Learning from Answer Sets via Single-Shot Disjunctive ASP Encoding: Technical Appendix

Roberto Borelli[1,2], **Agostino Dovier**[1]

[1]University of Udine, Dept of Mathematics, Computer Science, and Physics
[2]University of Padua, Dept of Mathematics
roberto.borelli@phd.unipd.it, agostino.dovier@uniud.it

This Technical Appendix is organized as follows. Appendix A illustrates the single-shot encoding on a small toy example. Appendix B presents the proofs omitted from Section 4. In particular, we provide detailed proofs of the correctness and completeness of the proposed approach. Finally, Appendix C contains the proofs omitted from Section 5, which concern the grounding process.

A Running Example

In this section, we provide a complete encoding of a small task. We begin by introducing the task along with its properties, followed by a presentation of both the cautious and brave encodings. Finally, we conclude with a discussion of the full encoding.

Task definition

Let $T_e = \langle B, S, E^+, E^- \rangle$ be the ILP_{LAS} task such that:

- $B := \{$`p :- not q.`$\}$
- $S := \{h_1$ `q.` h_2 `q :- not p.`$\}$
- $E^+ := \{\langle \{p.\}, \{q.\} \rangle\}$
- $E^- := \{\langle \{q.\}, \{p.\} \rangle\}$

There are 4 possible subsets of the hypothesis space. For each subset H, we give the set of answer sets of $B \cup H$.

1. $AS(B \cup \emptyset) = \{1 : \{p\}\}$
2. $AS(B \cup \{h_1\}) = \{1 : \{q\}\}$
3. $AS(B \cup \{h_2\}) = \{1 : \{p\},\ 2 : \{q\}\}$
4. $AS(B \cup \{h_1, h_2\}) = \{1 : \{q\}\}$

The hypotheses $H = \{h_1\}$ and $H = \{h_1, h_2\}$ are not inductive solutions, as the only answer set of $B \cup H$ extends the negative example. Moreover, $H = \{h_2\}$ is also not an inductive solution: while the answer set containing only the atom `p` does not extend the negative example, the answer set containing only the atom `q` does, contradicting the definition of an inductive solution. Thus, we conclude that $H = \emptyset$ is the only inductive solution of T_e.

Cautious encoding

We now present the program $P^-_{dis}(T_e)$ for the task T_e, explicitly illustrating the intermediate steps leading to its construction. The program $B \cup S$ is tight, so the formula $\phi^*_{as}(P'(T_e))$ coincides with the completion formula $\phi^*_{iff}(P'(T_e))$:

$$\phi^*_{as}(P'(T_e)) := (p \leftrightarrow \neg q) \wedge (q \leftrightarrow ((\neg p \wedge h(2)) \vee h(1))).$$

The set of free variables of $\phi(T_e)$ is $V_{P'(T_e)} = V_{P(T_e)} \cup V^H_{P'(T_e)} = \{p, q\} \cup \{h(1), h(2)\}$.

In the task T_e, the only negative example is $\langle \{$`q.`$\}, \{$`p.`$\} \rangle$, hence:

$$\phi_{neg}(E^-) := (\neg q \vee p).$$

The formula $\psi(T_e)$, obtained by transforming into NNF the formula $\phi^*_{as}(P'(T_e)) \to \phi_{neg}(E^-)$, is:

$$((((p \wedge q) \vee (\neg q \wedge \neg p)) \vee ((q \wedge (\neg h(1) \wedge (\neg h(2) \vee p)))\vee ((h(1) \vee (h(2) \wedge \neg p)) \wedge \neg q))) \vee (\neg q \vee p))$$

To better visualize the structure of $\psi(T_e)$ we show the syntax tree. Each internal node is annotated with a unique index:

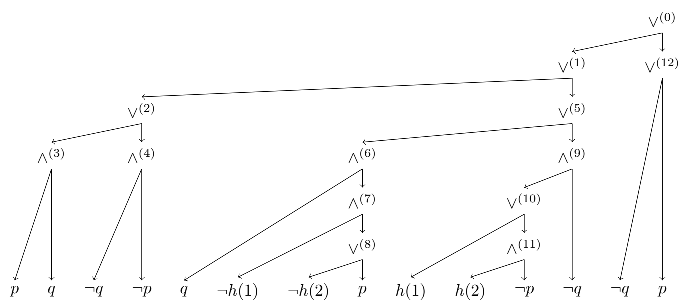

The final formula from Stage 1 is:

$$\phi(T_e) := \exists h(1)\ \exists h(2)\ \forall q\ \forall p\ \ \psi(T_e).$$

We now proceed with the encoding of $\phi(T_e)$ into the disjunctive ASP program $P^-_{dis}(T_e)$:

```
%% Existential variables
h(1) | nh(1). h(2) | nh(2).

%% Universal variables
p | np. q | nq.
p  :- w. np  :- w. q  :- w. nq  :- w.

%% Formuala encoding
w :- formula_0.
formula_0 :- formula_1.
formula_0 :- formula_12.
formula_1 :- formula_2.
formula_1 :- formula_5.
formula_2 :- formula_3.
formula_2 :- formula_4.
```

```
formula3 :- p,q.
formula4 :- nq,np.
formula5 :- formula6.
formula5 :- formula9.
formula6 :- q,formula7.
formula7 :- nh(1),formula8.
formula8 :- nh(2).
formula8 :- p.
formula9 :- formula10,nq.
formula10 :- h(1).
formula10 :- formula11.
formula11 :- h(2),np.
formula12 :- nq.
formula12 :- p.
%% Saturation predicate
w  :- not w.
```

To briefly explain the program *expansion*$(\psi(T_e))$, consider that the atom $\texttt{formula}_0$ represents the root of the syntax tree of $\psi(T_e)$. Since $\psi(T_e)$ is a disjunction, the program includes the rules

```
formula0 :- formula1.  formula0 :- formula12.
```

The atom $\texttt{formula}_{12}$ represents the disjunction $\neg q \vee p$, so it is encoded as:

```
formula12 :- nq.  formula12 :- p.
```

The only answer set of the program $P^-_{dis}(T_e)$ is:
`p np q nq nh(1) nh(2) w formula0 formula1 formula2 formula3 formula4 formula5 formula6 formula7 formula8 formula12`
This correctly indicates that every answer set of $B \cup \emptyset$ does not extend the negative example, satisfying the cautious entailment requirement.

