

345 **A Theorem 3.1**

346 Let us define the first-order approximation of ∇ as $\widehat{\nabla}_{1st-order} = \sum_i \sum_j \pi_j \frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} (\mathbf{I}_i - \mathbf{I}_j) \frac{d\pi_i}{d\boldsymbol{\theta}}$,
 347 which approximates $f(\mathbf{I}_i) - f(\mathbf{I}_j)$ in Equation 6 as $\frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} (\mathbf{I}_i - \mathbf{I}_j)$.

Theorem 3.1.

$$E[\widehat{\nabla}_{ST}] = \widehat{\nabla}_{1st-order}.$$

348 *Proof.* Based on the definition, we have:

$$\begin{aligned} \widehat{\nabla}_{1st-order} &= \sum_i \sum_j \pi_j \frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} (\mathbf{I}_i - \mathbf{I}_j) \frac{d\pi_i}{d\boldsymbol{\theta}} \\ &= \sum_j \pi_j \frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} \sum_i \mathbf{I}_i \frac{d\pi_i}{d\boldsymbol{\theta}} - \sum_j \pi_j \frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} \mathbf{I}_j \sum_i \frac{d\pi_i}{d\boldsymbol{\theta}} \end{aligned} \quad (9)$$

349 Since $\sum_i \pi_i = 1$, we have $\sum_i \frac{d\pi_i}{d\boldsymbol{\theta}} = 0$. Also, since $\boldsymbol{\pi} = \sum_i \pi_i \mathbf{I}_i$, we have $\frac{d\boldsymbol{\pi}}{d\boldsymbol{\theta}} = \sum_i \mathbf{I}_i \frac{d\pi_i}{d\boldsymbol{\theta}}$. Thus,
 350 together with Equation 9, we have:

$$\begin{aligned} \widehat{\nabla}_{1st-order} &= \sum_j \pi_j \frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} \sum_i \mathbf{I}_i \frac{d\pi_i}{d\boldsymbol{\theta}} \\ &= E\left[\frac{\partial f(\mathbf{D})}{\partial \mathbf{D}} \frac{d\boldsymbol{\pi}}{d\boldsymbol{\theta}}\right] = E[\widehat{\nabla}_{ST}]. \end{aligned}$$

351 □

352 **B Theorem 3.2**

Theorem 3.2.

$$E[\widehat{\nabla}_{ReinMax}] = \widehat{\nabla}_{2rd-order}.$$

353 *Proof.* Here, we aim to proof, $\forall k \in [1, n]$, we have $E[\widehat{\nabla}_{ReinMax, k}] = \widehat{\nabla}_{2rd-order, k}$. As defined in
 354 Equation 8, we have

$$\begin{aligned} \widehat{\nabla}_{2rd-order, k} &= \sum_i \sum_j \frac{\pi_j}{2} \left(\frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} + \frac{\partial f(\mathbf{I}_i)}{\partial \mathbf{I}_i} \right) (\mathbf{I}_i - \mathbf{I}_j) \frac{d\pi_i}{d\boldsymbol{\theta}_k} \\ &= \sum_i \sum_j \frac{\pi_j \pi_i (\delta_{ik} - \pi_k)}{2} \left(\frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} + \frac{\partial f(\mathbf{I}_i)}{\partial \mathbf{I}_i} \right) (\mathbf{I}_i - \mathbf{I}_j) \\ &= \sum_j \frac{\pi_j \pi_k}{2} \left(\frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} + \frac{\partial f(\mathbf{I}_k)}{\partial \mathbf{I}_k} \right) (\mathbf{I}_k - \mathbf{I}_j) \\ &= \frac{\pi_k}{2} \frac{\partial f(\mathbf{I}_k)}{\partial \mathbf{I}_k} (\mathbf{I}_k - \sum_j \pi_j \mathbf{I}_j) + \sum_j \frac{\pi_j \pi_k}{2} \frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} (\mathbf{I}_k - \mathbf{I}_j) \\ &= \frac{1}{2} E[\delta_{Dk} \frac{\partial f(\mathbf{D})}{\partial \mathbf{D}} (\mathbf{I}_D - \sum_j \pi_j \mathbf{I}_j)] + \frac{1}{2} E[\pi_k \frac{\partial f(\mathbf{D})}{\partial \mathbf{D}} (\mathbf{I}_k - \mathbf{I}_D)] \\ &= \frac{1}{2} E\left[\frac{\partial f(\mathbf{D})}{\partial \mathbf{D}} (\pi_k (\mathbf{I}_k - \mathbf{I}_D) + \delta_{Dk} (\mathbf{I}_D - \sum_i \pi_i \mathbf{I}_i))\right] \end{aligned} \quad (10)$$

