COPYRIGHT NOTICE

DOI: `10.1109/ICDM.2009.127`.

Outlier Detection using Inductive Logic Programming

Fabrizio Angiulli and Fabio Fassetti
DEIS, University of Calabria
87036 Rende (CS), Italy
{*f.angiulli, f.fassetti*}*@deis.unical.it*

***Abstract*—We present a novel definition of outlier in the context of inductive logic programming. Given a set of positive and negative examples, the definition aims at singling out the examples showing anomalous behavior. We note that the task here pursued is different from noise removal, and, in fact, the anomalous observations we discover are different in nature from noisy ones. We discuss pecularities of the novel approach, present an algorithm for detecting outliers, discuss some examples of knowledge mined, and compare it with alternative approaches.**

I. Introduction

Traditional approaches to outlier detection model the normal behavior of individuals by performing some statistical kind of computation on the given data set and, then, single out those individuals whose behavior or characteristics significantly deviate from normal ones. However, a very interesting direction of research concerns the capability of exploiting *domain knowledge* in order to guide search for anomalous observations. Indeed, while looking over a set of observations to discover outliers, it often happens that there is some qualitative description of the domain of interest encoding what an expected normal behavior should be. This description might be, for instance, derived by an expert and might be formalized by means of a suitable language for knowledge representation. With this aim, in [1], [2] a concept of outlier in the context of default reasoning ad logic programming is presented.

Inductive Logic Programming (ILP) is an important field at the intersection of machine learning and logic programming which aims at inducing relation descriptions of data in the form of logic programs [3]. In the context above delineated, we present a definition of outlier in the field of ILP. Given a set of positive and negative examples from a concept, the definition aims at singling out the examples showing exceptional behavior. The method is unsupervised, since there are no examples of normal/abnormal behavior, even if it has connections with supervised learning, since it is based on induction from examples which are instances of a concept. In particular, the definition distinguishes among three kinds of abnormalities, that are *irregular*, *anomalous*, and *outlier* observations. This allows us to provide a finer characterization of the anomaly at hand and to single out more subtle forms of anomalies. Moreover, we are also able to provide *exaplanations* for the abnormality of the observation, in the form of a pair of logic programs, which make more intellegible the motivation underlying its exceptionality.

ILP learning systems usually have a single mechanism, called *noise-handling mechanism* [4], for dealing with noisy, incomplete and inexact data, which prevents the induced hypothesis from overfitting the data set. While the presence of noise in the examples has some relationship with the approach here pursued, the other kinds of imperfect data are orthogonal to the present task, in that they are problems of the learning task in its entirety.

There are major differences also between the approach presented here and those introduced in [1], [2]. First, our approach is based of *induction*, while the others are based on *deduction*. Indeed, the disagreement of the abnormal observations with the theory at hand is perceived here by means of a measure of the difference of the generalization of the hypotheses induced in presence/absence of the observations (the compliance relationship), while in [1], [2] it is perceived by means of the satisfaction of certain conditions involving the entailment operator. Moreover, importantly, the definitions proposed in [1], [2] strongly rely on the non-monotonicity of the employed formalisms, which use either *default rules* or *negation by default* (other than the *classical negation*). As a matter of fact, if the logic program under consideration is positive, that is under the formal framework considered here, according to the definition provided in [2] *there are no outliers in a logic program* (cf. Theorem 3.2 of [2]). This makes the two approaches for defining outliers incomparable even from a practical point of view.

The rest of the paper is organized as follows. Section II presents some preliminary definitions. Section III introduces the notion of abnormal example. Section IV provides the definition of ILP-based outlier mining problem. Section V reports experimental results. Finally, Section VI draws conclusions.

II. Preliminary definitions

The reader is referred to [5] for basics on Logic Programming and to [3] for Inductive Logic Programming. Next we provide some definitions useful to define the concept of outlier in ILP.

The set of ground atoms covered by a clause c is denoted as *covers*(c). If C is a set of clauses, *covers*(C) is the union of the set of ground atoms covered by the clauses in C. Let E be a set of ground atoms. In the following, $covers_E(C)$ denotes the set $covers(C) \cap E$. By definition, we assume that $covers_E(\emptyset) = E$.

Given a logic program P and a set of ground atoms E, the *restriction* $P(E)$ of P to E is the logic program $P(E) = \{c \in P \mid covers_E(\{c\}) \neq \emptyset\}$.

Let U be an universal set of *observations*, also called *objects* or *instances*. A (*direct*) *concept* $\mathcal{C}$ is a subset of U. The *dual concept* $\overline{\mathcal{C}}$ of $\mathcal{C}$ is the concept $U \setminus \mathcal{C}$.

A *set of examples* $\mathcal{E}$ is a set of ground atoms that can be partitioned in two subsets, that are $\mathcal{E}^+$, the *set of positive examples*, and $\mathcal{E}^-$, the *set of negative examples*.

The problem that ILP is interested in solving can be stated as follows: *Given a set of examples $\mathcal{E}$, find a hypothesis $\mathcal{H}_\mathcal{B}^\mathcal{E}$ such that $\mathcal{H}_\mathcal{B}^\mathcal{E} \cup \mathcal{B}$ entails the examples in $\mathcal{E}$; namely:*

1) *for each $e \in \mathcal{E}^+$, $e \in covers(\mathcal{H}_\mathcal{B}^\mathcal{E} \cup \mathcal{B})$ (completeness),*
2) *for each $e \in \mathcal{E}^-$, $e \notin covers(\mathcal{H}_\mathcal{B}^\mathcal{E} \cup \mathcal{B})$ (consistency).*

Next we provide the concepts of *coverage*, *gain*, and *compliance*.

Definition 1 (Coverage): Let C be a set of clauses and let $\mathcal{E}$ be a set of examples. Then the *coverage* $cov_\mathcal{E}(C)$ of C in $\mathcal{E}$ is the following function:

$$cov_\mathcal{E}(C) = \frac{1}{|\mathcal{E}|} \left(\prod_{c \in C} |covers_\mathcal{E}(c)| \right)^{\frac{1}{|C|}}.$$

Intuitively, the coverage of a set of clauses in a set of examples measures how many examples of the set are covered in average by the clauses. In particular, the definition of coverage here provided is based on the geometric mean in order to penalize the presence of rules covering few examples. We will employ the coverage as a measure of the *generalization* of a set of clauses.

