

Sets of Robust Rules, and How to Find Them

Jonas Fischer (✉) and Jilles Vreeken

¹ Max Planck Institute for Informatics, and Saarland University, Germany

² CISPA Helmholtz Center for Information Security, Germany

Abstract. Association rules are among the most important concepts in data mining. Rules of the form $X \rightarrow Y$ are simple to understand, simple to act upon, yet can model important local dependencies in data. The problem is, however, that there are so many of them. Both traditional and state-of-the-art frameworks typically yield millions of rules, rather than identifying a small set of rules that capture the most important dependencies of the data. In this paper, we define the problem of association rule mining in terms of the Minimum Description Length principle. That is, we identify the best *set of rules* as the one that most succinctly describes the data. We show that the resulting optimization problem does not lend itself for exact search, and hence propose GRAB, a greedy heuristic to efficiently discover good sets of noise-resistant rules directly from data. Through extensive experiments we show that, unlike the state-of-the-art, GRAB does reliably recover the ground truth. On real world data we show it finds reasonable numbers of rules, that upon close inspection give clear insight in the local distribution of the data.

1 Introduction

Association rules are perhaps the most important primitive in data mining. Rules of the form $X \rightarrow Y$ are not only simple to understand, but they are also simple to act upon, and, most importantly, can express important *local* structure in the data. The problem is, however, that there are so many of them, and that telling the interesting from the uninteresting rules has so far proven impossible. Both traditional algorithms based on support and confidence [1], as well as modern approaches based on statistical tests [7] typically discover orders of magnitude more rules than the data has rows – even when the data consists of nothing but noise. In this paper we show how to discover a small, non-redundant set of noise-resistant rules that together describe the data well.

To succinctly express subtly different structures in data, we allow *multiple* items in the consequent of a rule. To illustrate, while rule sets $R_1 = \{A \rightarrow B, A \rightarrow C\}$ and $R_3 = \{A \rightarrow BC\}$ both express that B and C appear frequently in the context of A , the former states they do so independently, while the latter expresses a dependency between B and C . We additionally allow for *patterns*, which are simply rules like $R_4 = \{\emptyset \rightarrow ABC\}$ and express unconditional dependencies. Real data is often noisy, and hence we can allow rules to hold *approximately*. That is, for a transaction $t = ABC$, our models may infer




Fig. 1: *Five toy databases with corresponding rule sets* 1) B and C occur in the context of A but independently of each other, 2) C occurs in the context of B , which in turn occurs in the context of A , 3) B and C show strong joint dependence in the context of A , 4) A, B, C show strong unconditional dependence, and 5) a rule with noise, BCD occurring jointly in the context of A .

that rule $R_5 = \{A \rightarrow BCD\}$ holds, even though item D is not present in t . We call these noise-resistant, or *robust* rules. To determine the quality of a rule set for given data, we rely on information theory.

In particular, we define the rule set mining problem in terms of the Minimum Description Length (MDL) principle [6]. Loosely speaking, this means we identify the best rule set as that one that compresses the data best. This set is naturally non-redundant, and neither under- nor over-fitting, as we have to pay for every additional rule we use, as well as for every error we make. We formally show that the resulting problem is neither submodular, nor monotone, and as the search space is enormous, we propose GRAB, an efficient any-time algorithm to heuristically discover good rule sets directly from data. Starting from a singleton-only model, we iteratively refine our model by considering combinations of rules in the current model. Using efficiently computable tight estimates we minimize the number of candidate evaluations, and as the experiments show, GRAB is both fast in practice, and yields high quality rule sets. On synthetic data, GRAB recovers the ground truth, and on real-world data it recovers succinct models of meaningful rules. In comparison, state of the art methods discover up to several millions of rules for the same data, and are hence hardly useful.

2 Related Work

Pattern mining is arguably one of the most important and well-studied topics in data mining. We aim to give a succinct overview of the work most relevant to ours. The first, and perhaps most relevant proposal is that of association rule mining [1], where in an unsupervised manner the goal is to find all rules of the form $X \rightarrow Y$ from the data that have high frequency and high confidence. As it turns out to be straightforward to distill the high-confidence rules

from a given frequent itemset, research attention shifted to discovering frequent itemsets efficiently [14,31,8], and non-redundantly [2,3,17]. Frequency alone is a bad measure of interestingness, however, as it leads to spurious patterns [26]. To alleviate this, statistically sound measures were proposed that can mine patterns with frequencies that deviate significantly from our expectation based on margins [29,21], or richer background knowledge [9,24,4]. Perhaps because it is already difficult enough to determine the interestingness of a pattern, let alone a rule, most proposals restrict themselves to patterns. The key exception is KINGFISHER, which proposes an upper bound for Fisher’s exact test that allows to efficiently mine significant dependency rules using the branch-and-bound framework [7]. Notably, however, KINGFISHER can only discover exact rules with a single item consequent. In addition, all these approaches suffer from the problems of multiple test correction, and return all patterns they deem significant, rather than a small non-redundant set.

Less directly related to our problem setting, but still relevant, are supervised methods that discover rules that explain a given target variable Y . Zimmermann and Nijssen [32] give a good general overview. However, unlike Wang et al. [28] and Papaxanthos et al. [19], we are not interested in rules that explain only Y , but rather aim for a set of rules that together explains all of the data well.

Our approach is therewith a clear example of pattern set mining [26]. That is, rather than measuring the quality of individual patterns, we measure quality over a set of patterns [27,5]. Information theoretic approaches, such as MDL and the Maximum Entropy principle, have proven particularly successful for measuring the quality of sets of patterns [27,13]. Most pattern set approaches do not account for noise in the data, with ASSO [15], HYPER+ [30], and PANDA [12] as notable exceptions. However, extending any of the above from patterns to rules turns out to be far from trivial, because rules have different semantics than patterns. PACK [25] uses MDL to mine a small decision tree per item in the data, and while not technically a rule-mining method, we can interpret the paths of these trees as rules. In our experiments we will compare to KINGFISHER as the state-of-the-art rule miner, HYPER+ as a representative of noise resilient pattern miner, and PACK as a pattern miner, which output can be translated into rules.

