

Multi-dimensional Classification via Selective Feature Augmentation

Bin-Bin Jia^{1,2,3} Min-Ling Zhang^{1,3}

¹School of Computer Science and Engineering, Southeast University, Nanjing 210096, China

²College of Electrical and Information Engineering, Lanzhou University of Technology, Lanzhou 730050, China

³Key Laboratory of Computer Network and Information Integration, Southeast University, Ministry of Education, China

Abstract: In multi-dimensional classification (MDC), the semantics of objects are characterized by multiple class spaces from different dimensions. Most MDC approaches try to explicitly model the dependencies among class spaces in output space. In contrast, the recently proposed feature augmentation strategy, which aims at manipulating feature space, has also been shown to be an effective solution for MDC. However, existing feature augmentation approaches only focus on designing holistic augmented features to be appended with the original features, while better generalization performance could be achieved by exploiting multiple kinds of augmented features. In this paper, we propose the selective feature augmentation strategy that focuses on synergizing multiple kinds of augmented features. Specifically, by assuming that only part of the augmented features is pertinent and useful for each dimension's model induction, we derive a classification model which can fully utilize the original features while conduct feature selection for the augmented features. To validate the effectiveness of the proposed strategy, we generate three kinds of simple augmented features based on standard k NN, weighted k NN, and maximum margin techniques, respectively. Comparative studies show that the proposed strategy achieves superior performance against both state-of-the-art MDC approaches and its degenerated versions with either kind of augmented features.

Keywords: Machine learning, multi-dimensional classification, feature augmentation, feature selection, class dependencies.

Citation: B. B. Jia, M. L. Zhang. Multi-dimensional classification via selective feature augmentation. *Machine Intelligence Research*, vol.19, no.1, pp.38–51, 2022. <http://doi.org/10.1007/s11633-022-1316-5>

1 Introduction

Traditional supervised learning tasks usually characterize the semantics of objects with one output variable, i.e., single-output learning, among which multi-class classification is one of the most important learning frameworks. However, in some real-world applications, it is better to use multiple output variables to characterize the rich semantics of objects, which results in the problem of multi-output learning^[1]. Here, when the type of each output variable is restricted to discrete-valued, the multi-dimensional classification (MDC) framework is obtained^[2, 3]. Under the MDC setting, each object is represented by a single instance while associated with multiple class variables, each corresponding to a specific class space characterizing the object's semantics along one specific dimension. Specifically, the MDC problem widely exists in many application scenarios, such as bioinformatics^[4, 5], text classification^[6, 7], computer vision^[8–10], re-

source allocation^[11], etc. Fig.1 shows an illustrative example of MDC on vehicle classification.

Formally speaking, let $\mathcal{X} \in \mathbf{R}^d$ be the d -dimensional feature space and $\mathcal{Y} = C_1 \times C_2 \times \cdots \times C_q$ be the output space. Here, \mathcal{Y} corresponds to the Cartesian product of q class spaces $C_j = \{c_1^j, c_2^j, \cdots, c_{K_j}^j\}$ ($1 \leq j \leq q$) which consists of K_j possible classes respectively. Given a set of MDC training examples $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i) \mid 1 \leq i \leq m\}$, where $\mathbf{x}_i = [x_{i1}, x_{i2}, \cdots, x_{id}]^T \in \mathcal{X}$ is a d -dimensional feature vector and $\mathbf{y}_i = [y_{i1}, y_{i2}, \cdots, y_{iq}]^T \in \mathcal{Y}$ is the q -dimensional class vector associated with \mathbf{x}_i with each element $y_{ij} \in C_j$, the MDC task aims to learn a predictive model $f: \mathcal{X} \mapsto \mathcal{Y}$ from \mathcal{D} which can return a proper class vector $f(\mathbf{x}_*) \in \mathcal{Y}$ for unseen instance \mathbf{x}_* .

It is obvious that the MDC problem can be solved dimension by dimension, i.e., training a multi-class classifier for each class space. However, this independent decomposition strategy does not consider potential dependencies among class spaces which might impact the generalization performance of the resulting model. The MDC problem can also be solved by a single multi-class classifier, where each distinct class combination is regarded as a new class. However, this powerset-like strategy cannot consider class combinations not appearing in the training

Research Article

Manuscript received April 17, 2021; accepted September 8, 2021

Recommended by Associate Editor Hui Yu

Colored figures are available in the online version at <https://link.springer.com/journal/11633>

© The Author(s) 2022



Fig. 1 An illustrative example of multi-dimensional classification: Vehicle classification. For a vehicle, it can be classified from the type dimension (with possible classes car, SUV, bus, truck, etc.), from the brand dimension (with possible classes Audi, Benz, YUTONG, JAC, etc.), and from the color dimension (with possible classes black, white, red, blue, etc.). Here, (a) is a red Audi car, (b) is a black Benz SUV, (c) is a red YUTONG bus, and (d) is a white JAC truck.

set and usually suffers high computational complexity due to a possibly large number of classes. In fact, one of the key challenges for MDC studies is how to model dependencies among class spaces in appropriate ways. Existing works mainly focus on modeling class dependencies in output space, such as capturing pairwise class dependencies^[12–14], specifying chaining order over class spaces^[15, 16], learning a directed acyclic graph (DAG) structure for class spaces^[17–19], and partitioning class spaces into groups^[2], etc. Recently, feature augmentation strategy, which aims at manipulating feature space, has been shown as an effective solution for MDC. This strategy enriches the original feature space with a set of new features that are generated by making use of some well-established techniques, e.g., k NN^[20] or deep learning^[21]. Existing works only focus on how to design more informative augmented features, while it might be beneficial to exploit multiple kinds of augmented features generated by making use of different techniques. In this paper, we propose the selective feature augmentation strategy, which makes the first attempt to synergize multiple kinds of augmented features. The strategy is abbreviated as SFAM, i.e., selective feature augmentation for multi-dimensional classification, in the following parts of this paper for brevity. Specifically, SFAM assumes only part of augmented features are pertinent and useful for each dimension's model induction. In order to validate the effectiveness of SFAM, three simple kinds of augmented features are generated by making use of standard k NN, weighted k NN, and maximum margin techniques, respectively. After that, for each dimension, SFAM derives a classification model which can take full advantage of the original features via ℓ_2 regularization and conduct feature selection for the augmented features via ℓ_1 regularization (i.e., selective feature augmentation). Experimental

results demonstrate that SFAM achieves superior performance against both state-of-the-art MDC approaches and its degenerated versions with either kind of augmented features.

The rest of this paper is organized as follows. Firstly, related works on multi-dimensional classification are briefly discussed. Secondly, technical details of SFAM are introduced. Thirdly, experimental results of comparative studies are reported. Finally, we conclude this paper.

2 Related work

The most related learning framework to multi-dimensional classification is the widely studied multi-label classification (MLC)^[22–24], which can be regarded as a special case of MDC when the type of class variable in each dimension is restricted to binary-valued. However, MDC usually assumes heterogeneous class spaces, which are used to characterize the rich semantics of objects from different dimensions, while MLC usually assumes homogeneous class space in which multiple concepts are relevant to the polysemous objects.

The MDC problem can be solved via an independent decomposition strategy, where a total of q multi-class classifiers are learned independently, one per dimension. However, this intuitive strategy ignores possible dependencies among class spaces, and the induced model would be suboptimal. An improved strategy is learning the q multi-class classifiers in a chaining order, where predictions of preceding classifiers are used as extra features by the subsequent ones^[15, 16]. However, the chaining order would largely affect the generalization performance while determining an optimal one is NP-hard. The MDC problem can also be solved via a powerset transformation strategy where a single multi-class classifier is learned by regarding all distinct class combinations in the training set as new classes. However, this intuitive strategy cannot consider class combinations not appearing in the training set and usually suffers high computational complexity due to a large number of new classes. An improved strategy is partitioning the class spaces into groups according to conditional dependencies^[2]. However, the combinatorial nature still exists, which leads to that the deficiencies cannot be fully addressed. A family of MDC models called multi-dimensional Bayesian network classifier^[25] aims at learning different kinds of DAG structures over class spaces to explicitly model the class dependencies. However, determining DAG structures is computationally demanding, and only nominal features can be tackled generally. The class dependencies can also be modeled in a two-level strategy^[12–14], where pairwise dependencies are captured in the first level, and then high-order dependencies are captured in the second-level based on the predictions from the first level. However, capturing pairwise dependencies needs $O(q^2)$ complexity, which is very time-consuming.

The afore-mentioned strategies mainly focus on directly modeling class dependencies in the output space, while the KRAM approach^[20] attempts to manipulate the feature space of MDC examples via feature augmentation by making use of k NN techniques. Helpful discriminative information is expected to be brought into feature space which would facilitate the subsequent MDC model induction. Based on deep learning techniques, the LEFA approach^[21] further generates better-augmented features which can depict the inter-class dependencies and the intra-class exclusiveness simultaneously. However, these approaches simply treat the original and augmented features equally, which might be less reasonable due to various characteristics of different features. Moreover, it is usually easier to design multiple kinds of simple augmented features than a terrific one, and it might be beneficial to consider synergizing the discriminative information residing in different kinds of augmented features. Fig. 2 shows an intuitive comparison between existing feature augmentation techniques and the proposed one in this paper. Existing works usually employ general MDC algorithms to accomplish the training phase, while the proposed one designs a novel training algorithm that can accomplish the selective feature augmentation phase.

3 Technical details of SFAM

This section presents how we implement the selective

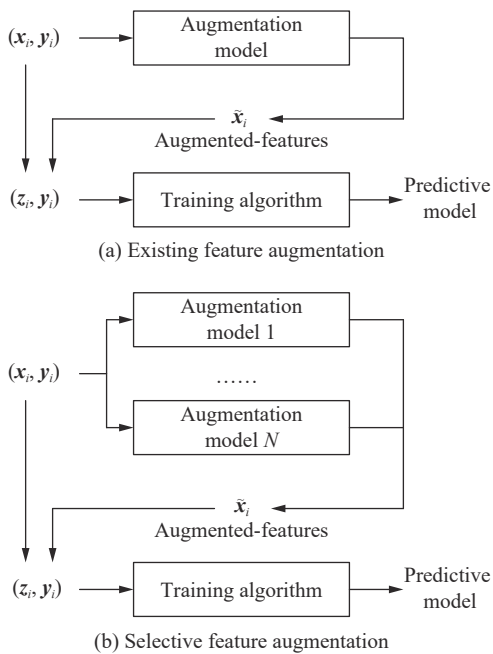


Fig. 2 An intuitive comparison of the training phase between existing feature augmentation techniques and the proposed one in this paper. Here, z_i is the concatenation of x_i and \tilde{x}_i . Specifically, existing feature augmentation techniques focus on designing a better augmentation model to generate more informative augmented features, while this paper focuses on designing a novel training algorithm that aims at synergizing multiple kinds of augmented features.

feature augmentation strategy. In other words, the technical details in this section correspond to the part of the training algorithm in Fig. 2(b). For any instance x_i , let \tilde{x}_i be the concatenation of all the corresponding augmented features generated by the N augmentation models, and z_i corresponds to the concatenation of x_i and \tilde{x}_i , i.e., $z_i = [x_i; \tilde{x}_i]$, in the following, we derive a regularized classification model which fully utilizes original features x_i and employs feature selection mechanism over augmented features \tilde{x}_i .

