
Conditional Bernoulli Mixtures for Multi-Label Classification

Cheng Li
Bingyu Wang
Virgil Pavlu
Javed Aslam

CHENGLI@CCS.NEU.EDU
RAINICY@CCS.NEU.EDU
VIP@CCS.NEU.EDU
JAA@CCS.NEU.EDU

College of Computer and Information Science, Northeastern University, Boston, MA 02115, USA

Abstract

Multi-label classification is an important machine learning task wherein one assigns a subset of candidate labels to an object. In this paper, we propose a new multi-label classification method based on *Conditional Bernoulli Mixtures*. Our proposed method has several attractive properties: it captures label dependencies; it reduces the multi-label problem to several standard binary and multi-class problems; it subsumes the classic independent binary prediction and power-set subset prediction methods as special cases; and it exhibits accuracy and/or computational complexity advantages over existing approaches. We demonstrate two implementations of our method using logistic regressions and gradient boosted trees, together with a simple training procedure based on Expectation Maximization. We further derive an efficient prediction procedure based on dynamic programming, thus avoiding the cost of examining an exponential number of potential label subsets. Experimental results show the effectiveness of the proposed method against competitive alternatives on benchmark datasets.

1. Introduction

Multi-label classification is an important machine learning task wherein one assigns a *subset* of candidate labels to an object. Such problems arise naturally in the context of medical diagnostics, text classification, image tagging, and a host of other areas. For example, in the medical domain, a patient may present multiple illnesses or undergo a procedure with multiple billing codes; in news categorization, an article can be associated with multiple categories; in image tagging, a photo may have multiple tags (see Figure 1).

Proceedings of the 33rd International Conference on Machine Learning, New York, NY, USA, 2016. JMLR: W&CP volume 48. Copyright 2016 by the author(s).



Figure 1. An image from the NUS-WIDE multi-label dataset. Independent binary logistic regressions predict the labels {clouds, lake, sky, sunset, water}. Our proposed method correctly predicts {clouds, lake, sky, sunset, water, reflection}, because it captures the dependencies between reflection and the other labels.

Multi-label classification is particularly difficult in the common case when the labels are *many* and *dependent*. In this case, simple approaches such as independent binary label prediction and power-set subset prediction fail due to label dependencies and combinatorial explosion, respectively. More sophisticated approaches such as classifier chains, conditional dependency networks, and conditional random fields suffer due to accuracy/computational-complexity trade-offs in training and/or prediction. We introduce *Conditional Bernoulli Mixtures* for multi-label classification to address these challenges, yielding multi-label classifiers that are flexible, accurate and efficient.

Formally, in a multi-label classification problem, we are given a set of label candidates $\mathcal{Y} = \{1, 2, \dots, L\}$. Every data point $\mathbf{x} \in \mathbb{R}^D$ matches a subset of labels $\mathbf{y} \subseteq \mathcal{Y}$, which can be equivalently written in the form of a binary vector $\mathbf{y} \in \{0, 1\}^L$, with each bit y_ℓ indicating the presence or absence of the corresponding label. The goal of learning is to build a classifier $h : \mathbb{R}^D \mapsto \{0, 1\}^L$ which maps an in-

stance to a subset of labels. The predicted label subset can be of arbitrary size; when the size is restricted to be 1, the problem is called multi-class classification. Furthermore, if the total number of label candidates L is 2, the problem becomes binary classification.

Various measures can be used for evaluating multi-label classifiers. Among them, three commonly used measures are *subset accuracy*, *Jaccard index* and *Hamming loss*. Let $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$ be a multi-label dataset with ground truth labels, and $\{\hat{\mathbf{y}}_n\}_{n=1}^N$ be the predictions made by a classifier. Subset accuracy generalizes the conventional multi-class accuracy notion:

$$\frac{1}{N} \sum_{n=1}^N \mathbb{1}[\mathbf{y}_n = \hat{\mathbf{y}}_n], \quad (1)$$

where $\mathbb{1}[\cdot]$ is the indicator function. In computing subset accuracy, a predicted subset is considered correct only when it matches the true subset exactly. Admittedly, this metric is very stringent in evaluation. However, optimizing for subset accuracy is a very interesting algorithmic design and research problem, as it encourages the classifiers to output coherent and complete predictions. The optimal prediction for subset accuracy is achieved by outputting the mode of the conditional joint distribution:

$$h^*(\mathbf{x}) = \underset{\mathbf{y}}{\operatorname{argmax}} p(\mathbf{y}|\mathbf{x}). \quad (2)$$

Jaccard index (also called *overlap*) measures the overlap between the true subset and the predicted subset using the size of their intersection divided by the size of their union, and Hamming loss counts the number of bit-wise mistakes. The theoretical results in (Dembczyński et al., 2012; Koyejo et al., 2015) show that a multi-label classifier designed for optimizing one measure may be suboptimal under other measures. In this paper, we focus on the multi-label classification task with subset accuracy as the target evaluation measure, but for completeness we also report the Jaccard index and Hamming loss results in the supplementary material.

According to (2), optimizing subset accuracy requires estimating the conditional joint distribution $p(\mathbf{y}|\mathbf{x})$, which captures conditional label dependencies given features. One naive approach is to assume conditional independence among labels $p(\mathbf{y}|\mathbf{x}) = \prod_{\ell=1}^L p(y_\ell|\mathbf{x})$, which reduces a multi-label problem into L binary classification problems. This approach is called *Binary Relevance* (Bin-Rel) (Tsoumakas & Katakis, 2007), and is widely used due to its simplicity. One obvious disadvantage is that the individual label predictions can often be conflicting, such as tagging `cat` but not `animal` for an image.

At the other extreme is the *Power-Set* (PowSet) approach (Tsoumakas & Katakis, 2007), treating each label

subset as a class, and trains a multi-class classifier. As a consequence, one would be restricted in practice to predicting only label subsets seen in the training data and always assigning zero probabilities to unseen subsets; even for many of the subsets observed there would likely be a scarcity of training data. Power-Set is handling label dependencies and its predictions are coherent, but the method is often infeasible on the exponential number of label sets.

