

Label Embedding with Partial Heterogeneous Contexts

Yaxin Shi,[†] Donna Xu,[†] Yuangang Pan,[†] Ivor W. Tsang,[†] Shirui Pan[†]

[†]Centre for Artificial Intelligence(CAI), University of Technology Sydney, Australia
 {Yaxin.Shi, Donna.Xu, Yuangang.Pan}@student.uts.edu.au, {Ivor.Tsang, Shirui.Pan}@uts.edu.au

Abstract

Label embedding plays an important role in many real-world applications. To enhance the label relatedness captured by the embeddings, multiple contexts can be adopted. However, these contexts are heterogeneous and often partially observed in practical tasks, imposing significant challenges to capture the overall relatedness among labels. In this paper, we propose a general Partial Heterogeneous Context Label Embedding (PHCLE) framework to address these challenges. Categorizing heterogeneous contexts into two groups, relational context and descriptive context, we design tailor-made matrix factorization formula to effectively exploit the label relatedness in each context. With a shared embedding principle across heterogeneous contexts, the label relatedness is selectively aligned in a shared space. Due to our elegant formulation, PHCLE overcomes the partial context problem and can nicely incorporate more contexts, which both cannot be tackled with existing multi-context label embedding methods. An effective alternative optimization algorithm is further derived to solve the sparse matrix factorization problem. Experimental results demonstrate that the label embeddings obtained with PHCLE achieve superb performance in image classification task and exhibit good interpretability in the downstream label similarity analysis and image understanding task.

Introduction

Label embedding, providing representations for labels, has been widely used in object classification (Akata et al. 2016), image retrieval (Siddiquie, Feris, and Davis 2011) and novelty detection (Wah and Belongie 2013) tasks. Context information, such as label hierarchy (Miller et al. 1990), class co-occurrence statistics (Mensink, Gavves, and Snoek 2014), semantic attributes (Lampert, Nickisch, and Harmeling 2009), tags and text descriptions (Mikolov et al. 2013) have all been exploited to learn label embeddings. As these contexts provide label relatedness in different aspects, they are in good complement to each other for overall understanding of the labels. For example, weasels, a mammal of the genus *Mustela*, are considered to be related to cats as they share similar visual attributes. While from the perspective of label hierarchy, weasels should be much more related to the skunks as they belong to the same animal family. Therefore, it is necessary to leverage multiple contexts

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

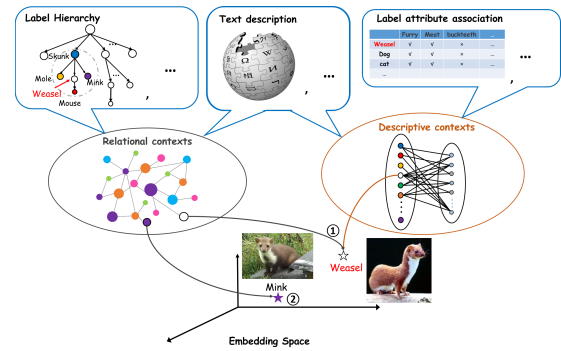


Figure 1: Motivation figure of PHCLE. Stars represent different label embeddings. ① denotes the label embedding for full contexts label; ② presents that of partial context label.

to learn the label embeddings so that it can well capture the label relatedness in multiple aspects.

However, learning label embedding from multiple contexts is challenging, due to the heterogeneous nature of the contexts. Generally, based on the relational properties, aforementioned label contexts can be grouped into two basic but heterogeneous categories: relational context and descriptive context. Examples for the heterogeneous contexts is given in Figure 1. The relational context conveys direct information on label relations. Label hierarchy, class co-occurrence statistics can be grouped into this category. The descriptive context, e.g. attributes and tags, provides the associations between labels and semantic descriptions. Label relatedness is reflected by the sharing descriptions. These two basic categories are universal for existing label contexts. For example, in word2vec (Mikolov et al. 2013), word-context pairs of noun-noun belong to relational context; pairs of noun-adjective and noun-verb belong to descriptive contexts. Universality of the two heterogeneous categories makes it reasonable to formulate multi-context label embedding learning as a general heterogeneous contexts embedding problem.

To solve the heterogeneous contexts embedding problem, we will have the following challenges:

- *Challenge 1:* How to effectively exploit the label relatedness conveyed in each type of the heterogeneous contexts?
- *Challenge 2:* How to align the label relatedness reflected in the heterogeneous contexts?

- *Challenge 3*: How can we overcome the partial context problem? As it is difficult to obtain label descriptions (Akata et al. 2016), the heterogeneous contexts are often partially observed in practical label embedding tasks. As shown in Figure 1, *mink* lacks the descriptive context.

These challenges cannot be comprehensively solved with existing studies. Former multi-contexts label embedding are mostly obtained via contexts fusion conducted by simple concatenation (Akata et al. 2015) or Canonical Correlation Analysis (Fu et al. 2014). In these approaches, the heterogeneous contexts are treated indiscriminately, so the label relatedness conveyed by the intrinsic property of each context is not fully exploited. Furthermore, as the dependencies between the contexts are not properly captured by their formulation, the contexts are not well aligned. Another representative work for attributed graph embedding, Text-associated Deep Walk (TADW) (Yang et al. 2015), deals with two heterogeneous contexts within one matrix factorization formula. One severe drawback of TADW is that it requires full correspondence between the contexts, which limits its ability in handling real-world label embedding tasks. In addition, only two contexts can be incorporated in the aforementioned works, which limits their ability to capture more complete relatedness among labels. These two limitations make it desirable to propose a model that can learn label embedding with more partial heterogeneous contexts.

