

Domain Generalization with MixStyle


Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang

ICLR 2021




中国科学院深圳先进技术研究院
SHENZHEN INSTITUTES OF ADVANCED TECHNOLOGY
CHINESE ACADEMY OF SCIENCES

Problem: CNNs do not work well on out-of-distribution (OOD) data




cat
dog ✓
⋮
person
car

train distribution



test distribution

also known as the domain shift problem



cat
dog
⋮
person
car





We focus on domain generalization (DG)

- setup


train a model using multiple source domains

e.g., 3 source domains: photo, cartoon & art painting



test on an unseen domain

e.g., sketch




- problem

- DG is challenging (without accessing the target data)
- intuitively, more source domains -> more generalizable -> better performance
- however, collecting data of a large variety of domains is often costly or even impossible

Motivation: to increase the diversity of source domains in the feature space


- Visual domain is closely related to image style



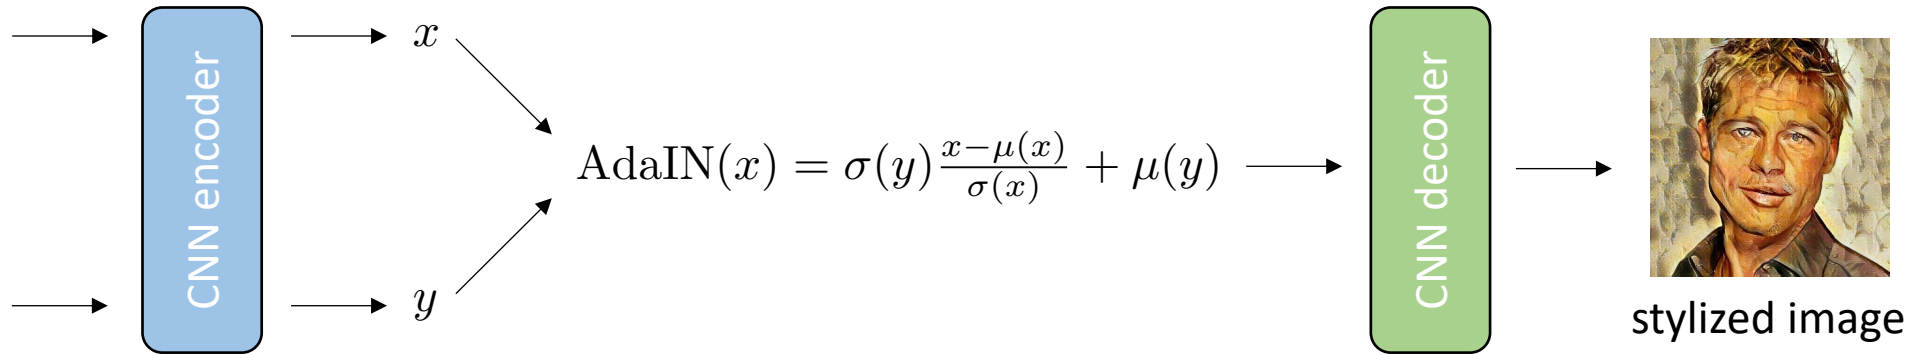
Motivation: to increase the diversity of source domains in the feature space

- CNN feature statistics (i.e., mean & std) can be used to manipulate image style (inspired by neural style transfer)