Bernoulli Mixtures. Mixture models offer a flexible and powerful framework for many multivariate density estimation problems. A mixture generally has the form

$$p(\mathbf{y}) = \sum_{k=1}^K \pi^k p(\mathbf{y}; \boldsymbol{\beta}^k), \quad (3)$$

which approximates a complex joint $p(\mathbf{y})$ by a weighted combination of K component densities $p(\mathbf{y}; \boldsymbol{\beta}^k)$, each of which typically takes some simple density form parametrized by $\boldsymbol{\beta}^k$. The Expectation Maximization (EM) algorithm can be employed to train such mixture models by iterating between estimating the mixture coefficients and fitting component densities.

Bernoulli Mixtures (BM) are classic models for multi-dimensional binary variable density estimation (Lazarsfeld et al., 1968; McLachlan & Peel, 2004; Bishop, 2006), where the learnability is realized by assuming independence of variables within each mixture component. Thus each component density $p(\mathbf{y}; \boldsymbol{\beta}^k)$ is simply a product of Bernoulli densities and the overall model has the form

$$p(\mathbf{y}) = \sum_{k=1}^K \pi^k \prod_{\ell=1}^L \operatorname{Ber}(y_\ell; \mu_\ell^k), \quad (4)$$

where $\operatorname{Ber}(y_\ell; \mu_\ell^k)$ denotes the Bernoulli distribution with head probability μ_ℓ^k . BM has two attractive properties. First, *dependencies*: although the L variables are assumed to be independent inside each component, they are in general dependent in the mixture (they are forced to be independent only when $K = 1$). This can be verified by computing the following covariance matrix and observing that it is non-diagonal for $K \geq 2$ (Bishop, 2006).

$$\operatorname{Cov}[\mathbf{y}] = \sum_{k=1}^K \pi^k [\boldsymbol{\Sigma}^k + \boldsymbol{\mu}^k (\boldsymbol{\mu}^k)^\top] - \mathbb{E}(\mathbf{y}) \mathbb{E}(\mathbf{y})^\top, \quad (5)$$

where $\mathbb{E}(\mathbf{y}) = \sum_{k=1}^K \pi^k \boldsymbol{\mu}^k$, $\boldsymbol{\Sigma}^k = \operatorname{diag}\{\mu_\ell^k(1 - \mu_\ell^k)\}$, and $\boldsymbol{\mu}^k = (\mu_1^k, \mu_2^k, \dots, \mu_L^k)^\top$.

Second, *efficiency*: the full factorization of each component makes fitting BM efficient via the EM algorithm.

In this paper, we propose a new multi-label classification method which approximates the conditional joint $p(\mathbf{y}|\mathbf{x})$ based on Bernoulli Mixtures. The method simultaneously *learns* the label dependencies from data and *trains* classifiers to account for such dependencies.

2. Conditional Bernoulli Mixtures

For multi-label classification, we model the conditional joint $p(\mathbf{y}|\mathbf{x})$ with a discriminative extension of BM, capturing conditional dependencies among binary labels given features. The analysis in (Dembczyński et al., 2012) shows that labels could be largely conditionally independent given features (i.e., $p(\mathbf{y}|\mathbf{x}) \approx \prod_{\ell=1}^L p(y_\ell|\mathbf{x})$), even when labels are strongly marginally dependent (i.e., $p(\mathbf{y}) \neq \prod_{\ell=1}^L p(y_\ell)$), as long as each label is highly predictable from features. Therefore it is not necessary to capture correlations among easy-to-predict labels; it is sufficient to only capture those correlations which involve some hard-to-predict labels. Thus conditioning on features greatly reduces the need to estimate label correlations and makes learning much easier; *the conditional on \mathbf{x} is essential for good approximation and effective training*, without making prior assumptions about the form of label dependencies (as many other methods do).

Making both mixture coefficients and Bernoulli distributions conditional on \mathbf{x} , we obtain our proposed model, Conditional Bernoulli Mixtures (CBM):

$$p(\mathbf{y}|\mathbf{x}) = \sum_{k=1}^K \pi(z = k|\mathbf{x}; \boldsymbol{\alpha}) \prod_{\ell=1}^L b(y_\ell|\mathbf{x}; \boldsymbol{\beta}_\ell^k), \quad (6)$$

where $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}_\ell^k$ ($\ell = 1, 2, \dots, L; k = 1, 2, \dots, K$) are model parameters to be learned, and z is a hidden categorical indicator variable, such that $z = k$ if the data point is assigned to component k . Here we use π and b to represent the conditional mixture membership distribution and the conditional binary label distribution, respectively.

CBM reduces a multi-label problem to a multi-class problem and several binary problems, approaching $p(\mathbf{y}|\mathbf{x})$ akin to *divide and conquer*: the categorical distribution $\pi(z|\mathbf{x}; \boldsymbol{\alpha})$ assigns each instance \mathbf{x} to 1 out of K components probabilistically. Its goal is to divide the feature space into several regions such that each region only has weak conditional label correlations and can be approximated by a simple component. This gating function $\pi(z|\mathbf{x}; \boldsymbol{\alpha})$ can be instantiated by any multi-class classifier which provides probability estimations, such as a multinomial logistic regression with parameters $\boldsymbol{\alpha}$.

Inside each region k , the local conditional joint density is approximated by a product of conditional marginal densities. Every local binary label predictor $b(y_\ell|\mathbf{x}; \boldsymbol{\beta}_\ell^k)$ estimates the probability of getting label y_ℓ from component k for data point \mathbf{x} , and can be instantiated by any binary classifier which provides probability estimations, such as a binary logistic regression with parameters $\boldsymbol{\beta}_\ell^k$. All K components together with the gating function are learned jointly in order to break the global label correlation into simple forms.

