

Knowledge Refinement via Rule Selection

Phokion G. Kolaitis,^{1,2} Lucian Popa,² Kun Qian²

¹UC Santa Cruz, ²IBM Research – Almaden

kolaitis@ucsc.edu, lpopa@us.ibm.com, qian.kun@ibm.com

Abstract

In several different applications, including data transformation and entity resolution, rules are used to capture aspects of knowledge about the application at hand. Often, a large set of such rules is generated automatically or semi-automatically, and the challenge is to refine the encapsulated knowledge by selecting a subset of rules based on the expected operational behavior of the rules on available data. In this paper, we carry out a systematic complexity-theoretic investigation of the following rule selection problem: given a set of rules specified by Horn formulas, and a pair of an input database and an output database, find a subset of the rules that minimizes the total error, that is, the number of false positive and false negative errors arising from the selected rules. We first establish computational hardness results for the decision problems underlying this minimization problem, as well as upper and lower bounds for its approximability. We then investigate a bi-objective optimization version of the rule selection problem in which both the total error and the size of the selected rules are taken into account. We show that testing for membership in the Pareto front of this bi-objective optimization problem is DP-complete. Finally, we show that a similar DP-completeness result holds for a bi-level optimization version of the rule selection problem, where one minimizes first the total error and then the size.

Rules, typically expressed as Horn formulas, are ubiquitous in several different areas of computer science and artificial intelligence. For example, rules are the basic construct of (function-free) logic programs. In data integration (Lenzerini 2002) and data exchange (Fagin et al. 2005), rules are known as GAV (global-as-view) constraints and are used to specify data transformations between a local (or source) schema and a global (or target) schema. In data mining, rules have many uses, including the specification of contextual preferences (Agrawal, Rantzaou, and Terzi 2006; de Amo et al. 2015). In entity resolution, rules have been used to specify blocking functions (Bilenko, Kamath, and Mooney 2006) and entity resolution algorithms (Qian, Popa, and Sen 2017).

Often, a large set of rules is generated automatically or semi-automatically, and the challenge is to refine the encapsulated knowledge by selecting a subset of rules based on the expected operational behavior of the rules on available

data. Rule selection arises naturally in all aforementioned contexts and, in fact, in most contexts involving reasoning about data. Here, we present an example motivated by a real-life application in which we are building a knowledge base of experts in the medical domain, based on public data, and where entity resolution is one of the crucial first steps.

In entity resolution, the aim is to identify references of the same real-world entity across multiple records or datasets. Consider the scenario depicted in Figure 1, where the aim is to identify occurrences of the same author across research publications from PubMed¹. This entity resolution task

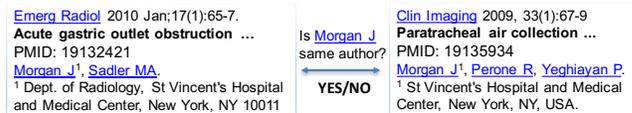


Figure 1: An entity resolution task

can be modeled using a source schema that includes a relation `Author` and a link schema that consists of a relation `SameAuthor`. Sample facts (records) over the source and the link relations are given in Figure 2. As in the frameworks of Markov Logic Networks (Singla and Domingos 2006), Dedupalog (Arasu, Re, and Suciu 2009), and declarative entity linking (Burdick et al. 2016), explicit link relations are used to represent entity resolution inferences. In particular, the `SameAuthor` fact in Figure 2 represents that an inference was made to establish that the author in position 1 of publication with `pmid 19132421` is the same as the author in position 1 of publication with `pmid 19135934`.

For a given entity resolution task, there is typically a large set of candidate rules that may apply on the input data to form matches among the entities. For our concrete scenario, Figure 3 gives a sample of candidate matching rules. These rules involve the alignment of relevant attributes (e.g., `lastname` with `lastname`, `affiliation` with `affiliation`) and the subsequent application of similarity predicates, filters, and thresholds. The challenge is to find a subset of the candidate rules with the “right” combinations of predicates and thresholds that will lead to high

¹<https://www.ncbi.nlm.nih.gov/pubmed>

Author (lastname, firstname, affiliation, pmid, auth_pos, title, coauthors)
 SameAuthor (pmid1, pos1, pmid2, pos2)

Author ("Morgan", "J", "Dept of Radiology ...", 19132421, 1, "Acute ...", "Morgan J; Sadler MA")
 SameAuthor (19132421, 1, 19135934, 1)

Figure 2: Schemas and facts

	Rule 1	Author (lastname, firstname, affiliation, pmid, auth_pos, title, coauthors) \wedge Author (lastname, firstname, affiliation', pmid', auth_pos', title', coauthors') \wedge commonCoAuthors (coauthors, coauthors', 2) \rightarrow SameAuthor (pmid, auth_pos, pmid', auth_pos')
✓	Rule 2	Author (lastname, firstname, affiliation, pmid, auth_pos, title, coauthors) \wedge Author (lastname, firstname, affiliation', pmid', auth_pos', title', coauthors') \wedge commonCoAuthors (coauthors, coauthors', 3) \rightarrow SameAuthor (pmid, auth_pos, pmid', auth_pos')
	...	
	Rule 5	Author (lastname, firstname, affiliation, pmid, auth_pos, title, coauthors) \wedge Author (lastname, firstname, affiliation', pmid', auth_pos', title', coauthors') \wedge jaccardSim (affiliation, affiliation', 20.0) \rightarrow SameAuthor (pmid, auth_pos, pmid', auth_pos')
	...	
✓	Rule 8	Author (lastname, firstname, affiliation, pmid, auth_pos, title, coauthors) \wedge Author (lastname, firstname, affiliation', pmid', auth_pos', title', coauthors') \wedge jaccardSim (affiliation, affiliation', 50.0) \rightarrow SameAuthor (pmid, auth_pos, pmid', auth_pos')
	...	

Figure 3: Candidate rules for an entity resolution task

precision and recall with respect to a given set of ground truth data. As an example, both Rule 1 and Rule 2 in Figure 3 generate a SameAuthor link between two author occurrences on two different publications, provided that the last names and first names are identical and provided that there is a sufficient number of common coauthors on the two publications. However, Rule 1, which checks for at least two common coauthors, may turn out to be too imprecise (i.e., may yield too many false positives), while Rule 2, which checks for at least three common coauthors, may result into fewer errors. Different rules may use different predicates in their premises. For example, Rules 5–8 exploit the Jaccard similarity of affiliation, but have different similarity thresholds. Only one of them (Rule 8, with similarity threshold of 50%) may achieve high enough precision.

Thus, the problem becomes how to select a set of rules that achieve high precision (i.e., minimize the number of false positives) and high recall (i.e., minimize the number of false negatives) with respect to a given set of ground truth data; furthermore, one would also like to select a compact (in terms of size) such set of rules.

Similar rule selection problems have been studied in several different contexts. In data exchange, (Kimmig et al. 2017) have investigated the *mapping selection* problem: given a set \mathcal{C} of rules expressing data transformations between a source and a target schema, and a pair (I, J) of a source database I and a target database J , find a subset \mathcal{C}' of the rules that minimizes the sum of the false positive errors, the false negative errors, and the sizes of the rules in \mathcal{C}' . (Sarma et al. 2010) have investigated the *view selection* problem: given a materialized view V , a database D , and

a collection \mathcal{S} of sets of rules on D , find a set of rules in \mathcal{S} that is as “close” to the view V and as compact as possible. In data mining, (Agrawal, Rantzaou, and Terzi 2006; Davis, Schwarz, and Terzi 2009) investigated the problems of selecting contextual preference rules and association rules; these problems can be cast as variants of the rule selection problem considered here.