3 Preliminaries

In this section we discuss preliminaries and introduce notation.

3.1 Notation

We consider binary transaction data D of size n -by- m , with $n = |D|$ transactions over an alphabet \mathcal{I} of $m = |\mathcal{I}|$ items. In general, we denote sets of items as $X \subseteq \mathcal{I}$. A transaction t is an itemset, e.g. the products bought by a customer. We write $\pi_X(D) := \{t \cap X \mid t \in D\}$ for the projection of D on itemset X . The transaction set, or selection, T of itemset X is the multiset of all transactions $t \in D$ that contain X , i.e. $T_X = \{t \in D \mid X \subseteq t\}$. We write $n_X = |T_X|$ to denote

the cardinality of a transaction multiset. The *support* of an itemset X is then simply the number of transactions in D that contain X , i.e. $\text{support}(X) = |T_X|$.

An association rule $X \rightarrow Y$ consists of two non-intersecting itemsets, the antecedent or head X , and consequent or tail Y . A rule makes a statement about the conditional occurrence of Y in the data where X holds. If $X = \emptyset$, we can interpret a rule as a pattern, as it makes a statement on where in the whole data the consequent holds. Throughout this manuscript, we will use A, B, C to refer to sets of single items and X, Y, Z for itemsets of larger cardinality.

3.2 Minimum Description Length

The Minimum Description Length (MDL) principle [22] is a computable and statistically well-founded approximation of Kolmogorov Complexity [11]. For given data D , MDL identifies the best model M^* in a given model class \mathcal{M} as that model that yields the best lossless compression. In one-part, or, *refined* MDL we consider the length in bits of describing data D using the entire model class, $L(D | \mathcal{M})$, which gives strong optimality guarantees [6] but is only feasible for certain model classes. In practice we hence often use two-part, or, *crude* MDL, which is defined as $L(M) + L(D | M)$. Here $L(M)$ is the length of the description of the model, and $L(D | M)$ the length in bits of the description of the data using M . We will use two-part codes where we have to, and one-part codes where we can. Note that in MDL we are only concerned with code *lengths*, not materialized codes. Also, as we are interested in measuring lengths in bits, all logarithms are to base 2, and we follow the convention $0 \log 0 = 0$.

4 Theory

To use MDL in practice, we first need to define our model class \mathcal{M} , how to describe a model M in bits, and how to describe data D using a model M . Before we do so formally, we first give the intuitions.

4.1 The Problem, Informally

Our goal is to find a set of rules that together succinctly describe the given data. Our models M hence correspond to sets R of rules $X \rightarrow Y$. A pattern ABC is simply a rule with an empty head, i.e. $\emptyset \rightarrow ABC$. A rule *applies* to a transaction $t \in D$ if the transaction supports its head, i.e. $X \subseteq t$. For each transaction to which the rule applies, the model specifies whether the rule *holds*, i.e. whether Y is present according to the model. We can either be strict, and require that rules only hold when $Y \subseteq t$, or, be more robust to noise and allow the rule to hold even when not all items of Y are part of t , i.e. $Y \setminus t \neq \emptyset$. In this setting, the model may state that rule $A \rightarrow BCD$ holds for transaction $t = ABC$, even though $D \notin t$ (see Fig. 1.5). A model M hence needs to specify for every rule $X \rightarrow Y \in R$ a set of transactions ids $T_{Y|X}^M$ where it asserts that Y holds in the context of X , and, implicitly also $T_{\bar{Y}|X}^M$, the set of transactions where it asserts

Y does not hold. Last, for both these we have to transmit which items of Y are actually in the data; the fewer errors we make here, the cheaper it will be to transmit. To ensure that we encode any data D over \mathcal{I} , we require that a model M contains at least singleton rules, i.e. $\emptyset \rightarrow A$ for all $A \in \mathcal{I}$. Cyclic dependencies would prevent us from decoding the data without loss. Any valid model M can hence be represented as a directed acyclic graph (DAG), in which the vertices of the graph correspond to rules in R , where vertex $r = X \rightarrow Y$ has incoming edges from all vertices $r' = X' \rightarrow Y'$ for which $X \cap Y'$ is non-empty.

We explicitly allow for rules with non-singleton tails, as this allows us to succinctly describe subtly different types of structure. When B happens independently of C in T_A (Fig. 1.1), rule set $R_1 = \{A \rightarrow B, A \rightarrow C\}$ is a good description of this phenomenon. In turn, when C occurs often – but not always – in T_B , which in turn happens often in T_A (Fig. 1.2) rule set $R_2 = \{A \rightarrow B, B \rightarrow C\}$ is a good description. To succinctly describe that B and C are statistically dependent in T_A (Fig. 1.3) we need rules with multiple items in its tail, i.e. $R_3 = \{A \rightarrow BC\}$. Finally, if A, B , and C frequently occur jointly, but conditionally independent of any other variable, we need patterns to express this, which are just consequents in the context of the whole database $R_4 = \emptyset \rightarrow ABC$.

4.2 MDL for Rule Sets

Next, we formalize an MDL score for the above intuition. We start by defining the cost of the data given a model, and then define the cost of a model.