For simplicity, we employ the one-vs-rest decomposition strategy for each dimension where a classification model with both ℓ_1 and ℓ_2 regularization is derived to solve the decomposed binary classification problems. Specifically, for the a -th decomposed binary classification problem in the j -th dimension, we determine the optimal model (w_{ja}^*, b_{ja}^*) as follows:

$$(w_{ja}^*, b_{ja}^*) = \arg \min_{w_{ja}, b_{ja}} L(w_{ja}, b_{ja}) + \lambda R(w_{ja}) \quad (1)$$

where λ is a trade-off parameter. The first term $L(w_{ja}, b_{ja})$ denotes the empirical loss function. In this paper, we employ the cross-entropy loss, which is defined as follows:

$$L(w_{ja}, b_{ja}) = - \sum_{i=1}^m \left[l_i^{ja} \times \ln h_{w_{ja}, b_{ja}}(z_i) + (1 - l_i^{ja}) \times \ln(1 - h_{w_{ja}, b_{ja}}(z_i)) \right].$$

Here, $l_i^{ja} = \mathbf{1}_{y_{ij}=c_a^j}$, the predicate $\mathbf{1}_\pi$ returns 1 if π holds and 0 otherwise, and $h_{w_{ja}, b_{ja}}(z_i)$ is the logistic function which is defined as follows:

$$h_{w_{ja}, b_{ja}}(z_i) = \frac{1}{1 + e^{-\langle w_{ja}, z_i \rangle + b_{ja}}}$$

where $\langle \cdot, \cdot \rangle$ returns the inner product of two vectors. The second term $R(w_{ja})$ denotes the regularization term, which is defined as follows:

$$R(w_{ja}) = \|\theta_{ja}\|_2^2 + \|\tilde{\theta}_{ja}\|_1$$

where θ_{ja} and $\tilde{\theta}_{ja}$ denote the first d elements and the last remaining elements of w_{ja} , respectively, i.e., $w_{ja} = [\theta_{ja}; \tilde{\theta}_{ja}]$. It is worth noting that the ℓ_2 regularization corresponds to the original features while the ℓ_1 regularization corresponds to the augmented features. By doing this, we employ a feature selection mechanism over the augmented features to synergize multiple kinds of augmented features in a better way while the original features are still fully utilized. In other words, the selective feature augmentation strategy is implemented here.

To optimize problem (1), we solve one of the three

sets of parameters $\{\theta_{ja}\}, \{\tilde{\theta}_{ja}\}$ and $\{b_{ja}\}$ alternately, while the remaining parameters are fixed.

1) Optimizing w.r.t. $\{\theta_{ja}\}$ when $\{\tilde{\theta}_{ja}\}$ and $\{b_{ja}\}$ are fixed: When $\{\tilde{\theta}_{ja}\}$ and $\{b_{ja}\}$ are fixed, the optimization problem (1) can be equivalently reformulated as follows:

$$\min_{\theta_{ja}} L_1(\theta_{ja}) + \lambda \|\theta_{ja}\|_2^2 \quad (2)$$

where

$$L_1(\theta_{ja}) = - \sum_{i=1}^m \left[l_i^{ja} \times \ln h_{\theta_{ja}}(z_i) + (1 - l_i^{ja}) \times \ln(1 - h_{\theta_{ja}}(z_i)) \right]$$

and

$$h_{\theta_{ja}}(z_i) = \frac{1}{1 + e^{-(\langle \theta_{ja}, x_i \rangle + C_1^i)}}.$$

Here, $C_1^i = \langle \tilde{\theta}_{ja}, \tilde{x}_i \rangle + b_{ja}$ is a constant which is not dependent on the variables θ_{ja} . In this paper, we use gradient descent to solve the optimization problem (2). Specifically, let $\Omega_1(\theta_{ja})$ be the objective function, the gradient is given as follows:

$$\frac{\partial \Omega_1(\theta_{ja})}{\partial \theta_{ja}} = - \sum_{i=1}^m (l_i^{ja} - h_{\theta_{ja}}(z_i)) x_i + 2\lambda \theta_{ja}.$$

2) Optimizing w.r.t. $\{\tilde{\theta}_{ja}\}$ when $\{\theta_{ja}\}$ and $\{b_{ja}\}$ are fixed: When $\{\theta_{ja}\}$ and $\{b_{ja}\}$ are fixed, the optimization problem (1) can be equivalently reformulated as follows:

$$\min_{\tilde{\theta}_{ja}} L_2(\tilde{\theta}_{ja}) + \lambda \|\tilde{\theta}_{ja}\|_1 \quad (3)$$

where

$$L_2(\tilde{\theta}_{ja}) = - \sum_{i=1}^m \left[l_i^{ja} \times \ln h_{\tilde{\theta}_{ja}}(z_i) + (1 - l_i^{ja}) \times \ln(1 - h_{\tilde{\theta}_{ja}}(z_i)) \right]$$

and

$$h_{\tilde{\theta}_{ja}}(z_i) = \frac{1}{1 + e^{-(\langle \tilde{\theta}_{ja}, \tilde{x}_i \rangle + C_2^i)}}.$$

Here, $C_2^i = \langle \theta_{ja}, x_i \rangle + b_{ja}$ is a constant which is not dependent on the variables $\{\tilde{\theta}_{ja}\}$. In this paper, we use the accelerated proximal gradient method^[26] to solve it.

Theorem 1. For the derivable function $L_2(\tilde{\theta}_{ja})$, $\nabla L_2(\tilde{\theta}_{ja})$ is Lipschitz continuous, and the Lipschitz constant is

$$L_f = \frac{1}{4} \sum_{i=1}^m \|\tilde{x}_i\|_2^2 \quad (4)$$

where ∇ denotes the differential operator.

Proof. For $\nabla L_2(\tilde{\theta}_{ja})$, it can be calculated as

$$\nabla L_2(\tilde{\theta}_{ja}) = \frac{\partial L_2(\tilde{\theta}_{ja})}{\partial \tilde{\theta}_{ja}} = - \sum_{i=1}^m (l_i^{ja} - h_{\tilde{\theta}_{ja}}(z_i)) \tilde{x}_i.$$

Given any $\tilde{\theta}_{ja}$ and $\tilde{\theta}_{ja}'$, we have

$$\begin{aligned} \|\nabla L_2(\tilde{\theta}_{ja}') - \nabla L_2(\tilde{\theta}_{ja})\|_2 &= \left\| \sum_{i=1}^m (h_{\tilde{\theta}_{ja}'}(z_i) - h_{\tilde{\theta}_{ja}}(z_i)) \tilde{x}_i \right\|_2 \leq \\ &\sum_{i=1}^m \|h_{\tilde{\theta}_{ja}'}(z_i) - h_{\tilde{\theta}_{ja}}(z_i)\|_2 \|\tilde{x}_i\|_2 \leq \\ &\sum_{i=1}^m \frac{1}{4} \|\langle \tilde{\theta}_{ja}', \tilde{x}_i \rangle - \langle \tilde{\theta}_{ja}, \tilde{x}_i \rangle\|_2 \|\tilde{x}_i\|_2 \leq \\ &\sum_{i=1}^m \frac{1}{4} \|\langle \tilde{\theta}_{ja}' - \tilde{\theta}_{ja}, \tilde{x}_i \rangle\|_2 \|\tilde{x}_i\|_2 \leq \\ &\frac{1}{4} \sum_{i=1}^m \|\tilde{\theta}_{ja}' - \tilde{\theta}_{ja}\|_2 \|\tilde{x}_i\|_2^2. \end{aligned}$$

Here, the second “ \leq ” is due to the truth that the Lipschitz constant of logistic function equals $\frac{1}{4}$. So then, it is easy to know:

$$\frac{\|\nabla L_2(\tilde{\theta}_{ja}') - \nabla L_2(\tilde{\theta}_{ja})\|_2}{\|\tilde{\theta}_{ja}' - \tilde{\theta}_{ja}\|_2} \leq \frac{1}{4} \sum_{i=1}^m \|\tilde{x}_i\|_2^2$$

which completes the proof. \square

According to Theorem 1, given any initial value $\tilde{\theta}_{ja}^{(t)}$ of $\tilde{\theta}_{ja}$, let $\Delta \tilde{\theta}_{ja} = \tilde{\theta}_{ja} - \tilde{\theta}_{ja}^{(t)}$, the following inequation always holds:

$$\|\nabla L_2(\tilde{\theta}_{ja}) - \nabla L_2(\tilde{\theta}_{ja}^{(t)})\|_2 \leq L_f \|\Delta \tilde{\theta}_{ja}\|_2.$$

Then, the quadratic approximation of $L_2(\tilde{\theta}_{ja})$ around $\tilde{\theta}_{ja}^{(t)}$ can be given as follows:

$$\begin{aligned} \hat{L}_2(\tilde{\theta}_{ja}) &\simeq L_2(\tilde{\theta}_{ja}^{(t)}) + \langle \nabla L_2(\tilde{\theta}_{ja}^{(t)}), \Delta \tilde{\theta}_{ja} \rangle + \frac{L_f}{2} \|\Delta \tilde{\theta}_{ja}\|_2^2 = \\ &\frac{L_f}{2} \|\tilde{\theta}_{ja} - \mathbf{u}^{(t)}\|_2^2 + C_{L_f} \end{aligned}$$

where $C_{L_f} = -\frac{1}{2L_f} \|\nabla L_2(\tilde{\theta}_{ja}^{(t)})\|_2^2 + L_2(\tilde{\theta}_{ja}^{(t)})$ is a constant which is not dependent on variables $\tilde{\theta}_{ja}$, and

$$\mathbf{u}^{(t)} = \tilde{\theta}_{ja}^{(t)} - \frac{1}{L_f} \nabla L_2(\tilde{\theta}_{ja}^{(t)}). \quad (5)$$

