Learning from Noisy Web Data with Category-level Supervision

Li Niu

Our method is built upon Variational Autoencoder (VAE) [1] .

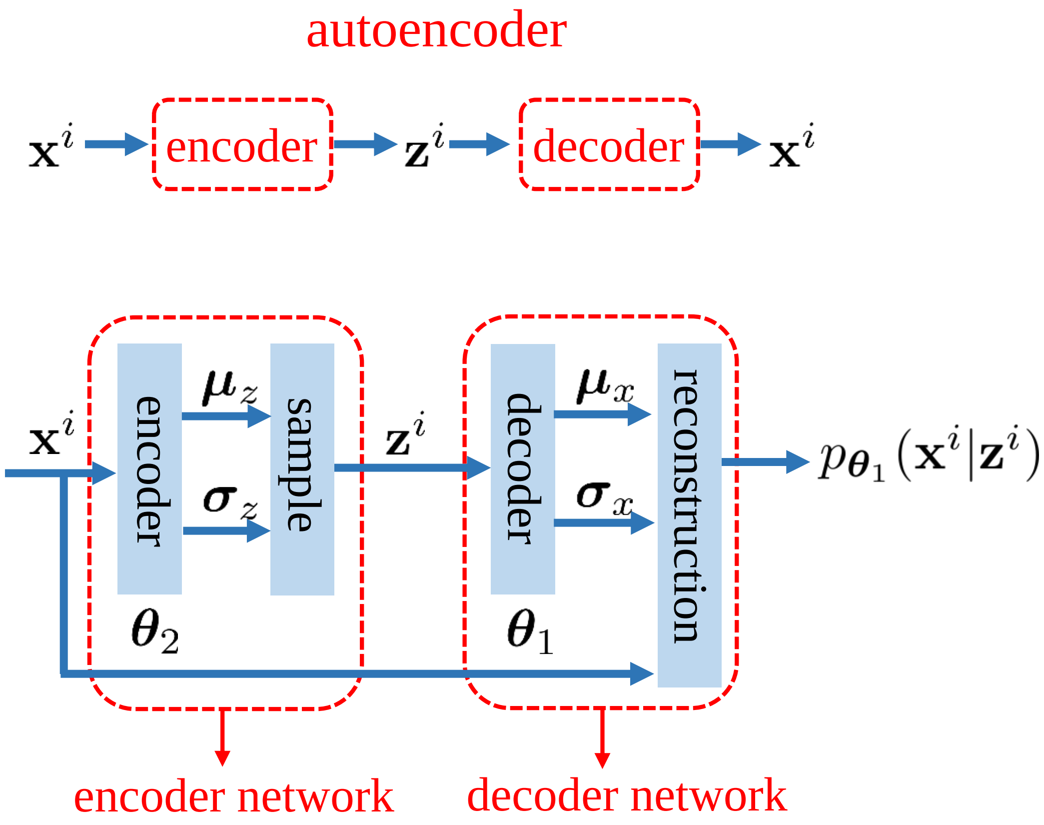

[1] Kingma and Welling, ICLR 2014

Our Method is Based on VAE

Our method is built upon Variational Autoencoder (VAE) [1] .