Cost of the data We start with the cost of the data described by an individual rule $X \rightarrow Y$. For now, assume we know $\pi_X(D)$ and T_X . We transmit the data over Y in the context of X , i.e. $D_{Y|X} = \pi_Y(T_X)$, in three parts. First, we transmit the transaction ids where model M specifies that both X and Y hold, $T_{Y|X}^M$, which implicitly gives $T_{\bar{Y}|X}^M = T_X \setminus T_{Y|X}^M$. We now, in turn transmit that part of $D_{Y|X}$ corresponding to the transactions in $T_{Y|X}^M$, resp. that part corresponding to $T_{\bar{Y}|X}^M$. We do so using optimal data-to-model codes, i.e. indices over canonically ordered enumerations,

$$L(D_{Y|X} | M) = \log \binom{|T_X|}{|T_{Y|X}^M|} + \log \binom{|T_{Y|X}^M| \times |Y|}{\mathbf{1}(T_{Y|X}^M)} + \log \binom{|T_{\bar{Y}|X}^M| \times |Y|}{\mathbf{1}(T_{\bar{Y}|X}^M)},$$

where we write $\mathbf{1}(T_{Y|X}^M)$ for the number of 1s in $T_{Y|X}^M$, i.e.

$$\mathbf{1}(T_{Y|X}^M) = \sum_{t \in T_{Y|X}^M} |t \cap Y| \leq |T_{Y|X}^M| \times |Y|.$$

We define $\mathbf{1}(T_{\bar{Y}|X})$ analogue.

When the model makes exact assertions on Y holding when X is present, i.e. when $T_{Y|X}^M = T_{Y|X}$, the second term vanishes, and analogously for the third term when $T_{\bar{Y}|X}^M = T_{\bar{Y}|X}$. Both terms vanish simultaneously only when $D_{Y|X} \in \{\emptyset, Y\}^{|D_X|}$. This is trivially the case when Y is a singleton.

The overall cost of the data given the model simply is the sum of the data costs per rule,

$$L(D | M) = \sum_{X \rightarrow Y \in M} L(D_{Y|X} | M)$$

To decode the data, the recipient will of course need to know each rule $X \rightarrow Y$. These are part of the model cost.

Cost of the Model To encode a rule, we first encode the cardinalities of X and Y using $L_{\mathbb{N}}$, the MDL-optimal code for integers $z \geq 1$, which is defined as $L_{\mathbb{N}}(z) = \log^* z + \log c_0$, where $\log^* z = \log z + \log \log z + \dots$, and c_0 is a normalization constant such that $L_{\mathbb{N}}$ satisfies the Kraft-inequality [23]. We can now encode the items of X , resp. Y , one by one using optimal prefix codes, $L(X) = -\sum_{x \in X} \log \frac{s_x}{\sum_{i \in \mathcal{I}} s_i}$. Last, but not least we have to encode its parameters, $|T_{Y|X}^M|$, $\mathbb{1}(T_{Y|X}^M)$, and $\mathbb{1}(T_{Y|X})$. These we encode using a refined, mini-max optimal MDL code. In particular, we use the regret of the Normalized Maximum Likelihood code length [10] for the class of binomials,

$$L_{pc}(n) = \log \left(\sum_{k=0}^n \frac{n!}{(n-k)!k!} \left(\frac{k}{n}\right)^k \left(\frac{n-k}{n}\right)^{n-k} \right),$$

which is also known as the parametric complexity of a model class. Kontkanen and Myllymäki [10] showed that this term can be computed in time $O(n)$ in a recursive manner. We obtain the model cost $L(X \rightarrow Y)$ for a rule $X \rightarrow Y$ by

$$\begin{aligned} L(X \rightarrow Y) &= L_{\mathbb{N}}(|X|) + L(X) + L_{\mathbb{N}}(|Y|) + L(Y) + \\ &L_{pc}(|T_X|) + L_{pc}(|T_{Y|X}^M| \times |Y|) + L_{pc}(|T_{Y|X}^M| \times |Y|). \end{aligned}$$

From how we encode the data we can simply ignore the last two terms for rules with $|Y| = 1$. The overall cost of a model M then amounts to

$$L(M) = L_{\mathbb{N}}(|R|) + \sum_{X \rightarrow Y \in R} L(X \rightarrow Y),$$

where we first send the size of rule set R , and then each of the rules in order defined by the spanning tree of the dependency graph.

4.3 The Problem, formally

We can now formally define the problem in terms of MDL.

Definition 1 (Minimal Rule Set Problem). *Given data D over items \mathcal{I} , find that rule set R and that set of \mathcal{T} of transaction sets $T_{Y|X}^M$ for all $X \rightarrow Y \in R$, such that for model $M = (R, T)$ the total description length,*

$$L(D, M) = L(M) + L(D | M)$$

is minimal.

Solving this problem involves enumerating all possible models $M \in \mathcal{M}$. There exist $\sum_{i=0}^{|\mathcal{I}|} \binom{|\mathcal{I}|}{i} \times 2^i = 3^{|\mathcal{I}|}$ possible rules – where the second term in the sum describes all possible partitions of i items into head and tail, and the equality is given by the binomial theorem. Assuming that the optimal $T_{Y|X}^M$ are given, there are generally $2^{3^{|\mathcal{I}|}}$ possible models. The search space does not exhibit any evident structure that can be leveraged to guide the search, which is captured by the following two theorems. We postpone the proofs to the online Appendix.³

Theorem 1 (Submodularity). *The search space of all possible sets of association rules 2^Ω , when fixing a dataset and using the description length $L(D, M)$ as set function, is not submodular. That is, there exists a data set D s.t. $\exists X \subset Y \subseteq \Omega, z \in \Omega. L(D, X \cup \{z\}) - L(D, X) \leq L(D, Y \cup \{z\}) - L(D, Y)$.*

Theorem 2 (Monotonicity). *The description length $L(D, M)$ on the space of all possible sets of association rules 2^Ω is not monotonously decreasing. That is, there exists a data set D s.t. $\exists X \subset Y \subseteq \Omega. f(X) \leq f(Y)$.*

Hence, we resort to heuristics.