In this paper, we formulate the *multi-context label embedding task* as a general heterogeneous context embedding problem, and propose a novel Partial Heterogeneous Context Label Embedding (PHCLE) approach to solve the challenges. To fully exploit the label relatedness in each context (*Challenge 1*), we tailor-make different matrix factorization formulas to learn the embedding from each type of contexts. To align the labeled relatedness in the heterogeneous contexts (*Challenge 2*), we adopt a shared embedding principle to capture their dependency. A sparsity constraint is further imposed on descriptive context embeddings to select the most discriminative descriptions for better contexts alignment. Due to the adopted shared embedding principle, the proposed PHCLE can handle partial context problem (*Challenge 3*) with an indicator matrix to indicate the missing entries. Furthermore, due to the additive property of the matrix factorization, the proposed PHCLE can be easily generalized to incorporate more contexts.

Contributions of this work can be summarized as follows:

- We propose a new framework for label embedding, which captures label contexts from two aspects—relational context and descriptive context. Our proposed model is flexible to be generalized to incorporate other label contexts.
- We study the new problem of partial correspondence in heterogeneous context embedding and present a PHCLE model as a solution. Our model captures the label relatedness of each context and aligns them in a shared embedding space via a joint matrix factorization framework, based on which an alternative optimization approach is derived to solve the problem effectively.
- Label embedding obtained with PHCLE achieves superior performance in image classification. Furthermore, the

superb interpretability of the obtained label embedding makes PHCLE promising in image understanding tasks.

Related Work

Label contexts Context information, such as label hierarchy, class co-occurrence statistics, semantic descriptions and text descriptions, have been widely adopted to learn the label embeddings. Label hierarchy, such as WordNet (Miller et al. 1990), defines the intrinsic structure of labels. Class co-occurrence statistics (Mensink, Gavves, and Snoek 2014) reflects label relations based on the label occurrence rate. Semantic descriptions, i.e. attributes (Lampert, Nickisch, and Harmeling 2009; Akata et al. 2016), provide descriptive information for labels. Relatedness between labels is implied by their common characteristics in the embeddings. Semantic text representations, such as Word2Vec and GloVe, preserve the semantic relatedness between labels based on the text information. As these contexts provide label relatedness in different aspects, it is promising to leverage multiple contexts to learn the label embeddings to capture the overall relatedness among labels. However, as those label contexts are heterogeneous, it is challenging to align multiple heterogeneous contexts in label embedding learning.

Multi-context label embedding methods Multiple contexts have been adopted in former embedding works. However, heterogeneity of adopted contexts is not considered. In (Akata et al. 2015), multiple label embeddings are fused through simple concatenation (CNC). As the embeddings for each context is independently learned, dependencies among the multiple contexts are not captured in CNC. Canonical Correlation Analysis (CCA) or its nonlinear variants (Andrew et al. 2013) can be adopted to fuse multiple contexts with the consideration of dependency (Fu et al. 2014). However, due to the intrinsic property of CCA, only the principal component variances of each context is preserved in the common latent space, while context information orthogonal to the principal directions are all lost. Consequently, the relative position of these labels are not well preserved in the obtained label embeddings. For representative multi-contexts network embeddings works, heterogeneity of the contexts is also overlooked. For example, in (Liao et al. 2017), the embeddings are obtained as early fusion for each context conducted via deep neural network. In (Tu et al. 2017), the incorporated contexts are uniformly modeled with softmax formulation. Other attributed network embedding approaches such as TADW (Yang et al. 2015) and AANE (Huang, Li, and Hu 2017), may also be adapted to learn the heterogeneous label embedding. However, these methods typically require full correspondence between different context. This drawback limits their ability in handling real-world label embedding tasks.

Partial Heterogeneous Context Label Embedding

In this section, we present the formulation of the proposed Partial Heterogeneous Context Label Embedding (PHCLE).

Notations

Let V_W and V_C be the label vocabulary and the context vocabulary in collection \mathcal{D} . And $D \in \mathbb{R}^{|V_C| \times |V_W|}$ is the co-occurrence matrix constructed via \mathcal{D} . $W \in \mathbb{R}^{n \times |V_W|}$ and $C \in \mathbb{R}^{n \times |V_C|}$ are the label embedding and the context embedding matrices, where n denotes the dimension of the label embedding. $A \in \mathbb{R}^{|V_W| \times m}$ is the label-attribute association matrix, where m is the number of attributes. $U \in \mathbb{R}^{n \times m}$ denotes the attribute embedding matrix. $\|\cdot\|_F$ denotes the Frobenius norm of the matrix and λ is the harmonic factor to balance the components in the formulation.

Formulation

In this paper, we tackle the challenges in label embedding with partial heterogeneous contexts, with three strategies.

