Imperial College London

The Complexity and Generality of Learning Answer Set Programs

(AIJ 2018)

Mark Law, Alessandra Russo and Krysia Broda

September 2, 2018

ILP under the Answer Set Semantics

- Several ILP frameworks have been proposed to learn ASP:
  - In ILP_b (resp ILP_c) *at least one* (resp *every*) answer set of $B \cup H$ must cover the (atom) examples.
  - In ILP_{LAS} examples are partial interpretations and a combination of ILP_b and ILP_c can be expressed.

ILP under the Answer Set Semantics

- Several ILP frameworks have been proposed to learn ASP:
  - In ILP_b (resp ILP_c) *at least one* (resp *every*) answer set of $B \cup H$ must cover the (atom) examples.
  - In ILP_{LAS} examples are partial interpretations and a combination of ILP_b and ILP_c can be expressed.
- This paper asks two fundamental questions:
  - What class of ASP programs can each framework learn?
  - Is there any (complexity) price paid by the more general frameworks?

ILP under the Answer Set Semantics

- Several ILP frameworks have been proposed to learn ASP:
  - In ILP_b (resp ILP_c) *at least one* (resp *every*) answer set of $B \cup H$ must cover the (atom) examples.
  - In ILP_{LAS} examples are partial interpretations and a combination of ILP_b and ILP_c can be expressed.
- This paper asks two fundamental questions:
  - What class of ASP programs can each framework learn?
  - Is there any (complexity) price paid by the more general frameworks?
- In the paper we also consider ILP_{sm}, ILP_{LOAS} and $ILP_{LOAS}^{context}$.

Imperial College London

One-to-one Distinguishability

Definition 1

A learning framework $\mathcal{F}$ can *distinguish* H_1 from H_2 wrt B iff there is at least one task $T_{\mathcal{F}} = \langle B, E_{\mathcal{F}} \rangle$ such that $H_1 \in \mathcal{F}(T_{\mathcal{F}})$ and $H_2 \notin \mathcal{F}(T_{\mathcal{F}})$.

One-to-one Distinguishability

Definition 1

A learning framework $\mathcal{F}$ can *distinguish* H_1 from H_2 wrt B iff there is at least one task $T_{\mathcal{F}} = \langle B, E_{\mathcal{F}} \rangle$ such that $H_1 \in \mathcal{F}(T_{\mathcal{F}})$ and $H_2 \notin \mathcal{F}(T_{\mathcal{F}})$.

- $\mathcal{D}_1^1(\mathcal{F})$ is the set of tuples $\langle B, H_1, H_2 \rangle$ such that $\mathcal{F}$ can distinguish H_1 from H_2 wrt B.

One-to-one Distinguishability

Definition 1

A learning framework $\mathcal{F}$ can *distinguish* H_1 from H_2 wrt B iff there is at least one task $T_{\mathcal{F}} = \langle B, E_{\mathcal{F}} \rangle$ such that $H_1 \in \mathcal{F}(T_{\mathcal{F}})$ and $H_2 \notin \mathcal{F}(T_{\mathcal{F}})$.

- $\mathcal{D}^1_1(\mathcal{F})$ is the set of tuples $\langle B, H_1, H_2 \rangle$ such that $\mathcal{F}$ can distinguish H_1 from H_2 wrt B.

Let $B = \emptyset$, $H_1 = \{\texttt{p.}\}$ and $H_2 = \left\{\texttt{0}\{\texttt{p}\}\texttt{1.}\right\}$.

- $\langle B, H_1, H_2 \rangle$ is not in $\mathcal{D}^1_1(ILP_b)$.

Imperial College London

One-to-one Distinguishability

Definition 1

A learning framework $\mathcal{F}$ can *distinguish* H_1 from H_2 wrt B iff there is at least one task $T_{\mathcal{F}} = \langle B, E_{\mathcal{F}} \rangle$ such that $H_1 \in \mathcal{F}(T_{\mathcal{F}})$ and $H_2 \notin \mathcal{F}(T_{\mathcal{F}})$.

- $\mathcal{D}^1_1(\mathcal{F})$ is the set of tuples $\langle B, H_1, H_2 \rangle$ such that $\mathcal{F}$ can distinguish H_1 from H_2 wrt B.