Definition 2 (Gain): Given two sets of clauses C_1 and C_2 and a set of examples $\mathcal{E}$, the gain $gain_\mathcal{E}(C_1, C_2)$ of C_1 over C_2 in $\mathcal{E}$ is defined as

$$gain_\mathcal{E}(C_1, C_2) = cov_\mathcal{E}(C_1) - cov_\mathcal{E}(C_2).$$

A positive gain means that the clauses in C_1 averagely cover a larger number of examples in $\mathcal{E}$ than the clauses in C_2. Intuitively, this means that the clauses in C_1 can be considered *more general* than those in C_2.

Given a set of examples $\mathcal{E}$ and a nonempty subset $\mathcal{O}$ of $\mathcal{E}$ we say that $\mathcal{O}$ is *pure* if either $\mathcal{O} \subseteq \mathcal{E}^+$ or $\mathcal{O} \subseteq \mathcal{E}^-$ hold.

Definition 3 (Compliance): Given a background knowledge $\mathcal{B}$, a set of examples $\mathcal{E}$, and a pure subset $\mathcal{O}$ of $\mathcal{E}$, we say that $\mathcal{O}$ α-complies (or, simply, complies) with $\mathcal{E} \cup \mathcal{B}$, written $\mathcal{E} \cup \mathcal{B} \rightsquigarrow \mathcal{O}$, if

$$gain_{\mathcal{E}^+ \setminus \mathcal{O}} \left(\mathcal{H}_\mathcal{B}^{\mathcal{E} \setminus O}, \mathcal{H}_\mathcal{B}^\mathcal{E} \right) < \alpha,$$

where α is an user-provided parameter in $[0, 1]$. Otherwise, $\mathcal{O}$ does not comply with $\mathcal{E} \cup \mathcal{B}$, written $\mathcal{E} \cup \mathcal{B} \not\rightsquigarrow \mathcal{O}$.

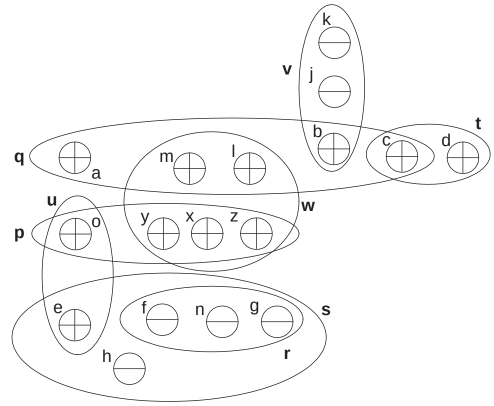

Figure 1. Example theory.

Intuitively, if a subset of examples does not comply with a background theory and the whole set of examples, then this means that the hypothesis induced in absence of this subset is significatively more general than the hypothesis induced when the examples are seen.

Given a set of examples $\mathcal{E}$, the *dual set* $\overline{\mathcal{E}}$ of $\mathcal{E}$ is the set of examples $\overline{\mathcal{E}}$ such that $\overline{\mathcal{E}}^+ = \mathcal{E}^-$ and $\overline{\mathcal{E}}^- = \mathcal{E}^+$. Note that by using $\overline{\mathcal{E}}$ as set of examples, the dual concept $\overline{\mathcal{C}}$ of $\mathcal{C}$ is learned, where $\mathcal{C}$ is the concept of which the examples $\mathcal{E}$ are instances. Let p denote the target predicate symbol. When the dual concept is learned we will employ in the induced hypothesis the predicate symbol not_p instead of p.

III. Abnormal examples

Given a pure subset of examples $\mathcal{O}$ of $\mathcal{E}$, we argue that the compliance of these examples with $\mathcal{E} \cup \mathcal{B}$ and $\overline{\mathcal{E}} \cup \mathcal{B}$ can be exploited in order to understand if the set $\mathcal{O}$ contains *abnormal* observations. In order to illustrate the concepts that will be defined in this section, we make use of an example.

EXAMPLE 1. Let

$$\mathcal{E}^+ = \{t_p(a), t_p(b), t_p(c), t_p(d), t_p(e), t_p(l), t_p(m)\},$$
$$\mathcal{E}^- = \{t_p(f), t_p(g), t_p(h), t_p(i), t_p(j), t_p(k), t_p(n)\}$$

be a given set of examples, and

$$\mathcal{B} = \{p(o),\ p(x),\ p(y),\ p(z),\ q(a),\ q(b),\ q(c),\ q(l),\ q(m), r(f),\ r(g),\ r(n),\ s(e),\ s(f),\ s(g),\ s(h),\ s(n), t(c),\ t(d), u(a),\ u(e),\ u(o),\ v(b),\ v(j),\ v(k), w(m),\ w(l),\ w(x),\ w(y),\ w(z)\}$$

be a given background knowledge. Figure 1 shows the examples in $\mathcal{E}$ and the subsets covered by each predicate in the background theory.

Let the induced hypothesis $\mathcal{H}_\mathcal{B}^\mathcal{E}$ be (throughout the paper, between angle brackets we report the number of examples covered by each clause):

$$\mathcal{H}_\mathcal{B}^\mathcal{E} = \left\{ \begin{array}{ll} c_1 \equiv t_p(X) \leftarrow q(X) & \langle 5 \rangle \\ c_2 \equiv t_p(X) \leftarrow w(X) & \langle 5 \rangle \\ c_3 \equiv t_p(X) \leftarrow u(X) & \langle 3 \rangle \\ c_4 \equiv t_p(X) \leftarrow t(X) & \langle 2 \rangle \end{array} \right.$$

and the induced dual hypothesis $\mathcal{H}_{\mathcal{B}}^{\overline{\mathcal{E}}}$ be:

$$\mathcal{H}_{\mathcal{B}}^{\overline{\mathcal{E}}} = \begin{cases} \overline{c}_1 \equiv not\_t_p(X) \leftarrow r(X) & \langle 3 \rangle \\ \overline{c}_2 \equiv not\_t_p(h) & \langle 1 \rangle \\ \overline{c}_3 \equiv not\_t_p(j) & \langle 1 \rangle \\ \overline{c}_4 \equiv not\_t_p(k) & \langle 1 \rangle \end{cases}$$