1. Autoencoder can be used for outlier detection.

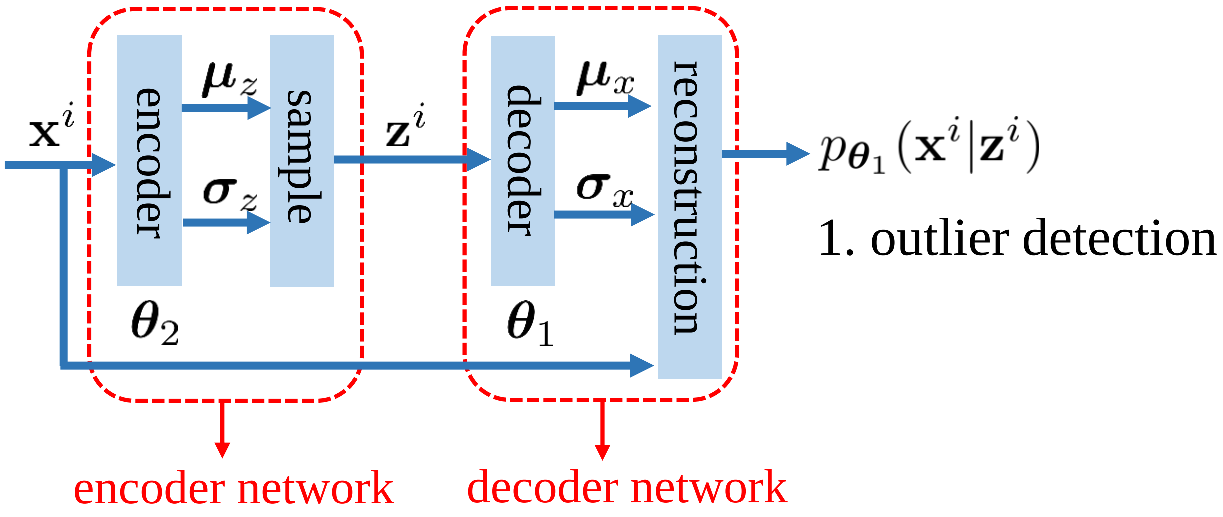

[1] Kingma and Welling, ICLR 2014

Our Method is Based on VAE

Our method is built upon Variational Autoencoder (VAE) [1] .

1. Autoencoder can be used for outlier detection.
2. Hidden layer of autoencoder can be injected with category-level information.

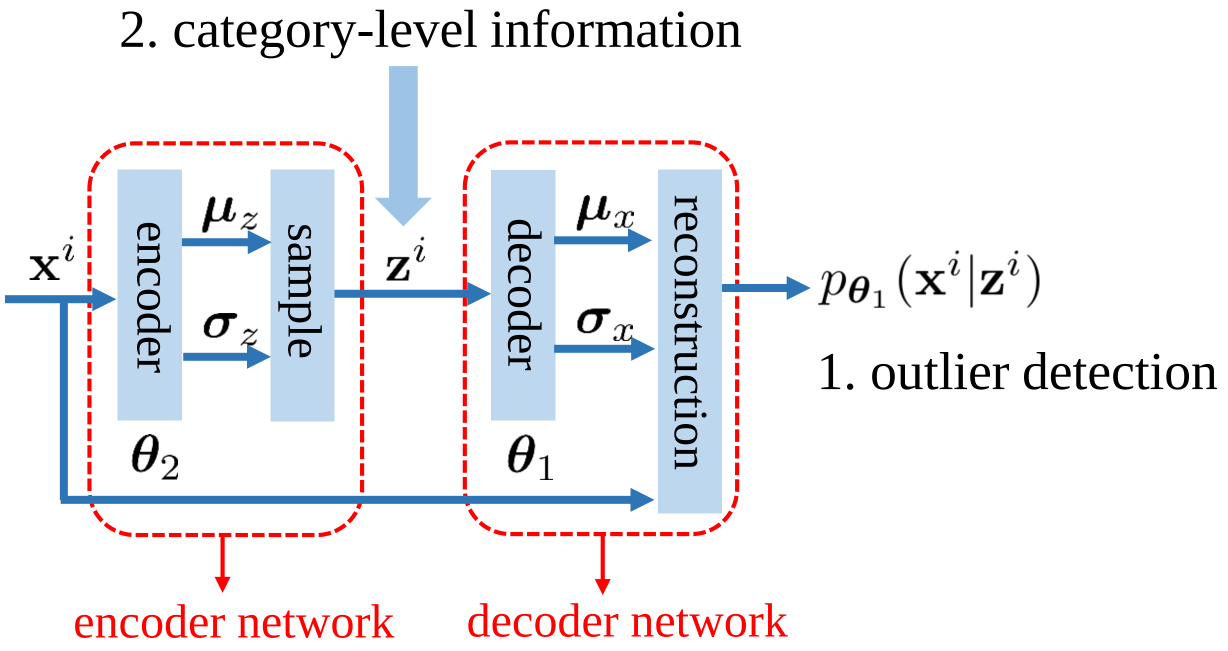

[1] Kingma and Welling, ICLR 2014

Our Method is Based on VAE

Our method is built upon Variational Autoencoder (VAE) [1] .

1. Autoencoder can be used for outlier detection.
2. Hidden layer of autoencoder can be injected with category-level information.

