RQL: An SQL-like Query Language for Discovering Meaningful Rules

Brice Chardin, Emmanuel Coquery, Marie Pailloux, Jean-Marc Petit

► To cite this version:

Brice Chardin, Emmanuel Coquery, Marie Pailloux, Jean-Marc Petit. RQL: An SQL-like Query Language for Discovering Meaningful Rules. ICDM 2014, Dec 2014, Shenzhen, China. pp.1203-1206, 10.1109/ICDMW.2014.50 . hal-01301091

HAL Id: hal-01301091

<https://hal.science/hal-01301091>

Submitted on 18 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

RQL: An SQL-like Query Language for Discovering Meaningful Rules

Brice Chardin*, Emmanuel Coquery†, Marie Pailloux‡ and Jean-Marc Petit§

*LIAS, ISAE-ENSMA, France

†Université Lyon 1, France

‡Université Blaise Pascal, France

§INSA Lyon, France

Abstract—The Rule Query Language (RQL) is an SQL-like pattern mining language that extends and generalizes functional dependencies to new and unexpected rules. It brings to the data analysts’ desktop a convenient tool to discover logical implications between attributes of the database. Such implications may reveal data quality problems or surprising correlations between attributes over some part of the database. The computation of RQL queries is based on a query rewriting technique that pushes as much processing as possible to the underlying DBMS. This contribution is an attempt to bridge the gap between pattern mining and databases and facilitates the use of data mining techniques by SQL-aware analysts and students.

I. INTRODUCTION

Pattern mining can be seen as an automated part of data exploration. For instance, functional dependencies or conditional functional dependencies are definitely useful to understand the data and to identify data quality problems [1]. However, pattern mining techniques are rarely usable directly by data analysts. Most of the time, they have to perform some data pre-processing between different systems and formats. The pattern mining codes themselves often require to be compiled from some specific programming languages. All these steps are out of reach of many data analysts, rendering round-trip engineering into a nightmare. Automated rule generation can also flood the analyst with huge amounts of patterns, and make it difficult to extract useful information. Other techniques have to be provided to interact with the data and give useful feedback to the analysts.

Demo contribution To improve pattern mining usability for data exploration, we introduce a Rule Query Language (RQL) that allows SQL-aware analysts to use pattern mining techniques with an interactive, user-friendly interface. In this demonstration, we show the usability of this web interface from the point of view of a data analyst. We focus on the expressive power of RQL through various examples, showing how easy it is to devise new and surprising rules with a very simple language derived from SQL. We also introduce how the data analysts can interact with the system through RQL queries and counterexamples taken from the database. During the demo, participants will be invited to formulate their own queries on predefined databases to discover attribute relationships through generated rules and counterexamples.

Figure 1 gives a preview of the web interface for RQL, made available¹ for research and educational purposes. This

Fig. 2. RQL queries processing overview

interface provides a unified access to the user’s data and pattern mining techniques using declarative languages: SQL and RQL.

From previous works [2], [3], [4], RQL is compliant with Armstrong’s axioms, i.e. the language generalizes functional dependencies to a new class of dependencies based on logical implications (if ... then statements). To the best of our knowledge, this class of dependencies has not been studied before. We have proven that these dependencies can be efficiently computed from a database using a two step technique. First, a non trivial SQL subquery is generated to compute a *base* of the associated closure system – a base is also called a context in formal concept analysis terminology [5]. Then, a state of the art algorithm [6] is used to generate a canonical cover of rules from this base. This approach allows RQL to benefit from DBMSs’ query optimization to access the data, and keep the data where they are.

Figure 2 gives an overview of this architecture with respect to RQL query processing. As for SQL queries, the application simply forwards them to the underlying DBMS, which makes the transition between SQL and RQL transparent to the user. The ultimate goal of this work is to integrate pattern mining techniques into core DBMS technologies [7].

Related works Defining specific languages for pattern mining is a long standing goal [8], for example using constraint programming techniques [9]. Nevertheless, we argue that pattern mining languages should benefit from direct extensions of the SQL language, since data are often stored in DBMSs. Several other practical propositions follow this principle and interact more directly with DBMSs query engines [10], [11], [12].

¹<http://rql.insa-lyon.fr>

Fig. 1. Web interface for RQL

II. THE RQL QUERY LANGUAGE

To make things concrete, let us consider the running example given in Figure 3 with the *EMP* table. The attribute *Educlevel* represents the number of years of formal education, *Sal* the yearly salary, *Bonus* the yearly bonus and *Comm* the yearly commission. The meaning of other attributes is straightforward.

