

VISION-LANGUAGE ASSISTED ATTRIBUTE LEARNING

Kongming Liang¹, Xinran Wang¹, Rui Wang², Donghui Gao², Ling Jin², Weidong Liu²,
Xiatian Zhu³, Zhanyu Ma^{1*}, Jun Guo¹

¹School of Artificial Intelligence, Beijing University of Posts and Telecommunications

²China Mobile Research Institute

³University of Surrey, UK

ABSTRACT

Attribute labeling at scale is typically incomplete and partial, posing significant challenges to model optimization. Existing attribute learning methods often treat the missing labels as negative or simply ignore them all during training, either of which could hamper the model performance to great extent. To overcome these limitations, in this paper we leverage the available vision-language knowledge to explicitly disclose the missing labels for enhancing model learning. Given an image, we predict the likelihood of each missing attribute label assisted by an off-the-shelf vision-language model, and randomly select to ignore those with high scores in training. Our strategy strikes a good balance between fully ignoring and negatifying the missing labels, as these high scores are found to be informative on revealing label ambiguity. Extensive experiments show that our proposed vision-language assisted loss can achieve state-of-the-art performance on the newly cleaned VAW dataset. Qualitative evaluation demonstrates the ability of the proposed method in predicting more complete attributes.

Key Words— Attribute Learning, Object Understanding, Partial Annotations

1. INTRODUCTION

Visual attributes (e.g., color, shape, texture, part) [1, 2, 3] are an important means of describing the fine-grained differences between objects. Learning to predict attributes underpins many downstream applications such as image captioning [4], visual question answering (VQA) [5], image generation [6, 7], and compositional zero-shot learning [8, 9]. Since the space of visual attributes is often large, it is exhaustive for the annotators to label out all the attributes per image in practice. For example, Pham et al. [10] constructed a large-scale attribute dataset (VAW), with 2260 objects and 620 attributes. On average, only 3.56 attributes are annotated for each image, largely incomplete and partial (see Fig. 1). This problem poses significant challenges to model optimization, which however is still understudied in the literature [11].



Fig. 1: Illustration of typical partial attribute labeling.

Typically, the number of missing/unannotated attribute labels is even more than that of labeled ones. Thus, dealing with the missing attributes is critical. There are three existing representative strategies. The first is discarding the loss terms with missing attributes, termed as *ignoring* [12]. As only a small portion of labels are used for model training, the performance is naturally limited. The second simply treats all the missing labels as negative, termed as *negatifying* [13]. This however would introduce a plenty of false labels to those missing attributes. The third aims to combine the above two strategies by *selecting* one for each attribute [14, 15]. Nonetheless, this method falls into *the chicken or the egg causality dilemma* given that predicting the attributes and training the attribute model (the two key components) are mutually depending on each other.

To overcome all the aforementioned limitations, we propose to leverage the available vision-language knowledge for enhancing attribute learning under partial labeling. Instead of struggling to predict the missing labels accurately, we set up to find out what attributes are ambiguous and thus need to be ignored during training. This is because, achieving this objective is less demanding and likely more achievable, without the notorious dilemma as suffered by [14]. Specifically, we first estimate all the missing attributes of a given training image assisted by an off-the-shelf vision-language model (e.g., CLIP [16]). Considering that such predictions serve as an indicator of the presence ambiguity, we take those with high probability

* indicates the corresponding author.

