

DLGP: An extended Datalog Syntax for Existential Rules and Datalog \pm *Version 2.0*

GraphIK Team,* LIRMM / Inria, Montpellier

June 9, 2015

Abstract

This document specifies the version 2.0 of *dlgp*, a textual format for the existential rule / Datalog \pm framework. This format is meant to be an exchange format at once human-friendly, concise and easy to parse. It can be seen as an extension of the commonly used format for plain Datalog. It is called “dlgp” for “Datalog Plus”. A file may contain four kinds of knowledge elements: facts, existential rules, negative constraints and conjunctive queries.

The main improvement in this version is the introduction of Web notions (IRI and literal, following Turtle format) to ensure compatibility with Semantic Web languages.

1 Introduction

The *dlgp* format encodes existential rules (and other constructs that can be seen as special kinds of existential rules: facts, negative constraints and conjunctive queries). A basic logical notion is that of an *atom*, which is composed of a *predicate* (or relation name) and arguments called *terms*. Terms can be variables or constants. Predicates can be of any arity greater or equal to one.

In their simplest form, predicates and terms are encoded by string identifiers built with *letters* from the Latin alphabet, *digits* and the *underscore* character (`_`). As usually in Datalog, variables begin with an uppercase letter, while constants and predicates begin with a lowercase letter. To make the *dlgp* format compatible with data from the Semantic Web, more elaborate kinds of identifiers are mandatory. This new version introduces the notions of IRI and literal, according to Turtle format (<http://www.w3.org/TR/turtle/>). Specifically:

*Jean-François Baget, Alain Gutierrez, Michel Leclère, Marie-Laure Mugnier, Swan Rocher, Clément Sipieter.

- a predicate or a constant can be denoted by an *IRI* in Turtle format, which can be written in three forms: a relative IRI, an absolute IRI or a prefixed name. Relative and absolute IRIs are strings of the form `<iri>` (see the symbol **IRIREF** in Turtle grammar). A prefixed name is a string prefixed by a namespace (see the symbol **PrefixName** in Turtle grammar); moreover, the classical Datalog format (a string beginning with a lowercase letter) is allowed and understood as a relative IRI;
- a constant can also be a literal (see the symbol **literal** in Turtle grammar).

For instance, here are different ways of writing a predicate:

- `pred` (classical logic programming notation);
- `<pred>` (a relative IRI, whose base is specified elsewhere);
- `<Pred>` (a relative IRI as well, starting with an uppercase letter);
- `<http://exemple.org/pred>` (an absolute IRI)
- `ex:pred` (prefixed name); in this case the IRI associated with the prefix `ex` has to be introduced by a specific annotation in the file header, e.g., `@prefix ex: <http://exemple.org/>`

All these forms are usable to encode a constant. Moreover, constants can be described by *literals*, e.g., `-5.1`, `true`, `"constant"`, or any other Turtle literal. Note that the tokens `true` and `false` are interpreted as Boolean literals and not as IRIs.

The special predicate `=` encodes equality between terms (e.g., `X = Y`).

Sets of such atoms are logically interpreted as *conjunctions* of atoms. The following four kinds of knowledge elements are built upon sets of atoms:

- *facts* : a fact is a set of atoms interpreted as an existentially closed conjunction, i.e., of the form $\exists X \mathcal{F}[X]$, where X denote the set of variables occurring in \mathcal{F} ; a fact can also be seen as an existential rule with an empty body;
- *(pure) existential rules* : an existential rule is an ordered pair of atom sets $(\mathcal{B}, \mathcal{H})$ interpreted as $\forall X (\exists Y \mathcal{B}[X, Y] \rightarrow \exists Z \mathcal{H}[X, Z])$, or, equivalently, as $\forall X \forall Y (\mathcal{B}[X, Y] \rightarrow \exists Z \mathcal{H}[X, Z])$, where X , Y and Z denote sets of variables (\mathcal{B} is built on X and Y , while \mathcal{H} is built on X and Z , hence X denotes the set of variables shared by \mathcal{B} and \mathcal{H});
- *negative constraints* : a negative constraint is a set of atoms interpreted as the negation of the corresponding fact, i.e., $\neg(\exists X \mathcal{C}[X])$; equivalently, it can be seen as a rule with a head restricted to the always-false symbol \perp : $\forall X (\mathcal{C}[X] \rightarrow \perp)$;

- *conjunctive queries* : a conjunctive query is a set of atoms provided with special variables (called the answer variables), interpreted as an existentially quantified conjunction where answer variables are kept free. It can also be seen as a rule of the form $\forall X \forall Y (\mathcal{B}[X, Y] \rightarrow ans(X))$, where *ans* is a special predicate and *X* is the set of answer variables.