To begin with, let us extract functional dependencies (FD) from the relation *Emp*. Recall that a FD $X \rightarrow Y$ holds in r if for all tuples $t_1, t_2 \in r$, and for all attribute $A \in X$ such that $t_1[A] = t_2[A]$ then for all $A \in Y$, $t_1[A] = t_2[A]$. With RQL, FDs are expressed in a similar way.

Example 1. Q_1 discovers FDs from *Emp* over a subset of attributes.

```
Q1: FINDRULES
OVER Empno, Lastname, Workdept, Job,
     Sex, Bonus, Mgrno
SCOPE t1, t2 Emp
CONDITION ON $A IS t1.$A = t2.$A
```

Note how the `CONDITION` clause matches the previous logical implication. We have also restricted FDs discovery to a subset of seven attributes in the `OVER` clause. In this example, a canonical cover of FDs that hold in *Emp* is generated (composed of twenty-four FDs), including FDs such as *Empno* $\rightarrow$ *Lastname* or *Workdept* $\rightarrow$ *Job*.

Overall, a RQL query has the following general form:

```
FINDRULES
OVER [set of attributes: A1, ..., An]
SCOPE [tuple variables: t1, ..., tn]
WHERE [condition on (t1, ..., tn)]
CONDITION ON [attribute variable: $A]
IS [condition on ($A, t1, ..., tn)]
```

The `FINDRULES` keyword identifies a RQL query, which generates rules of the form $X \rightarrow Y$ with X and Y disjoint attribute sets taken from the `OVER` clause. The `SCOPE` clause defines tuple-variables over some tables obtained by classical SQL queries. An optional `WHERE` clause defines relationships between tuple-variables, similar to the SQL `WHERE` clause. The `CONDITION ON $A` clause defines the predicate to be

satisfied by each attribute $\$A$ occurring in the left- and right-hand sides of the rule.

To illustrate the expressiveness of RQL queries, we now provide several examples.

Example 2. Let us consider null values known to be common in real-life databases. With RQL, the data analyst has the opportunity to discover rules between attributes with respect to null values as shown with query Q_2 .

```
Q2: FINDRULES
OVER Empno, Lastname, Workdept, Job,
     Sex, Bonus, Mgrno
SCOPE t1 Emp
CONDITION ON $A IS t1.$A IS NULL
```

The rule $Mgrno \rightarrow Workdept$ holds in *Emp* since each time the attribute *Mgrno* is null in a tuple, then *Workdept* is also null for the same tuple (only employee No. 20 in this example).

Note that the difference between Q_1 and Q_2 naturally lies on the predicate to be evaluated, but also on the number of tuple variables required. The predicate of Q_1 is evaluated on pairs of tuples, while Q_2 considers a single tuple.

Example 3. The following query Q'_1 restricts the scope of Q_1 , leading to the notion of conditional functional dependencies [1]. For example, we consider only employees with a level of qualification above 16.

```
Q'1: FINDRULES
OVER Empno, Lastname, Workdept, Job,
     Sex, Bonus
SCOPE t1, t2 (SELECT * FROM Emp
              WHERE Educlevel > 16)
CONDITION ON $A IS t1.$A = t2.$A
```

Interestingly, $Sex \rightarrow Bonus$ holds with this restriction, meaning that above a certain level of qualification (16), the gender determines the bonus.

Example 4. Query Q''_1 is an approximation of Q_1 for numeric values similar to Metric Functional Dependencies [13], where strict equality is discarded to take into account variations under 10%. For instance, salaries 41250 and 38250 are considered

EMP	Empno	Lastname	Workdept	Job	Educllevel	Sex	Sal	Bonus	Comm	Mgrno
	10	SPEN	C01	FINANCE	18	F	52750	500	4220	20
	20	THOMP	-	MANAGER	18	M	41250	800	3300	-
	30	KWAN	-	FINANCE	20	F	38250	500	3060	10
	50	GEYER	-	MANAGER	16	M	40175	700	3214	20
	60	STERN	D21	SALE	14	M	32250	500	2580	30
	70	PULASKI	D21	SALE	16	F	36170	700	2893	100
	90	HENDER	D21	SALE	17	F	29750	500	2380	10
	100	SPEN	C01	FINANCE	18	M	26150	800	2092	20

Fig. 3. Running example

close (7.5% difference), but not salaries 41250 and 36170 (13.1% difference).

Q'_1 : FINDRULES
OVER Educlevel, Sal, Bonus, Comm
SCOPE t1, t2 Emp
CONDITION ON \$A IS
 $2 * \text{ABS}(t1.\$A - t2.\$A) / (t1.\$A + t2.\$A) < 0.1$

In that case, $Sal \rightarrow Comm$ holds, meaning that employees earning similar salaries receive similar commissions.

