Question 1. (10 points)

Let h, h' be two propositional clauses or conjunctions. Show that $\mathsf{lgg}(h, h') = \mathsf{Lits}(h) \cap \mathsf{Lits}(h')$ is a least general generalization of h, h'.

Answer:

For shorter notation, we apply set operations on clauses or conjunctions as if they were sets of literals, so for example $\mathcal{L} \in h$ means $\mathcal{L} \in \mathsf{Lits}(h)$. The question assumes that $g = \mathsf{lgg}(h, h') = h \cap h'$, i.e.,

$$\forall \mathcal{L} : \mathcal{L} \in g \leftrightarrow \mathcal{L} \in h \wedge \mathcal{L} \in h' \tag{1}$$

from which we deduce

$$\vdash \forall \mathcal{L} : \mathcal{L} \in g \rightarrow \mathcal{L} \in h \wedge \mathcal{L} \in h'$$
$$\vdash \forall \mathcal{L} : \mathcal{L} \in g \rightarrow \mathcal{L} \in h \text{ and } \vdash \forall \mathcal{L} : \mathcal{L} \in g \rightarrow \mathcal{L} \in h' \tag{2}$$

The two formulas we deduced in (2) define the respective subset relations $g \subseteq h$ and $g \subseteq h'$ which we were to prove according to the definition of least general generalization. In addition, we are supposed to prove that there is no g' such that $g \subset g'$, $g' \subseteq h$, $g' \subseteq h'$. Assume for contradiction, that such a g' exists, i.e.,

$$\forall \mathcal{L} : \mathcal{L} \in g \rightarrow \mathcal{L} \in g' \tag{3}$$
$$\exists \mathcal{L} : \mathcal{L} \in g' \wedge \neg(\mathcal{L} \in g) \tag{4}$$
$$\forall \mathcal{L} : \mathcal{L} \in g' \rightarrow \mathcal{L} \in h \tag{5}$$
$$\forall \mathcal{L} : \mathcal{L} \in g' \rightarrow \mathcal{L} \in h' \tag{6}$$

We rewrite formulas 3 - 6 as clauses, *Skolemizing* formula 4.

$$\mathcal{L} \in g \rightarrow \mathcal{L} \in g' \tag{7}$$
$$\mathsf{a} \in g' \tag{8}$$
$$\neg(\mathsf{a} \in g) \tag{9}$$
$$\mathcal{L} \in g' \rightarrow \mathcal{L} \in h \tag{10}$$
$$\mathcal{L} \in g' \rightarrow \mathcal{L} \in h' \tag{11}$$

We also rewrite assumption (1) as a clause set

$$\mathcal{L} \in g \rightarrow \mathcal{L} \in h \tag{12}$$
$$\mathcal{L} \in g \rightarrow \mathcal{L} \in h' \tag{13}$$
$$\mathcal{L} \in h \wedge \mathcal{L} \in h' \rightarrow \mathcal{L} \in g \tag{14}$$

We are supposed to find a contradiction in the clause set 7 - 14. We proceed by resolution:

- Resolve (10) ($\mathcal{L} \mapsto \mathsf{a}$) with (8) obtaining

$$\mathsf{a} \in h \tag{15}$$

- Resolve (11) ($\mathcal{L} \mapsto \mathsf{a}$) with (8) obtaining

$$\mathsf{a} \in h' \tag{16}$$

- Resolve (14) ($\mathcal{L} \mapsto \mathsf{a}$) with (15) obtaining

$$\mathsf{a} \in h' \rightarrow \mathsf{a} \in g \tag{17}$$

- Resolve (17) with (16) obtaining

$$\mathsf{a} \in g \tag{18}$$

- Resolve (18) with (9) obtaining the contradiction $\square$.

