## A Appendix



Figure 4: Training a linear VAE with 32 latent variables with 64 categories on dynamically binarized MNIST, FashionMNIST, and Omniglot datasets by maximizing the ELBO. We plot the train ELBO (left column), the test 100-sample bound (middle column), and the variance of the gradient estimator (right column). For a fair comparison, the variance of all the gradient estimators was computed along the training trajectory of the RLOO estimator. We plot the mean and one standard error based on 5 runs from different random initializations.


Figure 5: The effect of logit ordering on the performance of DisARM-SB. We sort the encoder logits in the ascending (Green) or descending (Red) order, and compare against the default ordering (Blue).

## A. 1 Experiments with linear categorical VAEs

We evaluate the three proposed gradient estimators, DisARM-IW, DisARM-SB, and DisARM-Tree, by training linear variational auto-encoders with categorical latent variables on dynamically bina-


Figure 6: Training non-linear categorical VAEs with different model sizes on dynamically binarized FashionMNIST. Left: 5 latent variables of 64 categories. Middle: with 32 latent variables with 64 categories. Right: 32 latent variables with 16 categories.

Table 3: Training a non-linear VAE with categorical latents using multi-sample estimators. Mean variational lower bounds and the standard error of the mean computed based on 5 runs of $5 \times 10^{5}$-steps training from different random initializations. The best performing method (up to the standard error) for each task is in bold.

| \#Pairs | Dataset | DisARM-Tree | DisARM-SB | DisARM-IW | RLOO |
| :---: | :--- | :---: | :---: | :---: | :---: |
| 5 | Dynamic MNIST | $-\mathbf{9 1 . 3 6} \pm \mathbf{0 . 1 1}$ | $-91.67 \pm 0.08$ | $-91.85 \pm 0.04$ | $-92.35 \pm 0.10$ |
|  | Fashion MNIST | $\mathbf{- 2 3 1 . 0 1} \pm \mathbf{0 . 1 4}$ | $-231.38 \pm 0.14$ | $\mathbf{- 2 3 1 . 1 4} \pm \mathbf{0 . 1 7}$ | $-231.78 \pm 0.16$ |
|  | Omniglot | $\mathbf{- 1 0 9 . 0 0} \pm \mathbf{0 . 1 0}$ | $\mathbf{- 1 0 8 . 9 0} \pm \mathbf{0 . 0 9}$ | $\mathbf{- 1 0 8 . 8 3} \pm \mathbf{0 . 1 1}$ | $-109.73 \pm 0.07$ |
| 10 | Dynamic MNIST | $\mathbf{- 9 0 . 8 1} \pm \mathbf{0 . 0 6}$ | $\mathbf{- 9 1 . 0 1} \pm \mathbf{0 . 0 6}$ | $\mathbf{- 9 0 . 9 5} \pm \mathbf{0 . 0 8}$ | $-91.86 \pm 0.15$ |
|  | Fashion MNIST | $\mathbf{- 2 3 0 . 9 5} \pm \mathbf{0 . 2 1}$ | $-231.21 \pm 0.18$ | $-231.25 \pm 0.29$ | $-231.59 \pm 0.21$ |
|  | Omniglot | $-108.27 \pm 0.03$ | $-108.19 \pm 0.07$ | $-\mathbf{1 0 8 . 0 8} \pm \mathbf{0 . 0 6}$ | $-108.95 \pm 0.17$ |

rized MNIST, FashionMNIST, and Omniglot datasets. As in (Dong et al. 2020), we benchmark the proposed estimators against the 2-sample REINFORCE estimator with the leave-one-out baseline (RLOO; Kool et al. 2019). The linear model has a single layer of 32 categorical latent variables, each with 64 categories. We find no significant difference in performance between the proposed estimators and the RLOO baseline in this setting (Appendix Table 4 and Appendix Figure 4). However, as we noted in the maintext, Dong et al. (2020) found that for multi-layer linear models in the binary case, DisARM showed increasing improvement over RLOO for models with deeper hierarchies. So it would be interesting to see whether this holds for the categorical case in future work.

Table 4: Training a linear VAE with categorical latents using the proposed estimators and the RLOO baseline. Mean variational lower bounds and the standard error of the mean computed based on 5 runs of $5 \times 10^{5}$-steps training from different random initializations. The best performing method (up to the standard error) for each task is in bold.

