Mining Multi-Label Samples from Single Positive Labels

Youngin Cho*, Daejin Kim*, Mohammad Azam Khan, and Jaegul Choo (*: equal contributions)

Multi-Label Data Generation in the Single Positive Setting

Multi-label dataset: each data instance can be specified with multiple attributes.
Single Positive Setting: each data instance has a label indicating only the presence of one attribute (i.e., a single positive label), and the presence of the rest of the attributes remains unknown.
Single positive setting allows us to reduce the annotation cost!

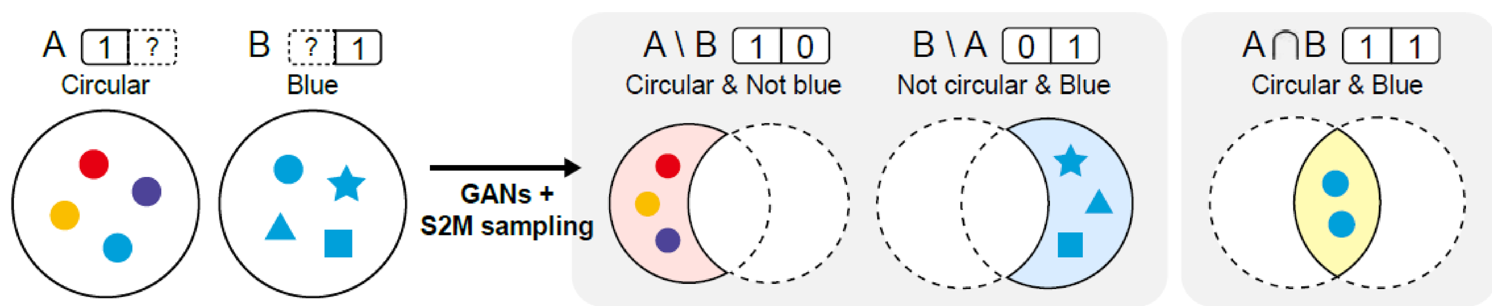

(a) Single positive setting (b) Non-overlapping classes (c) Overlapping class

In this setting, existing models cannot draw multi-label samples or tend to draw samples on a limited region of the data space.
In contrast, we propose a sampling method that draws multi-label samples without sacrificing diversity.

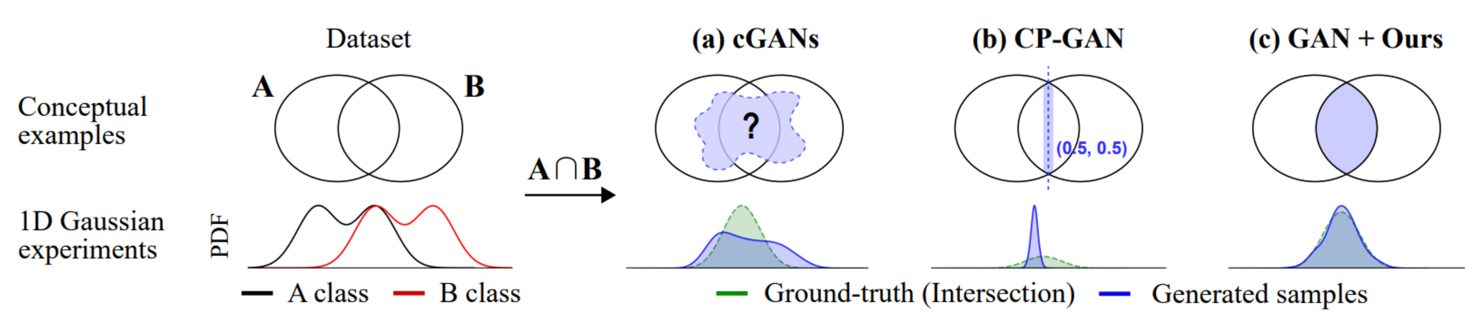

Our contributions:

- Introducing the single positive setting in the conditional generation task
- Providing the theoretical framework for estimating conditional densities of joint classes
- Proposing a sampling framework based on the Markov chain Monte Carlo method for generating multi-label data from single positive labels.

