^i = 1}} p_i \bigwedge_{\substack{i \in I \\ o^i = 0}} \neg p_i \quad (56)$$

So the `delete` function keeps exactly those literals from h_{k-1} which are consistent with o_{k-1} .

How do we know that through such an update rule the agent will indeed improve its guessing so that eventually it will only be receiving non-negative rewards? First, we need to assume that there indeed exists a 'correct' conjunction h^* . It is correct in the sense that if $h_k = h^*$, then (48) holds. In other words,

$$\pi(h^*, o_k) = e_k, \quad \forall o_k \in O \quad (57)$$

To resolve the question, we need a few lemmas.

Lemma 2.1. *$e_k = 1$ if and only if all literals of h^* are consistent with o_k .*

The above lemma follows directly from (53) and (57).

Lemma 2.2. *Whenever `delete`(h_{k-1}, o_{k-1}) is called, $e_{k-1} \neq y_{k-1}$, and if $e_{k-1} = 0$, then all literals of h_{k-1} are consistent with o_{k-1} .*

Proof. To see why Lemma 2.2 is true, note that according to (55), $r_k \neq 0$ when `delete` is called. Due to (41), this means that $e_{k-1} \neq y_{k-1}$. So if $e_{k-1} = 0$ then $y_{k-1} = 1$, but then due to (53), $o_{k-1} \models h_{k-1}$ and so all literals of h_{k-1} are indeed consistent with o_{k-1} . \square

Lemma 2.3. *`delete`(h_{k-1}, o_{k-1}) never removes a literal $l \in h_{k-1}$ which is also in h^* .*

Proof. Assume for contradiction that it does remove a literal $l \in h^*$. First assume $e_{k-1} = 0$. By Lemma 2.2, all literals of h_{k-1} are consistent with o_{k-1} . But because $\text{delete}(h_{k-1}, o_{k-1})$ keeps all literals of h_{k-1} consistent with o_{k-1} , it does not delete l , which is a contradiction. Now consider $e_{k-1} = 1$. Then by Lemma 2.1 all literals of h^* including l must be consistent with o_{k-1} . Again, since delete keeps all consistent literals, it does not delete l , which is a contradiction. \square

The starting hypothesis (54) of the designed agent is set to contain all possible literals, so $h_1 \supseteq h^*$, where the inclusion is with respect to the sets of literals in h_1 and h^* . Furthermore, due to Lemma 2.3, we have

$$h_k \supseteq h^*, \forall k \in N \quad (58)$$

Given the above, the agent makes mistakes only on positive examples, and the mistakes are corrected by removing at least one inconsistent literal, as the following lemma formalizes.

Lemma 2.4. *Assuming (54), whenever $\text{delete}(h_{k-1}, o_{k-1})$ is called, $e_{k-1} = 1$, and the function deletes at least one literal from h_{k-1} .*

Proof. Due to Lemma 2.2, $e_{k-1} \neq y_{k-1}$. If $e_{k-1} = 0$ and $y_{k-1} = 1$ then by the same lemma, all literals of h_{k-1} are consistent with o_{k-1} . According to Lemma 2.1, $e_{k-1} = 0$ means that there is a literal in h^* inconsistent with o_{k-1} . But due to (58), this inconsistent literal would also be contained in h_{k-1} , which is a contradiction. So we know that $e_{k-1} = 1$ and $y_{k-1} = 0$. According to (53), this means that h_{k-1} contains a literal inconsistent with o_{k-1} . Since delete , by (56), keeps exactly all consistent literals, the inconsistent literal is removed. \square

Theorem 2.5. *The generalizing agent makes at most $2n$ mistakes, i.e. the cumulative reward is*

$$\sum_{k=1}^{k'} r_k \geq -2n \quad (59)$$

for an arbitrary horizon $k' \in N$.

Proof. Since the first agent's conjunction has $2n$ literals by (54) and upon each mistake, at least one literal is removed from the conjunction by Lemma 2.4, the maximum number of mistakes is $2n$. \square

While the agent's strategy has been designed to learn conjunctions, it can be also made to learn disjunctions due to the equality

$$\neg(p_1 \vee p_2 \vee \dots \vee p_n) = \neg p_1 \wedge \neg p_2 \wedge \dots \wedge \neg p_n \quad (60)$$

So the only required change is that the agent replaces observations o_k with $\overline{o_k} = (1 - o_k^1, 1 - o_k^2, \dots, 1 - o_k^n)$ and its actions y_k with $1 - y_k$.

Other classes of logical formulas can also be reduced to conjunction and disjunction learning. Consider e.g. s -CNF ($s \in N$). These are conjunctions of s -clauses. An s -clause is a disjunction of at most s -literals. For a finite set of propositional variables, here is a finite number of s -clauses so the agent can simply establish one new propositional variable for each possible s -clause a learn a conjunction with these new variables. This reduction would even be efficient if s is a small constant. Indeed, if n is the number of original variables, then the number of possible clauses is $\binom{2^n}{s}$, i.e., the number of s -combinations of literals chosen from the set of n variables and their n negations. This number grows exponentially with s and polynomially with n . A similar reduction can be used to learn s -DNF.

2.2 The Subsumption Relation

It is instructive to view the generalization process as a path in the *subsumption lattice* of conjunctions shown for two propositional symbols in Fig. 11. A *lattice* is a partially ordered set where each two elements have their unique least upper bound and the greatest lower bound. The subsumption order is given by the subset relation

$$h_1 \subseteq h_2 \tag{61}$$

This means that conjunction h_1 precedes conjunction h_2 if the latter contains all literals of the former.

Recall from logic that a formula h_1 *entails* another formula h_2 if any model of h_1 is also a model of h_2 . We denote this as

$$h_1 \vdash h_2 \tag{62}$$

It is obvious that $h_1 \subseteq h_2$ implies $h_2 \vdash h_1$ if h_1 and h_2 are conjunctions. However, the inverse implication does not hold. For example (observe Fig. 11), we have both $p_1 \wedge \neg p_1 \vdash p_2 \wedge \neg p_2$ and $p_2 \wedge \neg p_2 \vdash p_1 \wedge \neg p_1$ simply because both of the formulas are non-satisfiable and thus neither has a model. However, they do not share any literal so the subset relation does not hold either way. Nevertheless, for satisfiable conjunctions (i.e., conjunctions other than ‘contradictions’) $h_1, h_2, h_1 \subseteq h_2$ is equivalent to $h_2 \vdash h_1$.

While so far, we considered subsumption only conjunctions, the literal subset relation (61) is obviously defined as well for disjunctions, i.e. clauses. However, the relationship to logical entailment becomes inverted. More precisely, for two clauses $h_1, h_2, h_1 \subseteq h_2$ implies $h_1 \vdash h_2$. Just like in the case of conjunctions, we cannot claim equivalence between the two latter relations. For example $p_1 \vee$




Figure 11: Subsumption lattice for conjunctions. The conjunction symbols \wedge are omitted for brevity. The curved arrows show how the agent generalizes its initial conjunction in two steps following the successive observations $(1, 0)$ and $(1, 1)$ carrying the respective truth values for p_1 and p_2 . All conjunctions below the dashed line are non-satisfiable.

$\neg p_1 \vdash p_2 \vee \neg p_2$. Again, the problem is with the atoms included both as a positive and a negative literals. While in conjunctions they produced contradictions, their presence in clauses make the latter tautologies, i.e. formulas true in any interpretation. But analogically to conjunctions, $h_1 \subseteq h_2$ is equivalent to $h_2 \vdash h_1$ if h_1, h_2 are not tautologies.

Contradictory conjunctions and tautological clauses have one property in common. They contain a positive literal as well as the negation of the same literal. Clauses, which have this property, are called *self-resolving*.¹

2.3 Separating agent

In Section 2.1 we have designed the generalization agent able to learn a conjunction while making only a finite number of mistakes. We have also seen that through efficient conversions, such an agent can also learn disjunction, *s*-DNF's and *s*-CNF's.

As an alternative example of a learning agent, we will demonstrate one using a different strategy to achieve the same goal. This time, the agent's hypothesis h_k will be represented by non-logical means. In particular, h_k will define a

¹As the adjective *self-resolving* originates from the resolution principle, which is applied on clauses and not on conjunctions, it is usually not associated with conjunctions.

hyperplane in the $O = \{0, 1\}^n$ space (51) so $C(h_k)$ (47) will include exactly those observations lying above the hyperplane.

Formally, h_k is an n -tuple of integer values bounded by some constant $q \in N$, i.e. $h_k \in [0, 1, \dots, q]^n$, so

$$h_k = [h_k^1, h_k^2, \dots, h_k^n] \quad (63)$$

The agent's decision policy (46) is given by a threshold function applied on a dot product

$$y_k = \pi(h_k, o_k) = \begin{cases} 1 & \text{if } h_k \cdot o_k > n/2 \\ 0 & \text{otherwise} \end{cases} \quad (64)$$

The initial hypothesis is

$$h_1 = (1, 1, \dots, 1) \quad (65)$$

And the update rule (50) is instantiated according to

$$h_k = \begin{cases} h_{k-1} & \text{if } r_k = 0 \\ \text{update}(2, h_{k-1}, o_{k-1}) & \text{if } h_{k-1} \cdot o_{k-1} \leq n/2 \\ \text{update}(0, h_{k-1}, o_{k-1}) & \text{if } h_{k-1} \cdot o_{k-1} > n/2 \end{cases} \quad (66)$$

wherein the function **update** is defined such that for $h_k = \text{update}(\alpha, h_{k-1}, o_{k-1})$ and each $i = 1, 2, \dots, n$,

$$h_k^i = \begin{cases} \alpha \cdot h_{k-1}^i & \text{if } o_{k-1}^i = 1 \\ h_{k-1}^i & \text{otherwise} \end{cases} \quad (67)$$

This agent, which can be considered an integer counterpart of the popular perceptron algorithm, learns a hyperplane. On the other hand, the generalization agent from the previous section was designed to learn logical formulas, namely conjunctions, disjunctions, s -DNF's, and s -CNF's. So how can we compare the two agents?

Assume that the target concept C corresponds to a disjunction c consisting of s literals made out of the variables p_1, \dots, p_n . That is to say, $\mu_O(o|1) > 0$ if and only if $o \models c$. It is well known that disjunctions are linearly separable, so for a sufficiently large q , there is a hyperplane h^* such that (57) holds. This means, that the agent can identify a target disjunction through its hypothesis, although the latter is a hyperplane rather than a disjunction. The theorem below states that it does so with a finite number of mistakes.

Theorem 2.6. *The agent makes at most $2+2s \lg n$ mistakes, i.e. the cumulative reward is*

$$\sum_{k=1}^{k'} r_k \geq -2 - 2s \lg n \quad (68)$$

for any horizon $k' \in N$.