5 Algorithm

In this section we introduce GRAB, an efficient heuristic for discovering good solutions to the Minimal Rule Set Problem. GRAB consists of two steps, candidate generation and evaluation, that are executed iteratively until convergence of $L(D, M)$, starting with the singleton-only rule set $R_0 = \{\emptyset \rightarrow A \mid A \in \mathcal{I}\}$.

Candidate generation From the current rule set R we iteratively discover that refined rule set R' that minimizes the gain $\Delta L = L(D, M') - L(D, M)$. As refinements we consider the combination of two existing rules into a new rule.

We generate candidate refinements by considering all pairs $r_1 = X \rightarrow Y, r_2 = X \rightarrow Z \in R$, assuming w.l.o.g. $n_{XY} \geq n_{XZ}$, and merging the tails of r_1 and r_2 to obtain candidate rule $r'_1 = X \rightarrow YZ$, and merging the tail of r_1 with the head to obtain candidate rule $r'_2 = XY \rightarrow Z$. We now construct refined rule sets R'_1 and R'_2 by adding rule r'_1 resp. r'_2 . To reduce redundancy, we remove r_2 from both R'_1 and R'_2 , and r_1 from R'_1 , taking care not to remove singleton rules. We only evaluate those refined rule sets R' for which the corresponding dependency graph is acyclic, and select the one with minimal gain $\Delta L < 0$. For completeness we give the pseudocode in the online Appendix.

Gain estimation To avoid naively evaluating the gain ΔL of every candidate, we rely on accurate gain estimations. In particular, we consider two different estimates; the first estimate is very inexpensive to compute, but overly optimistic as it assumes a perfect overlap between the two rules. The second estimate is computationally more costly, as it requires us to compute the intersection of the selections of the two original rules. In practice, however, it is exact (see Fig. 2b).

³ <http://eda.mmci.uni-saarland.de/grab/>




Fig. 2: GRAB *searches efficiently and estimates accurately*. For DNA we show (left) the convergence of the relative compression of model M at iteration i against the singleton model M_s , $\%L = \frac{L(D, M) \times 100}{L(D, M_s)}$, and (right) the correlation between estimated and actual gain of all evaluated candidates in real data.

Depending on how we combine two rules r_1 and r_2 , we need different estimate definitions. In the interest of space, we here consider one case in detail: that of combining singleton rules $r_1 = \emptyset \rightarrow A$ and $r_2 = \emptyset \rightarrow B$ into $r = A \rightarrow B$. For the remaining definitions we refer to the online Appendix.

Following the general scheme described above, for the first estimate $\hat{\Delta}_1$ we assume that $T_B \subseteq T_A$. With singleton tails we do not transmit any errors. Thus, we only subtract the old costs for r_2 and add the estimated cost of sending where the new rule r holds, as well as the estimated regret for the new matrices,

$$\hat{\Delta}_1(r) = -\log \binom{n}{n_B} + \log \binom{n_A}{n_B} + L_{pc}(n_A) + L_{pc}(n_B) + L_{pc}(n_A - n_B) - L_{pc}(n).$$

For the tighter, second estimate $\hat{\Delta}_2$ we instead need to retrieve the exact number of usages of the rule by intersecting the transaction sets of merged rules. The change in model costs $L(M)$ by introducing r appearing in $L(M)$ is trivially computable and thus abbreviated by $\hat{L}(M)$. For formerly covered transactions that are not covered by the new rule, we need to send singleton rules with adapted costs, which is estimated through simple set operations on the transaction sets. Additionally, we need to subtract the model costs for r_2 , in case B is completely covered by r , ensured by the indicator variable I . We hence have

$$\begin{aligned} \hat{\Delta}_2(r) = & -\log \binom{n}{n_B} + \log \binom{n_A}{|T_A \cap T_B|} + \log \binom{n}{|T_B \setminus T_A|} + \hat{L}(M) + L_{pc}(n_A) \\ & + L_{pc}(|T_A \cap T_B|) + L_{pc}(n_A - |T_A \cap T_B|) - I(T_B \subseteq T_A) \times L_{pc}(n). \end{aligned}$$

GRAB first computes the first order estimate $\hat{\Delta}_1$ per candidate, and only if this shows potential improvement, it computes the second order estimate $\hat{\Delta}_2$. Out of those, it evaluates all candidates that have the potential to improve over

the best refinement found so far. In the next paragraph we describe how to efficiently compute the overall score $L(D, M)$.

Efficiently computing $L(D, M)$ To get the code length of a rule set with a new candidate, two steps are carried out, which we summarize in Alg. 1. First, the data is covered with the new rule to determine where the rule holds and what error matrices to send. Covering the data is straightforward, but to find the error matrices we have—unless we rely on a user-defined threshold—to optimize for the best point to split between additive and destructive noise. We observe that each rule encoding is independent of every other rule (except singletons), that is, changing the error matrices for one rule does not change the matrices for any other rule as we always encode all transactions where the antecedent is fulfilled.