On one extreme, if CBM has only one component (and hence $\pi(z|\mathbf{x}; \boldsymbol{\alpha})$ plays no role), all labels are conditionally independent and CBM degenerates to Binary Relevance. On the other extreme, if CBM assigns one component to each unique label subset and fixes each $b(y_\ell = 1|\mathbf{x}; \boldsymbol{\beta}_\ell^k)$ to be the corresponding binary constant (1 if label ℓ is in the subset and 0 otherwise), then the overall CBM model simply selects one label subset from all possible subsets, which is conceptually the same as the Power-Set approach. By varying the number of components and the complexity of each component, CBM can provide a continuous spectrum between these two extremes. Just as Binary Relevance and Power-Set, CBM is purely a reduction method. The main advantage of reduction methods compared with new models specifically designed for multi-label problem is that the reduction approach makes it easier to incorporate and reuse many well-developed multi-class and binary classifiers. In Section 3, we will demonstrate two different instantiations of CBM, one with logistic regressions, and the other with gradient boosted trees.

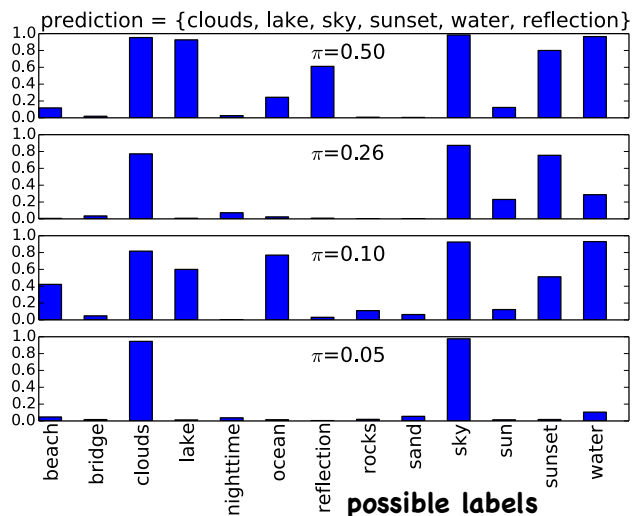


Figure 2. The top 4 most influential components for the test image in Figure 1. π values indicate component mixing coefficients. Each bar represents an individual label probability in one component. Labels with near zero probabilities are not displayed here.

An Illustrative Example. Before diving into the model training details, we illustrate how CBM performs joint label classifications with an example. Figure 1 shows a test image in the NUS-WIDE dataset, for which the matched label *reflection* is missed by Binary Relevance but is captured by CBM. For this image, the most influential components produced by CBM are shown in Figure 2. Simply averaging individual label probabilities by component mixing weights gives the conditional marginals $p(\text{lake}|\mathbf{x}) = 0.56$, $p(\text{water}|\mathbf{x}) = 0.69$, $p(\text{sunset}|\mathbf{x}) = 0.66$, and $p(\text{reflection}|\mathbf{x}) = 0.32$, which indicate that unlike

lake, water and sunset, the label reflection by itself is not deemed as probable by CBM. However from the CBM joint density $p(\mathbf{y}|\mathbf{x})$ we can also infer Pearson correlation coefficients $\rho_{\text{reflection, lake}} = 0.50$, $\rho_{\text{reflection, water}} = 0.40$, $\rho_{\text{reflection, sunset}} = 0.17$ and observe that reflection is positively correlated with lake, water, and sunset. In fact, the joint probability asserts that the subset {clouds, lake, sky, sunset, water, reflection} is the most probable one, with probability 0.09. By contrast, the subset {clouds, lake, sky, sunset, water} has a lower probability 0.06. Therefore, although individually unlikely, the label reflection is correctly added to the mostly likely subset due to label correlations.

3. Training CBM with EM

CBM can be trained by maximum likelihood estimation on a given dataset $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$. Below we first derive a generic EM algorithm that works for any instantiation of the components, and then consider two concrete instantiations.

To make the mathematical derivation clear, we use different symbols for random variables and values of random variables. We use \mathbf{Y}_n when treating the labeling of \mathbf{x}_n as an unknown random variable and \mathbf{y}_n when referring to its specific labeling assignment given in the training set. The likelihood for the dataset is

$$\prod_{n=1}^N \left\{ \sum_{k=1}^K [\pi(z_n = k | \mathbf{x}_n; \boldsymbol{\alpha}) \prod_{\ell=1}^L b(y_{n\ell} | \mathbf{x}_n; \boldsymbol{\beta}_\ell^k)] \right\}.$$

Since the model contains hidden variables, we use EM to minimize an upper bound of the negative log likelihood. Denoting the posterior membership distribution $p(z_n | \mathbf{x}_n, \mathbf{y}_n)$ as $\Gamma(z_n) = (\gamma_n^1, \gamma_n^2, \dots, \gamma_n^K)$, the upper bound can be written as

$$\sum_{n=1}^N \mathbb{KL}(\Gamma(z_n) || \pi(z_n | \mathbf{x}_n; \boldsymbol{\alpha})) + \sum_{k=1}^K \sum_{\ell=1}^L \sum_{n=1}^N \gamma_n^k \mathbb{KL}(\text{Ber}(Y_{n\ell}; y_{n\ell}) || b(Y_{n\ell} | \mathbf{x}_n; \boldsymbol{\beta}_\ell^k)), \quad (7)$$

where $\text{Ber}(Y_{n\ell}; y_{n\ell})$ is the Bernoulli distribution with head probability $y_{n\ell}$.