Alternative Formulation

Joint class: a class where data instances belong to certain classes but not to the rest.
We prove that the joint class density can be derived from the class-conditional densities of the single positive labels.

Assumption 1. *For every* $i, j \in N$ *such that* $i \neq j$, *if* $p(y_i = 1, y_j = 0) > 0$ *and* $p(y_j = 1, y_i = 0) > 0$, *then* $\operatorname{supp} p(x|y_i = 1, y_j = 0) \cap \operatorname{supp} p(x|y_j = 1, y_i = 0) = \emptyset$.

$$p_{(I,J)}(x) = \pi_{(I,J)}^{-1}(\min\{\pi_i p(x|y_i = 1) : i \in I\} - \max\{\pi_j p(x|y_j = 1) : j \in J\} \cup \{0\})^{+}.$$

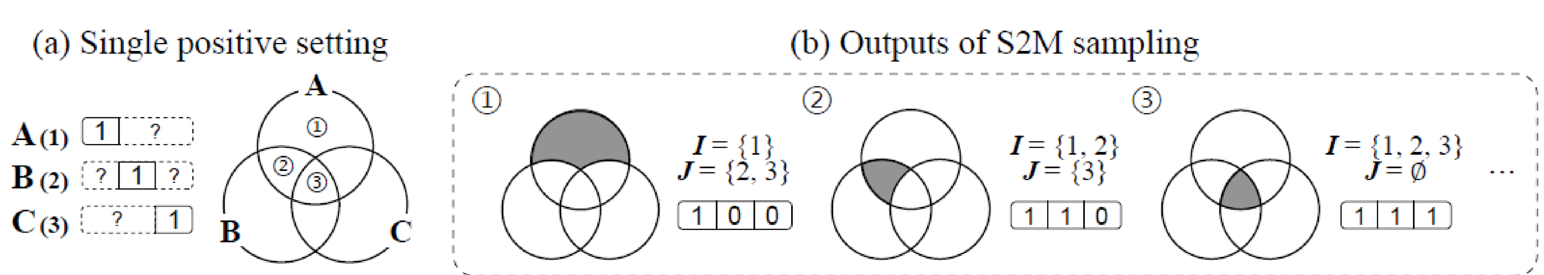

S2M Sampling

Training a generator that can model the alternative target density is a non-trivial problem since the formula consists of several implicitly defined conditional densities, class priors, and variable sets.
To address this issue, we propose the application of sampling approaches upon existing GANs. We utilize several classification networks to compute implicitly defined density ratios, which are used to run the independent Metropolis-Hastings algorithm.

$$\mathcal{L}_v = -\mathbb{E}_{(x,c)\sim p_{data}(x,c)}[\log D_v(x)] - \mathbb{E}_{x\sim p_G(x)}[\log(1 - D_v(x))],$$
$$\mathcal{L}_r = -\mathbb{E}_{(x,c)\sim p_{data}(x,c)}[\log D_r(c|x)]$$
$$D_v^*(x) = \frac{p_{data}(x)}{p_{data}(x) + p_G(x)}, D_r^*(c|x) = \frac{p(x|y_c = 1)p_{data}(c)}{p_{data}(x)}$$

Algorithm 1 *S2M Sampling for GANs*

Input: generator G, classifiers D_v^*, D_r^*, intersection index set I, difference index set J, and class prior ratios $\gamma_{1:N}$
Output: filtered sample x

Choose any $x \in \operatorname{supp} p_{(I,J)}$.
for $k = 1$ to K **do**
  Draw x' from G.
  Draw u from Uniform(0,1).
  $r_i \leftarrow \gamma_i D_r^*(i|x)$ for every $i \in I \cup J$
  $r'_i \leftarrow \gamma_i D_r^*(i|x')$ for every $i \in I \cup J$
  $\alpha \leftarrow \min\left(1, \frac{(\min\{r'_i : i\in I\} - \max\{r'_j : j\in J\}\cup\{0\})^{+}(D_v^*(x)^{-1}-1)}{(\min\{r_i : i\in I\} - \max\{r_j : j\in J\}\cup\{0\})^{+}(D_v^*(x')^{-1}-1)}\right)$
  if $u \leq \alpha$ **then**
    $x \leftarrow x'$
  end if
end for

Latent Adaptation

The sample efficiency of independent Metropolis-Hastings algorithm is low if the target distribution is significantly different from the generator distribution. To alleviate this issue, we update the latent distribution of GANs using target joint class samples obtained from past sampling attempts. The sample efficiency can be further improved since the newly obtained proposal distribution is much close to the target distribution.