Question 2. (15 points)

Determine if

1. $h \subseteq_\theta h'$
2. $h' \models h$

for

$$h = \mathrm{p}(x, y) \wedge \mathrm{p}(y, z) \wedge \neg\mathrm{p}(x, z)$$
$$h' = \mathrm{p}(\mathsf{a}, \mathsf{b}) \wedge \mathrm{p}(\mathsf{b}, \mathsf{c}) \wedge \mathrm{p}(\mathsf{c}, \mathsf{d}) \wedge \neg\mathrm{p}(\mathsf{a}, \mathsf{d})$$

Answer:

1. $h \subseteq_\theta h'$ does not hold because for no substitution θ, $\mathsf{Lits}(h\theta) \subseteq \mathsf{Lits}(h')$. Indeed, for $\neg\mathrm{p}(x, z)\theta$ to be in $\mathsf{Lits}(h')$, necessarily $\{ x \mapsto \mathsf{a}, z \mapsto \mathsf{d} \} \subseteq \theta$ as $\neg\mathrm{p}(\mathsf{a}, \mathsf{d})$ is the only negative literal in h'. In turn, $\mathrm{p}(x, y)\theta = \mathrm{p}(\mathsf{a}, y)$ can be in $\mathsf{Lits}(h')$ only if $(y \mapsto \mathsf{b}) \in \theta$. But then, $\mathrm{p}(y, z)\theta = \mathrm{p}(\mathsf{b}, \mathsf{d})$ is not in $\mathsf{Lits}(h')$.

2. $h' \models h$ is true. This is due to the soundness theorem and the fact that $h' \vdash h$. To prove the latter, we prove the equivalent relation $h' \wedge \neg h \vdash \Box$, i.e., iff the clause set

$$\{ \neg h \} \cup \mathsf{Lits}(h') = \left\{ \begin{array}{l} \neg\mathrm{p}(x, y) \vee \neg\mathrm{p}(y, z) \vee \mathrm{p}(x, z) \\ \mathrm{p}(\mathsf{a}, \mathsf{b}) \\ \mathrm{p}(\mathsf{b}, \mathsf{c}) \\ \mathrm{p}(\mathsf{c}, \mathsf{d}) \\ \neg\mathrm{p}(\mathsf{a}, \mathsf{d}) \end{array} \right\}$$

   is contradictory. We show the contradiction as follows

   (a) Resolve $\neg h\theta$, $\theta = \{ x \mapsto \mathsf{a}, z \mapsto \mathsf{d} \}$ with $\neg\mathrm{p}(\mathsf{a}, \mathsf{d})$ obtaining the resolvent clause

   $$c_1 = \neg\mathrm{p}(\mathsf{a}, y) \vee \neg\mathrm{p}(y, \mathsf{d})$$

   (b) Resolve $c_1\theta$, $\theta = \{ y \mapsto \mathsf{b} \}$ with $\mathrm{p}(\mathsf{a}, \mathsf{b})$ obtaining

   $$c_2 = \neg\mathrm{p}(\mathsf{b}, \mathsf{d})$$

   (c) Resolve $\neg h\theta$, $\theta = \{ x \mapsto \mathsf{b}, z \mapsto \mathsf{d} \}$ with c_2 obtaining

   $$c_3 = \neg\mathrm{p}(\mathsf{b}, y) \vee \neg\mathrm{p}(y, \mathsf{d})$$

   (d) Resolve $c_3\theta$, $\theta = \{ y \mapsto \mathsf{c} \}$ with $\mathrm{p}(\mathsf{b}, \mathsf{c})$ obtaining

   $$c_4 = \neg\mathrm{p}(\mathsf{c}, \mathsf{d})$$

   (e) Finally, resolve c_4 with $\mathrm{p}(\mathsf{c}, \mathsf{d})$ obtaining the empty clause $\Box$.

Note that there is a simple intuition behind the formal proof: $\neg h$ can be written also as $\mathrm{p}(x, y) \wedge \mathrm{p}(y, z) \rightarrow \mathrm{p}(x, z)$ expressing the transitivity property of $\mathrm{p}/2$. That property is in turn in obvious contradiction with all of $\mathrm{p}(\mathsf{a}, \mathsf{b}), \mathrm{p}(\mathsf{b}, \mathsf{c}), \mathrm{p}(\mathsf{c}, \mathsf{d})$ being true while $\mathrm{p}(\mathsf{a}, \mathsf{d})$ being false.