Brave encoding

After presenting the program $P^-_{dis}(T_e)$ in the last paragraph, we now see the program $P^+_{dis}(T_e)$.

```
%% Background knowledge B = {p :- not q}
inAs(p,M) :- not inAs(q,M),
interpretation(M).
%% Hypothesis S={h1: q, h2: q :- not p}
inAs(q,M) :- interpretation(M),h(1).
inAs(q,M) :- not inAs(p,M),
interpretation(M),h(2).
%% Positive example ({p.}, {q.})
interpretation(1).
cov(1) :- inAs(p,1),not inAs(q,1).
 :- not cov(1).
```

The program $P^+_{dis}(T_e)$ is not self-contained since atoms `h(1)` and `h(2)` do not occur in any rule head. We thus consider the program $P^{+'}_{dis}(T_e) := P^+_{dis}(T_e) \cup$ {`h(1) | nh(1). h(2) | nh(2).`}. The program $P^{+'}_{dis}(T_e)$ has two answer sets:

A_1: `interpretation(1) nh(1) nh(2) inAs(p,1) cov(1)`

A_2: `interpretation(1) nh(1) h(2) inAs(p,1) cov(1)`

This correctly indicates that:

- there exists an answer set of $B \cup \emptyset$ which extends the positive example ;
- there exists an answer set of $B \cup \{h_2\}$ which extends the positive example.

In fact, as seen in the presentation of the task T_e, it holds that $AS(B \cup \emptyset) = \{\{p\}\}$ and $AS(B \cup \{h_2\}) = \{\{p\}, \{q\}\}$.

Overall encoding

The only answer set A_1 of the full program $P_{dis}(T_e) = P^-_{dis}(T_e) \cup P^+_{dis}(T_e)$ is:
`interpretation(1) nh(1) inAs(p,1) cov(1) p np q nq nh(2) w formula0 formula1 formula2 formula3 formula4 formula5 formula6 formula7 formula8 formula12`
This indicates that the only inductive solution of T_e is $H = \{h_i \mid \texttt{h(i)} \in A_1\} = \emptyset$. This is the intersection between the set of solutions found by $P^-_{dis}(T_e)$ and $P^{+'}_{dis}(T_e)$ In particular, the set of inductive solutions is the set $\{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in AS(P^-_{dis}(T_e))\} \cap \left\{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in AS(P^{+'}_{dis}(T_e))\right\}$ which evaluates to $\{\emptyset\} \cap \{\emptyset, h_2\} = \{\emptyset\}$.

B Proofs from Section 4

We now provide proofs regarding correctness, completeness and complexity for the disjunctive encoding $P_{dis}(T)$. We begin by analyzing properties of the cautious entailment module $P^-_{dis}(T)$.

First we formally state that the formula $\phi_{neg}(E^-)$ ensures that, for each negative example, at least one inclusion atom is absent or one exclusion atom is present - i.e., an interpretation satisfying $\phi_{neg}(E^-)$ fails to extend every example in E^-.

Fact B.1 *Let A be an interpretation of $B \cup S$. Then A satisfies $\phi_{neg}(E^-)$ if and only if for every $e \in E^-$, A does not extend e.*

Proof. We have that $A \models \phi_{neg}(E^-)$ if and only if the formula $\bigvee_{a \in e^{incl}} \neg a \vee \bigvee_{a \in e^{excl}} a$ is satisfied for every $e = \langle e^{incl}, e^{excl} \rangle \in E^-$. Let $a \in e^{incl}$; if $A \models \neg a$, then it does not hold that $e^{incl} \subseteq A$. Let $a \in e^{excl}$; if $A \models a$, then it does not hold that $e^{excl} \cap A = \emptyset$. In either case, we conclude that A does not extend e. $\square$

Lemma 4.1, follows directly from Theorem 2.2 and the definition of the program $P'(T)$ and establishes the correspondence between models of $\phi^*_{as}(P'(T))$ and the answer sets of $P'(T)$.

We now show that given a model of the formula $\forall V_{P(T)}\ \psi(T)$ we can compute an hypothesis satisfying cautious requirements. The vice-versa also holds: given such an hypothesis we are able to find a model of $\forall V_{P(T)}\ \psi(T)$.

Lemma B.2 *Let T be a ground ILP_{LAS} task, let $M \subseteq V^H_{P'(T)}$ be an interpretation of $\forall V_{P(T)}\ \psi(T)$, and let $H := \{h_i \mid h(i) \in M\}$. Then the following are equivalent:*

1. $M \models \forall V_{P(T)}\ \psi(T)$ *;*
2. *For every $A \in AS(B \cup H)$ and every $e^- \in E^-$, it holds that A does not extend e^-.*

Proof. We have to prove two implications.

$\boxed{1 \rightarrow 2}$

Since M is a model of $\forall V_{P(T)}\ \psi(T)$, for every assignment $\sigma : V_{P(T)} \rightarrow \{True, False\}$ to the universal variables, following the definition of $\psi(T)$, it must hold $M, \sigma \models \phi^*_{as}(P'(T)) \rightarrow \phi_{neg}(E^-)$.