content image

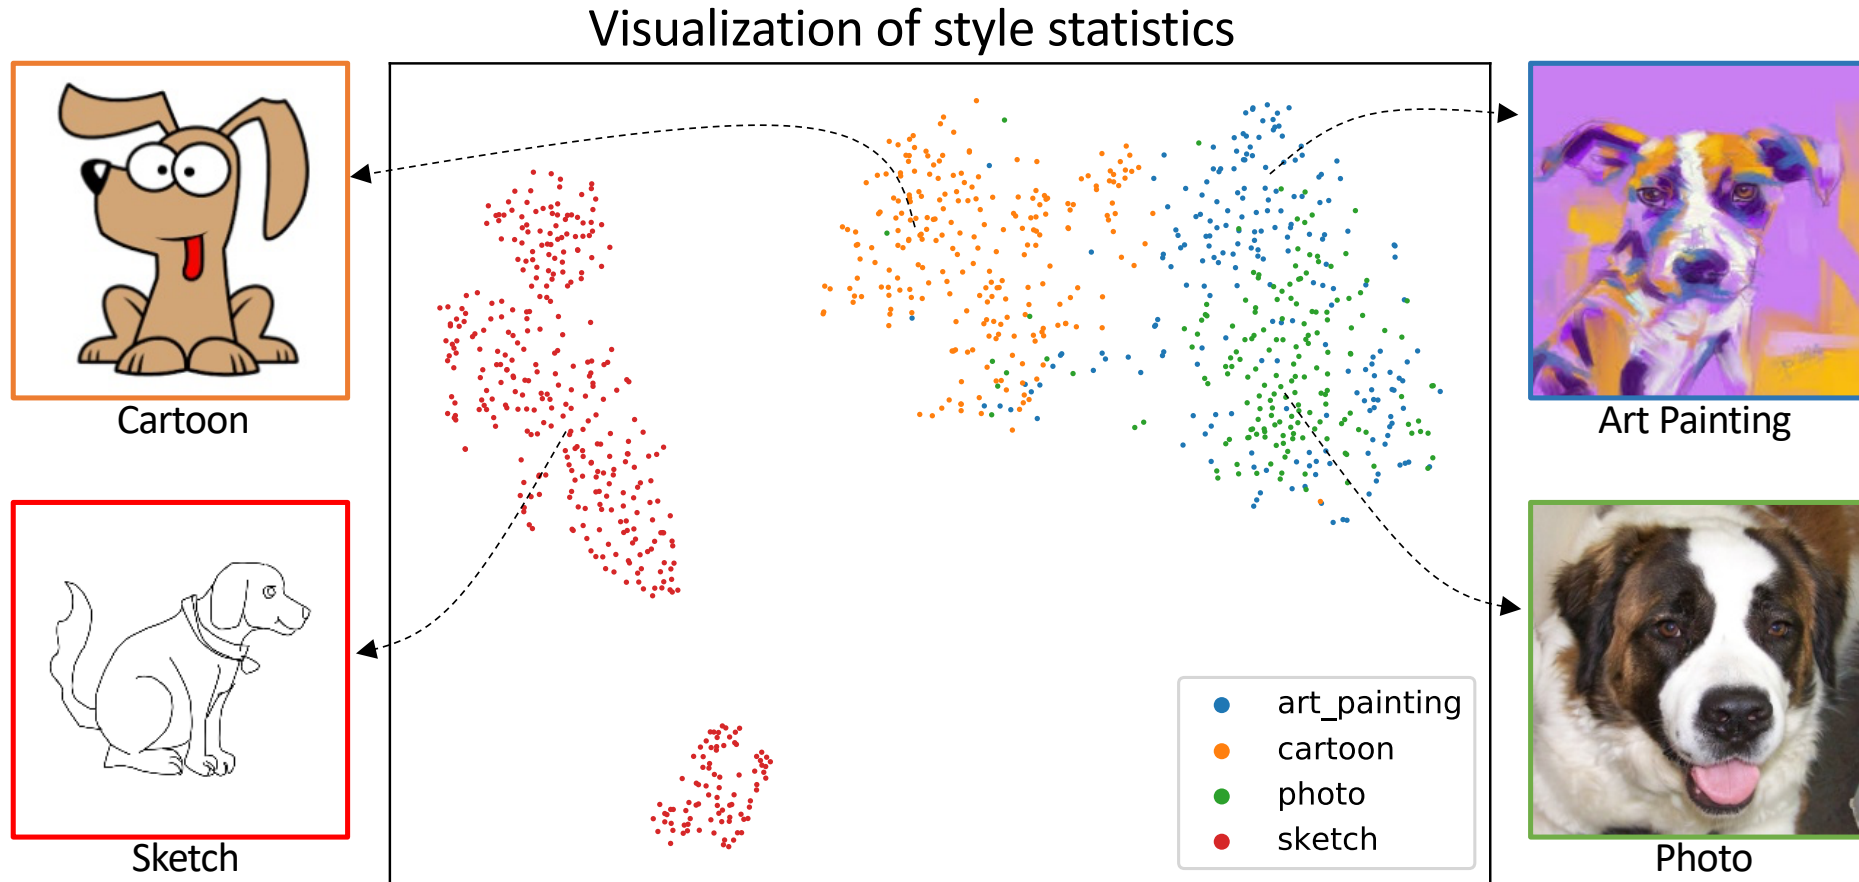


style image



Motivation: to increase the diversity of source domains in the feature space

- t-SNE visualization of feature statistics (a.k.a. style statistics)



Our method: MixStyle

$$\begin{aligned}\lambda &\sim \text{Beta}(\alpha, \alpha) \\ \gamma_{mix} &= \lambda \sigma(x) + (1 - \lambda) \sigma(\tilde{x}) \\ \beta_{mix} &= \lambda \mu(x) + (1 - \lambda) \mu(\tilde{x}) \\ \text{MixStyle}(x) &= \gamma_{mix} \frac{x - \mu(x)}{\sigma(x)} + \beta_{mix}\end{aligned}$$

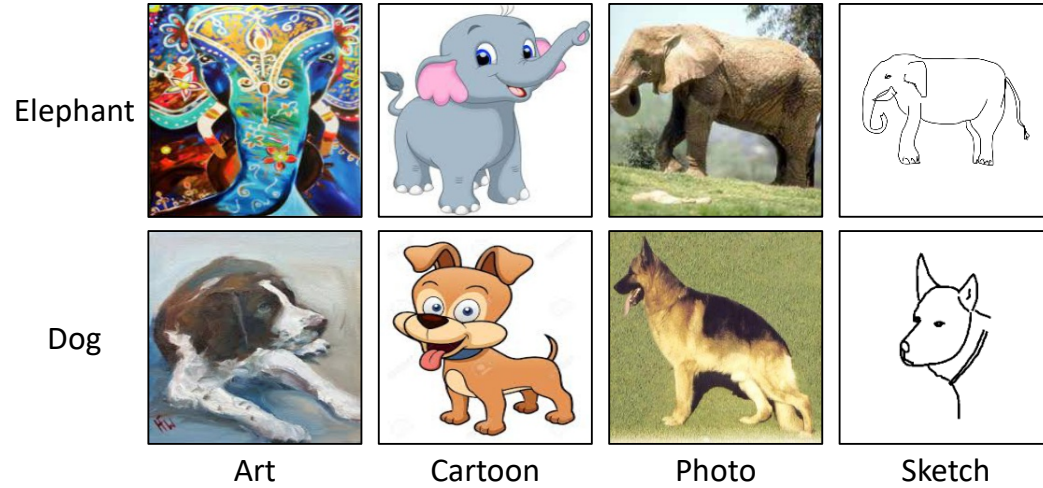
a random instance from the same mini-batch; or
an instance from a different domain (if domain
labels are provided)

MixStyle is inserted to multiple shallow layers in a CNN


```
def forward(self, x):
    x = self.conv1(x) # 1st convolution layer
    x = self.res1(x) # 1st residual block
    x = self.mixstyle(x)
    x = self.res2(x) # 2nd residual block
    x = self.mixstyle(x)
    x = self.res3(x) # 3rd residual block
    x = self.res4(x) # 4th residual block
    ...
```


MixStyle improves OOD generalization on these tasks

1. category classification



2. instance retrieval



3. reinforcement learning



more details about the results can be found in the paper:
<https://openreview.net/forum?id=6xHJ37MVxxp>

Thanks for your attention

interested in knowing more about the topic of domain generalization?

check out our latest survey paper at

<https://arxiv.org/abs/2103.02503> (Domain Generalization: A Survey)