Auto-regressive Image Synthesis with Integrated Quantization

Fangneng Zhan^{1,2}, Yingchen Yu¹, Rongliang Wu¹, Jiahui Zhang¹, Kaiwen Cui¹, Changgong Zhang³, and Shijian Lu¹[™]

¹ Nanyang Technological University, Singapore

² Max Planck Institute for Informatics, Germany

 3 Amazon

fzhan@mpi-inf.mpg.de, shijian.lu@ntu.edu.sg, cgzhang@amazon.com
{yingchen001,ronglian001,jiahui003,kaiwen001}@e.ntu.edu.sg

Abstract. Deep generative models have achieved conspicuous progress in realistic image synthesis with multifarious conditional inputs, while generating diverse yet high-fidelity images remains a grand challenge in conditional image generation. This paper presents a versatile framework for conditional image generation which incorporates the inductive bias of CNNs and powerful sequence modeling of auto-regression that naturally leads to diverse image generation. Instead of independently quantizing the features of multiple domains as in prior research, we design an integrated quantization scheme with a variational regularizer that mingles the feature discretization in multiple domains, and markedly boosts the auto-regressive modeling performance. Notably, the variational regularizer enables to regularize feature distributions in incomparable latent spaces by penalizing the intra-domain variations of distributions. In addition, we design a Gumbel sampling strategy that allows to incorporate distribution uncertainty into the auto-regressive training procedure. The Gumbel sampling substantially mitigates the exposure bias that often incurs misalignment between the training and inference stages and severely impairs the inference performance. Extensive experiments over multiple conditional image generation tasks show that our method achieves superior diverse image generation performance qualitatively and quantitatively as compared with the state-of-the-art.

1 Introduction

Conditional image generation aims to generate photorealistic images conditioning on certain guidance which can be semantic segmentation [35], key points [43], layout [18] as well as heterogeneous guidance such as text [39] and audio [2]. It has been widely formulated as one-to-one mapping tasks [48], though it is essentially one-to-many mappings since one conditional input could correspond to multiple images. Targeting to mimic the true conditional image distribution, diverse yet high-fidelity image synthesis remains a great challenge in conditional

^{*} Corresponding author

image generation, especially when the conditional inputs come from different visual domains or even heterogeneous domains.

A typical approach to model diverse mapping is to employ extra style exemplars to guide the generation process. For example, [63] build dense correspondences between conditional inputs and style exemplars to transfer textures for diverse generation, while building semantic correspondences essentially requires the exemplars to have similar semantics as the conditional inputs. Without requiring extra exemplars, Variational Autoencoders (VAEs) [5] aim to regularize the latent distribution of encoded features, thus diverse generation can be achieved by directly sampling from the latent distribution. However, VAEs inevitably suffer from *posterior collapse* phenomenon [24] which leads to degraded diverse generation performance. Instead of regularizing the latent feature distributions of image feature sequences. [6] further introduce transformers in VQ-VAE to achieve high-resolution image synthesis. Nevertheless, above auto-regressive generation methods discretize relevant features independently, neglecting the potential association among multi-domain features in latent spaces.

This paper presents an Integrated Quantization Variational Auto-Encoder (IQ-VAE) that inherits the merits of CNNs (locality and spatial invariance) for high-fidelity image generation and the powerful sequence modeling of autoregressive transformer for diverse image generation. Instead of quantizing multidomain features independently as in [6], we introduce an integrated quantization scheme to quantize the involved features collaboratively in the latent spaces. The integrated quantization scheme provides a sound way to regularize the latent structure of multi-domain distributions, which can facilitate the ensuing auto-regressive modeling of sequence distributions. However, as the conditional inputs and real images often have heterogeneous features with incomparable latent spaces, KL-divergence or Wasserstein distance cannot directly measure their feature discrepancy for regularization. Inspired by the differential circuit which takes the variation between two signals as valid input, we introduce a variational regularizer which penalizes the intra-domain variation between distributions to regularize their structural discrepancy.

In addition, most auto-regressive models are trained with a so-called "teacher forcing" framework where the ground truth of target sequence (i.e., gold sequence) is provided at the training stage. However, such framework is susceptible to exposure bias, i.e., the misalignment between the training stage and the inference stage where the gold target sequence is not available and decisions are conditioned on previous model prediction. We design a Gumbel sampling strategy that greatly mitigates the exposure bias by incorporating the uncertainty of sequence distributions in training stage. Specifically, we adopt a reparameterization trick with Gumbel softmax to samples tokens from the predicted distributions and then mixes them with the gold sequence according to a reliabilitybased scheduling to make the final prediction. The Gumbel sampling also serves as data augmentation strategy that helps to avoid overfitting and improve the auto-regression performance substantially. The contributions of this work can be summarized in three aspects. First, we introduce a versatile auto-regression framework with an integrated quantization scheme for conditional image generation. Second, we propose a variational regularizer that exploits intra-domain variations to regularize heterogeneous features in latent spaces. Third, we design a Gumbel sampling strategy with a reliability-based scheduling to mitigate the misalignment between the training and inference stages of auto-regressive models.