355 At the same time, based on the definition of $\widehat{\nabla}_{ReinMax}$, we have:

$$\begin{aligned} E[\widehat{\nabla}_{ReinMax, k}] &= E\left[\frac{\partial f(\mathbf{D})}{\partial \mathbf{D}} \left(2 \cdot \frac{\pi_k + \delta_{Dk}}{2} (\mathbf{D}_k - \sum_i \frac{\pi_i + \delta_{Dk}}{2} \mathbf{I}_i) - \frac{\pi_k}{2} (\mathbf{D}_k - \sum_i \pi_i \mathbf{I}_i) \right)\right] \\ &= \frac{1}{2} E\left[\frac{\partial f(\mathbf{D})}{\partial \mathbf{D}} (\pi_k (\mathbf{I}_k - \mathbf{I}_D) + \delta_{Dk} (\mathbf{I}_k - \sum_i \pi_i \mathbf{I}_i))\right] \end{aligned} \quad (11)$$

356 Since $\delta_{Dk}(\mathbf{I}_k - \sum_i \pi_i \mathbf{I}_i) = \delta_{Dk}(\mathbf{I}_D - \sum_i \pi_i \mathbf{I}_i)$, together with Equation 10 and 11, we have:

$$E[\widehat{\nabla}_{\text{ReinMax},k}] = \widehat{\nabla}_{2\text{rd-order},k}$$

357

□

358 C Remark 4.1

359 **Remark 3.1.** When $\sum_i \phi_i f(\mathbf{I}_i)$ is used as the baseline and $f(\mathbf{I}_i) - f(\mathbf{I}_j)$ is approximated as
 360 $\frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j}(\mathbf{I}_i - \mathbf{I}_j)$, we mark the resulting first-order approximation of ∇ as $\widehat{\nabla}_{1\text{st-order-avg-baseline}}$.
 361 Then, we have:

$$E\left[\frac{\phi_D}{\pi_D} \widehat{\nabla}_{\text{ST}}\right] = \widehat{\nabla}_{1\text{st-order-avg-baseline}}$$

362 *Proof.* Using $\sum_i \phi_i f(\mathbf{I}_i)$ as the baseline, we have:

$$\nabla = \sum_i (f(\mathbf{I}_i) - \sum_j \phi_j f(\mathbf{I}_j)) \frac{d\pi_i}{d\theta} = \sum_i \sum_j \phi_j (f(\mathbf{I}_i) - f(\mathbf{I}_j)) \frac{d\pi_i}{d\theta}$$

363 Approximating $f(\mathbf{I}_i) - f(\mathbf{I}_j)$ as $\frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j}(\mathbf{I}_i - \mathbf{I}_j)$, we have:

$$\begin{aligned} \widehat{\nabla}_{1\text{st-order-avg-baseline}} &= \sum_i \sum_j \phi_j \frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j}(\mathbf{I}_i - \mathbf{I}_j) \frac{d\pi_i}{d\theta} \\ &= \sum_j \frac{\phi_j}{\pi_j} \cdot \pi_j \frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} \sum_i \mathbf{I}_i \frac{d\pi_i}{d\theta} \\ &= E\left[\frac{\phi_D}{\pi_D} \widehat{\nabla}_{\text{ST}}\right] \end{aligned}$$

364

□

365 D Remark 4.2

366 **Remark 3.2.** In Equation 8, we approximate $f(\mathbf{I}_k) - f(\mathbf{I}_i)$ as $\frac{1}{2}(\frac{\partial f(\mathbf{I}_i)}{\partial \mathbf{I}_i} + \frac{\partial f(\mathbf{I}_k)}{\partial \mathbf{I}_k})(\mathbf{I}_k - \mathbf{I}_i)$, and mark
 367 the resulting second-order approximation of ∇_k as $\widehat{\nabla}_{2\text{rd-order-wo-baseline},k} = \pi_k \sum_i \pi_i \frac{1}{2}(\frac{\partial f(\mathbf{I}_i)}{\partial \mathbf{I}_i} +$
 368 $\frac{\partial f(\mathbf{I}_k)}{\partial \mathbf{I}_k})(\mathbf{I}_k - \mathbf{I}_i)$, Then, we have:

$$E[\widehat{\nabla}_{\text{ReinMax}}] = \widehat{\nabla}_{2\text{rd-order-wo-baseline}}$$

369 *Proof.* Here, we aim to proof, $\forall k \in [1, n]$, we have $E[\widehat{\nabla}_{\text{ReinMax},k}] = \widehat{\nabla}_{2\text{rd-order-wo-baseline},k}$.

$$\begin{aligned} \widehat{\nabla}_{2\text{rd-order-wo-baseline},k} &= \pi_k \sum_i \pi_i \frac{1}{2} \left(\frac{\partial f(\mathbf{I}_i)}{\partial \mathbf{I}_i} + \frac{\partial f(\mathbf{I}_k)}{\partial \mathbf{I}_k} \right) (\mathbf{I}_k - \mathbf{I}_i) \\ &= \pi_k \sum_i \pi_i \frac{1}{2} \frac{\partial f(\mathbf{I}_i)}{\partial \mathbf{I}_i} (\mathbf{I}_k - \mathbf{I}_i) + \pi_k \sum_i \pi_i \frac{1}{2} \frac{\partial f(\mathbf{I}_k)}{\partial \mathbf{I}_k} (\mathbf{I}_k - \mathbf{I}_i) \\ &= E\left[\frac{\partial f(\mathbf{D})}{\partial \mathbf{D}} \frac{\pi_k(\mathbf{I}_k - \mathbf{I}_D) + \delta_{Dk}(\mathbf{I}_k - \sum_i \pi_i \mathbf{I}_i)}{2}\right] = E[\widehat{\nabla}_{\text{ReinMax},k}] \end{aligned}$$