Let $B = \emptyset$, $H_1 = \{\texttt{p.}\}$ and $H_2 = \big\{\texttt{0}\{\texttt{p}\}\texttt{1.}\big\}$.

- $\langle B, H_1, H_2 \rangle$ is not in $\mathcal{D}^1_1(ILP_b)$.

$E^+ = \{\texttt{p}\}$ $\qquad$ $E^- = \emptyset$

$H_2 \in ILP_b(\langle B, \{\texttt{p}\}, \emptyset \rangle)$.

One-to-one Distinguishability

Definition 1

A learning framework $\mathcal{F}$ can *distinguish* H_1 from H_2 wrt B iff there is at least one task $T_{\mathcal{F}} = \langle B, E_{\mathcal{F}} \rangle$ such that $H_1 \in \mathcal{F}(T_{\mathcal{F}})$ and $H_2 \notin \mathcal{F}(T_{\mathcal{F}})$.

- $\mathcal{D}^1_1(\mathcal{F})$ is the set of tuples $\langle B, H_1, H_2 \rangle$ such that $\mathcal{F}$ can distinguish H_1 from H_2 wrt B.

Let $B = \emptyset$, $H_1 = \{\texttt{p.}\}$ and $H_2 = \left\{\texttt{0}\{\texttt{p}\}\texttt{1.}\right\}$.

- $\langle B, H_1, H_2 \rangle$ is not in $\mathcal{D}^1_1(ILP_b)$.
- $\langle B, H_2, H_1 \rangle$ is in $\mathcal{D}^1_1(ILP_b)$.

One-to-one Distinguishability

Definition 1

A learning framework $\mathcal{F}$ can *distinguish* H_1 from H_2 wrt B iff there is at least one task $T_{\mathcal{F}} = \langle B, E_{\mathcal{F}} \rangle$ such that $H_1 \in \mathcal{F}(T_{\mathcal{F}})$ and $H_2 \notin \mathcal{F}(T_{\mathcal{F}})$.

- $\mathcal{D}^1_1(\mathcal{F})$ is the set of tuples $\langle B, H_1, H_2 \rangle$ such that $\mathcal{F}$ can distinguish H_1 from H_2 wrt B.

Let $B = \emptyset$, $H_1 = \{\texttt{p.}\}$ and $H_2 = \{\texttt{0}\{\texttt{p}\}\texttt{1.}\}$.

- $\langle B, H_1, H_2 \rangle$ is not in $\mathcal{D}^1_1(ILP_b)$.
- $\langle B, H_2, H_1 \rangle$ is in $\mathcal{D}^1_1(ILP_b)$.

$E^+ = \emptyset$ $\qquad$ $E^- = \{\texttt{p}\}$

$H_2 \in ILP_b(\langle B, \emptyset, \{\texttt{p}\} \rangle)$ but $H_1 \notin ILP_b(\langle B, \emptyset, \{\texttt{p}\} \rangle)$.

One-to-one Distinguishability

Definition 1

A learning framework $\mathcal{F}$ can *distinguish* H_1 from H_2 wrt B iff there is at least one task $T_{\mathcal{F}} = \langle B, E_{\mathcal{F}} \rangle$ such that $H_1 \in \mathcal{F}(T_{\mathcal{F}})$ and $H_2 \notin \mathcal{F}(T_{\mathcal{F}})$.

- $\mathcal{D}^1_1(\mathcal{F})$ is the set of tuples $\langle B, H_1, H_2 \rangle$ such that $\mathcal{F}$ can distinguish H_1 from H_2 wrt B.

Let $B = \emptyset$, $H_1 = \{\texttt{p.}\}$ and $H_2 = \big\{\texttt{0}\{\texttt{p}\}\texttt{1.}\big\}$.

- $\langle B, H_1, H_2 \rangle$ is not in $\mathcal{D}^1_1(ILP_b)$.
- $\langle B, H_2, H_1 \rangle$ is in $\mathcal{D}^1_1(ILP_b)$.
- $\langle B, H_1, H_2 \rangle$ is in $\mathcal{D}^1_1(ILP_c)$.

One-to-one Distinguishability

Definition 1

A learning framework $\mathcal{F}$ can *distinguish* H_1 from H_2 wrt B iff there is at least one task $T_{\mathcal{F}} = \langle B, E_{\mathcal{F}} \rangle$ such that $H_1 \in \mathcal{F}(T_{\mathcal{F}})$ and $H_2 \notin \mathcal{F}(T_{\mathcal{F}})$.