With this in mind, it is clear that we can optimize the split point for each rule $X \rightarrow Y$ separately. Thus, we find a partitioning of T_X into $T_{Y|X}^M$ and $T_{\bar{Y}|X}^M$ that minimizes the contribution of this rule to the overall costs:

$$\begin{aligned} \Delta_{T_X, T_{Y|X}^M, T_{\bar{Y}|X}^M, \mathbb{1}(T_{Y|X}^M), \mathbb{1}(T_{\bar{Y}|X}^M)} &= L_{pc}(|T_X|) + L_{pc}(|T_{Y|X}^M| \times |Y|) \\ &+ L_{pc}(|T_{\bar{Y}|X}^M| \times |Y|) + \log \left(\frac{|T_{Y|X}^M| \times |Y|}{\mathbb{1}(T_{Y|X}^M)} \right) + \log \left(\frac{|T_{\bar{Y}|X}^M| \times |Y|}{\mathbb{1}(T_{\bar{Y}|X}^M)} \right). \end{aligned}$$

We can also view the problem from a different angle, namely, for each transaction $t \in T_X$ we count how many items of Y are present, which yields a vector of counts B , $B[i] = |\{t \in T_X \mid |t \cap Y| = i\}|$. For fixed split point k , we get the additive and destructive matrix sizes $\mathbb{1}(\cdot)^k$ and transaction set sizes $|\cdot|^k$:

$$\begin{aligned} |T_{Y|X}^M|^k &:= \sum_{i=k}^{|B|+1} B[i] & |T_{\bar{Y}|X}^M|^k &:= \sum_{i=1}^{k-1} B[i] \\ \mathbb{1}(T_{Y|X}^M)^k &:= \sum_{i=k}^{|B|} B[i] \times i & \mathbb{1}(T_{\bar{Y}|X}^M)^k &:= \sum_{i=0}^{k-1} B[i] \times i. \end{aligned}$$

To find the best split k^* we optimize along k using the two equation sets above, which is in time linear in the size of the consequent,

$$k^* = \operatorname{argmin}_{k=1 \dots |B|} \left(\Delta_{T_X, T_{Y|X}^M, T_{\bar{Y}|X}^M, \mathbb{1}(T_{Y|X}^M), \mathbb{1}(T_{\bar{Y}|X}^M)} \right). \quad (1)$$

This yields the best splitpoint k^* for how many items of the consequent are required for a rule to *hold* in terms of our MDL score and thus implicitly gives the error matrices.

Putting everything together, we have GRAB, given in pseudo-code as Alg. 2.

Complexity In the worst case we generate all pairs of combinations of rules, and hence at each step GRAB evaluates a number of candidates quadratic in the size of the rule table. Each evaluation of the $O(3^{2^m})$ candidates requires a database cover which costs time $O(n \times m)$, and singleton transaction set update,

Algorithm 1: COVER

```

input : Database  $D$ , model  $M = (R, \mathcal{T})$ , refined rule set  $R'$ 
output : Model  $M' = (R', \mathcal{T}')$ 
1  $T_{Y|X}^{M'} \leftarrow$  according to Equation (1); // Initialize where new rule holds
2 for  $I \in \mathcal{I}$  do // For each singleton
3    $T_I^{M'} \leftarrow T_I$ ; // Reset usage to baseline model
4   for  $\{X \rightarrow Y \in R' \mid I \in Y\}$  do // For each rule tail containing  $I$ 
5      $T_I^{M'} \leftarrow T_I^{M'} \setminus T_X$ ; // Remove these transactions from list
6 return  $(R', \{T_I^{M'} \mid I \in \mathcal{I}\} \cup \{T_{U|V}^M \in \mathcal{T} \mid U \rightarrow V \in R \cap R'\} \cup \{T_{Y|X}^{M'}\})$ ;

```

Algorithm 2: GRAB

```

input : Dataset  $D$ 
output : Heuristic approximation to  $M$ 
1  $M \leftarrow \{\emptyset \rightarrow \{A\} \mid A \in \mathcal{I}\}$ ; // Initialize model with singletons
2 do
3    $C \leftarrow \text{GENERATECANDIDATES}(D, M)$ ;
4    $M^* \leftarrow M$ ;  $\Delta^* \leftarrow 0$ ;
5   while  $C$  contains a refinement  $R$  with  $\widehat{\Delta}_2 < \Delta^*$  do
6      $R' \leftarrow$  refinement  $R \in C$  with best  $\widehat{\Delta}_2$ ;
7      $M' \leftarrow \text{COVER}(D, M, R')$ ; // Construct model  $M'$ 
8      $\Delta' \leftarrow L(D, M') - L(D, M)$ ; // Compute exact gain
9     if  $\Delta' < \Delta^*$  then
10       $M^* \leftarrow M'$ ;  $\Delta^* \leftarrow \Delta'$ ;
11   if  $M^* \neq M$  then // Update best model
12      $M \leftarrow M^*$ ;
13 while  $L(D, M) < L(D, M^*)$ ;
14 return  $M$ 


```

thus giving an overall time in $O(3^{2^m} \times m \times n)$. However, MDL ensures that the number of rules is small, and hence a more useful statement about runtime is given in the following theorems that are based on the size of the output or in other words the number of mined rules. For the proofs, see the online Appendix.

Theorem 3 (Grab candidate evaluations). *Given that we mine k rules for a given dataset D , GRAB evaluates $O((m+k)^3)$ candidates.*

This theorem gives us insight in how many times GRAB calls COVER. For the runtime analysis, we know that in each step i our rule table has size $m+i$ and GRAB has to compute the cover of the newest rule in time $O(n \times m)$ and update the singleton costs in time $O((m+i) \times m \times n)$.

Theorem 4 (Grab runtime). *Given that we mine k rules for a given dataset D , the overall runtime of GRAB is $O((m+k)^4 \times m \times n)$.*



(a) Rules found in pure noise data of different sparsity. GRAB and PACK do not find any rule at all, whereas KINGFISHER and HYPER+ pick up random noise. (b) Rules found in synthetic data. GRAB closely approximates the ground truth number of planted rules (black, dashed).

Fig. 3: GRAB *recovers ground truth*. Results on random data (left) and synthetic data with planted rules for different data dimensionalities (right).

In practice, however, this runtime is never reached both due to our gain estimates and because we only allow to merge rules with the same head.

6 Experiments

In this section we empirically evaluate GRAB quantitatively and qualitatively on both synthetic and real-world data. We implemented GRAB in C++. We make all code and data available for research purposes.⁴ All experiments were executed single-threaded on Intel Xeon E5-2643 v3 machines with 256 GB memory running Linux. We report the wall-clock running times.