For a specific assignment σ, let us consider A'_σ as the interpretation of $\psi(T)$ defined as $A'_\sigma := M \cup A_\sigma$, where A_σ is defined as:

$$A_\sigma := \{\mathtt{y} \mid y \in V_{P(T)}, \sigma(y) = True\} \tag{1}$$

By lemma 4.1, it holds $A'_\sigma \models \phi^*_{as}(P'(T))$ if and only if A'_σ is an answer set of $P'(T) \cup facts(guards \cap A'_\sigma) = B \cup S' \cup \bigcup_{h(i) \in A'_\sigma} \{\mathtt{h(i).}\}$. By definition of the guarded hypothesis space S' we have that A'_σ is also an answer set of $B \cup H' \cup facts(guards \cap A'_\sigma)$ where $H' = \{h'_i \mid \sigma(h(i)) = True\}$. This means that A_σ is an answer set of $B \cup H$ if and only if $A'_\sigma \models \phi^*_{as}(P'(T))$. For a specific σ, suppose that it holds $M, \sigma \models \phi^*_{as}(P'(T))$, then, it must also hold $M, \sigma \models \phi_{neg}(E^-)$. This implies that $A_\sigma \models \phi_{neg}(E^-)$. By fact B.1, we conclude that for every $e^- \in E^-$, A_σ does not extend e^-. Since σ was arbitrary, this holds for every answer set $A_\sigma \in AS(B \cup H)$, completing this part of the proof.

$\boxed{2 \rightarrow 1}$

Let A be an interpretation of $B \cup H$ and let us define $A' := M \cup A$. We now divide our discussion into two cases:

- If $A \notin AS(B \cup H)$, by theorem 2.2, $A \not\models \phi^*_{as}(P(T))$ which implies $A' \not\models \phi^*_{as}(P'(T))$.
- If $A \in AS(B \cup H)$, by theorem 2.2, $A \models \phi^*_{as}(P(T))$ and so $A' \models \phi^*_{as}(P'(T))$. Furthermore, by assumption, for every $e^- \in E^-$, it holds that A does not extend e^-. By fact B.1 $A \models \phi_{neg}(E^-)$ and thus $A' \models \phi_{neg}(E^-)$.

In either case, we obtain $A' \models \psi(T)$. Since $A \subseteq HB_{B \cup S}$ was chosen arbitrary, it follows that $M \models \forall V_{P(T)}\ \psi(T)$. $\square$

The following corollary reformulates lemma B.2 in terms of the validity of the formula $\phi(T) = \exists V^H_{P'(T)}\ \forall V_{P(T)}\ \psi(T)$.

Corollary B.3 *The formula $\phi(T)$ is valid if and only if there exists a model satisfying the formula $\forall V_{P(T)}\ \psi(T)$.*

Our goal is now to reformulate lemma B.2 in terms of the ASP^D program $P^-_{dis}(T)$. To this end, we define two notions of *consistency* of an interpretation of $P^-_{dis}(T)$.

Definition B.4 *Let T be a ground ILP_{LAS} task and consider the program $P^-_{dis}(T)$. Consider the following conditions:*

1. *for every Boolean formula F such that F is in* NNF*and expansion(F) is included in $P^-_{dis}(T)$:*
   - *if F is the conjunction $F = F_1 \wedge F_2$, then, $\mathtt{formula}_F \in A$ if and only if $\mathtt{formula}_{F_1} \in A$ and $\mathtt{formula}_{F_2} \in A$;*
   - *if F is the disjunction $F = F_1 \vee F_2$, then, $\mathtt{formula}_F \in A$ if and only if $\mathtt{formula}_{F_1} \in A$ or $\mathtt{formula}_{F_2} \in A$;*
   - *if F is an atomic formula p, then, $\mathtt{formula}_F \in A$ if and only if $\mathtt{p} \in A$;*
   - *if F is a negated literal $\neg p$, then, $\mathtt{formula}_F \in A$ if and only if $\mathtt{np} \in A$;*
2. *for every $y \in V_{P(T)}$ it holds that $\mathtt{y} \in A'$ and $\mathtt{ny} \in A'$;*
3. *for every $h(i) \in V^H_{P'(T)}$ it holds that exactly one of $\mathtt{h(i)}$ and $\mathtt{nh(i)}$ is included in A' ;*
4. $\mathtt{w} \in A$ *;*
5. $\mathtt{formula}_{\psi(T)} \in A$ *;*
6. *for every $y \in V_{P'(T)}$ it holds that exactly one of $\mathtt{y}$ and $\mathtt{ny}$ is included in A' .*

We say that an interpretation A of $P^-_{dis}(T)$ is:

- *formula-consistent if satisfies satisfies conditions 1 and 6;*
- *consistent if satisfies conditions 1,2,3,4 and 5.*

The formula-consistency property allows to threat interpretations of $P^-_{dis}(T)$ as interpretations of Boolean formulae.

Lemma B.5 *Let M be a formula-consistent interpretation of $P^-_{dis}(T)$. It holds that $\mathtt{formula}_F \in M$ if and only if $M \models F$.*

Proof. Point 6 of definition B.4 allow us to use M as an interpretation for Boolean formulas: for every variable $p \in V_{P'(T)}$ exactly one of $\mathtt{p}$ and $\mathtt{np}$ belongs to M. In the case $\mathtt{p} \in M$ we write $M \models p$. In the case $\mathtt{np} \in M$ we write $M \models \neg p$.

The proof of the claim is a very trivial induction on F by using point 1 of definition B.4. As an example, consider the case $F = F_2 \wedge F_2$. By inductive hypothesis we have (i) $\mathtt{formula}_{F_1} \in M \leftrightarrow M \models F_1$ and (ii) $\mathtt{formula}_{F_2} \in M \leftrightarrow M \models F_2$. Suppose $\mathtt{formula}_F \in M$, by point 1 of definition B.4 we have (iii) $\mathtt{formula}_{F_1} \in M$ and (iv) $\mathtt{formula}_{F_2} \in M$. From points (i) and (iii) we obtain (v) $M \models F_1$; from points (ii) and (iv) we obtain (vi) $M \models F_2$. From (v) and (vi) we obtain $M \models F_1 \wedge F_2 = F$. Conversely, starting from $M \models F$ we can easily establish that $\mathtt{formula}_F \in M$. $\square$

We also prove that the consistency property is a necessary condition for an interpretation A to be an answer set of $P^-_{dis}(T)$. This reflects the proof of theorem 2.1.