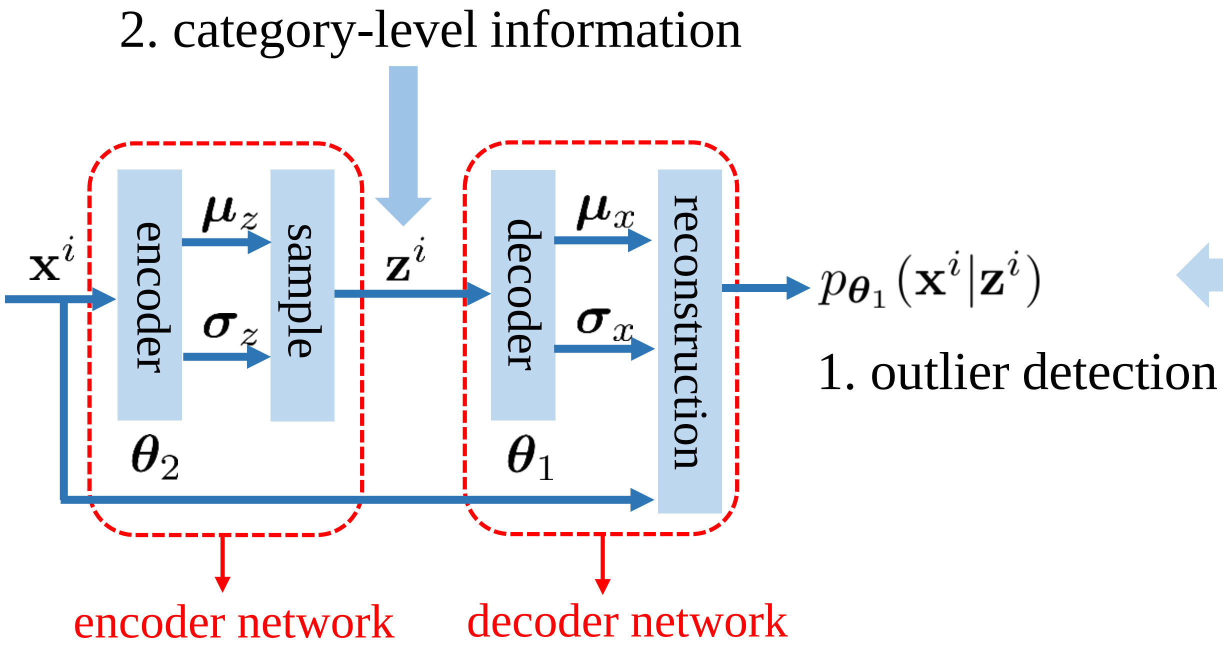

VAE

$$\min_{\boldsymbol{\theta}_1,\boldsymbol{\theta}_2} \underbrace{\boxed{\mathrm{KL}[q_{\boldsymbol{\theta}_2}(\mathbf{z}|\mathbf{x}^i)||p(\mathbf{z})]}}_{\text{regulate } \mathbf{z}} - \underbrace{\boxed{\log p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i)}}_{\text{reconstruct } \mathbf{x}}$$

[1] Kingma and Welling, ICLR 2014

Our Method to Handle Label Noise

VAE

$$\min_{\boldsymbol{\theta}_1,\boldsymbol{\theta}_2} \boxed{\mathrm{KL}[q_{\boldsymbol{\theta}_2}(\mathbf{z}|\mathbf{x}^i)||p(\mathbf{z})]} - \log p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i)$$

replace KL
with classifier

Our Method

$$\min_{\boldsymbol{\theta}_0,\boldsymbol{\theta}_1,\boldsymbol{\theta}_2} \boxed{-\log p(y^i = c^i|\mathbf{z}^i)} - \log p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i)$$

classification loss (1) — (2)