Summary of Results We formalize the rule selection problem with rules specified by Horn formulas of first-order logic; the relation symbols in the premises of the Horn formulas come from a *premise* schema, while those in the conclusions come from a *conclusion* schema that is disjoint from the premise schema. This formalization captures rule selection problems in a variety of contexts.

An input to the rule selection problem consists of a finite set \mathcal{C} of rules and a pair (I, J) of a premise database I and a conclusion database J that represents ground truth. When a subset \mathcal{C}' of \mathcal{C} is evaluated on the premise instance I , it produces a conclusion instance $\text{Eval}(\mathcal{C}', I)$. The set $\text{Eval}(\mathcal{C}', I) \setminus J$ is the set of the false positive errors, while the set $J \setminus \text{Eval}(\mathcal{C}', I)$ is the set of false negative errors. We study the optimization problem $\text{MIN RULE-SELECT}_{\text{FP+FN}}$ in which, given \mathcal{C} , (I, J) as above, the goal is to find a subset \mathcal{C}' of \mathcal{C} so that the number of false positive and false negative errors is minimized. We also study the optimization problem $\text{MIN RULE-SELECT}_{\text{FP}}$ in which the goal is to find a subset \mathcal{C}' of \mathcal{C} so that the number of false positive errors is minimized and there are no false negative errors (this is meaningful when $J \subseteq \text{Eval}(\mathcal{C}, I)$).

To gauge the difficulty of these two optimization problems, we first examine their underlying decision problems. We show that the decision problems involving a bound on the error are NP-hard. We also show that the *exact* decision problems asking if the error is equal to a given value are DP-hard; in particular, they are both NP-hard and coNP-hard (thus, unlikely to be in $\text{NP} \cup \text{coNP}$). In view of these hardness results, we focus on the approximation properties of the two rule selection optimization problems. We show that, in a precise sense, $\text{MIN RULE-SELECT}_{\text{FP}}$ has the same approximation properties as the RED-BLUE SET COVER problem, while $\text{MIN RULE-SELECT}_{\text{FP+FN}}$ has the same approximation properties as the POSITIVE-NEGATIVE PARTIAL SET COVER problem. These results yield both polynomial-time approximation algorithms and lower bounds for the approximability of our problems.

The preceding results focus on the minimization of the error produced by the selected rules. What if one wants to also take the size of the selected rules into account? Since error and size are qualitatively incomparable quantities, it is not meaningful to add them or even to take a linear combination of the two. Instead, we consider pairs of values (e, s) of error and size that are *Pareto optimal*, that is, neither of these values can be decreased without increasing the other value at the same time. The *Pareto front* of an instance is the set of all Pareto optimal pairs. Even though the study of Pareto optimality has been a central theme of multi-objective optimization for decades, it appears that no such study has been

carried out for rule selection problems in any of the contexts discussed earlier. Here, we initiate such a study and show that the following problem is DP-hard: given a set \mathcal{C} of rules, a pair (I, J) of a premise database and a conclusion database, and a pair (e, s) of integers, does (e, s) belong to the Pareto front of $\text{MIN RULE-SELECT}_{\text{FP}+\text{FN}}$? We also show that a similar DP-hardness result holds for $\text{MIN RULE-SELECT}_{\text{FP}}$.

Finally, we investigate a bi-level optimization version of $\text{MIN RULE-SELECT}_{\text{FP}+\text{FN}}$, where one minimizes first the total error and then the size. We show that the following problem is DP-hard: given a set \mathcal{C} of rules, a pair (I, J) of a premise database and a conclusion database, and a pair (e, s) of integers, is e the minimum possible error and is s the minimum size of subsets of rules having the minimum error? We also show a similar DP-hardness result holds for the bi-level optimization version of $\text{MIN RULE-SELECT}_{\text{FP}}$.

The main results of this paper are summarized in Table 1, which can be found in a subsequent section.

Related Work

We already mentioned that (Kimmig et al. 2017) studied the *mapping selection* problem in the context of data exchange. In addition to considering rules specified by Horn formulas (GAV constraints in data exchange), they also considered richer rules in which the conclusion involves existential quantification over a conjunction of atoms (GLAV - global and local as view - constraints in data exchange). They established that the mapping selection problem is NP-hard even for GAV constraints, but did not explore approximation algorithms for this optimization problem; instead, they designed an algorithm that uses probabilistic soft logic (Bach et al. 2017) to solve a relaxation of the mapping selection problem and then carried out a detailed experimental evaluation of this approach. Its many technical merits notwithstanding, the work of (Kimmig et al. 2017) suffers from a serious drawback, namely, the objective function of the mapping selection problem is defined to be the sum of the size of the rules and the error (the number of the false positives and the false negatives). As stated earlier, however, size and error are qualitatively different quantities, thus it is simply not meaningful to add them, the same way it is not meaningful to add dollars and miles if one is interested in a hotel room near the White House and is trying to minimize the cost of the room and the distance from the White House. This is why, to avoid this pitfall here, we first focus on error minimization alone and then study the Pareto optimality of pairs of size and error values.

In the rule selection problem, the aim is to select a set of rules from a larger set of candidate rules based on some given data. There is a large body of work on the problem of deriving a set of rules in data exchange and data integration from just one or more given *data examples*. There are several different approaches to this problem, including casting it as an optimization problem (Gottlob and Senellart 2010; ten Cate et al. 2017), as a “fitting” problem (Alexe et al. 2011a; 2011b), as an interactive derivation problem (Bonifati et al. 2017), or as a learning problem (ten Cate, Dalmau, and Kolaitis 2013; ten Cate et al. 2018). Clearly, this is a

related but rather different problem because, in contrast to the rule selection problem, no candidate rules are part of the input.

Basic Concepts and Algorithmic Problems

Schemas and Instances A *schema* \mathbf{R} is a set $\{R_1, \dots, R_k\}$ of relation symbols, each with a specified *arity* indicating the number of its arguments. An *\mathbf{R} -instance* I is a set $\{R_1^I, \dots, R_k^I\}$ of relations whose arities match those of the corresponding relation symbols. An *\mathbf{R} -instance* can be identified with the set of all *facts* $R_i(a_1, \dots, a_m)$, such that R_i is a relation symbol in \mathbf{R} and (a_1, \dots, a_m) is a tuple in the relation R_i^I of I interpreting the relation symbol R_i .

Rules. Let \mathbf{S} and \mathbf{T} be two disjoint relational schemas. In the rest of the paper, we will refer to \mathbf{S} as the *premise* schema and to \mathbf{T} as the *conclusion* schema. A *rule* over \mathbf{S} and \mathbf{T} is a Horn formula of first-order logic of the form

$$\forall \mathbf{x} (\psi(\mathbf{x}) \rightarrow P(\mathbf{x})),$$

where the premise $\psi(\mathbf{x})$ is a conjunction of atoms over \mathbf{S} and the conclusion $T(\mathbf{x})$ is a single atom over \mathbf{T} with variables among those in \mathbf{x} . For example, the rule

$$\forall x, y, z (E(x, z) \wedge E(z, y) \rightarrow F(x, y))$$

asserts that F contains all pairs of nodes connected via an E -path of length 2. For simplicity, we will be dropping the universal quantifiers \forall , so that the preceding rule about paths of length 2 will be written as $(E(x, z) \wedge E(z, y) \rightarrow F(x, y))$.