Lemma B.6 *Let T be a ground ILP_{LAS} task. Let $A \subseteq HB_{P^-_{dis}(T)}$ be an interpretation. If $A \in AS(P^-_{dis}(T))$ then A is consistent.*

Proof. Let us analyze the structure of a generic answer set $A \in AS(P^-_{dis}(T))$ using the intuitions of the proof of theorem 2.1. First, the predicate `w` (the saturation predicate), must belong to A (point 4 of definition B.4) otherwise the rule `w :- not w` would not be satisfied. In order for `w` to be included in A, it must be supported by the rule `w :-` $\mathtt{formula}_{\psi(T)}$, which in turn requires that $\mathtt{formula}_{\psi(T)} \in A$ (point 5 of definition B.4). For every $y \in V_{P(T)}$ it must hold that $\mathtt{y} \in A$ and $\mathtt{ny} \in A$ (point 2 of definition B.4), otherwise the corresponding rules `y :- w. ny :- w.` would not be satisfied. Moreover, for every $h(i) \in V^H_{P'(T)}$ it must hold that exactly one of `h(i)` and `nh(i)` is included in A (point 3 of definition B.4). If neither is included, then the rule `h(i) | nh(i)` is violated, and if both of them are included, then A is not an answer set because it contradicts the minimality requirement. Regarding point 1 of definition B.4, we analyze the case such that *expansion*(F) is in $P^-_{dis}(T)$ where $F = F_1 \wedge F_2$ (the other cases are addressed in a similar manner). By definition of the *expansion* procedure we have that the rule R_F defined as $\mathtt{formula}_F$ `:-` $\mathtt{formula}_{F_1}$, $\mathtt{formula}_{F_2}$ is included in the program $P_{dis}(T)$. If it holds $\mathtt{formula}_F \in A$, then, the only possibility is that it is supported by the rule R_F and so it must be $\mathtt{formula}_{F_1} \in A$ and $\mathtt{formula}_{F_2} \in A$. Conversely, if it holds $\mathtt{formula}_{F_1} \in A$ and $\mathtt{formula}_{F_2} \in A$, then the atom $\mathtt{formula}_F$ is supported by the rule R_F and so $\mathtt{formula}_F \in A$.

Since A satisfies conditions 1,2,3,4 and 5 of definition B.4, we conclude that A is consistent. □

Notice that the opposite direction of lemma B.6 does not hold. In particular, not every consistent interpretation is an answer set of $P^-_{dis}(T)$.

We now reformulate lemma B.2 in terms of the ASP^D program $P^-_{dis}(T)$. In particular, given a stable model of the program $P^-_{dis}(T)$ we can compute an hypothesis satisfying cautious conditions and vice-versa.

Lemma B.7 *Let T be a ground ILP_{LAS} task, let $A' \subseteq HB_{P^-_{dis}(T)}$ be a consistent Herbrand interpretation of $P^-_{dis}(T)$, and define $H \coloneqq \{h_i \mid \mathtt{h(i)} \in A'\}$. Then the following are equivalent:*

1. *$A' \in AS(P^-_{dis}(T))$;*
2. *For every $A \in AS(B \cup H)$ and every $e^- \in E^-$, it holds that A does not extend e^- .*

Proof. We have to prove two implications. Again, the proof uses concepts from theorem 2.1 by Eiter and Gottlob (1995).

$\boxed{1 \to 2}$

Suppose that $A' \in AS(P^-_{dis}(T))$. Consider any formula-consistent interpretation I that agrees with A' on the assignment of the atoms `h(i)` and `nh(i)` for every i, does not include `w`, and contains exactly one of `y` and `ny` for every $y \in V_{P(T)}$. Such an interpretation I cannot be a stable model, as otherwise it would contradict the minimality of A'. Since I is not a stable model, the only possibility is that $\mathtt{formula}_{\psi(T)} \in I$. By lemma B.5 it holds that $I \models \psi(T)$. Since the choice over variables `y` and `ny` for every $y \in V_{P(T)}$ was arbitrary, this proves that $M \models \forall\, V_{P(T)}\ \psi(T)$ where M is defined as $M \coloneqq \{h(i) \mid \mathtt{h(i)} \in A'\}$. By lemma B.2 it follows directly that for every $A \in AS(B \cup H)$ and every $e^- \in E^-$, it holds that A does not extend e^- .

$\boxed{2 \to 1}$

From lemma B.2 it follows that (i) $M \models \forall\, V_{P(T)}\ \psi(T)$ where $M = \{\mathtt{h(i)} \mid h(i) \in A'\}$. Since A' is consistent, and point (i) holds, it follows that A' is a model of $P^-_{dis}(T)^{A'}$. We have to prove that A' is minimal. Notice that $P^-_{dis}(T)^{A'}$ is the program $P^-_{dis}(T)$ without the rule `w :- not w`. Suppose by contradiction that there exists a model J of $P^-_{dis}(T)^{A'}$ such that $J \subseteq A'$. J must coincide with A' on atoms `h(i)` and `nh(i)` for every $h(i) \in V^H_{P'(T)}$. It must be $\mathtt{w} \notin J$ and for every $y \in V_{P(T)}$ exactly one atom between `y` and `ny` is included in J. Notice that such a model J is by definition formula-consistent. By point (i) it must be $J \models \psi(T)$ and by lemma B.5 it follows that $\mathtt{formula}_{\psi(T)} \in J$. However this implies $\mathtt{w} \in J$ which leads to a contradiction. This proves that A' must be a minimal model of $P^-_{dis}(T)^{A'}$ and so $A' \in AS(P^-_{dis}(T))$. □

From lemma B.7, it follows that the program $P^-_{dis}(T)$ is correct and complete for tasks T with only negative examples.