370

□

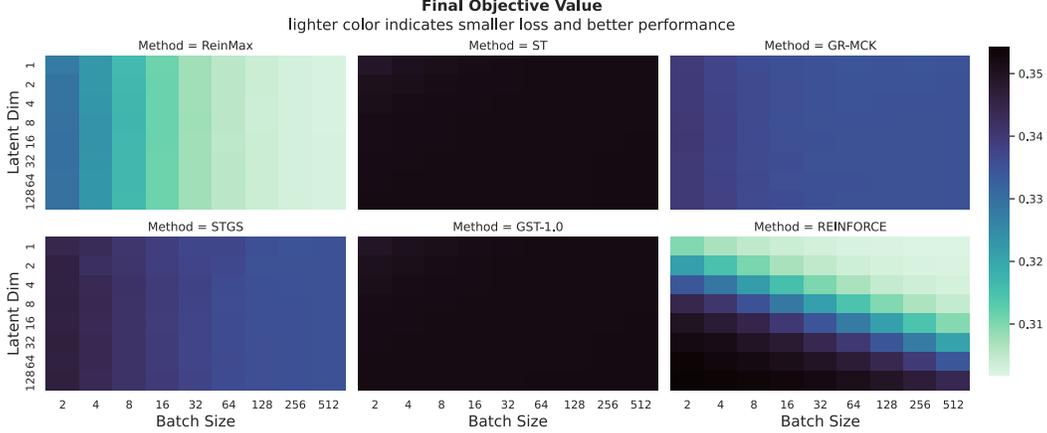


Figure 7: Polynomial programming loss after 40 epochs, with different batch sizes and random variable counts (L), i.e., $\min_{\theta} E[\frac{\|\mathbf{X}-c\|_{1.5}^5}{L}]$, where $\theta \in \mathcal{R}^{L \times 2}$, $\mathbf{X} \in \{0, 1\}^L$, and $\mathbf{X}_i \stackrel{\text{iid}}{\sim} \text{Multinomial}(\text{softmax}(\theta_i))$. More details are elaborated in Section 6.

371 E Forward Euler Method and Heun’s Method

372 For simplicity, we consider a simple function $g(x) : \mathcal{R} \rightarrow \mathcal{R}$ that is three times differentiable on
 373 $[t_0, t_1]$. Now, we proceed to a simple introduction to approximate $\int_{t_0}^{t_1} g'(x)dx$ with the Forward
 374 Euler Method and the Heun’s Method. For a detailed introduction to numerical ODE methods, please
 375 refer to Ascher & Petzold (1998).

376 **Forward Euler Method.** Here, we approximate $g(t_1)$ with the first-order Taylor expansion, i.e.,
 377 $g(t_1) = g(t_0) + g'(t_0) \cdot (t_1 - t_0) + O((t_1 - t_0)^2)$, then we have $\int_{t_0}^{t_1} g'(x)dx \approx g'(t_0)(t_1 - t_0)$.
 378 Since we used the first-order Taylor expansion, this approximation has first-order accuracy.

379 **Heun’s Method.** First, we approximate $g(t_1)$ with the second-order Taylor expansion:

$$g(t_1) = g(t_0) + g'(t_0) \cdot (t_1 - t_0) + \frac{g''(t_0)}{2} \cdot (t_1 - t_0)^2 + O((t_1 - t_0)^3). \quad (12)$$

380 Then, we show that we can match this approximation by combining the first-order derivatives of two
 381 samples. Taylor expanding $g'(t_1)$ to the first-order, we have:

$$g'(t_1) = g'(t_0) + g''(t_0) \cdot (t_1 - t_0) + O((t_1 - t_0)^2)$$

382 Therefore, we have:

$$g(t_0) + \frac{g'(t_0) + g'(t_1)}{2} (t_1 - t_0) = g(t_0) + g'(t_0) \cdot (t_1 - t_0) + \frac{g''(t_0)}{2} \cdot (t_1 - t_0)^2 + O((t_1 - t_0)^3).$$

383 It is easy to notice that the right-hand side of the above equation matches the second-order Taylor
 384 expansion of $g(t_1)$ as in Equation 12. Therefore, the above approximation (i.e., approximating
 385 $g(t_1) - g(t_0)$ as $\frac{g'(t_0) + g'(t_1)}{2} (t_1 - t_0)$) has second-order accuracy.