| Training set ELBO |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: |
|  | DisARM-IW | DisARM-Tree | DisARM-SB | RLOO |
| Dynamic MNIST | $-115.64 \pm 0.09$ | $\mathbf{- 1 1 5 . 5 9} \pm \mathbf{0 . 1 6}$ | $\mathbf{- 1 1 5 . 4 8} \pm \mathbf{0 . 1 4}$ | $\mathbf{- 1 1 5 . 5 0} \pm \mathbf{0 . 1 8}$ |
| Fashion MNIST | $\mathbf{- 2 5 4 . 2 3} \pm \mathbf{0 . 1 9}$ | $-254.56 \pm 0.22$ | $\mathbf{- 2 5 4 . 0 5} \pm \mathbf{0 . 2 0}$ | $\mathbf{- \mathbf { 2 5 4 . 1 8 } \pm \mathbf { 0 . 1 5 }}$ |
| Omniglot | $-124.64 \pm 0.02$ | $\mathbf{- 1 2 4 . 5 2} \pm \mathbf{0 . 1 1}$ | $-124.82 \pm 0.21$ | $-125.10 \pm 0.11$ |



Figure 7: Training a non-linear categorical VAE with 128 latent variables and 16 categories, using multi-sample objectives, on dynamically binarized MNIST, FashionMNIST, Omniglot datasets by maximizing the ELBO. The examples in the top 2 rows are using 5 pairs of samples, while the ones in the bottom 2 rows are using 10 pairs.

## A. 2 Experimental Details

We use the same model structure as in (Yin et al. 2019). The model has a single layer of categorical latent variables which are mapped to Bernoulli logits using an MLP with two hidden layers of 256 and 512 of LeakyReLU units (Xu et al., 2015) with negative slope of 0.2 . The encoder mirrors the structure, having two hidden layers of 512 and 256 LeakyReLU units.
For a fair comparison of the variance of the gradient estimators, we train a model with the RLOO estimator and evaluate the variance of all the estimators at each step along the training trajectory of this model. Based on preliminary experiments, the results were independent of the gradient estimator used to generate the model trajectory. We report the average per-parameter variance based on the parameter moments estimated with an exponential moving average with decay rate 0.999 .
Each experiment run takes around 12 hours on an NVIDIA Tesla P100 GPU. Our implementation was biased towards readability instead of computational efficiency, so we expect that significant improvements in runtime could be achieved.

## A. 3 Importance Weighting Derivation

We consider proposal distributions that factorize across dimensions $p(z, \tilde{z})=\prod_{k} p\left(z_{k}, \tilde{z}_{k}\right)$ and that are couplings such that the marginals are maintained $p\left(z_{k}\right)=p\left(\tilde{z_{k}}\right)=q\left(z_{k}\right)$. In general, $q_{\theta}(z) q_{\theta}(\tilde{z})$ has full support, but we know that the integrand $q_{\theta}(z) q_{\theta}(\tilde{z})(f(z)-$ $f(\tilde{z}))\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right)$ vanishes when $z_{k}=\tilde{z}_{k}$ by inspection, so we allow $p\left(z_{k}, \tilde{z}_{k}\right)$ to put zero mass on $z_{k}=\tilde{z}_{k}$ configurations, and require $p\left(z_{k}, \tilde{z}_{k}\right)>0$ otherwise.

Then with $z, \tilde{z} \sim p$, we show that $g_{\text {DisARM-Iw }}$ (Eq. 3) is an unbiased estimator. First,

$$
\begin{aligned}
\mathbb{E}_{p(z, \tilde{z})} & {\left[\frac{1}{2} \frac{q_{\theta}\left(z_{k}\right) q_{\theta}\left(\tilde{z_{k}}\right)}{p_{\theta}\left(z_{k}, \tilde{z_{k}}\right)}(f(z)-f(\tilde{z}))\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right)\right] } \\
& =\mathbb{E}_{p\left(z_{-k}, \tilde{z}_{-k}\right)} \mathbb{E}_{p\left(z_{k}, \tilde{z}_{k}\right)}\left[\frac{1}{2} \frac{q_{\theta}\left(z_{k}\right) q_{\theta}\left(\tilde{z_{k}}\right)}{p_{\theta}\left(z_{k}, \tilde{z_{k}}\right)}(f(z)-f(\tilde{z}))\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right)\right] \\
& =\mathbb{E}_{p\left(z_{-k}, \tilde{z}_{-k}\right)} \sum_{z_{k}, \tilde{z}_{k} \in \operatorname{supp} p\left(z_{k}, \tilde{z}_{k}\right)} \frac{1}{2} q_{\theta}\left(z_{k}\right) q_{\theta}\left(\tilde{z_{k}}\right)(f(z)-f(\tilde{z}))\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right) \\
& =\mathbb{E}_{p\left(z_{-k}, \tilde{z}_{-k}\right)} \sum_{z_{k}, \tilde{z_{k}}} \frac{1}{2} q_{\theta}\left(z_{k}\right) q_{\theta}\left(\tilde{z_{k}}\right)(f(z)-f(\tilde{z}))\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right) \\
& =\mathbb{E}_{p\left(z_{-k}, \tilde{z}_{-k}\right)} \mathbb{E}_{q_{\theta}\left(z_{k}\right) q_{\theta}\left(\tilde{z}_{k}\right)}\left[\frac{1}{2}(f(z)-f(\tilde{z}))\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right)\right]
\end{aligned}
$$

