

# Explainable Data Decompositions

Sebastian Dalleiger and Jilles Vreeken

CISPA Helmholtz Center for Information Security  
{sebastian.dalleiger, jv}@cispa.de

## Abstract

Our goal is to discover the components of a dataset, characterize *why* we deem these components, explain *how* these components are different from each other, as well as identify what properties they *share* among each other. As is usual, we consider regions in the data to be components if they show significantly different distributions. What is not usual, however, is that we parameterize these distributions with patterns that are informative for one or more components. We do so because these patterns allow us to characterize what is going on in our data as well as explain our decomposition.

We define the problem in terms of a regularized maximum likelihood, in which we use the Maximum Entropy principle to model each data component with a set of patterns. As the search space is large and unstructured, we propose the deterministic DISC algorithm to efficiently discover high-quality decompositions via an alternating optimization approach. Empirical evaluation on synthetic and real-world data shows that DISC efficiently discovers meaningful components and accurately characterises these in easily understandable terms.

## Introduction

Suppose we are analysing the sales data from a supermarket. Likely, our clientele consists of different groups of customers, each of which have their own buying behaviour. Students, for example, often buy *pasta* and *ketchup*, post-docs buy ready-made *sauce* with their *pasta*, whereas professors can afford to make the sauce themselves out of fresh ingredients. That is, the data consists of different *components*, parts of the data that show significantly different pattern distributions.

Certain patterns may be characteristic for more than just one component; both students and post-docs often consume *pizza* and *beer*, after all. That is, the set of patterns that characterise the data can also be partitioned: each such *pattern component* consists of those patterns that are characteristic for a distinct set of *data components*. Together, the pattern components give detailed yet easily interpretable insight in why there are components, how they are different from each other, and what properties are shared among them.

We propose to discover the *pattern composition* of a given database. That is, our goal is to jointly discover the components of the data as well as those pattern components that

optimally characterise their similarities and differences—and do so both efficiently and in a statistically well-founded manner where we only have to set a significance threshold  $\alpha$ . To this end, we model a component given a set of patterns using the Maximum Entropy principle (Jaynes 1982). That is, we use a maximum likelihood estimator that satisfies the empirically observed frequencies of the given patterns, but otherwise makes no further assumptions. We can then formulate the problem in terms of a likelihood maximization problem, where we after that composition that achieves the highest overall likelihood. To avoid overfitting we rely on the BIC model selection criterion. In other words, we are after the most succinct way to summarise the data, by partitioning, such that parts exhibit significantly different distributions, and describing these distributions non-redundantly using only small and interpretable pattern sets.

Clearly, the search space for this problem is enormous: for a given dataset there exist an exponential number of patterns, an exponential number of pattern sets, and an exponential number of partitions. Moreover, this space is not structured, barring efficient search for the optimum. We therefore introduce DISC, a deterministic method that heuristically discovers a good pattern composition. The main idea is that we split the problem into two parts, and iterate between them until convergence. That is, for a given data decomposition we propose to approximate the pattern components using DESC, and then given those pattern components to DISC to discover refined data components. In both steps we rely on statistical tests to both prune the search space, as well as make sure we do not discover any spurious patterns or components.

Extensive experiments show that DESC and DISC work well in practice. DESC outperforms the state of the art in pattern set mining, highly efficiently discovering succinct models, while DISC recovers meaningful components. Case studies confirm that the components and their characterisations make sense: the ecological niches and commonalities that DISC discovers correspond to the ground truth.

In sum, the contributions of this paper are as follows, **(I)** defining the *pattern composition* problem, **(II)** a fast method for discovering pattern components, **(III)** a fast method for discovering pattern compositions, and **(IV)** extensive validation on synthetic and real data. We provide the proofs, additional experiments, and details for reproducibility in the supplementary material.

## Preliminaries

In this work, we consider binary transaction data. Let  $\mathcal{I}$  be a set of items, e.g. products for sale in a supermarket. A transaction  $t \in \Omega$  is a set of items, e.g. the products a customer bought, where  $\Omega = 2^{\mathcal{I}}$  denotes the space of all possible transactions over  $\mathcal{I}$ . A dataset  $D$  over  $\mathcal{I}$  is then a bag of  $n$  transactions, e.g. the sales transactions on a given day. We write  $\pi(D) \in \Pi(D)$  for a partitioning of data  $D$  into  $k$  non-empty subsets,  $\pi(D) = \{D_1, \dots, D_k\}$  s.t.  $\bigcup_{D_j \in \pi(D)} D_j = D$ . Whenever  $D$  is clear from context, we simply write  $\pi$ .

As patterns we are interested in itemsets  $x \subseteq \mathcal{I}$ . A transaction  $t \in D$  supports an itemset  $x \subseteq \mathcal{I}$  iff  $x \subseteq t$ . The empirical frequency of an itemset  $x$  in  $D$  is denoted by  $q(x | D) = \frac{1}{n} |\{t \in D \mid x \subseteq t\}|$ , using the shorthands  $q(x)$  and  $q_j(x)$  to denote the empirical frequency of  $x$  in  $D$  resp.  $D_j$ . A *pattern set*  $S$  is simply a set of itemsets. Combined with empirical frequencies, a set of patterns  $S$  defines a probability distribution  $p$  over  $\Omega$ . For any  $n \in \mathbb{N}$  we write  $[n] = \{1, 2, \dots, n\}$ . All logarithms are to base 2, and by convention we use  $0 \log 0 = 0$ .

## The Pattern Composition Problem

In this section we define our problem. Starting informally, our goal is to characterize the composition of a given binary dataset  $D$ . That is, we aim to decompose the dataset  $D$  into disjoint components  $D_1, \dots, D_k$ , such that every component  $D_j$  has a significantly different pattern distribution  $p_j$ , while at the same time characterizing each of these distributions in terms of informative patterns, such that we can identify the similarities and differences between the components.

In other words, we our goal is to discover a partitioning  $\pi$  of  $D$ , a succinct, non-redundant set of patterns  $S$ , and an assignment  $A$  that associates patterns to components, such that we maximize the regularized likelihood

$$\ell(\pi, S, A) = - \sum_{D_j \in \pi(D)} \log p(D_j | S_j) + r(\pi, S, A) .$$

where  $S_j \subseteq S$  is the subset of  $S$  that  $A$  indicates to be relevant for  $D_j$ . Last,  $r(\pi, S, A)$  is a regularization term that steers the problem away from trivial solutions, such as decomposing the data into singleton components, or to including every possible pattern.

We discuss and define each of these terms in detail below, after which we give a formal definition of the problem.

## The Probability Distribution

We start by defining the probability distribution  $p$  over  $\Omega$ . It is easy to see that maximum likelihood estimator (MLE) of  $p$  depends on the patterns in  $S$  and their empirical frequencies in  $D$ . That is,  $p$  will only achieve maximum likelihood if for any pattern  $x \in S$  the estimated frequency  $p(x | S) = E_f[x | S]$  corresponds to the observed frequency  $q(x | D)$ . We define the set of feasible distributions as the polytope

$$\mathbb{P}_S^D \equiv \{f \in \Omega \rightarrow [0, 1] \mid E_f[x | S] = q(x | D) \forall x \in S\} ,$$

that contains all, infinitely many distributions consistent with the empirical observations  $q$  of patterns  $x \in S$  in  $D$ .