386 **Connection to $f(\mathbf{I}_i) - f(\mathbf{I}_j)$ in Equation 6.** By setting $g(x) = f(x \cdot \mathbf{I}_i + (1 - x) \cdot \mathbf{I}_j)$, we have
 387 $g(1) - g(0) = f(\mathbf{I}_i) - f(\mathbf{I}_j)$. Then, it is easy to notice that the forward Euler Method approximates
 388 $f(\mathbf{I}_i) - f(\mathbf{I}_j)$ as $\frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j} (\mathbf{I}_i - \mathbf{I}_j)$ and has first-order accuracy. Also, the Heun’s Method approximates
 389 $f(\mathbf{I}_i) - f(\mathbf{I}_j)$ as $\frac{1}{2} (\frac{\partial f(\mathbf{I}_i)}{\partial \mathbf{I}_i} + \frac{\partial f(\mathbf{I}_j)}{\partial \mathbf{I}_j}) (\mathbf{I}_i - \mathbf{I}_j)$ and has second-order accuracy.

390 F Experiment Details

391 F.1 Baselines

392 Here, we consider four methods as our major baselines:

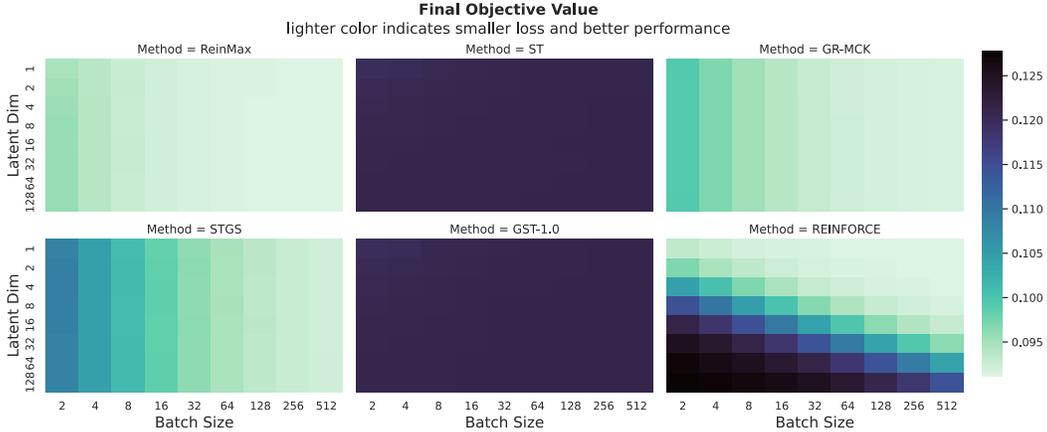


Figure 8: Polynomial programming loss after 40 epochs, with different batch sizes and random variable counts (L), i.e., $\min_{\theta} E[\frac{\|\mathbf{X}-\mathbf{c}\|_3^3}{L}]$ where $\theta \in \mathcal{R}^{L \times 2}$, $\mathbf{X} \in \{0, 1\}^L$, and $\mathbf{X}_i \stackrel{\text{iid}}{\sim} \text{Multinomial}(\text{softmax}(\theta_i))$. More details are elaborated in Section 6.

- 393 • Straight-Through (ST; Bengio et al., 2013) backpropagate through the sampling function as if it
394 had been the identity function.
- 395 • Straight-Through Gumbel-Softmax (STGS; Jang et al., 2017) integrates the Gumbel reparameteri-
396 zation trick to approximate the gradient.
- 397 • Gumbel-Rao Monte Carlo (GR-MCK; Paulus et al., 2021) leverages the Monte Carlo method to
398 reduce the variance introduced by the Gumbel noise in STGS. To obtain the optimal performance
399 for this baseline, we set the number of Monte Carlo samples to 1000 in most experiments. Except
400 in our discussions of efficiency, we set the number of Monte Carlo samples to 100, 300, and 1000
401 for a more comprehensive comparisons.
- 402 • Gapped Straight-Through (GST-1.0; Fan et al., 2022) aims to reduce the variance of STGS and
403 constructs a deterministic term to replace the Monte Carlo samples used in GR-MCK. Here, as
404 suggested in (Fan et al., 2022), we set the gap (a hyper-parameter) as 1.0.

405 **GST-1.0 Performance.** Despite GST-1.0 achieving good performance on most settings of MNIST-
406 VAE, it fails to maintain this performance on polynomial programming and unsupervised parsing, as
407 discussed before. At the same time, a different variant of GST (i.e., GST-p) achieves a significant
408 performance boost over GST-1.0 on polynomial programming. However, on MNIST-VAE and
409 ListOps, GST-p achieves an inferior performance. Upon discussing with the author of the GST-1.0,
410 we suggest that this phenomenon is caused by different characteristics of GST-1.0 and GST-p.

411 This observation verifies our intuition that, without understanding the mechanism of ST, different
412 applications have different preferences on its configurations. Meanwhile, ReinMax achieves consistent
413 improvements in all settings, which greatly simplifies future algorithms developments.

414 F.2 Hyper-Parameters

415 Without specifically, we conduct full grid search for all methods in all experiments, and report the
416 best performance (averaged with 10 random seeds on MNIST-VAE and 5 random seeds on ListOps).
417 The hyper-parameter search space is summarized in Table 5.

Table 5: Hyper-parameter search space.