using the fact that the integrand vanishes outside the support of $p\left(z_{k}, \tilde{z}_{k}\right)$. Then,

$$
\begin{aligned}
\mathbb{E}_{p\left(z_{-k}, \tilde{z}_{-k}\right)} & \mathbb{E}_{q_{\theta}\left(z_{k}\right) q_{\theta}\left(\tilde{z}_{k}\right)}\left[\frac{1}{2}(f(z)-f(\tilde{z}))\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right)\right] \\
& =\mathbb{E}_{q_{\theta}\left(z_{k}\right) q_{\theta}\left(\tilde{z}_{k}\right)}\left[\frac{1}{2}\left(\mathbb{E}_{p\left(z_{-k}, \tilde{z}_{-k}\right)}[f(z)-f(\tilde{z})]\right)\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right)\right] \\
& =\mathbb{E}_{q\left(z_{k}\right) q\left(\tilde{z}_{k}\right)}\left[\frac{1}{2}\left(\mathbb{E}_{p\left(z_{-k}, \tilde{z}_{-k}\right)}[f(z)]-\mathbb{E}_{p\left(z_{-k}, \tilde{z}_{-k}\right)}[f(\tilde{z})]\right)\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right)\right] \\
& =\mathbb{E}_{q\left(z_{k}\right) q\left(\tilde{z}_{k}\right)}\left[\frac{1}{2}\left(\mathbb{E}_{q\left(z_{-k}\right)}[f(z)]-\mathbb{E}_{q\left(\tilde{z}_{-k}\right)}[f(\tilde{z})]\right)\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right)\right] \\
& =\mathbb{E}_{q(z) q(\tilde{z})}\left[\frac{1}{2}(f(z)-f(\tilde{z}))\left(\nabla_{\theta_{k}} \log q_{\theta}\left(z_{k}\right)-\nabla_{\theta_{k}} \log q_{\theta}\left(\tilde{z}_{k}\right)\right)\right]
\end{aligned}
$$

following from the linearity of the expectation and the fact that the coupling preserves marginals.

## A. 4 Stick-Breaking Coupling

## A.4.1 Computing Importance Weights

To simplify the notation in this subsection, we consider a single dimension at a time and omit the dimension index. To compute $\frac{q_{\theta}(z) q_{\theta}(\tilde{z})}{p_{\theta}(z, \tilde{z})}$, it is helpful to work with the logits of the binary variables $\alpha_{i}=\operatorname{logit} \frac{q(z=i)}{\sum_{j=i+1}^{C} q(z=j)}$. First, we know that $q(z=i)=\prod_{j=1}^{i-1} \sigma\left(-\alpha_{j}\right) \sigma\left(\alpha_{i}\right)$. For a pair of antithetic binary variables $\left(b_{i}, \tilde{b}_{i}\right)$, the coupling joint probability is

$$
p\left(b_{i}, \tilde{b}_{i}\right)= \begin{cases}\max \left(1-2 \sigma\left(\alpha_{i}\right), 0\right) & b_{i}=\tilde{b}_{i}=0 \\ \max \left(2 \sigma\left(\alpha_{i}\right)-1,0\right) & b_{i}=\tilde{b}_{i}=1 \\ \sigma\left(-\left|\alpha_{i}\right|\right) & \text { o.w. }\end{cases}
$$

Since the categories are arranged in the ascending order of probability, $\alpha_{i} \leq 0$ for $i<C$, so the joint probability simplifies to

$$
p\left(b_{i}, \tilde{b}_{i}\right)= \begin{cases}1-2 \sigma\left(\alpha_{i}\right) & b_{i}=\tilde{b}_{i}=0 \\ 0 & b_{i}=\tilde{b}_{i}=1 \\ \sigma\left(\alpha_{i}\right) & \text { o.w. }\end{cases}
$$

for $i<C$. We do not need to compute the entries for $z=\tilde{z}$ as the integrand already vanishes. Because of symmetry, without loss of generality, assume $z<\tilde{z}$. From the construction of $z$ and $\tilde{z}$ in terms of the binary variables, we can reason about their values. We know that for $i<z$, we must have $b_{i}=\tilde{b}_{i}=0$. Then, for $i=z$, we must have $b_{i}=1$ and for $z \leq i<\tilde{z}, \tilde{b}_{i}=0$, and finally, for $i=\tilde{z}, \tilde{b}=1$. Putting this together yields

$$
p(z, \tilde{z})=\left[\prod_{i=1}^{z-1}\left(1-2 \sigma\left(\alpha_{i}\right)\right)\right] \sigma\left(\alpha_{z}\right)\left[\prod_{j=z+1}^{\tilde{z}-1} \sigma\left(-\alpha_{j}\right)\right] \sigma\left(\alpha_{\tilde{z}}\right) .
$$