Lemma B.8 *Let T be a ground ILP_{LAS} task. Let T' be the task T without positive examples. It holds that $ILP_{LAS}(T') = \big\{\{h_i \mid \mathtt{h(i)} \in A\} \mid A \in AS(P^-_{dis}(T))\big\}$.*

Proof. Let $T = \langle B, S, E^+, E^- \rangle$ be a ground ILP_{LAS} task, and let $T' = \langle B, S, \emptyset, E^- \rangle$ denote the corresponding task with all positive examples removed. By the definition of the encoding $P^-_{dis}(\cdot)$, it holds that $P^-_{dis}(T) = P^-_{dis}(T')$, since the construction only depends on negative examples. Let $A \in AS(P^-_{dis}(T))$ be an answer set of the encoding. Define the hypothesis $H \coloneqq \{h_i \mid \mathtt{h(i)} \in A\}$. By lemma B.7 (using the implication $1 \to 2$), we know that for every $A' \in AS(B \cup H)$ and for every $e^- \in E^-$, it holds that A' does not extend e^-. Therefore, by definition of T', H is an inductive solution for T'. Since this holds for every $A \in AS(P^-_{dis}(T))$, it follows that $ILP_{LAS}(T') \supseteq \big\{\{h_i \mid \mathtt{h(i)} \in A\} \mid A \in AS(P^-_{dis}(T))\big\}$. Furthermore, following an analogous reasoning using the implication $2 \to 1$ of lemma B.7, we prove that $ILP_{LAS}(T') \subseteq \big\{\{h_i \mid \mathtt{h(i)} \in A\} \mid A \in AS(P^-_{dis}(T))\big\}$. This concludes the proof. □

We now consider properties of the brave entailment module $P^+_{dis}(T)$. First notice that the program $P^+_{dis}(T)$ is not self-contained: atoms of the form `h(i)` do not appear in the head of any rule. Therefore, in order to use the $P^+_{dis}(T)$ encoding as a standalone module, it is necessary to add rules that explicitly encode the fact that each `h(i)` may or may not belong to an answer set. We define the *completed version* of

$P^+_{dis}(T)$ (denoted $P^{+'}_{dis}(T)$) as the program:

$$P^{+'}_{dis}(T) \coloneqq P^+_{dis}(T) \cup \{\texttt{h(i) | nh(i).} \mid h_i \in S\}$$

Lemma B.9 *Let $A \in AS(P^{+'}_{dis}(T))$, let $A_i = \{q \mid \texttt{inAs}(q,i) \in A\}$ and let $H = \{h_i \mid \texttt{h}(i) \in A\}$. It holds that:*

1. *If $\texttt{interpretation}(i) \in A$, then $A_i \in AS(B \cup H)$;*
2. *A_i extends the example e_i if and only if it holds $\texttt{cov}(i) \in A$.*

Proof. The proof is trivial and similar to the ones by Law (Law 2018) for the ILASP1 algorithm. Point 1 follows since each atom $\texttt{q}$ in a rule $R \in B \cup S$ is encoded by the (non-ground) atom $\texttt{inAs}(\texttt{q}, M)$ and the encoding of R contains the atom $\texttt{interpretation}(M)$ in its body. Moreover each rule h_i in S is encoded by adding the atom $\texttt{h(i)}$ in its body.

We now prove point 2. Consider an example e_i. By applying point 1 to rule $Ex_1(i)$ we have that $A_i \in AS(B \cup H)$. It holds $\texttt{cov}(i) \in A$ if and only if for every atom e^{incl}_j in the inclusion set of e_i it holds $\texttt{inAs}(e^{incl}_j, i) \in A$ and for every atom $\texttt{inAs}(e^{excl}_j, i)$ in the exclusion set of e_i it holds $e^{excl}_j \notin A_i$. This holds if and only if A_i extends the example e_i. □

Lemma B.8 establishes correctness and completeness for the module $P^-_{dis}(\cdot)$ when used with tasks with only negative examples. We now establish the counterpart for the module $P^+_{dis}(\cdot)$. In particular, from lemma B.9 it follows that the program $P^+_{dis}(T)$ is correct and complete for tasks T with only positive examples.

Lemma B.10 *Let T be a ground ILP_{LAS} task. Let T' be the task T without negative examples. It holds that $ILP_{LAS}(T') = \Big\{\{h_i \mid \texttt{h}(i) \in A\} \mid A \in AS(P^{+'}_{dis}(T))\Big\}$.*

Proof. Let $H \in ILP_{LAS}(T')$ be an inductive solution of T'. We construct an interpretation A of $P^{+'}_{dis}(T)$ as follows:

1. $\texttt{h(i)} \in A$ if $h_i \in H$;
2. $\texttt{nh(i)} \in A$ if $h_i \notin H$;
3. $\texttt{cov(i)} \in A$ for every $e_i \in E^+$;
4. $\texttt{inAs(p,i)} \in A$ for every atom p in the inclusion set of $e_i \in E^+$;
5. $\texttt{inAs(q,i)} \notin A$ for every q in the exclusion set of e_i.

It can be verified that A is a stable model of $P^{+'}_{dis}(T)$. In particular, note that for each $h_i \in S$, exactly one of $\texttt{h(i)}$ and $\texttt{nh(i)}$ must belong to A, otherwise, A would not satisfy minimality.

Conversely, suppose $A \in AS(P^{+'}_{dis}(T))$, and define the hypothesis $H = \{h_i \mid \texttt{h(i)} \in A\}$. By construction of the encoding (specifically, rule $Ex_1(i)$, we have $\texttt{interpretation}(i) \in A$ for every $e_i \in E^+$. Rule $Ex_3(i)$ then ensures that $\texttt{cov}(i) \in A$. By lemma B.9, the interpretation $A_i = \{p \mid \texttt{inAs(p,i)} \in A\}$ is a stable model of $B \cup H$ and extends the example e_i. Since this holds for every $e_i \in E^+$, we conclude that H is an inductive solution of T'. □

Lemmas B.8 and B.10 establish the correctness and completeness of the two submodules that compose $P_{dis}(\cdot)$. We now show that their combination yields a full encoding that is correct and complete with respect to the inductive solutions of ground ILP_{LAS} tasks.