The atoms in the premises may contain constants or they may be built-in predicates, such as **jaccardSim**. However, none of the lower-bound complexity results established here uses such atoms, while the upper-bound complexity results hold true even in the presence of such atoms, provided the built-in predicates are polynomial-time computable.

The *size* of a rule ρ , denoted by $|\rho|$, is the number of atoms in the premise of ρ . The *size* of a collection \mathcal{C} of rules, denoted by $|\mathcal{C}|$, is the sum of the sizes of the rules in \mathcal{C} .

Data example. A *data example* is a pair (I, J) , where I is an instance over the premise schema \mathbf{S} and J is an instance over the conclusion schema \mathbf{T} .

Rule evaluation. Given a rule ρ and an instance I , we write $\text{Eval}(\rho, I)$ to denote the result of evaluating the premise of ρ on I and then populating the conclusion of ρ accordingly. For example, if ρ is the rule $(E(x, z) \wedge E(z, y) \rightarrow F(x, y))$ and E is a graph, then $\text{Eval}(\rho, E)$ is the set F consisting of all pairs of nodes of E connected via a path of length 2. If \mathcal{C} is a set of rules, then $\text{Eval}(\mathcal{C}, I)$ is the set of facts $\bigcup_{\rho \in \mathcal{C}} \text{Eval}(\rho, I)$. In data exchange, computing $\text{Eval}(\mathcal{C}, I)$ amounts to running the *chase* procedure (Fagin et al. 2005).

In general, given a collection \mathcal{C} of rules and an instance I , computing $\text{Eval}(\mathcal{C}, I)$ is an exponential-time task; the source of the exponentiality is the maximum number of atoms in the premises of the rules in \mathcal{C} and the maximum arity of the relation symbols in the conclusions of the rules. If, however, both these quantities are bounded by constants, then $\text{Eval}(\mathcal{C}, I)$ is computable in polynomial time, according to the following fact (e.g., see (Fagin et al. 2005)).

Proposition 1. *Let a and r be two fixed positive integers. Then the following problem is solvable in polynomial time: given an instance I and a collection \mathcal{C} of rules such that the maximum number of atoms in the premises of rules in \mathcal{C} is at most a and the maximum arity of the relation symbols in the conclusions of these rules is at most r , compute $\text{Eval}(\mathcal{C}, I)$.*

False positive errors and false negative errors. Given a collection \mathcal{C} of rules and a data example (I, J) , a *false positive error* is a fact f in $\text{Eval}(\mathcal{C}, I)$ that is not in J , while a *false negative error* is a fact f in J that is not in $\text{Eval}(\mathcal{C}, I)$. We write $\text{FP}(\mathcal{C}, (I, J))$ and $\text{FN}(\mathcal{C}, (I, J))$ (or, simply, FP and FN) for the set of false positive and false negative errors with of \mathcal{C} with respect to (I, J) , that is,

$$\text{FP} = \text{Eval}(\mathcal{C}, I) \setminus J \text{ and } \text{FN} = J \setminus \text{Eval}(\mathcal{C}, I).$$

We will focus on the following two optimization problems concerning the minimization of the number of errors.

Definition 1. $[\text{MIN RULE-SELECT}_{\text{FP+FN}}]$

Input: A set \mathcal{C} of rules and a data example (I, J) .

Goal: Find a subset $\mathcal{C}^* \subseteq \mathcal{C}$ such that the sum of the number of the false positive errors and the number of false negative errors of \mathcal{C}^* with respect to (I, J) , is minimized.

A *feasible* solution of a given instance $\mathcal{C}, (I, J)$ of $\text{MIN RULE-SELECT}_{\text{FP+FN}}(a, r)$ is a subset \mathcal{C}' of \mathcal{C} . We write $\text{ERROR}_{\text{FP+FN}}(\mathcal{C}', (I, J))$ to denote the *error* of \mathcal{C}' with respect to (I, J) , i.e., the sum of the number of the false positive errors and the number of false negative errors.

Definition 2. $[\text{MIN RULE-SELECT}_{\text{FP}}]$

Input: A set \mathcal{C} of rules and a data example (I, J) such that $J \subseteq \text{Eval}(\mathcal{C}, I)$.

Goal: Find a subset $\mathcal{C}^* \subseteq \mathcal{C}$ such that the number of false negative errors is zero and the number of false positive errors of \mathcal{C}^* with respect to (I, J) is minimized.

A *feasible* solution of a given instance $\mathcal{C}, (I, J)$ of $\text{MIN RULE-SELECT}_{\text{FP}}(a, r)$ is a subset \mathcal{C}' of \mathcal{C} such that $J \subseteq \text{Eval}(\mathcal{C}', I)$. Feasible solutions always exist because \mathcal{C} is one. We write $\text{ERROR}_{\text{FP}}(\mathcal{C}', (I, J))$ to denote the *error* of \mathcal{C}' with respect to (I, J) , i.e., the number of the false positive errors.

We do not consider $\text{MIN RULE-SELECT}_{\text{FN}}$, i.e., the optimization problem that aims to minimize the number of false negative errors. The reason is that there is a trivial solution to this problem, namely, we can select all the rules (if the number of false positive errors is not required to be zero) or select all the rules that produce no false positive errors (if the number of false positive errors is required to be zero).

To gauge the difficulty of solving an optimization problem, one often studies two decision problems that underlie the optimization problem at hand: a decision problem about bounds on the optimum value and a decision problem about the exact optimum value. We introduce these two problems for each of the optimization problems in Definitions 1 and 2.

Definition 3. Given a set \mathcal{C} of rules, a data example (I, J) , and an integer k ,

- $\text{RULE-SELECT}_{\text{FP+FN}}$ asks: is there a subset \mathcal{C}' of \mathcal{C} such that $\text{ERROR}_{\text{FP+FN}}(\mathcal{C}', (I, J)) \leq k$?
- $\text{EXACT RULE-SELECT}_{\text{FP+FN}}$ asks: is the optimum value of $\text{MIN RULE-SELECT}_{\text{FP+FN}}$ on \mathcal{C} and (I, J) equal to k ?

Definition 4. Given a set \mathcal{C} of rules, a data example (I, J) such that $J \subseteq \text{Eval}(\mathcal{C}, I)$, and an integer k ,

- $\text{RULE-SELECT}_{\text{FP}}$ asks: is there a subset $\mathcal{C}' \subseteq \mathcal{C}$ such that the number of false negative errors of \mathcal{C}' with respect to (I, J) is zero and $\text{ERROR}_{\text{FP}}(\mathcal{C}', (I, J)) \leq k$?
- $\text{EXACT RULE-SELECT}_{\text{FP}}$ asks: is the optimum value of $\text{MIN RULE-SELECT}_{\text{FP}}$ on \mathcal{C} and (I, J) equal to k ?

The Complexity of Error Minimization

We will investigate the computational complexity of the decision problems introduced in Definitions 3 and 4 by considering parameterized versions of these problems with parameters the maximum number of atoms in the premises of the rules and the maximum arity of the relation symbols in the conclusion of the rules.

Definition 5. Let a and r be two fixed positive integers.

- $\text{RULE-SELECT}_{\text{FP+FN}}(a, r)$ is the restriction of $\text{RULE-SELECT}_{\text{FP+FN}}$ to inputs in which the maximum number of atoms in the premises of rules in the given set \mathcal{C} of rules is at most a and the maximum arity of the relation symbols in the conclusions of these rules is at most r .
- The decision problems $\text{RULE-SELECT}_{\text{FP}}(a, r)$, $\text{EXACT RULE-SELECT}_{\text{FP+FN}}(a, r)$, $\text{EXACT RULE-SELECT}_{\text{FP}}(a, r)$, $\text{MIN RULE-SELECT}_{\text{FP+FN}}(a, r)$, and $\text{MIN RULE-SELECT}_{\text{FP}}(a, r)$ are defined in an analogous way.