According to the descent lemma^[27], the approximation is an upper bound of the original function, i.e., $L_2(\tilde{\theta}_{ja}) \leq \hat{L}_2(\theta_{ja})$ always holds. Therefore, we can minimize the original function by iteratively minimizing the approximation. Plugging the above approximation into the optimization problem (3), we can obtain the following iterative equation:

$$\tilde{\theta}_{ja}^{(t+1)} = \arg \min_{\tilde{\theta}_{ja}} \frac{L_f}{2} \left\| \tilde{\theta}_{ja} - \mathbf{u}^{(t)} \right\|_2^2 + \lambda \left\| \tilde{\theta}_{ja} \right\|_1 = \text{soft} \left(\mathbf{u}^{(t)}, \frac{\lambda}{L_f} \right). \quad (6)$$

Here, $\text{soft}(\cdot, \cdot)$ is the (element-wise) soft-thresholding function which is defined as follows:

$$\text{soft}(x, \mu) = \begin{cases} x - \mu, & \text{if } x > \mu \\ x + \mu, & \text{if } x < -\mu \\ 0, & \text{otherwise.} \end{cases}$$

In [28], it is shown that the convergence rate of the iterative equation in (6) can be improved to $O(t^{-2})$ from $O(t^{-1})$ if we replace $\tilde{\theta}_{ja}^{(t)}$ in (5) with the following $\mathbf{v}_{ja}^{(t)}$:

$$\mathbf{v}_{ja}^{(t)} = \tilde{\theta}_{ja}^{(t)} + \frac{r_{t-1} - 1}{r_t} (\tilde{\theta}_{ja}^{(t)} - \tilde{\theta}_{ja}^{(t-1)}) \quad (7)$$

where $r_0 = r_1 = 1$ and $r_t = \frac{1 + \sqrt{1 + 4r_{t-1}^2}}{2}$ when $t > 1$.

3) Optimizing w.r.t. $\{b_{ja}\}$ when $\{\theta_{ja}\}$ and $\{\tilde{\theta}_{ja}\}$ are fixed: When $\{\theta_{ja}\}$ and $\{\tilde{\theta}_{ja}\}$ are fixed, the optimization problem (1) can be equivalently reformulated as follows:

$$\min_{b_{ja}} L_3(b_{ja}) \quad (8)$$

where

$$L_3(b_{ja}) = - \sum_{i=1}^m \left[l_i^{ja} \times \ln h_{b_{ja}}(\mathbf{z}_i) + (1 - l_i^{ja}) \times \ln(1 - h_{b_{ja}}(\mathbf{z}_i)) \right]$$

and

$$h_{b_{ja}}(\mathbf{z}_i) = \frac{1}{1 + e^{-(b_{ja} + C_3^i)}}.$$

Here, $C_3^i = \langle \mathbf{w}_{ja}, \mathbf{z}_i \rangle$ is a constant which is not dependent on variable $\{b_{ja}\}$. In this paper, we use gradient descent to solve it. Specifically, the gradient of the objective function (i.e., $L_3(b_{ja})$) is given as follows:

$$\frac{\partial L_3(b_{ja})}{\partial b_{ja}} = - \sum_{i=1}^m (l_i^{ja} - h_{b_{ja}}(\mathbf{z}_i)).$$

Algorithm 1. The proposed SFAM approach.

Input: \mathcal{D} : MDC training set $\{(\mathbf{x}_i, \mathbf{y}_i) \mid 1 \leq i \leq m\}$

λ : trade-off parameter in (1)

\mathcal{A}_n : augmentation model n ($1 \leq n \leq N$)

\mathbf{x}_* : unseen instance

Output: \mathbf{y}_* : predicted class vector for \mathbf{x}_*

- 1) $\tilde{\mathcal{D}} = \emptyset$;
- 2) **for** $i = 1$ to m **do**
- 3) **for** $n = 1$ to N **do**
- 4) Generate the n -th augmented features $\Delta_n^{\mathbf{x}_i}$ for the i -th training example \mathbf{x}_i via \mathcal{A}_n ;
- 5) **end for**
- 6) $\tilde{\mathcal{D}} = \tilde{\mathcal{D}} \cup \{(\mathbf{z}_i, \mathbf{y}_i) \mid \mathbf{z}_i = [\mathbf{x}_i; \tilde{\mathbf{x}}_i], \tilde{\mathbf{x}}_i \text{ is the concatenation of } \Delta_1^{\mathbf{x}_i}, \dots, \Delta_N^{\mathbf{x}_i}\}$;
- 7) **end for**
- 8) **for** $j = 1$ to q **do**
- 9) **for** $a = 1$ to K_j **do**
- 10) Initialize $\mathbf{w}_{ja} = [\theta_{ja}; \tilde{\theta}_{ja}] = \mathbf{0}, b_{ja} = 0$;
- 11) **Repeat**
- 12) Update θ_{ja} by solving the optimization problem (2) via gradient descent;
- 13) Initialize $\tilde{\theta}_{ja}^{(0)} = \tilde{\theta}_{ja}^{(1)} = \tilde{\theta}_{ja}, r_0 = r_1 = 1, t = 1$;
- 14) **Repeat**
- 15) Obtain $\mathbf{v}_{ja}^{(t)}$ according to (7);
- 16) Compute $\mathbf{u}^{(t)} = \mathbf{v}_{ja}^{(t)} - \frac{1}{L_f} \nabla L_2(\mathbf{v}_{ja}^{(t)})$;
- 17) Obtain $\tilde{\theta}_{ja}^{(t+1)}$ according to (6);
- 18) Compute $r_{t+1} = \frac{1 + \sqrt{1 + 4r_t^2}}{2}$;
- 19) $t = t + 1$;
- 20) **Until** Convergence
- 21) Update $\tilde{\theta}_{ja}$ with $\tilde{\theta}_{ja}^{(t)}$
- 22) Update b_{ja} by solving the optimization problem (8) via gradient descent;
- 23) **Until** Convergence
- 24) **end for**
- 25) **end for**
- 26) Obtain \mathbf{x}'_* 's augmented features $\tilde{\mathbf{x}}_* = [\Delta_1^{\mathbf{x}_*}; \dots; \Delta_N^{\mathbf{x}_*}]$
- 27) **for** $j = 1$ to q **do**
- 28) Determine the class y_{*j} according to (9);
- 29) **end for**
- 30) Return $\mathbf{y}_* = [y_{*1}, y_{*2}, \dots, y_{*q}]^T$.

As the above three alternating optimizing steps converge, we can obtain the optimal values of \mathbf{w}_{ja} and b_{ja} . For unseen instance \mathbf{x}_* , let $\tilde{\mathbf{x}}_*$ be its augmented features, then its class label in the j -th dimension can be determined based on the augmented instance $\mathbf{z}_* = [\mathbf{x}_*; \tilde{\mathbf{x}}_*]$ as follows:

$$y_{*j} = c_{\hat{a}}^j, \text{ where } \hat{a} = \arg \max_{1 \leq a \leq K_j} \langle \mathbf{w}_{ja}, \mathbf{z}_* \rangle + b_{ja}. \quad (9)$$

The complete procedure of the proposed SFAM approach is summarized in Algorithm 1. Firstly, SFAM transforms the original MDC training set \mathcal{D} into $\tilde{\mathcal{D}}$ by augmenting each instance's feature space (Steps 1)–7)).

After that, the predictive model is induced via a classification model with both ℓ_2 and ℓ_1 regularization (Steps 8)–25)), where the ℓ_2 regularization and the bias term are updated via gradient descent, and the ℓ_1 regularization is updated via an accelerated proximal gradient method. Finally, the class vector of unseen instance is predicted based on the augmented features as well (Steps 26)–30)). As shown in Algorithm 1, it is worth noting that SFAM should be regarded as a general framework and can be coupled with any kind of augmented-features, while this paper only aims at investigating the feasibility of synergizing different kinds of augmented features.

4 Experiments

This section conducts comparative studies, and the obtained experimental results clearly validate the superiority and effectiveness of SFAM. Firstly, Section 4.1 introduces the experimental setup, including the employed benchmark data sets, the evaluation metrics, and the compared approaches. Then, Section 4.2 reports the detailed experimental results with statistical tests. Finally, Section 4.3 further investigates SFAM's algorithmic design and parameter sensitivity.

4.1 Experimental setup

1) Benchmark data sets

In this paper, a total of ten benchmark data sets are collected for comparative studies. Table 1 summarizes the detailed characteristics of all benchmark data sets, including number of examples (#Exam.), number of dimensions (#Dim.), number of class labels per dimension (#Labels/Dim.), and number of features (#Features). Here, for the “#Labels/Dim.” column, if all dimensions contain the same number of class labels, then only this number is recorded. Otherwise, the number of class labels per dimension is recorded in turn. For the “#Features” column, n and x denote numeric and nominal features, respectively.

Table 1 Characteristics of benchmark data sets

Data Set	#Exam.	#Dim.	#Labels/Dim.	#Features
Flare1	323	3	3,4,2	10x
Enb	768	2	2,4	6n
WQplants	1 060	7	4	16n
WQanimals	1 060	7	4	16n
WaterQuality	1 060	14	4	16n
BeLaE	1 930	5	5	1n, 44x
Voice	3 136	2	4,2	19n
TIC2000	9 822	3	6,4,2	83x
Adult	18 419	4	7,7,5,2	5n, 5x
Default	28 779	4	2,7,4,2	14n, 6x

2) Evaluation metrics

In this paper, a total of three widely-used evaluation metrics are employed for performance evaluation, including Hamming score (HS), exact match (EM), and sub-exact match (SEM) [2, 13, 14, 20, 29]. Specifically, given the test set $\mathcal{S} = \{(\mathbf{x}_i, \mathbf{y}_i) \mid 1 \leq i \leq p\}$, for the MDC model $f: \mathcal{X} \mapsto \mathcal{Y}$ to be evaluated, let $\hat{\mathbf{y}}_i = f(\mathbf{x}_i) = [\hat{y}_{i1}, \hat{y}_{i2}, \dots, \hat{y}_{iq}]^T$ be the predicted class vector for \mathbf{x}_i while the ground-truth one is $\mathbf{y}_i = [y_{i1}, y_{i2}, \dots, y_{iq}]^T$, then the number of correctly predicted dimensions corresponds to $r^{(i)} = \sum_{j=1}^q \mathbf{1}_{y_{ij}=\hat{y}_{ij}}$. The detailed definitions of the metrics are given as follows:

$$\begin{aligned} \text{HS}_{\mathcal{S}}(f) &= \frac{1}{p} \sum_{i=1}^p \frac{1}{q} \times r^{(i)} \\ \text{EM}_{\mathcal{S}}(f) &= \frac{1}{p} \sum_{i=1}^p \mathbf{1}_{r^{(i)}=q} \\ \text{SEM}_{\mathcal{S}}(f) &= \frac{1}{p} \sum_{i=1}^p \mathbf{1}_{r^{(i)} \geq q-1}. \end{aligned}$$

For the three metrics, it is easy to know that the larger the values, the better the performance. Ten-fold cross-validation is conducted over each benchmark data set, where both the mean metric value as well as standard deviation are recorded for comparative studies.