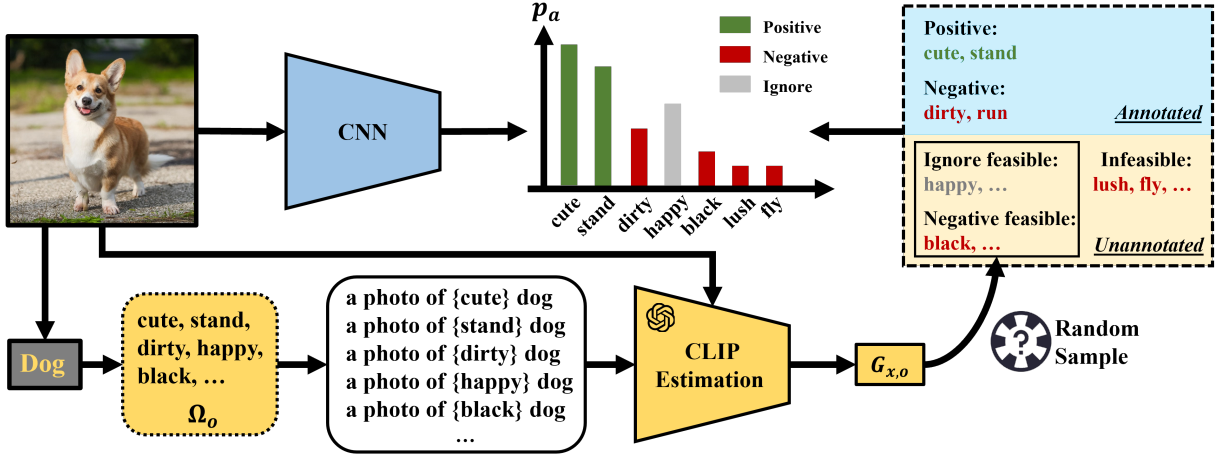


Fig. 2: Overview of the proposed vision-language assisted attribute learning. Given a training image with partial labeling, we exploit a vision-language model to estimate the presence probability of each attribute (Ω_o). As this prediction could be largely unreliable, we treat those attributes with high probabilities as ambiguity. For model optimization, we randomly ignore some of the ambiguous attributes and negate the rest missing ones.

scores as *ambiguous attributes*. To further tackle the predicting uncertainty, we randomly select some of these ambiguous attributes and ignore their loss contributions during training, whilst negating the remaining missing attributes. This strategy could yield a better trade-off between the ignoring and negating strategies. Experimental results demonstrate that the proposed vision-language assisted attribute learning method achieves the new state-of-the-art on the standard large attribute recognition benchmark.

2. METHOD

In this section, we start with the formalization of partial attribute learning. Then we leverage an off-the-shelf vision-language model to estimate the presence probability for each unannotated attribute. Finally, we design a selective loss by randomly ignoring the loss of the ambiguous attributes. The overall method is summarized in Fig. 2.

2.1. Attribute Learning

Attribute learning aims to predict the presence or absence of each pre-defined attribute for an input visual instance. Considering the limited annotation cost, the majority of the attributes are left unannotated which may hinder the learning process. For the visual instance x , we denote its corresponding object label as o and attribute vector as $y = \{y_a\}_{a=1}^A$, where A denotes the number of attributes. If the attribute a is present, absent or unannotated in the visual instance, y_a is represented as 1, -1, 0 respectively. Then the positive and negative attribute sets of x are constructed as $\mathcal{P}_x = \{a|y_a = 1\}$ and $\mathcal{N}_x = \{a|y_a = -1\}$. And the unannotated attribute set is denoted as $\mathcal{U}_x = \{a|y_a = 0\}$.

For an image x with its object label o , an attribute classifier f_a is trained to predict whether x possesses the attribute a or not. The prediction probability is denoted as $p_a = f_a(x, o; \theta_a)$. Then the general form of attribute learning loss can be formulated as follows,

$$\mathcal{L}(x, o) = \sum_{a \in \mathcal{P}_x} \mathcal{L}_a^+(x, o) + \sum_{a \in \mathcal{N}_x} \mathcal{L}_a^-(x, o) + \sum_{a \in \mathcal{U}_x} \mathcal{L}_a^u(x, o), \quad (1)$$

where $\mathcal{L}_a^+(x, o)$, $\mathcal{L}_a^-(x, o)$ and $\mathcal{L}_a^u(x, o)$ are the loss terms calculated by positive, negative, and unannotated attributes. In this paper, we use the asymmetric loss (ASL) [17] to mitigate the negative-positive imbalance problem. We denote the positive loss term as $\mathcal{L}_a^+(x, o) = (1 - p_a)^{\gamma^+} \log(p_a)$ and the negative loss term as $\mathcal{L}_a^-(x, o) = (p_a)^{\gamma^-} \log(1 - p_a)$. ASL can focus on the hard samples dynamically while controlling the contribution propagated from the positive and negative samples at the same time. For positive and negative loss, we use γ^+ and γ^- to denote their focusing parameters respectively.