Theorem 4.4 (Correctness and completeness) *Let $T = \langle B, S, E^+, E^- \rangle$ be a ground ILP_{LAS} task. It holds that $ILP_{LAS}(T) = \{\{h_i \mid \texttt{h}(i) \in A\} \mid A \in AS(P_{dis}(T))\}$.*

Proof. Let $T^+ = \langle B, S, E^+, \emptyset \rangle$ and $T^- = \langle B, S, \emptyset, E^- \rangle$ denote the positive-only and negative-only variants of T, respectively. Then:

$$ILP_{LAS}(T) = ILP_{LAS}(T^+) \cap ILP_{LAS}(T^-) \tag{2}$$

$$= \Big\{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in AS(P^{+'}_{dis}(T^+))\Big\} \tag{3}$$

$$\cap \ \Big\{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in AS(P^-_{dis}(T^-))\Big\}$$

$$= \Big\{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in AS(P^{+'}_{dis}(T^+)) \tag{4}$$

$$\cap\, AS(P^-_{dis}(T^-))\Big\}$$

$$= \Big\{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in AS(P^{+'}_{dis}(T^+) \cup P^-_{dis}(T^-))\Big\} \tag{5}$$

$$= \Big\{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in AS(P^{+'}_{dis}(T) \cup P^-_{dis}(T))\Big\} \tag{6}$$

$$= \Big\{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in AS(P^+_{dis}(T) \cup P^-_{dis}(T))\Big\} \tag{7}$$

$$= \{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in AS(P_{dis}(T))\} \tag{8}$$

Equation (2) follows directly from the definition of $ILP_{LAS}(T)$ as requiring both brave coverage of positive examples and cautious rejection of negative ones. Equations (3) and (4) follow from lemmas B.8 and B.10, which establish the correctness of the individual positive and negative encodings. To justify Equation (5), observe that: (i) $atoms(P^{+'}_{dis}(T^+)) \cap atoms(P^-_{dis}(T^-)) = \{\texttt{h(i), nh(i)} \mid h_i \in S\}$ and (ii) every answer set of each module contains exactly one of $\texttt{h(i)}$ or $\texttt{nh(i)}$ for each $h_i \in S$. As a consequence, the intersection of answer sets from the two modules determines the same hypothesis (i.e., the same assignment to $\texttt{h(i)}$ atoms) as any answer set of their union. Equation (6) holds because the modules $P^{+'}_{dis}(T^+)$ and $P^-_{dis}(T^-)$ are constructed to handle only positive and negative examples, respectively. Therefore, their semantics remain unaffected when both modules are used with the full task $T = \langle B, S, E^+, E^- \rangle$. Equation (7) holds because the rules of the form $\texttt{h(i) | nh(i)}$ do not need to appear in both modules. Since they are already included in $P^-_{dis}(T)$, they can be safely omitted from $P^{+'}_{dis}(T)$, allowing us to use $P^+_{dis}(T)$ in its place. Finally, equation (8) holds by definition of $P_{dis}(T)$ as the union $P^+_{dis}(T) \cup P^-_{dis}(T)$. This concludes the proof. □

So far, we have established the correctness and completeness of $P_{dis}(\cdot)$. We now focus on the structural complexity of the encoding.

Lemma 4.5 *Let $T = \langle B, S, E^+, E^- \rangle$ be a ground ILP_{LAS} task. Then the following holds:*

1. $|P_{dis}(T)| \in \Theta\left(|T| + |\phi^*_{loop}(P'(T))|\right)$
2. $|ground(P_{dis}(T))| \in \Theta(\,|B \cup S| \times |E^+| + |E^-| + |B| + |S| + |\phi^*_{loop}(P'(T))|\,)$
3. *If $B \cup S$ is tight then $|P_{dis}(T)| \in \Theta(|T|)$ and $|ground(P_{dis}(T))| \in \Theta(|B \cup S| \times |E^+| + |E^-| + |B| + |S|)$.*

Proof. The total size of $P_{dis}(T)$ is the sum of the sizes of $P^-_{dis}(T)$ and $P^+_{dis}(T)$. The number of rules in $P^+_{dis}(T)$ is in $\Theta(|B| + |S| + |E^+|)$, since each rule and each positive example contributes a constant number of clauses. The module $P^-_{dis}(T)$ is obtained by translating the formula

$$\phi(T) = \exists V^H_{P'(T)}\, \forall V_{P(T)} \left(\mathsf{NNF}\left(\phi^*_{as}(P'(T)) \rightarrow \phi_{neg}(E^-)\right)\right)$$

into a ground $\mathsf{ASP^D}$ program.

Observe that the NNF transformation is linear in the size of its argument. However, the presence of bi-implications may cause an exponential blow-up during the transformation. We therefore examine the size of $\phi(T)$ more carefully:

$$\begin{aligned}
|\phi(T)| &= |\mathsf{NNF}(\phi^*_{as}(P'(T)))| + |\mathsf{NNF}(\phi_{neg}(E^-))| && (9)\\
&= |\mathsf{NNF}(\phi^*_{iff}(P'(T)))| + |\mathsf{NNF}(\phi^*_{loop}(P'(T)))| && (10)\\
&\quad + |\mathsf{NNF}(\phi_{neg}(E^-))| \\
&= |\mathsf{NNF}(\phi^*_{iff}(P'(T)))| + \Theta(|\phi^*_{loop}(P'(T))|) && (11)\\
&\quad + \Theta(|\phi_{neg}(E^-)|) \\
&= |\mathsf{NNF}(\phi^*_{iff}(P'(T)))| + \Theta(|\phi^*_{loop}(P'(T))|) + \Theta(|E^-|) && (12)\\
&= \Theta(|\phi^*_{iff}(P'(T))|) + \Theta(|\phi^*_{loop}(P'(T))|) + \Theta(|E^-|) && (13)\\
&= \Theta(|B \cup S|) + |\phi^*_{loop}(P'(T))| + \Theta(|E^-|) && (14)
\end{aligned}$$