2 Related Work

2.1 Conditional Image Generation

Conditional image generation has achieved remarkable progress by learning the mapping among data of different domains. To achieve high-fidelity yet flexible image generation, various conditional inputs have been adopted including semantic segmentation [12,48,35,59,62], scene layouts [42,65,18], key points [26,29,61,57], edge maps [12,55,56], etc. Recently, several studies explored to generate images with cross-modal guidance [58,53]. For example, Qiao *et al.* [37] propose a novel global-local attentive and semantic-preserving text-to-image-totext framework based on the idea of redescription. Ramesh *et al.* [39] handle text-to-image generation by using a transformer that auto-regressively models the text and image tokens. Chen *et al.* [2] investigated audio-to-visual generation with a conditional GANs. Nevertheless, the aforementioned methods all focus on deterministic image generation with a single generated image.

As an ill-posed problem, conditional image generation is a naturally a one-tomany mapping task as one conditional input could map to multiple diverse and faithful images. Earlier studies [15] manipulate latent feature codes to control the generation outcome, but they struggle to capture complex textures. With the emergence of GANs [7,67,34,60,54], style code injection has been designed to address this issue. For example, Zhu *et al.* [69] design semantic region-adaptive normalization (SEAN) to control the style of each semantic region individually. Choi *et al.* [4] employ a style encoder for style consistency between exemplars and the translated images. Huang *et al.* [11] and Ma *et al.* [25] transfer style codes from exemplars to source images via adaptive instance normalization (AdaIN) [10]. Recently, Zhang *et al.* [63] learn dense semantic correspondences between conditional inputs and exemplars, but require the exemplars to have similar semantics with the conditional input.

The aforementioned methods all suffer from low performance in diverse generation or require extra guidance for decent diverse generation. In this work, we propose a versatile auto-regressive framework that introduces a joint quantization scheme to achieve conditional image generation, and it inherently allows to generate diverse yet high-fidelity images as well.

2.2 Auto-regression in Image Generation

Different from VAE or GANs in image generation, auto-regressive models treat image pixels as a sequence and generate pixels one by one conditioning on the

previously generated pixels by modeling their conditional distributions. With the recent advance of deep learning, a number of studies explored to use deep auto-regressive models to generate image pixels sequentially. For instance, PixelRNN and PixelCNN [44] utilize LSTM [9] layers and masked convolutions to capture pixel inter-dependencies in a fixed order. Gated PixelCNN [32] describes a gated convolution to improve the generation quality with lower computational cost. However, deep auto-regressive models still struggle to generate high-fidelity images due to the limitation of sequential prediction of pixels. To address this issue, VQ-VAE [33] adapts an encoder-decoder structure to learns discrete latent representations for autoregressive modeling, which enables high fidelity image synthesis.

Leveraging their powerful attention mechanisms, transformers [45] allow to establish long-range dependencies effectively and have been adopted in various computer vision tasks. In image generation, Chen *et al.* [3] introduce a sequence Transformer to generate low-resolution images auto-regressively. Based on VQ-VAE [33], Esser *et al.* [6] propose a VQ-GAN to learn a discrete codebook and utilize the transformers to efficiently model sequence distributions for high-resolution images synthesis. Nevertheless, the aforementioned methods all neglect exposure bias which often introduces clear misalignment between the training and the inference. The proposed Gumbel sampling strategy introduces uncertainty in training stage which mitigates the misalignment greatly.

3 Proposed Method

3.1 Overall Framework

The framework of the proposed IQ-VAE is illustrated in Fig. 1. The IQ-VAE is first trained to learn discrete feature representations of the real image and conditional input with learnable codebook as shown in Fig. 2 (a). With the learnt IQ-VAE and codebook, the conditional input and real image can be quantized into discrete sequences by IQ-VAE encoders E_x and E_c . The transformer then auto-regressively models the distribution of the image sequences with a given sequence of conditional input. With the sequence distributions predicted by the transformer, diverse sequences can be sampled and inversely quantized into feature vectors based on the learnt codebook. Finally, the inversely quantized feature vectors are concatenated with the conditional features and fed into the IQ-VAE decoder D_x to achieve diverse image generation. Details of IQ-VAE and auto-regressive transformer will be discussed in the ensuing subsections.

3.2 Integrated Quantization

For the task of conditional image generation, [6] employ two VQ-VAEs [33] to quantize the features of conditional inputs and real images independently. However, this naive quantization approach neglects the potential coupling between conditional inputs and real images in the latent spaces. Intuitively, as conditional



Fig. 1. The framework of the proposed auto-regressive image generation with integrated quantization: We design an integrated quantization VAE (IQ-VAE) with E_x and E_c to encode the *Image* and *Condition* into discrete representation sequences \mathcal{X} and \mathcal{C} concurrently. The distribution $p(x_t|\mathcal{C}, x_{< t})$ of sequence \mathcal{X} conditioned on \mathcal{C} is modeled by an auto-regressive *Transformer*. Finally, diverse sequences are sampled from the predicted distribution $p(x_t|\mathcal{C}, x_{< t})$ which are further inversely quantized and concatenated with the encoded condition features for diverse generation via the IQ-VAE decoder D_x .

inputs imply certain information (e.g., edges) of the corresponding images, certain coupling or correlation should exist between their latent feature spaces. Explicitly regularizing such coupling between images and conditional inputs will be beneficial for the modeling of image distribution from the given conditional inputs.