For the sake of simplicity, assume that $\mathcal{O}$ is a subset of $\mathcal{E}^+$. As already said, if the set $\mathcal{O}$ does not comply with $\mathcal{E} \cup \mathcal{B}$, then the description of the concept would be significantly more concise if each example in $\mathcal{O}$ were not observed. Hence, intuitively, we can say that the examples in $\mathcal{O}$ are likely to do not match regularities joining the remaining instances of the concept. In order to better understand the kind of irregularity represented by the examples in $\mathcal{O}$, the compliance of $\overline{\mathcal{O}}$ with $\overline{\mathcal{E}} \cup \mathcal{B}$ has to be investigated. In particular, if $\mathcal{O}$ does not comply with $\mathcal{E} \cup \mathcal{B}$, but $\overline{\mathcal{O}}$ complies with $\overline{\mathcal{E}} \cup \mathcal{B}$, then the description of the concept would be significantly more concise if each example in $\mathcal{O}$ were not observed, whereas the dual description of the concept is not affected by the set of examples $\mathcal{O}$. Thus, in this case the examples are hard to be covered since we can imagine they are "far away" the majority of the positive examples, but anyway "far" form the negative ones. We identify these examples as *irregular*.

Definition 4 (Irregular set): Given a background knowledge $\mathcal{B}$, a set of examples $\mathcal{E}$, and a subset $\mathcal{O}$ of $\mathcal{E}^+$ *(*$\mathcal{E}^-$, resp.*)*, we say that $\mathcal{O}$ is *irregular* in $\mathcal{E} \cup \mathcal{B}$ if $\mathcal{E} \cup \mathcal{B} \not\rightsquigarrow \mathcal{O}$ *(*$\mathcal{E} \cup \mathcal{B} \rightsquigarrow \mathcal{O}$, resp.*)* and $\overline{\mathcal{E}} \cup \mathcal{B} \rightsquigarrow \overline{\mathcal{O}}$ *(*$\overline{\mathcal{E}} \cup \mathcal{B} \not\rightsquigarrow \overline{\mathcal{O}}$, resp.*)*.

Example 1 (continued). Assume α is set to 0.04 and consider the set $\mathcal{O} = \{t_p(d)\}$. We note that the set $\mathcal{O}$ is irregular. Indeed, $\mathcal{E} \cup \mathcal{B} \not\rightsquigarrow \{t_p(d)\}$, since if $\{t_p(d)\}$ were not seen then the induced theory $\mathcal{H}_{\mathcal{B}}^{\mathcal{E}\setminus\mathcal{O}}$ would not contain c_4 being, hence, more concise. In particular, $cov_{\mathcal{E}^+\setminus\mathcal{O}}(\mathcal{H}_{\mathcal{B}}^{\mathcal{E}}) = \frac{\sqrt[4]{5\cdot5\cdot2\cdot1}}{10} = 0.27$, while $cov_{\mathcal{E}^+\setminus\mathcal{O}}(\mathcal{H}_{\mathcal{B}}^{\mathcal{E}\setminus\mathcal{O}}) = \frac{\sqrt[3]{5\cdot5\cdot2}}{10} = 0.37$, and the gain is 0.10. Conversely, $\overline{\mathcal{E}} \cup \mathcal{B} \rightsquigarrow \{t_p(d)\}$ since $\mathcal{H}_{\mathcal{B}}^{\overline{\mathcal{E}}\setminus\overline{\mathcal{O}}}$ is not affected by the absence of $\{t_p(d)\}$ (the gain is zero).

A similar line of reasoning can be employed if $\overline{\mathcal{O}}$ does not comply with $\overline{\mathcal{E}} \cup \mathcal{B}$. In particular, if $\overline{\mathcal{O}}$ does not comply with $\overline{\mathcal{E}} \cup \mathcal{B}$ and $\mathcal{O}$ complies with $\mathcal{E} \cup \mathcal{B}$, then the examples in $\mathcal{O}$ well fit the concept to be learned, but they also have some commonalities with the dual concept, so that it is very difficult to discriminate them from a non-instance. Hence, we call these examples *anomalous*.

Definition 5 (Anomalous set): Given a subset $\mathcal{O}$ of $\mathcal{E}^+$ *(*$\mathcal{E}^-$, resp.*)*, we say that $\mathcal{O}$ is *anomalous* in $\mathcal{E} \cup \mathcal{B}$ if $\mathcal{E} \cup \mathcal{B} \rightsquigarrow \mathcal{O}$ *(*$\mathcal{E} \cup \mathcal{B} \not\rightsquigarrow \mathcal{O}$, resp.*)* and $\overline{\mathcal{E}} \cup \mathcal{B} \not\rightsquigarrow \overline{\mathcal{O}}$ *(*$\overline{\mathcal{E}} \cup \mathcal{B} \rightsquigarrow \overline{\mathcal{O}}$, resp.*)*.

Example 1 (continued). Consider now the set $\mathcal{O} = \{t_p(b)\}$. This set is anomalous. Indeed, $\mathcal{E} \cup \mathcal{B} \rightsquigarrow \{t_p(b)\}$ since the hypothesis $\mathcal{H}_{\mathcal{B}}^{\mathcal{E}\setminus\mathcal{O}}$ induced in absence of $\mathcal{O}$ coincides with $\mathcal{H}_{\mathcal{B}}^{\mathcal{E}}$. Conversely, $\overline{\mathcal{E}} \cup \mathcal{B} \not\rightsquigarrow \{t_p(b)\}$ since if $\{t_p(b)\}$ were not seen then the induced dual theory $\mathcal{H}_{\mathcal{B}}^{\overline{\mathcal{E}}\setminus\overline{\mathcal{O}}}$ would contain the clause $t_p(X) \leftarrow v(X)$ instead of the facts $\overline{c}_3$ and $\overline{c}_4$. In fact, $cov_{\overline{\mathcal{E}}^+\setminus\overline{\mathcal{O}}}(\mathcal{H}_{\mathcal{B}}^{\overline{\mathcal{E}}\setminus\overline{\mathcal{O}}}) = \frac{\sqrt[3]{3\cdot2\cdot1}}{6}$, $cov_{\overline{\mathcal{E}}\setminus\overline{\mathcal{O}}}(\mathcal{H}_{\mathcal{B}}^{\overline{\mathcal{E}}}) = \frac{\sqrt[4]{3\cdot1\cdot1\cdot1}}{6}$, and the gain is 0.08.