2.2. Estimating the Unannotated Attributes

For the unannotated attribute labels, previous methods may ignore their loss by setting the $\mathcal{L}_a^u(x) = 0$ or treat them as negative by setting the $\mathcal{L}_a^u(x) = \log(1 - p_a)$. However, both methods above are inefficient in model training since ignoring mode only includes a portion of the labels, and negating mode introduces extra false labels as some missing attributes may be present. Another solution combines the above two strategies by selecting one for each attribute. Nonetheless, this method falls into the chicken or the egg causality dilemma as predicting the attributes and training the attribute model are mutually dependent on each other. In this paper, we use CLIP [16], a powerful vision-language model,

```

# NUM_IGNORE           = the size of ignore attribute set
# feasible_attributes  = the feasible attributes of o
# G_xo                 = the likelihood vector of x and o

ignore_attribute_set = np.random.choice(
    a=feasible_attributes,
    size=NUM_IGNORE,
    replace=False,
    p=G_xo
)

```

Fig. 3: Numpy-like pseudocode for the implementation of ignored attribute sampling.

to estimate the probability of each attribute being present in a given object instance. Specifically, we first calculate the object-attribute co-occurrence on the training data. Then, if there is any sample belonging to the object o and annotated with attribute a simultaneously, a will be appended to the feasible set Ω_o . Only the attributes contained in the feasible set are further processed by CLIP.

To use CLIP for unannotated attribute estimation, the object and attribute names are transformed into natural language prompts such as A photo of [attribute][object]. The prompts are further processed to get the embeddings for each word. Then, all the embeddings are passed through the text encoder to obtain the text representations. For the visual counterpart, we can extract the visual representations based on the image encoder. Given a visual instance x , we denote its visual representations from CLIP as x_v . And the text representations for attribute a and object o can be denoted as $t_{a,o}$. Then we compute the cosine similarities between the visual representations and the text representations to compute the presence probability of the attribute a :

$$g(a|x, o) = \frac{\exp(x_v \cdot t_{a,o}/\tau)}{\sum_{\hat{a} \in \Omega_o} \exp(x_v \cdot \hat{t}_{\hat{a},o}/\tau)}, \quad (2)$$

where τ is a temperature parameter and only the feasible attributes of the object o are included for probability estimation. The likelihood vector of feasible attributes can be denoted as $\mathbf{G}_{x,o} = \{g(a|x, o) \mid \forall a \in \Omega_o\}$.

2.3. Vision-language Guided Selective Loss

For partially labeled attribute learning, the number of unannotated attributes is much more than the number of positive and negative attributes (e.g. $|\mathcal{P}_x \cup \mathcal{N}_x| \ll |\mathcal{U}_x|$). Therefore, how to process the unannotated attributes is of great importance. With the aid of attribute likelihood provided by CLIP, we can estimate the probability of an unannotated attribute being present in a visual instance. The attributes with high probability scores are considered as ambiguous attributes. To further tackle the predicting uncertainty, we randomly select some of these ambiguous attributes and ignore their loss contributions during training, whilst negatifying the remaining

missing attributes. We formalize the above vision-language guided selective loss as follows,

$$\mathcal{L}_a^u(x, o) = \begin{cases} \omega_a(p_a)^{\gamma^\oplus} \log(1 - p_a) & \text{if } a \in \Omega_o \\ (p_a)^{\gamma^\ominus} \log(1 - p_a) & \text{if } a \notin \Omega_o \end{cases}, \quad (3)$$

where γ^\oplus and γ^\ominus denote the feasible and infeasible focusing parameters of the object o respectively. We set $\gamma^\oplus < \gamma^\ominus$ to decay the feasible term with a lower rate than the infeasible one since most infeasible attributes are easy negative samples. ω_a is used to ignore the suspected positive attributes. We generate random samples from the non-uniform distribution $\mathbf{G}_{x,o}(a)$ of feasible attributes to construct the ignore attribute set Ω_{Ignore} . In Fig.3, we include pseudocode for the core implementation of ignored attribute sampling. The parameter ω_a in Eqn. (3) can be defined as follows,

$$\omega_a = \begin{cases} 0 & \text{if } a \in \Omega_{Ignore} \\ 1 & \text{if } a \notin \Omega_{Ignore} \end{cases}. \quad (4)$$

The loss contributed by the ignored attribute is discarded during training.

3. EXPERIMENT

In this section, we first present the dataset and evaluation metrics for large-scale attribute learning. Next, we compare the proposed method with existing state-of-the-arts. Finally, we conduct ablation studies to show the hyperparameter effect and demonstrate the qualitative results.