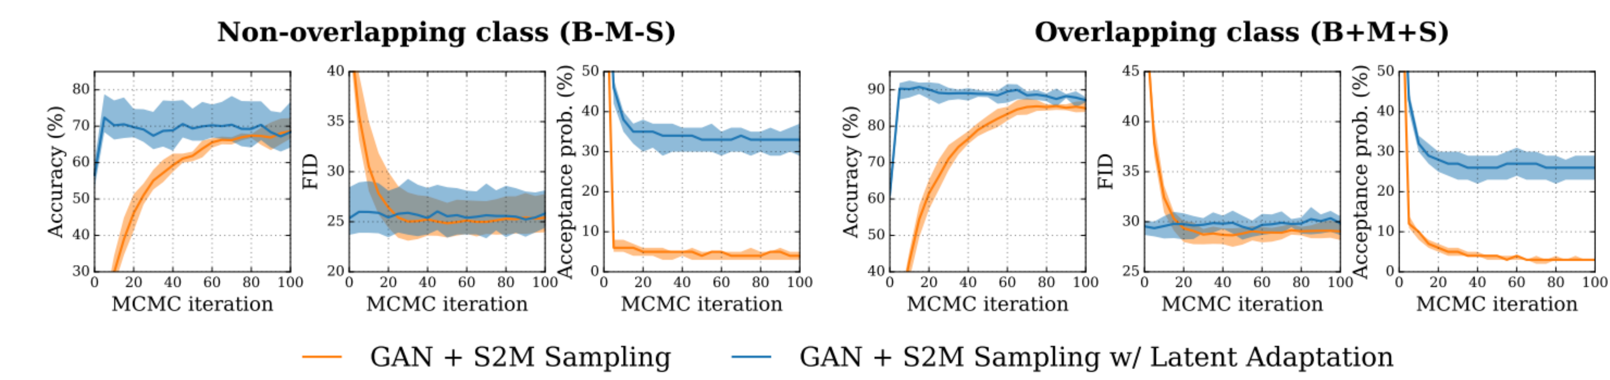

If the target class samples rarely appear in the generated samples, one can consider applying latent adaptation repeatedly. This technique reduces the sampling time complexity to be proportional to the number of attributes.

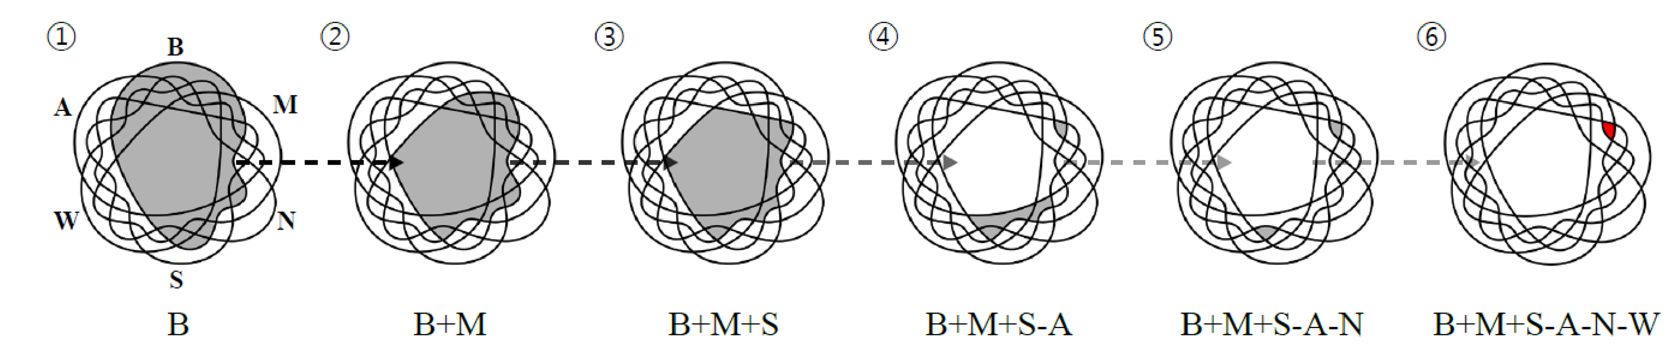

Multi-Label Image Generation Results in the Single Positive Setting

We conduct experiments using multi-class and multi-label image datasets.