(proof omitted)

Just like the generalizing agent designed to learn conjunctions could easily be modified to learn disjunctions, s -CNF's, and s -DNF's, also the separating agent can be altered to learn conjunctions as well as the latter two classes by means of the same reduction principles. So the two agents can in principle learn the same concept classes. The difference is in the mistake bound. The latter agent performs better when the number of variables n is larger than the number of relevant variables s .

2.4 Hypothesis and Concept Classes

So far we have designed two exemplary learning agents, each with a different set of hypotheses h it could express. We shall call the set of all hypotheses an agent can express its *hypothesis class* denoted with the letter \mathcal{H} and we will assume \mathcal{H} to be finite. For the generalizing agent, \mathcal{H} consisted of all conjunctions made of at most n variables. For the separating agent, \mathcal{H} was the set of (q -bounded) n -tuples of integers.

Considering (47), each hypothesis class induces a set of concepts

$$\mathcal{C}(\mathcal{H}) = \{ C(h) \mid h \in \mathcal{H} \} \quad (69)$$

called a *concept class*.

The two agents exemplified so far used their update rules specifically designed for their respective hypothesis classes. Can we design a more general learning agent which could work with an arbitrary hypothesis class \mathcal{H} ?

Assume that \mathcal{H} is rich enough to contain a hypothesis h matching the unknown target concept $C(h) = C$ (39). This assumption can be written as

$$C \in \mathcal{C}(\mathcal{H}) \quad (70)$$

Under such an assumption, since \mathcal{H} is finite, the agent can always try successively each element $h \in \mathcal{H}$, discarding it as soon as a mistake is made (negative

reward received) using that hypothesis. In the worst case, the last hypothesis remaining will match the target concept. This means that the maximum number of mistakes made before identifying the target concept is

$$|\mathcal{H}| - 1 \tag{71}$$

This is a first indication of a dilemma we are going to face repeatedly in different forms. In particular, given that C is unknown, the agent should possess a large hypothesis space to maximize chances that (70) is satisfied. On the other hand, a large hypothesis space entails a loose bound on the number of mistakes made according to (71).

2.5 Version Space Agent

The general mistake bound (71) can be readily improved to $\lg |C|$ using the *version space* strategy. Informally, its main idea is that on each observation, the agent discards all hypotheses from the hypothesis class which are inconsistent with the observation.

Before formalizing the principle, we will explain it by contrasting it to the generalization agent from Section 2.1. The latter agent worked with the hypothesis class \mathcal{H} of conjunctions made of at most n variables, i.e., at each time k , hypothesis $h_k \in \mathcal{H}$. The main idea of the version-space agent is to keep at each time k the subset of all hypotheses from \mathcal{H} which are consistent with all the observations received so far.

So instead of (72), the version-space agent’s state is given as

$$a_k = (o_k, \mathcal{H}_k) \tag{72}$$

where o_k is again the memorized observation and \mathcal{H}_k is a set of hypotheses, also called the *version space*.

Decisions are determined by *voting* among all hypotheses in \mathcal{H}_k . Assuming that \mathcal{H} consists of logical formulas, this is formalized as

$$y_k = \pi(\mathcal{H}_k, o_k) = \begin{cases} 1 & \text{if } |\{ h_k \in \mathcal{H}_k \mid o_k \models h_k \}| > |\mathcal{H}_k|/2 \\ 0 & \text{otherwise} \end{cases} \tag{73}$$

The agent starts with the full hypothesis space

$$\mathcal{H}_1 = \mathcal{H} \tag{74}$$

The hypothesis update step takes a form slightly different from (50), in particular

$$\mathcal{A}(a_{k-1}, x_k) = \mathcal{A}((o_{k-1}, \mathcal{H}_{k-1}), (o_k, r_k)) = (o_k, \mathcal{H}_k) \tag{75}$$

where \mathcal{H}_k is obtained by deleting all hypotheses inconsistent with o_{k-1} , i.e.

$$\mathcal{H}_k = \begin{cases} \{ h \in \mathcal{H}_{k-1} \mid o_{k-1} \models h \} & \text{if } e_{k-1} = 1 \\ \{ h \in \mathcal{H}_{k-1} \mid o_{k-1} \not\models h \} & \text{if } e_{k-1} = 0 \end{cases} \quad (76)$$

where e_{k-1} is determined as $e_{k-1} = |y_{k-1} + r_k|$ (check that this equality is true) and $y_{k-1} = \pi(\mathcal{H}_{k-1}, o_{k-1})$. To see that the latter expression can be evaluated at the hypothesis update step, note that both \mathcal{H}_{k-1} and o_{k-1} are indeed available at that step according to (76).

In (73) and (76) we used the relation \models assuming that h_k is a logical formula and o_k provides the truth-values for its atoms (variables). If this was not the case, we could assume more generally that hypotheses provide a mapping $h_k : O \rightarrow \{0, 1\}$. Then in the two formulas, $o_k \models h_k$ ($o_k \not\models h_k$, respectively) would be replaced by $h_k(o_k) = 1$ ($h_k(o_k) = 0$).

Assume that \mathcal{H} is rich enough so that (70) is true. Then the following theorem holds.

Theorem 2.7. *The agent makes at most $\lg |\mathcal{H}|$ mistakes, i.e. the cumulative reward is*

$$\sum_{k=1}^{k'} r_k \geq -\lg |\mathcal{H}| \quad (77)$$

for any horizon $k' \in \mathbb{N}$.

Proof. To see why the theorem holds note that the agent decides by the majority of hypotheses from \mathcal{H}_k . So if a mistake is made, at least half of the hypotheses in \mathcal{H}_k are deleted. In the worst case, the last remaining hypothesis is correct. \square

Once again, a dilemma is observed in that \mathcal{H} should be large enough so that (70) is satisfied. However, the size $|\mathcal{H}|$ also increases the mistake bound (77). The latter mistake bound is logarithmic, which is certainly a significant improvement over (71) good but the computational demands for storing \mathcal{H}_k (containing a potentially large number of hypotheses) can be prohibitive.

2.6 The Mistake Bound Learning Model

The linear mistake bounds we obtained for the generalizing and separating agents indicate that these agents are indeed able to learn well the conjunctive and disjunctive concepts but also other kinds of concepts (namely, *s*-DNF and *s*-CNF) that can be reduced to the latter. We will now generalize the notion of ‘good on-line learning.’ We say that an agent *learns class \mathcal{H} on-line (in*

the mistake-bound model) if it makes at most $\text{poly}(n)$ of mistakes in the on-line scenario provided that (70) holds. Here, poly is a polynomial and n is the size of observations. With our setting (51), the size of observations is the number n of binary values making up the observations.

By Theorems 2.5 and 2.6, the generalizing and separating agents learn \mathcal{H} on-line when \mathcal{H} are conjunctions, disjunctions, s -DNF or s -CNF (for a constant s). By Theorem 2.7, for an arbitrary class \mathcal{H} , the version-space algorithm has a mistake bound $\lg |\mathcal{H}|$ as long as (70) holds true. If furthermore $|\mathcal{H}|$ is at most exponential in n , the agent necessarily learns \mathcal{H} on-line, because the mistake bound $\lg |\mathcal{H}|$ is then polynomial. The condition ‘at most exponential’ above seems rather permissive. But note that $|\mathcal{H}|$ may easily be super-exponential. The extreme example of such a situation is the class \mathcal{H} such that (70) is true for any possible concept C , $C \subseteq O$. Since $|O| = |\{0,1\}^n| = 2^n$, we have $|\mathcal{H}| \geq 2^{|O|} = 2^{2^n}$, so $|\mathcal{H}|$ is super-exponential.

Furthermore, we refine the definition into a stricter form. An agent that learns hypothesis class \mathcal{H} on-line is said to learn it *efficiently* if it spends at most polynomial time (in n) between the receipt of a percept and the generation of the next action.

It is quite easy to verify that both the generalizing and separating agents learn their hypothesis classes efficiently. However, the version-space agent is evidently not efficient if $|\mathcal{H}|$ is super-polynomial in n as it has to visit each element of \mathcal{H} during the update (76). Say \mathcal{H} are conjunctions. There are $2n$ literals, each of which can be present or absent in a conjunction so there are $|\mathcal{H}| = 2^{2n}$ conjunctions and the version-space agent does not learn conjunctions efficiently. We could reduce the size of \mathcal{H} realizing that conjunctions containing the same atom both in a positive literal and a negative literal are unsatisfiable and can be discarded. Then we would derive $|\mathcal{H}|$ as follows. We have n atoms, each of which may be positive, negative, or absent in a conjunction. Thus there are 3^n conjunctions in total, which grows slower than 2^{2n} but still exponentially.

What about a *lower bound* on mistakes? The latter can be established using the notion called *VC-dimension* of a hypothesis class. We say that a set of observations $O' \subseteq O$ is *shattered* by hypothesis class \mathcal{H} if

$$\{ O' \cap C(h) \mid h \in \mathcal{H} \} = 2^{O'} \quad (78)$$

which means that the set of observations can be partitioned in all possible ways into two classes by the hypotheses from \mathcal{H} . The *Vapnik-Chervonenkis Dimension* (or VC-dimension) of \mathcal{H} , written $\text{VC}(\mathcal{H})$, is the cardinality of the largest set $O' \subseteq O$ that is shattered by \mathcal{H} .

Theorem 2.8. *Assume an agent deciding by (46) where $h_k \in \mathcal{H}$. No upper bound on the number of mistakes made by the agent is smaller than $\text{VC}(\mathcal{H})$.*

Proof. For any sequence of agent’s decisions $y_1, y_2, \dots, y_{\text{VC}(\mathcal{H})}$ there exists a $h \in \mathcal{H}$ according to which all these decisions are wrong. \square

The theorem does not directly apply to the version-space agent, which decides according to (73) rather than (46). This significance of $\text{VC}(\mathcal{H})$ for this agent is that it will not identify the target hypothesis (i.e. reduce $|\mathcal{H}_k|$ to 1) in fewer than $\text{VC}(\mathcal{H})$ steps. Indeed, $\mathcal{H}_{\text{VC}(\mathcal{H})}$ will contain two hypotheses, one deciding 1 for $o_{\text{VC}(\mathcal{H})}$ and another deciding otherwise.

3 Batch Concept Learning

We will maintain the assumption of concept learning (38), that is, no observation can be generated from two different states of the environment. For simplicity, we will continue working in the binary setting (40), i.e. the environment has only two different states and the agent can make only two different actions. Also, the observations are still assumed to be binary (51). Finally, rewards are still determined according to (41), i.e., -1 is a punishment for a wrong guess.