Thus, the importance weights are

$$
\frac{q_{\theta}(z) q_{\theta}(\tilde{z})}{p_{\theta}(z, \tilde{z})}=\frac{\left[\prod_{i=1}^{z-1} \sigma\left(-\alpha_{i}\right)\right] \sigma\left(\alpha_{z}\right)\left[\prod_{j=1}^{\tilde{z}-1} \sigma\left(-\alpha_{j}\right)\right] \sigma\left(\alpha_{\tilde{z}}\right)}{\left[\prod_{i=1}^{z-1}\left(1-2 \sigma\left(\alpha_{i}\right)\right)\right] \sigma\left(\alpha_{z}\right)\left[\prod_{j=z+1}^{\tilde{z}-1} \sigma\left(-\alpha_{j}\right)\right] \sigma\left(\alpha_{\tilde{z}}\right)}=\frac{\left[\prod_{i=1}^{z-1} \sigma\left(-\alpha_{i}\right)^{2}\right] \sigma\left(-\alpha_{z}\right)}{\prod_{i=1}^{z-1}\left(1-2 \sigma\left(\alpha_{i}\right)\right)} .
$$

## A.4.2 Unbiasedness of $g_{\text {DisARM-SB }}$

Recall the $g_{\text {DisARM-SB }}$ estimator. We have binary variables $b_{k, 1} \sim \operatorname{Bernoulli}\left(\sigma\left(\alpha_{k, 1}\right)\right), \ldots, b_{k, C} \sim$ Bernoulli $\left(\sigma\left(\alpha_{k, C}\right)\right)$ and independently sampled antithetic pairs $\left\{\tilde{b}_{k, c}\right\}$ such that $z=h(b)$ and $\tilde{z}=h(\tilde{b})$. The estimator is

$$
g_{\text {DisARM-SB }_{k, c}}=\left\{\begin{array}{ll}
\frac{1}{2}(f(z)-f(\tilde{z}))\left((-1)^{\tilde{b}_{k, c}} \mathbb{1}_{b_{k, c}, \tilde{b}_{k, c}} \sigma\left(\mid\left(\alpha_{k, c} \mid\right)\right)\right. & c \leq \min \left(z_{k}, \tilde{z}_{k}\right) \\
\frac{1}{2}(f(\tilde{z})-f(z)) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right) & z_{k}<c \leq \tilde{z}_{k} \\
\frac{1}{2}(f(z)-f(\tilde{z})) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(b_{k, c}\right) & \tilde{z}_{k}<c \leq z_{k} \\
0 & c>\max \left(z_{k}, \tilde{z}_{k}\right)
\end{array} .\right.
$$

We claim that for any $b_{-k, c}$ and $\tilde{b}_{-k, c}$,
$\mathbb{E}_{b_{k, c}, \tilde{b}_{k, c}}\left[g_{\text {DisARM-SB } k, c}\right]=\frac{1}{2} \mathbb{E}_{b_{k, c}}\left[f(h(b)) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(b_{k, c}\right)\right]+\frac{1}{2} \mathbb{E}_{\tilde{b}_{k, c}}\left[f(h(\tilde{b})) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right)\right]$,
which immediately implies unbiasedness. Importantly, the conditions defining the estimator can be determined solely based on $b_{-k, c}$ and $\tilde{b}_{-k, c}$, so it suffices to verify that the estimator is unbiased for each case separately. In the first case, $g_{\text {DisARM-SB }}$ is the DisARM estimator from (Dong et al. 2020) which was previously shown to be unbiased. The second and third cases are reminiscent of the 2-sample RLOO estimator, however, in the coupled case, such an estimator must be justified as unbiased. In the second case,

$$
\begin{aligned}
\mathbb{E}_{b_{k, c}, \tilde{b}_{k, c}} & {\left[\frac{1}{2}(f(\tilde{z})-f(z)) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right)\right] } \\
& =\frac{1}{2} \mathbb{E}_{\tilde{b}_{k, c}}\left[f(\tilde{z}) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right)\right]-\frac{1}{2} \mathbb{E}_{\tilde{b}_{k, c}}\left[\mathbb{E}_{b_{k, c} \mid \tilde{b}_{k, c}}[f(z)] \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right)\right] \\
& =\frac{1}{2} \mathbb{E}_{\tilde{b}_{k, c}}\left[f(\tilde{z}) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right)\right]-\overbrace{\frac{1}{2} \mathbb{E}_{\tilde{b}_{k, c}}\left[f(z) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right)\right]}^{0} \\
& =\frac{1}{2} \mathbb{E}_{\tilde{b}_{k, c}}\left[f(\tilde{z}) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right)\right]+\overbrace{\frac{1}{2} \mathbb{E}_{b_{k, c}}\left[f(z) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(b_{k, c}\right)\right]}^{0} .
\end{aligned}
$$

$f(z)$ depends on $\tilde{b}_{k, c}$ through the coupled sample $b_{k, c}$, however the condition $c>z_{k}$ implies that $z$ does not depend on $b_{k, c}$, so $f(z)$ is a constant with respect to both $b_{k, c}$ and $\tilde{b}_{k, c}$, resulting in the vanishing terms. The third case follows by symmetry. In the fourth case, the condition $c>\max \left(z_{k}, \tilde{z}_{k}\right)$ implies that neither $z$ nor $\tilde{z}$ depend on $b_{k, c}$ or $\tilde{b}_{k, c}$ hence the gradient vanishes.