- $\mathcal{D}^1_1(\mathcal{F})$ is the set of tuples $\langle B, H_1, H_2 \rangle$ such that $\mathcal{F}$ can distinguish H_1 from H_2 wrt B.

Let $B = \emptyset$, $H_1 = \{\mathtt{p.}\}$ and $H_2 = \big\{\mathtt{0\{p\}1.}\big\}$.

- $\langle B, H_1, H_2 \rangle$ is not in $\mathcal{D}^1_1(ILP_b)$.
- $\langle B, H_2, H_1 \rangle$ is in $\mathcal{D}^1_1(ILP_b)$.
- $\langle B, H_1, H_2 \rangle$ is in $\mathcal{D}^1_1(ILP_c)$.

$E^+ = \{\mathtt{p}\}$ $\qquad$ $E^- = \emptyset$

$H_1 \in ILP_c(\langle B, \emptyset, \{\mathtt{p}\} \rangle)$ but $H_2 \notin ILP_c(\langle B, \emptyset, \{\mathtt{p}\} \rangle)$.

Imperial College London

One-to-one Distinguishability Conditions

Framework $\mathcal{F}$	Sufficient/necessary condition for $\langle B, H_1, H_2 \rangle$ to be in $\mathcal{D}_1^1(\mathcal{F})$
ILP_b	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_{sm}	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_c	$AS(B \cup H_1) \neq \emptyset \wedge (AS(B \cup H_2) = \emptyset \vee (\mathcal{E}_c(B \cup H_1) \not\subseteq \mathcal{E}_c(B \cup H_2)))$
ILP_{LAS}	$AS(B \cup H_1) \neq AS(B \cup H_2)$
ILP_{LOAS}	$(AS(B \cup H_1) \neq AS(B \cup H_2)) \vee (ord(B \cup H_1) \neq ord(B \cup H_2))$
$ILP_{LOAS}^{context}$	$(B \cup H_1 \not\equiv^s B \cup H_2) \vee$ $(\exists C \in \mathcal{ASP}^{ch}$ st $ord(B \cup H_1 \cup C) \neq ord(B \cup H_2 \cup C))$

One-to-one Distinguishability Conditions

Framework $\mathcal{F}$	Sufficient/necessary condition for $\langle B, H_1, H_2\rangle$ to be in $\mathcal{D}_1^1(\mathcal{F})$
ILP_b	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_{sm}	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_c	$AS(B \cup H_1) \neq \emptyset \wedge (AS(B \cup H_2) = \emptyset \vee (\mathcal{E}_c(B \cup H_1) \not\subseteq \mathcal{E}_c(B \cup H_2)))$
ILP_{LAS}	$AS(B \cup H_1) \neq AS(B \cup H_2)$
ILP_{LOAS}	$(AS(B \cup H_1) \neq AS(B \cup H_2)) \vee (ord(B \cup H_1) \neq ord(B \cup H_2))$
$ILP_{LOAS}^{context}$	$(B \cup H_1 \not\equiv^s B \cup H_2) \vee$ $(\exists C \in \mathcal{ASP}^{ch}$ st $ord(B \cup H_1 \cup C) \neq ord(B \cup H_2 \cup C))$

- Neither ILP_b of ILP_{sm} can distinguish $H \cup C$ from H for any constraint C and any H – in practice, neither ILP_b nor ILP_{sm} can learn constraints.

One-to-one Distinguishability Conditions

Framework $\mathcal{F}$	Sufficient/necessary condition for $\langle B, H_1, H_2\rangle$ to be in $\mathcal{D}^1_1(\mathcal{F})$
ILP_b	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_{sm}	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_c	$AS(B \cup H_1) \neq \emptyset \wedge (AS(B \cup H_2) = \emptyset \vee (\mathcal{E}_c(B \cup H_1) \not\subseteq \mathcal{E}_c(B \cup H_2)))$
ILP_{LAS}	$AS(B \cup H_1) \neq AS(B \cup H_2)$
ILP_{LOAS}	$(AS(B \cup H_1) \neq AS(B \cup H_2)) \vee (ord(B \cup H_1) \neq ord(B \cup H_2))$
$ILP^{context}_{LOAS}$	$(B \cup H_1 \not\equiv^s B \cup H_2) \vee$ $(\exists C \in \mathcal{ASP}^{ch}$ st $ord(B \cup H_1 \cup C) \neq ord(B \cup H_2 \cup C))$

- ILP_{LAS} can distinguish any two hypotheses, so long as they have different answer sets (when combined with B).