We compare to state of the art methods for mining statistically significant patterns and association rules. In particular, we compare to HYPER+ [30], which mines noise-resistant patterns, KINGFISHER [7], which is arguably the current state of the art for mining statistically significant rules under the Fisher-exact-test⁵, and PACK [25], an MDL-based method that yields a binary tree per item $A \in \mathcal{I}$ of which we can interpret the paths to leaves as rules $X \rightarrow A$.

Synthetic data First, we consider data with known ground truth. As a sanity check, we start our experiments on data without any structure. We draw datasets of 10000-by-100 of $d\%$ 1s, and report for each method the average results over 10 independent runs in Fig. 3a. We find that both KINGFISHER and HYPER+ quickly discover up to millions of rules. This is easily explained, as the former

⁴ <http://eda.mmci.uni-saarland.de/grab/>

⁵ No relation to the first author.

relies on statistical significance only, and lacks a notion of support, whereas the latter does have a notion of support, but lacks a notion of significance. PACK and GRAB, however, retrieve the ground truth in all cases.

Next, we consider synthetic data with planted rules. We generate datasets of $n = 20000$ transactions, and vary m from 10 to 1000 items. We generate rules that together cover all features. We sample the cardinality of the heads and tails from a Poisson with $\lambda = 1.5$. To avoid convoluting the ground truth via overlap, or by rules forming chains, we ensure that every item A is used in at most one rule. Per rule, we choose confidence c uniformly at random between 50 and 100%. We then randomly partition the n transactions into as many parts as we have rules, and per part, set the items of the corresponding rule head X to 1, and set Y to 1 for $c\%$ of transactions within the part. Finally, we add noise by flipping 1% of the items in the data – we use this low noise level to allow for a fair comparison to the competitors that do not explicitly model noise.

We provide the results in Fig. 3b. We observe that unlike in the previous experiment, here PACK strongly overestimates the number of rules – it runs out of memory for data of more than 92 features. KINGFISHER and HYPER+ both discover over an order of magnitude more rules than the ground truth. GRAB, on the other hand, is the only one that reliably retrieves the ground truth.

Real-World Data Second, we verify whether GRAB also yields meaningful results on real data. To this end we consider 8 data sets over a variety of domains. In particular, from the UCI repository we consider *Mushroom*, *Adult*, *Covtype*, and *Plants*. In addition we use data of Belgium traffic *Accidents*, *DNA* amplification [18], *Mammals* [16], and *ICDM Abstracts* [25]. We give basic statistics in Table 1, and provide more detailed information in the online Appendix.

We run each of the methods on each data set, and report the number of discovered non-singleton rules for all methods and the average number of items in head and tail for GRAB in Table 1. We observe that GRAB retrieves much more succinct sets of rules than its competitors, typically in the order of tens, rather than in the order of thousands to millions. The rules that GRAB discovers are also more informative, as it is not constrained to singleton-tail rules. This is also reflected by the number of items in the consequent, where the average tail size is much larger than 1 for e.g. *Mammals* and *Plants*, where we find multiple rules with more than 10 items in the consequent.

To qualitatively evaluate the rules that GRAB discovers, we investigate the results on *Abstracts* and *Mammals* in closer detail. For *Abstracts* we find patterns such as $\emptyset \rightarrow \{\textit{naive}, \textit{bayes}\}$, $\emptyset \rightarrow \{\textit{nearest}, \textit{neighbor}\}$, $\emptyset \rightarrow \{\textit{pattern}, \textit{frequency}\}$, and, notably, $\emptyset \rightarrow \{\textit{association}, \textit{rule}\}$. Further, we find meaningful rules, including $\{\textit{high}\} \rightarrow \{\textit{dimension}\}$, $\{\textit{knowledge}\} \rightarrow \{\textit{discovery}\}$, $\{\textit{ensembl}\} \rightarrow \{\textit{bagging}, \textit{boosting}\}$, and $\{\textit{support}\} \rightarrow \{\textit{vector}, \textit{machin}, \textit{SVM}\}$. All patterns and rules correspond to well-known concepts in the data mining community.

On *Mammals*, GRAB finds large patterns such as $\emptyset \rightarrow \{\textit{red deer}, \textit{European mole}, \textit{European fitch}, \textit{wild boar}, \textit{marten}, \textit{mice}^*\}$, and $\emptyset \rightarrow \{\textit{common squirrel}, \textit{deer}, \textit{ermine}, \textit{marten}, \textit{mice}^*\}$, that correspond to animals that commonly occur across Europe, with multiple mouse species (items) indicated by \textit{mice}^* . In addi-

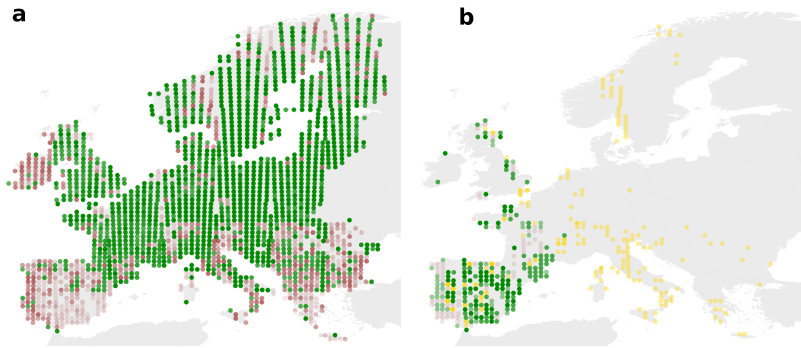


Fig. 4: *Example Rules for Mammals* Shown are the inferred presence (green) and absence (red) of **a** pattern $\emptyset \rightarrow \{\textit{common squirrel}, \textit{deer}, \textit{ermine}, \textit{marten}, \textit{mice}^*\}$ and **b** rule $\{\textit{Southwest European cat}\} \rightarrow \{\textit{Mediterranean mice}^*, \textit{Iberian rabbit}\}$. The intensity of the colour indicates how many items of the tail hold – the ideal result is hence dark green and light red. Yellow dots indicate presence of animals from tail of rule where animals of head of rule were not sighted.