Tailor-made formulas for heterogeneous contexts For relational context, label relatedness is directly implied by the relationship among labels. Defining label-context as the label relation, Skip-Gram Negative Sampling (SGNS) (Mikolov et al. 2013) is representative work that effectively learns label embeddings from label relations. To enable the incorporation of multiple relational contexts, we adopt an Explicit Matrix Factorization (EMF) (Li et al. 2015b) formulation of SGNS, which is given as

$$\begin{aligned} \min_{C,W} EMF(D, C^T W) \\ = -tr(D^T C^T W) + \sum_{w \in V_W} \log\left(\sum_{d'_w \in \mathcal{S}_w} e^{d'^T_w C^T w}\right), \end{aligned} \quad (1)$$

where W and C is the label embedding and the context embedding, respectively. $w \in V_W$ represents a label, $d_w \in \mathbb{R}^{|V_C|}$ denotes the explicit word vector of w , which also corresponds to the w^{th} column of the co-occurrence matrix D . \mathcal{S}_w is the Cartesian product of $|V_C|$ subsets, which represents w ' candidate set of all possible explicit word vectors. $\mathcal{S}_{w,c} = \{0, 1, \dots, Q_{w,c}\}$, where $Q_{w,c}$ is an upper bound of the co-occurrence count for the label w and context $c \in V_C$. $Q_{w,c}$ is set to be $k \frac{\sum_i |V_W| d_{i,c} \sum_j |V_C| d_{w,j}}{\sum_i |V_W| \sum_j |V_C| d_{i,j}} + d_{w,c}$, where k is the number of negative context samples for each label.

Specifically, as the label relatedness is conveyed by the label co-occurrence matrix D , multiple relational contexts can be exploited by defining multiple label relations (Omer and Yoav 2014) to construct D for specific label embedding task. Particularly, classic SGNS defines linear label-context relation in a large text corpus to learn label representations. However, this context definition fails to consider the structural nature of the labels, thus influence the label relatedness captured by the label embedding. Therefore, label relations that reveal the intrinsic structure of the labels, e.g. label hierarchy, are more suitable to be exploited to construct the co-occurrence matrix D in PHCLE. Consequently, label relatedness conveyed by relational contexts can be properly exploited with the proposed PHCLE.

For descriptive context, label relatedness is reflected by sharing descriptions. As the descriptive contexts A encodes the label-description associations, it can be effectively modeled with a traditional matrix factorization formula in Eq. (2) (Koren, Bell, and Volinsky 2009), with the two matrices representing the label embedding W and the description embedding U , respectively.

$$\|A - W^T U\|_F^2, \quad (2)$$

Align partially observed contexts via shared embedding

As both types of the contexts are properly modeled, we further adopt a shared embedding principle for the alignment of the heterogeneous contexts, which also works with partial contexts. Specifically, as shown in Eq. (3), we formulate the label embedding W to be shared by the two formulas (Eq. (1), Eq. (2)) to capture the dependencies of the heterogeneous label contexts. The adopted principle also enables the proposed PHCLE to handle partial context problem with the matrix I indicating the missing entries.

$$\min_{C,W,U} EMF(D, C^T W) + \frac{\lambda_1}{2} \|I \odot (A - W^T U)\|_F^2. \quad (3)$$

Enhance alignment via discriminative contexts selection

Furthermore, as label descriptions are either manually defined or learned by classifiers, they are often noisy or redundant (Akata et al. 2016). To better align the heterogeneous contexts, we impose a sparsity constraint on descriptive contexts embeddings U to select the most discriminative descriptions for contexts alignment. Consequently, the proposed Partial Heterogeneous Context Label Embedding (PHCLE) model can be formulated as

$$\begin{aligned} \min_{C,W,U} EMF(D, C^T W) + \frac{\lambda_1}{2} \|I \odot (A - W^T U)\|_F^2 + \lambda_2 \|U\|_1 \\ + \frac{\lambda_3}{2} (\|W\|_F^2 + \|U\|_F^2), \end{aligned} \quad (4)$$

In Eq. (4), heterogeneous label contexts are jointly modeled within a unified matrix factorization framework. This formulation can be detailed explained as follows. (1). For relational context, the EMF preserves the label proximity with the replicated softmax loss (Hinton and Salakhutdinov 2009); (2). For descriptive context, the matrix factorization preserves the label relations implied by the label-description associations. (3). A shared label embedding variable is introduced to achieve consistency for the joint factorization, which contributes to the alignment of the contexts. (4). An indicator matrix is adopted to indicate the partial contexts in PHCLE. Based on matrix completion theory, the factorization can handle missing values (Hu et al. 2013). Consequently, partially observed contexts can still be aligned with the proposed PHCLE framework.

Conclusively, PHCLE exploits the label relatedness conveyed in heterogeneous context by tailor-made formulas and achieves contexts alignment via a shared embedding principle which also works in partial contexts setting. In this way, the proposed PHCLE framework simultaneously tackles the




Figure 2: Comparison between TADW and PHCLE. D and A are heterogeneous contexts, W is the learned label embedding, and the others (H , C and U) are auxiliary matrices.

three challenges in the multi-context label embedding problem, resulting in label embeddings that well preserves the label relatedness conveyed in the heterogeneous contexts.