We propose an integrated quantization scheme to regularize the discretization of the image and conditional input as illustrated in Fig. 2 (a). Specially, two VQ-VAEs are employed to encode the image and conditional input to a pair of feature distributions as denoted by $\mathcal{X} = [x_1, x_2, \dots, x_n]$ and $\mathcal{C} = [c_1, c_2, \dots, c_n]$. An intuitive method to regularize the feature distributions is to employ KL divergence to measure and minimize their inter-domain discrepancy, namely KL(\mathcal{C}, \mathcal{X}). However, this approach fails when a meaningful cost across the distributions cannot be defined. This is especially true for heterogeneous conditional inputs (e.g., texts and audios) that have incomparable latent spaces with respect to the image. Under such context, the KL divergence is ill-suited and inapplicable to capture the discrepancy between distributions. We thus design a novel variational regularizer that leverages the intra-domain variations of distributions to adaptively regularize their latent distributions.

Variational Regularizer Inspired by the differential circuit which takes the variation of two signals as the valid input, we propose a variational regularizer that penalizes the inter-domain discrepancy via the intra-domain variations as illustrated in Fig. 2 (b). Although the discrepancy between incomparable domain features $\mathcal{C} = [c_1, c_2, \cdots, c_n]$ and $\mathcal{X} = [x_1, x_2, \cdots, x_n]$ cannot be duly measured, the distance (or variation) among samples in the same domain can be effectively



Fig. 2. (a) illustrates the framework of the proposed integrated quantization scheme. We introduce a variational regularizer to regularize their feature distributions in latent spaces. As shown in (b), the variational regularizer employs the intra-domain variations to penalize the structural inter-domain discrepancy, and it is optimized through a sliced projection.

measured with some simple metric \mathcal{M} (Euclidean distances is adopted in this work). We thus first compute the distances among intra-domain samples for the conditioned input and real image as denoted by $\mathcal{M}(c_i, c_k)$ and $\mathcal{M}(x_j, x_l)$. The discrepancy between intra-domain variations $\{\mathcal{M}(c_i, c_k)\}, i, k \in [1, n]$ and $\{\mathcal{M}(x_j, x_l)\}, j, l \in [1, n]$ can then serve as a proxy to indicate the inter-domain discrepancy between the conditional input and real image.

To regularize the structural difference between two latent distributions effectively, we adopt the discrete optimal transport (OT) [36,41] with a 2^{th} Euclidean distance cost as the discrepancy metric which naturally induces the intrinsic geometries of distributions and can measure the discrepancy between intra-domain variations as follows:

$$OT(\mathcal{C}, \mathcal{X}) = \min_{\Gamma \in \prod(\alpha, \beta)} \sum_{i, j, k, l} \left| \mathcal{M}(c_i, c_k) - \mathcal{M}(x_j, x_l) \right|^2 \Gamma_{ij} \Gamma_{kl}$$
(1)

where Γ_{ij} and Γ_{kl} are entries of coupling matrice Γ , $\prod(\alpha, \beta) = \{\Gamma \in \mathbb{R}^{n \times n} | \Gamma \vec{1}_n = \alpha, \Gamma^T \vec{1}_n = \beta\}$, $\vec{1}_n$ is a n-dimensional all-one vector, $\alpha = \{\alpha_i\}$ and $\beta = \{\beta_j\}, i, j \in [1, n]$ are vectors of probability weights associated with c_i and x_j ($\alpha_i = 1/n, \beta_j = 1/n$). The formulation in Eq. (1) is often referred as Gromov Wasserstein (GW) distance [28] between distributions \mathcal{C} and \mathcal{X} .

With GW distance as the metric in variational regularizer, we impose a constraint on the posterior distributions defined in different latent spaces which encourages structural similarity between them [51]. This regularizer helps avoid over-regularization as it does not enforce a shared latent distribution across different or heterogeneous domains. In addition, the GW distance is invariant to translations, permutations or rotations on both distributions when Euclidean distances are used, which allows to capture discrepancy between complex latent distributions effectively.

Optimization. The solution of the variational regularizer in Eq. (1) is a nonconvex optimization problem. Grounded in the well-studied theory of Wasser-



Fig. 3. (a) illustrates the framework of the proposed Gumbel sampling with twice executions. In the first forward pass, token distribution $P_{\theta}(X)$ is predicted from the gold sequence with network parameters θ . A sample x is sampled from $P_{\theta}(X)$ according to a reliability-based scheduling and is mixed with the gold sequence for the second pass (namely final pass). (b) compares the gradient flows of direct sampling and Gumbel sampling. The presence of stochastic node x in direct sampling precludes the backpropagation of gradient from x to $P_{\theta}(X)$. Gumbel sampling allows gradient flow from x to $P_{\theta}(X)$ through a reparameterization trick which transfers the stochasticity to a Gumbel distribution.

stein disance [46], Eq. (1) can be solved through sliced Gromov Wasserstein (sliced GW) distance [46]. Specifically, the original metric measure spaces are projected to 1D spaces with random directions, and the sliced GW corresponds to the expectation of the GW distances in these projected 1D spaces. In this case, the sliced GW is approximated based on sample observations from the distributions shown in Fig. 2 (b).

In particular, given $[c_1, c_2, \cdots, c_n]$ from \mathcal{C} and $[x_1, x_2, \cdots, x_n]$ from \mathcal{X} and L projection vectors $\{\gamma_m\}_{m=1}^L$, the empirical sliced GW can be formulated by:

$$\frac{1}{L} \sum_{m=1}^{L} \min_{\Gamma_{ij}, \Gamma_{kl} \in \prod(p,q)} \sum_{i,j,k,l} \left| \mathcal{M}(\langle c_i, \gamma_m \rangle, \langle c_j, \gamma_m \rangle) - \mathcal{M}(\langle x_k, \gamma_m \rangle, \langle x_l, \gamma_m \rangle) \right|^2 \Gamma_{ij} \Gamma_{kl}.$$
(2)

where $\langle c_i, \gamma_m \rangle$ denotes the projection of c_i on direction γ_m . Compared with direct computation via proximal gradient optimization [52], the sliced GW has much lower computational complexity of $\mathcal{O}(nd)$, where n and d denote the sample number and sample dimension, respectively.