However, instead of the on-line setting, we will now study concept learning in the batch framework we defined in Section 1.7. Since observations are i.i.d. here, we can express the probability distribution of reward r_{k+1} by (35) as $\mu_r(r_{k+1}|a_k)$, only depending on the current agent’s state. Considering that the decision for an observation o_k is made only using the hypothesis component h_k of the state a_k (46), this is the same as writing $\mu_r(r_{k+1}|h_k)$. As follows from (41), $\mu_r(-1|h_k)$ is the probability that the decision $y_k = \pi(h_k, o_k)$ is wrong, i.e., $y_k \neq e_k$. We will call this probability the *error* of h_k

$$\text{err}(h_k) = \mu_r(-1|h_k) \tag{79}$$

Remind that the agent’s goal in the batch setting is to reach at the end of the training phase (i.e., at time K), a state a_K which maximizes the expected reward (36). Again, since only the hypothesis component of the agent’s state determines the decision, the expected reward is $\sum_{r_k \in R} \mu_R(r_k|h_K)r_k$. This expectation is equal to $-\text{err}(h_K)$ as can be easily seen realizing that rewards can be only 0 or -1.

So the goal of batch learning can be equivalently stated as arriving at a hypothesis h_K minimizing the error. A natural question of interest is how the algorithms we designed for on-line concept learning in the sequential scenario would perform in terms of the error. Evidently, the bounds on the number of mistakes we established in Theorems 2.5, 2.6, and 2.7 do not translate to any

bound on $\text{err}(h_K)$ as there is no guarantee that the mistakes will happen in the learning phase ($k \leq K$) where the agent still can fix its hypothesis.

But unlike in the on-line learning case, the batch case inherits the non-sequential assumptions (30) and (32), meaning that states and observations are sampled i.i.d. according to distributions that do not change with k . They prevent the environment from ‘adversarial’ behavior, for example, one where the training phase would only contain ‘easy’ examples and the ‘hard’ ones would be kept for the testing phase. As we will see, in this scenario we can indeed bound $\text{err}(h_K)$ for particular learning agents, although we will be able to do it only with certain probability smaller than 1.

3.1 Batch Learning with the Generalizing Agent

Consider the generalizing agent as described in Section 2.1 working in the learning phase ($k \leq K$) of the batch scenario just as it worked in the on-line scenario.

Denote by $\Pr(l)$ the probability that literal $l \in h_k$ is deleted by (56) while updating to h_{k+1} . Since the agent makes mistakes only on positive examples, this means that o_k is positive and l is inconsistent with it. Observations are i.i.d. in the non-sequential batch setting so $\Pr(l)$ does not depend on k .

According to (55), deletion from h_k takes place if and only if $r_{k+1} = -1$. So considering (79), $\text{err}(h_k)$ is equal to the probability that *some* of the literals in h_k get deleted, which in turn can be bounded by the sum

$$\text{err}(h_k) \leq \sum_{l \in h_k} \Pr(l) \tag{80}$$

We have no more than $2n$ literals in h_k so if $\Pr(l) \leq \epsilon/2n$ ($\epsilon \in \mathbf{R}$) for each of them then $\text{err}(h_k) \leq \epsilon$. Let us fix an arbitrary ϵ and let us call a literal *bad* if $\Pr(l) > \epsilon/2n$.

Consider a bad literal l from h_1 (54). The probability that it will still be in h_2 (i.e., that it will ‘survive’ the first observation) is $1 - \Pr(l)$. Appealing again to the i.i.d. character of observations, the probability of $l \in h_{k+1}$ is then

$$(1 - \Pr(l))^k < \left(1 - \frac{\epsilon}{2n}\right)^k \tag{81}$$

There are at most $2n$ bad literals so the probability that *some* of them is in h_{k+1} is at most

$$2n \left(1 - \frac{\epsilon}{2n}\right)^k \tag{82}$$

To work with this upper bound easily, we make use of the inequality $1 - x \leq e^{-x}$ which holds for $x \in [0; 1]$, to obtain

$$2n \left(1 - \frac{\epsilon}{2n}\right)^k \leq 2ne^{-k\frac{\epsilon}{2n}} \quad (83)$$

So hypothesis h_{k+1} of the generalizing agent in the learning phase ($k < K$) has $\text{err}(h_{k+1}) \leq \epsilon$ with probability at least $1 - 2ne^{-k\frac{\epsilon}{2n}}$, and specifically for h_K we can state the following theorem.

Theorem 3.1. *The probability that the generalizing agent's hypothesis h_K has error $\text{err}(h_K) \leq \epsilon$ is at least $1 - 2ne^{-m\frac{\epsilon}{2n}}$ where $m = K - 1$.*

3.2 Batch Learning with General On-line Agents

We define a *standard* on-line agent as one that uses a single hypothesis to determine its decisions (46), and changes the hypothesis if and only if a mistake has been made by the previous hypothesis. This includes the generalizing and separating agents as follows from the update rules (55) and (66). On the other hand, the version-space agent is not standard for obvious reasons.

The next lemma will enable us to accommodate any standard on-line learning agent for the batch learning scenario with a probabilistic bound on the error of the learned hypothesis.

Lemma 3.2. *If a standard on-line agent retains a hypothesis h_k for q steps ($h_k = h_{k+1} = \dots h_{k+q}$), then $\text{err}(h_k) \leq \epsilon$ with probability at least $1 - e^{-q\epsilon}$.*

Proof. The probability that the standard agent retains a bad hypothesis (i.e., one for $\text{err}(h_k) > \epsilon$) on receiving an observation is the probability $1 - \text{err}(h_k)$ that the bad hypothesis produces a correct decision for that observation. Since $\text{err}(h_k) > \epsilon$, the probability is at most $1 - \epsilon$. The probability of keeping the hypothesis over q i.i.d. observations is thus at most $(1 - \epsilon)^q$, and we already know that $(1 - \epsilon)^q \leq e^{-q\epsilon}$. Otherwise, i.e. with probability at least $1 - e^{-q\epsilon}$, the hypothesis was not bad, i.e $\text{err}(h_k) \leq \epsilon$. \square

So the rule is: wait until $h_k = h_{k+1} = \dots h_{k+q}$ happens and then keep h_k with the probabilistic error bound. The question is how to guarantee that the event indeed happens within the learning phase, i.e. $k + q \leq K$. If we have a mistake bound M for the agent, we know that the standard agent makes at most M hypothesis changes. In this case we set the learning phase long enough, in particular $K = Mq$, to guarantee that one of the hypotheses in the learning phase survives at least q observations.

3.3 Consistent Agent

Here we design a general agent working with an arbitrary hypothesis space. This is analogical to the version space agent we studied in the on-line setting.

We first adapt the version space agent from on-line learning to batch learning. In the learning phase, the agent works just as in the on-line setting. When the phase ends, i.e. $k = K$, the agent updates its state for the last time according to (76), and then selects an *arbitrary* hypothesis h_K from \mathcal{H}_K .

Once again, let us call a hypothesis $h \in \mathcal{H}$ *bad* if $\text{err}(h) > \epsilon$. The probability that a particular bad hypothesis $h \in \mathcal{H}$ is in h_K is at most $(1 - \epsilon)^m \leq e^{-\epsilon m}$ where² $m = K - 1$. The probability that *some* bad hypothesis from \mathcal{H} (74) is in h_K is at most $|\mathcal{H}|e^{-\epsilon m}$. So the probability that no bad hypothesis survives and thus $\text{err}(h_K) \leq \epsilon$ whichever h_K the agent has picked from \mathcal{H}_K , is at least $1 - |\mathcal{H}|e^{-\epsilon m}$.

Maintaining the sets $\mathcal{H}_1, \mathcal{H}_2, \dots$ along the course of learning can certainly be prohibitive as the latter can be very large sets. A behavior equivalent to the described version-space strategy can however be obtained without maintaining the version spaces.

In particular, the *consistent agent* collects all observations seen during the training phase along with their true classes in its memory without updating its hypothesis. Only at the end of the training phase, it picks from its hypothesis class a hypothesis consistent with all the collected observations.

Formally, the agent has state

$$a_k = ((o_1, e_1, \dots, o_{k-1}, e_{k-1}), h_k) \quad (84)$$

for time $k = 1, 2, \dots, K - 1$. Here, h_k is just any ‘dummy’ hypothesis that makes $y_k = \pi(h_k, o_k) = 0$ for any observation $o_k \in O$ (e.g., an unsatisfiable formula). With $y_k = 0$ the agent can indeed determine the true classes in (84) as $e_k = -r_{k+1}$. Only at time K , the agent finds any hypothesis $h_K \in \mathcal{H}$ consistent with the collected set, i.e. $\pi(h_K, o_k) = e_k$ for all $k < K$. Analogically, to the reasoning above, we have that

Theorem 3.3. *The probability that the consistent agent’s hypothesis h_K has error $\text{err}(h_K) \leq \epsilon$ is at least $1 - |\mathcal{H}|e^{-\epsilon m}$ where $m = K - 1$.*

²Remind that o_K is the first testing observation (its true class e_k cannot be determined at time K). So $m = K - 1$ is the number of observations with known class (label) received in the training phase. In other words, m is the number of labeled training examples received by the agent.

The first component of the agent’s state at time K , i.e.

$$o_1, e_1, \dots, o_m, e_m \tag{85}$$

is called the *training set*. Unlike the version space agent, the consistent agent does not store large hypothesis spaces and so it obviously consumes less memory. On the other hand, finding a hypothesis h_K consistent with the training set may be computationally hard, and potentially intractable.

The generalizing agent from Section 2.1 can be formally viewed as a consistent agent simply by letting it delete inconsistent literals after all training observations have been received, rather than deleting them on-line. Thus the bound $1 - |\mathcal{H}|e^{-\epsilon m}$ from Theorem 3.3 applies to this agent, wherein $|\mathcal{H}| = 2^{2n}$ is the number of conjunctions which can be made out of $2n$ literals. The latter bound is asymptotically (for growing m) higher (better) than $1 - 2ne^{-m\frac{\epsilon}{2n}}$ from Theorem 3.1. This is interesting given that the latter theorem was derived for the specific agent rather than for a class of agents as in the case of Theorem 3.3. The reason for that is the rather loose bound (80) we adopted to derive Theorem 3.1. Note that the formula in the latter theorem is ‘stronger’ in the sense that it is a lower bound on the probability that *the learned conjunction will contain no bad literal*, which is a sufficient but not necessary condition for $\text{err}(h_K) \leq \epsilon$.

3.4 The PAC Learning Model

Agent *probably approximately learns class \mathcal{H} (in the batch setting)* if at the end of the training phase it produces h_K such that $\text{err}(h_K) \leq \epsilon$ with probability at least $1 - \delta$, and $K - 1 = m \leq \text{poly}(n, 1/\delta, 1/\epsilon)$, where *poly* is a polynomial.