One-to-one Distinguishability Conditions

Framework $\mathcal{F}$	Sufficient/necessary condition for $\langle B, H_1, H_2\rangle$ to be in $\mathcal{D}^1_1(\mathcal{F})$
ILP_b	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_{sm}	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_c	$AS(B \cup H_1) \neq \emptyset \wedge (AS(B \cup H_2) = \emptyset \vee (\mathcal{E}_c(B \cup H_1) \not\subseteq \mathcal{E}_c(B \cup H_2)))$
ILP_{LAS}	$AS(B \cup H_1) \neq AS(B \cup H_2)$
ILP_{LOAS}	$(AS(B \cup H_1) \neq AS(B \cup H_2)) \vee (ord(B \cup H_1) \neq ord(B \cup H_2))$
$ILP^{context}_{LOAS}$	$(B \cup H_1 \not\equiv^s B \cup H_2) \vee$ $(\exists C \in \mathcal{ASP}^{ch} \text{ st } ord(B \cup H_1 \cup C) \neq ord(B \cup H_2 \cup C))$

- $ILP^{context}_{LOAS}$ can distinguish any two hypotheses, so long as they are not strongly equivalent (when combined with B).

One-to-one Distinguishability Conditions

Framework $\mathcal{F}$	Sufficient/necessary condition for $\langle B, H_1, H_2 \rangle$ to be in $\mathcal{D}_1^1(\mathcal{F})$
ILP_b	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_{sm}	$AS(B \cup H_1) \not\subseteq AS(B \cup H_2)$
ILP_c	$AS(B \cup H_1) \neq \emptyset \wedge (AS(B \cup H_2) = \emptyset \vee (\mathcal{E}_c(B \cup H_1) \not\subseteq \mathcal{E}_c(B \cup H_2)))$
ILP_{LAS}	$AS(B \cup H_1) \neq AS(B \cup H_2)$
ILP_{LOAS}	$(AS(B \cup H_1) \neq AS(B \cup H_2)) \vee (ord(B \cup H_1) \neq ord(B \cup H_2))$
$ILP_{LOAS}^{context}$	$(B \cup H_1 \not\equiv^s B \cup H_2) \vee$ $(\exists C \in \mathcal{ASP}^{ch}$ st $ord(B \cup H_1 \cup C) \neq ord(B \cup H_2 \cup C))$

$$\mathcal{D}_1^1(ILP_b) = \mathcal{D}_1^1(ILP_{sm}) \subset \mathcal{D}_1^1(ILP_{LAS}) \subset \mathcal{D}_1^1(ILP_{LOAS}) \subset \mathcal{D}_1^1(ILP_{LOAS}^{context})$$
$$\mathcal{D}_1^1(ILP_c) \subset \mathcal{D}_1^1(ILP_{LAS})$$

Imperial College London

One-to-many Distinguishability

Definition 2

For a framework $\mathcal{F}$, $\mathcal{D}^1_m(\mathcal{F})$ is the set of tuples $\langle B, H, \{H_1, \ldots, H_n\}\rangle$ st there is a task $T_{\mathcal{F}}$ which distinguishes H from each H_i with respect to B.

One-to-many Distinguishability

Definition 2

For a framework $\mathcal{F}$, $\mathcal{D}^1_m(\mathcal{F})$ is the set of tuples $\langle B, H, \{H_1, \ldots, H_n\}\rangle$ st there is a task $T_\mathcal{F}$ which distinguishes H from each H_i with respect to B.

Let $B = \emptyset$, $H = \{1\{\texttt{heads}, \texttt{tails}\}1.\}$, $H'_1 = \{\texttt{heads.}\}$, $H'_2 = \{\texttt{tails.}\}$

- $\langle B, H, H'_1\rangle \in \mathcal{D}^1_1(ILP_b)$ and $\langle B, H, H'_2\rangle \in \mathcal{D}^1_1(ILP_b)$

Imperial College London

One-to-many Distinguishability

Definition 2

For a framework $\mathcal{F}$, $\mathcal{D}^1_m(\mathcal{F})$ is the set of tuples $\langle B, H, \{H_1, \ldots, H_n\}\rangle$ st there is a task $T_{\mathcal{F}}$ which distinguishes H from each H_i with respect to B.