Besides the loss of variational regularizer (namely sliced GW) as denoted by \mathcal{L}_{reg} for the optimization of IQ-VAE, we also include reconstruction loss \mathcal{L}_{recon} and quantization loss \mathcal{L}_{quan} of the conditional input and real image. To further improve the image quality, a perceptual loss \mathcal{L}_{perc} and discriminator loss \mathcal{L}_{dis} are also included. Thus, the overall objective for the IQ-VAE network is:

$$\mathcal{L}_{IQ-VAE} = \lambda_1 \mathcal{L}_{reg} + \lambda_2 \mathcal{L}_{recon} + \lambda_3 \mathcal{L}_{quan} + \lambda_4 \mathcal{L}_{perc} + \lambda_5 \mathcal{L}_{dis}.$$
 (3)

where λ balances the loss terms.

3.3 Auto-Regression

Auto-regressive (AR) modeling is representative objective to accommodate sequence dependencies in a raster scan order. The probability of each position in the sequence is conditioned on all previously prediction and the joint distribution of sequences is modeled as the product of conditional distributions: $p(x) = \prod_{t=1}^{n} p(x_t | x_1, x_2, \dots, x_{t-1}) = \prod_{t=1}^{n} p(x_t | x_{<t})$. Under the context of conditional image generation, a conditional auto-regression is actually adopted for the modeling of image distribution. For clarity, we still denote the discrete image sequence as $\mathcal{X} = [x_1, x_2, \dots, x_n]$, the conditional sequence as $\mathcal{C} = [c_1, c_2, \dots, c_n]$. Then the joint distribution of image sequence conditioned on \mathcal{C} can be formulated as:

$$p(x|\mathcal{C}) = \prod_{t=1}^{n} p(x_t|c_1, c_2, \cdots, c_n, x_1, x_2, \cdots, x_{t-1}) = \prod_{t=1}^{n} p(x_t|\mathcal{C}, x_{< t}).$$
(4)

Auto-regressive models factorize the predicted tokens with chain rule of probability, which establishes the output dependency effectively for yielding better predictions. During inference, each token is predicted auto-regressively in a raster-scan order. A top-k (k is 100 in this work) sampling strategy is adopted to randomly sample from the k most likely next tokens, which naturally enables diverse sampling results. The predicted tokens are then concatenated with the previous sequence as conditions for the prediction of next token. This process repeats iteratively until all the tokens are sampled.

Gumbel sampling. Auto-regressive models are trained using the ground truth sequence (i.e., gold sequence). This framework leads to quick convergence during training, but it is misaligned with the inference stage where gold sequence is not available and decisions are purely conditioned on previous predictions. This phenomenon is typically referred as exposure bias [40]. Intuitively, this problem can be tackled by using the previous predictions as conditions with certain probability in training stage as mentioned in [30].

Specially, in order to conduct sampling from previous predictions, the autoregression process is executed twice in training stage as illustrated in Fig. 3. In the first execution, the predictions are conditioned on the gold sequence and yield discrete distribution $p_{\theta}(X) = [p_1, \dots, p_l]$ for each token (θ is network parameter, l is the number of codebook embedding). In the second execution, we aim to sample tokens according to the discrete distributions. However, direct sampling from a distribution will preclude the gradient backpropagation as shown in Fig. 3 (b). A Gumbel sampling strategy is thus introduced with a reparameterization trick [13] to enable gradient backpropagation in discrete distribution sampling. Specially, the sampling operation is conducted on a Gumbelsoftmax distribution [13] which is defined by: $\operatorname{softmax}(1/\tau(p_{\theta}(X) + g))$, where $g \sim \operatorname{Gumbel}(0, 1) = -\log(-\log U), U \sim \operatorname{Uniform}(0, 1)$. A sample x_i drawn from the Gumbel-softmax distribution can be denoted by:

$$x_{i} = \frac{\exp((\log(p_{i})) + g_{i})/\tau)}{\sum_{j=1}^{n} \exp((\log(p_{j}) + g_{j})/\tau)} \quad \text{for } i = 1, 2, \cdots, n.$$

where τ is an annealing parameter. The sampling from a Gumbel-softmax distribution exactly approximates the sampling from the categorical distribution $p_{\theta}(X)$ as proved in [27]. In forward pass of network training, sampling is actually conducted on the Gumbel(0,1) distribution which is independent of the network parameter θ . In backpropagation, the sampling operation is not involved in the gradient flow, which means that the stochasticity of sampling operation is transferred from $p_{\theta}(X)$ to the Gumbel(0,1) distribution.

To schedule the sampling in accordance with the training process, we design a *Gumbel sampling* strategy based on the prediction reliability. Considering sampled tokens are more difficult to learn than the ground truth especially at the early training stage, we only sample tokens for positions with high prediction reliability as denoted by R [20]. For a ground truth embedding γ_i and predicted distributions $[p_1, \dots, p_l]$ associated with normalized codebook embeddings $[\gamma_1, \dots, \gamma_l]$, the prediction reliability R_i can be quantified by the weighted summation of the inner products of embeddings:

$$R_i = \sum_{j=1}^{l} p_j * \gamma_j \cdot \gamma_i , \quad i \in [1, n] .$$
(5)

 $R_i \in [0, 1]$ accurately indicates the similarity between the predicted token distribution and the ground truth token, and measures whether the prediction reliability reaches the threshold (0.9 by default) to conduct token sampling.