Not all distributions in  $\mathbb{P}_S^D$  suit our needs: we need a distribution that does not introduce additional assumptions beyond the information that  $S$  specifies. From an information theoretic point of view, additional assumptions correspond to additional information. We can measure the amount of information in a distribution using Shannon entropy,  $H(p) = - \sum_x p(x) \log p(x)$ . The lower the information content of a distribution  $p$ , the higher its entropy. We can uniquely identify the feasible distribution that makes the least additional assumptions as the one with the highest entropy

$$f \equiv \arg \max_{f \in \mathbb{P}} H(f) ,$$

which is known as the Principle of Maximum Entropy (Jaynes 1982). In general this does not immediately provide a family of distributions to use. In our case, however, as the constraints of  $\mathbb{P}_S^D$  are linear, we know that distribution  $p$  over transactions  $t \in \Omega$  takes an exponential form

$$f(t | S) = \theta_0 \prod_{x_i \in S} \theta_i^{\mathbf{1}[x_i \subseteq t]} ,$$

that for appropriately chosen coefficients  $\theta \in \mathbb{R}^{|S|+1}$  factorizes into marginals (Csiszár 1975). Conveniently, optimizing  $\theta$  is a convex problem, and hence we can employ standard convex optimizers such as iterative scaling (Darroch and Ratcliff 1972). We are specifically interested in inferring the expected frequency  $p(x | S)$  for arbitrary itemsets  $x \in \Omega$ —the frequencies of  $x \in S$  are given, after all. To infer  $p$  we have to sum the probabilities of *every possible transaction*  $t \in \Omega$  that supports  $x$ ,

$$p(x | S) = E_f[x | S] = \sum_{t \in \Omega} f(t | S) \mathbf{1}[x \subseteq t] ,$$

which is known to be **PP**-hard (Tatti 2006).

Usually, the assumption is that the full dataset is generated from *i.i.d.* samples, however, a component  $D_j \in \pi(D)$  is defined to be a subset of  $D$  with a unique distribution. Within a component, samples are independently drawn. The likelihood of  $D$  is the product of component likelihoods, given sets  $S_j$ , where the likelihood of  $D_j$  is

$$p(D_j | S_j) = \prod_{t \in D_j} p(t | S_j) .$$

Wherever clear from context, we will write  $p_S$  for  $p(\cdot | S)$ , and  $p_j$  for  $p(\cdot | S_j)$ .

## Informative Patterns

Our goal is to discover that maximally succinct, maximally non-redundant pattern set  $S$  that maximizes the likelihood of the data,  $p_S(D)$ . Likelihood  $\ell$  is a monotonic function: if we insert any pattern  $x \notin S$  into  $S$  it almost always increases, and at worst stays the same.

We say that a pattern  $x$  is *informative* for  $D$  with respect to  $S$  if we see a *significant* increase in likelihood if we include  $x$  in  $S$ . To determine whether a pattern or component is informative, we use a model selection criterion. We use BIC, as it is simple, efficiently computable, and as we will see,

works well in practice (Schwarz 1978). For a single component, we have  $|S| = m$  degrees-of-freedom (dof) and hence  $r(S) = m/2 \log |D|$ .

It is straightforward to generalize the above to multiple components. Given a partitioning  $\pi$  and a pattern set  $S$ , we need to determine which patterns  $x \in S$  are informative for which component  $D_i \in \pi$ . This is what our assignment matrix  $A \in \{0, 1\}^{m \times k}$  is for. It is a binary matrix over components and patterns where  $A_{ij} = 1$  if pattern  $x_i \in S$  is informative for component  $D_j$ . The set of patterns that are informative for a component  $D_j$  is defined as  $S_j = \{x_i \in S \mid A_{ij} = 1\}$ . If a pattern is informative for multiple components, we call it common or shared among those components.

Taking assignment matrix  $A$  into account, the BIC regularization cost  $r$  is

$$r(\pi, S, A) = \frac{1}{2} [k \cdot (2m + d) + n] \log n ,$$

where **(i)** the assignment matrix for patterns is accounted for by having  $m \times k$  dof. Next, **(ii)**, the number of coefficients,  $\theta$  used by  $k$  distributions amounts to  $(m + d) \times k$  additional dof, where  $d = \dim D$  correspond to the additional dof used for singletons from  $\mathcal{I}$ . Thirdly, **(iii)** we have additional  $|D|$  dof that can encode any partitioning, by assigning a label to each row. However, the latter is constant for any partitioning.

## The Problem, formally

Combining the above, we can now formally state the problem.

**Problem 1 (The Pattern Composition Problem)** *Given a transactional dataset  $D$  over items  $\mathcal{I}$ , our goal is to jointly discover that **(i)** partitioning  $\pi \in \Pi(D)$  of  $D$  into fewest parts, that **(ii)** smallest pattern set  $S \subseteq \mathcal{T}_{\mathcal{I}}$ , and that **(iii)** assignment matrix  $A \in \{0, 1\}^{|S| \times |\pi|}$  such that*

$$\ell(\pi, S, A) = - \sum_{D_j \in \pi(D)} \log p(D_j \mid S_j) + r(\pi, S, A)$$

is minimal.

Unsurprisingly, this is a difficult problem with a very large search space. First of all, there exist a Bell-number  $B_{|D|}$  of possible partitionings  $\pi$ . Secondly, the number of possible pattern sets is doubly exponential in the number of unique items in  $D$ , as  $S \in 2^{2^{\mathcal{I}}}$ . Finally, the objective function does not exhibit any structure that we can exploit for efficient search, i.e. it is neither (anti-)monotone, nor submodular.

## Algorithms

To efficiently discover good solutions in practice, we separate the problem into two parts and take an alternating optimization approach. That is, starting from partitioning  $\pi_0$  where all of  $D$  is in one part, we iterate between the following two steps until convergence. First, given a partitioning  $\pi(D)$  we efficiently discover a high quality pattern set  $S$  and assignment matrix  $A$ . Second, given a pattern set  $S$ , assignment matrix  $A$ , and partitioning  $\pi$  we discover a refined partitioning  $\pi'$  such that we minimize our objective function.

Below we discuss these two steps in turn.

## Discovering Patterns given a Partitioning

The first problem we consider is that of discovering a high-quality set of informative patterns  $S \in 2^{\Omega}$  and assignment matrix  $A$  for a given partitioning  $\pi \in \Pi(D)$ . Like our overarching problem, this problem is also too hard to solve exactly; there exist doubly exponentially many pattern sets, and our hard-to-compute score does not show any structure we can exploit. We are hence going to approximate the optimal result through an iterative greedy approach.

For a single dataset, Mampaey, Vreeken, and Tatti (2012) showed that finding the set  $S$  with minimal  $\ell$  is equivalent to finding the set that has the highest gain in likelihood in comparison to the empty pattern set  $S_{\emptyset}$ . That is, minimizing

$$\ell_j(S) = - \sum_{t \in D_j} \log p_j(t) ,$$

is equivalent to maximizing the KL-divergence

$$\arg \max_S KL[p_S \parallel p_{\emptyset}] \quad (1)$$

and the greedy solution to Eq. (1) is equivalent to iteratively minimizing  $\ell_j(S)$  directly. In fact, we can guarantee the quality of the greedy solution.

**Lemma 1** *Eq. (1) is a Submodular Function Maximization problem. The greedy solution  $S$  is in the  $e^{-m/m^*}$ -radius of the optimal solution  $S^*$ , where  $m = |S|$  and  $m^* = |S^*|$*

For conciseness, we provide the proof in the supplementary.

The greedy algorithm to minimize Eq. (1) is still prohibitively slow in practice: it repeatedly evaluates KL-divergences  $KL[p_S \parallel p_{S \cup \{x\}}]$  to measure the information gain of adding itemset  $x$  to  $S$ , and this computation relies on computationally costly inference of frequencies using  $p$ . Nonetheless, this formulation does allow us to derive a computationally efficient admissible heuristic.