Equation (9) reflects the definition of $\phi(T)$. Equation (10) follows form the definition of $\phi^*_{as}(P'(T))$ as the conjunction of $\phi^*_{iff}(P'(T))$ and $\phi^*_{loop}(P'(T))$. Equation (11) holds because both $\phi^*_{loop}(P'(T))$ and $\phi_{neg}(E^-)$ do not contain bi-implications, and their NNF transformation is therefore linear in size. Equation (12) trivially follows by the definition of $\phi_{neg}(E^-)$. We now consider equation (13). Notice that $\phi^*_{iff}(P'(T))$ contains exactly one bi-implication for every atom $a \in atoms(P'(T)) \setminus S$ where $S = \{\mathtt{h(i)} \mid h_i \in S\}$. The structure of $\phi^*_{iff}(P'(T))$ is as follows:

$$\bigwedge_{a \in atoms(P'(T)) \setminus S} lhs(a) \leftrightarrow rhs(a)$$

Each bi-implication can be rewritten as the conjunction of two implications:

$$\bigwedge_{a \in atoms(P'(T)) \setminus S} (lhs(a) \rightarrow rhs(a)) \wedge (rhs(a) \rightarrow lhs(a))$$

This transformation only doubles the size of the formula, and since NNF of an implication is linear, the transformation is still linear in total size. Finally, equation (14) follows from the observation that each rule $R \in B \cup S$ contributes a bounded number of atoms to the completion $\phi^*_{iff}(P'(T))$. This analysis shows that the size of $P^-_{dis}(T)$ is in $\Theta(|B| + |S| + |\phi^*_{loop}(P'(T))| + |E^-|)$. By summing this with the size of $P^+_{dis}(T)$, we obtain the claimed bound for $|P_{dis}(T)|$.

To prove point 2, observe that $P^-_{dis}(\cdot)$ is a ground program. In $P^+_{dis}(T)$, each rule in $B \cup S$ is instantiated once for each interpretation label `interpretation(i)`, that is, once per positive example.

For point 3, recall that the formula $\phi^*_{loop}(P'(T))$ is empty if and only if the program $B \cup S$ is tight, in which case the loop formulas are not required. □

C Proofs from Section 5

We now provide proofs regarding the grounding process. We prove that every safe ILP_{LAS} task T is equivalent to its grounded version.

Theorem 5.2 *Let T be a (possibly non-ground) safe ILP_{LAS} task. Let $T' = ground(T)$. It holds that T and T' are equivalent.*

Proof. Let $T = \{B, S = \{h_1, \ldots, h_s\}, E^+, E^-\}$ and let $T' = \{B', S', E^+, E^-\}$ be the ILP^{agg}_{LAS} task $ground(T)$ where:

- U is the Herbrand universe of $B \cup S$.
- $B' := \bigcup_{R_i \in B} ground_U(R_i)$
- $S' := \{h'_i \mid h_i \in S, h'_i = ground_U(h_i)\}$

Let J be a set of indexes $\{j_1, \ldots, j_k\}$ such that $1 \leq j_i \leq s$ for every $j_i \in J$. Define the corresponding hypothesis for both T and T':

$$H := \{h_{j_1}, \ldots, h_{j_k}\} \subseteq S \qquad H' := \{h'_{j_1}, \ldots, h'_{j_k}\} \subseteq S'$$

To prove the theorem we have to prove that H is an inductive solution of T if and only if H' is an inductive solution of T', in symbols:

$$H \in ILP_{LAS}(T) \leftrightarrow H' \in ILP^{agg}_{LAS}(T') \qquad (15)$$

We begin by studying answer sets of $B \cup H$ and $B' \cup \bigcup_{h'_i \in H'} h'_j$. For simplicity we define the following abbreviations:

$$P = B \cup S \qquad P_H = B \cup H \qquad U = HU_P$$
$$U_H = HU_{P_H} \qquad P'_H = B' \cup \bigcup_{h'_i \in H'} h'_i$$

Consider the case in which P_H is non-ground. We have that $A \in AS(P_H)$ if and only if $A \in AS(ground_{U_H}(P_H))$. The notation highlights the fact that the ground instance of each rule in P_H is computed considering the Herbrand universe of P_H (i.e. U_H). We now prove the following fact:

$$A \in AS(ground_{U_H}(P_H)) \rightarrow A \in AS(ground_U(P_H)) \quad (16)$$

To prove equation (16), we proceed by induction. In particular, let $F \subseteq P$ and let $F^+ = F \cup \{h\}$ where $h \in P \setminus F$, we have to prove that if $A \in AS(ground_{HU_F}(P_H))$ then $A \in AS(ground_{HU_{F^+}}(P_H))$. If it holds $HU_F = HU_{F^+}$, i.e. h does not introduce new terms, then trivially we have $A \in AS(ground_{HU_{F^+}}(P_H))$. Consider now the case where $HU_F \subsetneq HU_{F^+}$ and in particular let $V = HU_{F^+} \setminus HU_F$ the set of terms in HU_{F^+} but not in HU_F. Let W denote the set of non-ground atoms appearing in P_H. In program $ground_{HU_{F^+}}(P_H)$, atoms in W can be instantiated also with terms in V. Denote with Z the ground instances of atoms in W such that, each $z_i \in Z$ contains at least one term in V, i.e, z_i is of type $z_i(\dots, v_j, \dots)$ for some $v_j \in V$. Suppose to prove that (i) each atom in Z cannot be supported by $ground_{HU_{F^+}}(P_H)$. Consider a rule $R \in ground_{HU_{F^+}}(P_H) \setminus ground_{HU_F}(P_H)$. By definition R must contain an atom in $z_i(\dots, v_j, \dots) \in Z$ and since the task is safe it must be that (ii) there must exists $z'_i(\dots, v_j, \dots) \in Z \cap body^+(R)$, i.e., there exists a positive literal $z'_i \in Z$ occurring in the body of R and such that z'_i contains the term v_j. By combining (i) and (ii) we obtain that R can never be activated and hence $A \in AS(ground_{HU_{F^+}}(P_H))$. To prove point (i), notice that, trivially, (iii) $ground_{HU_F}(P_H)$ cannot support any atom in Z. Interestingly, also every rule in $K := ground_{HU_{F^+}}(P_H) \setminus ground_{HU_F}(P_H)$ cannot support an atom in Z. Consider $G(K, Z)$ as the following graph (which is a slight modification of the positive atom dependency graph defined in section 2):

$$G(K, Z) = (Z, \{(b, h) \mid R \in K, b \in body^+(R),$$
$$h = head(R), b \in Z, h \in Z\})$$

