Semi-Supervised Review-Aware Rating Regression (Student Abstract)

Xiangkui Lu, Jun Wu†

School of Computer and Information Technology, Beijing Jiaotong University, Beijing 100044, China
{luxkui, wuj}@bjtu.edu.cn

Abstract

Semi-supervised learning is a promising solution to mitigate data sparsity in review-aware rating regression (RaRR), but it bears the risk of learning with noisy pseudo-labelled data. In this paper, we propose a paradigm called co-training-teaching (CoT^2), which integrates the merits of both co-training and co-teaching towards the robust semi-supervised RaRR. Concretely, CoT^2 employs two predictors and each of them alternately plays the roles of "labeler" and "validator" to generate and validate pseudo-labelled instances. Extensive experiments show that CoT^2 considerably outperforms state-of-the-art RaRR techniques, especially when training data is severely insufficient.

Introduction

Review-aware Rating Regression (RaRR) plays a key role in capturing users' preferences by analyzing ratings accompanied by reviews that provide explanations of users' scores. However, such multi-modality interactions are extremely sparse. Thus, RaRR suffers from a more severe challenge of data sparsity as compared to conventional recommendation tasks. The observed user-item ratings can be regarded as labelled data while the unobserved ones are unlabelled data. Although labelled data is expensive to obtain, unlabelled data is abundantly available and can be explored for recommendation. Towards this direction, some studies (Zhang et al. 2014; Huang, Luo, and Wu 2021) are conducted on developing semi-supervised rating regression techniques based on the co-training paradigm (Blum and Mitchell 1998). Despite enjoying many advantages, current co-training style schemes bear the risk of learning from noisy labels introduced by the pseudo-labelling operation.

Fortunately, co-teaching (Han et al. 2018) has shown the advantage in learning from noisy labels. We argue that co-training and co-teaching are mutually beneficial. For one thing, co-training fails to handle noisy pseudo-labels, while co-teaching can serve as a supplementary to clean the pseudo-labelled data. For another thing, the performance of co-teaching is bounded by the amount of reliable labelled data, and such a bottleneck might be broken by exploiting unlabelled data if cooperated with co-training. More importantly, co-training and co-teaching share the same data processing flow, which makes it possible to unify them towards robust semi-supervised learning. Specifically, both co-training and co-teaching maintain two predictors which communicate with each other in an iterative learning process. The major difference between them is that, co-training exchanges pseudo-labelled data between two predictors, while co-teaching exchanges those reliable ones sampled from noisy labelled data.

Motivated by the complementarity between co-training and co-teaching, we design a simple yet effective paradigm dubbed **Co-T**raining-**T**eaching (CoT^2) by integrating the merits of co-training and co-teaching towards robust semi-supervised RaRR. Benefiting from the capability of exploiting unlabelled data safely, CoT^2 considerably outperforms several state-of-the-art review-aware recommendation schemes on three benchmarks.

The Proposed Paradigm

Let $\mathcal{L} = \{(\mathbf{x}_l, y_l) | y_l \in [0,1]^1\}_{l=1}^{\Omega}$ be the set of labelled data, while $\mathcal{U} = \{(\mathbf{x}_l, y_l) | y_l = null\}_{l=\Omega+1}^{m\times n}$ denotes the set of unlabelled data, where Ω is the number of the observed ratings between m users and n items. A feature vector $\mathbf{x}_l = [\mathbf{o}_u, \mathbf{o}_i, \mathbf{f}_u, \mathbf{f}_i] \in \mathbb{R}^d$, where $\mathbf{o}_u$ (or $\mathbf{o}_i$) and $\mathbf{f}_u$ (or $\mathbf{f}_i$) denote the one-hot ID and the textual features of user u (or item i).

Our idea is to maintain two predictors simultaneously, both of which are initially learned from $\mathcal{L}$ and then reinforced to each other by seeking reliable pseudo-labelled instances from $\mathcal{U}$ during the subsequent iterations. Concretely, each predictor (labeler) first labels unlabelled data for its peer predictor (validator); then the validator assesses the reliability of pseudo-labelled data and samples a set of reliable instances to the labeler for updating the parameters. Repeat the above labelling and validating processes until the stop conditions are satisfied. The final prediction is made by averaging the outputs of both refined predictors. There are two key questions for designing CoT^2, which are not well considered by the conventional co-training paradigm.

How do we initialize two predictors with large diversity? This question is connected with a key requirement of co-training style approaches, that is, the initial predictors

† Corresponding author.

[1] Assume that all ratings are rescaled in the interval $[0, 1]$.