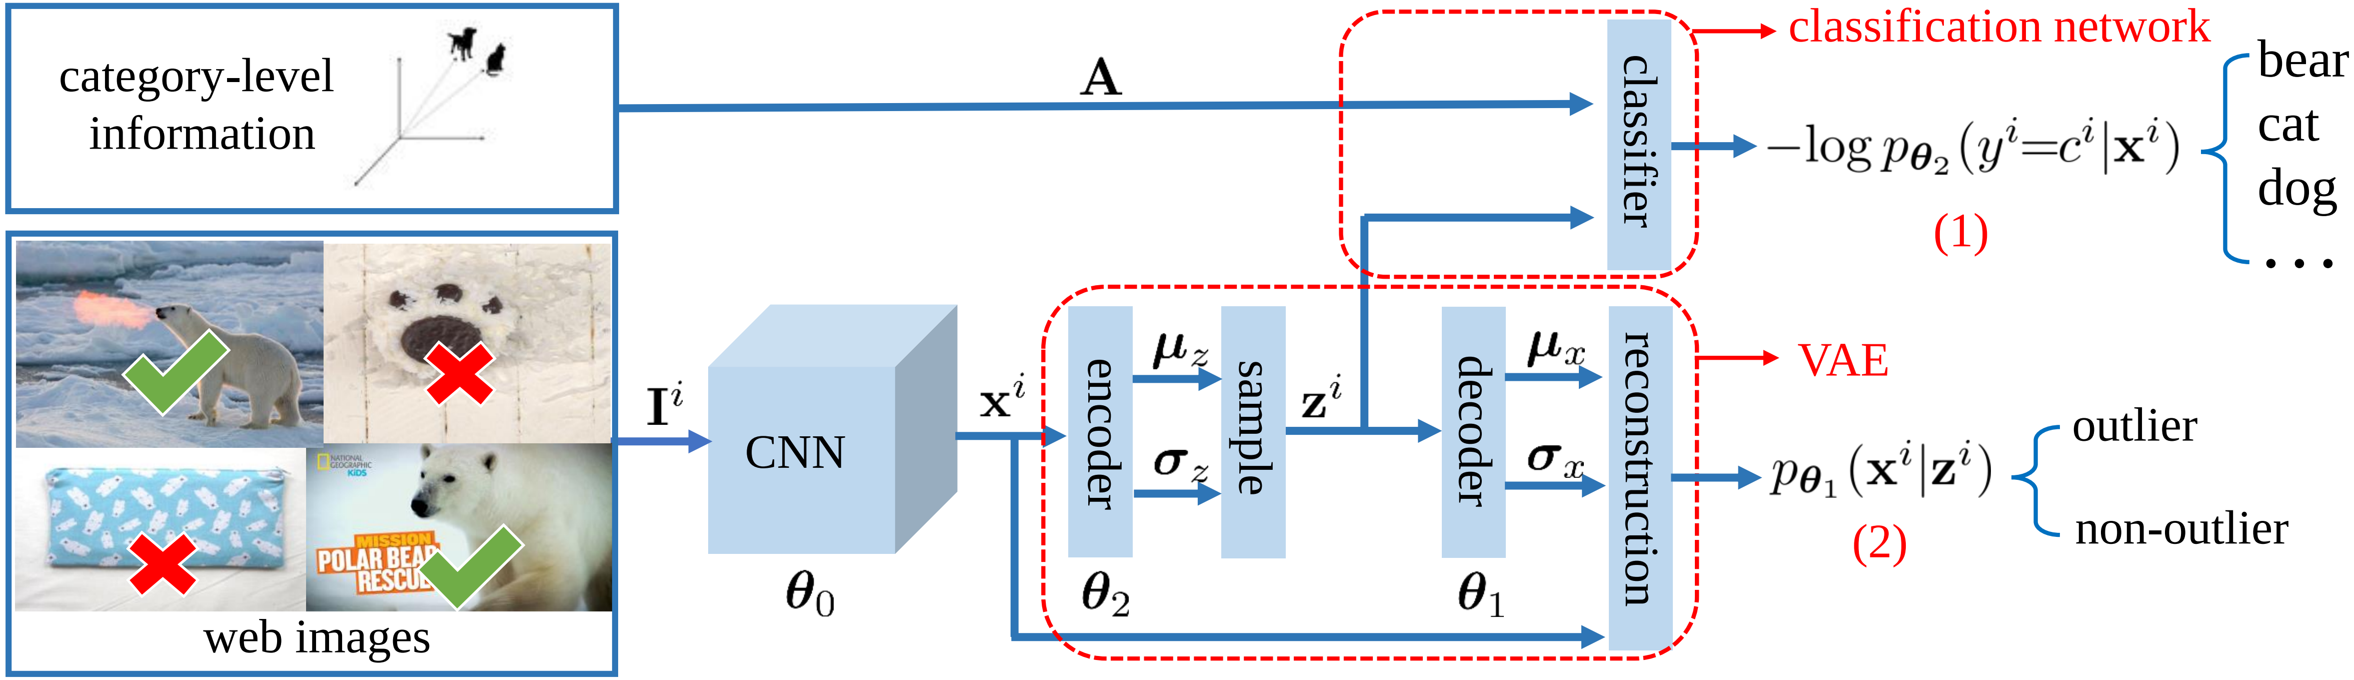

Our Method to Handle Label Noise

VAE | Our Method

$$\min_{\boldsymbol{\theta}_1,\boldsymbol{\theta}_2} \mathrm{KL}[q_{\boldsymbol{\theta}_2}(\mathbf{z}|\mathbf{x}^i)||p(\mathbf{z})] - \log p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i) \xrightarrow[\text{with classifier}]{\text{replace KL}} \min_{\boldsymbol{\theta}_0,\boldsymbol{\theta}_1,\boldsymbol{\theta}_2} -\log p(\boxed{y^i} = c^i|\mathbf{z}^i) - \log p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i)$$

However, we only know the noisy labels $\tilde{y}^i$ instead of the ground-truth labels y^i .