Theorem 1. *Let a and r be two fixed positive integers. The decision problems $\text{RULE-SELECT}_{\text{FP}}(a, r)$ and $\text{RULE-SELECT}_{\text{FP+FN}}(a, r)$ are NP-complete.*

Proof. (Sketch) Membership in NP follows easily from Proposition 1. The NP-hardness of $\text{RULE-SELECT}_{\text{FP}}(a, r)$ and $\text{RULE-SELECT}_{\text{FP+FN}}(a, r)$ is proved using a polynomial-time reduction from SET COVER, a problem which was among the twenty one NP-complete problems in Karp's seminal paper (Karp 1972). The same reduction is used for both $\text{RULE-SELECT}_{\text{FP}}(a, r)$ and $\text{RULE-SELECT}_{\text{FP+FN}}(a, r)$; however, the argument for the correctness of the reduction is different in each case. We give the argument for $\text{RULE-SELECT}_{\text{FP+FN}}(a, r)$.

The SET COVER problem asks: given a finite set $U = \{u_1, \dots, u_m\}$, a collection $\mathcal{S} = \{S_1, \dots, S_p\}$ of subsets of U whose union is U , and an integer k , is there a cover of U of size at most k ? (i.e., a subset $\mathcal{S}' \subseteq \mathcal{S}$ such that the union of the members of \mathcal{S}' is equal to U and $|\mathcal{S}'| \leq k$.)

Let $U = \{u_1, \dots, u_m\}$, $\mathcal{S} = \{S_1, \dots, S_p\}$, k be an instance of SET COVER. For every i with $1 \leq i \leq p$, we introduce the rule $\text{Set}_i(x) \rightarrow B(x)$ and we let \mathcal{C} be the set of all these rules. Thus, the premise schema consists of the unary relation symbols Set_i , $1 \leq i \leq p$, and the conclusion schema consists of the unary relation symbol B . We introduce p new elements a_1, \dots, a_p and construct the premise instance I with $\text{Set}_i^I = S_i \cup \{a_i\}$, $1 \leq i \leq p$ (intuitively, each new element a_i encodes the index of the set S_i). We also construct the conclusion instance J with $B^J = U$.

We claim that there is a cover of U of size at most k if and only if there is a subset \mathcal{C}' of \mathcal{C} such that $\text{ERROR}_{\text{FP+FN}}(\mathcal{C}', (I, J)) \leq k$. If there is a cover of

		False Positive Errors	False Positive + False Negative Errors
RULE-SELECT(a,r)		NP-complete	NP-complete
EXACT RULE-SELECT(a,r)		DP-complete	DP-complete
MIN RULE-SELECT(a,r)	approx. upper bound	$2\sqrt{ \mathcal{C} \log J }$	$2\sqrt{(\mathcal{C} + J) \log J }$
	approx. lower bound	$2^{\log^{1-\epsilon}(\mathcal{C})}$, for every $\epsilon > 0$	$2^{\log^{1-\epsilon}(J)}$, for every $\epsilon > 0$
PARETO OPT SOLUTION(a,r)		coNP-complete	coNP-complete
PARETO FRONT MEMBERSHIP(a,r)		DP-complete	DP-complete
BI-LEVEL OPT SOLUTION(a,r)		coNP-complete	coNP-complete
BI-LEVEL OPT VALUE(a,r)		DP-complete	DP-complete

Table 1: Summary of Results (\mathcal{C} : set of input rules; J : input conclusion instance; approx. lower bounds assume that $P \neq NP$)

U of size at most k , then there is a subset \mathcal{C}' of \mathcal{C} that has no false negative errors with respect to (I, J) and is such that $\text{ERROR}_{\text{FP}}(\mathcal{C}', (I, J)) \leq k$. Therefore, $\text{ERROR}_{\text{FP+FN}}(\mathcal{C}', (I, J)) = \text{ERROR}_{\text{FP}}(\mathcal{C}', (I, J)) \leq k$. Conversely, assume that there is a subset \mathcal{C}' of \mathcal{C} such that $\text{ERROR}_{\text{FP+FN}}(\mathcal{C}', (I, J)) \leq k$. If there are no false negative errors, then \mathcal{C}' corresponds to a cover of U of size at most k . So, assume there are t false negative errors, for some $t > 0$. These must result from elements u_{i_1}, \dots, u_{i_t} of U that are not in any of the sets S_j that produce rules $\text{Set}_j(x) \rightarrow B(x)$ in \mathcal{C}' . For each such element, there is a set in \mathcal{S} containing it; thus, the elements u_{i_1}, \dots, u_{i_t} can be covered using at most t sets in \mathcal{S} . Let \mathcal{C}'' be the subset of \mathcal{C} obtained by adding to \mathcal{C}' the rules arising from these sets. Clearly, \mathcal{C}'' has no false negative errors, so it has t fewer negative errors than \mathcal{C}' does and at most t more positive errors than \mathcal{C}' does with respect to (I, J) . Thus, $\text{ERROR}_{\text{FP+FN}}(\mathcal{C}'', (I, J)) \leq \text{ERROR}_{\text{FP+FN}}(\mathcal{C}', (I, J)) \leq k$. It follows that the subset \mathcal{C}'' gives rise to a cover of U of size at most k . This completes the proof that $\text{RULE-SELECT}_{\text{FP+FN}}(a,r)$ is NP-hard. \square

Next, we consider the problems $\text{EXACT RULE-SELECT}_{\text{FP}}(a,r)$ and $\text{EXACT RULE-SELECT}_{\text{FP+FN}}(a,r)$. The class DP consists of all decision problems that can be written as the conjunction of a problem in NP and a problem in coNP (Papadimitriou and Yannakakis 1982). The prototypical DP-complete problem is SAT-UNSAT: given two Boolean formulas φ and ψ , is φ satisfiable and ψ unsatisfiable? Furthermore, EXACT CLIQUE is DP-complete: given a graph G and an integer k , is the size of the largest clique in G exactly k ? Note that a DP-complete problem is both NP-hard and coNP-hard.

Theorem 2. *Let a and r be two fixed positive integers. The decision problems $\text{EXACT RULE-SELECT}_{\text{FP}}(a,r)$ and $\text{EXACT RULE-SELECT}_{\text{FP+FN}}(a,r)$ are DP-complete.*

Proof. (Hint) Using Proposition 1, it is easy to see that both these problems are in the class DP. The problem EXACT SET COVER asks: given a set U , a collection \mathcal{S} of subsets of U , and an integer k , is the size of smallest cover of U exactly k ? Using the DP-hardness of EXACT CLIQUE and reductions in (Karp 1972), one can show that EXACT SET COVER is DP-hard. Finally, the reduction in the proof of Theorem 1 is also a polynomial-time reduction of EXACT SET COVER to both $\text{EXACT RULE-SELECT}_{\text{FP}}(a,r)$ and $\text{EXACT RULE-SELECT}_{\text{FP+FN}}(a,r)$. Consequently, these two problems are

DP-complete. \square

In Theorems 1 and 2, the restriction to sets of rules with a bound a on the number of premise atoms and a bound r on the arity of relation symbols in the conclusion schema was used to obtain the complexity-theoretic upper bounds (membership in NP and, respectively, membership in DP). The matching lower bounds hold true even if $a = 1$ and $r = 1$, because the rules used to prove NP-hardness in Theorem 1 and DP-hardness in Theorem 2 have the form $\text{Set}_i(x) \rightarrow B(x)$.