To reduce the complexity of computing  $KL$ , we want to reduce the number of queries it makes to  $p$ . In its full computation, it both considers the frequencies of pattern  $x$  itself, as well as that of its exponentially many subsets  $y \subset x$ . Ignoring these subsets permits the following lower-bound

$$h(x \mid S) = q(x) \log q(x) / p_S(x) .$$

Due to the decomposition of the data, this admissible heuristic easily generalizes to an admissible normalized *lower-bound information gain* over multiple components  $h(x \mid S)$  is  $E_{S_j} [h(x \mid S_j)]$ . In general, our pattern set discovery strategy is hence as follows. In the current iteration  $i$ , the last pattern set  $S^{i-1}$  is known and fixed. We use  $h$  to select that  $x$  in candidate set  $F \subseteq \Omega$  that has the highest marginal gain. That is, until convergence of  $\ell$ , we iterate

$$S^i \leftarrow S^{i-1} \cup \arg \max_{x \in F} h(x \mid S^{i-1}) .$$

This leaves us to specify the candidate set  $F$ . Naively, we could set  $F = \Omega$ . Clearly, this is not practical: besides that  $\Omega$  is typically prohibitively large, it contains exponentially many candidates that will be uninformative with regard to patterns in  $S$ . We hence propose a more effective search strategy, in which we take into account what  $S$  can already

---

**Algorithm 1: DESC for Describing the Composition**

---

**Input:** Data  $D$ , partitioning  $\pi$   
**Output:** Distributions  $p_j$ , pattern set  $S$ , assignment  $A$

- 1  $S \leftarrow \{x \in \mathcal{I}\}$
- 2  $A \leftarrow$  initialize with all 1s
- 3  $p \leftarrow$  infer  $p(\cdot | S_j)$  for each component  $D_j$
- 4  $F \leftarrow \{z = x \cup y \mid x, y \in S, r(z) < h(z)\}$
- 5 **while**  $F \neq \emptyset$  **do**
- 6      $z \leftarrow \arg \max_{x \in F} h(x | S)$
- 7      $A' \leftarrow$  according Eq. (2) wrt  $z$
- 8      $S' \leftarrow S \cup \{z\}$  if  $z$  assigned to a component
- 9      $p' \leftarrow$  infer  $p(\cdot | S'_j)$  for each component  $D_j$
- 10    **if**  $\ell(\pi, S', A') < \ell_j(\pi, S, A)$  **then**
- 11        $A \leftarrow A'; S \leftarrow S'; p \leftarrow p'$
- 12        $F \leftarrow \{z = x \cup y \mid x, y \in S, r(z) < h(z | S)\}$
- 13    **else**
- 14        $F \leftarrow F \setminus z$
- 15 **return**  $(p, S, A)$

---

explain well. In a nutshell, we iteratively generate candidates by merging pairs of patterns  $x, y \in S \cup \mathcal{I}$  into a candidate  $x \cup y \in F$ . However, we only want to consider the subset of candidates that will surely reduce our objective  $\ell$ . Those are candidates  $z \in F$  for which  $h(z | S) > r(z)$ , where  $r(z) = r(\pi, S \cup \{x\}) - r(\pi, S)$ . Similarly, we assign a candidate  $i$  to a component  $j$  if it yields a gain in  $\ell_j$ , i.e.

$$A_{ij} = 1 \iff h(x_i | S_j) < r_j(x_i), \quad (2)$$

where the cost  $r_j(z)$  is  $r(\pi, S, A') - r(\pi, S, A)$ . Here  $A'$  is equivalent to  $A$ , but with  $A_{ij} = 1$ .

Putting the above together, we have algorithm DESC, for which we give the pseudo-code as Algorithm 1. In short, starting with the singleton only model (line 1–3) we generate our initial batch of candidates  $F$  (ln. 4). We consider these candidates descending on  $h$  (ln. 6) and evaluate each  $z \in F$  (ln. 7–9). If the objective improves, we keep the candidate (11–12), and otherwise reject it (14).

The computational complexity of DESC depends on the number of candidates in  $F$ , which is quadratic in the number of patterns in  $S$  and in the worst case can grow up to  $|\Omega|$ . In practice, however, the properties of the maximum entropy distribution together with the BIC regularizer keep the size of  $S$  small, in the order of tens to hundreds of patterns, say  $S_{\max}$ . The worst case complexity of DESC is dominated by the inference of the distributions  $p_j$ , and hence in PP. The average complexity,  $\gamma$ , of  $p$  is much lower (Mampaey, Vreeken, and Tatti 2012), however, and hence the average complexity of DESC is  $\mathcal{O}(\gamma \cdot |S_{\max}|^2)$ .

## Discovering the Composition

Next, we consider the orthogonal problem of discovering a high-quality partitioning  $\pi(D)$  given a pattern set  $S$  and assignment matrix  $A$ . As there is no effective exact search for the optimal partitioning, we again rely on heuristics. In particular, we take a top-down approach where we iteratively refine the current partitioning  $\pi$  using the patterns in  $S$ .

Our strategy is based on the idea that a significantly different distributions of patterns are an indicator for the presence of latent factors of unknown components. In other words, we say that a component was generated using a latent data source that left a distinctly distributed trail of patterns in the data. By narrowing down a given component to a subset with a distribution that stands out from the rest of the data, we can refine the current partitioning to identify these latent parts as separate components. We can also use this observation also in reverse: when we narrow down a component and find that the pattern distribution we so obtain is not significantly different from the remainder or the other components, we do not want this candidate component to be part of our solution.

We write  $p_j^x$  for the pattern distribution we infer on that part of component  $D_j$  where pattern  $x$  occurs. Likewise, we consider  $p_j^{x'}$  over that part of  $D_j$  where  $x$  does not occur. We measure the divergence between two distributions with same support using the Jensen-Shannon divergence  $JS(P, Q) = KL[P \| M] + KL[Q \| M]$ , where  $M = (P + Q)/2$ . The scaled  $JS(p_j^x, p_j^{x'})$  statistics is asymptotically  $\chi^2$  distributed with  $|S| - 1$  dof (Menéndez et al. 1997). From this we get a  $p$ -value for a single test. However, as we test many hypotheses, i.e. candidate refinements, we hence correct for the familywise error-rate (FWER) by adjusting the significance level  $\alpha$  using Bonferroni correction (Bonferroni 1936).

For a given  $S$  and for any partitioning  $\pi \in \Pi(D)$ , we write  $A(\pi, S)$  as the assignment matrix that characterizes the partitioning  $\pi$  with  $S$  by minimizes our objective function with respect to Eq. (2). Formally, for a given  $S$ , the problem of discovering components is as follows

$$\arg \min_{\pi \in \Pi(D)} \ell(\pi, S, A(\pi, S))$$

subject to  $p_i, p_j$  significantly JS-divergent  $\forall i \neq j$

This is, again, a hard problem, and again, the search space is large and unstructured. We therefore employ a greedy top-down approach. Starting with a single component  $D_j \in \pi$ , we decompose  $D_j$  into two sub-components  $D_j^1$  and  $D_j^2$  such that these are significantly differently distributed from each other, as well as from the rest of other components in  $\pi$ . Following the notion that latent factors are identifiable by distinct pattern distributions, we start the refinement process of a given component  $D_j \in \pi$  with a pattern  $x \in S \times \mathcal{I}$  by separating a component into two children  $D_j^x \equiv \{t \in D_j \mid x \subseteq t\}$  and  $D_j^{x'} \equiv D_j \setminus D_j^x$ . The corresponding refinement of  $\pi$  is written as

$$\text{refine}_{\pi}(x, j) \equiv \{\pi', A(\pi')\},$$

where the new partitioning  $\pi'$  is  $\pi \setminus \{D_j\} \cup \{D_j^x, D_j^{x'}\}$ .