3) Compared approaches

In this paper, a total of six state-of-the-art MDC approaches are employed as compared approaches, including BR, CP, ECC [16], ESC [2], gMML [29], and SEEM [13]:

i) BR solves the MDC problem via training a number of multi-class classifiers independently, one per dimension. BR serves as the baseline when all possible class dependencies are ignored.

ii) CP transforms the MDC problem into a single multi-class classification problem by treating the whole output space as a compound one. Each distinct class combination in the training set is regarded as a new class. CP serves as the baseline when all possible class dependencies in the training set are considered, but overfitting might occur because CP cannot return class combinations not appearing in the training set.

iii) ECC solves the MDC problem via training a chain of multi-class classifiers, one per dimension, where predictions of preceding classifiers on the chain are used as extra features by the subsequent ones.

iv) ESC preprocesses the MDC problem via partitioning the class variables into super-class, where each super-class is used as a compound class variable.

v) gMML works by learning a regression model for each class label as well as a Mahalanobis metric which can shorten the distance between the regression outputs and ground-truth label vector.

vi) SEEM models the class dependency via a two-level strategy, where the pairwise and high-order class depend-

encies are modeled in the first and second levels, respectively.

For BR, CP, ECC, ESC, and SEEM, the multi-class base learner is implemented via LIBLINEAR^[30] with the parameter setting “L2-regularized logistic regression (primal)” for a fair comparison. Following [2], for ensemble approaches ECC and ESC, a total of 10 base models are trained over 67% examples randomly selected from the training set, and the predictive results are combined via majority voting. For gMML and SEEM, the recommended parameters are used according to respective literature.

In order to validate the effectiveness of the selective feature augmentation strategy for the proposed SFAM approach, a total of three simple kinds of augmented features are generated. Two of them are generated using standard and weighted k NN techniques, respectively, and the remaining one is generated using maximum margin techniques. Specifically, the two kinds of k NN-based augmented features are generated by KRAM^[20]. To be more specific, for each instance \mathbf{x} , let $\mathcal{N}_k(\mathbf{x}) = \{i_r \mid 1 \leq r \leq k\}$ be the set of indices for the k nearest neighbors of \mathbf{x} identified in the training set \mathcal{D} , we can define an indicating vector $\mathbf{v}_{ja}^{\mathbf{x}} = [v_{ja}^{\mathbf{x}}(1), v_{ja}^{\mathbf{x}}(2), \dots, v_{ja}^{\mathbf{x}}(k)]^T \in \{0, 1\}^k$ which is defined as follows:

$$v_{ja}^{\mathbf{x}}(r) = \mathbf{1}_{y_{i_r j} = c_a^j} \quad (1 \leq r \leq k, i_r \in \mathcal{N}_k(\mathbf{x})).$$

Here, $1 \leq a \leq K_j, 1 \leq j \leq q$. $\mathbf{y}_{i_r} = [y_{i_r 1}, \dots, y_{i_r q}]^T$ corresponds to the class vector of the neighboring MDC example \mathbf{x}_{i_r} for \mathbf{x} . Based on $\mathbf{v}_{ja}^{\mathbf{x}}$, the following discrete version of statistics $\delta_j^{\mathbf{x}} = [\delta_{j1}^{\mathbf{x}}, \delta_{j2}^{\mathbf{x}}, \dots, \delta_{jK_j}^{\mathbf{x}}]^T$ can be defined w.r.t. the j -th class space:

$$\delta_{ja}^{\mathbf{x}} = \langle \mathbb{I}_k, \mathbf{v}_{ja}^{\mathbf{x}} \rangle \quad (1 \leq a \leq K_j)$$

where \mathbb{I}_k is a column vector of all ones with length k . By concatenating all the q counting statistics vectors, the first k NN-augmented feature vector $\Delta_{k\text{NN}1}^{\mathbf{x}}$ based on standard k NN techniques for \mathbf{x} can be obtained:

$$\Delta_{k\text{NN}1}^{\mathbf{x}} = [\delta_1^{\mathbf{x}}, \delta_2^{\mathbf{x}}, \dots, \delta_q^{\mathbf{x}}]^T. \quad (10)$$

Moreover, let $\beta = [1, 1/\sqrt{2}, \dots, 1/\sqrt{k}]^T$, a bias vector $\zeta_j^{\mathbf{x}} = [\zeta_{j1}^{\mathbf{x}}, \zeta_{j2}^{\mathbf{x}}, \dots, \zeta_{jK_j}^{\mathbf{x}}]^T$ is defined as follows:

$$\zeta_{ja}^{\mathbf{x}} = \frac{\langle \beta, \mathbf{v}_{ja}^{\mathbf{x}} \rangle - \min(\mathbf{v}_{ja}^{\mathbf{x}})}{\max(\mathbf{v}_{ja}^{\mathbf{x}}) - \min(\mathbf{v}_{ja}^{\mathbf{x}})} (\zeta_{\max} - \zeta_{\min}) + \zeta_{\min}.$$

Here, ζ_{\max} and ζ_{\min} are two hyper-parameters and set as 0.5 and 0 respectively, and $\max(\mathbf{v}_{ja}^{\mathbf{x}}) = \sum_{r=1}^{\delta_{ja}^{\mathbf{x}}} \beta(r)$, $\min(\mathbf{v}_{ja}^{\mathbf{x}}) = \sum_{r=r_0}^k \beta(r)$ where $r_0 = k - \delta_{ja}^{\mathbf{x}} + 1$ and $\beta(r)$ denotes the r -th element of the weight vector β . Then, the second k NN-augmented feature vector $\Delta_{k\text{NN}2}^{\mathbf{x}}$ based on weighted k NN techniques for \mathbf{x} can be obtained:

$$\Delta_{k\text{NN}2}^{\mathbf{x}} = \Delta_{k\text{NN}1}^{\mathbf{x}} + [\zeta_1^{\mathbf{x}}, \zeta_2^{\mathbf{x}}, \dots, \zeta_q^{\mathbf{x}}]^T. \quad (11)$$

For the maximum margin-augmented features, SFAM employs the real-valued predictions, which are returned by the multi-class support vector machine to generate the maximum margin-augmented features. Specifically, SFAM solves the following maximum margin formulation^[31] for the j -th dimension ($1 \leq j \leq q$):

$$\begin{aligned} \min_{\mathbf{P}_j, \xi_j} \quad & \frac{1}{2} \sum_{a=1}^{K_j} \|\mathbf{p}_{ja}\|_2^2 + \gamma \times \sum_{i=1}^m \xi_{ji} \\ \text{s.t.} \quad & \langle \mathbf{p}_{ja_i}, \mathbf{x}_i \rangle - \langle \mathbf{p}_{ja}, \mathbf{x}_i \rangle \geq e_{ji}^a - \xi_{ji}, \\ & 1 \leq a \leq K_j, 1 \leq i \leq m \end{aligned}$$

where $\mathbf{z}_i = [\mathbf{x}_i; \tilde{\mathbf{x}}_i]$, $\mathbf{P}_j = [\mathbf{p}_{j1}, \mathbf{p}_{j2}, \dots, \mathbf{p}_{jK_j}] \in \mathbf{R}^{d \times K_j}$ is the weight matrix to be determined, and $\xi_j = [\xi_{j1}, \xi_{j2}, \dots, \xi_{jm}]^T \in \mathbf{R}^{m \times 1}$ is the slack variable vector. Suppose $y_{ij} = c_{a_i}^j$, then $e_{ji}^a = 1$ if $y_{ij} \neq c_a^j$ (i.e., $a_i \neq a$) and 0 otherwise. Furthermore, γ is a trade-off parameter. Let $\mathbf{P} = [\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_q]$, the maximum margin-augmented feature vector $\Delta_{\text{SVM}}^{\mathbf{x}}$ for each instance \mathbf{x} can be defined as follows:

$$\Delta_{\text{SVM}}^{\mathbf{x}} = \mathbf{P}^T \mathbf{x}. \quad (12)$$

Then, the MDC training set \mathcal{D} can be transformed into:

$$\tilde{\mathcal{D}} = \{(\mathbf{z}_i, \mathbf{y}_i) \mid 1 \leq i \leq m\}$$

where $\mathbf{z}_i = [\mathbf{x}_i; \tilde{\mathbf{x}}_i]$, and $\tilde{\mathbf{x}}_i = [\Delta_{k\text{NN}1}^{\mathbf{x}_i}, \Delta_{k\text{NN}2}^{\mathbf{x}_i}, \Delta_{\text{SVM}}^{\mathbf{x}_i}]$ denotes the augmented features of \mathbf{x}_i . Here, we reiterate that we make use of standard k NN, weighted k NN, and maximum margin techniques to generate the augmented feature only for the purpose of simplicity. The experiments in this paper mainly aims at validating the effectiveness of the selective feature augmentation strategy. In the future, it is interesting to further investigate synergizing multiple kinds of augmented features which are generated by making use of more advanced techniques such as deep learning^[21].

4.2 Experimental results

Table 2 reports the detailed experimental results with the performance rank shown in the parentheses. Moreover, the Wilcoxon signed-ranks test^[32] (at 0.05 significance level) serves as the statistical tool to show whether SFAM achieves better performance against the compared approaches over the whole benchmark data sets. The corresponding test results are summarized in Table 3.

According to the reported experimental results, we can make the following observations:

1) Among all the 30 cases (10 data sets \times 3 evaluation

Table 2 Experimental results (mean \pm std. deviation) of each MDC approach. In addition, the performance rank on each data set is also shown in the parentheses.