tion, it also discovers specific patterns, e.g. $\emptyset \rightarrow \{\textit{snow rabbit}, \textit{elk}, \textit{lynx}, \textit{brown bear}\}$, which are mammals that appear almost exclusively in northeastern Europe. We visualized the second rule in Figure 4a to show that the consequent should hold in most of the cases, but not necessarily need to be always present. Moreover, GRAB is able to find meaningful rules in the presence of noise, e.g. $\{\textit{Southwest European cat}\} \rightarrow \{\textit{Mediterranean mice}^*, \textit{Iberian rabbit}\}$, where the rule should only hold in southwest Europe. For the rule that GRAB discovers this is indeed the case, although the data contains (likely spurious) sightings of Iberian rabbits or Mediterranean mice in Norway (see Fig. 4b) and some sightings of mice alone, along the Mediterranean sea.

Runtime and Scalability Last, but not least, we investigate the runtime of GRAB. We first consider scalability with regard to number of features. For this, in Fig. 5a we give the runtimes for the synthetic datasets we used above. From the figure we see that while GRAB is not as fast as KINGFISHER and HYPER+, it scales favourably with regard to the number of features. Although it considers a much larger search space, GRAB only needs seconds to minutes. On real data GRAB is the fastest method for five of the data sets, and only requires seconds for the other datasets, whereas the other methods take up to hours for particular instances (compare Figure 5b).

7 Discussion

The experiments show that GRAB is fast and returns crisp, informative rule sets. On synthetic data it recovers the ground truth, without picking up noise. On

Dataset	n	m	GRAB			HYPERT+	KINGFISHER	PACK
			$ \overline{X} $	$ \overline{Y} $	$ R $	Number of Discovered Rules		
Abstracts	859	3933	0.9	1.2	29	1508	42K	334
Accidents	339898	468	1	1.1	138	65M	<i>mem</i>	69M
Adult	10830	97	1	1.1	27	26K	9K	68M
Covtype	581012	105	1.3	1.1	41	13K	43K	286M
DNA	1316	392	1	1.7	147	49	140K	451
Mammals	2183	121	1.5	2	38	<i>mem</i>	$\geq 10M$	2K
Mushroom	8124	119	1.6	1.5	65	13K	81K	7K
Plants	34781	69	1.2	3.2	20	6M	<i>mem</i>	910

Table 1: For GRAB, the size of the rule set and average size of head $|\overline{X}|$ and tail $|\overline{Y}|$ are given. For the other methods, number of found rules are given, *mem* indicates an aborted run due to memory usage $> 256\text{GB}$.

real world data, it retrieves concise and easily interpretable rule sets, as opposed to the state of the art that discovers thousands, up to millions of rules.

The results on the *Mammals* data clearly show GRAB recovers known population structures, even in the presence of noise. The results on the ICDM *Abstracts* data are equally good, with rule $\{\textit{support}\} \rightarrow \{\textit{vector}, \textit{machin}, \textit{svm}\}$ as a notable example. In contrast to machine learning, in data mining “support” is ambiguous. In the ICDM abstracts it means the support of a pattern, as well as support vector machines, and the rule expresses this. To verify this, we additionally ran GRAB on abstracts from the Journal of Machine Learning Research (JMLR), where it instead recovers the pattern $\emptyset \rightarrow \{\textit{support}, \textit{vector}, \textit{machin}, \textit{svm}\}$.

Thanks to careful implementation and accurate gain estimates, GRAB scales very well in the number of transactions, as well as in the number of features. In practice, GRAB can consider up to several thousand features in reasonable time. Ultimately, we are interested in bioinformatics applications, and are hence interested in rule set search strategies that scale up to millions of features or more. For similar reasons we are interested in extending GRAB towards continuous-valued, and mixed-type data. This we also leave for future work.

Whereas the rules GRAB discovers provide useful insight, they are not necessarily actionable; that is only the case when X causes Y . Currently GRAB can only reward correlation, and we are interested in extending it towards additionally identifying causal rules from observational data [20].

8 Conclusion

We considered the problem of non-parametrically discovering sets of association rules for a given dataset. We proposed to mine small, non-redundant sets of highly informative noise-resistant rules and patterns, that together succinctly describe the data at hand. To do so, we defined a score based on solid information theoretic grounds, showed the problem does not lend itself for efficient

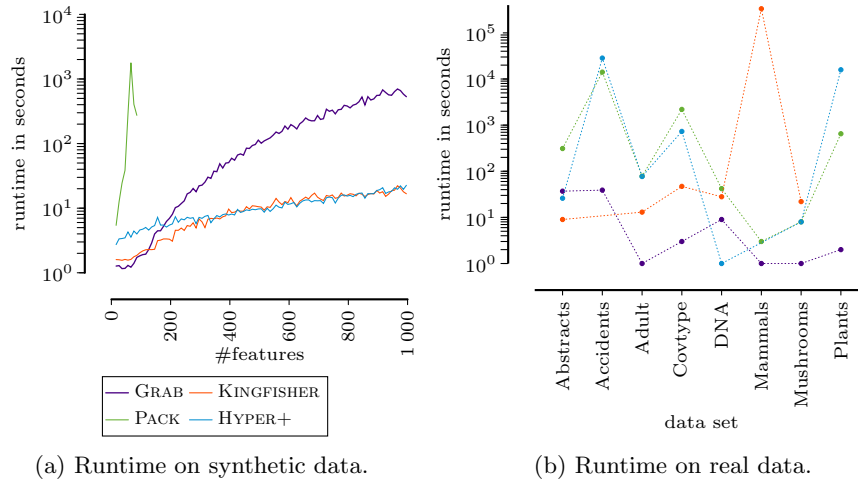


Fig. 5: *Scalability* On the left side, runtimes are visualized on a logarithmic y-axis for synthetic data of varying number of features (x-axis). On the right, runtimes (logarithmic y-axis) are depicted for 8 real world data sets (x-axis). KINGFISHER did not finish on *Accident* and *Plants*, HYPER+ did not finish on *Mammals*.