E Step: Re-estimate the posterior membership probability of each data point belonging to each component:

$$\gamma_n^k = \frac{\pi(z_n = k | \mathbf{x}_n; \boldsymbol{\alpha}) \prod_{\ell=1}^L b(y_{n\ell} | \mathbf{x}_n; \boldsymbol{\beta}_\ell^k)}{\sum_{k=1}^K \pi(z_n = k | \mathbf{x}_n; \boldsymbol{\alpha}) \prod_{\ell=1}^L b(y_{n\ell} | \mathbf{x}_n; \boldsymbol{\beta}_\ell^k)}. \quad (8)$$

M Step: Update all model parameters. The bound (7) shows that the optimization of all parameters can be nicely

decomposed into a series of separate optimization problems:

$$\boldsymbol{\alpha}_{\text{new}} = \underset{\boldsymbol{\alpha}}{\text{argmin}} \sum_{n=1}^N \mathbb{KL}(\Gamma(z_n) || \pi(z_n | \mathbf{x}_n; \boldsymbol{\alpha})), \quad (9)$$

$$\boldsymbol{\beta}_{\ell}^k_{\text{new}} = \underset{\boldsymbol{\beta}_\ell^k}{\text{argmin}} \sum_{n=1}^N \gamma_n^k \mathbb{KL}(\text{Ber}(Y_{n\ell}; y_{n\ell}) || b(Y_{n\ell} | \mathbf{x}_n; \boldsymbol{\beta}_\ell^k)). \quad (10)$$

The optimization problem defined in (9) is a multi-class classification problem with target class distribution (soft labels) $\Gamma(z_n) = (\gamma_n^1, \gamma_n^2, \dots, \gamma_n^K)$ for \mathbf{x}_n . The training goal for $\pi(z | \mathbf{x}; \boldsymbol{\alpha})$ is to assign (probabilistically) each data point to a component based only on its features, such that the assignment matches the posterior component membership determined by both features and true labels.

The optimization problem defined in (10) is a weighted binary classification problem. \mathbf{x}_n has target label $y_{n\ell}$ and weight γ_n^k . The prediction component $b(y_\ell | \mathbf{x}; \boldsymbol{\beta}_\ell^k)$ is trained as a binary classifier for label ℓ in component k , based only on training data within (soft) component k .

To train CBM, we iterate between the E step and the M step, until the upper bound (7) converges (see Algorithm 1). In practice, the EM algorithm can get stuck in local op-

Algorithm 1 Generic Training for CBM

- 1: **repeat**
 - 2: E Step
 - 3: **for** $n = 1, 2, \dots, N; k = 1, 2, \dots, K$ **do**
 - 4: update γ_n^k as in (8)
 - 5: M Step
 - 6: update $\boldsymbol{\alpha}$ as in (9)
 - 7: **for** $k = 1, 2, \dots, K; \ell = 1, 2, \dots, L$ **do**
 - 8: update $\boldsymbol{\beta}_\ell^k$ as in (10)
 - 9: **until** convergence
-

tima, so careful initializations are often necessary for good performance. Due to the natural connection between CBM and BM, one simple way of initializing CBM is to first fit a BM just for label density estimation without looking at features. BM can be trained very quickly by a simpler EM algorithm (Bishop, 2006). We can use several random starts for a set of BM models, and select the one with the best training objective score to initialize the CBM training procedure.

3.1. CBM with Logistic Regression Learners

In the simplest form all underlying models in CBM are linear. We employ a multinomial logistic regression for the gating function π , and a binary logistic regression for each predictor b . The outer loop of the training procedure is the

EM algorithm described above. In each M step, all objectives defined in (9) and (10) are convex (although the overall EM optimization is non-convex). To optimize each logistic regression, we compute the gradient of the corresponding objective w.r.t. its parameters and update all parameters using gradient based optimization methods. Here we choose L-BFGS method (Nocedal & Wright, 2006), which is the state-of-the-art optimization method for large scale logistic regressions. We can start with the existing model, and perform a few incremental update steps until it converges. To avoid over-fitting, we also add L_2 regularizations (Gaussian priors) to all parameters.

CBM+LR Complexity. The overall complexity of the training procedure for CBM is dominated by the updates in M steps and depends on the complexities of the binary and multi-class learning algorithms used. For CBM with logistic regression (LR) learners, we can measure the complexity of each L-BFGS (or gradient descent) update for all model parameters. If N is the number of instances, D the number of features, L the number of labels, C the number of unique label subsets in the training set, and K the number of CBM components, then each CBM+LR update takes $\mathcal{O}(KLDN)$. For comparison, each update takes $\mathcal{O}(LDN)$ in Binary Relevance with LR, and takes $\mathcal{O}(CDN)$ in Power-Set with LR. For large datasets, typically $KL < C$, and CBM is slower than Binary Relevance but faster than Power-Set.

3.2. CBM with Gradient Boosting Learners

Many datasets require non-linear decision boundaries, in which case the CBM model with logistic regressions may not have enough explanation power. To make CBM non-linear, we use gradient boosted trees (GB) for both π and b . The original gradient boosted trees algorithm described in (Friedman, 2001) is designed for standard multi-class problems, and does not take label distributions or instance weights as inputs, thus some modifications are necessary. The target label distribution is easy to deal with: one still computes the functional gradient of the objective w.r.t. each ensemble scoring function, where the objective is defined by the target distribution. To handle instance weights, one ignores them when calculating functional gradients, and performs a weighted least squares fit when fitting each regression tree to the gradients. Unlike logistic regression, where all parameters can be easily updated, gradient boosting introduces new trees to the ensemble while keeping old trees untouched. Thus the M step here is slightly different from before: rather than re-adjusting all parameters to optimize the objective, boosting improves the objective by adding a few more trees. This amounts to a partial M step, and the resulting EM algorithm is usually called the generalized EM algorithm (Gupta & Chen, 2011). We emphasize that the use of boosting as an underlying non-

linear classifier here is just for demonstration purposes. Other classifiers such as neural networks could also be used. This is in direct contrast to AdaBoost.MH and Adaboost.MR (Schapire & Singer, 2000) which specifically employ boosting to optimize Hamming loss and rank loss.

3.3. Validation of Training Procedure on Artificial Data

To validate the training procedure, we fit a CBM model on artificial data generated by a “true CBM” and see whether the training procedure could recover the true parameters. The “true CBM” has $K = 3$ components and uses logistic regressions for both π and b . Each of the $N = 15,000$ data points \mathbf{x} has $D = 7$ features generated independently from Gaussian distributions. We test two different ways of generating the ground-truth label subset \mathbf{y} : *argmax*, where the most likely \mathbf{y} is always selected; and *sampling*, where \mathbf{y} is sampled from the true CBM distribution. For each of these two datasets, we also make a *noise* variant with Gaussian noise added to logistic regression scores. Test

Table 1. Test subset accuracy on artificial data.