3.1. Dataset and Evaluation Metric

We use the currently released VAW dataset [14] in the experiments. To overcome the negative impact of tiny objects, we discard the object instances that are smaller than 50×50 . Then we merge the objects and attributes based on their plurals and synonyms (e.g. *plane*, *airplane*, and *airplanes*). The cleaned version of VAW contains 170,407 visual instances, 1763 object categories, and 591 attributes. For evaluation, we follow the previous multi-label learning and attribute prediction methods to report the mAP for all attributes. Since the dataset is partially labeled, only the annotated attributes are included for evaluation [10].

3.2. Implementation Details

ResNet-50 is adopted as the backbone for attribute prediction. During training, we select Adam as the optimizer and set the initial learning rate as 10^{-5} for the pre-trained parameters and 7×10^{-4} for the newly added parameters. The model is optimized for 12 epochs with a batch size of 64. We set $\gamma^+ = 1, \gamma^- = 2, \gamma^\oplus = 4, \gamma^\ominus = 7$. For the unannotated label estimation, ViT-B/32 is set as the vision transformer of CLIP.

Table 1: Comparison with the state-of-the-art methods of partially labeled multi-label learning on the cleaned VAW dataset. ignoring means discarding the loss of all the unannotated labels. And *negatifying* means considering all the unannotated labels to be negative.

Methods	Overall (mAP)	Class imbalance (mAP)			Attribute types (mAP)							
		Head	Medium	Tail	Color	Material	Shape	Size	Texture	Action	State	Others
CE-ignoring	61.18	68.19	58.31	44.88	52.99	56.90	61.87	59.69	61.61	53.49	58.49	65.95
CE-negatifying [2]	61.37	69.61	58.41	40.91	60.82	66.42	61.03	58.60	62.42	53.43	57.60	62.87
ASL-ignoring	61.29	68.45	58.49	44.18	52.76	57.94	61.92	59.51	62.39	53.88	58.45	65.93
ASL-negatifying [17]	62.01	70.02	58.64	43.61	61.96	65.42	62.44	60.91	62.33	52.73	58.18	63.77
PSL [14]	62.28	69.83	59.47	43.82	57.80	64.91	62.74	60.35	63.55	52.87	58.57	65.36
Ours	64.06	71.38	61.67	45.16	63.87	66.36	63.55	60.46	64.18	57.30	61.01	65.73



Fig. 4: Qualitative results. We show the ground truth positive attribute labels and the Top-10 attributes predicted by the proposed method. The correct predictions are highlighted in green, the incorrect predictions are highlighted in red, and the correct predictions without annotation are highlighted in blue.

3.3. Comparison with State-of-the-art Methods

The proposed method is compared with multi-label cross entropy (CE) [2], asymmetric loss (ASL) [17], and Partial-ASL (PSL) [14] which are the state-of-the-art methods of partially labeled multi-label classification. We use the same backbone ResNet-50 for the above methods to make a fair comparison. Both *ignore* and *negative* modes for CE and ASL are reported in our experiments. The overall experiments are shown in Table 1. Compared to the other methods, our proposed vision-language guided selective loss achieves 64.06 mAP score and outperforms the state-of-the-art method (PSL). To consider the positive and negative label imbalance, we report the mAP scores of head, medium, and tail attributes. Similarly, our method can achieve better performance under different imbalance conditions. In addition, our proposed loss outperforms the other methods in terms of most of the attribute types. The above experiments consistently demonstrate the effectiveness of the proposed vision-language assisted attribute estimation procedure.

3.4. Ablation Study and Qualitative Results

The effectiveness of the number of ignored attributes is investigated on partially labeled attribute learning. As demonstrated in Table 2, the performance firstly increases as more attributes are involved in ignore set. Then the performance keeps stabilizing when the size of ignore set is larger than 30. If we regard all the attributes to be ignored, the overall perfor-

Table 2: The ablation study of the parameters in the vision-language guided selective loss. The results are shown for different values of NUM_IGNORE.

NUM_IGNORE	Overall (mAP)
10	63.70
20	63.96
30	64.06
40	63.81
50	63.83
ALL	59.41

mance decreases to 59.41, which is even lower than the baseline methods. We show some qualitative results in Fig. 4. The results show the proposed method can predict more complete attributes, though only part of them are manually labeled.