Model Comparison and Generalization

One nice property of PHCLE is that it overcomes the partial context problem which can not be solved by existing methods. Furthermore, it is also flexible to incorporate more contexts into our PHCLE model.

Comparison with existing methods

There are some attributed graph embedding approaches (Yang et al. 2015; Huang, Li, and Hu 2017; Pan et al. 2018; Shen et al. 2018) that can be adapted to handle heterogeneous contexts. However these methods cannot handle the partial context problem due to the full context correspondence requirement in their formulation.

TADW (Yang et al. 2015), for example, learns the embedding W from two heterogeneous contexts D and A within a single matrix factorization formula in Eq. (5).

$$\min_{W, H} \|D - W^T H A\|_F^2 + \frac{\lambda}{2} (\|W\|_F^2 + \|H\|_F^2), \quad (5)$$

Obviously, for a label without descriptive contexts, the relational contexts for that label are also dropped. In contrast, in PHCLE, heterogeneous contexts are modeled with a tailor-made formula, with a shared label embedding matrix to capture their dependency (as shown in Figure 2). Consequently, the two contexts are independent given the label embedding matrix. Thus, our method is more general and flexible to handle partial heterogeneous contexts settings.

Model generalization

Due to the additive property of matrix factorization, our PHCLE model can be easily generalized to jointly embed multiple label contexts which belongs to these two heterogeneous categories. Assume there are a total of n relational label contexts, with the constructed co-occurrence matrix denoted as

$D^{(i)}$, and m descriptive label contexts, each denoted as $A^{(j)}$. $C^{(i)}$ and $U^{(j)}$ denotes the corresponding label context embedding and descriptive contexts embedding, respectively. The Generalized Partial Heterogeneous-Context Label Embedding model (GPHCLE) can be formulated as Eq. (6).

$$\begin{aligned} \min_{C^{(i)}, W, U^{(j)}} & \sum_{i=1}^n \alpha^{(i)} EMF(D^{(i)}, C^{(i)T} W) \\ & + \sum_{j=1}^m \beta^{(j)} \|A^{(j)} - W^T U^{(j)}\|_F^2 + \Omega(W, C^{(i)}, U^{(j)}) \\ \text{s.t. } & \alpha^{(i)} \geq 0, \alpha^T \mathbf{1}_n = 1, \beta^T \mathbf{1}_m = 1, \end{aligned} \quad (6)$$

where $\alpha = [\alpha^{(1)}, \alpha^{(2)}, \dots, \alpha^{(i)}]$ controls the weight for relational contexts and $\beta = [\beta^{(1)}, \beta^{(2)}, \dots, \beta^{(j)}]$ is the importance factor for the descriptive contexts, $\Omega(W, C^{(i)}, U^{(j)})$ represents the regularizers.

Optimization Algorithm

Given the above formulation, we propose a gradient descent based alternating minimization algorithm to optimize the proposed PHCLE model. For simplicity, we denote the objective in Eq. (4) as $\mathcal{L}(C, W, U)$ in the following part.

Specifically, C and W are optimized through SGD. The gradients of $\mathcal{L}(C, W, U)$ with respect to C and W are

$$\frac{\partial \mathcal{L}}{\partial C} = (\mathbb{E}_{D'|C^T W} D' - D) W^T, \quad (7)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial W} &= C (\mathbb{E}_{D'|C^T W} D' - D) + \lambda_1 (-U A^T \odot I + U U^T W \odot I) \\ &+ \lambda_3 W, \end{aligned} \quad (8)$$

where $\mathbb{E}_{D'|C^T W} D' = Q^T \cdot \frac{1}{1 + \exp(-C^T W)}$ (Li et al. 2015b).

The subproblem with respect to U is an Elastic Net (Zou and Hastie 2005) problem defined as

$$\min_U \lambda_1 \|I \odot (A - W^T U)\|_F^2 + \frac{\lambda_3}{2} \|U\|_F^2 + \lambda_2 \|U\|_1, \quad (9)$$

Due to the sparsity constraint on U , Eq. (9) can not be directly optimized through SGD. Instead, we adopt FISTA (Beck and Teboulle 2009) algorithm to update U . Let $g(U) = \frac{\lambda_3}{2} \|U\|_F^2 + \lambda_2 \|U\|_1$ and $f(U) = \lambda_1 \|I \odot (A - W^T U)\|_F^2$, then quadratic approximation of Eq. (9) at given point Z is

$$Q_\tau(U, Z) := f(Z) + \langle U - Z, \nabla f(Z) \rangle + \frac{\tau}{2} \|U - Z\|_F^2 + g(U),$$

which admits a unique minimizer as

$$p_\tau(Z) = \arg \min_U \left\{ \frac{\tau}{2} \|U - K\|_F^2 + g(U) \right\}$$

where $\tau > 0$ is a constant as stepsize and $K = Z - \frac{1}{\tau} \nabla f(Z)$.