Methods	$\mathbf{y}=\text{argmax}$	$\mathbf{y}=\text{argmax}$ +noise	$\mathbf{y}=\text{sample}$	$\mathbf{y}=\text{sample}$ +noise
Upper bound	100	76.4	65.2	48.1
CBM+LR:K=2	89.9	69.3	60.2	45.8
CBM+LR:K=3	98.4	75.8	65.0	48.0
BinRel+LR	82.0	49.9	48	40.0
PowSet+LR	97.5	74.1	64.6	47.3

subset accuracy results are presented in Table 1. It shows that if the true distribution is indeed generated by a CBM, then the proposed training procedure is quite effective in recovering the distribution.

4. Fast Prediction by Dynamic Programming

According to (2), making the optimal prediction in terms of subset accuracy for a given \mathbf{x} requires finding the most probable label subset $\mathbf{y}^* = \text{argmax}_{\mathbf{y}} p(\mathbf{y}|\mathbf{x})$. There are 2^L label subset candidates, and it is intractable to evaluate the probability for each of them. Many multi-label methods suffer from this intractability for exact inference (see Section 5). Fortunately for CBM, its special structure allows it to make predictions efficiently, with either *sampling* or *dynamic programming*. As a common preprocessing step, for a given \mathbf{x} , we can first compute $\pi^k = \pi(z = k|\mathbf{x}; \boldsymbol{\alpha})$ and $\mu_\ell^k = b(y_\ell = 1|\mathbf{x}; \beta_\ell^k)$, for all $k = 1, 2, \dots, K$ and $\ell = 1, 2, \dots, L$.

Prediction by Sampling. The CBM density form (6) suggests a natural sampling strategy for a label subset \mathbf{y} . We first sample a component k according to the mixture coefficients π^1, \dots, π^K . Then from this component, we sample each label y_ℓ independently with probability μ_ℓ^k . The procedure can be repeated multiple times to generate a set of

y candidates, from which we pick the most probable one. The sampling strategy works best when the most probable label subset has a high probability. Sampling is easy to implement, but does not guarantee that the predicted y is indeed the optimal one.

Prediction by Dynamic Programming and Pruning. In order for the overall probability $p(\mathbf{y}|\mathbf{x})$ to be high, there must exist a component k for which the component probability $\prod_{\ell=1}^L b(y_\ell|\mathbf{x}; \beta_\ell^k)$ is high. On the other hand, one can show that the \mathbf{y}^* maximizing the overall probability does not necessarily maximize any component probability. To find \mathbf{y}^* , we design a dynamic programming procedure FIND-NEXT-HIGHEST() that enumerates label subsets in a decreasing probability order in each component, and then we iterate round-robin across components until we are certain that the unchecked subsets will never produce a high overall probability.

Algorithm 2 Prediction by Dynamic Prog. and Pruning

```

1: Input:  $\pi^k, \mu_\ell^k, k = 1, 2, \dots, K, \ell = 1, 2, \dots, L$ 
2: Initialize candidate component set  $\mathcal{S} = \{1, 2, \dots, K\}$ 
3: Initialize the maximum overall probability  $M = -\infty$ 
4: for  $k = 1, 2, \dots, K$  do
5:   Initialize the maximum component probability  $G^k$ 
6:   Initialize the priority queue  $Q^k$ 
7: while  $\mathcal{S} \neq \Phi$  do
8:   for  $k \in \mathcal{S}$  do
9:      $\mathbf{y} = Q^k.\text{FIND-NEXT-HIGHEST}()$ 
10:    Let  $p = \sum_{m=1}^K \pi^m \prod_{\ell=1}^L \text{Ber}(y_\ell; \mu_\ell^m)$  and  $q = \prod_{\ell=1}^L \text{Ber}(y_\ell; \mu_\ell^k)$ 
11:    if  $p > M$  then
12:      Set  $M = p$  and  $\mathbf{y}^* = \mathbf{y}$ 
13:      if  $\pi^k q \leq M/K$  or  $\pi^k q + \sum_{r \neq k} \pi^r G^r \leq M$  then
14:        Remove  $k$  from  $\mathcal{S}$ 
15: Output:  $\mathbf{y}^*$ 

```

```

16: function FIND-NEXT-HIGHEST()
17:    $\mathbf{y} = Q^k.\text{deque}()$ 
18:   for  $\ell = 1, 2, \dots, L$  do
19:     Generate  $\mathbf{y}'$  by flipping the  $\ell$ -th bit of  $\mathbf{y}$ 
20:     if  $\mathbf{y}'$  is unseen then
21:        $Q^k.\text{enqueue}(\mathbf{y}')$ 
22: Output:  $\mathbf{y}$ 

```

For a component k , the \mathbf{y} with the highest component probability is the set containing precisely all the ℓ with $\mu_\ell^k \geq 1/2$. To produce a ranked list, we use a max priority queue Q^k to store candidate label subsets and their associated component probabilities. Initially, Q^k contains the \mathbf{y} with the highest component probability. The t -th call to $Q^k.\text{FIND-NEXT-HIGHEST}()$ returns the \mathbf{y} with the t -th highest component probability. The overall prediction al-

gorithm (Algorithm 2) iterates over all components, checks the next candidate (Lines 9-10), updates the best label subset found so far (Lines 11-12), and prunes the remaining subset candidates in a component's ranked list (Lines 13-14). The algorithm terminates when all components' ranked lists have been pruned, and surely returns the most probable \mathbf{y}^* . In practice, we observe that the algorithm rarely visits elements deeper than rank 10 in the component ranked lists.

5. Related Work

Modeling label dependencies has been long regarded as a central theme in multi-label classification. Estimating the high-dimensional conditional joint $p(\mathbf{y}|\mathbf{x})$ is very challenging and various approaches have been proposed to tackle this problem based on different approximations.