## A. 5 Tree-Structured Coupling

We construct a categorical sample based on a binary sequence arranged as a balanced binary tree. Considering a binary sequence $b=\left[b_{0}, b_{1}, b_{2}, \cdots, b_{C-1}\right]$, we interpret as a binary tree recursively with root $b[0]$ and left subtree $b[1: \operatorname{len}(b) / / 2+1]$ and right subtree $b[\operatorname{len}(b) / / 2+1:]$. The binary variables correspond to internal routing decisions in the binary tree with the categories as the leaves, so that $T(b)$ is defined by the following recursive function (assuming $C$ is a power of 2 )

```
def T(b):
    half = len(b)//2 + 1
    if not b:
            return 1
    elif b[0] == 0:
```

```
    return T(b[1:half])
else:
    return half + T(b[half:])
```

For each binary variable $b_{i}$, we can find all the categories residing in its left subtree (Left(i)) and all the categories in its right subtree Right(i). We would like the probability of traversing the right subtree $\sigma\left(\alpha_{i}\right)$ to be such that

$$
\sigma\left(\alpha_{i}\right)=\frac{\sum_{j \in \operatorname{Right}(\mathrm{i})} q(z=j)}{\sum_{j \in \operatorname{Right}(\mathrm{i})} q(z=j)+\sum_{j \in \operatorname{Left}(\mathrm{i})} q(z=j)}
$$

Hence,

$$
\alpha_{i}=\log \frac{\sum_{j \in \operatorname{Right}(\mathrm{i})} q(z=j)}{\sum_{j \in \operatorname{Left}(\mathrm{i})} q(z=j)}
$$

Recall, the estimator from the maintext. Given binary variables $b_{1}, \ldots, b_{C-1}$, let $I\left(b_{1}, \ldots, b_{C-1}\right)$ be the set of variables used in routing decisions $\left(\left|I\left(b_{1}, \ldots, b_{C-1}\right)\right|=\log _{2} C\right)$. With a pair of antithetically sampled binary sequences $(b, \tilde{b})$, the following is an unbiased estimator:

$$
g_{\text {DisARM-Tree } k, c}=\left\{\begin{array}{ll}
\frac{1}{2}(f(z)-f(\tilde{z}))\left((-1)^{\tilde{b}_{k, c}} \mathbb{1}_{b_{k, c} \neq \tilde{b}_{k, c}} \sigma\left(\mid\left(\alpha_{k, c} \mid\right)\right)\right. & c \in I\left(b_{k, .}\right) \cap I\left(\tilde{b}_{k, \cdot}\right)  \tag{6}\\
\frac{1}{2}(f(\tilde{z})-f(z)) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right) & c \in I\left(\tilde{b}_{k, .}\right)-I\left(b_{k, \cdot}\right) . \\
\frac{1}{2}(f(z)-f(\tilde{z})) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(b_{k, c}\right) & c \in I\left(b_{k, .}\right)-I\left(\tilde{b}_{k, .}\right) \\
0 & c \notin I\left(b_{k, .}\right) \cup I\left(\tilde{b}_{k, \cdot}\right)
\end{array} .\right.
$$

Closely following the argument from the previous section, we claim that for any $b_{-k, c}$ and $\tilde{b}_{-k, c}$, $\mathbb{E}_{b_{k, c}, \tilde{b}_{k, c}}\left[g_{\text {DisARM-Tree } k, c}\right]=\frac{1}{2} \mathbb{E}_{b_{k, c}}\left[f(h(b)) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(b_{k, c}\right)\right]+\frac{1}{2} \mathbb{E}_{\tilde{b}_{k, c}}\left[f(h(\tilde{b})) \nabla_{\alpha_{k, c}} \log q_{\theta}\left(\tilde{b}_{k, c}\right)\right]$, which immediately implies unbiasedness. Importantly, the conditions defining the estimator can be determined solely based on $b_{-k, c}$ and $\tilde{b}_{-k, c}$, so it suffices to verify that the estimator is unbiased for each case separately. In the first case, $g_{\text {DisARM-Tree }}$ is the DisARM estimator from (Dong et al. 2020) which was previously shown to be unbiased. The second and third cases are reminiscent of the 2 -sample RLOO estimator, however, in the coupled case, such an estimator must be justified as unbiased. In the second case, $f(z)$ depends on $\tilde{b}_{k, c}$ through the coupled sample $b_{k, c}$, however the condition $c \in I\left(\tilde{b}_{k, .}\right)-I\left(b_{k, \cdot}\right)$ implies that $z$ does not depend on $b_{k, c}$, so $f(z)$ is a constant with respect to both $b_{k, c}$ and $\tilde{b}_{k, c}$, resulting in the vanishing terms. Thus, we can apply the same argument as in the previous section. The third case follows by symmetry. In the fourth case, the condition $c \notin I\left(b_{k, .}\right) \cup I\left(\tilde{b}_{k, .}\right)$ implies that neither $z$ nor $\tilde{z}$ depend on $b_{k, c}$ or $\tilde{b}_{k, c}$ hence the gradient vanishes.