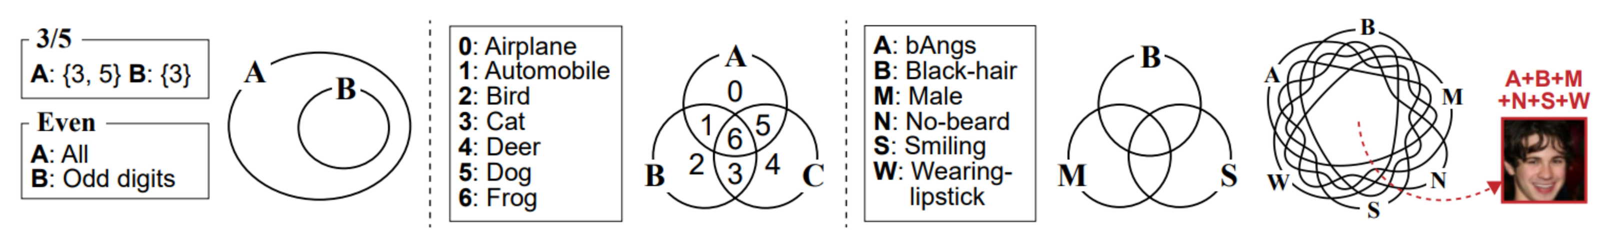

(a) (F)MNIST (b) CIFAR-7to3 (c) CelebA-BMS (d) CelebA-ABMNSW

CIFAR-7to3 has three classes (A, B, C), each of which contains four original classes of CIFAR-10, i.e. A={Airplance, Automobile, Dog, Frog}, B={Automobile, Bird, Cat, Frog}, C={Cat, Deer, Dog, Frog}.