“probably approximately learns” = “PAC-learns” (C for correctly)

It PAC-learns the class *efficiently* if it spends at most polynomial (in the same variables) time between the receipt of a percept and the generation of the next action in the training phase.

Theorem 3.4. *The generalizing agent efficiently PAC-learns conjunctions.*

Proof. Efficiency is obvious: at most $2n$ unit steps (going over literals) for each of $m = K - 1$ observations. From Theorem 3.1, the probability that $\text{err}(h_K) > \epsilon$ is at most $2ne^{-m\frac{\epsilon}{2n}}$. It remains to determine how many observations m are needed to make the probability smaller than a given δ and see if the result is

polynomial.

$$\begin{aligned} \delta &> 2ne^{-m\frac{\epsilon}{2n}} \\ \frac{\delta}{2n} &> e^{-m\frac{\epsilon}{2n}} \\ \ln \frac{\delta}{2n} &> -m\frac{\epsilon}{2n} \\ \frac{2n}{\epsilon} \ln \frac{2n}{\delta} &\leq m \end{aligned}$$

So the required $m = K - 1$ is indeed polynomial in n , $1/\epsilon$ and $1/\delta$. \square

Theorem 3.5. *Any standard agent learning (efficiently) a concept class C online, has a counterpart which (efficiently) PAC-learns C .*

Proof. The agent makes at most $M < \text{poly}(n)$ updates, i.e. max number of mistakes.

Its batch counterpart works as follows.

$$\text{Set } q = \frac{1}{\epsilon} \ln \frac{1}{\delta}$$

If before M updates have been made, each hypothesis survived for less than q steps, then the last one (which makes no mistakes) is found in at most Mq steps, and is kept as h_K . Both M and q are polynomial.

If some of them survived for at least q steps, than according to lemma (3.2), its error is less than ϵ with probability at least

$$1 - e^{-q\epsilon} = 1 - e^{-\epsilon \frac{1}{\epsilon} \ln \frac{1}{\delta}} = 1 - \delta \tag{86}$$

This hypothesis found with less than Mq (poly) steps, will be kept as h_K . \square

So a negative batch (PAC) result also means a negative on-line result.

Theorem 3.6. *If (70) holds and $|\mathcal{H}|$ is at most exponential in n then the consistent agent PAC-learns \mathcal{H} .*

Proof. By Theorem (3.3), probability that $\text{err}(h) > \epsilon$ is at most $|\mathcal{H}|e^{-\epsilon m}$ and we need this to be smaller than δ . The required m is then set as

$$|\mathcal{H}|e^{-\epsilon m} < \delta$$

$$e^{-\epsilon m} < \frac{\delta}{|\mathcal{H}|}$$


$$m \geq \frac{1}{\epsilon} \ln \frac{|\mathcal{H}|}{\delta}$$

Since $|\mathcal{H}|$ is at most exponential in n , $\ln |\mathcal{H}|$ is at most polynomial in it, so setting m to the right-hand side of the above inequality means that $m < \text{poly}(1/\epsilon, 1/\delta, n)$. \square

Also, s -CNF and s -DNF learnable by poly reduction to conjunctions.

4 Learning First-Order Logic Concepts

We now revisit the agent-scientist from Section 2, although the agent will now be in a slightly different situation. In particular, it will investigate chemical compounds, that is, structures such as



and learn to predict for each compound whether it is toxic or safe. The important distinction from the story captured in Table 2 is that there is now no obvious way to encode structure such as the above through tuples $o = (o^1, o^2, \dots)$ of truth values. Here, observations are graphs and we need a language more expressive than propositional logic to describe such graphs, and also to form hypotheses about them. To this end, we will use the language of first-order predicate logic.

First, we will simplify the situation by abstracting from the parts not important for studying the learning principles. In particular, we will ignore the types of chemical elements in the vertices and also the bond types (single, double). We will simply assume that observations are oriented³ graphs, that is, directed graphs in which no two vertices are connected in both directions. We will assign unique numbers to vertices so that the latter can be addressed later. This is exemplified in Table 2.

³Orientation of edges may e.g. correspond to charge distribution along the bonds; we assume oriented edges as their representation is simpler in logic than that of non-oriented edges.

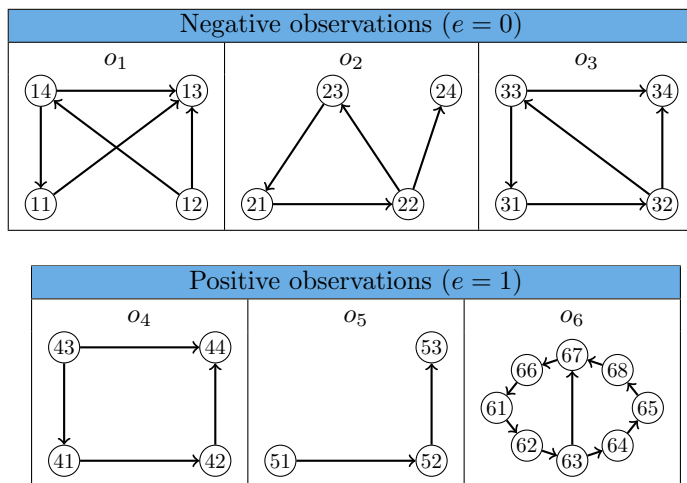


Table 2: Graphical observations from which the agent should learn to classify new observations as negative or positive.

Encoding the graphs shown in Table 2 through the language of predicate logic is straightforward. For each graph, we will simply list all of its edges as ground facts of the binary predicate `edge`. So, for example, the second negative observation will be represented as

$$o_2 = \{ \text{edge}(21, 22), \text{edge}(22, 23), \text{edge}(23, 21), \text{edge}(23, 24) \} \quad (87)$$

and the second positive observed compound will be encoded as

$$o_5 = \{ \text{edge}(51, 52), \text{edge}(52, 53) \} \quad (88)$$

Note that this representation is perfectly analogical to the one used in Section 2.1. In the latter, the observations (51) were truth values assigned to propositional symbols. In (87) and (88) we implicitly assign truth values to all possible ground facts of `edge/2` by including exactly those ground facts which hold true, i.e. listing all edges actually in the graph. In both cases, the truth value assignment is called an *interpretation*. In the present first-order context, interpretations such as those shown above, which assign truth values directly to ground atoms of the logical language, are called *Herbrand interpretations*.

Naturally, for more complex problems, the vocabulary of predicates would include more predicates than just `edge/2`, and an interpretation would define the truth values of all ground facts of all the predicates.

Returning to the example in Table 2, we would like to design an agent able to learn what patterns are common for only the positive observations (safe compounds) so that such patterns can be used for classifying compounds observed




Figure 12: Each negative and no positive observation from Table 2 contains one of these two kinds of triangles.

in the future. As we are also intelligent agents, we observe that none of the safe compounds (and each of the toxic ones) contains a triangle. As edges are directed, a triangle may take one of the two forms shown in Fig. 12.

Any other oriented triangle is isomorphic to one of the two shown in the figure. The following formula γ_1 in predicate logic expresses that a graph does not contain the first (left) kind of triangle

$$\gamma_1 = \forall x, y, z : \neg \text{edge}(x, y) \vee \neg \text{edge}(y, z) \vee \neg \text{edge}(z, x) \quad (89)$$

In the logical representation of graphs, this means that e.g. the negative interpretation o_2 above will not be a *model* of this formula. Indeed $o_2 \not\models \gamma_1$ since there exists a substitution θ , namely

$$\theta = \{ x \mapsto 21, y \mapsto 22, z \mapsto 23 \} \quad (90)$$

making all the literals in $\gamma\theta$ false with respect to o_2 , since all of $\text{edge}(21, 22)$, $\text{edge}(22, 23)$, and $\text{edge}(23, 21)$ are in o_2 . Analogously, $o_3 \not\models \gamma_1$, so the third negative example is also ‘eliminated’ by γ_1 . However, the formula does not eliminate o_1 – indeed $o_1 \models \gamma_1$. This observation does not contain the first kind of triangle shown in Fig. 12. However, it contains the second one, which is in turn eliminated by the formula

$$\gamma_2 = \forall x, y, z : \neg \text{edge}(x, y) \vee \neg \text{edge}(z, x) \vee \neg \text{edge}(z, y) \quad (91)$$

As before, we can check easily that $o_1 \not\models \gamma_2$ using either the substitution $\theta = \{ x \mapsto 14, y \mapsto 11, z \mapsto 13 \}$ or the substitution $\theta = \{ x \mapsto 12, y \mapsto 14, z \mapsto 13 \}$.

In summary, the conjunction

$$\gamma_1 \wedge \gamma_2 \quad (92)$$

eliminates all negative observations. On the other hand, both γ_1 and γ_2 are true in all the positive observations (there is no substitution to x, y, z making all literals of γ_1 or γ_2 false in them), so indeed the above conjunction perfectly discriminates between the positive and negative observations.

Note that in a substitution, different variables can map to the same term. So e.g. γ_1 would also be false in interpretations which we did not intend such as

$$o' = \{ \text{edge}(1, 1), \text{edge}(1, 2), \text{edge}(2, 1) \} \quad (93)$$

as with the substitution

$$\theta = \{ x \mapsto 1, y \mapsto 1, z \mapsto 2 \} \tag{94}$$

all literals $\gamma_1\theta = \neg\text{edge}(1, 1) \vee \neg\text{edge}(1, 2) \vee \neg\text{edge}(2, 1)$ are false with respect to o' . Similarly, one can check easily that γ_2 is false e.g. in interpretation $o'' = \{ \text{edge}(1, 2), \text{edge}(2, 1) \}$. This would be a problem if o' or o'' were positive observations, since they would be eliminated incorrectly by $\gamma_1 \wedge \gamma_2$. There is an apparent reason why o' or o'' should be classified as positive (safe): these observations do not contain a proper triangle. However, note that neither o' or o'' represent an *oriented* graph as they contain mutually inverse edges and thus they cannot come as observations at all.

4.1 Generalizing Agent

In the illustrative example above, the formulas γ_1 and γ_2 were first-order logic clauses, that is, universally quantified disjunctions of first-order logic literals. Thus (92) is a first-order logic CNF. We will design an agent able to learn such CNF's in the on-line scenario. A CNF is a logical formula whose truth-value for a given observation can be determined, we naturally adopt the decision policy as defined in (53).

We will keep the basic concept-learning assumptions, i.e. observations uniquely map (39) to binary states (37). Note that these assumptions fit well the illustrative example above. Also, a reward will punish an incorrect prediction just as was done through (41). Observations, however, will no longer be tuples of truth values as in (51). They will take the form of (Herbrand) interpretations of maximum cardinality n . In the illustrative graph example above, n is the number of edges in the largest observed graph.