Details of the alternating minimization of Multi-context Label Embedding are summarized in Algorithm 1. FISTA optimization for U is shown in Algorithm 2. Specifically,

$L(f)$ is the Lipschitz constant of the gradient ∇f , and is set as $L(f) = 2\lambda_{\max}(W_i^T W_i)$ in our algorithm. Stopping condition of the algorithm is set as: $\frac{F(\hat{U}) - F(\hat{U}^*)}{F(\hat{U}^*)} < \epsilon$, where ϵ is a small tolerance value ($\epsilon = 0.0001$ in our experiment).

Convergence Analysis The proposed PHCLE is formulated as joint matrix factorization. Regarding to each variable, C , W and U , the sub-optimization problem are all convex. Fixing the other variables, convergence of the alternating optimization for each variable is guaranteed (Li et al. 2015a; Wang et al. 2017). Therefore, the objective function will converge to the local minimum accordingly.

Optimization for GPHCLE For the optimization of the generalized model in Eq. (6), if the weights for the contexts in each category are all predefined, GPHCLE can be directly optimized with the proposed algorithm. If α and β are to be learned, the self-weighted mechanism in (Nie, Li, and Li 2017) can be adopted in each sub-optimization problem in Algorithm 1. As it is not the focus, we omit the details here.

Algorithm 1: Alternating Minimization for PHCLE

Require: co-occurrence matrix D , label-attribute association matrix A , step-size η , Maximum iteration number K , trade-off factors $\lambda_1, \lambda_2, \lambda_3$ and λ_4
Ensure: C_K, W_K, U_K
Initialize C_0, W_0, V_0 to matrix of ones;
while $i < K$ **do**
 repeat $C_i = C_{i-1} - \eta \frac{\partial L}{\partial C}$ (See Eq. (7))
 repeat $W_i = W_{i-1} - \eta \frac{\partial L}{\partial W}$ (See Eq. (8))
 Update U using FISTA (Algorithm 2)
 $i = i + 1$;
end

Experiments

In this section, we first evaluate the label embeddings with zero-shot image classification task. Then, we demonstrate the label interpretability of PHCLE with two tasks: label similarity analysis and novel image understanding.

Experiment setup

Setup for PHCLE We learn task free label embeddings for the 1000 labels of ImageNet 2012 dataset (Russakovsky et al. 2015). Incorporated contexts are constructed as follows.

Relational contexts: We leverage the neighborhood structure characterized by the label hierarchy in WordNet (Miller et al. 1990) to construct the co-occurrence matrix D .

Descriptive contexts: We adopt the attributes given in the two commonly used attributed image datasets: Animals with Attributes (AWA) and Attribute Pascal and Yahoo (aPY), as the descriptive contexts. As not all of the 1000 labels have attributes, PHCLE is set with partial contexts.

Parameter setting: For PHCLE, we set $K = 50$, $d = 100$, $InnerMaxIter = 50$, $stepsize = 10^{-5}$. The trade-off parameters are set via “grid search” (Wang et al. 2017) from

Algorithm 2: FISTA for updating U

Require: A : the label-description association matrix;
 $I \odot W_{i-1}$: label embedding obtained in last iteration;
 U_{i-1} : description embedding obtained in last iteration;
 L : the Lipschitz constant of ∇f ;
 $InnerMaxIter$: maximum iteration of the algorithm
Ensure: The optimal solution of \hat{U}
Initialize $\hat{U}_0 = U_{i-1}$, $Z_1 = \hat{U}_0$, $t_1 = 1$;
while $j < MaxIter$ **do**
 $\hat{U}_j = p_\tau(Z_j)$;
 $t_{j+1} = \frac{1 + \sqrt{1 + 4t_j^2}}{2}$;
 $Z_{j+1} = \hat{U}_j + (\frac{t_j - 1}{t_{j+1}})(\hat{U}_j - \hat{U}_{j-1})$;
 $j = j + 1$;
 if $\frac{F(\hat{U}) - F(\hat{U}^*)}{F(\hat{U}^*)} < \epsilon$ **break**;
end

$\{10^{-2}, 10^{-1}, 1, 10^1, 10^2\}$, for each baseline. The number of negative samples k is set to be 10 for EMF.

Baselines We select three single context label embeddings and four multi-context label embeddings as the baselines: (1) Attribute label embedding (**ALE**). We use the attribute annotations released with the datasets; (2) Word2Vec Label Embedding (**WLE**). We use the 500 dimensional word embedding vector trained on 5.4 billion words Wikipedia; (3) Hierarchy Label Embedding (**HLE**). We follow the setting in (Akata et al. 2015), and construct an 1000-dimension embedding for each word accordingly; (4) Concatenation (**CNC**). We use simple concatenation of ALE, WLE and HLE, as in (Akata et al. 2015); (5) CCA-fused label embedding (**CCA**). (6) **TADW**. We obtain TADW label embedding for labels with full context correspondence with the released codes. As TADW requires fully observed contexts, we specifically compare it to PHCLE with full context correspondence (**PHCLE_FC**) in the image classification experiment. Furthermore, for label similarity analysis, label embedding obtained by PHCLE without sparsity constraint (**PHCLE_NoSp**) is also compared to demonstrate the impact of discriminative contexts selection in PHCLE.

