Supplementary: Supervised Attribute Information Removal and Reconstruction for Image Manipulation

Nannan Li and Bryan A. Plummer

Boston University
{nnli,bplum}@bu.edu

1 Proof of the Upper Bound for Mutual Information

In Eq. (2) of the paper, we claimed that the mutual information between the source attributes $\mathbf{a}_i$ and the attribute excluded features $R(E_1(I_{cl}))$ is upper bounded by the maximum log probability in the attribute distribution. We prove this claim in the following.

Let $\mathrm{MI}(\mathbf{a}_i, R(E_1(I_{cl})))$ denote the mutual information. Replacing $R(E_1(I_{cl}))$ with r for convenience gives

$$\begin{aligned} \mathrm{MI}(\mathbf{a}_i, r) &= \sum\nolimits_r \sum\nolimits_{\mathbf{a}_i} p(\mathbf{a}_i, r) \log \frac{p(\mathbf{a}_i, r)}{p(\mathbf{a}_i)p(r)} \\ &= \sum\nolimits_r \sum\nolimits_{\mathbf{a}_i} p(\mathbf{a}_i, r) \log \frac{p(\mathbf{a}_i|r)}{p(\mathbf{a}_i)} \\ &= \sum\nolimits_r \sum\nolimits_{\mathbf{a}_i} p(\mathbf{a}_i, r)[\log p(\mathbf{a}_i|r) - \log p(\mathbf{a}_i)] \end{aligned} \tag{15}$$

Since the number of attribute values in $\mathbf{a}_i$ is finite, $-\log p(\mathbf{a}_i)$ can be upper bounded by a constant $c, c > 0$:

$$\begin{aligned} \mathrm{MI}(\mathbf{a}_i, r) &\leq \sum\nolimits_r \sum\nolimits_{\mathbf{a}_i} p(\mathbf{a}_i, r) \log p(\mathbf{a}_i|r) + c \sum\nolimits_r \sum\nolimits_{\mathbf{a}_i} p(\mathbf{a}_i, r) \\ &= \sum\nolimits_r \sum\nolimits_{\mathbf{a}_i} p(\mathbf{a}_i, r) \log p(\mathbf{a}_i|r) + c \end{aligned} \tag{16}$$

In the r.h.s., we can continue upper bounding $p(\mathbf{a}_i|r)$ with the maximum probability in the distribution to make it independent of $\mathbf{a}_i$:

$$\begin{aligned} \mathrm{MI}(\mathbf{a}_i, r) &\leq \sum\nolimits_r \max\nolimits_{\mathbf{a}_i} \log p(\mathbf{a}_i|r) \sum\nolimits_{a_i} p(\mathbf{a}_i, r) + c \\ &= \sum\nolimits_r p(r) \max\nolimits_{\mathbf{a}_i} \log p(\mathbf{a}_i|r) + c \\ &= \mathbb{E}_{r \sim p(r)}[\max\nolimits_{\mathbf{a}_i} \log p(a_i|r)] + c, \end{aligned} \tag{17}$$

where c is a constant. Note that we can not minimize the mutual information itself because the joint distribution $p(\mathbf{a}_i, r)$ is intractable. The tightness of this upper bound depends on the distribution $p(\mathbf{a}_i)$ and $p(\mathbf{a}_i|r)$. More specifically, larger $\min_{\mathbf{a_i}} p(\mathbf{a_i})$ gives smaller constant c, and smaller $\max_{\mathbf{a_i}} p(\mathbf{a_i}|r)$ reduces the gap. The equality is reached when $p(\mathbf{a_i}|r)$ is an uniform distribution.

To conclude, using an attribute classifier to estimate the above conditional probability $p(a_i|r)$, we prove that the upper bound is the maximum log probability in the attribute distribution as in Eq. (2).

2 Ablations on Hyperparameters

In Figure 1, we provide the experimental results for setting different values of the hyperparameters in Eq. (13) and (14) on CelebA. λ_1 to λ_4 denotes the trade-off parameter for disentanglement, image attribute prediction, image reconstruction and perceptual loss, respectively. Figure 1a shows the manipulation accuracy, top-5 retrieval and top-20 retrieval rates for each parameter. The reconstruction error has a different unit of measurement, for which we show its corresponding graph in Figure 1b. It can be noticed that increasing the weight (*i.e.*, λ_2) for the image attribute loss improves the manipulation accuracy, whereas it can hurt the reconstruction performance. This indicates a trade-off between successful manipulation and qualitative reconstruction. In the paper, we chose the values of each trade-off parameter for a balance between these two aspects.

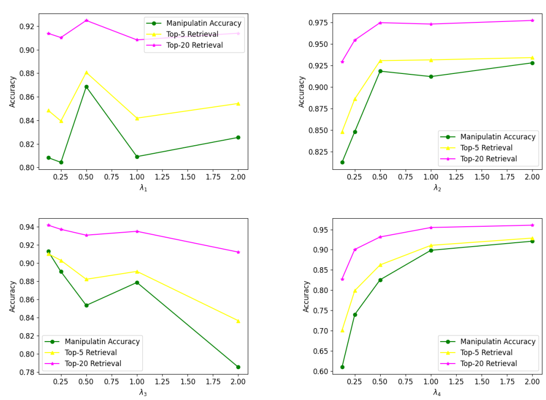

(a) Parameter value v.s. Accuracy. Higher is better

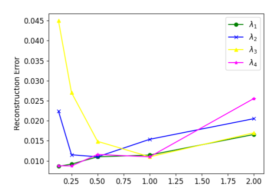

(b) Parameter Value v.s. Reconstruction Error. Lower is better

Fig. 1: Results on using different values of the hyperparameters. λ_1 to λ_4 denotes the trade-off parameters for disentanglement, image attribute prediction, image reconstruction and perceptual loss, respectively

Fig. 2: Additional examples on manipulating the attribute strength