After obtaining a sequence representing the model prediction for each position, we mix the gold tokens and predicted tokens with a given probability which is a function of the training step and is calculated with a selected schedule. We then pass the new mixed sequence to the transformer for the second execution to yield the final predictions. Note that only the gradient of the second execution is backpropagated in model training.

Computational cost. Twice execution for Gumbel sampling will increase the training time, which can be mitigated by reducing the frequency of applying Gumbel sampling. In our implementation, the Gumbel sampling is applied for every 4 iterations by default. The average speed of our model with Gumbel sampling is 2.8 iteration/s, and the model speed without Gumbel sampling is 3.0 iteration/s. Therefore, the increase of computational cost is very limited.

4 Experiments

4.1 Experimental Settings

Datasets. We benchmark our method over multiple public datasets in conditional image generation.

• ADE20k [66] has 20k training images associated with a 150-class segmentation mask. We use its semantic segmentation as conditional inputs in experiments.

• CelebA-HQ [22] has 30,000 high quality face images whose semantic maps and edges serve as the condition for image generation.

Table 1. Comparing IQ-VAE with state-of-the-art image generation methods over four conditional image generation tasks. The adopted evaluation metrics include FID, SWD and LPIPS.

Methods	ADE20K			CelebA-HQ(Edge)			DeepFashion		
	FID \downarrow	SWD \downarrow	LPIPS \uparrow	$FID \downarrow$	$\mathrm{SWD}\downarrow$	LPIPS \uparrow	FID \downarrow	$\mathrm{SWD}\downarrow$	LPIPS \uparrow
Pix2pixHD [48]	61.08	28.47	N/A	42.70	33.30	N/A	25.20	16.40	N/A
Pix2pixSC [47]	56.23	24.52	0.378	49.39	33.20	0.193	28.49	21.13	0.172
BicycleGAN [68]	62.52	33.27	0.405	44.63	31.96	0.224	29.82	22.74	0.251
StarGAN v2 $[4]$	98.72	65.47	0.451	48.63	41.96	0.214	43.29	30.87	0.296
DRIT++ [16]	105.1	81.82	0.432	50.31	47.21	0.313	52.67	42.34	0.281
SPADE [35]	33.90	19.70	0.344	31.50	26.90	0.207	36.20	27.80	0.231
SMIS [70]	42.17	22.67	0.416	23.71	22.23	0.201	26.23	23.73	0.240
VQ-GAN $[6]$	35.50	21.50	0.421	16.23	23.33	0.330	16.49	21.20	0.314
IQ-VAE	29.77	17.44	0.447	14.71	19.74	0.344	11.15	19.01	0.320



Fig. 4. Qualitative illustration of IQ-VAE and state-of-the-art image generation methods over four types of generation tasks. IQ-VAE is able to generate faithful images with high fidelity.

• DeepFashion [21] has 52,712 person images of different appearances and poses. We use the key points of the person images as conditional inputs in experiments.

• COCO-Stuff [1] augments COCO [19] with pixel-level stuff annotations. We use its layout as condition for image generation.

• CUB-200 [49] has 200 bird species with attribute labels and we use it for text-to-image generation.

• Sub-URMP [2] is a subset of URMP [17] and we use it for audio-to-image generation.

Evaluation Metrics. We evaluate the proposed IQ-VAE on the tasks of semanticto-image, edge-to-image and keypoint-to-image generation, as these tasks have rich prior studies for comprehensive yet fair benchmarking. We assess the compared methods with several widely adopted evaluation metrics. Specifically, *Fréchet Inception Score (FID)* [8] and *Sliced Wasserstein distance (SWD)* [14] are employed to evaluate the quality of generated images. *Learned Perceptual Image Patch Similarity (LPIPS)* [64] measures the distance between image patches,

Transformer		IQ-VAE		
Parameters	Setting	Parameters	Setting	
learning rate	1.5e-4	learning rate	1.5e-4	
batch size	32	batch size	32	
epoch	50	epoch	100	
vocabulary size	1024	codebook embedding number	1024	
embedding number	1024	codebook embedding dimension	256	
sequence length	512	feature number	256	
number of transformer block	24			

Table 2. The parameter setting in the proposed transformer and IQ-VAE.



Fig. 5. Illustration of diverse image generation by the proposed IQ-VAE: Faithful yet diverse images are successfully generated with different types of conditional inputs such as semantic maps, edge maps, key points, layout maps, as well as heterogeneous conditions like texts and audios.

which is employed to evaluate the diversity of generated images and reconstruction performance of auto-encoder.

Implementation Details. The proposed model is optimized with a learning rate of 1.5*e*-4. The auto-regressive transformer is implemented based on the GPT2 architecture [38] with a input size of 256. AdamW [23] solver is adopted with $\beta_1 = 0.9$ and $\beta_2 = 0.95$. All experiments are conducted on 4 Tesla V100 GPUs with a batch size of 32. The size of generated images is 256×256 for all evaluated generation tasks. The transformer is implemented based on minGPT⁴. Table. 2 shows parameter setting in the transformer and IQ-VAE.