## A. 6 Rao-Blackwellized ARS \& ARSM

First, we briefly review the Augment-REINFORCE-Swap (ARS) \& Augment-REINFORCE-SwapMerge (ARSM) estimators (Yin et al., 2019). Yin et al. (2019) use the fact that the discrete distribution can be reparameterized by an underlying continuous augmentation: if $\pi \sim \prod_{k} \operatorname{Dirichlet}\left(1_{C}\right)$ and $z_{k}:=\arg \min _{i} \pi_{k, i} e^{-\alpha_{k, i}}$, then $z_{k} \sim \operatorname{Cat}\left(\alpha_{k}\right)$; and show that $\nabla_{\alpha_{k, c}} \mathbb{E}_{q_{\theta}(z)}[f(z)]=$ $\mathbb{E}_{\pi}\left[f(z)\left(1-C \pi_{k, c}\right)\right]$. Furthermore, they define a swapped probability matrix $\pi_{k}^{i \leftrightharpoons j}$ by swapping the entries at indices $i$ and $j$ in $\pi_{k}$

$$
\pi_{k, c}^{i \leftrightharpoons j}:= \begin{cases}\pi_{k, i} & c=j \\ \pi_{k, j} & c=i \\ \pi_{k, c} & \text { o.w. }\end{cases}
$$

and $z_{k}^{i \leftrightharpoons j}:=\arg \min _{c} \pi_{k, c}^{i \leftrightharpoons j} e^{-\alpha_{k, c}}$. Using these constructions, they show an important identity

$$
\nabla_{\alpha_{k, c}} \mathbb{E}_{q_{\theta}(z)}[f(z)]=\mathbb{E}_{\pi}\left[g_{\mathrm{ARS}_{k, c}}:=\left[f\left(z^{c \leftrightharpoons j}\right)-\frac{1}{C} \sum_{m=1}^{C} f\left(z^{m \leftrightharpoons j}\right)\right]\left(1-C \pi_{k, j}\right)\right]
$$

which shows that $g_{\text {ARS }}^{k, c}$ is an unbiased estimator. To further improve the estimator, Yin et al. (2019) average over the choice of the reference $j$, resulting in the ARSM estimator

$$
g_{\mathrm{ARSM}_{k, c}}:=\frac{1}{C} \sum_{j=1}^{C}\left[f\left(z^{c \leftrightharpoons j}\right)-\frac{1}{C} \sum_{m=1}^{C} f\left(z^{m \leftrightharpoons j}\right)\right]\left(1-C \pi_{k, j}\right) .
$$

Notably, both ARS and ARSM only evaluate $f$ at discrete values, and thus do not rely on a continuous relaxation.

## A.6.1 Rao-Blackwellization

Motivated by the approach of Dong et al. (2020), we can derive improved versions of ARS and ARSM by integrating out the extra randomness due to the continuous variables. ARS and ARSM heavily rely on a continuous reparameterization of the problem, yet the original problem only depends on the discrete values. Ideally, we would integrate out $\pi \mid z^{1 \leftrightharpoons j}, \ldots, z^{C \leftrightharpoons j}$, however, unlike in the binary case, computing the expectation analytically appears infeasible. Instead, we analytically integrate out two dimensions of $\pi$ and use Monte Carlo sampling to deal with the rest. This is a straightforward albeit tedious calculation

Starting with ARS, ideally, we would like to compute

$$
\begin{aligned}
& \mathbb{E}_{\pi \mid z^{1} \leftrightharpoons j, \ldots, z^{C \leftrightharpoons j}}\left[g_{\text {ARS }}^{k, c}\right]=\left[f\left(z^{c \leftrightharpoons j}\right)-\frac{1}{C} \sum_{m=1}^{C} f\left(z^{m \leftrightharpoons j}\right)\right]\left(1-C \mathbb{E}_{\pi \mid z^{1 \leftrightharpoons j}, \ldots, z^{C \leftrightharpoons j}}\left[\pi_{k, j}\right]\right) \\
& =\left[f\left(z^{c \leftrightharpoons j}\right)-\frac{1}{C} \sum_{m=1}^{C} f\left(z^{m \leftrightharpoons j}\right)\right]\left(1-C \mathbb{E}_{\pi_{k, j} \mid z_{k}^{1 \leftrightharpoons j}, \ldots, z_{k}^{C \leftrightharpoons j}}\left[\pi_{k, j}\right]\right),
\end{aligned}
$$