$\mathcal{O} \subset \mathcal{E}$ pure	$\overline{\mathcal{E}} \cup \mathcal{B} \rightsquigarrow \overline{\mathcal{O}}$	$\overline{\mathcal{E}} \cup \mathcal{B} \not\rightsquigarrow \overline{\mathcal{O}}$
$\mathcal{E} \cup \mathcal{B} \rightsquigarrow \mathcal{O}$	*normal*	*anomalous* (positive) *irregular* (negative)
$\mathcal{E} \cup \mathcal{B} \not\rightsquigarrow \mathcal{O}$	*irregular* (positive) *anomalous* (negative)	*outlier*

Table I
THE DIFFERENT KINDS OF ABNORMAL EXAMPLES.

Assume now that $\mathcal{O}$ does not comply with $\mathcal{E} \cup \mathcal{B}$ and also $\overline{\mathcal{O}}$ does not comply with $\overline{\mathcal{E}} \cup \mathcal{B}$. In this case also the description of the dual concept would be significantly more concise if each example in $\overline{\mathcal{O}}$ were not observed. Intuitively, this means that the examples in $\mathcal{O}$ present some commonalities with the non-instances of the concept to be learned. In this case, both the description of the concept $\mathcal{C}$ and the description of the dual concept $\overline{\mathcal{C}}$ would benefit if the examples in $\mathcal{O}$ and in $\overline{\mathcal{O}}$, respectively, were not observed. We can imagine that these examples are hard to be covered, since they lie either very close or even within the "shape" of the dual concept, and we identify these examples as *outliers*.

Definition 6 (Outlier set): Given a pure subset $\mathcal{O}$ of $\mathcal{E}$, we say that $\mathcal{O}$ is *outlier* in $\mathcal{E} \cup \mathcal{B}$ if $\mathcal{E} \cup \mathcal{B} \not\rightsquigarrow \mathcal{O}$ and $\overline{\mathcal{E}} \cup \mathcal{B} \not\rightsquigarrow \overline{\mathcal{O}}$.

Example 1 (continued). The set $\mathcal{O} = \{t_p(e)\}$ is an outlier. Indeed, $\mathcal{E} \cup \mathcal{O} \not\rightsquigarrow \{t_p(e)\}$, since if $\{t_p(e)\}$ were not seen then the induced theory $\mathcal{H}_{\mathcal{B}}^{\mathcal{E}\setminus\mathcal{O}}$ would not contain c_3, being, hence, more concise. In particular, $cov_{\mathcal{E}^+\setminus\mathcal{O}}(\mathcal{H}_{\mathcal{B}}^{\mathcal{E}\setminus\mathcal{O}}) = \frac{\sqrt[3]{5\cdot4\cdot2}}{10}$, since $\mathcal{H}_{\mathcal{B}}^{\mathcal{E}\setminus\mathcal{O}}$ would contain $c_5 \equiv t_p(X) \leftarrow p(X)$ instead of c_2 and c_3, $cov_{\mathcal{E}^+\setminus\mathcal{O}}(\mathcal{H}_{\mathcal{B}}^{\mathcal{E}}) = \frac{\sqrt[4]{5\cdot5\cdot2\cdot1}}{10}$ and the gain is 0.08. Moreover, $\overline{\mathcal{E}} \cup \mathcal{B} \not\rightsquigarrow \{t_p(e)\}$, since if $\{t_p(e)\}$ were not seen then the induced dual theory $\mathcal{H}_{\mathcal{B}}^{\overline{\mathcal{E}}\setminus\overline{\mathcal{O}}}$ would contain the clause $\overline{c}_6 \equiv not\_t_p(X) \leftarrow s(X)$ instead of $\overline{c}_1$ and $\overline{c}_2$. In particular, $cov_{\overline{\mathcal{E}}^+\setminus\overline{\mathcal{O}}}(\mathcal{H}_{\mathcal{B}}^{\overline{\mathcal{E}}\setminus\overline{\mathcal{O}}}) = \frac{\sqrt[3]{4\cdot1\cdot1}}{6}$ and $cov_{\overline{\mathcal{E}}^+\setminus\overline{\mathcal{O}}}(\mathcal{H}_{\mathcal{B}}^{\overline{\mathcal{E}}}) = \frac{\sqrt[4]{3\cdot1\cdot1\cdot1}}{6}$, and the gain is 0.045.

A pure set of examples $\mathcal{O}$ such that both $\mathcal{E} \cup \mathcal{B} \rightsquigarrow \mathcal{O}$ and $\overline{\mathcal{E}} \cup \mathcal{B} \rightsquigarrow \overline{\mathcal{O}}$ hold is said to be *normal*, otherwise it is said to be *abnormal*. Abnormal set of examples can be partitioned into outlier, irregular and anomalous examples, according to what aforesaid. Table I summarizes the different kinds of abnormal example sets that have been defined.

IV. STATEMENT OF THE PROBLEM AND ALGORITHM

In this section we define the *outlier detection problem* in the context of ILP. First of all, we introduce an alternative notion of compliance which is based on the previous one, but presents some advantages which will be discussed next. Intuitively, the novel definition of compliance focuses on the portion of the theory involving only the examples in $\mathcal{O}$.

For the sake of simplicity, let $\mathcal{O}$ be a set of positive examples and consider the theory $\mathcal{H}_{\mathcal{B}}^{\mathcal{E}}$ induced in presence of $\mathcal{O}$. The set of clauses in $\mathcal{H}_{\mathcal{B}}^{\mathcal{E}}$ can be partitioned in two groups, that are the clauses in $\mathcal{H}_{\mathcal{B}}^{\mathcal{E}}(\mathcal{O})$ that cover some of the examples in $\mathcal{O}$, and the remaining ones, that are the clauses in $\mathcal{H}^{\mathcal{E}}_{\mathcal{B}} \setminus \mathcal{H}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$. We call *starting theory* the former theory, since it builds on the examples in $\mathcal{O}$.