Image classification

We first apply PHCLE to zero-shot image classification task, where the label relatedness is critical for the performance. Specifically, we conduct experiments on the data that overlaps the 1000 ImageNet labels for the AWA (26) and aPY (22) datasets. Details of the adopted datasets are presented in Table 1. We adopt the ResNet features (Xian, Schiele, and Akata 2017) as the image features and apply three representative zero-shot learning methods, ESZSL (Romera-Paredes et al. 2015), ConSE (Norouzi et al. 2013) and SJE (Akata et al. 2015), (all with their default parameters) to all the embedding methods. We adopt the average per-class top-1 accuracy (Akata et al. 2016) as the performance metric.

Table 2 shows the result of this image classification task. PHCLE consistently outperforms baseline embeddings with

Table 1: Statistics for the adopted AWA and aPY datasets.

Dataset	Y	Y ^{tr}	Y ^{te}	Att	Training	Test	Dim
AWA	26	20	6	85	11090	4019	2048
aPY	22	17	5	32	6925	1333	2048

Table 2: The classification accuracy obtained with different label embeddings on AWA and aPY dataset.

Methods	Datasets					
	AWA			aPY		
	ESZSL	ConSE	SJE	ESZSL	ConSE	SJE
ALE	63.66	50.83	61.48	52.60	52.87	49.75
WLE	56.98	45.92	38.63	45.53	39.13	38.04
HLE	53.76	58.08	54.18	49.79	43.35	52.18
CNC	63.66	57.97	74.27	48.71	50.78	33.34
CCA	47.05	35.46	47.05	41.94	34.34	39.45
PHCLE	69.11	58.43	77.47	54.61	52.87	50.91

all three zero-shot learning methods on both datasets. It demonstrates that PHCLE captures better label relatedness compared with other methods. Specifically, CCA performs the worst among all the baselines. It verifies that the label relatedness is not well preserved in the CCA, as the context information apart from the principal directions is lost. In addition, CNC performs comparably to single context label embeddings on AWA dataset, its results on aPY are even worse than other methods. This indicates that the concatenation of multiple contexts may counteract the influence of each of them. The result of CNC on aPY dataset also indicates that context alignment is difficult in this task. The superior performance of PHCLE verifies its robustness in handling heterogeneous context embedding problem.

Table 3 compares our method (PHCLE.FC) with TADW in the full context setting. We can see that PHCLE.FC outperforms TADW on both datasets. In particular, it achieves almost twice the accuracy to TADW on the aPY dataset. It indicates that PHCLE.FC is superior in heterogeneous contexts alignment with the shared embedding principle. Furthermore, comparing PHCLE and PHCLE.FC in these two tables, we can see that PHCLE achieves higher accuracy than that of PHCLE.FC for most of the settings. It clearly indicates that in PHCLE, the labels with partial contexts help to improve the label embedding for labels with full contexts.

Label similarity and interpretability

To assess the efficiency of PHCLE in preserving label relatedness, we analyze the label similarity and interpretability of PHCLE and other baselines embeddings.

Label retrieval We first conduct label retrieval task for the labels with partial contexts in PHCLE. Specifically, given each query, we retrieve the top 5 labels regarding to the cosine similarity (Omer and Yoav 2014). As compared labels are lack of attribute annotations, ALE is not compared here.

As shown in Table 4, the retrieved labels of PHCLE are all highly relevant to the query *coffeepot*, as they share the functionality to contain drinks. For WLE, “chiffonier” and “washbasin” can be considered to be relevant to *coffeepot*, as they are all household items. But it is hard to explain the relevance of *fire screen* and *coffeepot* based on human

Table 3: The comparison of label embeddings with full contexts.

Methods	Datasets					
	AWA			aPY		
	ESZSL	ConSE	SJE	ESZSL	ConSE	SJE
TADW	50.18	56.62	40.84	28.13	25.21	23.48
PHCLE.FC	69.08	56.77	53.33	50.56	53.04	42.52

Table 4: Label retrieval results. The retrieved labels are listed in descending order. Highly relevant labels are marked with bold, weakly relevant labels are in normal font, and irrelevant ones are in italics.

Query label	PHCLE	WLE	HLE
coffeepot	teapot	chiffonier	cauldron
	cauldron	<i>fire screen</i>	teapot
	beaker	washbasin	barrel
	vase	chocolate sauce	bathtub
	coffee mug	<i>window shade</i>	bucket

interpretability. The superior performance of PHCLE over WLE indicates that PHCLE better captures the label relatedness owing to the label hierarchy adopted for the relational context. For HLE, *calderon* and *teapot* are also retrieved, but the other three labels are less relevant to *coffeepot* compared with PHCLE. This demonstrates the tailor-made formula performs better to capture label relatedness from relational contexts. The overall superior results of PHCLE indicate that it successfully captures label relatedness from multiple aspects and even labels with partial contexts benefit from the alignment of two heterogeneous contexts.

Clustering visualization We further test label similarities for labels with full contexts, based on the cosine similarity. Cluster visualizations for different label embeddings for AWA dataset are shown in Figure 3. The results are quite revealing in several ways.