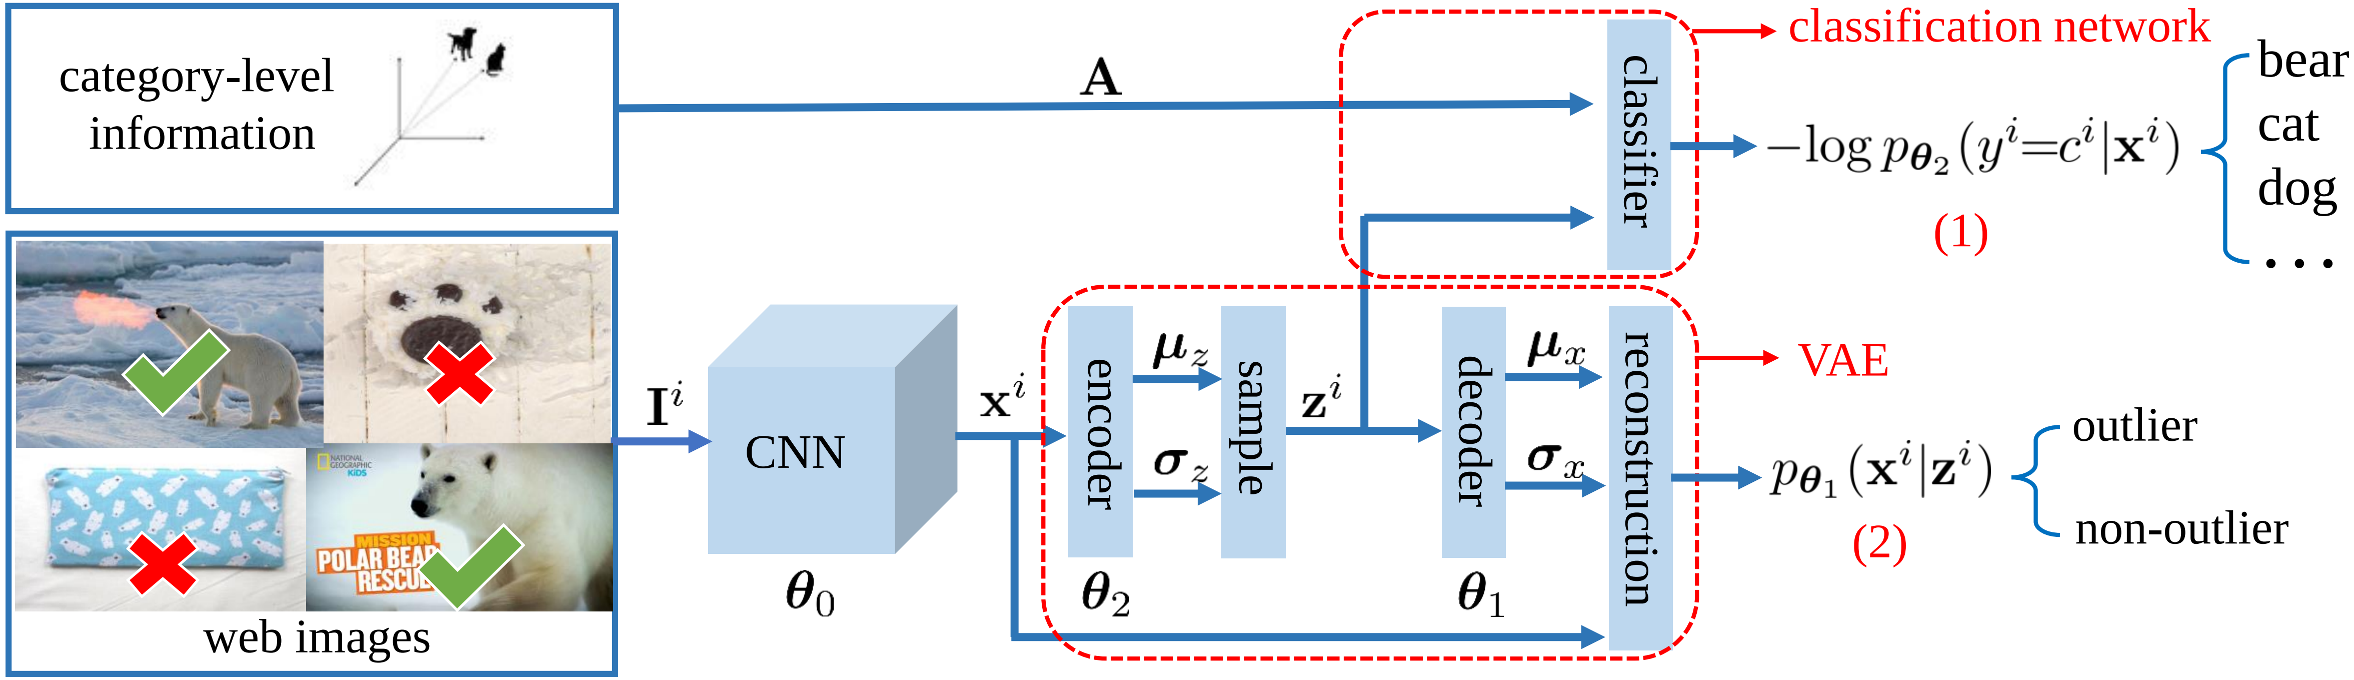

Our Method to Handle Label Noise

Our Method

$$\min_{\boldsymbol{\theta}_0,\boldsymbol{\theta}_1,\boldsymbol{\theta}_2} -\log p(\boxed{y^i} = c^i|\mathbf{z}^i) - \log p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i) \qquad (1)$$