As real data is noisy and distributions are complex, it is unlikely that an individual pattern  $x$  perfectly identifies a latent component. That is, after splitting a component it may be that the overall assignment of transactions to components may be suboptimal with regard to likelihood. Just like in the EM algorithm we therefore iteratively reassign transactions that components where they achieve the highest likelihood. That is, in each iteration we ensure for every  $t \in D$  that

$$t \in D_i \iff \hat{i} = \arg \max_{j \in [k]} p(t | S_j), \quad (3)$$

---

**Algorithm 2: DISC** for Discovering the Composition

---

**Input:** Data  $D$ , significance threshold  $\alpha$ **Output:** Partitioning  $\pi$ , pattern set  $S$ , assignment  $A$ 

```

1  $\pi \leftarrow \{D\}$ 
2  $S, A \leftarrow \text{DESC}(D, \pi)$ 
3  $G \leftarrow$  according Eq. (4)
4 while  $G \neq \emptyset$  and  $\ell$  has not converged do
5    $S, A \leftarrow \text{DESC}(D, \pi)$ 
6    $c \leftarrow \arg \min_{c \in G} \ell(\text{refine}_\pi(c), S)$  cf. Eq. (5)
7    $\pi', A' \leftarrow \text{refine}_\pi(c)$ 
8   while  $\ell$  has not converged do
9     let  $\pi'$  satisfy Eq. (3)
10     $A' \leftarrow A(\pi')$  according to Eq. (2)
11     $p \leftarrow \text{infer } p(\cdot | S_j)$  for each component  $D_j$ 
12    if  $\ell(\pi', S, A') < \ell(\pi, S, A)$  then
13       $\pi \leftarrow \pi', A \leftarrow A'$ 
14       $G \leftarrow$  according Eq. (4)
15    else
16       $G \leftarrow G \setminus c$ 
17 return  $(\pi, S, A)$ 

```

---

re-estimate the distribution  $p$ , re-compute  $A(\pi, S)$  and repeat until convergence. Starting with  $\pi = \{D\}$ , we iteratively refine the current partitioning by selecting the JS-significance refinement of  $\pi$  with highest marginal gain, until convergence of  $\ell$ . Formally, out of the set  $G$  of candidates

$$\{(j, x) \in [k] \times S_j \cup \mathcal{I} \mid \text{refine}_\pi(x, j) \text{ significant}\}, \quad (4)$$

we select the refinement candidate that reduces  $\ell$  most

$$\arg \min_{x, j \in G} \ell(\text{refine}_\pi(x, j), S). \quad (5)$$

Putting all the above together, we have the DISC algorithm. We give the pseudo-code as Algorithm 2. In a nutshell, starting from the trivial partitioning  $\pi(D) = \{D\}$ , we iteratively use DESC to discover the pattern components, use these patterns to find the best refinement of  $\pi$ , reassign the rows to optimize the likelihood, and only accept this refinement if it is significant. We repeat this until convergence. Disregarding the complexity of querying our distribution, DISC scales linearly with the size of  $G$ . In the worst case this means  $|\Omega| \times |D|$ . However, since in practice both  $S$  and  $\pi$  tend to be small, DISC is feasible on real world data.

## Related Work

The vast majority of literature has been devoted on either finding clusters of transactions or finding patterns that characterise a dataset. Surprisingly, there exist no technique to discover the *pattern composition* of the data.

Our problem obviously relates to mixture modelling (Dempster, Laird, and Rubin 1977) where data is modelled as a mixture of several probability distributions. Mixture modelling, however, requires us to assume a probability distribution, whereas the true distribution is unknown. Similarly, clustering (MacQueen 1967) is related, as it groups data points, but relies on an assumed distance measure. Additionally in contrast to our approach, many of these approaches

are stochastic, require us to choose the number of components up front, and none characterise the commonalities and differences between the components in interpretable terms.

In this sense, co-clustering, also known as bi-clustering, is more closely aligned to our goal. In co-clustering we simultaneously cluster rows and columns and can interpret the column-clustering as an implicit characterisation of the row-cluster. Moreover, there exist parameter-free methods like Information Co-clustering (Dhillon, Mallela, and Modha 2003), Cross-associations (Chakrabarti et al. 2004). These techniques, only discover non-overlapping rectangles in the data that are exceptionally dense or sparse, rather than data components with significantly different pattern distributions.

There also exist methods that can provide post-hoc explanations, for example using a consistent set of decision rules, which lead to a prediction (Lakkaraju, Bach, and Leskovec 2016) or a clustering (Kim, Shah, and Doshi-Velez 2015; Chen et al. 2016). These rules together characterise the decision boundary for a cluster, whereas we are interested those patterns that characterise the similarities and differences between components. In other words, rather than explaining the clustering after the fact, our models directly explain *why* there is a clustering.

Pattern mining methods are obviously strongly related to DISC. Frequent pattern mining (Agrawal and Srikant 1994; Nijssen, Guns, and De Raedt 2009) is well-known to discover far too many patterns for the result to be interpretable. OPUS (Webb and Vreeken 2014; Webb 2010) curbs the pattern explosion through the use of statistical tests. SLIM (Smets and Vreeken 2012), IIM (Fowkes and Sutton 2016) and MTV (Mampaey, Vreeken, and Tatti 2012) are examples of techniques that discover concise and non-redundant pattern sets. These methods all only provide a single pattern set for a single database. DIFFNORM (Budhathoki and Vreeken 2015) is the only method we know that for a given data partitioning can characterise differences and similarities between the pattern distributions.

## Experiments

In the experiments, we evaluate DISC on synthetic, as well as 17 real-world datasets that together span a wide variety of domains, sizes, and dimensionalities. We implemented DISC in C++, ran experiments on a 12-Core Intel Xeon E5-2643 CPU, and report wall-clock time. We provide the source code, datasets, synthetic dataset generator, and additional information needed for reproducibility.<sup>1</sup>

In many of the following experiments we compare likelihood  $\ell$  of the estimated model with the likelihood  $\ell^0$  of the initial model, that is  $S^0 = \mathcal{I}$  for a single component  $\pi(D) = \{D\}$ . We measure the likelihood ratio  $\ell/\ell^0$ , in percent, where a lower values corresponds to a higher regularized likelihood of the data under the model. In all experiments we have used the same significance level  $\alpha = 0.01$ .

**Describing Components** In this section, we study our pattern set miner DESC on real-world datasets. Before we characterize datasets for a given partitioning, we start with the

<sup>1</sup><https://eda.mmci.uni-saarland.de/disc/>

special case of discovering a pattern set for a given composition. In this set-up, we compare against SLIM (Smets and Vreeken 2012), MTV (Mampaey, Vreeken, and Tatti 2012) and DIFFNORM (Budhathoki and Vreeken 2015). For efficiency reasons they only consider frequent patterns, i.e. for which  $q(z) > \kappa$  according to a user defined minimum frequency threshold  $\kappa$ . It is trivial to constraint DESC to consider frequent patters only and to compare fairly, we use the same thresholds for all methods in the following experiments.

DIFFNORM characterizes a pre-partitioning dataset  $\pi(D)$ . To this end we consider 9 labeled datasets, which we partition based on their class labels. Neither SLIM, MTV make use of any partitioning and can only be applied on the complete dataset, while DESC can do both and is applied to partitioned data where available and otherwise to the complete data.