Observe also that, in the proofs of Theorems 1 and 2, there is no fixed bound on the number of relation symbols in the premise schema; instead, this number depends on the size of the given instance of SET COVER and EXACT SET COVER . If we also impose a fixed bound on the number of relation symbols, then the rule selection problems trivializes because, in this case, there is a fixed number of rules. Our next result shows that the rule selection problems become intractable if we fix the premise schema (hence, we also fix the number of relation symbols occurring in it), but impose no bounds on the number of atoms in the premises of the rules.

Theorem 3. *Let \mathbf{S} be a premise schema consisting of one unary and four binary relation symbols, and let \mathbf{T} be a target schema consisting of a single unary relation symbol.*

If the rules contain an arbitrary finite number of atoms in their premises, then $\text{RULE-SELECT}_{\text{FP}}$ is NP-hard and $\text{EXACT RULE-SELECT}_{\text{FP}}$ is DP-hard. Similar results hold for $\text{RULE-SELECT}_{\text{FP+FN}}$ and $\text{EXACT RULE-SELECT}_{\text{FP+FN}}$.

Proof. (Hint) Assume that the premise schema \mathbf{S} consists of the unary relation symbol *One* and binary relation symbols *S*, *Bit*₀, *Bit*₁, *Succ*; assume also that the conclusion schema consists of the unary relation symbol *B*. We simulate each rule $\text{Set}_i(x) \rightarrow B(x)$, $1 \leq i \leq p$, via a rule $\sigma_{i,p}$ that, intuitively, asserts that “if $z = i$ and $x \in \text{Set}_i$, then $x \in B$ ”. Since $i \leq p$, we can write i in binary notation using $\lceil \log p \rceil$ bits. We encode this binary representation of i via a conjunction of premise atoms involving the binary relation symbols *Bit*₀, *Bit*₁, *Succ* and the unary relation symbol *One*. A typical rule $\sigma_{i,p}$ is of the form

$$\text{Bit}_0(i_1, z) \wedge \text{Bit}_1(i_2, z) \wedge \dots \wedge \text{Bit}_1(i_{\lceil \log p \rceil}, z) \wedge \text{One}(i_1) \\ \wedge \text{Succ}(i_1, i_2) \wedge \dots \wedge \text{Succ}(i_{\lceil \log p \rceil - 1}, i_{\lceil \log p \rceil}) \wedge S(x, z) \rightarrow B(x),$$

where the binary representation of i using $\lceil \log p \rceil$ bits gives rise to the sequence of atoms Bit_0 and Bit_1 in the premise of the rule. If an instance of SET COVER contains a set U and a collection $\mathcal{S} = \{S_1, \dots, S_p\}$ of subsets of U , then we construct a premise instance I as follows:

- $S^I = \{(x, i) : x \in S_i \cup \{a_i\}, 1 \leq i \leq p\}$, where each a_i is a new element;
- $Bit_k^I = \{(j, i) : i \leq p \wedge \text{the } j\text{-th bit of } i \text{ is } k\}, k = 0, 1;$
- $One^I = \{1\}; Succ^I = \{(j, j+1) : 1 \leq j \leq \lceil \log p \rceil - 1\}.$

As a conclusion instance, we construct J with $B^J = U$. \square

The Approximation of Error Minimization

The hardness results in Theorems 1 and 2 imply that, unless $P=NP$, there is no polynomial-time algorithm for solving exactly the optimization problems $\text{MIN RULE-SELECT}_{\text{FP}(a,r)}$ and $\text{MIN RULE-SELECT}_{\text{FP+FN}(a,r)}$. As detailed in (Garey and Johnson 1979), approximation algorithms offer a way to cope with the computational hardness of optimization problems. In the case of minimization problems, the goal is to design a polynomial-time algorithm that, given an instance of the optimization problem at hand, returns a feasible solution such that the approximate value is less than or equal to a certain factor of the optimum value; in general, the factor depends on the size of the instance. For example, there is a polynomial-time algorithm that, given an instance (U, \mathcal{S}) of the MIN SET COVER problem, the algorithm returns a cover whose size is less than or equal to $\ln(|\mathcal{S}|)$ times the size of the optimal cover of the instance (U, \mathcal{S}) .

It is well known that different optimization problems may have widely different approximation properties; for example, KNAPSACK is ϵ -approximable, for every $\epsilon > 0$, while MIN SET COVER is not ϵ -approximable for any $\epsilon > 0$, unless $P=NP$ (see (Arora and Barak 2009) for an exposition).

Let a and r be two fixed positive integers. What are the approximation properties of $\text{MIN RULE-SELECT}_{\text{FP}(a,r)}$ and $\text{MIN RULE-SELECT}_{\text{FP+FN}(a,r)}$? At first sight and in view of the reduction from SET COVER in Theorems 1 and 2, one may guess that they ought to be the same as those of MIN SET COVER. As we shall see, though, the approximation properties of $\text{MIN RULE-SELECT}_{\text{FP+FN}(a,r)}$ appear to be worse than those of $\text{MIN RULE-SELECT}_{\text{FP}(a,r)}$, which, in turn, appear to be worse than those of MIN SET COVER.

Note that polynomial time reductions between two decision problems need not preserve the approximation properties of the corresponding optimization problems. For example, SET COVER is polynomial-time reducible to NODE COVER, yet MIN NODE COVER is constant-approximable, but MIN SET COVER is not. There is, however, a finer notion of polynomial-time reduction, called \mathcal{L} -reduction, that preserves approximation properties between optimization problems; this notion, which was introduced in (Papadimitriou and Yannakakis 1991), is defined next.

Let Q and Q' be two optimization problems. We say that Q \mathcal{L} -reduces to Q' if there are two polynomial-time algorithms f, g , and two positive constants α, β such that for every instance I of Q , the following hold:

- Algorithm f produces an instance $I' = f(I)$ of Q' such that the optimum values $\text{opt}_Q(I)$ of Q on I and $\text{opt}_{Q'}(I')$ of Q' on I' satisfy $\text{opt}_Q(I) \leq \alpha \text{opt}_{Q'}(I')$.
- For every solution I' of Q' with value c' , algorithm g produces a solution I of Q with value c such that $|c - \text{opt}_Q(I)| \leq \beta |c' - \text{opt}_{Q'}(I')|$.

By Proposition 2 in (Papadimitriou and Yannakakis 1991), if Q reduces to Q' via an \mathcal{L} -reduction in which $\alpha = 1$ and $\beta = 1$, then every polynomial-time approximation algorithm for Q' with factor ϵ gives rise to a polynomial-time approximation algorithm for Q with the same factor ϵ . We say that two optimization problems Q and Q' have the same approximation properties if there are \mathcal{L} -reductions from Q to Q' and from Q' to Q in both of which $\alpha = 1$ and $\beta = 1$.

We now bring into the picture the RED-BLUE SET COVER problem, a variant of MIN SET COVER that will turn out to have the same approximation properties as MIN RULE-SELECT_{FP}(a,r). The properties of RED-BLUE SET COVER were studied in (Carr et al. 2000), where both upper and lower bounds for its approximability were obtained; improved upper bounds were obtained in (Peleg 2007).

Definition 6. [RED-BLUE SET COVER problem]

Input: Two disjoint sets $R = \{r_1, \dots, r_\alpha\}$ and $B = \{b_1, \dots, b_\beta\}$ of “red” and “blue” elements, and a collection $\mathcal{S} = \{S_1, \dots, S_m\}$ of subsets of $R \cup B$.