taking advantage of independence between dimensions (indexed by $k$ ). To reduce notational clutter, we omit the dimension index in the following derivation.
We have reduced the problem to computing

$$
\mathbb{E}_{\pi_{j} \mid z^{1 \leftrightharpoons j}, \ldots, z^{C \leftrightharpoons j}}\left[\pi_{j}\right],
$$

but were unable to compute the expectation analytically. Instead, we analytically integrate out some dimensions of $\pi$ and use Monte Carlo sampling to deal with the rest. First, we know that $\sum_{i} \pi_{i}=1$, so one variable is redundant (denote this choice by $l$ ). Next, we show how to integrate the reference index $j \neq l$ (i.e., compute $\mathbb{E}_{\pi_{j} \mid \pi_{-j, l}, z^{1 \leftrightharpoons j}, \ldots, z^{C \leftrightharpoons j}}\left[\pi_{j}\right]$, where $\pi_{-j, l}$ denotes $\pi$ excluding its $j$-th and $l$-th elements.).
The known values of $\pi_{-j, l}, z^{1 \leftrightharpoons j}, \ldots, z^{C \leftrightharpoons j}$ imply lower and upper bounds on $\pi_{j}$. First, because $1-\sum_{i \neq l} \pi_{i}=\pi_{l} \geq 0$, we conclude that $\pi_{j} \leq 1-\sum_{i \neq j, l} \pi_{i}$. To determine the implications of the configurations $z^{1 \leftrightharpoons j}, \ldots, z^{C \leftrightharpoons j}$, it is helpful to define some additional notation. Let $s^{c \leftrightharpoons j}:=\pi^{c \leftrightharpoons j} e^{-\alpha}$. Let's look at what the value of $z^{m \leftrightharpoons j}:=\arg \min _{i} s_{i}^{m \leftrightharpoons j}$ tells us about $\pi_{j}$. We need to consider two cases:

- $z^{m \leftrightharpoons j}=m$ : This means that $s_{m}^{m \leftrightharpoons j}=\pi_{j} e^{-\alpha_{m}}$ is the smallest entry in $s^{m \leftrightharpoons j}: \pi_{j} e^{-\alpha_{m}} \leq$ $\min _{i \neq m} s_{i}^{m \leftrightharpoons j}$, which implies that $\pi_{j} \leq \min _{i \neq m} e^{\alpha_{m}} s_{i}^{m \leftrightharpoons j}$.
$e^{\alpha_{m}} s_{i}^{m \leftrightharpoons j}$ contains $\pi_{l}$ when $m=l$ and $i=j$ or $m \neq l$ and $i=l$. When $m=l$ and $i=j$, we have that $e^{\alpha_{l}} s_{j}^{m \leftrightharpoons j}=e^{\alpha_{l}} \pi_{l} e^{-\alpha_{j}}=\left(1-\sum_{n \neq j, l} \pi_{n}-\pi_{j}\right) e^{\alpha_{l}-\alpha_{j}}$. Therefore,

$$
\pi_{j} \leq \frac{\left(1-\sum_{n \neq j, l} \pi_{n}\right) e^{-\alpha_{j}}}{e^{-\alpha_{j}}+e^{-\alpha_{l}}}
$$

A similar computation is required for the case $m \neq l$ and $i=l$.

- $z^{m \leftrightharpoons j} \neq m$ : This means that $\pi_{j} e^{-\alpha_{m}}$ is larger than the smallest entry in $s^{m \leftrightharpoons j}: \pi_{j} e^{-\alpha_{m}} \geq$ $\min _{i} s_{i}^{m \leftrightharpoons j}$ which implies that $\pi_{j} \geq \min _{i} e^{\alpha_{m}} s_{i}^{m \leftrightharpoons j}$. As above, we can eliminate $\pi_{l}$ from the bounds.

Finally, we aggregate the inequalities to compute the lower and upper bounds. Because $\pi \sim$ $\operatorname{Dirichlet}\left(1_{C}\right)$ is a uniform distribution over the simplex, $\pi_{j} \mid \pi_{-j, l}, z^{1 \leftrightharpoons j}, \ldots, z^{C \leftrightharpoons j}$ will be uniformly
distributed over an interval, which means that it suffices to compute the lower and upper bounds to compute the expectation.

We can apply the same ideas to ARSM, however, in preliminary experiments with ARS+, we found that leveraging the symmetry (described next) was responsible for most of the performance improvement for ARS+. So, for ARSM, we reduce the variance only by leveraging the symmetry and call the resulting estimator ARSM+.