We have shown so far query examples related to implications (in formal Concept Analysis) and functional dependencies (in databases). Nevertheless, RQL is not restricted to these types of queries at all and can express many more rules.

Example 5. Assume we are interested in a kind of sequential dependencies [14], i.e. dependencies showing similar behavior of attribute values. Q_3 discovers numerical attributes that vary together (i.e., $X \rightarrow Y$ means that if X increases then Y also increases).

Q_3 : FINDRULES
OVER Educlevel, Sal, Bonus, Comm
SCOPE t1, t2 Emp
CONDITION ON \$A IS $t1.\$A > t2.\A

$Sal \rightarrow Comm$ and $Comm \rightarrow Sal$ hold in *Emp*, which means that a higher salary is equivalent to a higher commission.

Example 6. Continuing the previous example, assume now the analyst wants to focus on male employees (see also Figure 1).

Q'_3 : FINDRULES
OVER Educlevel, Sal, Bonus, Comm
SCOPE t1, t2 (SELECT * FROM Emp
WHERE Sex='M')
CONDITION ON \$A IS $t1.\$A > t2.\A

In that case, $Educlevel \rightarrow Bonus$ also holds, which means that male employees with higher education levels receive higher bonuses.

Example 7. Instead of narrowing the scope of a query, user-defined conditions can bind different tuple variables together with a custom relationship specified in the WHERE clause of the RQL query. For example, Q_4 finds disparities between managers and managees, i.e. rules on attributes for which managers have values greater than or equal to their managees.

Q_4 : FINDRULES
OVER Educlevel, Sal, Bonus, Comm
SCOPE t1, t2 Emp
WHERE $t1.\text{Empno} = t2.\text{Mgrno}$
CONDITION ON \$A IS $t1.\$A \geq t2.\A

In this example, $\emptyset \rightarrow Bonus$ holds in *Emp*, which means that managers always earn a bonus greater than or equal to their managees'.

III. FEEDBACK THROUGH COUNTEREXAMPLES

Given a RQL query, the data analyst may also interact with the system to know whether or not a given rule holds. She can provide a rule to the system and two cases arise: either the rule holds and the analyst is notified that the rule is indeed valid; or the rule does not hold which means that at least one counterexample exists and one of them is provided by the system. This notion of counterexample is well known for functional dependencies, and provides very good feedback to the data analyst with her own data. We strongly believe that counterexamples are a great tool to help the analyst understand why a particular rule does not hold, and refine if necessary her analysis in an iterative process.

To illustrate counterexamples, suppose that a data analyst wants to explore her hypothesis that higher salaries and higher education levels yield higher bonuses ($Salary, Educlevel \rightarrow Bonus$), using Q_3 . Figure 4 gives an overview of what RQL provides as a counterexample for this rule, that is, two tuples among which one (employee No. 10) has a higher *Salary* and *Educlevel* than the other (employee No. 50), but not a higher *Bonus*.

With this counterexample as a starting point, and especially the SQL query generated to extract it from the database, the data analyst can quickly switch to SQL to get an idea of why this rule is not verified. For instance, employees that are either female or have a finance job are easily pointed out as having a higher salary and education level, but lower bonuses than others. The data analyst can then refine her RQL query, for example by narrowing the scope of the data, such as in Q'_3 where a higher *Educlevel* by itself implies a higher *Bonus*.

The number of tuples required to provide a counterexample depends on the number of tuple variables. Q_3 or FDs need at least two tuples. But with Q_2 , one tuple (for instance employee No. 30) is enough to prove that the rule $Workdept \rightarrow Mgrno$ does not hold.

Rule verification:

The rule `Sal Educlevel → bonus` is false

Counter-example:

EMPNO	LASTNAME	WORKDEPT	JOB	EDUCLEVEL	SEX	SAL	BONUS	COMM	MGRNO
10	SPEN	C01	FINANCE	18	F	52750	500	4220	20
50	GEYER	null	MANAGER	16	M	40175	700	3214	20

Generated query:

```
1. SELECT t1.*, t2.*
2. FROM Emp t1, Emp t2
3. WHERE (t1.Sal > t2.Sal AND t1.Educlevel > t2.Educlevel)
4. AND CASE WHEN (t1.bonus > t2.bonus) THEN 1 ELSE 0 END = 0
5. AND rownum <= 1
```

Fig. 4. Counterexample with RQL

IV. IMPLEMENTATION AND APPLICATION

The RQL web application has been implemented in Java with the Play Framework [15]. External tools have been used for the most expensive part of the rule generation process, i.e. the enumeration of minimal transversal of hypergraphs [6]. The chosen DBMS is Oracle 11g Release 2.