Question 3. (5 points)

Consider the following statements

1. X = non-self-resolving FOL clauses
2. X = contingent FOL clauses
3. There is no $k \in \mathbb{N}$, $x \in X$ such that $h_k \models x$ and $h_k \not\subseteq_\theta x$, where h_k ($k \in \mathbb{N}$) are the hypotheses of the generalization algorithm.

Decide for each of the implications $1 \to 2$, $1 \to 3$, $2 \to 3$, whether it is true. Change the relation $h_k \models x$ in (3) so that all the implications you decided true are true when (1) and (2) assume conjunctions instead of clauses.

Answer:

1. $1 \to 2$ is not true because an empty clause is not self-resolving but it is not contingent. The implication is however true for *non-empty* clauses. Assume for contradiction that a non-empty $x \in X$ is not self-resolving but not contingent, i.e. it is tautologically false or true. First assume the former case, i.e. $\models \neg x$. This means that for any formula x', it holds $x' \models \neg x$, i.e. $x \models \neg x'$. Since x is not self-resolving, this means $x \subseteq_\theta \neg x'$ for any clause $\neg x'$. This can only happen if $\mathsf{Lits}(x)$ is empty but then x is an empty clause, which contradicts the assumption. The case for $\models x$ is analogical (use the fact that the negation of a non-self-resolving clause is a non-self-resolving conjunction).
2. $1 \to 3$ is true. $h_k \models x$ and $h_k \not\subseteq_\theta x$ can happen simultaneously only if h_k is self-resolving and we will prove the implication by showing that for all $k \in \mathbb{N}$, h_k is not self-resolving.

   By the generalization algorithm, h_1 is the first positive example, which is from X so it is not self-resolving by the assumption of the implication. Assume for contradiction that for some $k \in \mathbb{N}$,

   $$h_{k+1} = \mathsf{lgg}(h_k, x),\ \ x \in X \tag{19}$$

   is self-resolving, i.e., contains both a positive literal and a negative literal with the same predicate $\mathcal{P}$. By definition of the lgg operator, $\mathsf{lgg}(h_k, x)$ contains a predicate only if it occurs in both of h_k, x with the same sign. Thus x must also contain both a positive literal and a negative literal with the predicate $\mathcal{P}$. Thus x is self-resolving, which contradicts the assumption.
3. $2 \to 3$ is not true. Let the first positive example and thus also h_1 be $\mathrm{p}(x) \to \mathrm{p}(\mathsf{f}(x))$, i.e., $\neg\mathrm{p}(x) \vee \mathrm{p}(\mathsf{f}(x))$. Clearly, it is contingent: it is false e.g. with Herbrand interpretation $\{\, \mathrm{p}(\mathsf{a}) \,\}$ and true with Herbrand the (infinite) interpretation $\{\, \mathrm{p}(\mathsf{a}), \mathrm{p}(\mathsf{f}(\mathsf{a})), \mathrm{p}(\mathsf{f}(\mathsf{f}(\mathsf{a}))), \dots \,\}$. But it is also self-resolving because it contains the predicate p/1 in both a negative and positive literal. Thus it can happen that $h_1 \models x$ and $h_1 \not\subseteq_\theta x$. Indeed, this happens e.g. with the contingent observation $x = \mathrm{p}(x) \to \mathrm{p}(\mathsf{f}(\mathsf{f}(x)))$.

   However, the implication is true in the special case of *propositional* clauses because a propositional self-resolving clause is necessarily non-contingent.

To make $1 \to 3$ true for conjunctions, change $h_k \models x$ to $x \models h_k$ in item 2 of the question and in item 2 of the answer.

Question 4. (1 points)

Find two different least general generalizations of $\mathrm{p}(\mathsf{a})$ and $\mathrm{p}(\mathsf{b}) \vee \mathrm{p}(\mathsf{c})$, prove that they are indeed generalizations of the two clauses and prove that they are mutually θ-equivalent. Explain why two least general generalizations of the same pair of clauses or conjunctions must be θ-equivalent.