Goal: Find a subset $\mathcal{S}^* \subseteq \mathcal{S}$ that covers all blue elements, but covers as few red elements as possible.

Theorem 4. *Let a and r be two fixed positive integers. MIN RULE-SELECT_{FP}(a,r) and RED-BLUE SET COVER have the same approximation properties. Consequently, the following hold true for MIN RULE-SELECT_{FP}(a,r).*

- MIN RULE-SELECT_{FP}(a,r) is approximable within a factor of $2\sqrt{|\mathcal{C}| \log |J|}$, where $|\mathcal{C}|$ is the number of input rules and $|J|$ is the size of the input conclusion instance J .
- Unless $P=NP$, for every $\epsilon > 0$, there is no polynomial time algorithm that approximates MIN RULE-SELECT_{FP}(a,r) within a factor of $2^{\log^{1-\epsilon}(|\mathcal{C}|)}$, where $|\mathcal{C}|$ is as above.

Proof. (Sketch) We show that MIN RULE-SELECT_{FP}(a,r) reduces to RED-BLUE SET COVER via a \mathcal{L} -reduction with $\alpha = 1$ and $\beta = 1$. Given an instance $K = (\mathcal{C}, (I, J))$ of MIN RULE-SELECT_{FP}(a,r), we construct the following instance $K' = (R, B, \mathcal{S})$ of RED-BLUE SET COVER.

We put $B = J$ and $R = \text{Eval}(\mathcal{C}, I) \setminus J$. Thus, the blue elements are the facts of J , while the red elements are the facts of $\text{Eval}(\mathcal{C}, I)$ that are not in J . We form the collection \mathcal{S} consisting of all sets $\text{Eval}(r, I)$, where r is a rule in \mathcal{C} .

It is easy to see that this is a \mathcal{L} -reduction with $\alpha = 1$ and $\beta = 1$. Thus, every approximation algorithm for RED-BLUE SET COVER gives rise to an approximation algorithm for MIN RULE-SELECT_{FP}(a,r) with the same approximation factor. In (Peleg 2007), a polynomial-time algorithm with approximation factor of $2\sqrt{|\mathcal{S}| \log \beta}$ for RED-BLUE SET COVER is described, where \mathcal{S} is the given collection of sets and β is the number of blue elements. This yields a polynomial-time algorithm for MIN RULE-SELECT_{FP}(a,r) with an approximation factor of $2\sqrt{|\mathcal{C}| \log |J|}$.

Next, we show that RED-BLUE SET COVER reduces to MIN RULE-SELECT_{FP(a,r)} via a \mathcal{L} -reduction with $\alpha = 1$ and $\beta = 1$. Given an instance $K' = (R, B, \mathcal{S})$ of RED-BLUE SET COVER, we construct the following instance $K = (\mathcal{C}, (I, J))$ of MIN RULE-SELECT_{FP(a,r)}.

Let $\mathcal{S} = \{S_1, \dots, S_m\}$. The premise schema consists of unary relation symbols Set_i , $1 \leq i \leq m$, and the conclusion schema consists of a unary relation symbol R . We put $\mathcal{C} = \{Set_i(x) \rightarrow R(x) : 1 \leq i \leq m\}$. We construct the premise instance I by putting $Set_i^I = S_i$, $1 \leq i \leq m$; we construct the conclusion instance J by putting $R^J = B$.

It is easy to see that this is a \mathcal{L} -reduction with $\alpha = 1$ and $\beta = 1$. Thus, inapproximability results for RED-BLUE SET COVER transfer to inapproximability results for MIN RULE-SELECT_{FP(a,r)}. In (Carr et al. 2000), it was shown that, unless P=NP, for every $\epsilon > 0$, there is no polynomial time algorithm that approximates RED-BLUE SET COVER within a factor of $2^{\log^{1-\epsilon}(|\mathcal{S}|)}$. Thus, unless P=NP, for every $\epsilon > 0$, there is no polynomial time algorithm that approximates MIN RULE-SELECT_{FP(a,r)} within a factor of $2^{\log^{1-\epsilon}(|\mathcal{C}|)}$. \square

Finally, we bring into the picture the POSITIVE-NEGATIVE PARTIAL SET COVER (\pm PSC) problem, a variant of MIN SET COVER that will turn out to have the same approximation properties as MIN RULE-SELECT_{FP+FN(a,r)}. The \pm PSC problem has been studied in (Miettinen 2008).

Definition 7. [POSITIVE-NEGATIVE PARTIAL SET COVER]

Input: Two disjoint sets P and N of “positive” and “negative” elements, and a collection $\mathcal{S} = \{S_1, \dots, S_m\}$ of subsets of $P \cup N$.

Goal: Find a subset $\mathcal{S}^* \subseteq \mathcal{S}$ that minimizes the sum of the number of uncovered positive elements and the number of covered negative elements, i.e., the quantity

$$\text{cost}(\mathcal{S}^*) = |P \setminus (\cup \mathcal{S}^*)| + |N \cap (\cup \mathcal{S}^*)|.$$

Theorem 5. *Let a and r be two fixed positive integers. MIN RULE-SELECT_{FP+FN(a,r)} and \pm PSC have the same approximation properties. Consequently, the following hold true for MIN RULE-SELECT_{FP+FN(a,r)}.*

- MIN RULE-SELECT_{FP+FN(a,r)} is approximable within a factor of $2\sqrt{(|\mathcal{C}| + |J|)\log|J|}$, where $|\mathcal{C}|$ is the number of input rules and $|J|$ is the size of the input conclusion instance J .
- Unless P=NP, for every $\epsilon > 0$, there is no polynomial time algorithm that approximates MIN RULE-SELECT_{FP+FN(a,r)} within a factor of $2^{\log^{1-\epsilon}(|J|)}$, where $|J|$ is as above.

Proof. (Hint) The \mathcal{L} -reductions between MIN RULE-SELECT_{FP+FN(a,r)} and \pm PSC are the same as those between MIN RULE-SELECT_{FP(a,r)} and RED-BLUE SET COVER in the proof of Theorem 4, but with the positive elements playing the role of the red elements and with the negative elements that of the negative elements. The upper and lower bounds for the approximability of MIN RULE-SELECT_{FP+FN(a,r)} follow from such bounds for \pm PSC. \square

Bi-Objective and Bi-Level Minimization

The optimization problems MIN RULE-SELECT_{FP(a,r)} and MIN RULE-SELECT_{FP+FN(a,r)} have a single objective, namely, the minimization of the error. Thus, all feasible solutions of minimum error are optimal, even though they may differ in size. What if one is also interested in taking the size of the solution into account? Since error and size are qualitatively incomparable quantities, it is not meaningful to combine them into a single objective function by taking, say, a linear combination of the two. Instead, we can cast the problem as a bi-objective optimization problem and study the *Pareto optimal solutions* that strike the right balance between error and size by capturing the trade-offs between these two quantities.

If (a, b) and (c, d) are two pairs of integers, then $(a, b) \leq (c, d)$ denotes that $a \leq c$ and $b \leq d$, while $(a, b) < (c, d)$ denotes that $(a, b) \leq (c, d)$ and $(a, b) \neq (c, d)$.

Definition 8. [Pareto Optimal Solution and Pareto Front]

Let a and r be two fixed positive integers and let $K = (\mathcal{C}, (I, J))$ be an instance of MIN RULE-SELECT_{FP+FN(a,r)}.