Importantly, the web interface discloses the SQL code generated by the system to help the user build her own queries. For instance, the query used to retrieve an arbitrary counterexample from the database can be extended so that the analyst identifies more insightful ones.

RQL can be interacted with in two modes: (i) a Sample DB is provided with selected examples to offer a quick way into RQL (ii) a Sandbox allows users to upload and query their own data (currently limited to 3 tables and 200 kB). RQL has been used by 120 undergraduate students to play with functional dependencies and other constraints in a database course at INSA Lyon. Not surprisingly, the notion of counterexamples has been widely used by students. RQL has been appreciated for its ability to bridge the gap between the SQL language and functional dependencies in database design by providing a unified interface for both SQL and RQL queries.

Previous works [4] have highlighted the efficiency of RQL as a two step process, even on large databases.

V. CONCLUSION

RQL is introduced as a web interface to discover rule patterns over relational databases. RQL subsumes SQL statements by providing the opportunity to specify and get results as a set of rules or some counterexamples. The rule mining problem is seen as a query processing problem, for which we have proposed a query rewriting technique allowing the delegation of as much processing as possible to the underlying DBMS engine [4]. RQL allows SQL developers to extract precise information without any specific knowledge in data mining.

REFERENCES

- [1] P. Bohannon, W. Fan, F. Geerts, X. Jia, and A. Kementsietsidis, "Conditional functional dependencies for data cleaning," in *Proceedings of the 23rd International Conference on Data Engineering*, ser. ICDE '07, 2007, pp. 746–755.
- [2] M. Agier, C. Froidevaux, J.-M. Petit, Y. Renaud, and J. Wijsen, "On Armstrong-compliant logical query languages," in *Proceedings of the 4th International Workshop on Logic in Databases*, ser. LID '11, 2011, pp. 33–40.

- [3] M. Agier, J.-M. Petit, and E. Suzuki, "Unifying framework for rule semantics: Application to gene expression data," *Fundamenta Informaticae*, vol. 78, no. 4, pp. 543–559, 2007.
- [4] B. Chardin, E. Coquery, B. Gouriou, M. Pailloux, and J.-M. Petit, "Query Rewriting for Rule Mining in Databases," in *Languages for Data Mining and Machine Learning, in conjunction with ECML/PKDD*, 2013, pp. 35–49.
- [5] B. Ganter and R. Wille, *Formal Concept Analysis*. Springer, 1999.
- [6] K. Murakami and T. Uno, "Efficient algorithms for dualizing large-scale hypergraphs," *CoRR*, vol. 1102.3813, 2011.
- [7] A. Netz, S. Chaudhuri, J. Bernhardt, and U. M. Fayyad, "Integration of data mining with database technology," in *Proceedings of the 26th International Conference on Very Large Data Bases*, ser. VLDB '00, 2000, pp. 719–722.
- [8] H. Blockeel, T. Calders, E. Fromont, B. Goethals, A. Prado, , and C. Robardet, "A practical comparative study of data mining query languages," in *Inductive Databases and Constraint-Based Data Mining*, Springer, Ed., 2010, pp. 59–77.
- [9] T. Guns, S. Nijssen, and L. D. Raedt, "Itemset mining: A constraint programming perspective," *Artif. Intell.*, vol. 175, no. 12–13, pp. 1951–1983, 2011.
- [10] L. Fang and K. LeFevre, "Splash: ad-hoc querying of data and statistical models," in *Proceedings of the 13th International Conference on Extending Database Technology*, ser. EDBT '10, 2010, pp. 275–286.
- [11] C. Ordonez and S. K. Pitchaimalai, "One-pass data mining algorithms in a DBMS with UDFs," in *SIGMOD Conference*, 2011, pp. 1217–1220.
- [12] H. Blockeel, T. Calders, É. Fromont, B. Goethals, A. Prado, and C. Robardet, "An inductive database system based on virtual mining views," *Data Min. Knowl. Discov.*, vol. 24, no. 1, pp. 247–287, 2012.
- [13] N. Koudas, A. Saha, D. Srivastava, and S. Venkatasubramanian, "Metric functional dependencies," in *Proceedings of the 2009 IEEE International Conference on Data Engineering*, ser. ICDE '09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 1275–1278. [Online]. Available: <http://dx.doi.org/10.1109/ICDE.2009.219>
- [14] L. Golab, H. J. Karloff, F. Korn, A. Saha, and D. Srivastava, "Sequential dependencies," *PVLDB*, vol. 2, no. 1, pp. 574–585, 2009.
- [15] "Play framework," <http://www.playframework.com/>.