Data set	Hamming score						
	SFAM	BR	CP	ECC	ESC	gMML	SEEM
Flare1	0.925 \pm 0.035(2)	0.925 \pm 0.034(2)	0.923 \pm 0.033(7)	0.926 \pm 0.034(1)	0.925 \pm 0.035(2)	0.925 \pm 0.034(2)	0.925 \pm 0.032(2)
Enb	0.865 \pm 0.036(1)	0.774 \pm 0.023(3)	0.764 \pm 0.031(5)	0.773 \pm 0.034(4)	0.754 \pm 0.029(6)	0.742 \pm 0.027(7)	0.777 \pm 0.031(2)
WQplants	0.666 \pm 0.016(1)	0.658 \pm 0.014(3)	0.649 \pm 0.016(7)	0.654 \pm 0.016(5)	0.653 \pm 0.016(6)	0.655 \pm 0.015(4)	0.661 \pm 0.023(2)
WQanimals	0.641 \pm 0.012(1)	0.631 \pm 0.013(3)	0.628 \pm 0.013(7)	0.629 \pm 0.013(6)	0.631 \pm 0.014(3)	0.630 \pm 0.015(5)	0.635 \pm 0.015(2)
WQ	0.653 \pm 0.012(1)	0.644 \pm 0.011(3)	0.625 \pm 0.011(7)	0.642 \pm 0.012(5)	0.642 \pm 0.014(5)	0.643 \pm 0.013(4)	0.646 \pm 0.014(2)
BeLaE	0.441 \pm 0.013(1)	0.427 \pm 0.017(2)	0.383 \pm 0.023(7)	0.424 \pm 0.021(3)	0.420 \pm 0.022(4)	0.417 \pm 0.020(5)	0.416 \pm 0.020(6)
Voice	0.947 \pm 0.009(1)	0.900 \pm 0.012(3)	0.898 \pm 0.011(4)	0.896 \pm 0.012(6)	0.897 \pm 0.011(5)	0.842 \pm 0.009(7)	0.910 \pm 0.011(2)
TIC2000	0.946 \pm 0.003(1)	0.915 \pm 0.006(3)	0.905 \pm 0.006(6)	0.915 \pm 0.006(3)	0.915 \pm 0.006(3)	0.895 \pm 0.007(7)	0.917 \pm 0.005(2)
Adult	0.724 \pm 0.005(1)	0.721 \pm 0.004(2)	0.709 \pm 0.004(6)	0.720 \pm 0.003(4)	0.710 \pm 0.005(5)	0.705 \pm 0.004(7)	0.721 \pm 0.004(2)
Default	0.671 \pm 0.003(3)	0.669 \pm 0.003(5)	0.669 \pm 0.004(5)	0.670 \pm 0.003(4)	0.672 \pm 0.004(1)	0.666 \pm 0.004(7)	0.672 \pm 0.003(1)
Data set	Exact match						
	SFAM	BR	CP	ECC	ESC	gMML	SEEM
Flare1	0.824 \pm 0.073(1)	0.821 \pm 0.075(4)	0.817 \pm 0.068(7)	0.824 \pm 0.073(1)	0.824 \pm 0.073(1)	0.821 \pm 0.075(4)	0.818 \pm 0.075(6)
Enb	0.729 \pm 0.071(1)	0.548 \pm 0.045(3)	0.529 \pm 0.063(5)	0.546 \pm 0.069(4)	0.508 \pm 0.057(6)	0.483 \pm 0.053(7)	0.554 \pm 0.063(2)
WQplants	0.100 \pm 0.037(1)	0.092 \pm 0.033(5)	0.093 \pm 0.031(3)	0.092 \pm 0.034(5)	0.093 \pm 0.036(3)	0.092 \pm 0.035(5)	0.096 \pm 0.034(2)
WQanimals	0.058 \pm 0.013(5)	0.058 \pm 0.017(5)	0.065 \pm 0.018(1)	0.059 \pm 0.017(4)	0.064 \pm 0.019(2)	0.062 \pm 0.023(3)	0.049 \pm 0.022(7)
WQ	0.009 \pm 0.006(1)	0.005 \pm 0.008(4)	0.000 \pm 0.000(7)	0.005 \pm 0.008(4)	0.005 \pm 0.008(4)	0.006 \pm 0.008(3)	0.009 \pm 0.006(1)
BeLaE	0.028 \pm 0.011(1)	0.021 \pm 0.008(7)	0.026 \pm 0.014(3)	0.023 \pm 0.010(5)	0.027 \pm 0.009(2)	0.022 \pm 0.009(6)	0.026 \pm 0.011(3)
Voice	0.897 \pm 0.018(1)	0.809 \pm 0.023(3)	0.807 \pm 0.021(4)	0.802 \pm 0.022(6)	0.803 \pm 0.019(5)	0.699 \pm 0.017(7)	0.831 \pm 0.020(2)
TIC2000	0.846 \pm 0.009(1)	0.762 \pm 0.017(3)	0.739 \pm 0.017(6)	0.762 \pm 0.017(3)	0.762 \pm 0.016(3)	0.706 \pm 0.018(7)	0.770 \pm 0.015(2)
Adult	0.284 \pm 0.010(5)	0.275 \pm 0.008(6)	0.317 \pm 0.010(1)	0.287 \pm 0.007(4)	0.312 \pm 0.011(2)	0.230 \pm 0.009(7)	0.289 \pm 0.010(3)
Default	0.185 \pm 0.006(4)	0.181 \pm 0.007(6)	0.194 \pm 0.008(1)	0.185 \pm 0.006(4)	0.187 \pm 0.007(3)	0.177 \pm 0.007(7)	0.190 \pm 0.009(2)
Data set	Sub-exact match						
	SFAM	BR	CP	ECC	ESC	gMML	SEEM
Flare1	0.954 \pm 0.039(5)	0.957 \pm 0.039(2)	0.954 \pm 0.039(5)	0.957 \pm 0.039(2)	0.954 \pm 0.039(5)	0.957 \pm 0.039(2)	0.960 \pm 0.033(1)
Enb	1.000 \pm 0.000(1)	1.000 \pm 0.000(1)	1.000 \pm 0.000(1)	1.000 \pm 0.000(1)	1.000 \pm 0.000(1)	1.000 \pm 0.000(1)	1.000 \pm 0.000(1)
WQplants	0.292 \pm 0.035(1)	0.286 \pm 0.044(3)	0.285 \pm 0.052(5)	0.285 \pm 0.053(5)	0.282 \pm 0.049(7)	0.286 \pm 0.053(3)	0.287 \pm 0.042(2)
WQanimals	0.246 \pm 0.034(1)	0.229 \pm 0.030(5)	0.232 \pm 0.032(3)	0.226 \pm 0.026(7)	0.231 \pm 0.029(4)	0.227 \pm 0.033(6)	0.241 \pm 0.029(2)
WQ	0.061 \pm 0.022(1)	0.047 \pm 0.019(6)	0.034 \pm 0.017(7)	0.048 \pm 0.022(4)	0.048 \pm 0.019(4)	0.049 \pm 0.024(3)	0.050 \pm 0.025(2)
BeLaE	0.137 \pm 0.024(1)	0.134 \pm 0.025(2)	0.117 \pm 0.019(7)	0.130 \pm 0.025(3)	0.128 \pm 0.024(5)	0.130 \pm 0.020(3)	0.125 \pm 0.022(6)
Voice	0.998 \pm 0.003(1)	0.991 \pm 0.006(2)	0.989 \pm 0.006(5)	0.989 \pm 0.008(5)	0.991 \pm 0.007(2)	0.985 \pm 0.011(7)	0.990 \pm 0.006(4)
TIC2000	0.993 \pm 0.002(1)	0.983 \pm 0.003(4)	0.978 \pm 0.002(6)	0.984 \pm 0.003(2)	0.984 \pm 0.003(2)	0.978 \pm 0.003(6)	0.982 \pm 0.003(5)
Adult	0.690 \pm 0.007(1)	0.685 \pm 0.009(2)	0.637 \pm 0.007(7)	0.679 \pm 0.008(4)	0.644 \pm 0.007(6)	0.669 \pm 0.008(5)	0.680 \pm 0.006(3)
Default	0.604 \pm 0.007(1)	0.601 \pm 0.006(4)	0.594 \pm 0.008(6)	0.600 \pm 0.007(5)	0.604 \pm 0.008(1)	0.593 \pm 0.008(7)	0.604 \pm 0.007(1)

metrics), SFAM ranks first in 24 cases, second in 1 case, third in 1 case, fourth in 1 case, fifth in 3 cases, and never ranks last.

2) BR solves the MDC problem by dealing with each dimension independently, where potential class dependencies are fully ignored. Although SFAM also induces classification models for each dimension independently, the

class dependencies can be considered by the augmented features^[20]. It is shown that SFAM achieves superior performance against BR in terms of each metric, which reveals that considering class dependencies is important for learning MDC models.

3) Both ECC and gMML explicitly consider the class dependencies, where a chaining order over class spaces or

Table 3 Wilcoxon signed-ranks test for SFAM against each compared approach where the p-values at 0.05 significance level are also shown in the brackets.

SFAM against	Evaluation metric		
	HS	EM	SEM
BR	win [3.91e-03]	win [3.91e-03]	win [7.81e-03]
CP	win [1.95e-03]	tie [2.75e-01]	win [7.81e-03]
ECC	win [3.91e-03]	win [3.91e-02]	win [7.81e-03]
ESC	win [7.81e-03]	tie [2.62e-01]	win [1.56e-02]
gMML	win [3.91e-03]	win [9.77e-03]	win [7.81e-03]
SEEM	win [9.77e-03]	tie [7.42e-02]	win [3.91e-02]

a Mahalanobis metric is employed to accomplish this task. It is shown that SFAM also achieves superior performance against ECC and gMML in terms of each metric, which validates the superiority of SFAM's selective feature augmentation strategy.

4) CP solves the MDC problem by dealing with all dimensions jointly via powerset transformation, which can be viewed as optimizing exact match. ESC and SEEM can be regarded as two improved versions of CP, where class spaces are grouped into super-classes according to conditional dependencies, or each pair of class spaces are considered in the first level learning. It is shown that SFAM still achieves comparable performance against CP, ESC, and SEEM in terms of exact match and superior performance against CP, ESC, and SEEM in terms of Hamming score and sub-exact match.