In Fig. 1a we show that SLIM discovers pattern sets that consist of hundreds up to thousands of patterns, the results of MTV and DESC are in the order of tens of patterns. The pattern sets discovered by DISC, DESC and MTV are much more concise than the results from SLIM for a single component or DIFFNORM for a given composition.

We give the wall-clock runtime of the three methods in Fig. 1b. We see that DESC achieves these results in an order of magnitude faster than MTV and surpasses SLIM. On average DESC requires less than a quarter (16%) of the runtime of SLIM, just 13% of DIFFNORM and only 0.24% of the runtime of MTV. As MTV and DESC optimize the same score, we can fairly compare them. In Fig. 1c we show that DESC outperforms MTV in almost all cases.

**Discovering the Composition** Now, we study the full algorithm: simultaneously discover both the pattern sets and the partitioning of the dataset using DISC. First, we test and verify DISC on synthetic data with 128 items in  $\mathcal{I}$ . For this, we generate synthetic datasets such that we have access to the ground-truth. In each trial, we randomly sample a dataset  $D$ , containing 1, 2, 4, 8 components. For each component  $D_j \in \pi$  we randomly generate and insert 5 characteristic patterns into  $S_j$ . For any disjoint pair  $D_i, D_j$ , we generate 3 shared patterns with probability of 20%, that are inserted into both  $S_i$  and  $S_j$ . Every pattern has a randomly chosen frequency associated with it. Each component  $D_j$  consists of 256 rows. In each row, we uniformly at random insert each pattern from  $S_j$  with its corresponding frequency. Lastly, we introduce additive noise, by randomly insert items into each row independently with probability of 5%. For each  $k^*$ , we sample 20 datasets and compare the ground-truth with DISC and DESC. On average, DISC reaches a likelihood  $\ell^{\text{DISC}}$  within 2% of the ground-truth, i.e.  $\ell^{\text{DESC}_{\pi^*}} \pm 2\%$  and always recovers the ground-truth number of components  $k^*$ .

Now after we have verified that DISC works on synthetic examples, we study DISC on the real-world datasets. To do so, we measure how similar data within a component is, using  $\ell$ , and measure how differently distributed the components are, using the *pairwise symmetric KL-divergence* (PSKL)

$$\frac{1}{2 \binom{k}{2}} \sum_{ij \in \binom{[k]}{2}} KL[p_i || p_j] + KL[p_j || p_i],$$

which averages the divergence between pairs of distributions.

Next, we compare DISC to clustering. While options include  $k$ -means, and Expectation-Maximization, these are not a good fit for our case as they are stochastic, require a number of clusters, and neither is the concept of a centroid well defined for discrete spaces, nor is it clear what efficiently queryable distribution to use. DBSCAN (Ester et al. 1996) relies on a distance measure, that we define as

$$d(s, t) = 1 - |s \cap t| / \max(|s|, |t|).$$

Our approach is as follows: First, we cluster the dataset using DBSCAN and get  $\pi \in \Pi(D)$ . Next, we use DESC to describe the clusters  $\pi$  post-hoc by means of  $S, p$  and  $A(\pi, S)$ . Since DBSCAN relies on hyper-parameter, we optimize  $\ell$  using a grid-search over 7  $\epsilon$ -candidates and we do not constraint cluster-sizes. We call this algorithm DBDESC. Similarly, we define DBDESC2 that uses a different distance measure  $d'(s, t) = d(c(s), c(t))$ , where  $c(t)$  contains patterns from  $S$  that are subsets of  $t$ , i.e.  $\{s \in S \mid s \subseteq t\}$ .

We apply DISC, DBDESC and DBDESC2 on all 17 datasets without using any class-labels and summarize results in Fig. 2. First of all, we compare the composition of DISC, DBDESC and DBDESC2 with a class-based composition discovered by DESC on 9 labeled datasets. Almost always, we observed a significantly lower PSKL-divergence between classes than between components, which we show in Fig. 2b. Noticeable is the PSKL-divergence of *Pumsl Star*, as with only 10 components, the divergence is very high compared to DBDESC and DBDESC2.

Additionally, DISCs model usually has a significantly lower objective function  $\ell$  than the class-based result from DESC. This suggest that classes are not always good indicator for components. We see a higher likelihood for decomposed data in comparison to DESC without decomposing the data on class labels, in Fig. 2c. We note the much higher likelihood of the DISCs composition in comparison to DBDESC and DBDESC2. Overall we see that DISC discovers diverging components that have higher in likelihood in comparison to the cluster based composition from DBSCAN.

**Qualitative Study** Now, we study the interpretability and qualitative evaluation of the composition discovered by DISC, by manually inspecting the composition on two datasets. In supplementary material, we discuss the interpretability of the composition of ArXiv-abstracts.

First, we consider the *Mammals* dataset provided by the European Mammal Society. This dataset consists of presence records of 124 European mammals within areas of 50-by-50 kilometres. Additional geographical information were not used during the experiments.

DISC discovers 9 components with 64 patterns in total. We geographically depict the components DISC discovers in Fig. 3. Although it did not know the spatial locations of the data points, it discovers (almost complete) contiguous areas in Europe that correspond to ground truth habitats. Moreover, the patterns it discovers for these components are meaningful: for example, although the combination of species as *Wolverine* and *Norway Lemming* are highly characteristic for both ‘‘Scandinavian’’ components, their distribution differs between these components. For the Iberian Peninsula the




Figure 1: **DESC Efficiently Discovers Concise Pattern Sets.** From left to right, in Fig. 1a the number of discovered patterns (log-scale), in Fig. 1b the runtime (seconds, log-scale), and in Fig. 1c the likelihood ratio  $\ell/\ell_0$  (lower is better). Each figure consists of two groups for labeled (first group) and unlabeled (second group) data.




Figure 2: **DISC Discovers Informative and Interpretable Compositions.** The number of components and the in Fig. 2a the number of discovered components (log-scale), in Fig. 2b PSKL (higher is better), and in Fig. 2c the likelihood ratio  $\ell/\ell_0$  (lower is better) of MTV and DESC. Each figure consists of two groups for labeled (first group) and unlabeled (second group) data.

*Common Genet* and *Mediterranean Pine Vole* are discovered to be very characteristic. The habitation zone of the latter spreads to southern France, and this is reflected by this pattern being shared between these two components. Further, DISC discovers that the co-occurrence of *Eurasian beaver*, *Red squirrel* are descriptive across Europe. Last, but not least, DISC finds that the *Eurasian Harvest Mouse*, *European Mole*, *Eurasian Water and Pygmy Shrew*, *Stoat*, *Field Vole* are all very common across Europe, and include them in a single pattern shared amongst most components stretching Europe.

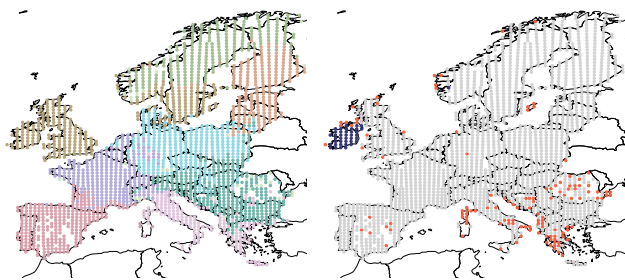


Figure 3: **Discovers Meaningful Partitions.** Results of DISC (**left**) on the *Mammals* dataset. The 9 components represent contiguous areas that correspond to known habitats. DBSCAN (**right**) essentially only discovers Ireland

## Discussion

DISC discovers meaningful, easily interpretable composition. The components are described concisely, by characteristic and shared patterns. DISC explains *why* there are components, what makes them special and what is common among them. Our manual inspection show that the results are easy to interpret and informative. In addition, DESC is a highly efficient pattern miner for single and multiple datasets and beats state-of-the-art in descriptiveness, conciseness and runtime. Furthermore we showed how DESC can be used to describe the result of a clustering algorithm. Furthermore, we observe that jointly optimizing for interpretability and likelihood is doable in practice and can outperform clustering algorithms like DBSCAN with a post-hoc explanation.