Classifier Chains. According to the chain rule, the joint probability $p(\mathbf{y}|\mathbf{x})$ can be decomposed into a product of conditionals $p(y_1|\mathbf{x})p(y_2|\mathbf{x}, y_1) \cdots p(y_L|\mathbf{x}, y_1, \dots, y_{L-1})$. This reduces a multi-label learning problem to L binary learning problems, each of which learns a new label given all previous labels. During prediction, finding the exact joint mode is intractable. Classifier Chains (CC) (Read et al., 2011) classify labels greedily in a sequence: label y_ℓ is decided by maximizing $p(y_\ell|\mathbf{x}, y_1, \dots, y_{\ell-1})$, and becomes a feature to be used in the prediction for label $y_{\ell+1}$. This greedy prediction procedure has three issues: 1) the predicted subset can be far away from the joint mode (Dembczyński et al., 2011); 2) errors in early label predictions propagate to subsequent label predictions; 3) the overall prediction depends on the chain order. To address the first two issues, Probabilistic Classifier Chains (PCC) replace the greedy search strategy with some more accurate search strategies, such as exhaustive search (Cheng et al., 2010), ϵ -approximate search (Dembczynski et al., 2012), Beam Search (Kumar et al., 2012; 2013), or A* search (Mena et al., 2015). To address the third issue, Ensemble of Classifier Chains (ECC) (Read et al., 2011) averages several predictions made by different chains, wherein the averaging can take place at either the individual label level (ECC-label) or the label subset level (ECC-subset). Using dynamic programming to find the optimal chain order (Liu & Tsang, 2015) has been proposed recently.

A similar reduction method named **Conditional Dependency Networks** (CDN) estimates $p(\mathbf{y}|\mathbf{x})$ based on full conditionals and Gibbs sampling (Guo & Gu, 2011). During learning, one binary classifier is trained for each full conditional $p(y_\ell|\mathbf{x}, y_1, \dots, y_{\ell-1}, y_{\ell+1}, \dots, y_L)$. During prediction, Gibbs sampling is used to find the mode of the joint. The method's major limitation is that it cannot handle perfectly correlated or anti-correlated labels: consider

binary classification as multi-label problem with only 2 exhaustive and exclusive labels. A perfect model for $p(y_1 = 1 | \mathbf{x}, y_2)$ is $1 - y_2$; in other words, the feature information is completely ignored. The same applies to $p(y_2 | \mathbf{x}, y_1)$. So the prediction will inevitably fail. In general, Gibbs sampling may fail in the presence of perfect correlations or anti-correlations since the resulting stochastic process is not ergodic; when relations are imperfect but very strong, Gibbs sampling may need a very long time to converge.

(Gasse et al., 2015) factorize labels into independent factors based on some statistical conditional independence tests and a Markov boundary learning algorithm. Its limitation, as pointed out by the authors, is that the statistical tests are ineffective when the ratio between samples and labels is not big enough. Their method outperforms the Power-Set baseline on synthetic datasets, but not on any real dataset. (Zhang & Zhang, 2010) propose to use a Bayesian Network to encode the conditional dependencies of the labels as well as the feature set, with the features as the common parent of all labels.

Conditional Random Fields (CRF) (Sutton & McCallum, 2006) offer a general framework for structured prediction problems based on undirected graphical models. In multi-label classifications labels can form densely connected graphs, so restrictions are imposed in order to make training and inference tractable. For problems involving only hard label relations (exclusive or hierarchical), a special CRF model is proposed (Deng et al., 2014); this works only when label dependencies are *strict* and *a priori* known.

pair-CRF limits the scope of potential functions to label pairs, and does not model higher order label interactions (Ghamrawi & McCallum, 2005). The exact inference and prediction in pair-wise CRF requires checking all possible label subsets, which is intractable for datasets with many labels. As shown in their experiment results, the most effective way of doing approximate inference and prediction is to consider only label subsets that occur in the training set. But this eliminates the possibility of predicting unseen subsets. Our CBM model does not have this limitation (see Section 4).

There are also algorithms which exploit label structures but do not estimate $p(\mathbf{y} | \mathbf{x})$ explicitly. Examples include Multi-label k-Nearest Neighbors (Zhang & Zhou, 2007), Compressed Sensing based method (Hsu et al., 2009), Principle Label Space Transformation (Tai & Lin, 2012), Conditional Principal Label Space Transformation (Chen & Lin, 2012) and Compressed Labeling (Zhou et al., 2012).

Various **mixture models** exist for multi-label document classification or supervised topical modeling (McCallum, 1999; Ueda & Saito, 2002; McAuliffe & Blei, 2008; Yang et al., 2009; Ramage et al., 2009; Puurula, 2011; Kim et al.,

2012). Our CBM model differs from these models in two aspects: 1) These models are generative in nature, i.e., they model $p(\mathbf{x})$, where \mathbf{x} is typically a bag-of-words representation of a document. By contrast, our CBM is purely discriminative, since it targets $p(\mathbf{y} | \mathbf{x})$ directly. The principled advantage of discriminative approach over generative approach, as stated in (Sutton & McCallum, 2011), is that the former does not make overly simplistic independence assumptions among features, and is thus better suited to including rich, overlapping features. 2) Most of these models are designed specifically for text data, while our method can be applied to any multi-label classification problem. A Conditional Multinomial Mixture model has been proposed for Superset Label Learning (Liu & Dietterich, 2012).

The architecture of CBM also mimics that of Mixture of Experts (ME) (Jacobs et al., 1991; Jordan & Jacobs, 1994), in which a gate model divides the input space into disjoint regions probabilistically and an expert model generates the output in each region. ME has been mainly used in regression and multi-class problems (Yuksel et al., 2012). CBM can be viewed as a multi-label extension of ME with a particular factorization of labels inside each expert.

6. Experiments

We perform experiments on five commonly used and relatively large multi-label datasets: SCENE, TMC2007, MEDIAMILL, NUS-WIDE from Mulan¹ and RCV1 (topics subset 1) from LIBSVM². For the sake of reproducibility, we adopt the train/test splits provided by Mulan and LIBSVM. Datasets details are shown at the top of Table 2.