As for the set of examples, it can be partitioned in three sets, that are the examples in $\mathcal{O}$, the examples $\widehat{\mathcal{O}}$ not in $\mathcal{O}$ and covered only by clauses in $\mathcal{H}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$, and the remaining ones. Consider now the theory $\mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}$ induced in absence of $\mathcal{O}$. In this case, the piece of theory which is affected by the absence of the examples in $\mathcal{O}$ is that composed of the clauses that cover some of the examples in $\widehat{\mathcal{O}}$ (recall that these examples are those not covered by the clauses in $\mathcal{H}^{\mathcal{E}}_{\mathcal{B}} \setminus \mathcal{H}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$). Hence, we call *ending theory* this theory.

We can now redefine the compliance by exploiting the gain of the ending theory over the starting theory in $\mathcal{E}^+ \setminus \mathcal{O}$.

Definition 7 (Compliance): Given a pure subset $\mathcal{O}$ of $\mathcal{E}$, let $\widehat{\mathcal{E}}$ denote $\mathcal{E}^+ \setminus \mathcal{O}$. We say that $\mathcal{O}$ α-complies (or, simply, complies) with $\mathcal{E} \cup \mathcal{B}$, written $\mathcal{E} \cup \mathcal{B} \rightsquigarrow \mathcal{O}$, if

$$gain_{\widehat{\mathcal{E}}}\left(\overleftarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O}), \overrightarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})\right) < \alpha,$$

where the theories $\overrightarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$ (the *starting theory*) and $\overleftarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$ (the *ending theory*), are defined as follows:

- if $\mathcal{O} \subseteq \mathcal{E}^+$, then $\overrightarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$ is $\mathcal{H}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$ and $\overleftarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$ is $\mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}(\widehat{\mathcal{O}})$, where $\widehat{\mathcal{O}}$ is

$$covers_{\widehat{\mathcal{E}}}\left(\mathcal{H}^{\mathcal{E}}_{\mathcal{B}}\right) \setminus covers_{\widehat{\mathcal{E}}}\left(\mathcal{H}^{\mathcal{E}}_{\mathcal{B}} \setminus \mathcal{H}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})\right),$$

  that is the set of examples in $\mathcal{E}^+ \setminus \mathcal{O}$ covered only by clauses of $\mathcal{H}^{\mathcal{E}}_{\mathcal{B}}$ that cover some examples in $\mathcal{O}$;
- if $\mathcal{O} \subseteq \mathcal{E}^-$, then $\overrightarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$ is $\mathcal{H}^{\mathcal{E}}_{\mathcal{B}}(\widehat{\mathcal{O}})$ and $\overleftarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$ is $\mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}(\mathcal{O})$, where $\widehat{\mathcal{O}}$ is

$$covers_{\widehat{\mathcal{E}}}\left(\mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}\right) \setminus covers_{\widehat{\mathcal{E}}}\left(\mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}} \setminus \mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}(\mathcal{O})\right),$$

  that is the set of examples in $\mathcal{E}^+ \setminus \mathcal{O} = \mathcal{E}^+$ covered only by clauses of $\mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}$ that cover some examples in $\mathcal{O}$.

Example 1 (continued). Consider the outlier set $\mathcal{O} = \{t_p(e)\}$. Then the starting theory $\overrightarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$ is $\mathcal{H}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O}) = \{c_3\}$ and, moreover, $\widehat{\mathcal{O}}$ is

$$covers_{\mathcal{E}^+\setminus\{t_p(e)\}}(\{c_1, c_2, c_3, c_4\}) \setminus covers_{\mathcal{E}^+\setminus\{t_p(e)\}}(\{c_1, c_2, c_4\})$$

that is $\{t_p(o)\}$. Let the theory $\mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}$ induced in absence of $\mathcal{O}$ be

$$\mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}} = \begin{cases} c'_1 \equiv t_p(X) \leftarrow q(X) & \langle 5 \rangle \\ c'_2 \equiv t_p(X) \leftarrow p(X) & \langle 4 \rangle \\ c'_3 \equiv t_p(X) \leftarrow t(X) & \langle 2 \rangle \end{cases}$$

Then the ending theory $\overleftarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O})$ is $\mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}(\widehat{\mathcal{O}}) = \{c'_2\}$, and the gain is $\frac{3}{10} - \frac{1}{10} = 0.20$.

As for the dual concept, $\overline{\mathcal{O}}$ is $\{t_p(e)\}$, let the theory $\mathcal{H}^{\overline{\mathcal{E}}\setminus\overline{\mathcal{O}}}_{\mathcal{B}}$ induced in absence of $\overline{\mathcal{O}}$ be

$$\mathcal{H}^{\overline{\mathcal{E}}\setminus\overline{\mathcal{O}}}_{\mathcal{B}} = \begin{cases} \overline{c}'_1 \equiv not\_t_p(X) \leftarrow s(X) & \langle 4 \rangle \\ \overline{c}'_2 \equiv not\_t_p(j) & \langle 1 \rangle \\ \overline{c}'_3 \equiv not\_t_p(k) & \langle 1 \rangle \end{cases}$$

Then the ending theory $\overleftarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}})$ is $\mathcal{H}^{\overline{\mathcal{E}}\setminus\overline{\mathcal{O}}}_{\mathcal{B}}(\overline{\mathcal{O}}) = \{\overline{c}'_1\}$. The set $\widehat{\overline{\mathcal{O}}}$ is

$$covers_{\mathcal{E}^-}(\{\overline{c}'_1, \overline{c}'_2, \overline{c}'_3\}) \setminus covers_{\mathcal{E}^-}(\{\overline{c}'_2, \overline{c}'_3\}),$$

that is $\{not\_t_p(f), not\_t_p(n), not\_t_p(g), not\_t_p(h)\}$. The starting theory $\overrightarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}})$ is $\mathcal{H}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\widehat{\overline{\mathcal{O}}}) = \{\overline{c}_1, \overline{c}_2\}$, and the gain is $\frac{4}{6} - \frac{\sqrt{3 \cdot 1}}{6} = 0.38$.

Now we discuss the advantages of the novel definition of compliance.