4.2 Quantitative Results

We compare the proposed IQ-VAE with several state-of-the-art conditional image generation methods including 1) Pix2pixHD [48]; 2) Pix2pixSC [47]; 3) BicycleGAN [68]; 4) StarGAN v2 [4]; 5) DRIT++ [16]; 6) SPADE [35]; 7) SMIS [70]; 8) Taming Transformer [6].

⁴ https://github.com/karpathy/minGPT

Table 3. Ablation study of IQ-VAE on ADE20k. VR and None denote the proposed variational regularizer and no regularization, respectively. GS denotes the proposed Gumbel sampling.

Models	$\mathbf{FID}\downarrow$	$\mathbf{SWD}\downarrow$	$\mathbf{LPIPS}\uparrow$
VQ-GAN	35.50	21.50	0.421
IQ-VAE(None)	31.88	19.14	0.441
IQ-VAE(VR)	31.41	18.71	0.450
IQ-VAE(VR) + GS	29.77	17.44	0.447

In the quantitative experiments, all compared methods generate diverse images except Pix2PixHD [48] which does not support diverse generation. Table 1 shows experimental results in FID, SWD and LPIPS. It can be observed that IQ-VAE outperforms all compared methods across most metrics and tasks consistently. DRIT++ [16] and StarGAN v2 [4] achieve relatively high LPIPS scores by sacrificing the image quality as measured by FID and SWD, while SPADE [35] and SMIS [70] achieve decent FID and SWD scores with degraded LPIPS scores. The proposed IQ-VAE employs powerful variational auto-encoders to achieve high-fidelity image synthesis and a auto-regressive model for faithful image diversity modeling, thus achieving superior performance in terms of image quality and diversity. Compared with Taming transformer [6], the proposed IQ-VAE allows to quantize the image sequences and conditional sequence jointly and boosts the auto-regressive modeling for better FID and SWD scores. In addition, the proposed Gumbel sampling introduces uncertainty of distribution sampling into the training process which mitigates the exposure bias and improves the inference performance clearly. As the mixed sequence serves as certain extra data augmentation, the Gumbel sampling also helps to alleviate the over-fitting of auto-regressive model effectively.

4.3 Qualitative Evaluation

We perform qualitative comparisons as shown in Fig. 4. The experiments are conducted over six datasets including ADE20k [66], CelebA-HQ [22], DeepFashion [21], COCO-Stuff [1], CUB-200 [49], and Sub-URMP [2]. The splits of training and testing sets on all above datasets follow the default split settings. In addition, the data used in the experiments do not contains person identity related information or offensive contents. It can be seen that IQ-VAE achieves the best visual quality and presents remarkable coherence with the condition. SPADE [35] and SMIS [70] adopt VAE to constraint the distribution of encoded features which cannot capture the complex distributions of real images. Star-GAN v2 [4] and DRIT++ [16] adopt single latent code to encode image styles, which tends to capture global styles but misses local details.

IQ-VAE also generalizes well and demonstrates superior synthesis quality and diversity in various generation tasks as illustrated in Fig. 5. It can be observed that IQ-VAE is capable of synthesizing high-fidelity images with various conditional inputs such as semantic maps, edge maps, keypoints, layout maps as well as heterogeneous conditions such as texts and audios.

4.4 Ablation Study

We conduct extensive ablation studies to evaluate IQ-VAE as shown in Table 3. The baseline is selected as VQ-GAN (namely Taming Transformer [6]). Replacing VQ-GAN with the proposed IQ-VAE without any regularization in IQ-VAE(None) brings in marginal improvement. The proposed variational regularizer with adaptive weights in IQ-VAE(VR) improves the generation performance, demonstrating the effectiveness of adaptive weights learning. Finally, including the Gumbel sampling remarkably boosts the performance as indicated in IQ-VAE(VR)+GS.

We study the effect of feature sizes for discrete representation in IQ-VAE and Fig. 6 shows experimental results on the CelebaHQ dataset. As Fig. 6 shows, we specify the size of representation features in terms of a factor F where Fx denotes a feature size of $x \times x$. Note the input size of transformer is always fixed at 16 \times 16. The horizontal axis of the graph shows reconstruction error as measured by LPIPS [64] which indicates the upper bound of generation quality (lower is better), while the vertical axis shows negative log-likelihood from the transformer which indicates the performance of the auto-regressive modeling (lower is better). We can see



Fig. 6. Trade-off between negative loglikelihood and reconstruction error with different sizes of encoded features on CelebaHQ [22].

that there is a trade-off between the negative log-likelihood and reconstruction error. Though an encoded feature of small size allows the transformer to better model the image distribution, the reconstruction deteriorates severely after a certain value (F16 in this case). The proposed integrated quantization and Gumbel sampling instead improve the negative log-likelihood remarkably without sacrificing the reconstruction performance clearly.

4.5 User Study

We conduct crowdsourcing user study to evaluate the quality of generated images as shown in Fig. 7. Specifically, 100 pairs of images generated by all compared methods are shown to 10 users who selected the image with the best visual quality. As shown in Fig. 7, we compared the proposed IQ-VAE with several state-of-the-art generation methods including BicycleGAN [68], SPADE [35],

SMIS [70], and Taming Transformer [6]. The images generated by the proposed IQ-VAE are much more realistic according to the user feedback.