- A *Pareto optimal solution* for K is a subset \mathcal{C}^* of \mathcal{C} for which there is no subset $\mathcal{C}' \subseteq \mathcal{C}$ such that $(e', |\mathcal{C}'|) < (e^*, |\mathcal{C}^*|)$, where $e' = \text{ERROR}_{\text{FP+FN}}(\mathcal{C}', (I, J))$ and $e^* = \text{ERROR}_{\text{FP+FN}}(\mathcal{C}^*, (I, J))$.
- The *Pareto front* of K is the set of all pairs (e^*, s^*) of integers such that there is a Pareto optimal solution \mathcal{C}^* for K with $\text{ERROR}_{\text{FP+FN}}(\mathcal{C}^*, (I, J)) = e^*$ and $|\mathcal{C}^*| = s^*$.

The preceding notions of Pareto optimal solution and Pareto membership front give rise to natural decision problems. In what follows, we give the definitions for the parameterized versions of these problems.

Definition 9. [Pareto Optimal Solution Problem]

Let a and r be two fixed positive integers.

The PARETO OPT SOLUTION RULE-SELECT_{FP+FN(a,r)} problem asks: given an instance $K = (\mathcal{C}, (I, J))$ of MIN RULE-SELECT_{FP+FN(a,r)}, and a subset \mathcal{C}^* of \mathcal{C} , is \mathcal{C}^* a Pareto optimal solution for K ?

Definition 10. [Pareto Front Membership Problem]

Let a and r be two fixed positive integers.

The PARETO FRONT MEMBERSHIP RULE-SELECT_{FP+FN(a,r)} problem asks: given an instance $K = (\mathcal{C}, (I, J))$ of MIN RULE-SELECT_{FP+FN(a,r)} and a pair (e, s) of integers, is (e, s) on the Pareto front of K ?

The decision problems PARETO OPT SOLUTION RULE-SELECT_{FP(a,r)} and PARETO FRONT MEMBERSHIP RULE-SELECT_{FP(a,r)} are defined in an analogous manner.

Theorem 6. *Let a and r be two fixed positive integers. The following statements are true.*

- The decision problems PARETO OPT SOLUTION RULE-SELECT_{FP(a,r)} and PARETO OPT SOLUTION RULE-SELECT_{FP+FN(a,r)} are coNP-complete.
- The decision problems PARETO FRONT MEMBERSHIP RULE-SELECT_{FP(a,r)} and PARETO FRONT MEMBERSHIP RULE-SELECT_{FP+FN(a,r)} are DP-complete.

Proof. (Sketch) The coNP-hardness of the two Pareto optimality problems is shown via reductions from the coNP-complete problem SET-COVER OPTIMALITY: given a set U , a collection $\mathcal{S} = \{S_1, \dots, S_p\}$ of subsets of U whose union is U , and a subset \mathcal{S}' of \mathcal{S} , is \mathcal{S}' an optimal cover of U ? The DP-hardness of the two Pareto front membership problems is shown via reductions from the EXACT SET COVER problem, encountered in the proof of Theorem 2. In what follows, we outline the reduction of EXACT SET COVER to PARETO FRONT MEMBERSHIP RULE-SELECT_{FP+FN}(a,r).

Let $U = \{u_1, \dots, u_m\}$, $\mathcal{S} = \{S_1, \dots, S_p\}$, k be an instance of EXACT SET COVER. For each i , $1 \leq i \leq p$, we introduce the rule $Set_i(x) \rightarrow B(x)$ and we let \mathcal{C} be the set of all these rules. We introduce p new elements a_1, \dots, a_p ; furthermore, for every element $u_j \in U$, $1 \leq j \leq m$, we introduce p new elements b_j^1, \dots, b_j^p , which can be thought of as “clones” of u_j . We construct the premise instance I with $Set_i^I = S_i \cup \{a_i\} \cup \bigcup_{u_j \in S_i} \{b_j^1, \dots, b_j^p\}$, $1 \leq i \leq p$, that is, Set_i^I consists of a_i , the elements of the set S_i , and all their clones. We also construct the conclusion instance J with $B^J = U \cup \bigcup_{u_j \in U} \{b_j^1, \dots, b_j^p\}$, that is B^J consists of the elements of U and all their clones. Let $K = (\mathcal{C}, (I, J))$ be the resulting instance of MIN RULE-SELECT_{FP+FN}(a,r).

We claim that the optimum size of the SET COVER instance (U, \mathcal{S}) is k if and only if the pair (k, k) is on the Pareto front of the instance K . This claim follows by combining the following two facts. First, there is a 1-1 correspondence between covers \mathcal{S}' of U and subsets \mathcal{C}' of \mathcal{C} of the same size that have no negative errors with respect to (I, J) ; moreover, for such subsets \mathcal{S}' , we have that $|\mathcal{S}'| = \text{ERROR}_{\text{FP+FN}}(\mathcal{C}', (I, J))$. Second, if a subset \mathcal{C}' of \mathcal{C} does not correspond to a cover of U , then $\text{ERROR}_{\text{FP+FN}}(\mathcal{C}', (I, J)) \geq p + 1 > k$, since $k \leq p$ and, in this case, there are at least $p + 1$ false positives arising from an uncovered element of U and its p -many clones. \square

For bi-objective optimization problems, one would ideally like to efficiently produce the Pareto front of a given instance. For some such problems, this is impossible because the size (i.e., the number of points) of the Pareto front can be exponential in the size of a given instance. In the case of the Pareto front of rule selection problems, the size of the Pareto front is polynomial in the size of any given instance; nonetheless, Theorem 6 implies the following result.

Corollary 1. *Let a and r be two fixed positive integers. Unless $P=NP$, there is no polynomial-time algorithm that, given an instance K of MIN RULE-SELECT_{FP}(a,r) (or of MIN RULE-SELECT_{FP+FN}(a,r)), constructs the Pareto front of K .*

We conclude this section by identifying the complexity of the following two bi-level minimization problems in which one minimizes first for error and then for size.

Definition 11. [Bi-level Optimal Solution Problem] Let a and r be two fixed positive integers.

The BI-LEVEL OPT SOLUTION RULE-SELECT_{FP+FN}(a,r) problem asks: given an instance $K = (\mathcal{C}, (I, J))$ of MIN RULE-SELECT_{FP+FN}(a,r), and a subset \mathcal{C}^* of \mathcal{C} , is \mathcal{C}^* a bi-level optimal solution? (i.e., is \mathcal{C}^* an optimal solution of

MIN RULE-SELECT_{FP+FN}(a,r) that also has minimal size among all such optimal solutions?)

Definition 12. [Bi-level Optimal Value Problem] Let a and r be two fixed positive integers.

The BI-LEVEL OPT VALUE RULE-SELECT_{FP+FN}(a,r) problem asks: given an instance $K = (\mathcal{C}, (I, J))$ of MIN RULE-SELECT_{FP+FN}(a,r) and a pair (e, s) of integers, is e the error of a bi-level optimal solution and is s its size?

The decision problems BI-LEVEL OPT SOLUTION RULE-SELECT_{FP}(a,r) and BI-LEVEL OPT VALUE RULE-SELECT_{FP}(a,r) are defined in an analogous manner. Note that if a pair (e, s) is a bi-level optimal value, then it is a point on the Pareto front. Moreover, (e, s) is a special point because e is the minimum possible error and s is the minimum size of subsets of the given set of rules having this minimum error e .

The following result can be obtained by analyzing the proof of Theorem 6.