First, we note that the *starting theory* and the *ending theory* associated with the abnormal set of examples play the role of *explanation* for its abnormality. Indeed, since they represent the portion of knowledge which is affected by the presence/absence of the abnormal set, by comparing them the analyst can understand the motivation underlying the abnormality of the example set. With this aim, if $\mathcal{O}$ ($\overline{\mathcal{O}}$, resp.) does not comply with $\mathcal{E} \cup \mathcal{B}$, then we call *direct* (*dual*, resp.) *explanation* the pair $(\overrightarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O}), \overleftarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O}))$ ($(\overrightarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}}), \overleftarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}}))$, resp.), also said the *direct* (*dual*, resp.) *starting/ending theories* associated with $\mathcal{O}$ ($\overline{\mathcal{O}}$, resp.) in $\mathcal{H}^{\mathcal{E}}_{\mathcal{B}}$ ($\mathcal{H}^{\overline{\mathcal{E}}}_{\mathcal{B}}$, resp.).

Second, comparing only the pieces of theories related to the abnormal set is more meaningful than comparing the two full theories, and, moreover, this kind of comparison makes the definition less sensitive to "global changes".

Now we are in the position of formally defining the *ILP-based Outlier Detection* problem.

Definition 8: Given a background knowledge $\mathcal{B}$, a set of examples $\mathcal{E}$, and a maximum size k_{max}, find the minimal abnormal subsets $\mathcal{O}$ of $\mathcal{E}$ of size not exceeding k_{max}, together with their associated explanations.

The pseudo-code of the ILP-*based-outlier-detector* algorithm for mining the ILP-based outliers is reported in Figure 2. The functions *Gain*$^+$ and *Gain*$^-$ are used to test compliance according to Definition 7.

V. Experiments

In this section we present some experiments conducted by using the proposed algorithm. We implemented the outlier detection algorithm in Yap Prolog on top of the P-Progol system[1] which is based on the PROGOL algorithm [6].

A. Zoo data set

In this experiment we considered the Zoo data set from the UCI Machine Learning Repository[2]. This database contains instances associated with animals. Each instance consists of the animal *name*, the *class* which it belongs to (amphibian, bird, fish, invertebrate, insect, mammal, reptile), the number of *legs* and other boolean attributes. We built a background theory consisting of one unary predicate for each boolean attribute, and of the binary predicate *legs*. We used as target

[1] `www.comlab.ox.ac.uk/oucl/research/areas/machlearn/PProgol/pprogol.pl`.

[2] `http://archive.ics.uci.edu/ml/datasets/Zoo`.

ILP-*based-outlier-detector*($\mathcal{B}$, $\mathcal{E}$, α, k_{max})

$\quad$ *Out* $\leftarrow \emptyset$
$\quad$ *Anom* $\leftarrow \emptyset$
$\quad$ *Irr* $\leftarrow \emptyset$
$\quad$ find the hypothesis $\mathcal{H}_0 \leftarrow \mathcal{H}^{\mathcal{E}}_{\mathcal{B}}$
$\quad$ find the hypothesis $\overline{\mathcal{H}}_0 \leftarrow \mathcal{H}^{\overline{\mathcal{E}}}_{\mathcal{B}}$
$\quad$ let $Cand_1 \leftarrow \{\{e\} \mid e \in \mathcal{E}\}$
$\quad$ **for** $k = 1$ **to** k_{max} **do**
$\qquad$ $NextCand_k \leftarrow \emptyset$
$\qquad$ **for each** $\mathcal{O}$ **in** $Cand_k$ **do**
$\qquad\quad$ **if** $\mathcal{O} \subseteq \mathcal{E}^+$ **then**
$\qquad\qquad$ let $\langle \mathcal{H}^+_s, \mathcal{H}^+_e, g^+ \rangle \leftarrow Gain^+(\mathcal{H}_0, \mathcal{B}, \mathcal{E}, \mathcal{O})$
$\qquad\qquad$ let $\langle \mathcal{H}^-_s, \mathcal{H}^-_e, g^- \rangle \leftarrow Gain^-(\overline{\mathcal{H}}_0, \mathcal{B}, \overline{\mathcal{E}}, \overline{\mathcal{O}})$
$\qquad\quad$ **else**
$\qquad\qquad$ let $\langle \mathcal{H}^-_s, \mathcal{H}^-_e, g^- \rangle \leftarrow Gain^-(\mathcal{H}_0, \mathcal{B}, \mathcal{E}, \mathcal{O})$
$\qquad\qquad$ let $\langle \mathcal{H}^+_s, \mathcal{H}^+_e, g^+ \rangle \leftarrow Gain^+(\overline{\mathcal{H}}_0, \mathcal{B}, \overline{\mathcal{E}}, \overline{\mathcal{O}})$
$\qquad\quad$ **if** $g^+ \geq \alpha$ **and** $g^- \geq \alpha$ **then**
$\qquad\qquad$ *Out* $\leftarrow$ *Out* $\cup \{\langle \mathcal{O}, (\mathcal{H}^+_s, \mathcal{H}^+_e), (\mathcal{H}^-_s, \mathcal{H}^-_e) \rangle\}$
$\qquad\quad$ **else if** $g^+ \geq \alpha$ **then**
$\qquad\qquad$ *Irr* $\leftarrow$ *Irr* $\cup \{\langle \mathcal{O}, (\mathcal{H}^+_s, \mathcal{H}^+_e) \rangle\}$
$\qquad\quad$ **else if** $g^- \geq \alpha$ **then**
$\qquad\qquad$ *Anom* $\leftarrow$ *Anom* $\cup \{\langle \mathcal{O}, (\mathcal{H}^-_s, \mathcal{H}^-_e) \rangle\}$
$\qquad\quad$ **else**
$\qquad\qquad$ let $NextCand_k \leftarrow NextCand_k \cup \{\mathcal{O}\}$
$\qquad$ let $Cand_{k+1} \leftarrow GenerateCand(C_k)$
$\quad$ **return** $\langle Out, Irr, Anom \rangle$

$Gain^+$($\mathcal{H}$, $\mathcal{B}$, $\mathcal{E}$, $\mathcal{O}$)