4.3 Further analysis

1) Effectiveness of algorithmic design

In this paper, SFAM generates three kinds of augmented features according to (10)–(12), respectively. To further investigate the effectiveness of SFAM's algorithmic design, we also compare SFAM with its three degenerated versions which generate either kind of augmented features. The one with the discrete version of k NN-augmented features in (10) is denoted as DeV1, which is also known as the KRAM_d approach^[20], the one with the continuous version of k NN-augmented features in (11) is denoted as DeV2, which is also known as the KRAM_c approach^[20], and another one with maximum margin-augmented features in (12) is denoted as DeV3. It is worth noting that the baseline BR actually serves as another degenerated version without any kind of augmented features, whose experimental results have been reported and analyzed in Section 4.2.

Detailed experimental results are shown in Table 4. Table 5 summarizes the test results of the Wilcoxon signed-ranks test (at 0.05 significance level). It is shown that SFAM achieves superior performance against DeV1 in terms of Hamming score and exact match, and DeV3 in terms of all metrics. For DeV2, although SFAM

achieves comparable performance against it in terms of all metrics, as shown in Table 4, among the 19 cases where the performance of SFAM is different from DeV2, there are a total of 14 cases where the performance of SFAM is better than DeV2. These results clearly validate that SFAM can identify the pertinent and useful features from the three kinds of augmented features. Besides, it is shown that both DeV1 and DeV2 achieve similar results compared to SFAM. A possible reason is that the two kinds of k NN-augmented features contain more useful discriminative information than the maximum margin-augmented features. Nonetheless, SFAM is able to utilize all the available information in hand to achieve better generalization performance.

Furthermore, Fig. 3 shows the weight matrix (absolute value) of the learned model w.r.t. the ℓ_1 regularization. Specifically, following the notations in Section 3, Fig. 3 shows the absolute value of the weight matrix $\tilde{\Theta} = [\tilde{\theta}_{11}, \dots, \tilde{\theta}_{1K_1}, \dots, \tilde{\theta}_{q1}, \dots, \tilde{\theta}_{qK_q}]^T$ for the data set Voice, TIC2000, Adult, and Default. For each figure, each row corresponds to the binary classification model of one class label. The first third of the columns correspond to the discrete version of k NN-augmented features in (10), the middle third of the columns to the continuous version of k NN-augmented features in (11), and the last third of the columns to the maximum margin-augmented features in (12). It is shown that, for each third of all columns, the diagonal element usually takes the largest value in its corresponding row. Note that each element in all the three kinds of augmented features (i.e., each column in Fig. 3) corresponds to one class label, and each binary classification model (i.e., each row in Fig. 3) also corresponds to one class label. In other words, the largest value corresponds to the augmented feature w.r.t. its own class label. It is also shown that each binary classification model is only related to a part of the augmented features, where the model weights w.r.t. the two kinds of k NN-augmented features are usually larger than the model weights w.r.t. the maximum margin-augmented features. This observation further supports the afore-mentioned conjecture that the two kinds of k NN-augmented features contain more useful discriminative information than the maximum margin-augmented features. Besides, we can also observe that the model weights w.r.t. the ℓ_1 regularization for some binary classification models are almost all zero, which means that not all binary classification models rely on augmented features.

2) Parameter sensitivity analysis

The regularized classification model (1) has one trade-off parameter, λ . In this section, we investigate how the performance of SFAM changes with different values of λ . Fig. 4 illustrates SFAM's performance fluctuation when λ ranges in $\{0.01, 0.1, 1, 10, 100\}$ over data sets Flare1, WQplants, WQanimals, and BeLaE. It is shown that the performance of SFAM degenerates with either a small or

Table 4 Experimental results (mean \pm std. deviation) of SFAM and its three degenerated versions. In addition, the performance rank on each data set is also shown in the parentheses.

Data set	Hamming score			
	SFAM	DeV1	DeV2	DeV3
Flare1	0.925 \pm 0.035(1)	0.924 \pm 0.036(4)	0.925 \pm 0.035(1)	0.925 \pm 0.035(1)
Enb	0.865 \pm 0.036(1)	0.852 \pm 0.036(2)	0.848 \pm 0.043(3)	0.807 \pm 0.029(4)
WQpla.	0.666 \pm 0.016(1)	0.664 \pm 0.016(2)	0.664 \pm 0.016(2)	0.659 \pm 0.015(4)
WQani.	0.641 \pm 0.012(1)	0.641 \pm 0.015(1)	0.641 \pm 0.015(1)	0.632 \pm 0.013(4)
WQ	0.653 \pm 0.012(1)	0.650 \pm 0.011(3)	0.652 \pm 0.010(2)	0.645 \pm 0.012(4)
BeLaE	0.441 \pm 0.013(1)	0.436 \pm 0.008(2)	0.434 \pm 0.011(4)	0.435 \pm 0.019(3)
Voice	0.947 \pm 0.009(1)	0.945 \pm 0.010(3)	0.946 \pm 0.010(2)	0.913 \pm 0.010(4)
TIC2000	0.946 \pm 0.003(2)	0.946 \pm 0.003(2)	0.947 \pm 0.003(1)	0.914 \pm 0.006(4)
Adult	0.724 \pm 0.005(1)	0.724 \pm 0.004(1)	0.724 \pm 0.005(1)	0.721 \pm 0.004(4)
Default	0.671 \pm 0.003(1)	0.671 \pm 0.003(1)	0.671 \pm 0.003(1)	0.669 \pm 0.003(4)
Data set	Exact match			
	SFAM	DeV1	DeV2	DeV3
Flare1	0.824 \pm 0.073(1)	0.821 \pm 0.079(4)	0.824 \pm 0.073(1)	0.824 \pm 0.073(1)
Enb	0.729 \pm 0.071(1)	0.705 \pm 0.072(2)	0.696 \pm 0.086(3)	0.613 \pm 0.058(4)
WQpla.	0.100 \pm 0.037(1)	0.097 \pm 0.038(2)	0.097 \pm 0.037(2)	0.095 \pm 0.034(4)
WQani.	0.058 \pm 0.013(3)	0.058 \pm 0.011(3)	0.061 \pm 0.016(1)	0.059 \pm 0.018(2)
WQ	0.009 \pm 0.006(1)	0.008 \pm 0.005(2)	0.008 \pm 0.005(2)	0.005 \pm 0.008(4)
BeLaE	0.028 \pm 0.011(1)	0.025 \pm 0.007(2)	0.024 \pm 0.007(3)	0.024 \pm 0.008(3)
Voice	0.897 \pm 0.018(1)	0.893 \pm 0.020(3)	0.894 \pm 0.021(2)	0.834 \pm 0.018(4)
TIC2000	0.846 \pm 0.009(3)	0.847 \pm 0.008(2)	0.849 \pm 0.008(1)	0.760 \pm 0.018(4)
Adult	0.284 \pm 0.010(1)	0.284 \pm 0.009(1)	0.284 \pm 0.010(1)	0.274 \pm 0.008(4)
Default	0.185 \pm 0.006(2)	0.185 \pm 0.006(2)	0.186 \pm 0.006(1)	0.181 \pm 0.006(4)
Data set	Sub-exact match			
	SFAM	DeV1	DeV2	DeV3
Flare1	0.954 \pm 0.039(1)	0.954 \pm 0.039(1)	0.954 \pm 0.039(1)	0.954 \pm 0.039(1)
Enb	1.000 \pm 0.000(1)	1.000 \pm 0.000(1)	1.000 \pm 0.000(1)	1.000 \pm 0.000(1)
WQpla.	0.292 \pm 0.035(3)	0.294 \pm 0.037(1)	0.294 \pm 0.036(1)	0.289 \pm 0.042(4)
WQani.	0.246 \pm 0.034(1)	0.244 \pm 0.031(2)	0.243 \pm 0.035(3)	0.231 \pm 0.032(4)
WQ	0.061 \pm 0.022(1)	0.059 \pm 0.019(2)	0.058 \pm 0.020(3)	0.048 \pm 0.020(4)
BeLaE	0.137 \pm 0.024(1)	0.133 \pm 0.021(4)	0.134 \pm 0.022(2)	0.134 \pm 0.023(2)
Voice	0.998 \pm 0.003(1)	0.998 \pm 0.003(1)	0.998 \pm 0.003(1)	0.993 \pm 0.006(4)
TIC2000	0.993 \pm 0.002(1)	0.992 \pm 0.002(3)	0.993 \pm 0.002(1)	0.983 \pm 0.003(4)
Adult	0.690 \pm 0.007(1)	0.690 \pm 0.007(1)	0.689 \pm 0.007(3)	0.685 \pm 0.008(4)
Default	0.604 \pm 0.007(1)	0.604 \pm 0.007(1)	0.604 \pm 0.006(1)	0.601 \pm 0.007(4)

large value of λ generally, and $\lambda = 1$ is usually a better choice. Therefore, we set λ to 1 in all the previous comparative studies.