optimization, and proposed GRAB, a highly efficient heuristic that greedily approximates the MDL optimal result. GRAB is unique in that it can discover both patterns and rules, is noise-resistant and allows rules and patterns to hold approximately, and, can discover rules with non-singleton consequents. Through thorough experiments we showed that unlike the state-of-the-art, GRAB is able to recover the ground truth in synthetic data, and discovers small sets of highly meaningful rules from real world data.

References

1. R. Agrawal and R. Srikant. Fast algorithms for mining association rules. In *VLDB*, pages 487–499, 1994.
2. R. Bayardo. Efficiently mining long patterns from databases. In *SIGMOD*, pages 85–93, 1998.
3. T. Calders and B. Goethals. Non-derivable itemset mining. *Data Min. Knowl. Disc.*, 14(1):171–206, 2007.
4. T. De Bie. Maximum entropy models and subjective interestingness: an application to tiles in binary databases. *Data Min. Knowl. Disc.*, 23(3):407–446, 2011.
5. J. Fowkes and C. Sutton. A subsequence interleaving model for sequential pattern mining. In *KDD*, 2016.
6. P. Grünwald. *The Minimum Description Length Principle*. MIT Press, 2007.
7. W. Hämaläinen. Kingfisher: an efficient algorithm for searching for both positive and negative dependency rules with statistical significance measures. *Knowl. Inf. Sys.*, 32(2):383–414, 2012.

8. J. Han, J. Pei, and Y. Yin. Mining frequent patterns without candidate generation. In *SIGMOD*, pages 1–12. ACM, 2000.
9. S. Jaroszewicz and D. A. Simovici. Interestingness of frequent itemsets using bayesian networks as background knowledge. In *KDD*, pages 178–186. ACM, 2004.
10. P. Kontkanen and P. Myllymäki. MDL histogram density estimation. In *AISTATS*, 2007.
11. M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and its Applications*. Springer, 1993.
12. C. Lucchese, S. Orlando, and R. Perego. Mining top-k patterns from binary datasets in presence of noise. In *SDM*, pages 165–176, 2010.
13. M. Mampaey, J. Vreeken, and N. Tatti. Summarizing data succinctly with the most informative itemsets. *ACM TKDD*, 6:1–44, 2012.
14. H. Mannila, H. Toivonen, and A. I. Verkamo. Efficient algorithms for discovering association rules. In *KDD*, pages 181–192, 1994.
15. P. Miettinen and J. Vreeken. MDL4BMF: Minimum description length for Boolean matrix factorization. *ACM TKDD*, 8(4):A18:1–31, 2014.
16. T. Mitchell-Jones. *Societas europaea mammalogica*. <http://www.european-mammals.org>, 1999.
17. F. Moerchen, M. Thies, and A. Ultsch. Efficient mining of all margin-closed itemsets with applications in temporal knowledge discovery and classification by compression. *Knowl. Inf. Sys.*, 29(1):55–80, 2011.
18. S. Myllykangas, J. Himberg, T. Böhling, B. Nagy, J. Hollmén, and S. Knuutila. DNA copy number amplification profiling of human neoplasms. *Oncogene*, 25(55):7324–7332, 2006.
19. L. Papaxanthos, F. Llinares-López, D. A. Bodenham, and K. M. Borgwardt. Finding significant combinations of features in the presence of categorical covariates. In *NIPS*, pages 2271–2279, 2016.
20. J. Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2nd edition, 2009.
21. L. Pellegrina and F. Vandin. Efficient mining of the most significant patterns with permutation testing. In *KDD*, pages 2070–2079, 2018.
22. J. Rissanen. Modeling by shortest data description. *Automatica*, 14(1):465–471, 1978.
23. J. Rissanen. A universal prior for integers and estimation by minimum description length. *Annals Stat.*, 11(2):416–431, 1983.
24. N. Tatti. Maximum entropy based significance of itemsets. *Knowl. Inf. Sys.*, 17(1):57–77, 2008.
25. N. Tatti and J. Vreeken. Finding good itemsets by packing data. In *ICDM*, pages 588–597, 2008.
26. J. Vreeken and N. Tatti. *Interesting Patterns*, pages 105–134. Springer, 2014.
27. J. Vreeken, M. van Leeuwen, and A. Siebes. KRIMP: Mining itemsets that compress. *Data Min. Knowl. Disc.*, 23(1):169–214, 2011.
28. F. Wang and C. Rudin. Falling rule lists. In *AISTATS*, 2015.
29. G. I. Webb. Discovering significant patterns. *Mach. Learn.*, 68(1):1–33, 2007.
30. Y. Xiang, R. Jin, D. Fuhry, and F. F. Dragan. Succinct summarization of transactional databases: an overlapped hyperrectangle scheme. In *KDD*, pages 758–766, 2008.
31. M. J. Zaki, S. Parthasarathy, M. Ogihara, and W. Li. New algorithms for fast discovery of association rules. In *KDD*, Aug 1997.
32. A. Zimmermann and S. Nijssen. Supervised pattern mining and applications to classification. In *Frequent Pattern Mining*, pages 425–442. Springer, 2014.