Corollary 2. *Let a and r be two fixed positive integers. The following statements are true.*

- *The decision problems BI-LEVEL OPT SOLUTION RULE-SELECT_{FP}(a,r) and BI-LEVEL OPT SOLUTION RULE-SELECT_{FP+FN}(a,r) are coNP-complete.*
- *The decision problems BI-LEVEL OPT VALUE RULE-SELECT_{FP}(a,r) and BI-LEVEL OPT VALUE RULE-SELECT_{FP+FN}(a,r) are DP-complete.*

Concluding Remarks

We carried out a systematic complexity-theoretic investigation of rule selection problems from several different angles, including an exploration of their complexity when they are cast as bi-objective optimization problems. A natural next step is the implementation and experimental evaluation of the approximation algorithms for MIN RULE-SELECT_{FP}(a,r) and MIN RULE-SELECT_{FP+FN}(a,r) based on corresponding approximation algorithms for RED-BLUE SET COVER and POSITIVE-NEGATIVE PARTIAL SET COVER. Note that the approximation algorithms for the latter problems have been used for blocking function selection in entity resolution (Bilenko, Kamath, and Mooney 2006) and in view selection (Sarma et al. 2010). A different next step is to leverage the large literature on bi-objective optimization, including (Legriel et al. 2010), and design heuristic algorithms for approximating the Pareto front of rule selection problems.

Acknowledgment The research of Phokion Kolaitis was partially supported by NSF Grant IIS:1814152.

References

Agrawal, R.; Rantzaou, R.; and Terzi, E. 2006. Context-sensitive ranking. In Chaudhuri, S.; Hristidis, V.; and Polyzotis, N., eds., *Proceedings of the ACM SIGMOD International Conference on Management of Data, Chicago, Illinois, USA, June 27-29, 2006*, 383–394. ACM.

- Alexe, B.; ten Cate, B.; Kolaitis, P. G.; and Tan, W. C. 2011a. Characterizing schema mappings via data examples. *ACM Trans. Database Syst.* 36(4):23:1–23:48.
- Alexe, B.; ten Cate, B.; Kolaitis, P. G.; and Tan, W. C. 2011b. Designing and refining schema mappings via data examples. In Sellis, T. K.; Miller, R. J.; Kementsietsidis, A.; and Velegrakis, Y., eds., *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2011, Athens, Greece, June 12-16, 2011*, 133–144. ACM.
- Arasu, A.; Re, C.; and Suciu, D. 2009. Large-Scale Deduplication with Constraints using Dedupalog. In *ICDE*, 952–963.
- Arora, S., and Barak, B. 2009. *Computational Complexity - A Modern Approach*. Cambridge University Press.
- Bach, S. H.; Broecheler, M.; Huang, B.; and Getoor, L. 2017. Hinge-loss markov random fields and probabilistic soft logic. *Journal of Machine Learning Research* 18:109:1–109:67.
- Bilenko, M.; Kamath, B.; and Mooney, R. J. 2006. Adaptive blocking: Learning to scale up record linkage. In *Proceedings of the 6th IEEE International Conference on Data Mining (ICDM 2006), 18-22 December 2006, Hong Kong, China*, 87–96. IEEE Computer Society.
- Bonifati, A.; Comignani, U.; Coquery, E.; and Thion, R. 2017. Interactive mapping specification with exemplar tuples. In *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD '17*, 667–682. New York, NY, USA: ACM.
- Burdick, D.; Fagin, R.; Kolaitis, P. G.; Popa, L.; and Tan, W. 2016. A Declarative Framework for Linking Entities. *ACM Trans. Database Syst.* 41(3):17. Preliminary version appeared in *ICDT*, pages 25–43, 2015.
- Carr, R. D.; Doddi, S.; Konjevod, G.; and Marathe, M. 2000. On the red-blue set cover problem. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '00*, 345–353. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.
- ten Cate, B.; Kolaitis, P. G.; Qian, K.; and Tan, W. 2017. Approximation algorithms for schema-mapping discovery from data examples. *ACM Trans. Database Syst.* 42(2):12:1–12:41.
- ten Cate, B.; Kolaitis, P. G.; Qian, K.; and Tan, W. 2018. Active learning of GAV schema mappings. In den Bussche, J. V., and Arenas, M., eds., *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Houston, TX, USA, June 10-15, 2018*, 355–368. ACM.
- ten Cate, B.; Dalmau, V.; and Kolaitis, P. G. 2013. Learning schema mappings. *ACM Trans. Database Syst.* 38(4):28:1–28:31.
- Davis, W. L.; Schwarz, P. M.; and Terzi, E. 2009. Finding representative association rules from large rule collections. In *Proceedings of the SIAM International Conference on Data Mining, SDM 2009, April 30 - May 2, 2009, Sparks, Nevada, USA*, 521–532. SIAM.
- de Amo, S.; Diallo, M. S.; Diop, C. T.; Giacometti, A.; Li, D. H.; and Soulet, A. 2015. Contextual preference mining for user profile construction. *Inf. Syst.* 49:182–199.
- Fagin, R.; Kolaitis, P. G.; Miller, R. J.; and Popa, L. 2005. Data exchange: semantics and query answering. *Theoretical Computer Science* 336(1):89 – 124. Database Theory.
- Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman.
- Gottlob, G., and Senellart, P. 2010. Schema mapping discovery from data instances. *J. ACM* 57(2):6:1–6:37.
- Karp, R. 1972. Reducibility among combinatorial problems. In Miller, R., and Thatcher, J., eds., *Complexity of Computer Computations*. Plenum Press. 85–103.
- Kimmig, A.; Memory, A.; Miller, R. J.; and Getoor, L. 2017. A collective, probabilistic approach to schema mapping. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, 921–932.
- Legriel, J.; Guernic, C. L.; Cotton, S.; and Maler, O. 2010. Approximating the pareto front of multi-criteria optimization problems. In Esparza, J., and Majumdar, R., eds., *Tools and Algorithms for the Construction and Analysis of Systems, 16th International Conference, TACAS 2010, Proceedings*, volume 6015 of *Lecture Notes in Computer Science*, 69–83. Springer.
- Lenzerini, M. 2002. Data integration: A theoretical perspective. In Popa, L.; Abiteboul, S.; and Kolaitis, P. G., eds., *Proceedings of the Twenty-first ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 3-5, Madison, Wisconsin, USA*, 233–246. ACM.
- Miettinen, P. 2008. On the positive-negative partial set cover problem. *Information Processing Letters* 108(4):219 – 221.
- Papadimitriou, C. H., and Yannakakis, M. 1982. The complexity of facets (and some facets of complexity). In Lewis, H. R.; Simons, B. B.; Burkhard, W. A.; and Landweber, L. H., eds., *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, May 5-7, 1982, San Francisco, California, USA*, 255–260. ACM.
- Papadimitriou, C. H., and Yannakakis, M. 1991. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences* 43(3):425 – 440.
- Peleg, D. 2007. Approximation algorithms for the label-covermax and red-blue set cover problems. *Journal of Discrete Algorithms* 5(1):55 – 64.
- Qian, K.; Popa, L.; and Sen, P. 2017. Active learning for large-scale entity resolution. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, 1379–1388. New York, NY, USA: ACM.
- Sarma, A. D.; Parameswaran, A. G.; Garcia-Molina, H.; and Widom, J. 2010. Synthesizing view definitions from data. In Segoufin, L., ed., *Database Theory - ICDT 2010, 13th International Conference, Proceedings*, ACM International Conference Proceeding Series, 89–103. ACM.
- Singla, P., and Domingos, P. 2006. Entity Resolution with Markov Logic. In *ICDM*, 572–582.