We studied the novel problem of discovering the composition, that is a partitioning of the dataset and its description using locally characteristic patterns, or patterns shared across sets of components. We formalized this problem in terms of the family of maximum entropy distributions over itemsets and defined the best composition as the one that gives the most succinct description of the data. We introduced an efficient pattern miner DESC for succinctly describing a single or multiple data components. Using this, we described an algorithm to discover the partitioning of a dataset. Both together discover the composition. Experimental evaluation showed that DISC efficiently discovers interesting and meaningful composition of the data.

## References

- Agrawal, R., and Srikant, R. 1994. Fast algorithms for mining association rules. In *VLDB*, 487–499.
- Bonferroni, C. E. 1936. Teoria statistica delle classi e calcolo delle probabilita. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze* 8:3–62.
- Budhathoki, K., and Vreeken, J. 2015. The difference and the norm – characterising similarities and differences between databases. In *ECML PKDD*. Springer.
- Chakrabarti, D.; Papadimitriou, S.; Modha, D. S.; and Faloutsos, C. 2004. Fully automatic cross-associations. In *KDD*, 79–88.
- Chen, J.; Chang, Y.; Hobbs, B.; Castaldi, P.; Cho, M.; Silverman, E.; and Dy, J. 2016. Interpretable Clustering via Discriminative Rectangle Mixture Model. In *ICDM*, 823–828.
- Csiszár, I. 1975. I-divergence geometry of probability distributions and minimization problems. *Annals Prob.* 3(1):146–158.
- Darroch, J., and Ratcliff, D. 1972. Generalized iterative scaling for log-linear models. *Annals Math. Stat.* 43(5):1470–1480.
- Dempster, A.; Laird, N.; and Rubin, D. 1977. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Statist. Soc. B* 39(1):1–38.
- Dhillon, I. S.; Mallela, S.; and Modha, D. S. 2003. Information-theoretic co-clustering. In *KDD*, 89–98.
- Ester, M.; Kriegel, H.; S, J.; and Xu, X. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. 226–231. AAAI Press.
- Fowkes, J., and Sutton, C. 2016. A Bayesian Network Model for Interesting Itemsets. In *ECML PKDD*, 410–425. Springer.
- Jaynes, E. 1982. On the rationale of maximum-entropy methods. *Proc. IEEE* 70(9):939–952.
- Kim, B.; Shah, J. A.; and Doshi-Velez, F. 2015. Mind the Gap: A Generative Approach to Interpretable Feature Selection and Extraction. In *NIPS*. Curran Associates, Inc. 2260–2268.
- Krause, A., and Golovin, D. 2012. Submodular function maximization. *Tractability: Practical Approaches to Hard Problems* 3.
- Lakkaraju, H.; Bach, S. H.; and Leskovec, J. 2016. Interpretable Decision Sets: A Joint Framework for Description and Prediction. In *KDD*, 1675–1684.
- MacQueen, J. 1967. Some methods for classification and analysis of multivariate observations. In *Berkeley Symp. Math. Stat. Prob.* Vol. I: Statistics, pp. 281–297.
- Mampaey, M.; Vreeken, J.; and Tatti, N. 2012. Summarizing data succinctly with the most informative itemsets. *ACM TKDD* 6:1–44.
- Menéndez, M. L.; Pardo, J. A.; Pardo, L.; and Pardo, M. C. 1997. The Jensen-Shannon divergence. *J Franklin Inst* 334(2):307–318.
- Nijssen, S.; Guns, T.; and De Raedt, L. 2009. Correlated itemset mining in ROC space: a constraint programming approach. In *KDD*, 647–656. Springer.
- Schwarz, G. 1978. Estimating the dimension of a model. *Annals Stat.* 6(2):461–464.
- Smets, K., and Vreeken, J. 2012. SLIM: Directly mining descriptive patterns. In *SDM*, 236–247. SIAM.
- Tatti, N. 2006. Computational complexity of queries based on itemsets. *Inf. Process. Lett.* 98(5):183–187.
- Webb, G., and Vreeken, J. 2014. Efficient discovery of the most interesting associations. *ACM TKDD* 8(3):1–31.
- Webb, G. I. 2010. Self-sufficient itemsets: An approach to screening potentially interesting associations between items. *ACM TKDD* 4(1):1–20.



## Appendix

### Submodularity

We want to minimize  $\ell$ , which we do by via a greedy algorithm that iteratively selects the highest marginal gain candidate  $x \in F$  from the set of candidates  $F$ . In particular, we are interested in bounding the error of the greedy solution with respect to the optimal solution. We start this proof by repeating Theorem 6.1 from Mampaey, Vreeken, and Tatti (2012), that states

$$\arg \min_{x \in \Omega} \ell(\{D\}, S_i \cup \{x\}, \mathbf{1}) = \arg \max_{x \in \Omega} KL(p_{S_i \cup x}, q),$$

where  $q$  is a consistent distribution with the same support than  $p^S$ . In other words, for solution  $S_i$  in iteration  $i$ , the candidates  $x$  with highest marginal gain of  $\ell$  or the candidate  $x'$  for  $KL$ , respectively, are identical. This means that greedy algorithms for both scores are interchangeable. Greedily minimizing  $\ell$  results in the same estimator than greedily maximizing  $KL$ . As a shorthand we write

$$f(S) = KL(p_S; q).$$

and in order to derive the error-bound to  $f$ , we proof monotonicity and submodularity of  $f$  first.

**Proof 1 (monotonicity)** We write the polytope of feasible distributions  $p$  as

$$\mathbb{P}_S \equiv \{p \in \Omega \rightarrow [0, 1] \mid \sum p = 1, p_x = q_x \forall x \in S\}.$$

By consistency of  $q$  we know that  $\mathbb{P}_{S \cup \{x\}} \subseteq \mathbb{P}_S$ . By tightening the constraints around  $p$  for any  $x \in 2^T$  we are reducing the radius between any  $p \in \mathbb{P}_S$  and  $q$  and hence  $f(S) \geq f(S \cup \{x\})$ .

Note, that  $f$  is not necessarily strictly monotonic, since  $x$  might not carry additional information at all and hence  $\mathbb{P}_{S \cup \{x\}} = \mathbb{P}_S$ . This is however not a problem and additionally, in practice, we would not consider adding these uninformative candidates into our solution anyway.

A function  $f$  is a submodular set function if for  $\forall S \subseteq T \subseteq 2^T$  it holds that

$$f(S \cup \{x\}) - f(S) \geq f(T \cup \{x\}) - f(T).$$

Equivalently  $\forall S, T \subseteq 2^T$ , it has to hold that

$$f(S \cap T) + f(S \cup T) \leq f(S) + f(T)$$

in order to be submodular (Krause and Golovin 2012).

**Proof 2 (submodularity of  $f$ )**

$$\begin{aligned} f(S \cup T) + f(S \cap T) &= KL(p_{S \cup T} \parallel q) + KL(p_{S \cap T} \parallel q) \\ &\leq KL(p_S \parallel q) + KL(p_T \parallel q) \\ &\quad - KL(p_{S \cap T} \parallel q) + KL(p_{S \cap T} \parallel q) \\ &= KL(p_S \parallel q) + KL(p_T \parallel q) \\ &= f(S) + f(T) \end{aligned}$$

Here, we factorized the divergence between  $p_{S \cup T}$  and  $q$  into the divergences of  $p_S$  and  $p_T$  and subtracted the divergence to the distribution  $p_{S \cap T}$  concerned with the intersection.