Answer:

$\mathrm{p}(x)$ and $\mathrm{p}(y) \vee \mathrm{p}(z)$

$\mathrm{p}(x)\theta \subseteq \mathrm{p}(y) \vee \mathrm{p}(z)$ for $\theta = \{\, x \mapsto y \,\}$

$(\mathrm{p}(y) \vee \mathrm{p}(z))\,\theta \subseteq \mathrm{p}(x)$ for $\theta = \{\, y \mapsto x, z \mapsto x \,\}$

So unlike in the propositional case, a least general generalization of FOL clauses (or conjunctions) is not unique!

Consider for contradiction two least general generalizations g, g' of the same pair of clauses or conjunctions such that $g \subset_\theta g'$. Then by the definition of least general generalization, g is not a least general generalization.

Question 5. (10 points)

Explain why the proof of the mistake bound n of the generalization algorithm is no longer valid when the assumption on X is changed to $X =$ non-self-resolving FOL conjunctions or non-self-resolving FOL conjunctions clauses, the relations $\subseteq, \subset$ are changed to $\subseteq_\theta, \subset_\theta$ (respectively), and we set $n = |\mathcal{P}|$. Show that the proof cannot be rectified, in particular that no finite mistake bound exists under said assumption even if $\mathcal{F} = \emptyset$.

Answer:

The proof uses the following argument

> *Since examples are contingent, they have at most n literals (each of the n atoms is included either as a positive or negative literal but not both) and since h_1 is the first positive example, it also has at most n literals*

which is only true for propositional clauses or conjunctions. The number of different atoms using $n = |\mathcal{P}|$ predicates may be larger than n since there may be multiple atoms with the same predicate, each one using different terms as arguments.

Also the inference

> *at least one literal is removed on each mistake, so the* maximum number of mistakes is n

is not true in FOL. While a strict generalization ($h_{k+1} \subset_\theta h_k$) is still made after each mistake, this does not mean that $|\mathsf{Lits}(h_{k+1})| < |\mathsf{Lits}(h_k)|$ (consider e.g. $h_{k+1} = \mathrm{p}(x), h_k = \mathrm{p}(\mathsf{a})$.)

We can even show that no mistake bound can be provided under the assumption above. Let us assume for contradiction that the bound is M.

Let the target hypothesis be $\tilde{h} = \mathrm{p}(z_1, z_2)$ and let the environment present the following positive examples x_k for $1 \leq k \leq M$

$$x_k = \bigvee_{1 \leq i,j \leq M-k+2, i \neq j} \mathrm{p}(z_i, z_j) \tag{20}$$

and

$$x_{M+1} = \mathrm{p}(z_1, z_2) \tag{21}$$

so e.g. for $M = 2$

$$x_1 = \mathrm{p}(z_1, z_2) \vee \mathrm{p}(z_2, z_1) \vee \mathrm{p}(z_1, z_3) \vee \mathrm{p}(z_3, z_1) \vee \mathrm{p}(z_2, z_3) \vee \mathrm{p}(z_3, z_2)$$
$$x_2 = \mathrm{p}(z_1, z_2) \vee \mathrm{p}(z_2, z_1)$$
$$x_3 = \mathrm{p}(z_1, z_2)$$

and since $x_3 \subset_\theta x_2 \subset_\theta x_1$ and the agent makes *least* generalizations, it will make $3 > M$ mistakes. In general, for any $M \in \mathbb{N}$, the sequence (20)-(21) is a chain of $M + 1$ strict generalizations forcing the agent to make $M + 1$ mistakes.

Note that for any $M \in \mathbb{N}$ and $k \in [1; M+1]$, $x_k \subseteq_\theta \mathrm{p}(z, z)$. Therefore, there is an infinite *chain of generalizations between two* finite *elements $\mathrm{p}(z_1, z_2)$ and $\mathrm{p}(z, z)$ in the θ-subsumption lattice!*