Recall the generalizing agent from Section 2.1, whose strategy was to gradually remove all constituents of it hypothesized conjunction inconsistent with the received observations. The agent's initial hypothesis (54) conjoined all possible atoms using the variables $p_1, p_2 \dots p_n$ fixed before the learning session. So these n variables determined the size of the the initial conjunction and n was a suitable measure of the learning-task complexity.

The present agent will follow an analogical strategy, except that *first-order clauses*—rather than *propositional literals*—will form the hypothesized conjunction. So the initial hypothesis should take the form

$$\phi_1 = \bigwedge_{\gamma \in \Gamma} \gamma \tag{95}$$

where Γ is the set of all clauses the agent can express. Here, the situation becomes more complex as Γ is not determined simply by a set of n variables,

but rather by a finite set of first-order predicates P and a finite set of functions F . So in the running example, P would contain only the binary predicate `edge`, and F would contain constants⁴ at least those occurring in the observations (see Table 2).

Although we requested that both P and F be finite, in general this does not guarantee Γ to be finite. In particular, Γ is infinite when functions of arity 1 or greater are in F . E.g., in the running example, we can consider a function `next` mapping a vertex onto another vertex. Its purpose is irrelevant, it could for example encode some linear order on vertices. Importantly, such a function would give rise to an infinite set of first-order atoms including `edge(1,next(1))`, `edge(next(1),next(next(1)))`, and so on. This of course means that also Γ is infinite.

We do need Γ finite for (95) to make sense. We will do it by requesting that the number of literals in any $\gamma \in \Gamma$ is no greater than $s \in \mathbb{N}$, and each of the literals contains at most $t \in \mathbb{N}$ occurrences of predicates, functions or variables. A clause complying with these constraints is called an *st-clause*. Clearly, Γ is finite under such a constraint.

Let us return to the notion of complexity of a particular learning task, which in the propositional-logic version of the generalizing agent was simply the number of propositional variables n determining the size of each observation. In the current context, n is again the size of the largest observation, which is now a Herbrand interpretation. However, besides n , the task size is now also determined by the number of predicates and functions available, i.e., $|P|$, $|F|$. So, when we say that the present agent learns in the mistake-bound learning model (c.f. Section 2.6), its mistake bound must be polynomial each of the three quantities n , $|P|$, and $|F|$.

Finally, we can elaborate how the present agent will update its hypothesis from observations. The hypothesis update function will be the same as in (55)-(56), except the `delete` function is different:

$$\text{delete} \left(\bigwedge_{i \in I} \gamma_i, o \right) = \bigwedge_{\substack{i \in I \\ o \models \gamma_i}} \gamma_i \quad (96)$$

In other words, on receiving an observation o , the agent keeps exactly all clauses consistent with that observation.

Does this agent learn conjunctions of *st*-clauses on-line? Since the agent follows the same strategy as the generalizing agent for propositional conjunctions (Sec-

⁴Recall from logic that constants are a special case of functions, namely functions of arity 0.

tion 2.1), i.e., deleting inconsistent clauses from the current conjunction, starting with the maximal conjunction, Lemma 2.4 applies. That is to say, the agent makes errors only on positive observations. Also, using Theorem 2.5 (in which we replace the number $2n$ of literals to the number $|\Gamma|$ of clauses in the conjunction), we get that the agent makes at most $|\Gamma|$ mistakes. Now the question is whether $|\Gamma|$ is polynomial in the size of the problem, i.e, in n , $|P|$, and $|F|$.

To determine $|\Gamma|$, we first examine the number of different atoms and literals, and then the number $|\Gamma|$ of different st -clauses. Each atom has exactly 1 predicate symbol chosen out of $|P|$ symbols, and at most $t - 1$ other symbols in arguments. Each can be chosen out of $|F|$ function symbols or it can be a variable. An st -clause can have at most st variables. So there are at most $|P|(|F| + st)^{t-1}$ different atoms, and $2|P|(|F| + st)^{t-1}$ literals. A clause combines at most s literals, so the number of all st -clauses can be estimated as

$$\sum_{i=1}^s \binom{2|P|(|F| + st)^{t-1}}{i} \leq \sum_{i=1}^s [2|P|(|F| + st)^{t-1}]^i = \text{poly}(|P|, |F|) \quad (97)$$

and so is polynomial in the size of the learning task. So indeed, the agent learns conjunctions of range-restricted st -clauses on-line.

Does it also learn efficiently, i.e. does it spend at most polynomial time at each hypothesis update step? That depends on whether the relation $o \models \gamma$ in (96) can be determined efficiently. Unfortunately, this is not the case for general st -clauses. However, the test can be made efficient if we adopt yet another restriction on clauses, in particular, the *range-restriction*.

A clause is range-restricted if any variable occurring in a positive literal of it, also occurs in a negative literal of it. So clauses γ_1 and γ_2 from the previous section are examples of range-restricted 3, 3-clauses. The range restriction here follows simply from the fact that the clauses have no positive literals at all.

We recall from elementary first-order logic that $o \models \gamma$ does *not* hold if and only if there is a ground instance⁵ $\gamma\theta$ of γ such that

1. atoms of all negative literals of $\gamma\theta$ are in o , and
2. no positive literal of $\gamma\theta$ is in o

Consider for example a clause $\gamma = \neg\text{edge}(x, y) \vee \neg\text{edge}(y, z) \vee \text{path}(x, z)$, which can be rewritten as

$$\text{edge}(x, y) \wedge \text{edge}(y, z) \rightarrow \text{path}(x, z) \quad (98)$$

⁵that is, a clause obtained from γ by substituting each of its variables to some ground term. A ground term does not contain variables. The substitution θ for which $\gamma\theta$ is ground, is called a *grounding substitution*.

Note that we did not write the quantification $\forall x, y, z$ in the prefix of γ . In first-order clauses, *all* variables are universally quantified, and therefore the quantifier part need not be marked explicitly. We will not indicate quantifiers in clauses further in the text, assuming their implicit universal quantification.

Interpretation $o = \{ \text{edge}(a, b), \text{edge}(b, c) \}$ is not a model of γ ($o \not\models \gamma$). Indeed the following ground instance $\gamma\theta$ of γ

$$\text{edge}(a, b) \wedge \text{edge}(b, c) \rightarrow \text{path}(a, c) \quad (99)$$

where $\theta = \{ x \mapsto a, y \mapsto b, z \mapsto c \}$ has both the atoms corresponding to its negative literals, i.e. $\text{edge}(a, b)$ and $\text{edge}(b, c)$, in o but the positive literal $\text{path}(a, c)$ is not in o . However, the interpretation

$$o = \{ \text{edge}(a, b), \text{edge}(b, c), \text{path}(a, c) \} \quad (100)$$

is a model of γ .

So to determine $o \models \gamma$, the agent first finds all substitutions θ satisfying condition 1, and then checks if each of them satisfies condition 2.

The first stage can be arranged as a tree search. The agent starts with the set $A = \{ a_1, a_2, \dots \}$ of atoms corresponding to γ 's negative literals. At level i of the tree, atom a_i is unified with some element of o in multiple possible ways corresponding to branches leading to level $i + 1$. The unification grounds a subset of variables present in A . When all variables in A have been grounded, the corresponding search node is a successful leaf representing a grounding substitution θ .

We illustrate this with the clause (98) and interpretation (100). The γ 's negative literals $A = \{ \text{edge}(x, y), \text{edge}(y, z) \}$ are unified with o as shown in Fig. 13. In this example, only one grounding substitution exists and is found. We observe that the search tree has, in general, at most s levels and branching factor at most n , so it has at most n^s vertices. The atom in each vertex has at most t arguments so the tree can be searched in at most tn^s time units, which is polynomial in n .

The second stage is easy due to range-restriction. Due to the latter assumption, any substitution θ resulting from the tree search, which makes all negative literals ground, also makes all positive literals ground. Thus checking for each ground positive literal whether or not it is in o (i.e. verifying condition 2) can be done in at most sn unit steps, which is therefore polynomial.

So the agent *efficiently* learns conjunctions of range-restricted st -clauses on-line.

Due to Theorem 3.5, the agent also efficiently PAC-learn this hypothesis class.




Figure 13: By searching this tree, the substitution $\{x \mapsto a, y \mapsto b, z \mapsto c\}$ is found unifying all the negative literals of (98) with elements of (100).

4.2 Generalization of Clauses

So far, observations $o \in O$ the agent received from the environment were ground data encoded through logical interpretations. The latter were simply truth-assignment to propositional variables (which can be perceived as data features), or, in the first-order case, sets of observed ground facts. Now we consider a more interesting situation where observations are akin more to *knowledge* than *data*. Informally, the difference is that knowledge captures patterns from which multiple ground data can be inferred.

A natural way to encode pieces of knowledge in logic is through formulas. In the previous section it was convenient to adhere to clauses as a specific type of formulas, to represent learned knowledge. To represent observations through formulas, we will again stick to clauses.

To exemplify the situation through a simple example, consider that the environment tells the agent two principles through two positive observations ($e_1 = e_2 = 1$)

$$o_1 = \text{male}(x) \wedge \text{female}(y) \wedge \text{parent}(x, y) \rightarrow \text{daughter}(y, x) \quad (101)$$

$$o_2 = \text{female}(x) \wedge \text{parent}(\text{ann}, x) \rightarrow \text{daughter}(x, \text{ann}) \quad (102)$$

Both observations represent a rule that is true in the natural real-world interpretation. But neither of them reflects the entire truth — the rules are not general enough. The agent’s goal is to find a joint generalization of the observations, which in this case would be

$$h^* = \text{female}(x) \wedge \text{parent}(y, x) \rightarrow \text{daughter}(x, y) \quad (103)$$

The fact that h^* is indeed more general than both o_1, o_2 are a logical consequence of h^* , i.e. they are entailed by h^* , which we write as

$$h^* \vdash o_1 \tag{104}$$

$$h^* \vdash o_2 \tag{105}$$

Obviously, there are multiple clauses h satisfying (104),(105). Another option would be e.g.

$$h = \text{parent}(y, x) \rightarrow \text{daughter}(x, y)$$

which in the real-world is an over-generalization. To prevent it, the agent should make sure that the hypothesized clause h does not entail *negative observations*. The environment could e.g. provide the negative example ($e_3 = 0$),

$$o_3 = \text{parent}(\text{jack}, \text{john}) \rightarrow \text{daughter}(\text{john}, \text{jack}) \tag{106}$$

It is indeed entailed by h , indicating the latter is an over-generalization.