Hyperparameters	Search Space
Optimizer	{Adam(Kingma & Ba, 2015), RAdam(Liu et al., 2020)}
Learning Rate	{0.001, 0.0007, 0.0005, 0.0003}
Temperature	{0.1, 0.3, 0.5, 0.7, 1.0, 1.1, 1.2, 1.3, 1.4, 1.5}

Table 6: Test –ELBO on MNIST. Hyper-parameters are chosen based on Train –ELBO.

	AVG	8 × 4	4 × 24	8 × 16	16 × 12	64 × 8	10 × 30
STGS	106.89	128.09±0.79	103.60±0.45	99.32±0.33	102.49±0.32	106.20±0.46	101.61±0.54
GR-MCK	109.03	127.90±0.71	102.76±0.33	102.12±0.29	104.23±0.65	113.54±0.50	103.62±0.13
GST-1.0	106.85	128.20±1.12	103.95±0.49	101.44±0.32	101.28±0.59	105.44±0.62	100.78±0.44
ST	118.85	137.06±0.51	113.41±0.49	114.25±0.29	114.48±0.56	115.43±0.29	118.46±0.18
ReinMax	105.74	126.89±0.79	102.40±0.43	100.63±0.41	100.85±0.50	102.91±0.67	100.75±0.50

Table 7: Test –ELBO on MNIST. Hyper-parameters are chosen based on Test –ELBO.

	AVG	8 × 4	4 × 24	8 × 16	16 × 12	64 × 8	10 × 30
STGS	107.15	128.09±0.79	103.25±0.22	101.44±0.32	102.29±0.39	106.20±0.46	101.61±0.54
GR-MCK	108.87	127.86±0.54	102.40±0.37	101.59±0.22	104.22±0.63	113.54±0.50	103.62±0.13
GST-1.0	106.55	128.03±1.02	103.63±0.24	100.67±0.34	101.04±0.39	105.44±0.62	100.51±0.37
ST	118.79	137.05±0.36	113.23±0.43	114.11±0.31	114.48±0.56	115.43±0.29	118.46±0.18
ReinMax	105.60	126.29±0.32	102.40±0.43	100.45±0.26	100.84±0.56	102.91±0.68	100.69±0.48

Table 8: Train –ELBO on MNIST. Hyper-parameters are chosen based on Test –ELBO.

	AVG	8 × 4	4 × 24	8 × 16	16 × 12	64 × 8	10 × 30
STGS	105.31	126.85±0.85	101.81±0.14	99.32±0.33	100.22±0.47	104.02±0.41	99.63±0.63
GR-MCK	107.37	126.53±0.55	100.47±0.31	99.75±0.29	103.11±0.58	112.34±0.48	102.02±0.18
GST-1.0	104.60	126.63±1.16	102.11±0.24	98.40±0.34	98.76±0.41	102.53±0.57	99.14±0.30
ST	117.76	136.75±0.22	112.09±0.50	113.06±0.26	113.31±0.43	113.90±0.28	117.46±0.09
ReinMax	103.40	124.92±0.38	99.77±0.45	98.06±0.31	98.51±0.54	100.71±0.70	98.40±0.48

418 **Polynomial Programming.** As this problem is relatively simple, we set the learning rate to 0.001
419 and the optimizer to Adam, and only tune the temperature hyper-parameter.

420 **MNIST-VAE.** Following the previous study (Dong et al., 2020, 2021; Fan et al., 2022), we used
421 2-layer MLP as the encoder and the decoder. We set the hidden state dimension of the first-layer
422 and the second-layer as 512 and 256 for the encoder, and 256 and 512 for the decoder. For our
423 experiments on MNIST-VAE with 32 latent dimensions and 64 categorical dimensions, we set the
424 batch size to 200, training steps to 5×10^5 , and activation function to LeakyReLU, in order to be
425 consistent with the literature. For other experiments, we set the batch size to 100, the activation
426 function to ReLU, and training steps to 9.6×10^4 (i.e., 160 epochs).

427 **ListOps.** We followed the same setting of Fan et al. (2022), i.e., used the same model configuration
428 as in Choi et al. (2017) and set the maximum sequence length to 100.

429 F.3 Hardware and Environment Setting

430 Most experiments (except efficiency comparisons) are conducted on Nvidia P40 GPUs. For efficiency
431 comparisons, we measured the average time cost per batch and peak memory consumption on
432 quadratic programming and MNIST-VAE on the same system with an idle A6000 GPU. Also, to
433 better reflect the efficiency of gradient estimators, we skipped all parameter updates in this experiment.

434 F.4 Additional Results on Polynomial Programming

435 Here, we visualized the heat map for polynomial programming with various batch sizes and latent
436 dimensions in Figure 7 (for $p = 1.5$) and Figure 8 (for $p = 3$). We visualized the training curve for
437 polynomial programming with various batch sizes and latent dimensions in Figure 9 (for $p = 1.5$),
438 Figure 10 (for $p = 2$), and Figure 11 (for $p = 3$).

439 F.5 Additional Results on MNIST-VAE

440 In our discussions in Section 6, we focused on the training ELBO only. Here, we provide a brief
441 discussion on the test ELBO.