$\quad$ let $\mathcal{H}_1 \leftarrow \mathcal{H}(\mathcal{O})$ *// starting theory*
$\quad$ find the hypothesis $\mathcal{H}_2 \leftarrow \mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}$
$\quad$ let $\widehat{\mathcal{E}} \leftarrow \mathcal{E}^+ \setminus \mathcal{O}$
$\quad$ let $\widehat{\mathcal{O}} \leftarrow covers_{\widehat{\mathcal{E}}}(\mathcal{H}) \setminus covers_{\widehat{\mathcal{E}}}(\mathcal{H} \setminus \mathcal{H}_1)$
$\quad$ let $\mathcal{H}_3 \leftarrow \mathcal{H}_2(\widehat{\mathcal{O}})$ *// ending theory*
$\quad$ **return** $\langle \mathcal{H}_1, \mathcal{H}_3, gain_{\widehat{\mathcal{E}}}(\mathcal{H}_3, \mathcal{H}_1) \rangle$

$Gain^-$($\mathcal{H}$, $\mathcal{B}$, $\mathcal{E}$, $\mathcal{O}$)

$\quad$ find the hypothesis $\mathcal{H}_2 \leftarrow \mathcal{H}^{\mathcal{E}\setminus\mathcal{O}}_{\mathcal{B}}$
$\quad$ let $\mathcal{H}_3 \leftarrow \mathcal{H}_2(\mathcal{O})$ *// ending theory*
$\quad$ let $\widehat{\mathcal{E}} \leftarrow \mathcal{E}^+$
$\quad$ let $\widehat{\mathcal{O}} \leftarrow covers_{\widehat{\mathcal{E}}}(\mathcal{H}_2) \setminus covers_{\widehat{\mathcal{E}}}(\mathcal{H}_2 \setminus \mathcal{H}_3)$
$\quad$ let $\mathcal{H}_1 \leftarrow \mathcal{H}(\widehat{\mathcal{O}})$ *// starting theory*
$\quad$ **return** $\langle \mathcal{H}_1, \mathcal{H}_3, gain_{\widehat{\mathcal{E}}}(\mathcal{H}_3, \mathcal{H}_1) \rangle$

Figure 2. The ILP-*based-outlier-detector* algorithm.

predicate the binary predicate *class*. The set of positive examples consists of one hundred facts. The set of negative examples, consisting of six hundreds facts, has been obtained by associating each animal with the classes it does not belong to.

We executed the algorithm with $\alpha = 0.05$ and $k_{max} = 1$. Besides the facts in the induced hypothesis and the induced dual hypothesis, which are classified as irregular sets, the algorithm reported the following abnormal sets: $\mathcal{O}_1 = \{class(amphibian, newt)\}$ as positive outlier, $\mathcal{O}_2 = \{class(insect, ladybird)\}$ as positive anomalous, and $\mathcal{O}_3 = \{class(mammal, platypus)\}$ as positive anomalous. Next we briefly comment on the knowledge discovered by the method.

Newt: The positive outlier set $\mathcal{O}_1 = \{class(amphibian, newt)\}$ is a fact in the direct theory, while it has as dual explanation the dual starting theory $\overrightarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}}_1)$:

$not\_class(amphibian, X) \leftarrow catsize(X) \;\; \langle 44 \rangle$
$not\_class(amphibian, X) \leftarrow legs(2, X) \;\; \langle 27 \rangle$
$not\_class(amphibian, tuatara) \;\; \langle 1 \rangle$
$not\_class(amphibian, scorpion) \;\; \langle 1 \rangle$,

and the dual ending theory $\overleftarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}}_1)$:

$not\_class(amphibian, X) \leftarrow tail(X) \;\; \langle 74 \rangle$,

with gain 0.11. From this explanation, it is clear that the newt is the only amphibian in the example set having the tail. As a matter of fact, it is the only amphibian of the *Caudata order* belonging to the set of examples, while all the other amphibians in the example set belong to the *Anura order*, which is characterized by the absence of tail.

Ladybird: The positive anomalous set $\mathcal{O}_2 = \{class(insect, ladybird)\}$ has as dual explanation the dual starting theory $\overrightarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}}_2)$:

$not\_class(insect, X) \leftarrow aquatic(X) \;\; \langle 35 \rangle$
$not\_class(insect, scorpion) \;\; \langle 1 \rangle$,

and the dual ending theory $\overleftarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}}_2)$:

$not\_class(insect, X) \leftarrow predator(X) \;\; \langle 54 \rangle$,

with gain 0.08. Indeed, among the insects present in the examples, that are the *flea*, *gnat*, *honeybee*, *housefly*, *moth*, *termite*, and *wasp*, the *ladybird* is the only predator.

Platypus: The positive anomalous set $\mathcal{O}_3 = \{class(mammal, platypus))\}$ has as dual explanation the dual starting theory $\overrightarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}}_3)$:

$not\_class(mammal, X) \leftarrow legs(6, X) \;\; \langle 10 \rangle$
$not\_class(mammal, X) \leftarrow feathers(X) \;\; \langle 20 \rangle$
$not\_class(mammal, X) \leftarrow eggs(X), toothed(X) \;\; \langle 19 \rangle$
$not\_class(mammal, X) \leftarrow eggs(X), legs(0, X) \;\; \langle 19 \rangle$
$not\_class(mammal, starfish) \;\; \langle 1 \rangle$
$not\_class(mammal, tortoise) \;\; \langle 1 \rangle$
$not\_class(mammal, crab) \;\; \langle 1 \rangle$
$not\_class(mammal, octopus) \;\; \langle 1 \rangle$,

and the dual ending theory $\overleftarrow{\mathcal{H}}^{\overline{\mathcal{E}}}_{\mathcal{B}}(\overline{\mathcal{O}}_3)$:

$not\_class(mammal, X) \leftarrow eggs(X) \;\; \langle 57 \rangle$,

with gain 0.09. The platypus is a well-known strange mammal, since the female lays eggs, although the newly hatched young are fed by the mother's milk.