There is however another way to prevent over-generalization. Recall that in Sections 2.1 and 4.1 we devised agents that did not need negative observations as they made the smallest possible generalization steps starting from the most specific hypothesis. If the smallest possible generalization had already been over-general, it simply meant that the target concept had no exact match in the agent's hypothesis space, i.e. (48) was not true. Here we will also design an agent making the smallest possible generalization steps.

To build the agent, we would like to adopt a policy based on entailment between clauses such as used in (104,105), which would thus be $\pi(h, o) = 1$ if and only if $h \vdash o$. Unfortunately, \vdash is *undecidable* for general clauses γ, o .

Recall that in Section 2.2 we discussed syntactic subsumption between propositional clauses as an ‘approximation’ of the entailment relation. Syntactic subsumption is also defined for first-order clauses, is decidable, but is more involved than in the propositional case.

In the first-order case, clause γ_1 is said to *θ -subsume* clause γ_2 if there is a substitution θ such that $\gamma_1\theta \subseteq \gamma_2$ where \subseteq is with respect to the sets of literals on either side. This is denoted as

$$\gamma_1 \subseteq_{\theta} \gamma_2$$

For example, $h^* \subseteq_{\theta} o_1$ with $\theta = \{ x \mapsto y, y \mapsto x \}$ as indeed

$$\begin{aligned} & \{ \neg\text{female}(y), \neg\text{parent}(x, y), \text{daughter}(x, y) \} \\ & \quad \subseteq \\ & \{ \neg\text{male}(x), \neg\text{female}(y), \neg\text{parent}(x, y), \text{daughter}(x, y) \} \end{aligned}$$

Similarly, $h^* \subseteq_{\theta} o_2$ with $\theta = \{ y \mapsto \text{ann} \}$.

Two clauses γ_1, γ_2 are said to be *subsume-equivalent*, denoted $\gamma_1 \approx_\theta \gamma_2$ if $\gamma_1 \subseteq_\theta \gamma_2$ and $\gamma_2 \subseteq_\theta \gamma_1$. Clause γ_1 *strictly subsumes* clause γ_2 , written as $\gamma_1 \subset_\theta \gamma_2$ if $\gamma_1 \subseteq_\theta \gamma_2$ but $\gamma_2 \not\subseteq_\theta \gamma_1$.

Lemma 4.1. *The relations $\subseteq_\theta, \approx_\theta, \subset_\theta$ are transitive, i.e. if $\gamma_1 \subseteq_\theta \gamma_2$ and $\gamma_2 \subseteq_\theta \gamma_3$, then $\gamma_1 \subseteq_\theta \gamma_3$ (and analogically for the other two relations).*

As in the propositional setting, subsumption implies entailment, that is to say $\gamma_1 \subseteq_\theta \gamma_2$ implies $\gamma_1 \vdash \gamma_2$, as discussed in Section 2.2. But again, the reverse implication holds only if γ_1 and γ_2 are not self-resolving. An important difference from the propositional setting is that a first-order self-resolving clause need not be a tautology. For instance, consider the non-tautological clause

$$\gamma_1 = \text{natural}(x) \rightarrow \text{natural}(s(x))$$

which can be interpreted to express that a successor of a natural number is also a natural number. Evidently, γ_1 logically entails

$$\gamma_2 = \text{natural}(x) \rightarrow \text{natural}(s(s(x)))$$

That is, $\gamma_1 \vdash \gamma_2$, however, $\gamma_1 \not\subseteq_\theta \gamma_2$ does not hold.

Note that in the above example, none of $\gamma_1, \gamma_2, \gamma$ are self-resolving so the decidable relation \subseteq_θ is indeed equivalent to \vdash for any pair of these clauses.

Clause γ_3 is a *generalization* of clauses γ_1 and γ_2 if $\gamma_3 \subseteq_\theta \gamma_1$ and $\gamma_3 \subseteq_\theta \gamma_2$.

Clause γ_3 is a *least general generalization* of clauses γ_1 and γ_2 if it is their generalization and there is no other generalization γ_4 of the same clauses such that $\gamma_3 \subset_\theta \gamma_4$.

There may be multiple least general generalizations of one pair of clauses, but such generalizations are mutually subsume-equivalent. For example $\text{p}(x)$ and $\text{p}(x) \vee \text{p}(y)$ are both least general generalizations of $\text{p}(a)$ and $\text{p}(b)$, and indeed $\text{p}(x) \approx_\theta \text{p}(x) \vee \text{p}(y)$.

We shall now discuss how to compute a least general generalization of two clauses. The core of the procedure is the *anti-unification* procedure which can be seen as a complement to the unification procedure known from the first-order logic resolution algorithm. Before exposing the algorithm, we need to establish a few notions.

Two *atoms* are *compatible* if they have the same predicate symbol and arity. Two *literals* are *compatible* if they have the same sign, predicate symbol, and arity.

Let a be an atom of arity n . The *position* $[i]$ in a is the i -th argument place in a . The position $[i_1, i_2, \dots, k, l]$ is the l -th argument place in the term occurring

Algorithm 1 Anti-unification of two compatible atoms

Require: Atoms a, b compatible with each other

```

1:  $i := 0; \theta := \emptyset; \sigma := \emptyset$   $\triangleright$  a counter and two substitutions
2:  $v_1, v_2, \dots$ : variables not appearing in  $a$  or  $b$ 
3: while  $a \neq b$  do
4:   Let  $p$  be the leftmost position where  $a$  and  $b$  differ and  $s$  and  $t$  be the terms at this
   position in  $a$  and  $b$ , respectively.
5:   if for some  $j$  ( $1 \leq j \leq i$ ),  $v_j \theta = s$  and  $v_j \sigma = t$  then  $\triangleright$  variable already assigned
6:     put  $v_j$  to position  $p$  in both  $a$  and  $b$   $\triangleright$  replace the terms with that variable
7:   else
8:      $i := i + 1$ 
9:     put  $v_i$  to position  $p$  in both  $a$  and  $b$   $\triangleright$  replace the terms with a new variable
10:     $\theta := \theta \cup \{v_i \mapsto s\}, \sigma := \sigma \cup \{v_i \mapsto t\}$   $\triangleright$  store assignment of  $v_i$ 
11:   end if
12: end while
13: return  $a$ 

```

i	a	θ	b	σ
0	$p(x, f(a, b, g(b, a)), h(a))$	\emptyset	$p(y, f(b, a, g(a, a)), s(a))$	\emptyset
1	$p(v_1, f(a, b, g(b, a)), h(a))$	$\{v_1 \mapsto x\}$	$p(v_1, f(b, a, g(a, a)), s(a))$	$\{v_1 \mapsto y\}$
2	$p(v_1, f(v_2, b, g(b, a)), h(a))$	$\{v_1 \mapsto x,$ $v_2 \mapsto a\}$	$p(v_1, f(v_2, a, g(a, a)), s(a))$	$\{v_1 \mapsto y,$ $v_2 \mapsto b\}$
3	$p(v_1, f(v_2, v_3, g(b, a)), h(a))$	$\{v_1 \mapsto x,$ $v_2 \mapsto a,$ $v_3 \mapsto b\}$	$p(v_1, f(v_2, v_3, g(a, a)), s(a))$	$\{v_1 \mapsto y,$ $v_2 \mapsto b,$ $v_3 \mapsto a\}$
4	$p(v_1, f(v_2, v_3, g(v_3, a)), h(a))$	$\{v_1 \mapsto x,$ $v_2 \mapsto a,$ $v_3 \mapsto b\}$	$p(v_1, f(v_2, v_3, g(v_3, a)), s(a))$	$\{v_1 \mapsto y,$ $v_2 \mapsto b,$ $v_3 \mapsto a\}$
5	$p(v_1, f(v_2, v_3, g(v_3, a)), v_4)$	$\{v_1 \mapsto x,$ $v_2 \mapsto a,$ $v_3 \mapsto b,$ $v_4 \mapsto h(a)\}$	$p(v_1, f(v_2, v_3, g(v_3, a)), v_4)$	$\{v_1 \mapsto y,$ $v_2 \mapsto b,$ $v_3 \mapsto a,$ $v_4 \mapsto s(a)\}$

Table 3: An example of the anti-unification steps as conducted by Alg. 1.

at position $[i_1, i_2, \dots, k]$ in a . For example, the variable x occurs at positions $[1]$ and $[2, 2, 1]$ in atom $a = p(x, f(c, g(x, d)))$. By putting variable y to position $[2, 2]$ of a , we change a into $p(x, f(c, y))$.

Position u is *left of* position v if u precedes v in the lexical order, e.g. $[1, 3, 5]$ is left of $[1, 4]$.

The *anti-unification* of two compatible atoms is an atom produced by Algorithm 1. An example of the steps conducted by the Algorithm is shown in Tab. 3.

Now we shall see how to use the anti-unification algorithm defined for atoms to get a least general generalization defined for clauses. We first define the *selection set*, which picks all pairs of compatible literals from two input clauses.

The *selection set* $\text{Sel}(\gamma_1, \gamma_2)$ of two clauses γ_1, γ_2 is thus defined as

$$\text{Sel}(\gamma_1, \gamma_2) = \{ (l, m) \mid l \in \gamma_1, m \in \gamma_2, l \text{ is compatible with } m \}$$

In the next theorem, clauses are converted into atoms using the selection set and using \bigvee formally as a predicate symbol so that the anti-unification procedure can be applied on such atoms. The resulting atom is then converted back to a clause.

Theorem 4.2. *Let γ_1, γ_2 be clauses and let $\bigvee(a_1, a_2, \dots, a_n)$ be the anti-unification of atoms $\bigvee(l_1, l_2, \dots, l_n)$ and $\bigvee(m_1, m_2, \dots, m_n)$ where*

$$\{(l_1, m_1), (l_2, m_2), \dots, (l_n, m_n)\} = \text{Sel}(\gamma_1, \gamma_2)$$

Then $a_1 \vee a_2 \vee \dots \vee a_n$, denoted $\text{lgg}(\gamma_1, \gamma_2)$, is a least general generalization of γ_1 and γ_2 .

(Proof omitted)

Consider an example using o_1, o_2 from the beginning of this section. Here, the selection set is

$$\text{Sel}(o_1, o_2) = \{(\neg\text{female}(y), \neg\text{female}(x)), \\ (\neg\text{parent}(x, y), \neg\text{parent}(\text{ann}, x)), \\ (\text{daughter}(y, x), \text{daughter}(x, \text{ann}))\}$$

The anti-unification of

$$\bigvee(\neg\text{female}(y), \neg\text{parent}(x, y), \text{daughter}(y, x))$$

and

$$\bigvee(\neg\text{female}(x), \neg\text{parent}(\text{ann}, x), \text{daughter}(x, \text{ann}))$$

is

$$\bigvee(\neg\text{female}(v_1), \neg\text{parent}(v_2, v_1), \text{daughter}(v_1, v_2))$$

So the $\text{lgg}(o_1, o_2)$ is

$$\neg\text{female}(v_1) \vee \neg\text{parent}(v_2, v_1) \vee \text{daughter}(v_1, v_2)$$

which can be transcribed into an equivalent but nicer form

$$\text{female}(x) \wedge \text{parent}(y, x) \rightarrow \text{daughter}(x, y)$$

The exemplified process can be visualized in a matrix form as in Table 4.