442 **Choosing Hyper-parameter Based on Training Performance.** Similar to Table 2, for each
 443 method, we select the hyper-parameter based on its training performance. The Test –ELBO in this
 444 setting is summarized in 6. Despite the model being trained without dropout or other overfitting
 445 reduction techniques, ReinMax maintained the best performance in this setting.

446 **Choosing Hyper-parameter Based on Test Performance.** We also conduct experiments by
 447 selecting hyper-parameters directly based on their test performance. In this setting, the test –ELBO
 448 is summarized in Table 7, and the training –ELBO is summarized in Table 8. ReinMax achieves
 449 the best performance in all settings except the test performance of the setting with 10 categorical
 450 dimensions and 30 latent dimensions.

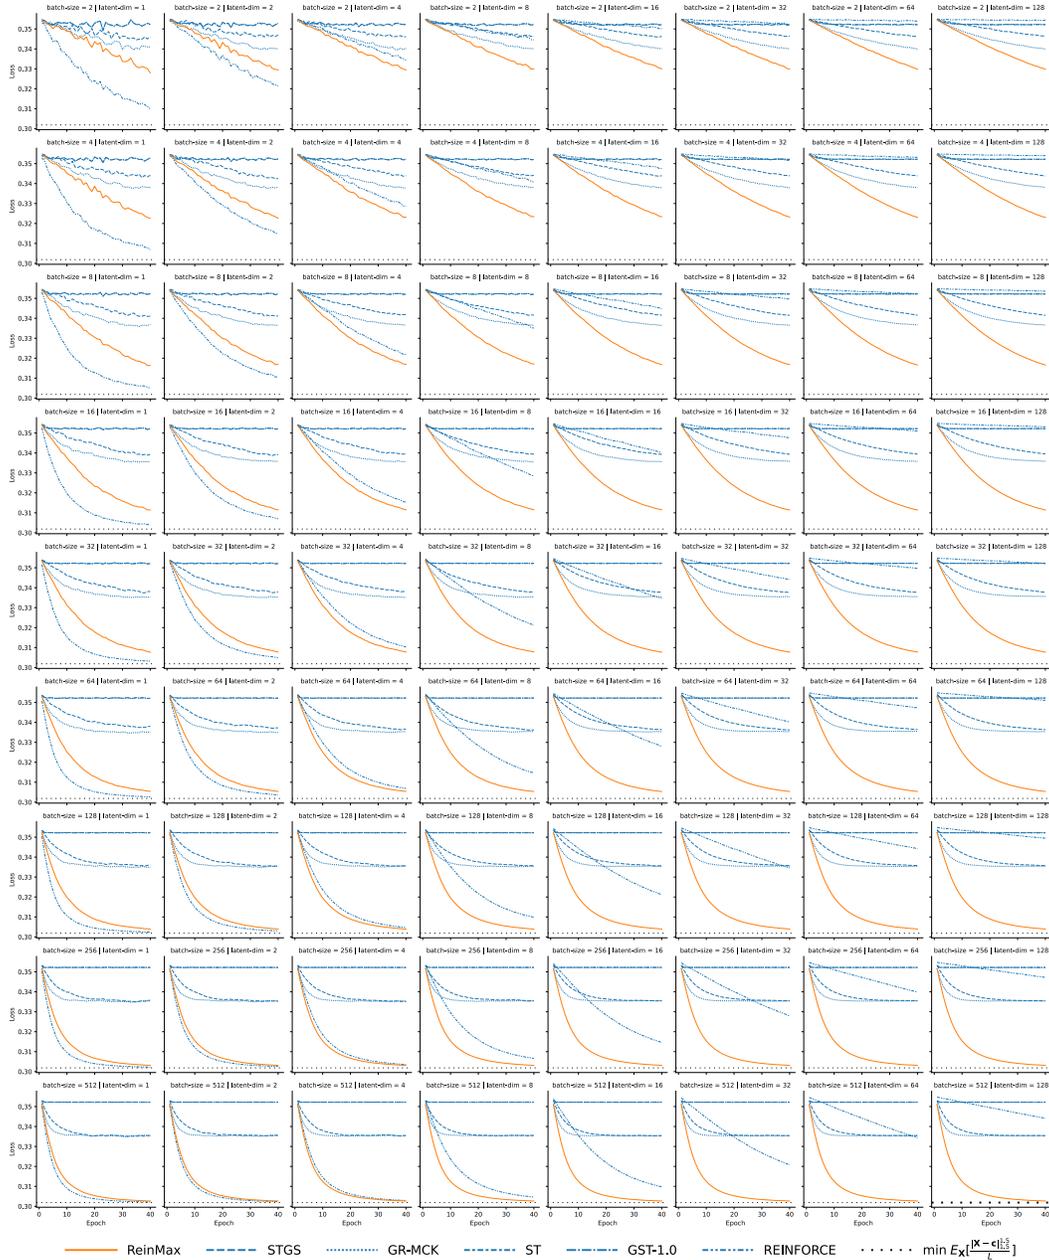


Figure 9: Polynomial programming training curve, with different batch sizes and random variable counts (L), i.e., $\min_{\theta} E[\frac{\|\mathbf{X}-\mathbf{c}\|_{1.5}}{L}]$, where $\theta \in \mathcal{R}^{L \times 2}$, $\mathbf{X} \in \{0, 1\}^L$, and $\mathbf{X}_i \stackrel{iid}{\sim} \text{Multinomial}(\text{softmax}(\theta_i))$. More details are elaborated in Section 6.