Let $B = \emptyset$, $H = \{1\{\texttt{heads}, \texttt{tails}\}1.\}$, $H'_1 = \{\texttt{heads.}\}$, $H'_2 = \{\texttt{tails.}\}$

- $\langle B, H, H'_1\rangle \in \mathcal{D}^1_1(ILP_b)$ and $\langle B, H, H'_2\rangle \in \mathcal{D}^1_1(ILP_b)$
- $\langle B, H, \{H'_1, H'_2\}\rangle \notin \mathcal{D}^1_m(ILP_b)$

One-to-many Distinguishability

Definition 2

For a framework $\mathcal{F}$, $\mathcal{D}^1_m(\mathcal{F})$ is the set of tuples $\langle B, H, \{H_1, \ldots, H_n\}\rangle$ st there is a task $T_{\mathcal{F}}$ which distinguishes H from each H_i with respect to B.

Let $B = \emptyset$, $H = \{\texttt{1\{heads, tails\}1.}\}$, $H'_1 = \{\texttt{heads.}\}$, $H'_2 = \{\texttt{tails.}\}$

- $\langle B, H, H'_1\rangle \in \mathcal{D}^1_1(ILP_b)$ and $\langle B, H, H'_2\rangle \in \mathcal{D}^1_1(ILP_b)$
- $\langle B, H, \{H'_1, H'_2\}\rangle \notin \mathcal{D}^1_m(ILP_b)$
- $\langle B, H, \{H'_1, H'_2\}\rangle \in \mathcal{D}^1_m(ILP_{sm})$

One-to-many Distinguishability

Definition 2

For a framework $\mathcal{F}$, $\mathcal{D}^1_m(\mathcal{F})$ is the set of tuples $\langle B, H, \{H_1, \ldots, H_n\}\rangle$ st there is a task $T_{\mathcal{F}}$ which distinguishes H from each H_i with respect to B.

$$\mathcal{D}^1_m(ILP_b) \subset \mathcal{D}^1_m(ILP_{sm}) \subset \mathcal{D}^1_m(ILP_{LAS}) \subset \mathcal{D}^1_m(ILP_{LOAS}) \subset \mathcal{D}^1_m(ILP^{context}_{LOAS})$$

$$\mathcal{D}^1_m(ILP_c) \subset \mathcal{D}^1_m(ILP_{LAS})$$

Many-to-many Distinguishability

Definition 3

For a framework $\mathcal{F}$, $\mathcal{D}^m_m(\mathcal{F})$ is the set of tuples $\langle B, S_1, S_2 \rangle$, st there is a task $T_{\mathcal{F}}$ with background B, st $S_1 \subseteq ILP_{\mathcal{F}}(T_{\mathcal{F}})$ and $S_2 \cap ILP_{\mathcal{F}}(T_{\mathcal{F}}) = \emptyset$.

Many-to-many Distinguishability

Definition 3

For a framework $\mathcal{F}$, $\mathcal{D}^m_m(\mathcal{F})$ is the set of tuples $\langle B, S_1, S_2 \rangle$, st there is a task $T_{\mathcal{F}}$ with background B, st $S_1 \subseteq ILP_{\mathcal{F}}(T_{\mathcal{F}})$ and $S_2 \cap ILP_{\mathcal{F}}(T_{\mathcal{F}}) = \emptyset$.

$\mathcal{D}^m_m(ILP_b) \subset \mathcal{D}^m_m(ILP_{sm}) \subset \mathcal{D}^m_m(ILP_{LAS}) \subset \mathcal{D}^m_m(ILP_{LOAS}) \subset \mathcal{D}^m_m(ILP^{context}_{LOAS})$

$\mathcal{D}^m_m(ILP_c) \subset \mathcal{D}^m_m(ILP_{LAS})$

Complexity

Framework	Verification	Satisfiablity
ILP_b	NP-complete	NP-complete
ILP_{sm}	NP-complete	NP-complete
ILP_c	DP-complete	Σ_2^P-complete
ILP_{LAS}	DP-complete	Σ_2^P-complete
ILP_{LOAS}	DP-complete	Σ_2^P-complete
$ILP_{LOAS}^{context}$	DP-complete	Σ_2^P-complete