We state some basic properties of lgg without a proof.

	$\neg\text{female}(x)$	$\neg\text{parent}(\text{ann}, x)$	$\text{daughter}(x, \text{ann})$	θ	σ	
$\neg\text{male}(x)$						
$\neg\text{female}(y)$	$\neg\text{female}(v_1)$			y	x	v_1
$\neg\text{parent}(x, y)$		$\neg\text{parent}(v_2, v_1)$		x	ann	v_2
$\text{daughter}(y, x)$			$\text{daughter}(v_1, v_2)$			

Table 4: Visualizing the lgg algorithm in a matrix form. The coordinates of the non-empty matrix entries correspond to the selections of compatible literals and the entries represent the anti-unified literals. The table to the right stores the substitutions established by successive anti-unifications of literals.

Lemma 4.3. *Let $\gamma_1, \gamma_2, \gamma_3$ be clauses. Then*

1. *If $\gamma_1 \subseteq_{\theta} \gamma_2$, then $\text{lgg}(\gamma_1, \gamma_2) \approx_{\theta} \gamma_1$*
2. *(commutativity) $\text{lgg}(\gamma_1, \gamma_2) \approx_{\theta} \text{lgg}(\gamma_2, \gamma_1)$*
3. *(associativity) $\text{lgg}(\text{lgg}(\gamma_1, \gamma_2), \gamma_3) = \text{lgg}(\gamma_1, \text{lgg}(\gamma_2, \gamma_3))$*

Property 1 means simply that the least general generalization of a clause with a more general clause is just the more general clause or its equivalent.

Due to the commutative and associative properties, an lgg of a finite set of clauses can be obtained by a repeated application of lgg to arbitrary clause-pairs from the set, always replacing the chosen pair with its lgg, until the set has only one element. Due to the two properties, the order of such lgg applications is irrelevant. Hence we can design an agent for learning from clauses which does not need to collect all clausal observations before jointly generalizing them. In other words, we can design an *on-line learning* agent for clausal observations, and need not resort to the batch-learning setting.

At each time instant k , the agent's hypothesis h_k is single clause. It makes a positive decision if and only if the current observation is subsumed by h_k ,

$$y_k = \pi(h_k, o_k) = \begin{cases} 1 & \text{if } h_k \subseteq_{\theta} o_k \\ 0 & \text{otherwise} \end{cases} \quad (107)$$

So e.g. if h_k is the clause γ from (103) then observations (101) and (102) are decided positively, whereas o_3 (106) or another negative observation such as

$$o_4 = \text{male}(x) \wedge \text{female}(y) \rightarrow \text{daughter}(y, x)$$

would be decided negatively since $h_k \not\subseteq_{\theta} o_3$, $h_k \not\subseteq_{\theta} o_4$.

The agent starts with the initial clause that just copies the first positive observation. For simplicity and without loss of generality we assume that the

first observation o_1 is positive (otherwise we would just let the agent discard observations until the first positive one comes). So

$$h_1 = o_1 \tag{108}$$

The hypothesis update rule is similar to those used by the generalizing agents considered in Sections 2.1 and 4.1, as prescribed by (55)-(56) and (96), respectively. However, the former two agents generalized by deleting literals or clauses (respectively) from their hypothesized conjunction, whereas the present agent will generalize the hypothesized clause using the lgg operator. So, referring to the general update rule (50), we let

$$h_k = \begin{cases} h_{k-1} & \text{if } r_k = 0 \\ \text{lgg}(h_{k-1}, o_{k-1}) & \text{otherwise} \end{cases} \tag{109}$$

For the current agent, we can show that the hypothesized clause is changed only after a making a wrong decision on a positive observation. In other words, the agent—just like both of the generalizing agents we have seen so far—never makes mistakes on negative observations.

To demonstrate this, we adopt again the assumption that there indeed exists a clause h^* , which is correct in the sense that if $h_k = h^*$, then (48) holds.

Now we realize a fact analogical to Lemma 2.3.

Lemma 4.4. *Assuming (108), the clause h_k produced by $\text{lgg}(h_{k-1}, o_{k-1})$ in (109) satisfies for all $k = 1, 2, \dots$*

$$h^* \subseteq_{\theta} h_k \tag{110}$$

The lemma says that the agent never ‘over-generalizes’ by skipping over the target clause in terms of generality. This can be shown by mathematical induction.

Proof. In the inductive step, we assume that for h_{k-1} in (109) it holds $h^* \subseteq_{\theta} h_{k-1}$, and will show that this implies $h^* \subseteq_{\theta} h_k$.

First consider that $e_{k-1} = 0$. If lgg was called in (109), then $r_k \neq 0$, so the previous decision must have been wrong, i.e. $0 = e_{k-1} \neq y_{k-1} = 1$. By (107), this in turn implies that $h_{k-1} \subseteq_{\theta} o_{k-1}$. From Lemma 4.3 (Item 1) we then have that $h_k = \text{lgg}(h_{k-1}, o_{k-1}) \approx_{\theta} h_{k-1}$, in other words h_k and h_{k-1} are subsume-equivalent. But then the assumption $h^* \subseteq_{\theta} h_{k-1}$ implies $h^* \subseteq_{\theta} h_k$.

Next consider that $e_{k-1} = 1$. Considering (107), $h^* \subseteq_{\theta} o_{k-1}$ since h^* is correct. Recall that for induction we assumed $h^* \subseteq_{\theta} h_{k-1}$. Assume now for contradiction

that $h^* \not\subseteq_{\theta} h_k = \text{lgg}(h_{k-1}, o'_{k-1})$. But then lgg would not be a least general generalization, as follows from the definition of the latter (see on page 47) by identifying $h_{k-1}, o_{k-1}, h_k, h^*$ (in this order) with $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ in the definition. But that contradicts Theorem 4.2. So we have again that $h^* \subseteq_{\theta} h_k$.

We have proven the inductive step and it remains to prove the base case, i.e. that $h^* \subseteq_{\theta} h_1$. But this follows immediately from (108), the fact that o_1 is assumed to be positive, and (107). \square

Now we are ready to show a lemma analogical to Lemma 2.4. In particular the current agent makes wrong decisions only for positive observations, and each such error leads to a new hypothesis that is strictly more general than (i.e., strictly subsumes) the previous one.

Lemma 4.5. *Assuming (108), if $r_k \neq 0$, then $e_{k-1} = 1$ and $h_k \subset_{\theta} h_{k-1}$.*

Proof. To see that $e_{k-1} = 1$, assume for contradiction that $e_{k-1} = 0$ (i.e., o_{k-1} was a negative observation). If $r_k \neq 0$ (i.e. y_{k-1} was a wrong decision) then $e_{k-1} \neq y_{k-1} = 1$. Then by (107), $h_{k-1} \subseteq_{\theta} o_{k-1}$. From Lemma 4.4, we also have $h^* \subseteq_{\theta} h_{k-1}$. Due to transitivity of subsumption (Lemma 4.1), this means $h^* \subseteq_{\theta} o_{k-1}$. But then according to (107), o_{k-1} was a positive observation, which is a contradiction.

Since $e_{k-1} = 1$ and $r_k \neq 0$, we have $y_{k-1} = 0$, so by (107), $h_{k-1} \not\subseteq_{\theta} o_{k-1}$. Because $h_k = \text{lgg}(h_{k-1}, o_{k-1})$ is a generalization of both of its arguments, $h_k \subseteq_{\theta} h_{k-1}$ and $h_k \subseteq_{\theta} o_{k-1}$. Given that $h_{k-1} \not\subseteq_{\theta} o_{k-1}$ and $h_k \subseteq_{\theta} o_{k-1}$, it cannot be that $h_{k-1} \approx_{\theta} h_k$. So the subsumption $h_k \subseteq_{\theta} h_{k-1}$ must be strict, i.e. $h_k \subset_{\theta} h_{k-1}$. \square

Lemma 4.4 established that the agent makes a strict generalization upon each mistake yet it never over-generalizes. It was exactly this reasoning that made us able to prove a mistake bound (Theorem 2.5) for the agent in Section 2.1, which also applied to the agent in Section 4.1. For these agents, the maximum number of generalization steps, and correspondingly the mistake bound, was $2n$ and $|\Gamma|$, respectively. These numbers were finite. Unfortunately, we cannot apply the same reasoning for the current agent, as the number of strict generalization steps from the initial clause towards the more general target clause does not have such a general finite bound. Indeed, consider for example the following infinite series of clauses⁶ for $n = 2, 3, \dots$

$$\gamma_n = \bigvee_{1 \leq i, j \leq n, i \neq j} \mathbf{p}(x_i, x_j)$$

⁶Example adopted from Nienhuys-Cheng, de Wolf: Foundation of Inductive Logic Programming, Springer 1998

So e.g.

$$\gamma_2 = \mathbf{p}(x_1, x_2) \vee \mathbf{p}(x_2, x_1)$$

$$\gamma_3 = \mathbf{p}(x_1, x_2) \vee \mathbf{p}(x_2, x_1) \vee \mathbf{p}(x_1, x_3) \vee \mathbf{p}(x_3, x_1) \vee \mathbf{p}(x_2, x_3) \vee \mathbf{p}(x_3, x_2)$$

and so on. We leave it to the reader to verify that $\gamma_2 \subset_{\theta} \gamma_3 \subset_{\theta} \dots$

Now let the target clause h^* be $h^* = \gamma_2$. For any finite number $M \in N$, the environment can present a sequence $o_1 = \gamma_{M+3}, o_2 = \gamma_{M+2}, o_3 = \gamma_{M+1}, \dots, o_{M+1} = \gamma_3$ of $M + 1$ positive observations to the agent, causing it to generalize after each observation, therefore making $M + 1$ mistakes. Thus no finite number M is a mistake bound for the current agent.