Let $G^c(K, Z)$ be the condensation of $G(K, Z)$, i.e. the graph in which each strongly connected component is contracted to a single vertex. Let $\pi = \langle C_1, \dots, C_k \rangle$ be a topological sorting of $G^c(K, Z)$. Notice that (by using property (iii)) each atom in C_1 cannot be supported by any rule in $ground_{HU_{F^+}}(P_H)$. By a simple inductive argument it holds that each atom in C_i cannot be supported in $ground_{HU_{F^+}}(P_H)$. This proves property (i) and equation (16).

The converse of equation (16) also holds. In particular, we have:

$$A \in AS(ground_{P_H}(P_H)) \leftarrow A \in AS(ground_P(P_H)) \quad (17)$$

Equation (17) is proved in a similar manner. We proceed by induction by removing rules from P and gradually arriving to P_H, in particular we prove that if $A \in AS(ground_F(P_H))$ then $A \in AS(ground_{F^-}(P_H))$ where $F \subseteq S$ and $F^- = F \setminus \{h\}$ for $h \in F \setminus P_H$. This is trivially verified if $HU_F = HU_{F^-}$. Consider the case $HU_F \supsetneq HU_{F^-}$ and let $V = HU_F \setminus HU_{F^-}$. This case is addressed exactly like in the proof of equation (16). The intuition is that rules in $K = ground_F(P_H) \setminus ground_{F^-}(P_H)$ can never be activated since (iv) $ground_{F^-}(P_H)$ cannot support any atom with a term in V and (v) at least one term in V appears in the body of every rule in K (since T is safe). This proves equation (17).

The key intuition behind equation (16) and equation (17) is that, due to the safety of the task, rules involving new terms introduced by a larger Herbrand universe are not activatable and thus do not influence the stable models.

We can now prove that an answer set of P_H is also an answer set of P'_H and vice-versa:

$$\begin{aligned}
&A \in AS(P_H) \\
\leftrightarrow\ &A \in AS(ground_{U_H}(P_H)) \\
\leftrightarrow\ &A \in AS(ground_U(P_H)) \\
\leftrightarrow\ &A \in AS(ground_U(B \cup H)) \\
\leftrightarrow\ &A \in AS(\bigcup_{R_i \in B} ground_U(R_i) \cup \bigcup_{h_i \in H} ground_U(h_i)) \\
\leftrightarrow\ &A \in AS(B' \cup \bigcup_{h'_i \in H'} h'_i) \\
\leftrightarrow\ &A \in AS(P'_H)
\end{aligned}$$

We now conclude the proof of the theorem by proving equation (15):

$$\begin{aligned}
&H \in ILP_{LAS}(T) \\
\leftrightarrow\ &\forall e^+ \in E^+\ \exists A \in AS(P_H)\ A \text{ extends } e^+ \text{ and} \\
&\forall e^- \in E^-\ \forall A \in AS(P_H)\ A \text{ does not extend } e^- \\
\leftrightarrow\ &\forall e^+ \in E^+\ \exists A \in AS(P'_H)\ A \text{ extends } e^+ \text{ and} \\
&\forall e^- \in E^-\ \forall A \in AS(P'_H)\ A \text{ does not extend } e^- \\
\leftrightarrow\ &H' \in ILP^{agg}_{LAS}(T')
\end{aligned}$$

□

As a direct consequence of theorem 5.2, we now show that, given a correct and complete encoding P_{enc} for the grounded framework ILP^{agg}_{LAS}, the composition $P_{enc}(ground(\cdot))$ yields a correct and complete encoding for every safe ILP_{LAS} task.

Theorem C.1 (Correctness and completeness) *Let $T = \langle B, \{h_1, \dots, h_s\}, E^+, E^- \rangle$ be a (possibly non-ground) safe ILP_{LAS} task. Let $P_{enc}(\cdot)$ be a correct and complete encoding for ground tasks in the ILP^{agg}_{LAS} framework, i.e., for every ground task $T' = \langle B', \{h'_1, \dots, h'_s\}, E'^+, E'^- \rangle$, it holds that: $ILP^{agg}_{LAS}(T') = \{\{h'_i \mid \texttt{h'(i)} \in A\} \mid A \in AS(P_{enc}(T'))\}$ Then, the composition $P_{enc}(ground(T))$ is a correct and complete encoding for T, that is the set $ILP_{LAS}(T)$ corresponds to:*

$$\{\{h_i \mid \texttt{h(i)} \in A\} \mid A \in\ AS(P_{enc}(ground(T)))\}$$

Proof. Let $T'' = \langle B'', \{h''_1, \dots, h''_n\}, E^+, E^- \rangle$ be the task ILP^{agg}_{LAS} $T'' = ground(T)$. Let . By theorem We have that:

$$H = \{h_{j_1}, \dots, h_{j_1}\} \in ILP_{LAS}(T) \quad (18)$$
$$\leftrightarrow H'' = \{h''_{j_1}, \dots, h''_{j_1}\} \in ILP^{agg}_{LAS}(ground(T)) \quad (19)$$
$$\leftrightarrow \exists A \in AS(P_{enc}(T'')) \text{ such that} \quad (20)$$
$$\{\texttt{h''(i)} \mid h''(i) \in A\} = H''$$

Equation 19 follows from theorem 5.2. Equation 20 follows from the fact that $P_{enc}(\cdot)$ is a correct and complete encoding for ground ILP^{agg}_{LAS} tasks. □