Imperial College London

Complexity

Framework	Verification	Satisfiablity
ILP_b	NP-complete	NP-complete
ILP_{sm}	NP-complete	NP-complete
ILP_c	DP-complete	Σ_2^P-complete
ILP_{LAS}	DP-complete	Σ_2^P-complete
ILP_{LOAS}	DP-complete	Σ_2^P-complete
$ILP_{LOAS}^{context}$	DP-complete	Σ_2^P-complete
ILP_{LOAS}^{noise}	DP-complete	Σ_2^P-complete

Conclusion

- We have introduced three new measures of the generality of a learning framework.
- For each of the three measures:

  $\mathcal{D}(ILP_b) \subseteq \mathcal{D}(ILP_{sm}) \subset \mathcal{D}(ILP_{LAS}) \subset \mathcal{D}(ILP_{LOAS}) \subset \mathcal{D}(ILP^{context}_{LOAS})$

  $\mathcal{D}(ILP_c) \subset \mathcal{D}(ILP_{LAS})$

- There is no price to be paid (in terms of complexity) for the gain in generality of $ILP^{context}_{LOAS}$ over ILP_c.
- ILP_b and ILP_{sm} are of lower complexity, but are less general than ILP_{LAS}.

Backup Slides

One-to-many Distinguishability

- In the paper, we proved that if for any two $\mathcal{F}$ tasks T_1, T_2 there is a task T_3 such that $ILP_{\mathcal{F}}(T_3) = ILP_{\mathcal{F}}(T_1) \cap ILP_{\mathcal{F}}(T_2)$ then:

$$\mathcal{D}^1_m(\mathcal{F}) = \left\{ \langle B, H, \{H_1, \ldots, H_n\}\rangle \;\middle|\; \begin{array}{c} \langle B, H, H_1\rangle \in \mathcal{D}^1_1(\mathcal{F}), \\ \ldots, \\ \langle B, H, H_n\rangle \in \mathcal{D}^1_1(\mathcal{F}) \end{array} \right\}.$$

- In ILP_{LAS}, T_3 can be constructed as $\langle B, E_1^+ \cup E_2^+, E_1^- \cup E_2^- \rangle$.
- This property holds for every framework (in the paper) other than ILP_b.

$$\mathcal{D}^1_m(ILP_b) \subset \mathcal{D}^1_m(ILP_{sm}) \subset \mathcal{D}^1_m(ILP_{LAS}) \subset \mathcal{D}^1_m(ILP_{LOAS}) \subset \mathcal{D}^1_m(ILP^{context}_{LOAS})$$

$$\mathcal{D}^1_m(ILP_c) \subset \mathcal{D}^1_m(ILP_{LAS})$$

Brave Induction cannot learn constraints

- Let H be a hypothesis and C be a constraint.
- For any $T = \langle B, E^+, E^- \rangle$ st $H \cup C \in ILP_b(T)$, there is an $A \in AS(B \cup H \cup C)$ st $E^+ \subseteq A$ and $E^- \cap A = \emptyset$.

  Any such A is also an answer set of $B \cup H$.
- Hence ILP_b cannot distinguish $H \cup C$ from H (wrt any background knowledge).
- In practice this means that ILP_b cannot learn constraints.

Other notion of generality

- (De Raedt 1997) defined generality in terms of reductions. $\mathcal{F}_1$ is said to be more general than $\mathcal{F}_2$ iff $\mathcal{F}_2 \rightarrow_r \mathcal{F}_1$ and $\mathcal{F}_1 \not\rightarrow_r \mathcal{F}_2$.
- These reductions allowed the background knowledge B to be modified in the reduction, whereas distinguishability does not.
- In the paper we define *strong reductions* which force the background knowledge to be the same and show that $\mathcal{F}_1 \rightarrow_{sr} \mathcal{F}_2$ if and only if $\mathcal{D}^m_m(\mathcal{F}_1) \subseteq \mathcal{D}^m_m(\mathcal{F}_2)$.
- Other than the restriction on the background knowledge, distinguishability also allows for fine grained comparisons of frameworks which are *incomparable* under reductions and strong reductions.

De Raedt, L. 1997.
Logical settings for concept-learning.
Artificial Intelligence 95, 1, 187–201.