4. CONCLUSION

A vision-language assisted selective loss is proposed for partially labeled attribute learning. The proposed method can well estimate the presence probability of unannotated attributes by an off-the-shelf vision-language model and further randomly select to ignore those with high scores in training. The proposed loss achieves the best performance on the newly cleaned VAW dataset. The qualitative results demonstrate the proposed vision-language assisted selective loss can predict the attributes more completely.

5. ACKNOWLEDGEMENTS

This work is supported in part by National Natural Science Foundation of China (NSFC) No. 62106022, 62225601, U19B2036, and in part by Beijing University of Posts and Telecommunications-China Mobile Research Institute Joint Innovation Center, and in part by Program for Youth Innovative Research Team of BUPT.

6. REFERENCES

- [1] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth, “Describing objects by their attributes,” in *CVPR*. IEEE, 2009, pp. 1778–1785. [1](#)
- [2] Genevieve Patterson and James Hays, “Coco attributes: Attributes for people, animals, and objects,” in *ECCV*. Springer, 2016, pp. 85–100. [1, 4](#)
- [3] Kongming Liang, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen, “Unifying visual attribute learning with object recognition in a multiplicative framework,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 7, pp. 1747–1760, 2019. [1](#)
- [4] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al., “From captions to visual concepts and back,” in *CVPR*, 2015, pp. 1473–1482. [1](#)
- [5] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh, “Vqa: Visual question answering,” in *ICCV*, 2015, pp. 2425–2433. [1](#)
- [6] Bing Li, Yuanlue Zhu, Yitong Wang, Chia-Wen Lin, Bernard Ghanem, and Linlin Shen, “Anigan: Style-guided generative adversarial networks for unsupervised anime face generation,” *IEEE Transactions on Multimedia*, vol. 24, pp. 4077–4091, 2021. [1](#)
- [7] Bing Li, Chia-Wen Lin, Cheng Zheng, Shan Liu, Junsong Yuan, Bernard Ghanem, and C-C Jay Kuo, “High quality disparity remapping with two-stage warping,” in *CVPR*, 2021, pp. 2269–2278. [1](#)
- [8] Tushar Nagarajan and Kristen Grauman, “Attributes as operators: factorizing unseen attribute-object compositions,” in *ECCV*, 2018, pp. 169–185. [1](#)
- [9] Massimiliano Mancini, Muhammad Ferjad Naeem, Yongqin Xian, and Zeynep Akata, “Open world compositional zero-shot learning,” in *CVPR*, 2021, pp. 5222–5230. [1](#)
- [10] Khoi Pham, Kushal Kaffle, Zhe Lin, Zhihong Ding, Scott Cohen, Quan Tran, and Abhinav Shrivastava, “Learning to predict visual attributes in the wild,” in *CVPR*, 2021, pp. 13018–13028. [1, 3](#)
- [11] Kongming Liang, Yuhong Guo, Hong Chang, and Xilin Chen, “Incomplete attribute learning with auxiliary labels,” in *IJCAI*, 2017, pp. 2252–2258. [1](#)
- [12] Thibaut Durand, Nazanin Mehrasa, and Greg Mori, “Learning a deep convnet for multi-label classification with partial labels,” in *CVPR*, June 2019. [1](#)
- [13] Kaustav Kundu and Joseph Tighe, “Exploiting weakly supervised visual patterns to learn from partial annotations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 561–572, 2020. [1](#)
- [14] Emanuel Ben-Baruch, Tal Ridnik, Itamar Friedman, Avi Ben-Cohen, Nadav Zamir, Asaf Noy, and Lihi Zelnik-Manor, “Multi-label classification with partial annotations using class-aware selective loss,” in *CVPR*, June 2022, pp. 4764–4772. [1, 3, 4](#)
- [15] Kongming Liang, Xinran Wang, Tao Wei, Wei Chen, Zhanyu Ma, and Jun Guo, “Attribute learning with knowledge enhanced partial annotations,” in *ICIP*. IEEE, 2023, pp. 1715–1719. [1](#)
- [16] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al., “Learning transferable visual models from natural language supervision,” in *ICML*. PMLR, 2021, pp. 8748–8763. [1, 2](#)
- [17] Tal Ridnik, Emanuel Ben-Baruch, Nadav Zamir, Asaf Noy, Itamar Friedman, Matan Protter, and Lihi Zelnik-Manor, “Asymmetric loss for multi-label classification,” in *CVPR*, 2021, pp. 82–91. [2, 4](#)