Methods	MusicalInstruments			OfficeProducts			VideoGames		
	8:2	6:4	4:6	8:2	6:4	4:6	8:2	6:4	4:6
HFT	0.842	0.917	1.102	0.778	0.826	0.856	1.127	1.159	1.226
AHN	0.758	0.784	0.828	0.719	0.744	0.764	1.115	1.183	1.225
CSEL	0.767	0.846	0.888	0.738	0.743	0.766	1.114	1.145	1.189
CoFM	0.759	0.808	0.798	0.733	0.742	0.763	1.109	1.132	1.188
CoT^2	**0.724**	**0.755**	**0.775**	**0.717**	**0.730**	**0.752**	**1.093**	**1.130**	**1.178**

Table 1: Performance of proposed method compared with the state-of-the-art methods, where the best performance is boldfaced and $\%$ indicates the relative improvement of our method over the others in terms of MSE metric.

should be different. In this paper, we diversify two predictors from three aspects. First, we construct different feature sets for two predictors by applying different sets of reviews to generate textual features. Second, we independently conduct bootstrap on dataset $\mathcal{L}$ twice to generate two different training sets. Third, we initialize two predictors by Factorization Machines (Rendle 2010) with different parameters. After that, we train two predictors on different features with different training sets respectively. In this way, we can generate two initial predictors with large diversity.

How do we exploit unlabelled data safely for refining such two predictors? In order to make CoT^2 robust against noisy pseudo labels, we employ two predictors to validate the pseudo labels mutually. Concretely, each predictor $H_s(s \in \{0,1\})$ labels a set of unlabelled instances according to the predictive function $\hat{y}^{(s)} = H_s(\mathbf{x})$, and then delivers the pseudo-labelled set to its peer predictor H_{1-s}. After receiving a pseudo-labelled instances $(\mathbf{x}, \hat{y}^{(s)})$, the validator also makes its own prediction $\hat{y}^{(1-s)} = H_{1-s}(\mathbf{x})$ and computes the squared loss between the pseudo label and the predictive label by $\Delta(\mathbf{x}) = (\hat{y}^{(s)} - \hat{y}^{(1-s)})^2$. If $\Delta(\mathbf{x})$ is small, we claim that $\mathbf{x}$ is reliable, and vice versa. Thus, the validator selects a ratio of instances with the smallest $\Delta(\mathbf{x})$ from the pseudo-labelled set as the reliable ones and returns them to the labeler for parameter updating.

Experiments

We conduct experiments on Amazon[2] MusicalInstruments, OfficeProducts, and VideoGames datasets, each of which is randomly split into training set and testing set at the ratio of 8:2, 6:4, and 4:6. We compare CoT^2 with two review-only models (**HFT** (McAuley and Leskovec 2013), and **AHN** (Dong et al. 2020)), as well as two semi-supervised models using both reviews and unlabelled data (**CSEL** (Zhang et al. 2014) and **CoFM** (Huang, Luo, and Wu 2021)).

The quantitative results in terms of MSE are shown in Table 1. Overall, three semi-supervised solutions outperform two review-only methods in most cases, especially when the dataset is small or the proportion of training instances is low. This observation indicates that exploiting unlabelled data has the potential to boost recommendation accuracy. Among the three semi-supervised solutions, CoFM and CoT^2 consistently outperform CSEL in all settings, which reveals that exploiting unlabelled data safely is significant to semi-supervised regression. Between the two semi-supervised solutions with confidence validation, CoT^2 demonstrates consistent improvements over CoFM. Such a considerable performance gap demonstrates the superiority of CoT^2 and the advantage of unifying co-training with co-teaching for robust semi-supervised regression. That is, the motivation of this work has been empirically verified.

[2]http://jmcauley.ucsd.edu/data/amazon/

Conclusion

By combining the dual advantages of co-training and co-teaching, we have proposed a novel CoT^2 paradigm. In essence, it is a generic framework of robust semi-supervised learning. Therefore, deploying CoT^2 in other applications is well worth studying in the future.

Acknowledgements

This work was supported in part by the Open Research Fund from the Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese University of Hong Kong, Shenzhen, under Grant B10120210117-OF01, and in part by the Fundamental Research Funds for the Central Universities, under Grant 2022JBZY019.

References

Blum, A.; and Mitchell, T. M. 1998. Combining Labeled and Unlabeled Data with Co-Training. In *COLT*, 92–100.

Dong, X.; Ni, J.; Cheng, W.; Chen, Z.; Zong, B.; Song, D.; and et. al. 2020. Asymmetrical Hierarchical Networks with Attentive Interactions for Interpretable Review-Based Recommendation. In *AAAI*, 7667–7674.

Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I. W.; and Sugiyama, M. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, 8536–8546.

Huang, J.; Luo, F.; and Wu, J. 2021. Semi-supervised Factorization Machines for Review-Aware Recommendation. In *DASFAA*, volume 12683, 85–99.

McAuley, J. J.; and Leskovec, J. 2013. Hidden factors and hidden topics: understanding rating dimensions with review text. In *RecSys*, 165–172.

Rendle, S. 2010. Factorization Machines. In *ICDM*, 995–1000.

Zhang, M.; Tang, J.; Zhang, X.; and Xue, X. 2014. Addressing cold start in recommender systems: a semi-supervised co-training algorithm. In *SIGIR*, 73–82.