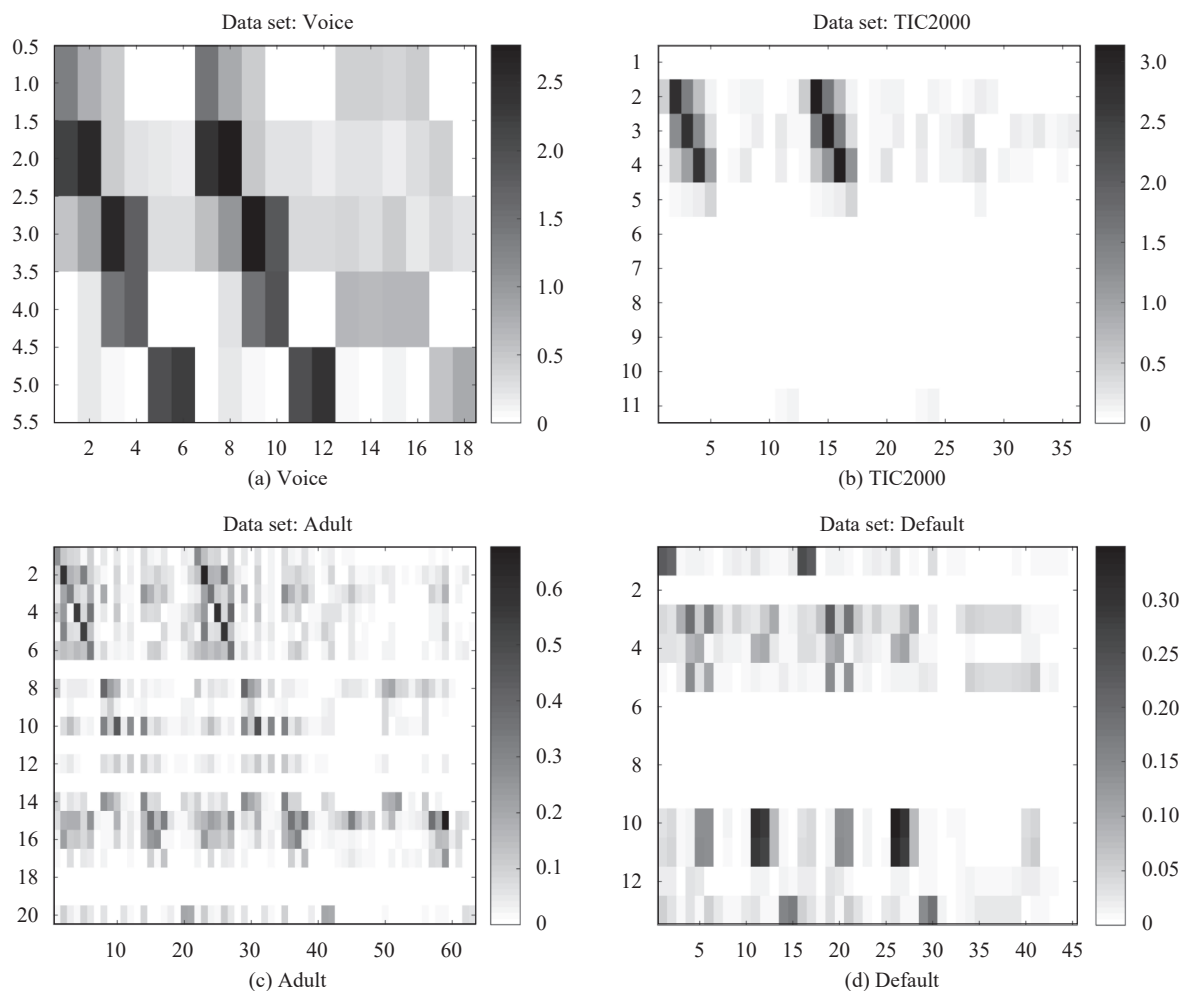
5 Conclusions

Feature augmentation has been shown as an effective

strategy for solving the MDC problem. Existing works only focus on generating better augmented features, while it might be beneficial to exploit multiple kinds of augmented features generated using different techniques. This paper makes a first attempt to synergize the discriminative information residing in multiple kinds of aug-

Table 5 Wilcoxon signed-ranks test for SFAM against its three variants where the p-values at 0.05 significance level are also shown in the brackets.

SFAM against	Evaluation metric		
	HS	EM	SEM
DeV1	win [3.13e-02]	win [4.69e-02]	tie [3.13e-01]
DeV2	tie [1.25e-01]	tie [3.52e-01]	tie [1.88e-01]
DeV3	win [3.91e-03]	win [7.81e-03]	win [7.81e-03]

Fig. 3 The weight matrix (absolute value) of the learned model w.r.t. the ℓ_1 regularization

mented features. Accordingly, a novel strategy named selective feature augmentation is proposed, which assumes that only part of the augmented features is pertinent and useful for each dimension's model induction. Comparative studies clearly validate the effectiveness of the proposed strategy.

Current feature augmentation works simply concatenate the original and augmented features, though the proposed SFAM has treated them differently via different regularization terms. In fact, the original and augmented features (even different kinds of augmented features) can be regarded as features from different views^[33]. In the future, other ensemble strategies borrowing from multi-view learning can also be used instead of merely using the con-

catenation operation. Besides, this paper only generates three simple kinds of augmented features to validate the proposed selective feature augmentation strategy, it is also worth investigating generating more kinds of augmented features.

Acknowledgements

This work was supported by National Science Foundation of China (No. 62176055) and China University S&T Innovation Plan Guided by the Ministry of Education. The authors wish to thank the associate editor and anonymous reviewers for their insightful comments and suggestions. The authors thank the Big Data Center of

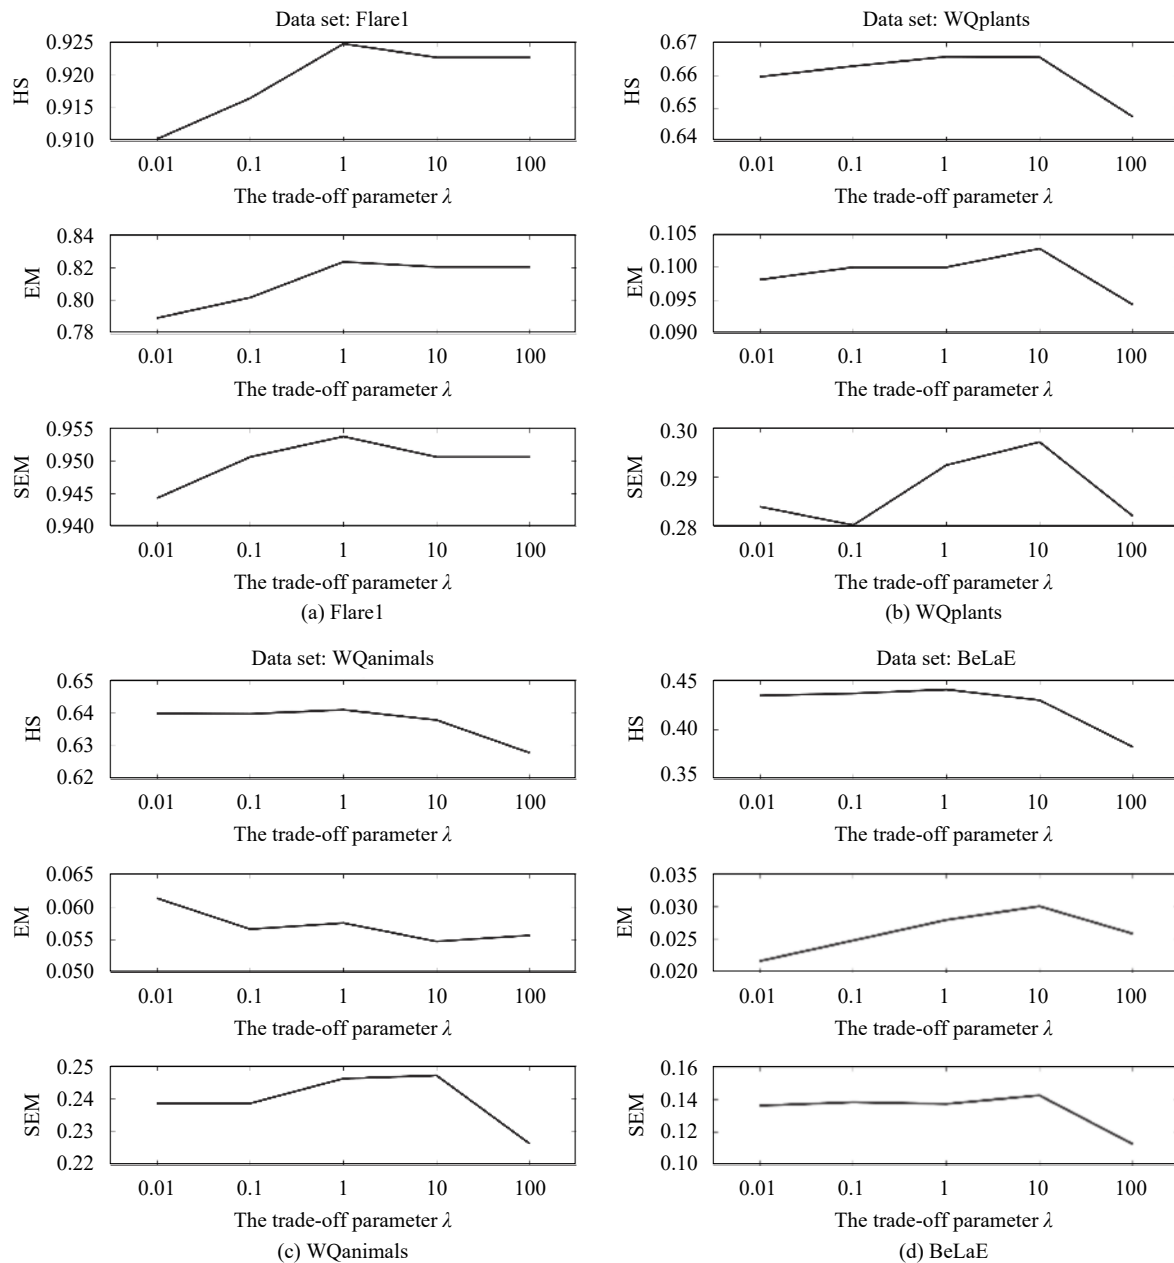


Fig. 4 Performance SFAM changes as λ ranges in $\{0.01, 0.1, 1, 10, 100\}$

Southeast University for providing the facility support on the numerical calculations in this paper.

Open Access

This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the

material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>

References

- [1] D. Xu, Y. X. Shi, I. W. Tsang, Y. S. Ong, C. Gong, X. B. Shen. Survey on multi-output learning. *IEEE Transactions on Neural Networks and Learning Systems*, vol.31, no. 7, pp. 2409–2429, 2020. DOI: [10.1109/TNNLS.2019.2945133](https://doi.org/10.1109/TNNLS.2019.2945133).