### Proof 3 (Approximation Bound of Greedy Algorithm)

First, we observe that we will not increase the information gain

$$f(S^*) \leq f(S^* \cup S_i),$$

by inserting patterns into the optimal sufficient statistics  $S^*$ . Lets unravel  $f(S^* \cup S_i)$  into the sum of residuals:

$$\begin{aligned} &= f(S^* \cup S_i) \\ &\quad + f(S_i \cup S^* \setminus \{x_k^*\}) - f(S_i \cup S^* \setminus \{x_k^*\}) \\ &= f(S^* \cup S_i) \\ &\quad + f(S_i \cup S^* \setminus \{x_{k-1}^*, x_k^*\}) - f(S_i \cup S^* \setminus \{x_k^*\}) \\ &\quad + f(S_i \cup S^* \setminus \{x_k^*\}) - f(S_i \cup S^* \setminus \{x_{k-1}^*, x_k^*\}) \\ &= \dots \end{aligned}$$

By continuing this process we have iteratively reduced the terms in  $S^*$

$$\begin{aligned} f(S^* \cup S_i) &= f(S^* \cup S_i) \\ &\quad + f(S_i) + \sum_{j \in [k]} f(S_i \cup S^* \setminus \{x_1^* \dots x_j^*\}) \\ &\quad - f(S_i \cup S^* \setminus \{x_1^* \dots x_{j-1}^*\}) \end{aligned}$$

First, we subtract  $f(S^* \cup S_i)$  from both sides and use the bound in Eq. (3). Now we use the sub-modularity of  $d$  to bound the residuals of the unravelled sum:

$$\begin{aligned} f(S^*) &\leq f(S_i) + \sum_{j \in [k]} f(S_i \cup S^* \setminus \{x_1^* \dots x_j^*\}) \\ &\quad - f(S_i \cup S^* \setminus \{x_1^* \dots x_{j-1}^*\}) \\ &\leq f(S_i) + \sum_{j \in [k]} f(S_i \cup \{x_j^*\}) - f(S_i) \end{aligned}$$

where we have used that  $f(S_i \cup S^* \setminus \{x_1^* \dots x_{j-1}^*\}) \leq f(S_i)$  and  $f(S_i \cup S^* \setminus \{x_1^* \dots x_j^*\}) \leq f(S_i \cup \{x_j^*\})$ .

In the  $i$ -th iteration, we denote  $\hat{x}_i$  for the candidate with highest marginal gain and by definition let  $S_{i+1} = S_i \cup \{\hat{x}_i\}$ . We can observe the trivial bound  $f(S_i \cup \{x_j^*\}) \leq f(S_{i+1})$ , which we use in the following

$$\begin{aligned} f(S^*) &\leq f(S_i) + \sum_{j \in [k]} f(S_i \cup \{x_j^*\}) - f(S_i) \\ &\leq f(S_i) + \sum_{j \in [k]} f(S_{i+1}) - f(S_i) \\ &= f(S_i) + k \cdot (f(S_{i+1}) - f(S_i)) \end{aligned}$$

Now, by together left and right hand side we get the following simple bound to the optimal solution.

$$f(S^*) \leq f(S_i) + k \cdot (f(S_{i+1}) - f(S_i)),$$

for which the following list of equivalent forms lead step-by-

step to the bound.

$$\begin{aligned}
1/k (f(S^*) - f(S_i)) &\leq f(S_{i+1}) - f(S_i) \\
\Leftrightarrow 1/k f(S^*) + (1 - 1/k)f(S_i) &\leq f(S_{i+1}) \\
\Leftrightarrow (1 - 1/k)f(S_i) &\leq f(S_{i+1}) - 1/k f(S^*) \\
\Leftrightarrow (1 - 1/k)f(S_i) &\leq f(S_{i+1}) - f(S^*) + (1 - 1/k)f(S^*) \\
\Leftrightarrow 0 \leq f(S_{i+1}) - f(S^*) &+ (1 - 1/k)(f(S^*) - f(S_i)) \\
\Leftrightarrow f(S^*) - f(S_{i+1}) &\leq (1 - 1/k)(f(S^*) - f(S_i)) \\
\Leftrightarrow g_{i+1} &\leq (1 - 1/k)g_i
\end{aligned}$$

By induction we have for step  $l$ :  $g_l \leq (1 - 1/k)^l g_0$ . Now, we use the fact that  $(1+x) \leq e^x$  and get  $g_l \leq e^{-l/k} g_0$ . Since  $g_0 \leq f(S^*)$  we finally get

$$f(S^*) - f(S_l) \leq e^{-l/k} f(S^*)$$

The solution  $S_l$  lies within the  $e^{-l/k}$  radius of  $f(S^*)$ . Set  $l = k$  then the radius is  $1/e \approx 0.368$  but setting  $l = 2k$  our solution lies within the 0.135 radius. For our practical purposes this bound is sufficiently good.  $\square$

We conclude that maximizing  $KL$  is a Submodular Function Maximization problem which has approximation guarantees and furthermore, we know that the estimates of greedily optimizing  $\ell$  and  $KL$  are equivalent.

## Experiments

All datasets we have used in our experiments are publicly available. We have taken *BMS VW*, *Adult*, *Page Blocks*, *DNA Ampl.*, *Letter Recog.*, *Anneal*, *MCADD*, *Led 7*, *Mammals*, *ICDM Abstracts*, *Waveform*, *Plants* from the UCI Machine Learning Repository and *Chess*, *Mushroom*, *Pumsb Star* from the Mining Dataset Repository.<sup>2</sup> The *DQ* dataset of lemmatized Deep-Learning and Quantum-Theory ArXiv abstracts can be found in the supplementary material.<sup>3</sup> In Table 1, we provide basic information about the datasets and the minimal support we have used in our experiments.

**Qualitative Study** Next, we study the composition of the *DQ* dataset. This is a dataset that consists of 10 000 abstracts crawled from arXiv. Half the abstracts are from papers on Deep Learning, and half from papers on Quantum Theory.

From these abstracts we remove stop words, extract and lemmatize nouns and verbs, erase words with a frequency lower than 0.05, and remove the class labels. Overall, we have a dataset over 433 items. The composition that DISC discovers on this data consists of 5 components and 224 patterns. The ‘‘Deep Learning’’ class is covered by 3 components, spanning resp. 36%, 53% and 10% of the papers on deep learning. The ‘‘Quantum’’ class consist of 2 components which cover 39.9% and 59.9% of the quantum papers. Overall, the components have a total purity of more than 99.5%. In Table 2 we give a number of exemplary patterns that DISC discovers. Common patterns across all papers include *experimental results*, *lower bound*, and *Hilbert space*, whereas a pattern such as *reinforcement agent environment* is only characteristic for one component of the deep learning papers.

<sup>2</sup><https://archive.ics.uci.edu/ml>, <http://fimi.ua.ac.be/data/>

<sup>3</sup><https://eda.mmci.uni-saarland.de/disc/>

Dataset	$ D $	dim $D$	$k^c$	Min Support
BMS WV 1	59,602	497	1	32
Mushroom	8,124	23	2	10
Adult	48,842	15	2	5
Page Blocks	5,473	11	5	1
DNA Amp.	4,590	392	1	5
Chess Big	3,196	37	18	319
Let. Recog.	20,000	17	26	1
DQ	9,993	433	1	99
Anneal	898	71	5	1
Pumsb Star	49,046	7,116	1	12,500
MCADD	31,924	198	2	50
Chess	28,056	7	2	5
Led 7	3,200	8	10	40
Mammals	2,183	121	1	5
ICDM	859	3,932	1	10
Waveform	5,000	22	3	5
Plants	34,781	68	1	5

Table 1: **Datasets** The sizes, dimensionality, number of classes, and the minimal support of patterns for datasets we have used in the experiments in Sec. 17 that compare DESC and DISC with MTV, SLIM and DIFFNORM.