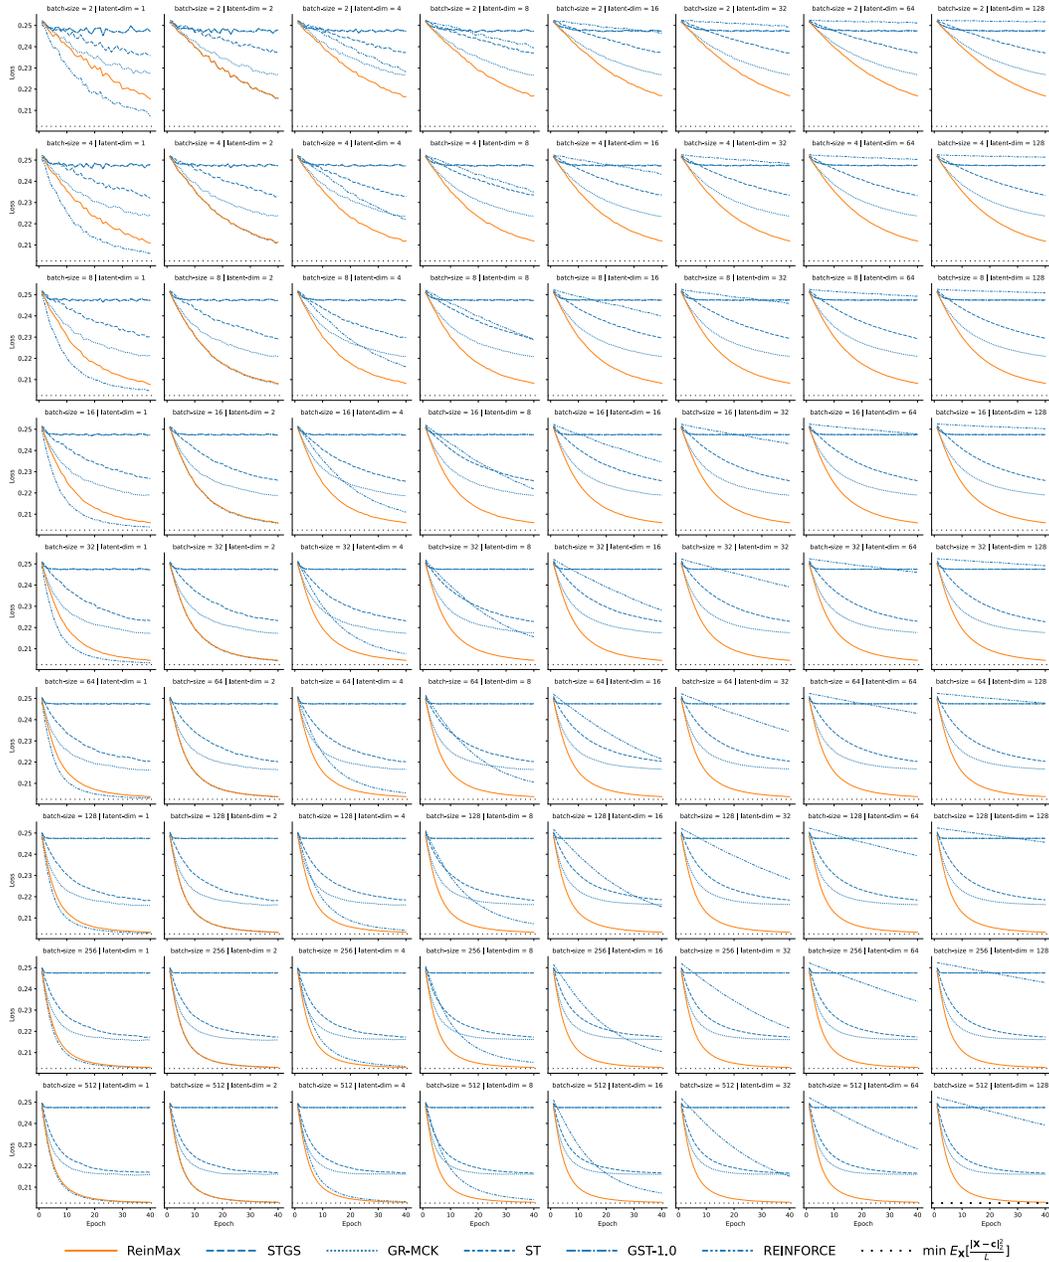


Figure 10: Quadratic programming training curve, with different batch sizes and random variable counts (L), i.e., $\min_{\theta} E[\frac{\|\mathbf{X} - \mathbf{c}\|_2^2}{L}]$, where $\theta \in \mathcal{R}^{L \times 2}$, $\mathbf{X} \in \{0, 1\}^L$, and $\mathbf{X}_i \stackrel{\text{iid}}{\sim} \text{Multinomial}(\text{softmax}(\theta_i))$. More details are elaborated in Section 6.