Furthermore, When all of the swapped $z$ s agree on a dimension (i.e., $z_{k}^{1 \leftrightharpoons j}=\cdots=z_{k}^{C \leftrightharpoons j}$ ), then we will show that both $g_{\text {ARS }}^{k, c}$ and $g_{\operatorname{ARSM}}^{k, c}$ vanish in expectation, so we can zero out these terms explicitly. The high level intuition is that even though they may disagree in other dimensions for a single sample because the other dimensions are independent and expectations are linear, in expectation they cancel out. Let $\delta_{k}=\mathbb{1}_{z_{k}^{1 \leftrightharpoons j}=\cdots=z_{k}^{C \leftrightharpoons j}}$. Then, we have

$$
\begin{aligned}
\mathbb{E}_{\pi \mid \delta_{k}=1}\left[g_{\mathrm{ARS}_{k, c}}\right] & =\mathbb{E}_{\pi \mid \delta_{k}}\left[\left[f\left(z^{c \leftrightharpoons j}\right)-\frac{1}{C} \sum_{m=1}^{C} f\left(z^{m \leftrightharpoons j}\right)\right]\left(1-C \pi_{k, j}\right)\right] \\
& =\mathbb{E}_{\pi_{k} \mid \delta_{k}=1}\left[\left(\mathbb{E}_{\pi_{-k}}\left[f\left(z^{c \leftrightharpoons j}\right)\right]-\frac{1}{C} \sum_{m=1}^{C} \mathbb{E}_{\pi_{-k}}\left[f\left(z^{m \leftrightharpoons j}\right)\right]\right)\left(1-C \pi_{k, j}\right)\right] .
\end{aligned}
$$

Now, we claim that inside the expectation $\mathbb{E}_{\pi_{-k}}\left[f\left(z^{m \leftrightharpoons j}\right)\right]$ is constant with respect to $m$. First, we know that inside the expectation $z_{k}^{1 \leftrightharpoons j}=\cdots=z_{k}^{C \leftrightharpoons j}$ and that the dimensions indexed by $k$ are independent. Because $\pi \sim \prod_{k}$ Dirichlet $\left(1_{C}\right)$, Dirichlet $\left(1_{C}\right)$ is symmetric, and we are taking an unconditional expectation over the remaining dimensions, the value is invariant to the swapping operation. As a result, the entire expression vanishes. Thus, we conclude that

$$
\left(1-\delta_{k}\right) g_{\text {ARS }_{k, c}}
$$

is still an unbiased estimator. A similar argument holds for $g_{\text {ARSM }}$. This is complementary to the approach in the previous subsection and can done in combination

$$
g_{\mathrm{ARS}+k, c}:=\mathbb{E}_{\pi_{k, j} \mid \pi_{k,-j l}, z_{k}^{1 \leftrightharpoons j}, \ldots, z_{k}^{C \leftrightharpoons j}}\left[\left(1-\delta_{k}\right) g_{\mathrm{ARS}_{k, c}}\right],
$$

where we choose $l \neq j$ uniformly at random. This is the estimator we use in our experiments.

## A.6.2 Evaluating Rao-Blackwellized ARS \& ARSM

We train models with 10/5/3/2-category latent variables on dynamically binarized MNIST. For comparison, we train models with ARS, ARSM, and an $n$-sample RLOO. To match computation, RLOO uses $C$ samples for comparing against ARS/ARS+, and uses $C(C-1) / 2+1$ samples for ARSM/ARSM + , where $C$ is the number of categories. Based on preliminary experiments with ARS+, we found that leveraging the symmetry led to most of the performance improvement for ARS+. So, for ARSM, we reduce the variance only by leveraging the symmetry and call the resulting estimator ARSM+.

As shown in Figure 8 and Appendix Figure 9 the proposed estimators, ARS+/ARSM+, significantly outperform ARS/ARSM. Surprisingly, we find that both ARS and ARSM underperform the simpler RLOO baseline in all cases. For $C=2$, ARS+ and ARSM+ reduce to DisARM/U2G and as expected, outperform REINFORCE LOO; however, for $C>2$, REINFORCE LOO is superior and the gap increases as $C$ does. This suggests that partially integrating out the randomness is insufficient to account for the variance introduced by the continuous augmentation.


Figure 8: Training a non-linear categorical VAE with latent variables with 5 categories on dynamically binarized MNIST dataset by maximizing the ELBO. We plot the train ELBO (left column), test 100 -sample bound (middle column), and the variance of gradient estimator (right column). We plot the mean and one standard error based on 5 runs from different random initializations.


Figure 9: Training a non-linear categorical VAE with latent variables with 10/5/3/2 categories on dynamically binarized MNIST dataset by maximizing the ELBO. We plot the train ELBO (left column), the test 100-sample bound (middle column), and the variance of gradient estimator (right column). We plot the mean and one standard error based on 5 runs from different random initializations.