The example above follows from the fact that we do not bound the maximum size of clauses included in the lattice of clauses induced by the **lgg** operator. The simple lattice we encountered earlier (Fig. 11) could also contain infinite paths if we alleviated the size bound $2n$ on the conjunctive elements in it. An intricacy distinguishing the subsumption lattice of size-unbounded first-order clauses from the latter lattice is that the infinite path $\gamma_2 \subset_{\theta} \gamma_3 \subset_{\theta} \dots$, which grows in size (number of literals in γ_n), in fact connects two *small* elements in the lattice. More precisely

$$\mathbf{p}(x_1, x_2) = \gamma_2 \subset_{\theta} \gamma_3 \subset_{\theta} \dots \subset_{\theta} \mathbf{p}(x_1, x_1)$$

Again, we leave it to the reader to verify that indeed $\gamma_n \subset_{\theta} \mathbf{p}(x_1, x_1)$ for any $n \geq 2$. So, speaking informally, the clause subsumption lattice is not just infinitely large but also infinitely *dense*.

Finally we will explore how the clausal formalism allows to build an agent that does not start learning from ‘scratch’, i.e. zero initial knowledge, but rather possesses some prior (‘background’) knowledge that just needs to be extended for making correct decisions. We motivate this situation through the following example where the environment provides two positive observations

$$o_1 = \mathbf{female}(x) \wedge \mathbf{father}(y, x) \rightarrow \mathbf{daughter}(x, y) \quad (111)$$

$$o_2 = \mathbf{female}(x) \wedge \mathbf{mother}(y, x) \rightarrow \mathbf{daughter}(x, y) \quad (112)$$

Because $\mathbf{father}(y, x)$ and $\mathbf{mother}(y, x)$ are not mutually compatible, the present agent would generalize o_1, o_2 using **lgg** into

$$\mathbf{female}(x) \rightarrow \mathbf{daughter}(x, y)$$

which is clearly unsatisfactory. Consider, however, that the agent has background knowledge in the form of a set of clauses (i.e., a *clausal theory*) B :

$$\mathbf{father}(x, y) \rightarrow \mathbf{parent}(x, y) \quad (113)$$

$$\mathbf{mother}(x, y) \rightarrow \mathbf{parent}(x, y) \quad (114)$$

Knowing B , the agent should be able to generalize o_1, o_2 into (103). More precisely, the agent starts with background knowledge only $h_1 = B$ and should extend this knowledge towards the hypothesis h_k by learning the additional clause h^* (103) such that

$$h_k = B \cup \{h^*\} \vdash o_k \quad (115)$$

for all *positive* examples o_k (i.e., for $\forall k$ such that $e_k = 1$) and

$$h_k = B \cup \{h^*\} \not\vdash o_k \quad (116)$$

for all *negative* examples o_k (i.e., for $\forall k$ such that $e_k = 0$).

Once again, the consequence relation \vdash in Equations (115) - (116) is generally undecidable and we thus need to reformulate the task using the approximation by θ -subsumption. To do so, note that if $\gamma_1 \subseteq_\theta \gamma_2$ then the formula $\gamma_1\theta \rightarrow \gamma_2$ is a tautology, i.e.

$$\vdash \gamma_1\theta \rightarrow \gamma_2 \quad (117)$$

For example $\mathbf{p}(x) \subseteq_\theta \mathbf{p}(a) \vee \mathbf{q}(y)$ so $\vdash \mathbf{p}(a) \rightarrow (\forall y : \mathbf{p}(a) \vee \mathbf{q}(y))$. We shall account for background knowledge B by making (117) relative to it, i.e.

$$B \vdash \gamma_1\theta \rightarrow \gamma_2 \quad (118)$$

So here $\gamma_1\theta \rightarrow \gamma_2$ is a tautological consequence of (is entailed by) B rather than being a tautology.

This leads us to the following definition. We say that clause γ_1 *θ -subsumes* clause γ_2 *relative to* clause set B , written $\gamma_1 \subseteq_\theta^B \gamma_2$ if there is a substitution θ such that (118) holds. Clauses γ_1, γ_2 are *subsume-equivalent relative to B* if $\gamma_1 \subseteq_\theta^B \gamma_2$ and $\gamma_2 \subseteq_\theta^B \gamma_1$; we denote this as $\gamma_1 \approx_\theta^B \gamma_2$.

Furthermore, we define *relative least general generalization* (with respect to B) just as least general generalization (page 47), except that we replace the relation \subseteq_θ with \subseteq_θ^B in the definition.

We wanted to base the notion of relative generalization on θ -subsumption to avoid the undecidable relation \vdash . It seems we have not succeeded in so doing as (118) still contains the latter relation. The good news is that according to the following theorem, in the special case that B contains only ground facts, relative subsumption can be easily converted to plain subsumption.

Theorem 4.6. *Let γ_1, γ_2 be clauses and B a finite set of ground facts. Then $\text{rlgg}_B(\gamma_1, \gamma_2)$ is a relative least general generalization of γ_1, γ_2 with respect to B , where*

$$\text{rlgg}_B(\gamma_1, \gamma_2) = \text{lbg}(\gamma_1 \vee_{l \in B} \neg l, \gamma_2 \vee_{l \in B} \neg l)$$

(Proof omitted)

	$d(c, b)$	$\neg m(a)$	$\neg p(a, b)$	$\neg f(b)$	$\neg p(b, c)$	$\neg f(c)$
$d(b, a)$	$d(v_1, v_2)$					
$\neg m(a)$		$\neg m(a)$				
$\neg p(a, b)$			$\neg p(a, b)$		$\neg p(v_2, v_1)$	
$\neg f(b)$				$\neg f(b)$		$\neg f(v_1)$
$\neg p(b, c)$			$\neg p(v_3, v_4)$		$\neg p(b, c)$	
$\neg f(c)$				$\neg f(v_4)$		$\neg f(c)$

Table 5: Computation of the lgg of clauses (122,123) visualized as in Table 4.

θ	σ	new variable
b	c	v_1
a	b	v_2
b	a	v_3
c	b	v_4

Table 6: Assignment of new variables to terms in clause (122) (θ column) and terms in clause (123) (σ column) to new variables during the lgg computation as shown in Table 5.

Note that this theorem excludes, for example, the clauses (113, 114) we used as background knowledge in the motivating example as they are not ground facts. Instead, we shall exemplify the theorem with simpler background knowledge describing ground family relationships for persons we identify for brevity with constants **a**, **b**, **c**, using predicates $m/1$, $f/1$, $p/1$ with the informal meaning *male*, *female*, and *parent of* (respectively).

$$B = \{ m(a), p(a, b), f(b), p(b, c), f(c) \} \quad (119)$$

The agent should again learn the definition of the daughter relation. Equipped with background knowledge B , it receives the first positive observation of the daughter relation expressed through predicate $d/2$

$$o_1 = d(b, a) \quad (120)$$

and thus forms its first hypothesized clause $h_1 = o_1 = d(b, a)$. Once the second positive observation

$$o_2 = d(c, b) \quad (121)$$

has been received, the agent should update its hypothesis as $h_2 = r\text{lgg}_B(h_1, o_2)$. According to Theorem 4.6, this can be computed as the lgg of the clauses

$$d(b, a) \vee \neg m(a) \vee \neg p(a, b) \vee \neg f(b) \vee \neg p(b, c) \vee \neg f(c) \quad (122)$$

$$d(c, b) \vee \neg m(a) \vee \neg p(a, b) \vee \neg f(b) \vee \neg p(b, c) \vee \neg f(c) \quad (123)$$

The lgg computation is shown in Table 5 in the same way as was done in Table 4. Table 6 shows the two substitutions created in the process. The resulting clause is

$$h_2 = \mathbf{d}(v_1, v_2) \leftarrow \mathbf{m}(\mathbf{a}) \wedge \mathbf{p}(\mathbf{a}, \mathbf{b}) \wedge \mathbf{p}(v_2, v_1) \wedge \mathbf{f}(\mathbf{b}) \wedge \mathbf{f}(v_1) \wedge \mathbf{p}(v_3, v_4) \wedge \mathbf{p}(\mathbf{b}, \mathbf{c}) \wedge \mathbf{f}(v_4) \wedge \mathbf{f}(\mathbf{c})$$

While this is a correct relative least general generalization, it is evidently redundant. Informally, the ground facts appearing in the body of the clause (the conjunctive part to the right of \leftarrow) can be deleted as they are true due to B . Also, the body literals whose all variables do not appear in the head of the clause (to the left of \leftarrow) are redundant. Speaking precisely, h_2 is subsume-equivalent (relative to B) to the clause

$$h'_2 = \mathbf{d}(v_1, v_2) \leftarrow \mathbf{p}(v_2, v_1) \wedge \mathbf{f}(v_1) \tag{124}$$

which indeed represents the desired learned hypothesis. In formal notation, $h_2 \approx_{\theta}^B h'_2$. To prove this equivalence relation, we need to prove $h_2 \subseteq_{\theta}^B h'_2$ and $h'_2 \subseteq_{\theta}^B h_2$.

The latter relation, which by definition (page 54) transcribes into

$$B \vdash (h'_2 \theta \rightarrow h_2)$$

is evident, because $h'_2 \subset h_2$, implying $h'_2 \vdash h_2$. So $h'_2 \theta \rightarrow h_2$ is a tautology for $\theta = \{\}$, meaning it is true in any model, not just in any model of B .

To demonstrate $h_2 \subseteq_{\theta}^B h'_2$, i.e.

$$B \vdash (h_2 \theta \rightarrow h'_2)$$

we set $\theta = \{v_3 \mapsto v_2, v_4 \mapsto v_1\}$. Then

$$h_2 \theta = \mathbf{d}(v_1, v_2) \leftarrow \mathbf{m}(\mathbf{a}) \wedge \mathbf{p}(\mathbf{a}, \mathbf{b}) \wedge \mathbf{p}(v_2, v_1) \wedge \mathbf{f}(\mathbf{b}) \wedge \mathbf{f}(v_1) \wedge \mathbf{p}(\mathbf{b}, \mathbf{c}) \wedge \mathbf{f}(\mathbf{c})$$

In any model of B , all of the ground literals of $h_2 \theta$ are true so by deleting them, we get a clause logically equivalent to $h_2 \theta$. But such a clause is exactly h'_2 so $h_2 \theta \rightarrow h'_2$ becomes $h'_2 \rightarrow h'_2$ which is satisfied trivially.

The conversion of h_2 into h'_2 as shown above is an example of *clause reduction*.

We say that a clause γ is *reduced* if for no clause γ' , $\gamma' \subset \gamma$, $\gamma' \approx_{\theta} \gamma$. A reduced clause γ' is a *reduction* of a clause γ if $\gamma' \approx_{\theta} \gamma$.

Similarly, a clause γ is *reduced with respect to B* if for no clause γ' , $\gamma' \subset \gamma$, $\gamma' \approx_{\theta}^B \gamma$. A clause γ' which is reduced with respect to B is a *reduction of γ with respect to B* if $\gamma' \approx_{\theta}^B \gamma$.