B. Student Loan

Here we consider the *Student Loan* relational domain from the UCI Machine Learning Repository[3]. The target unary predicate *no_payment_due*(*Person*) is true for those people who are not required to repay a student loan. Auxiliary relations can be used to fully discriminate positive from negative instances. We executed the algorithm with $\alpha = 0.05$ and $k_{max} = 1$, with 78 positive examples and 34 negative examples, consisting of the students whose identifier number starts with 1. Next we briefly commment on the negative outlier set $\mathcal{O}_1 = \{$*no_payment_due*(*student149*)$\}$. It has as direct explanation the direct starting theory $\overrightarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O}_1)$:

[3] `http://archive.ics.uci.edu/ml`.

$no\_payment\_due(X) \leftarrow enrolled(X, Y, 10) \quad \langle 12 \rangle$

and the direct ending theory $\overleftarrow{\mathcal{H}}^{\mathcal{E}}_{\mathcal{B}}(\mathcal{O}_1)$:

$no\_payment\_due(X) \leftarrow male(X), enrolled(X, Y, 10) \quad \langle 6 \rangle$

$no\_payment\_due(student165) \quad \langle 1 \rangle$

$no\_payment\_due(student112) \quad \langle 1 \rangle$

$no\_payment\_due(student196) \quad \langle 1 \rangle$,

with gain 0.13, while it is a fact in the dual theory. The *student149* is strange, since it is the only enrolled in ten units which is required to repay a student loan.

C. Comparison with noise-handling mechanisms

In order to compare our approach with noise-handling mechanisms implemented in ILP learning systems, we ran P-Progol by setting the parameter *noise* (representing an upper bound on the number of negative examples allowed to be covered by an acceptable clause) to 5 and the parameter *minpos* (representing a lower bound on the number of positive examples to be covered by an acceptable clause) to 2. According to the induced hypothesis the positive examples associated with $clam$, $crab$, $seawasp$, $slug$, $starfish$, $worm$, $pitviper$, $seasnake$, $slowworm$, $tortoise$, and $tuatara$ are classified as noise. Moreover, the unique negative example classified as noise is *class(amphibian,tuatara)*.

The positive examples classified as noise, were reported as irregular singleton sets by our method since they do not comply with the set of examples, being facts in the complete and consistent direct hypothesis. As for the negative example *class(amphibian,tuatara)*, it was also reported as an irregular singleton set by our method, since it does not comply with the dual set of examples, being a fact in the complete and consistent dual hypothesis. As for its compliance with the set of examples, the starting theory associated with the set $\{class(amphibian, tuatara)\}$ is $\{class(amphibian, toad)$, $class(amphibian, frog)$, $class(amphibian, newt)\}$, while the corresponding ending theory is:

$class(amphibian, X) \leftarrow eggs(X), toothed(X), legs(4, X) \ \langle 3 \rangle$

with associated gain 0.02. Hence, for $\alpha \leq 0.02$ it would be recognized as an outlier set by our method. We note that our method found the outlier set $\{class(amphibian, newt)\}$, and the anomalous sets $\{class(insect, ladybird)\}$ and $\{class(mammal, platypus)\}$, for which there is no counterpart in the noise returned by P-Progol.

D. Comparison with DB-Outliers

We compared the method here presented with the distance-based outlier definition [7], by searching the distance-based outliers in the original Zoo data set. We used as outlier score the sum of the distances to the k-nearest neighbors [8], the Hamming function as distance measure, and we set both k (the number of nearest neighbors to consider) and n (the number of outliers to return) to 5 (that corresponds to the 5% of the positive examples). The following table reports the top-n distance-based outliers:

Outlier (Score)	Nearest Neighbors
1. *scorpion* (25)	$worm, slug, pitviper, clam, crab$
2. *seasnake* (19)	$pitviper, stingray, chub, herring, bass$
3. *tortoise* (18)	$tuatara, ostrich, rhea, slowworm, wren$
4. *toad* (17)	$frog, newt, tuatara, worm, crab$
5. *pitviper* (15)	$slowworm, tuatara, seasnake, newt, kiwi$

As for the comparison with our method, the singleton sets of positive examples associated with the distance-based outliers 2, 3, 4, and 5, are returned as irregular sets by our method, while, as for the outlier 1, some singleton sets of negative examples involving it are returned as irregular sets. The example $class(invertebrate, scorpion)$ is not recognized as abnormal, since it shares with the octopus the property of having eight legs. It is clear, that the abnormal instances returned by the distance-based method are of different nature with respect to those returned by the approach here introduced. In particular, distance-based outliers are likely to correspond to irregular instances, since, intuitively, they are objects whose attribute-value pattern is shared less.

VI. Conclusions

In this work, we presented a novel approach to detect outliers which exploits domain knowledge. We introduced the novel definition of ILP-based outlier. We discussed some examples of knowledge mined, and compared it with alternative approaches. As future work, we are intended to investigate the computational complexity of the problem and to design more efficient mining algorithms for dealing with large data sets.

References

[1] F. Angiulli, R. Ben-Eliyahu-Zohary, and L. Palopoli, “Outlier detection using default reasoning,” *Artif. Intell.*, vol. 172, no. 16-17, pp. 1837–1872, 2008.

[2] F. Angiulli, G. Greco, and L. Palopoli, “Outlier detection by logic programming,” *ACM Trans. Comput. Log.*, vol. 9, no. 1, 2007.

[3] N. Lavrač and S. Džeroski, *Inductive Logic Programming, Techniques and Applications*. Ellis Horwood, 1994.

[4] N. Lavrač, S. Džeroski, and I. Bratko, “Handling imperfect data in inductive logic programming,” in *Advances in Inductive Logic Programming*, L. D. Raedt, Ed., 1996, pp. 48–64.

[5] J. W. Lloyd, *Foundations of Logic Programming*. Berlin: Springer-Verlag, 1987.

[6] S. Muggleton, “Inverse entailment and Progol,” *New Generation Computing, Special issue on Inductive Logic Programming*, vol. 13, no. 3-4, pp. 245–286, 1995.

[7] E. Knorr and R. Ng, “Algorithms for mining distance-based outliers in large datasets,” in *Proc. of the Int. Conf. on Very Large Data Bases (VLDB)*, 1998, pp. 392–403.

[8] F. Angiulli and C. Pizzuti, “Outlier mining in large high-dimensional data sets.” *IEEE Trans. Knowl. Data Eng.*, pp. 203–215, 2005.