- [2] J. Read, C. Bielza, P. Larrañaga. Multi-dimensional classification with super-classes. *IEEE Transactions on Knowledge and Data Engineering*, vol.26, no.7, pp.1720–1733, 2014. DOI: [10.1109/TKDE.2013.167](https://doi.org/10.1109/TKDE.2013.167).
- [3] B. B. Jia, M. L. Zhang. Maximum margin multi-dimensional classification. *IEEE Transactions on Neural Networks and Learning Systems*, published online. DOI: [10.1109/TNNLS.2021.3084373](https://doi.org/10.1109/TNNLS.2021.3084373).
- [4] J. D. Rodríguez, A. Pérez, D. Arteta, D. Tejedor, J. A. Lozano. Using multidimensional Bayesian network classifiers to assist the treatment of multiple sclerosis. *IEEE Transactions on Systems, Man, and Cybernetics – Part C*, vol.42, no.6, pp.1705–1715, 2012. DOI: [10.1109/TSMCC.2012.2217326](https://doi.org/10.1109/TSMCC.2012.2217326).
- [5] H. Borchani, C. Bielza, C. Toro, P. Larrañaga. Predicting human immunodeficiency virus inhibitors using multi-dimensional Bayesian network classifiers. *Artificial Intelligence in Medicine*, vol.57, no.3, pp.219–229, 2013. DOI: [10.1016/j.artmed.2012.12.005](https://doi.org/10.1016/j.artmed.2012.12.005).
- [6] H. Shatkay, F. X. Pan, A. Rzhetsky, W. J. Wilbur. Multi-dimensional classification of biomedical text: Toward automated, practical provision of high-utility text to diverse users. *Bioinformatics*, vol.24, no.18, pp.2086–2093, 2008. DOI: [10.1093/bioinformatics/btn381](https://doi.org/10.1093/bioinformatics/btn381).
- [7] F. Serafino, G. Pio, M. Ceci, D. Malerba. Hierarchical multidimensional classification of web documents with multi-webclass. In *Proceedings of the 18th International Conference on Discovery Science*, Springer, Banff, Canada, pp.236–250, 2015. DOI: [10.1007/978-3-319-24282-8_20](https://doi.org/10.1007/978-3-319-24282-8_20).
- [8] Z. Lian, Y. Li, J. H. Tao, J. Huang, M. Y. Niu. Expression analysis based on face regions in real-world conditions. *International Journal of Automation and Computing*, vol.17, no.1, pp.96–107, 2020. DOI: [10.1007/s11633-019-1176-9](https://doi.org/10.1007/s11633-019-1176-9).
- [9] Z. W. He, L. Zhang, F. Y. Liu. DiscoStyle: Multi-level logistic ranking for personalized image style preference inference. *International Journal of Automation and Computing*, vol.17, no.5, pp.637–651, 2020. DOI: [10.1007/s11633-020-1244-1](https://doi.org/10.1007/s11633-020-1244-1).
- [10] Y. Zhang, X. Y. Shi, S. Y. Mi, X. Yang. Image captioning with transformer and knowledge graph. *Pattern Recognition Letters*, vol.143, pp.43–49, 2021. DOI: [10.1016/j.patrec.2020.12.020](https://doi.org/10.1016/j.patrec.2020.12.020).
- [11] A. H. Al Muktaadir, T. Miyazawa, P. Martinez-Julia, H. Harai, V. P. Kaffle. Multi-target classification based automatic virtual resource allocation scheme. *IEICE Transactions on Information and Systems*, vol.E102-D, no.5, pp.898–909, 2019. DOI: [10.1587/transinf.2018NTP0016](https://doi.org/10.1587/transinf.2018NTP0016).
- [12] J. Arias, J. A. Gamez, T. D. Nielsen, J. M. Puerta. A scalable pairwise class interaction framework for multidimensional classification. *International Journal of Approximate Reasoning*, vol.68, pp.194–210, 2016. DOI: [10.1016/j.ijar.2015.07.007](https://doi.org/10.1016/j.ijar.2015.07.007).
- [13] B. B. Jia, M. L. Zhang. Multi-dimensional classification via stacked dependency exploitation. *Science China Information Sciences*, vol.63, no.12, Article number 222102, 2020. DOI: [10.1007/s11432-019-2905-3](https://doi.org/10.1007/s11432-019-2905-3).
- [14] B. B. Jia, M. L. Zhang. MD-KNN: An instance-based approach for multi-dimensional classification. In *Proceedings of the 25th International Conference on Pattern Recognition*, IEEE, Milan, Italy, pp.126–133, 2021. DOI: [10.1109/ICPR48806.2021.9412974](https://doi.org/10.1109/ICPR48806.2021.9412974).
- [15] J. H. Zaragoza, L. E. Sucar, E. F. Morales, C. Bielza, P. Larrañaga. Bayesian chain classifiers for multidimensional classification. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, Barcelona, Spain, pp.2192–2197, 2011. DOI: [10.5591/978-1-57735-516-8/IJCAI11-365](https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-365).
- [16] J. Read, L. Martino, D. Luengo. Efficient Monte Carlo methods for multi-dimensional learning with classifier chains. *Pattern Recognition*, vol.47, no.3, pp.1535–1546, 2014. DOI: [10.1016/j.patcog.2013.10.006](https://doi.org/10.1016/j.patcog.2013.10.006).
- [17] C. Bielza, G. Li, P. Larrañaga. Multi-dimensional classification with Bayesian networks. *International Journal of Approximate Reasoning*, vol.52, no.6, pp.705–727, 2011. DOI: [10.1016/j.ijar.2011.01.007](https://doi.org/10.1016/j.ijar.2011.01.007).
- [18] J. H. Bolt, L. C. Van Der Gaag. Balanced sensitivity functions for tuning multi-dimensional Bayesian network classifiers. *International Journal of Approximate Reasoning*, vol.80, pp.361–376, 2017. DOI: [10.1016/j.ijar.2016.07.011](https://doi.org/10.1016/j.ijar.2016.07.011).
- [19] M. Benjumea, C. Bielza, P. Larrañaga. Tractability of most probable explanations in multidimensional Bayesian network classifiers. *International Journal of Approximate Reasoning*, vol.93, pp.74–87, 2018. DOI: [10.1016/j.ijar.2017.10.024](https://doi.org/10.1016/j.ijar.2017.10.024).
- [20] B. B. Jia, M. L. Zhang. Multi-dimensional classification via kNN feature augmentation. *Pattern Recognition*, vol.106, Article number 107423, 2020. DOI: [10.1016/j.patcog.2020.107423](https://doi.org/10.1016/j.patcog.2020.107423).
- [21] H. B. Wang, C. Chen, W. W. Liu, K. Chen, T. L. Hu, G. Chen. Incorporating label embedding and feature augmentation for multi-dimensional classification. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, AAAI, New York, USA, pp.6178–6185, 2020. DOI: [10.1609/AAAI.V34I04.6083](https://doi.org/10.1609/AAAI.V34I04.6083).
- [22] M. L. Zhang, Z. H. Zhou. A review on multi-label learning algorithms. *IEEE Transactions on Knowledge and Data Engineering*, vol.26, no.8, pp.1819–1837, 2014. DOI: [10.1109/TKDE.2013.39](https://doi.org/10.1109/TKDE.2013.39).
- [23] E. Gibaja, S. Ventura. A tutorial on multilabel learning. *ACM Computing Surveys*, vol.47, no.3, Article number 52, 2015. DOI: [10.1145/2716262](https://doi.org/10.1145/2716262).
- [24] M. L. Zhang, Y. K. Li, X. Y. Liu, X. Geng. Binary relevance for multi-label learning: An overview. *Frontiers of Computer Science*, vol.12, no.2, pp.191–202, 2018. DOI: [10.1007/s11704-017-7031-7](https://doi.org/10.1007/s11704-017-7031-7).
- [25] S. Gil-Begue, C. Bielza, P. Larrañaga. Multi-dimensional Bayesian network classifiers: A survey. *Artificial Intelligence*

- gence Review*, vol.54, no.1, pp.519–559, 2021. DOI: [10.1007/s10462-020-09858-x](https://doi.org/10.1007/s10462-020-09858-x).
- [26] J. Huang, G. R. Li, Q. M. Huang, X. D. Wu. Learning label-specific features and class-dependent labels for multi-label classification. *IEEE Transactions on Knowledge and Data Engineering*, vol.28, no.12, pp.3309–3323, 2016. DOI: [10.1109/TKDE.2016.2608339](https://doi.org/10.1109/TKDE.2016.2608339).
- [27] H. H. Bauschke, J. Bolte, M. Teboulle. A descent lemma beyond Lipschitz gradient continuity: First-order methods revisited and applications. *Mathematics of Operations Research*, vol.42, no.2, pp.330–348, 2017. DOI: [10.1287/moor.2016.0817](https://doi.org/10.1287/moor.2016.0817).
- [28] A. Beck, M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, vol.2, no.1, pp.183–202, 2009. DOI: [10.1137/080716542](https://doi.org/10.1137/080716542).
- [29] Z. C. Ma, S. C. Chen. Multi-dimensional classification via a metric approach. *Neurocomputing*, vol.275, pp.1121–1131, 2018. DOI: [10.1016/j.neucom.2017.09.057](https://doi.org/10.1016/j.neucom.2017.09.057).
- [30] R. E. Fan, K. W. Chang, C. J. Hsieh, X. R. Wang, C. J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, vol.9, pp.1871–1874, 2008.
- [31] K. Crammer, Y. Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research*, vol.2, pp.265–292, 2001.
- [32] J. Demšar. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, vol.7, pp.1–30, 2006.
- [33] J. Zhao, X. J. Xie, X. Xu, S. L. Sun. Multi-view learning overview: Recent progress and new challenges. *Information Fusion*, vol.38, pp.43–54, 2017. DOI: [10.1016/j.inffus.2017.02.007](https://doi.org/10.1016/j.inffus.2017.02.007).



Bin-Bin Jia received the B.Sc. degree in electronic information science and technology from North China Electric Power University, China in 2010, and the M.Sc. degree in information and communication engineering from Beihang University, China in 2013. He joined College of Electrical and Information Engineering, Lanzhou University of Technology, China in

2013. Currently, he is a Ph.D. degree candidate in School of Computer Science and Engineering, Southeast University, China.

His research interests include machine learning and data mining, especially in multi-dimensional classification.

E-mail: jiabb@seu.edu.cn

ORCID iD: 0000-0003-3302-9398



Min-Ling Zhang received the B.Sc., M.Sc., and Ph.D. degrees in computer science from Nanjing University, China in 2001, 2004 and 2007, respectively. Currently, he is a professor at School of Computer Science and Engineering, Southeast University, China. In recent years, Dr. Zhang has served as the General Co-Chairs of ACML'18, Program Co-Chairs of PAK-

DD'19, CCF-ICAI'19, ACML'17, CCF-FAI'17, PRICAI'16, Senior PC member or Area Chair of AAAI 2017–2020, IJCAI 2017–2022, KDD 2021, ICDM 2015–2021, etc. He is also on the editorial board of *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *ACM Transactions on Intelligent Systems and Technology*, *Neural Networks*, *Science China Information Sciences*, *Frontiers of Computer Science*, etc. Dr. Zhang is the Steering Committee Member of ACML and PAKDD, Vice Chair of CAAI Machine Learning Society, standing committee member of the CCF Artificial Intelligence & Pattern Recognition Society. He is a Distinguished Member of CCF, CAAI, and Senior Member of ACM, IEEE.

His research interests include machine learning and data mining.

E-mail: zhangml@seu.edu.cn (Corresponding author)

ORCID iD: 0000-0003-1880-5918