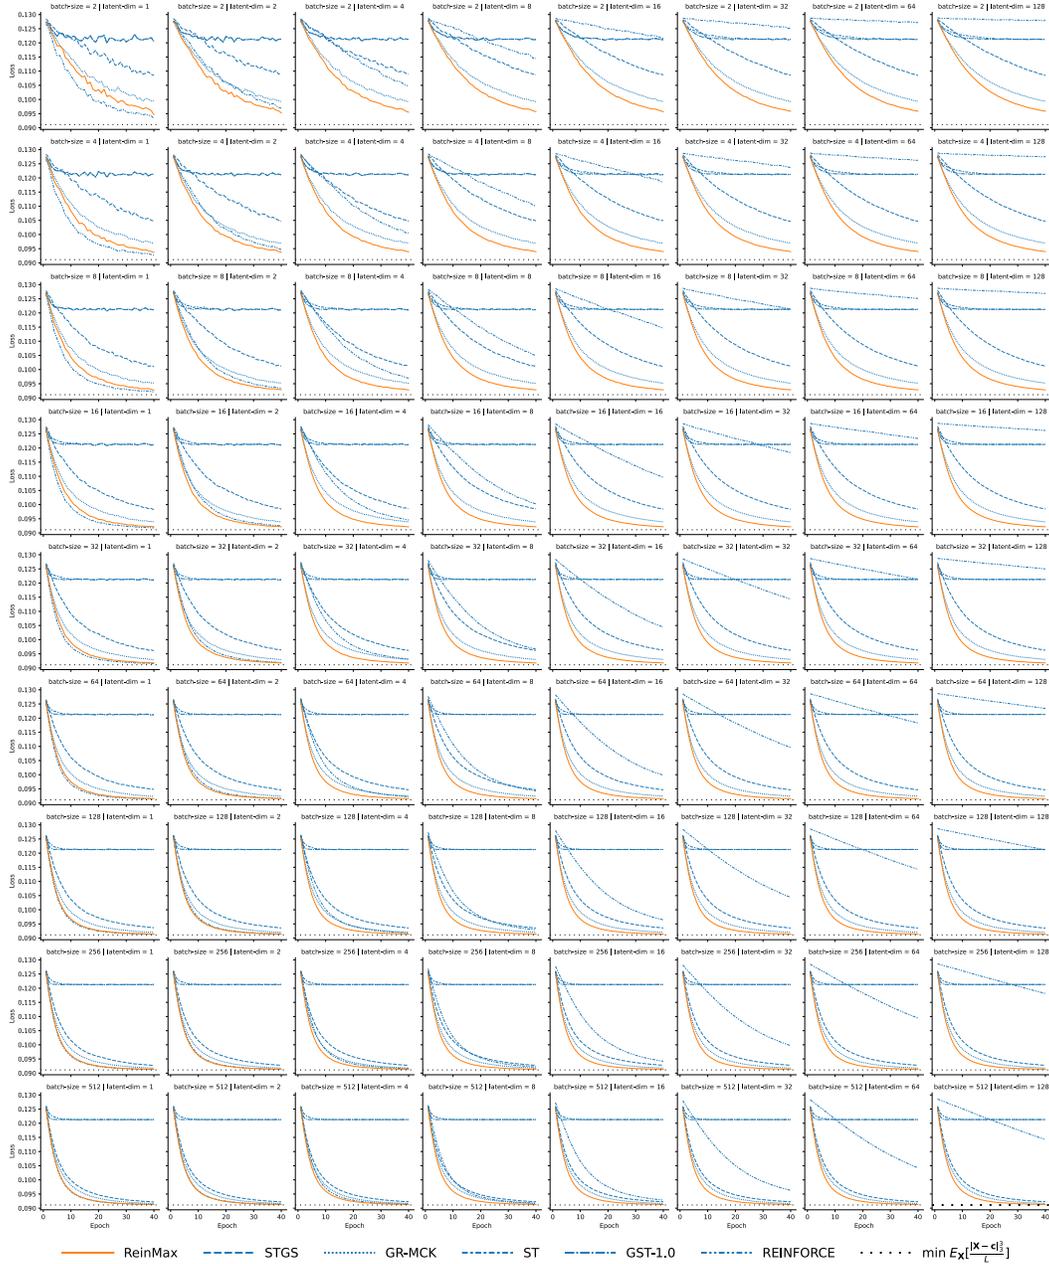


Figure 11: Polynomial programming training curve, with different batch sizes and random variable counts (L), i.e., $\min_{\theta} E[\frac{\|\mathbf{X}-\mathbf{c}\|_3^3}{L}]$, where $\theta \in \mathcal{R}^{L \times 2}$, $\mathbf{X} \in \{0, 1\}^L$, and $\mathbf{X}_i \stackrel{\text{iid}}{\sim} \text{Multinomial}(\text{softmax}(\theta_i))$. More details are elaborated in Section 6.