Fig. 7. User study over four datasets ADE20K [66], CelebA-HQ [22](Semantic), CelebA-HQ [22](Edge), DeepFashion [21]. The bars show the number of images that AMT users ranked with the best visual quality.

5 Conclusions

This paper presents IQ-VAE, an auto-regressive framework with integrated quantization for conditional image synthesis. We propose a novel variational regularizer to regularize the feature distribution structures of conditional inputs and real images, which boosts the auto-regressive modeling clearly. To mitigate the misalignment between training and inference of auto-regressive model, a Gumbel sampling strategy with a reliability-based scheduling is included in the training stage and improves the inference performance by a large margin. Quantitative and qualitative experiments show that IQ-VAE is capable of generating diverse yet high-fidelity images with multifarious conditional inputs.

Limitations. As auto-regression is adopted in the model to predict image sequence, the inference speed is inevitably constrained which may limit the application of the proposed model in time-critical tasks. Although some works [50,31] have been proposed to speed up the autoregressive sampling, the acceleration for the inference of auto-regressive model is still an open challenge.

Potential Negative Societal Impacts This work aims to synthesize diverse yet high-fidelity images with given conditional inputs. It could have negative impacts if it is used for certain illegal purpose such as image forgery and manipulation.

Acknowledgement. This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1209–1218 (2018)
- Chen, L., Srivastava, S., Duan, Z., Xu, C.: Deep cross-modal audio-visual generation. In: Proceedings of the on Thematic Workshops of ACM Multimedia 2017. pp. 349–357 (2017)
- Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Luan, D., Sutskever, I.: Generative pretraining from pixels. In: International Conference on Machine Learning. pp. 1691–1703. PMLR (2020)
- Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: Stargan v2: Diverse image synthesis for multiple domains. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8188–8197 (2020)
- 5. Doersch, C.: Tutorial on variational autoencoders. arXiv preprint arXiv:1606.05908 (2016)
- Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2021)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems 27 (2014)
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Advances in neural information processing systems. pp. 6626–6637 (2017)
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation 9(8), 1735–1780 (1997)
- Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1501–1510 (2017)
- Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised imageto-image translation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 172–189 (2018)
- Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)
- Jang, E., Gu, S., Poole, B.: Categorical reparameterization with gumbel-softmax. arXiv preprint arXiv:1611.01144 (2016)
- 14. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196 (2017)
- Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
- Lee, H.Y., Tseng, H.Y., Mao, Q., Huang, J.B., Lu, Y.D., Singh, M., Yang, M.H.: Drit++: Diverse image-to-image translation via disentangled representations. International Journal of Computer Vision 128(10), 2402–2417 (2020)
- Li, B., Liu, X., Dinesh, K., Duan, Z., Sharma, G.: Creating a multitrack classical music performance dataset for multimodal music analysis: Challenges, insights, and applications. IEEE Transactions on Multimedia **21**(2), 522–535 (2018)
- Li, Y., Cheng, Y., Gan, Z., Yu, L., Wang, L., Liu, J.: Bachgan: High-resolution image synthesis from salient object layout. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2020)

- 16 F. Zhan et al.
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
- Liu, Y., Meng, F., Chen, Y., Xu, J., Zhou, J.: Confidence-aware scheduled sampling for neural machine translation. arXiv preprint arXiv:2107.10427 (2021)
- 21. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1096–1104 (2016)
- Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: Proceedings of the IEEE international conference on computer vision. pp. 3730– 3738 (2015)
- Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
- Lucas, J., Tucker, G., Grosse, R., Norouzi, M.: Don't blame the elbo! a linear vae perspective on posterior collapse. arXiv preprint arXiv:1911.02469 (2019)
- Ma, L., Jia, X., Georgoulis, S., Tuytelaars, T., Van Gool, L.: Exemplar guided unsupervised image-to-image translation with semantic consistency. In: International Conference on Learning Representations (2018)
- Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Van Gool, L.: Pose guided person image generation. In: Advances in neural information processing systems. pp. 406–416 (2017)
- 27. Maddison, C.J., Tarlow, D., Minka, T.: A* sampling. In: NIPS (2014)
- 28. Mémoli, F.: Gromov–wasserstein distances and the metric approach to object matching. Foundations of computational mathematics **11**(4), 417–487 (2011)
- Men, Y., Mao, Y., Jiang, Y., Ma, W.Y., Lian, Z.: Controllable person image synthesis with attribute-decomposed gan. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5084–5093 (2020)
- Mihaylova, T., Martins, A.F.: Scheduled sampling for transformers. arXiv preprint arXiv:1906.07651 (2019)
- Oord, A., Li, Y., Babuschkin, I., Simonyan, K., Vinyals, O., Kavukcuoglu, K., Driessche, G., Lockhart, E., Cobo, L., Stimberg, F., et al.: Parallel wavenet: Fast high-fidelity speech synthesis. In: International conference on machine learning. pp. 3918–3926. PMLR (2018)
- Oord, A.v.d., Kalchbrenner, N., Vinyals, O., Espeholt, L., Graves, A., Kavukcuoglu, K.: Conditional image generation with pixelcnn decoders. arXiv preprint arXiv:1606.05328 (2016)
- Oord, A.v.d., Vinyals, O., Kavukcuoglu, K.: Neural discrete representation learning. arXiv preprint arXiv:1711.00937 (2017)
- Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: European Conference on Computer Vision. pp. 319–345. Springer (2020)
- Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2337–2346 (2019)
- Peyré, G., Cuturi, M., et al.: Computational optimal transport: With applications to data science. Foundations and Trends® in Machine Learning 11(5-6), 355–607 (2019)
- Qiao, T., Zhang, J., Xu, D., Tao, D.: Mirrorgan: Learning text-to-image generation by redescription. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1505–1514 (2019)

- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI blog 1(8), 9 (2019)
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. arXiv preprint arXiv:2102.12092 (2021)
- 40. Schmidt, F.: Generalization in generation: A closer look at exposure bias. arXiv preprint arXiv:1910.00292 (2019)
- 41. Solomon, J.: Optimal transport on discrete domains. AMS Short Course on Discrete Differential Geometry (2018)
- Sun, W., Wu, T.: Image synthesis from reconfigurable layout and style. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 10531–10540 (2019)
- Tang, H., Xu, D., Liu, G., Wang, W., Sebe, N., Yan, Y.: Cycle in cycle generative adversarial networks for keypoint-guided image generation. In: Proceedings of the 27th ACM International Conference on Multimedia. pp. 2052–2060 (2019)
- Van Oord, A., Kalchbrenner, N., Kavukcuoglu, K.: Pixel recurrent neural networks. In: International Conference on Machine Learning. pp. 1747–1756. PMLR (2016)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. arXiv preprint arXiv:1706.03762 (2017)
- Vayer, T., Flamary, R., Tavenard, R., Chapel, L., Courty, N.: Sliced gromovwasserstein. arXiv preprint arXiv:1905.10124 (2019)
- 47. Wang, M., Yang, G.Y., Li, R., Liang, R.Z., Zhang, S.H., Hall, P.M., Hu, S.M.: Example-guided style-consistent image synthesis from semantic labeling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1495–1504 (2019)
- Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: Highresolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8798–8807 (2018)
- Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., Perona, P.: Caltech-ucsd birds 200. California Institute of Technology (2010)
- Wiggers, A., Hoogeboom, E.: Predictive sampling with forecasting autoregressive models. In: International Conference on Machine Learning. pp. 10260–10269. PMLR (2020)
- Xu, H., Luo, D., Henao, R., Shah, S., Carin, L.: Learning autoencoders with relational regularization. In: International Conference on Machine Learning. pp. 10576– 10586. PMLR (2020)
- 52. Xu, H., Luo, D., Zha, H., Duke, L.C.: Gromov-wasserstein learning for graph matching and node embedding. In: International conference on machine learning. pp. 6932–6941. PMLR (2019)
- 53. Yu, Y., Zhan, F., Wu, R., Zhang, J., Lu, S., Cui, M., Xie, X., Hua, X.S., Miao, C.: Towards counterfactual image manipulation via clip. arXiv preprint arXiv:2207.02812 (2022)
- Zhan, F., Xue, C., Lu, S.: Ga-dan: Geometry-aware domain adaptation network for scene text detection and recognition. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9105–9115 (2019)
- 55. Zhan, F., Yu, Y., Cui, K., Zhang, G., Lu, S., Pan, J., Zhang, C., Ma, F., Xie, X., Miao, C.: Unbalanced feature transport for exemplar-based image translation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2021)

- 18 F. Zhan et al.
- Zhan, F., Yu, Y., Wu, R., Cui, K., Xiao, A., Lu, S., Shao, L.: Bi-level feature alignment for versatile image translation and manipulation. arXiv preprint arXiv:2107.03021 (2021)
- Zhan, F., Yu, Y., Wu, R., Zhang, C., Lu, S., Shao, L., Ma, F., Xie, X.: Gmlight: Lighting estimation via geometric distribution approximation. arXiv preprint arXiv:2102.10244 (2021)
- Zhan, F., Yu, Y., Wu, R., Zhang, J., Lu, S.: Multimodal image synthesis and editing: A survey. arXiv preprint arXiv:2112.13592 (2021)
- Zhan, F., Yu, Y., Wu, R., Zhang, J., Lu, S., Zhang, C.: Marginal contrastive correspondence for guided image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10663–10672 (2022)
- Zhan, F., Zhang, C.: Spatial-aware gan for unsupervised person re-identification. In: 2020 25th International Conference on Pattern Recognition (ICPR). pp. 6889– 6896. IEEE (2021)
- Zhan, F., Zhang, C., Yu, Y., Chang, Y., Lu, S., Ma, F., Xie, X.: Emlight: Lighting estimation via spherical distribution approximation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 3287–3295 (2021)
- Zhan, F., Zhang, J., Yu, Y., Wu, R., Lu, S.: Modulated contrast for versatile image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18280–18290 (2022)
- 63. Zhang, P., Zhang, B., Chen, D., Yuan, L., Wen, F.: Cross-domain correspondence learning for exemplar-based image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5143–5153 (2020)
- 64. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
- Zhao, B., Meng, L., Yin, W., Sigal, L.: Image generation from layout. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8584–8593 (2019)
- Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 633–641 (2017)
- 67. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)
- Zhu, J.Y., Zhang, R., Pathak, D., Darrell, T., Efros, A.A., Wang, O., Shechtman, E.: Toward multimodal image-to-image translation. Advances in Neural Information Processing Systems **2017**, 466–477 (2017)
- Zhu, P., Abdal, R., Qin, Y., Wonka, P.: Sean: Image synthesis with semantic regionadaptive normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5104–5113 (2020)
- Zhu, Z., Xu, Z., You, A., Bai, X.: Semantically multi-modal image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5467–5476 (2020)