Figure 4: **Distribution of Component Sizes** Shown are the relative sizes  $|D_j| / |D|$  of the components that DISC discovers




Figure 5: **Purity of the Partitioning** For 9 labeled datasets, we report purity of the partitionings result from DBDESC, DBDESC2 and DISC

Pattern Set	“Quantum” Patterns
$S_{1,2}$	local entanglement, Bell inequality, standard model
$S_1$	standard approach, learn data, research paper
$S_2$	computer computation, search algorithm
Pattern Set	“Deep” Patterns
$S_{3,4,5}$	neural networks, hidden layer, computer vision
$S_{3,4}$	information prediction, space representation
$S_3$	neural processing, reinforcement environment agent
$S_4$	feature representation, training optimization
Pattern Set	“Common” Patterns
$S_{1,2,3,4,5}$	experimental results, lower bound, Hilbert space

Table 2: **DISC Discovers Interpretable Compositions**  
 Shown is a selection from the pattern composition that DISC discovers on the *Deep Quantum* dataset. Overall it discovers 223 characteristic and common patterns and 5 components, two of which consist of mostly “quantum” papers, and three of which of “deep learning” papers.




Figure 6: **Runtime in Minutes** of DBDESC, DBDESC2 and DISC on all 17 datasets

Dataset	$k^c$	$\hat{k}$	S				$\ell/\ell_\emptyset$ [%]			Purity	PSKL		Time	
			DIFFNORM	DESC	DESC $_\pi$	DISC	DESC	DESC $_\pi$	DISC	DISC	DESC $_\pi$	DISC	DESC	DISC
Adult	2	56	719	28	28	29	88.9	83.1	56.5	0.8	8.3	63.3	899ms	11h36m20s
Anneal	5	4	95	17	17	17	91	83.9	78.5	0.8	31.1	30.1	101ms	17s
Chess	2	10	192	44	44	44	88.1	84.7	71.6	0.6	5.6	53.7	278ms	18m39s
Chess Big	18	68	869	13	13	13	98	91.4	69.3	0.3	9.2	26.9	55ms	52m19s
Led 7	10	20	30	10	10	10	89.6	65.5	54.2	0.7	17	30.6	45ms	1m54s
Let. Recog.	26	50	1398	47	47	49	88.9	79	67.6	0.3	20.5	65.5	710ms	15h52m2s
Mushroom	2	17	424	66	79	70	82	77	48.7	1	59.6	133.6	660ms	32m24s
Page Blocks	5	4	30	10	10	10	98.2	68	63.2	0.9	4.9	18.6	95ms	14s
Waveform	3	14	535	35	35	36	89.1	82.2	74.2	0.7	31.6	57.7	309ms	1h45m

Table 3: **DESC Discovers Informative Patterns, DISC Discovers Informative Components** For 9 datasets with class labels we compare DIFFNORM, DESC on the full data, DESC $_\pi$  given the class-label decomposition, and DISC. We give the true number of classes ( $k^c$ ), the number of components that DISC discovers ( $\hat{k}$ ), the number of patterns discovered ( $|S|$ ), the likelihood ratio ( $\ell/\ell_\emptyset$ , lower is better), the class-label purity of the components discovered by DISC (higher is better), the divergence (PSKL, higher is better) between the ground truth resp. discovered components, and the runtime (s, lower is better)

Dataset	DESC		$\hat{k}$	DISC			Time
	S	$\ell/\ell_\emptyset$		S	$\ell/\ell_\emptyset$	PSKL	
BMS WV 1	16	97.1	25	28	84.3	14.8	3h21m49s
DNA Amp.	76	73.1	3	84	73.6	50	1m6s
DQ	43	96.9	5	55	91.3	51.4	11m45s
ICDM	16	99.5	1	16	99.5	0	2m
MCADD	95	94	18	104	91.2	41.2	1h13m21s
Mammals	59	80.4	9	64	66.1	78.1	7m34s
Plants	58	61.2	52	60	42.4	125.3	13h5m3s
Pumsb Star	53	84.9	10	53	64.6	268.4	1h34m22s

Table 4: **DESC and DISC on Unlabelled Data** For 8 dataset without class labels, for DESC and DISC we report the number of patterns ( $|S|$ ), runtime (s), and for DISC we additionally report the number of components discovered ( $\hat{k}$ ), the gain in likelihood ( $\ell/\ell_\emptyset$ ) in percent, and the PSKL-divergence between the discovered components

Dataset	DBDESC					DBDESC2					DISC				
	k	$ S $	$\ell/\ell_0[\%]$	PSKL	run time	k	$ S $	$\ell/\ell_0[\%]$	PSKL	run time	k	$ S $	$\ell/\ell_0[\%]$	PSKL	run time
BMS WV 1	7	16	97.8	50.1	6m2s	2	16	96.6	14.9	13m33s	25	28	84.3	14.8	3h21m49s
Mushroom	23	66	47.7	123.4	7s	11	66	55.2	127.7	36s	17	70	48.7	133.6	32m24s
Adult	8	28	79.1	39.4	4m1s	154	28	80.5	47.6	4m41s	56	29	56.5	63.3	1h36m20s
Page Blocks	16	10	67.1	21.5	4s	7	10	63.4	27	5s	4	10	63.2	18.6	14s
DNA Amp.	27	76	109.2	64.3	15s	5	76	81.3	36.3	12s	3	84	73.6	50	1m6s
Chess Big	1	14	98.7	0	10s	42	13	78.7	33.5	38s	68	13	69.3	26.9	52m19s
Let. Recog.	2	48	88.2	16.5	2m27s	2	48	88.4	13.9	4m12s	50	49	67.6	65.5	15h52m2s
DQ	2	43	97.9	62.4	5s	2	43	93.7	64.6	1m5s	5	55	91.3	51.4	11m45s
Anneal	3	17	89	34.8	532ms	2	17	86.3	30.9	783ms	4	17	78.5	30.1	17s
Pumsh Star	23	53	77	303.6	6m54s	45	53	79.5	288.1	3h17m38s	10	53	64.6	268.4	1h34m22s
MCADD	1	118	95.3	0	19s	2	95	95.8	12.3	4m10s	18	104	91.2	41.2	1h13m21s
Chess	2	46	86.5	29.3	3s	9	44	86.7	82.1	16s	10	44	71.6	53.7	18m39s
Led 7	75	10	62.3	29.4	1s	47	10	58.7	28.7	1s	20	10	54.2	30.6	1m54s
Mammals	3	59	82.6	57.5	3s	3	59	84.7	54.2	4s	9	64	66.1	78.1	7m34s
ICDM	1	16	99.6	0	1s	2	16	105.6	36.5	52s	1	16	99.5	0	2m
Waveform	6	35	83.9	84.5	3s	4	36	88.5	38.2	14s	14	36	74.2	57.7	1h45m
Plants	61	58	64.9	10.3	1m17s	49	58	61.4	24.2	5m29s	52	60	42.4	125.3	13h5m3s

Table 5: **DBDESC, DBDESC2 and DISC** For all datasets in this experiment, we report the number of patterns ( $|S|$ ), runtime, the number of components discovered ( $k$ ), the gain in likelihood ( $\ell/\ell_0$ ) in percent, and the PSKL-divergence between the discovered components