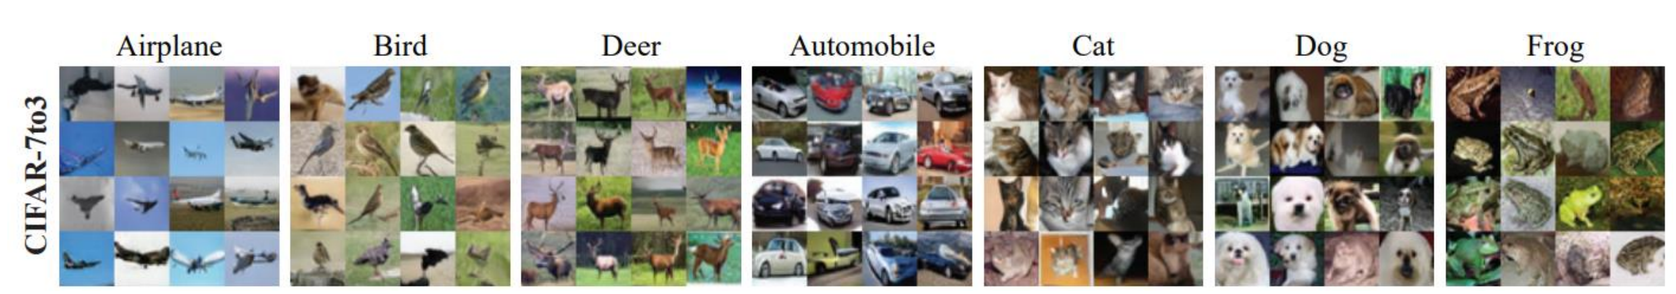

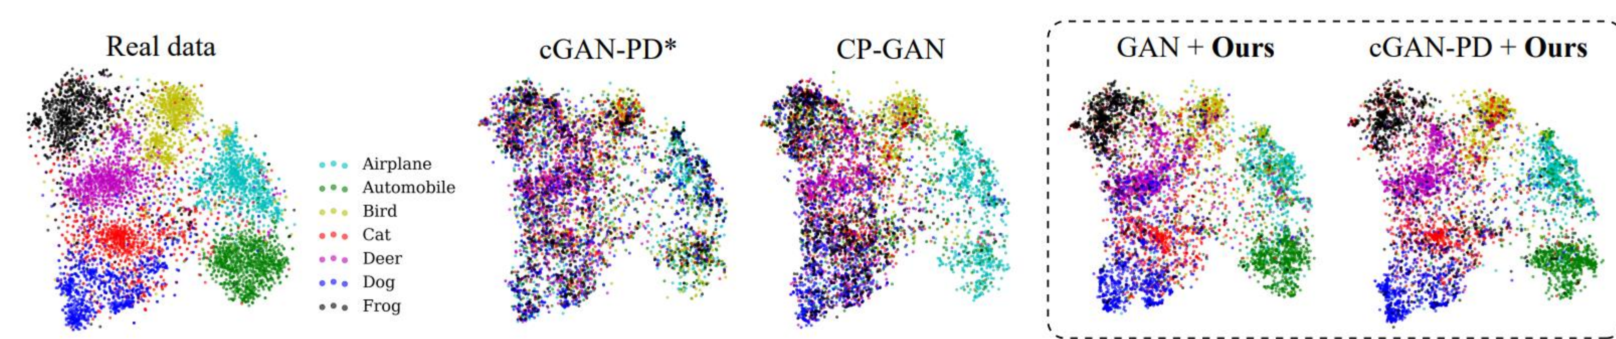

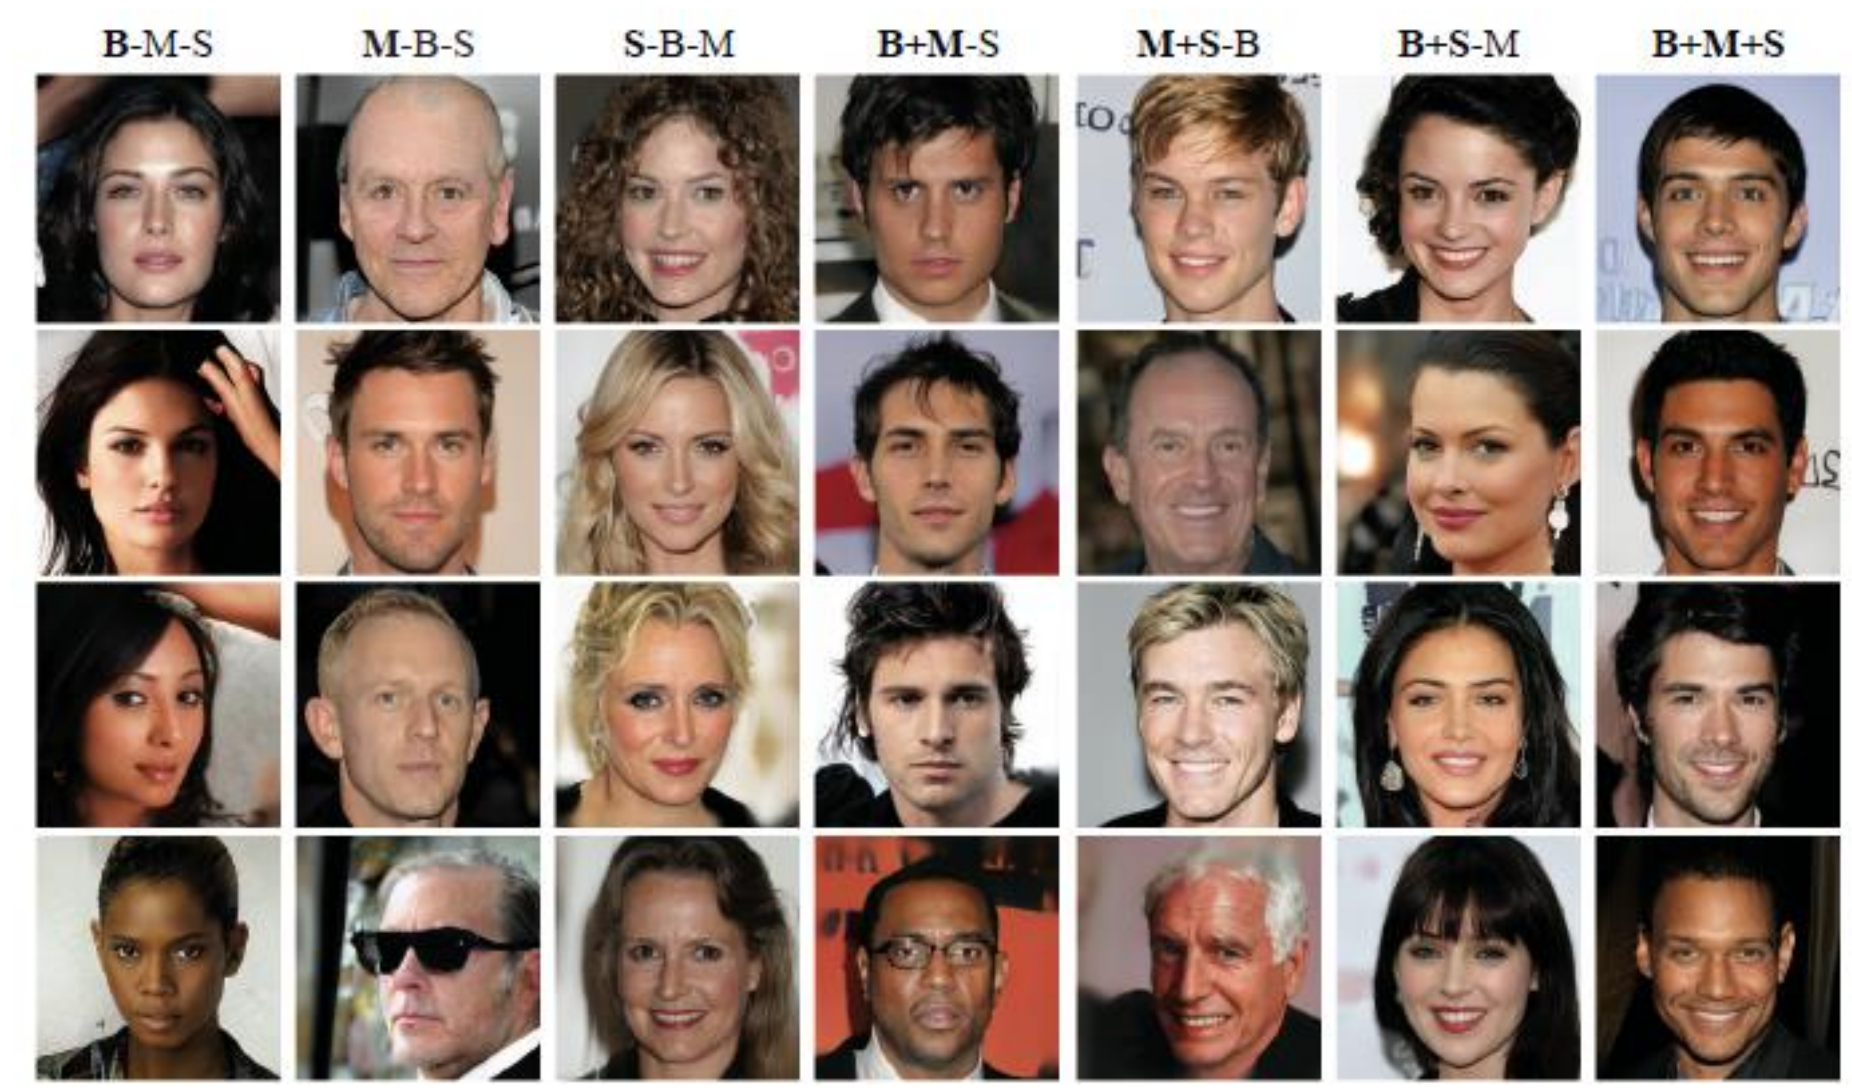

Given three attributes Black-hair (B), Male (M), Smiling (S), we can obtain Unsmiling black-haired Woman (B-M-S), Smiling black-haired Man (B+M+S) …