replace y^i with $\tilde{y}^i$

$$\min_{\boldsymbol{\theta}_0,\boldsymbol{\theta}_1,\boldsymbol{\theta}_2} -\log p(\boxed{\tilde{y}^i} = c^i|\mathbf{z}^i) - \log p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i) \qquad (2)$$

upper bound

$$\min_{\boldsymbol{\theta}_0,\boldsymbol{\theta}_1,\boldsymbol{\theta}_2} -\boxed{p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i)\left(\log p(y^i = c^i|\mathbf{z}^i)\right.} + \left.\log C\right) - \log p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i) \qquad (3)$$

weighted classification loss

Our Method to Handle Label Noise

Our Method

$$\min_{\boldsymbol{\theta}_0,\boldsymbol{\theta}_1,\boldsymbol{\theta}_2} \boxed{-p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i)\left(\log p(y^i=c^i|\mathbf{z}^i)+\log C\right)-\log p_{\boldsymbol{\theta}_1}(\mathbf{x}^i|\mathbf{z}^i)} \qquad (3)$$

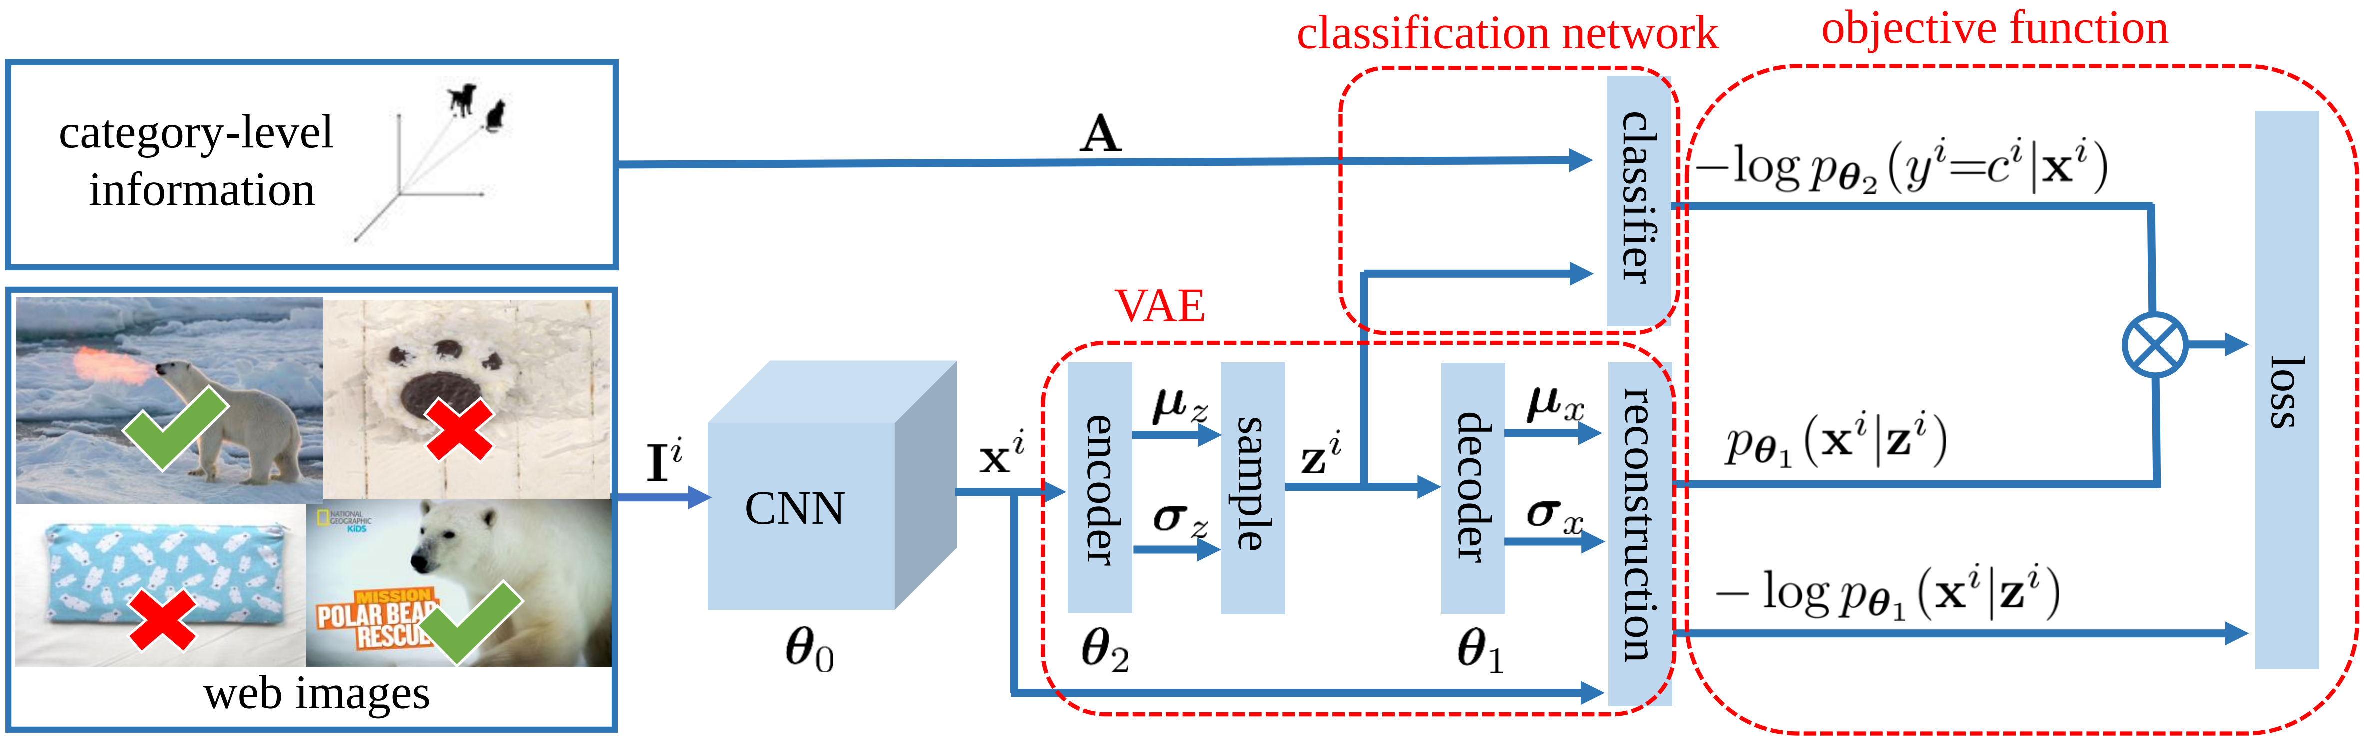

Extension to Handle Domain Shift

Extend our method to address the domain shift using two strategies.

Strategy 1: use the same VAE to reconstruct unlabeled test samples.