- 1). For embeddings of single context (*i.e.*, ALE, HLE and WLE), HLE fails to capture the difference among different embeddings, as the off-diagonal elements show high correlations. The clustering of WLE is not balanced in size, making it difficult to distinguish labels within the big clusters. ALE seems to show superior interpretability among others. However, ALE clusters the *weasel* together with dogs and cats erroneously (in the 3rd cluster), as these animals share the common attributes, such as “without buckteeth” and “eating meat”. Compared with ALE, PHCLE successfully groups *weasel* with its family *skunk* due to its consideration of label hierarchy. Furthermore, many of the off-diagonal elements in ALE are in red color (high correlation coefficients), which indicates its low inter-cluster similarities compared with PHCLE.
- 2). For fused embeddings (CCA, CNC and PHCLE.NoSp), CCA fails to capture the difference among different labels. For CNC, all the dogs and cats (5, 26, 3, 17, 4, 6) are clustered into three different clusters, with two cats separated. This indicates that multiple contexts aligned with simple concatenation. Our PHCLE produces humanly interpretable clusterings, with the second cluster groups all different species of dogs and cats. For com-




Figure 3: Cluster visualization of the correlation matrix constructed via PHCLE and contrastive label embeddings. **The first column** shows the correlation matrices of multi-context label embedding baselines: CNC and CCA and PHCLE_NoSp. **The second column** is the correlation matrix of our PHCLE. **The third column** presents the correlation matrices of single-context label embedding baselines: HLE, ALE and WLE. Each index corresponds to a label which is listed on the right-hand side.

parison with PHCLE_NoSp, it is easy to see that PHCLE and PHCLE_NoSp share similar intrinsic grouping structure, which demonstrates the efficacy of the shared embedding principle for the alignment of the heterogeneous contexts. Furthermore, it is interesting to note that the correlations between the label embedding in PHCLE_NoSp are very weak, which verifies that the discriminative attribute selection imposed by the sparsity constraint contribute to better alignment of heterogeneous contexts.

Image understanding

As PHCLE well aligns the label relations and label descriptions, our PHCLE can be adopted to handle the novel image understanding task. Specifically, for an image which does not belong to any existing classes, we can describe it with relative labels and specific semantic descriptions.

We conduct experiments on two typical novel classes, centaur and jetski, in aPY dataset. Specifically, we adopt the image-semantic mapping of ESZSL to obtain the semantic embedding of the image W^* . Then, the related labels are retrieved in the existing label set; the image description A^* is obtained with W^*TU . For related labels, top-ranked labels whose cosine similarity account for 80% of the overall similarity are selected. The obtained similarities are then normalized to get the similarity percentage for each related labels. For attribute description, regarding the predicted value, the top 6 attributes are selected. Figure 4 illustrates the image understanding results for two images. The results are quite revealing for its good human interpretability.

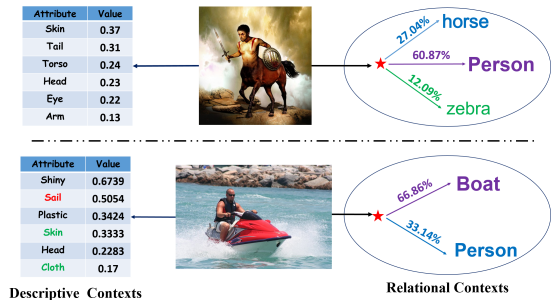


Figure 4: Image understanding with PHCLE. For relational contexts, the digits denotes the images' similarity with its related labels. For descriptive contexts, the attributes in green are creatively exploited, those in red are wrongly predicted.

- 1). For the image of a centaur (a mythological creature of half human and half horse), it is described as similar to human and horse, and have specific attributes, such as *skin*, *tail*, *Torso*, et.al. We can observe the result coincides with human's interpretation and the predicted attributes all overlaps the attributes ground truth of the image.
- 2). For the image of a person sitting on a Jetski (a recreational watercraft that the rider sits or stands on), it is described to be relative to boat and human. For the semantic description, attributes *Shiny* and *head*, are successfully predicted; *sail* is incorrectly predicted mainly

because the image regarded to be most similar to a boat, which also verifies the alignment for heterogeneous contexts in PHCLE. The most interesting is that based on human cognition, *skin* and *cloth* are reasonable to be attributes of the image, but it is not given in the ground truth. This verifies two points: a). human annotated attributes are noisy, making it reasonable to add the sparsity constraint for descriptive contexts selection for better heterogeneous contexts alignment. b). PHCLE achieves good attribute prediction ability due to the alignment of the heterogeneous contexts.

Conclusion and Future Work

In this paper, we provide a general Partial Heterogeneous Context Label Embedding (PHCLE) framework to solve the three challenges in multiple heterogeneous context label embedding problem. Specifically, we categorize the heterogeneous contexts into two groups, and tailor-make matrix factorization formulas to exploit the label relatedness for each group of contexts. Label relatedness conveyed in those contexts is selectively aligned in a shared space with a shared embedding principle. Due to this formulation, PHCLE can handle partial context problem with an indicator matrix to indicate the missing entries. It can also be easily generalized to incorporate more contexts. Experimental results demonstrate that label embedding obtained with PHCLE achieves superb performance in image classification task and exhibits good human interpretability.