We compare Conditional Bernoulli Mixtures (CBM) with the following methods which estimate $p(\mathbf{y} | \mathbf{x})$: Binary Relevance (BinRel), Power-Set (PowSet), Classifier Chains (CC), Probabilistic Classifier Chains with Beam Search (PCC), Ensemble of Classifier Chains with label voting (ECC-label) and subset voting (ECC-subset), Conditional Dependency Networks (CDN) and pair-wise Conditional Random Fields (pairCRF). CDN is taken from the package MEKA,³ the rest are based on our implementations.⁴ For all methods which require a base learner, we employ logistic regression as a common base learner. Additionally, we test gradient boosted trees (GB) as a non-linear base learner in conjunction with BinRel, PowSet and CBM. Both LR and pairCRF are L_2 regularized. Hyper parameter tuning is done by cross-validation on the training set

¹<http://mulan.sourceforge.net>

²<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/multilabel.html>

³<http://meka.sourceforge.net>

⁴Our implementations of CBM and several baselines (PowSet, PCC, CRF, etc.) are available at <https://github.com/cheng-li/pyramid>.

Table 2. Test subset accuracy of different methods on five datasets. All numbers are in percentages. Best performances are bolded.

dataset domain		SCENE image	RCV1 text	TMC2007 text	MEDIAMILL video	NUS-WIDE image
#labels / #label subsets	#features / #data points	6 / 15 294 / 2407	103 / 799 47K / 6000	22 / 1341 49K / 29K	101 / 6555 120 / 44K	81 / 18K 128 / 270K
Method	Learner					
BinRel	LR	51.5	40.4	25.3	9.6	24.7
PowSet	LR	68.1	50.2	28.2	9.0	26.6
CC	LR	62.9	48.2	26.2	10.9	26.0
PCC	LR	64.8	48.3	26.8	10.9	26.3
ECC-label	LR	60.6	46.5	26.0	11.3	26.0
ECC-subset	LR	63.1	49.2	25.9	11.5	26.0
CDN	LR	59.9	12.6	16.8	5.4	17.1
pairCRF	linear	68.8	46.4	28.1	10.3	26.4
CBM	LR	69.7	49.9	28.7	13.5	27.3
BinRel	GB	59.3	30.1	25.4	11.2	24.4
PowSet	GB	70.5	38.2	23.1	10.1	23.6
CBM	GB	70.5	43.0	27.5	14.1	26.5

(see the supplementary material for details). For methods involving random initializations or sampling, the reported results are averaged over 3 runs.

6.1. Results

Test subset accuracy on five datasets is shown in Table 2, grouped by the base learner. On four datasets, the highest subset accuracy is achieved by one of the CBM instantiations; on the other dataset, CBM is close to the best one. This demonstrates the effectiveness of CBM for optimizing subset accuracy. On the SCENE and MEDIAMILL datasets, CBM+GB performs better than CBM+LR, which shows the benefit of being able to incorporate non-linear components. By contrast, pairCRF is restricted to work with linear functions, lacking the flexibility of CBM. For completeness, we present in the supplementary material the Jaccard index and Hamming loss results. CBM also achieves the highest Jaccard index on three out of five datasets, which is reasonable given the fact that CBM only targets subset accuracy and multi-label optimality is measure-related. We also report all algorithms’ running times in the supplementary material. CBM requires significantly less training time than PowSet on datasets with a large number of possible label subsets.

6.2. Impact of Number of Components

Figure 3 shows how increasing the number of components K affects test performance for CBM on TMC dataset. When $K = 1$, CBM only estimates marginals and is equivalent to BinRel (the slight difference in performance is due to the implementation details of CBM; see the supplement-




Figure 3. Test subset accuracy on TMC dataset with varying number of components K for CBM+LR (left) and CBM+GB (right) compared with BinRel, PowSet, and pairCRF.

tary material). Increasing K from 1 to 15 makes CBM quickly become a better joint estimator and outperform BinRel. Increasing K further gives smaller improvement for CBM and the performance asymptotes as K reaches about 30. Other datasets show similar trends.

7. Conclusion

In this paper, we propose a new multi-label classification method using Conditional Bernoulli Mixtures. It captures label dependencies; it reduces a multi-label problem to binary and multi-class problems, and subsumes the classic independent binary prediction and power-set subset prediction methods as special cases. A simple Expectation Maximization training procedure is presented, together with an efficient prediction procedure based on dynamic programming. Experimental results show the effectiveness of the proposed method against competitive alternatives on benchmark multi-label datasets.

Acknowledgements

We thank reviewers for their helpful comments, Shivani Agarwal for her suggestions on related work, and Pavel Metrikov for discussions on the model. This work has been generously supported through a grant from the Massachusetts General Hospital Physicians Organization, MGH 507222.