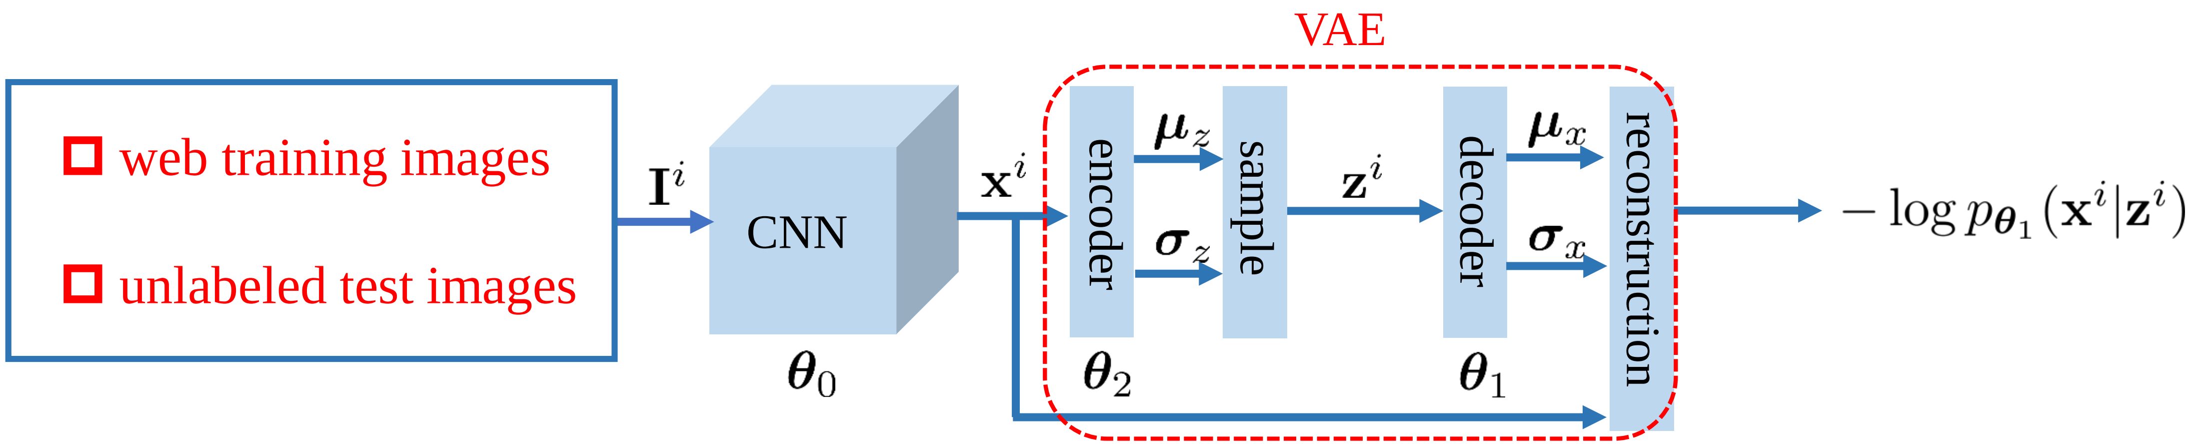

Extension to Handle Domain Shift

Extend our method to address the domain shift using two strategies.

Strategy 2: refine the latent variables $\mathbf{Z}^t$ of target domain samples.

low-rank representation [8]

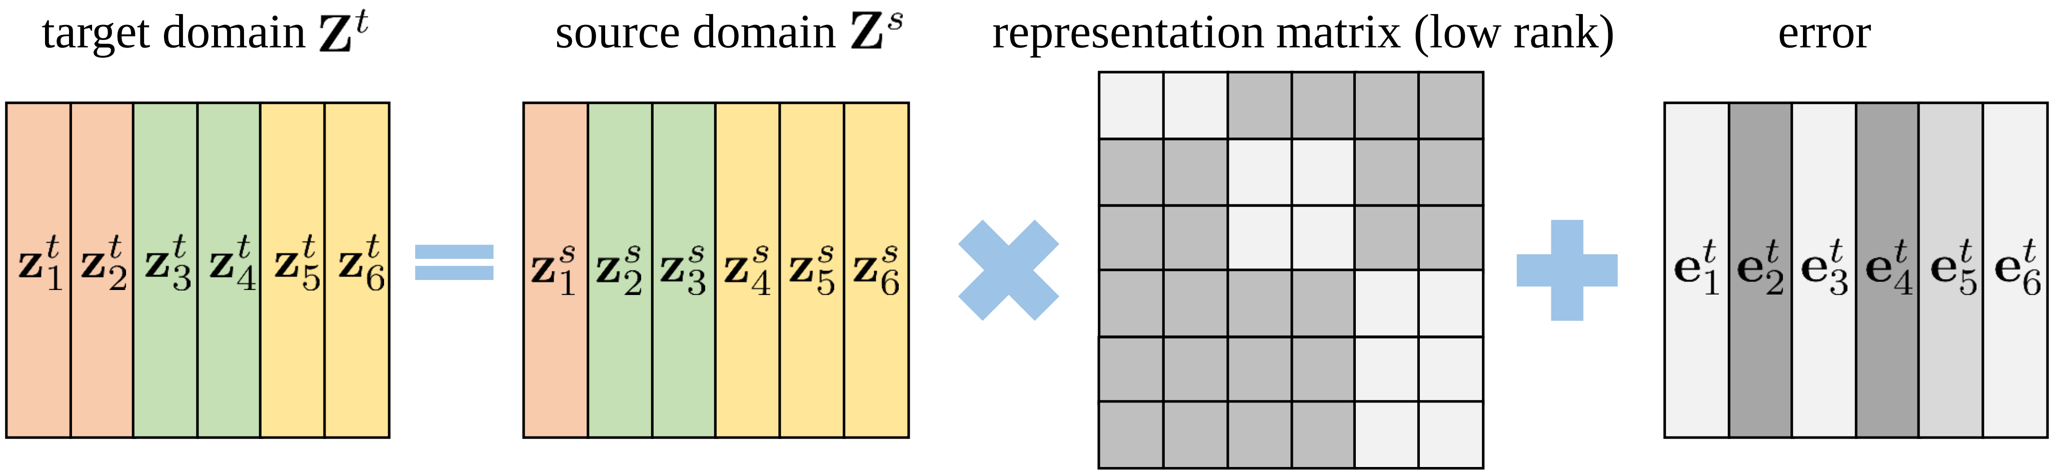

[2] Liu et al., ICML 2010

Experiments

Image classification: training: Google images → testing: AwA, CUB, SUN