A dlgp document is any sequence of such elements. The file name has extension `.dlp` or `.dlgp`. Characters are assumed to be encoded in UTF-8. Analysis directives are introduced by the symbol `@` and comments by the symbol `%`.

2 Syntax of the dlgp format

In the following the syntax is specified by a grammar in BNF style: non-terminal symbols are enclosed in angle brackets `<>`, terminal symbols are in bold font, the `|` symbol indicates a choice, parts enclosed in square brackets `[]` are optional, *choice₁..choice_n* indicates a choice within an interval; parts enclosed in braces can be repeated from 0 to *n* times (*{repeated-pattern}**) or from 1 to *n* times (*{repeated-pattern-at-least-once}+*).

2.1 Comments

Comments are introduced by the symbol `%` outside Turtle tokens (i.e., outside `<IRIREF>` and `<literal>`, which may themselves contain the symbol `%`). A comment ends at the end of the same line or at the end of the file. Moreover, comments introduced by `%%` can be interpreted in a specific way. Our parser generates an event when such a specific comment is read, which can be exploited by event listeners.

2.2 Parsing Information

Please refer to Turtle syntax for building an absolute IRI from a relative IRI (and `@base` directives) or from a prefixed name (and `@prefix` directives). Note that the directive base may occur at most once. Similarly, it is not possible to successively assign several IRIREF to the same prefix (`@prefix` directive). These two directives may occur only in the head of the file.

A `<l-ident>` token is seen as a relative IRI, and the corresponding absolute IRI is obtained by adding the IRI of the `@base` directive in front of the `<l-ident>` token.

2.3 Elements used to define tokens

<uppercase-letter> ::= **A..Z**
<lowercase-letter> ::= **a..z**
<digit> ::= **0..9**
<underscore> ::= **-**
<letter> ::= <uppercase-letter> | <lowercase-letter>
<simple-char> ::= <letter> | <digit> | <underscore>
<PN_CHARS> *see turtle grammar*
<space> ::= **#x20** /* #x20 = space character */

2.4 Tokens

<u-ident> ::= <uppercase-letter> {<simple-char>}*
<l-ident> ::= <lowercase-letter> {<simple-char>}*
<label> ::= {<PN_CHARS> | <space>}*

2.5 Global Grammar

```

<document> ::= <header> <body>
<header> ::= [<base>] {<prefix>}* [<top>] [<una>]
<base> ::= @base <IRIREF>
<prefix> ::= @prefix <PNAME_NS> <IRIREF>
<top> ::= @top <l-ident> |
        @top <IRIREF>
<una> ::= @una
<body> ::= {<statement>}* |
        {<section>}*
<section> ::= @facts {<fact>}* |
        @rules {<rule>}* |
        @constraints {<constraint>}* |
        @queries {<query>}*
<statement> ::= <fact> | <rule> | <constraint> | <query>
<fact> ::= [ [<label>] ] <conjunction>.
<rule> ::= [ [<label>] ] <head> :- <body>.
<head> ::= <conjunction>
<body> ::= <conjunction>
<constraint> ::= [ [<label>] ] ! :- <conjunction>.
<query> ::= [ [<label>] ] ? [( <term-list> ) ] :- <conjunction>.
<conjunction> ::= <atom> { , <atom> }*
<atom> ::= <std-atom> | <equality>
<equality> ::= <term> = <term>
<std-atom> ::= <predicate>( <term-list> )
<term-list> ::= <term> { , <term> }*
<term> ::= <variable> | <constant>
<predicate> ::= <l-ident> | <IRIREF> | <PrefixedName>
<variable> ::= <u-ident>
<constant> ::= <l-ident> | <IRIREF> | <PrefixedName> | <literal>

```

The symbol @ is used to introduce several kinds of annotations:

- In the header:
 - @base is a Turtle directive which defines a base to complement relative IRIs. This annotation is not mandatory even if relative IRIs are used in the document; a base IRI is then defined by the context of the application;
 - @prefix is a Turtle directive that defines a pair *prefix / namespace identifier* used to build an IRI from a prefixed name (the prefix is a string terminated by a colon and the namespace identifier is an <IRIREF>). The IRI is obtained by replacing the prefix in the prefixed name by the IRI identifying the namespace

associated with that prefix. All the prefixes occurring in the document must be defined by a `@prefix` annotation. All the prefixes occurring in `@prefix` annotations must be pairwise distinct;

- `@top` defines a predicate whose semantics is \top (i.e., the knowledge base satisfies the axiom $\forall x \top(x)$, where \top is replaced by this predicate);
 - `@una` states that this knowledge base makes the Unique Name Assumption, i.e., all constants are assumed to refer to distinct objects.
- In the body:
 - the directives `@facts`, `@rules`, `@constraints`, `@queries` can be used to declare the nature of the elements that follow (respectively: facts, rules, constraints and queries). Such a directive defines a section that spans to the next annotation or to the end of the file. The interest of this kind of directive is to inform sooner the parser of the nature of the next statements. For instance, without these directives, it is not possible to distinguish between a fact and a rule before the sequence of atoms composing the fact or the rule head has been completely analyzed. However, it may be interesting to know that a fact is coming to process it more efficiently, specially if large facts can be encountered. However, such a directive is only an information to the parser, which can choose to ignore it.
 - each element can be optionally preceded by a *label* enclosed in square brackets. This name is any string composed of `<PN_CHARS>` in Turtle format as well as the space character.