References

- Bishop, Christopher M. *Pattern recognition and machine learning*. Springer, 2006.
- Chen, Yao-Nan and Lin, Hsuan-Tien. Feature-aware label space dimension reduction for multi-label classification. In *Advances in Neural Information Processing Systems*, pp. 1529–1537, 2012.
- Cheng, Weiwei, Hüllermeier, Eyke, and Dembczynski, Krzysztof J. Bayes optimal multilabel classification via probabilistic classifier chains. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pp. 279–286, 2010.
- Dembczyński, Krzysztof, Waegeman, Willem, and Hüllermeier, Eyke. Joint mode estimation in multi-label classification by chaining. In *CoLISD 2011 (Collective Learning and Inference on Structured Data 2011)*, 2011.
- Dembczyński, Krzysztof, Waegeman, Willem, Cheng, Weiwei, and Hüllermeier, Eyke. On label dependence and loss minimization in multi-label classification. *Machine Learning*, 88(1-2):5–45, 2012.
- Dembczynski, Krzysztof, Waegeman, Willem, and Hüllermeier, Eyke. An analysis of chaining in multi-label classification. In *ECAI*, pp. 294–299, 2012.
- Deng, Jia, Ding, Nan, Jia, Yangqing, Frome, Andrea, Murphy, Kevin, Bengio, Samy, Li, Yuan, Neven, Hartmut, and Adam, Hartwig. Large-scale object classification using label relation graphs. In *Computer Vision—ECCV 2014*, pp. 48–64. Springer, 2014.
- Friedman, Jerome H. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pp. 1189–1232, 2001.
- Gasse, Maxime, Aussem, Alexandre, and Elghazel, Haytham. On the optimality of multi-label classification under subset zero-one loss for distributions satisfying the composition property. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pp. 2531–2539, 2015.
- Ghamrawi, Nadia and McCallum, Andrew. Collective multi-label classification. In *Proceedings of the 14th ACM international conference on Information and knowledge management*, pp. 195–200. ACM, 2005.
- Guo, Yuhong and Gu, Suicheng. Multi-label classification using conditional dependency networks. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, pp. 1300, 2011.
- Gupta, Maya R and Chen, Yihua. *Theory and use of the EM algorithm*. Now Publishers Inc, 2011.
- Hsu, Daniel, Kakade, Sham, Langford, John, and Zhang, Tong. Multi-label prediction via compressed sensing. In *NIPS*, volume 22, pp. 772–780, 2009.
- Jacobs, Robert A, Jordan, Michael I, Nowlan, Steven J, and Hinton, Geoffrey E. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- Jordan, Michael I and Jacobs, Robert A. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994.
- Kim, Dongwoo, Kim, Suin, and Oh, Alice. Dirichlet process with mixed random measures: a nonparametric topic model for labeled data. *arXiv preprint arXiv:1206.4658*, 2012.
- Koyejo, Oluwasanmi O, Natarajan, Nagarajan, Ravikumar, Pradeep K, and Dhillon, Inderjit S. Consistent multilabel classification. In *Advances in Neural Information Processing Systems*, pp. 3303–3311, 2015.
- Kumar, Abhishek, Vembu, Shankar, Menon, Aditya Krishna, and Elkan, Charles. Learning and inference in probabilistic classifier chains with beam search. In *Machine Learning and Knowledge Discovery in Databases*, pp. 665–680. Springer, 2012.
- Kumar, Abhishek, Vembu, Shankar, Menon, Aditya Krishna, and Elkan, Charles. Beam search algorithms for multilabel learning. *Machine learning*, 92(1):65–89, 2013.
- Lazarsfeld, Paul Felix, Henry, Neil W, and Anderson, Theodore Wilbur. *Latent structure analysis*. Houghton Mifflin Boston, 1968.
- Liu, Liping and Dietterich, Thomas G. A conditional multinomial mixture model for superset label learning. In *Advances in neural information processing systems*, pp. 557–565, 2012.
- Liu, Weiwei and Tsang, Ivor. On the optimality of classifier chain for multi-label classification. In *Advances in Neural Information Processing Systems*, pp. 712–720, 2015.

- McAuliffe, Jon D and Blei, David M. Supervised topic models. In *Advances in neural information processing systems*, pp. 121–128, 2008.
- McCallum, Andrew. Multi-label text classification with a mixture model trained by EM. In *AAAI99 workshop on text learning*, pp. 1–7, 1999.
- McLachlan, Geoffrey and Peel, David. *Finite mixture models*. John Wiley & Sons, 2004.
- Mena, Deiner, Montañés, Elena, Quevedo, José R, and Del Coz, Juan José. Using A* for inference in probabilistic classifier chains. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pp. 3707–3713. AAAI Press, 2015.
- Nocedal, Jorge and Wright, Stephen. *Numerical optimization*. Springer Science & Business Media, 2006.
- Puurula, Antti. Mixture models for multi-label text classification. In *10th New Zealand Computer Science Research Student Conference*, 2011.
- Ramage, Daniel, Hall, David, Nallapati, Ramesh, and Manning, Christopher D. Labeled lda: A supervised topic model for credit attribution in multi-labeled corpora. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1*, pp. 248–256. Association for Computational Linguistics, 2009.
- Read, Jesse, Pfahringer, Bernhard, Holmes, Geoff, and Frank, Eibe. Classifier chains for multi-label classification. *Machine learning*, 85(3):333–359, 2011.
- Schapire, Robert E and Singer, Yoram. Boostexter: A boosting-based system for text categorization. *Machine learning*, 39(2):135–168, 2000.
- Sutton, Charles and McCallum, Andrew. An introduction to conditional random fields for relational learning. *Introduction to statistical relational learning*, pp. 93–128, 2006.
- Sutton, Charles and McCallum, Andrew. An introduction to conditional random fields. *Machine Learning*, 4(4): 267–373, 2011.
- Tai, Farbound and Lin, Hsuan-Tien. Multilabel classification with principal label space transformation. *Neural Computation*, 24(9):2508–2542, 2012.
- Tsoumakas, Grigorios and Katakis, Ioannis. Multi-label classification: An overview. *Int J Data Warehousing and Mining*, 2007:1–13, 2007.
- Ueda, Naonori and Saito, Kazumi. Parametric mixture models for multi-labeled text. In *Advances in neural information processing systems*, pp. 721–728, 2002.
- Yang, Shuang-Hong, Zha, Hongyuan, and Hu, Bao-Gang. Dirichlet-bernoulli alignment: A generative model for multi-class multi-label multi-instance corpora. In *Advances in neural information processing systems*, pp. 2143–2150, 2009.
- Yuksel, Seniha Esen, Wilson, Joseph N, and Gader, Paul D. Twenty years of mixture of experts. *Neural Networks and Learning Systems, IEEE Transactions on*, 23(8): 1177–1193, 2012.
- Zhang, Min-Ling and Zhang, Kun. Multi-label learning by exploiting label dependency. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 999–1008. ACM, 2010.
- Zhang, Min-Ling and Zhou, Zhi-Hua. Ml-knn: A lazy learning approach to multi-label learning. *Pattern recognition*, 40(7):2038–2048, 2007.
- Zhou, Tianyi, Tao, Dacheng, and Wu, Xindong. Compressed labeling on distilled labelsets for multi-label learning. *Machine Learning*, 88(1-2):69–126, 2012.