As descriptive contexts exert huge impacts on relation analysis applications, such as social network analysis and recommendation, we will further study on exploiting label relations conveyed by the descriptive contexts with PHCLE in the future work.

Acknowledgments

This project is supported by the ARC Future Fellowship FT130100746, ARC LP150100671, DP180100106 and Chinese Scholarship Council (No.201706330075).

References

Akata, Z.; Reed, S.; Walter, D.; Lee, H.; and Schiele, B. 2015. Evaluation of output embeddings for fine-grained image classification. In *CVPR*, 2927–2936.

Akata, Z.; Perronnin, F.; Harchaoui, Z.; and Schmid, C. 2016. Label-embedding for image classification. *TPAMI* 38(7):1425–1438.

Andrew, G.; Arora, R.; Bilmes, J.; and Livescu, K. 2013. Deep canonical correlation analysis. In *ICML*, 1247–1255.

Beck, A., and Teboulle, M. 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *JIS* 2(1):183–202.

Fu, Y.; Hospedales, T.; Xiang, T.; Fu, Z.; and Gong, S. 2014. Transductive multi-view embedding for zero-shot recognition and annotation. In *ECCV*, 584–599.

Hinton, G., and Salakhutdinov, R. 2009. Replicated softmax: an undirected topic model. In *NIPS*, 1607–1614.

Hu, Y.; Zhang, D.; Ye, J.; Li, X.; and He, X. 2013. Fast and accurate matrix completion via truncated nuclear norm regularization. *TPAMI* 35(9):2117–2130.

Huang, X.; Li, J.; and Hu, X. 2017. Accelerated attributed network embedding. In *Proceedings of the 2017 SIAM International Conference on Data Mining*, 633–641. SIAM.

Koren, Y.; Bell, R. M.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer* 42(8):30–37.

Lampert, C.; Nickisch, H.; and Harmeling, S. 2009. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 951–958.

Li, X.; Liao, S.; Lan, W.; Du, X.; and Yang, G. 2015a. Zero-shot image tagging by hierarchical semantic embedding. In *SIGIR*, 879–882.

Li, Y.; Xu, L.; Tian, F.; Jiang, L.; Zhong, X.; and Chen, E. 2015b. Word embedding revisited: A new representation learning and explicit matrix factorization perspective. In *IJCAI*, 3650–3656.

Liao, L.; He, X.; Zhang, H.; and Chua, T. 2017. Attributed social network embedding. *arXiv:1705.04969*.

Mensink, T.; Gavves, E.; and Snoek, C. G. M. 2014. COSTA: co-occurrence statistics for zero-shot classification. In *CVPR*, 2441–2448.

Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, 3111–3119.

Miller, G.; Beckwith, R.; Fellbaum, C.; Gross, D.; and Miller, K. 1990. Introduction to wordnet: An on-line lexical database. *IJL* 3(4):235–244.

Nie, F.; Li, J.; and Li, X. 2017. Self-weighted multiview clustering with multiple graphs. In *IJCAI*, 2564–2570.

Norouzi, M.; Mikolov, T.; Bengio, S.; Singer, Y.; Shlens, J.; Frome, A.; Corrado, G.; and Dean, J. 2013. Zero-shot learning by convex combination of semantic embeddings. *arXiv:1312.5650*.

Omer, L., and Yoav, G. 2014. Dependency-based word embeddings. In *ACL*, 302–308.

Pan, S.; R.Hu; Long, G.; Jiang, J.; Yao, L.; and Zhang, C. 2018. Adversarially regularized graph autoencoder for graph embedding. In *IJCAI*, 2609–2615.

Romera-Paredes, B.; Torr, P.; ; and and. 2015. An embarrassingly simple approach to zero-shot learning. In *ICML*, 2152–2161.

Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A.; and Fei-Fei, L. 2015. Imagenet large scale visual recognition challenge. *IJCV* 115(3):211–252.

Shen, X.; Pan, S.; Liu, W.; Ong, Y.; and Sun, Q. 2018. Discrete network embedding. In *IJCAI*, 3549–3555.

Siddiquie, B.; Feris, R. S.; and Davis, L. S. 2011. Image ranking and retrieval based on multi-attribute queries. In *CVPR*, 801–808.

Tu, C.; Liu, H.; Liu, Z.; and Sun, M. 2017. CANE: context-aware network embedding for relation modeling. In *ACL*, 1722–1731.

Wah, C., and Belongie, S. 2013. Attribute-based detection of unfamiliar classes with humans in the loop. In *CVPR*, 779–786.

Wang, S.; Aggarwal, C.; Tang, J.; and Liu, H. 2017. Attributed signed network embedding. In *CIKM*, 137–146.

Xian, Y.; Schiele, B.; and Akata, Z. 2017. Zero-shot learning—the good, the bad and the ugly. *arXiv:1703.04394*.

Yang, C.; Liu, Z.; Zhao, D.; Sun, M.; and Chang, E. Y. 2015. Network representation learning with rich text information. In *IJCAI*, 2111–2117.

Zou, H., and Hastie, T. 2005. Regularization and variable selection via the elastic net. *JRSS* 67(2):301–320.