To encode other kinds of information about the knowledge base, it is recommended to use specific comments introduced by `%%`.

Note on the scope of variable identifiers. While the scope of a constant or a predicate identifier is the whole document, the scope of a variable is local to a `<statement>`. Thus two different facts, rules or constraints actually do not share any variable (more precisely, variables with the same name in different statements are each bound by their own quantifier).

Example:

`p(X,a), q(X,Y).`

`q(X,b).`

is logically interpreted as $\exists X \exists Y (p(X, a) \wedge q(X, Y)) \wedge \exists X q(X, b)$

while:

`p(X,a), q(X,Y), q(X,b).`

is logically interpreted as $\exists X \exists Y (p(X, a) \wedge q(X, Y) \wedge q(X, b))$.

3 Differences with dlgp 1.0

The only incompatibility with Version 1.0 is the impossibility to use a quoted string to denote a predicate. Now, predicates have to be `<l-ident>`, `<IRIREF>` or `<PrefixedName>`. Constants have been extended: they may also be `<IRIREF>`, `<PrefixedName>` or any `<literal>` (not only integer, float et quoted string).

4 Examples (dlgp 1.0, compatible with dlgp 2.0)

A syntactically correct but not human-friendly file:

```
[f1] p(X),r(X,b),q(b).
[f2] t(X,a), s(a,z).
[c1] ! :- r(X,X) .
% This is a comment
[q1]?(X) :- p(X),r(X,Z), t(a,Z).
[r1]r(X,Y):-p(X) , t(X,Z).
[constraint_2]! :- X = X.
[f3] q(a),r(a,X), % This is another comment
t(a,z).
s(X,Y),s(Y,Z)
:-
q(X),t(X,Z).
s(Z):-a=b,X=Y.
t(X,a), r(Y,z).
[rA_1]p(X):-q(X).
!:-p(X),q(X).
? :- p(X).
```

The same logical content organized in a more readable way:

```
% facts
[f1] p(X), r(X,b), q(b).
[f2] t(X,a), s(a,z).
[f3] q(a), r(a,X), t(a,z).
t(X,a), r(Y,z).
% constraints
[c1] ! :- r(X,X).
[constraint_2] ! :- X=X.
! :- p(X), q(X).
% rules
[r1] r(X,Y) :- p(X), t(X,Z).
s(X,Y), s(Y,Z) :- q(X),t(X,Z).
s(Z) :- a=b, X=Y.
```

```

[rA_1] p(X) :- q(X).
% queries
[q1] ? (X) :- p(X), r(X,Z), t(a,Z).
? :- p(X).

```

Again the same logical content while using annotations (directives) instead of simple comments to distinguish between the different kinds of knowledge:

```

@Facts
[f1] p(X), r(X,b), q(b).
[f2] t(X,a), s(a,z).
[f3] q(a), r(a,X), t(a,z).
t(X,a), r(Y,z).
@Constraints
[c1] ! :- r(X,X).
[contrainte_2] ! :- X=X.
! :- p(X), q(X).
@Rules
[r1] r(X,Y) :- p(X), t(X,Z).
s(X,Y), s(Y,Z) :- q(X),t(X,Z).
s(Z) :- a=b, X=Y.
[rA_1] p(X) :- q(X).
@Queries
[q1] ? (X) :- p(X), r(X,Z), t(a,Z).
? :- p(X).

```

5 Examples (dlgp 2.0)

This example illustrates the use of IRIs and literals. Note also that spaces are now allowed in labels.

```

@base <http://www.example.org/>
@prefix ex: <http://www.example.org/>
@prefix inria-team: <https://team.inria.fr/>
@prefix xsd: <http://www.w3.org/2001/XMLSchema#>

@facts

% use of @base
[f 1] <Pred>(1.5).

% use of @prefix
[f 2] ex:Pred("1.5"^^xsd:decimal).

```



```
% absolute IRIs
[f 3] <http://www.example.org/Pred>
      ("1.5"^^<http://www.w3.org/2001/XMLSchema#decimal>).

% use of @base for the predicate and @prefix for the argument
[f 4] team(inria-team:graphik).
```

The three first facts have the same interpretation.