

# Controllable Image Generation with Composed Parallel Token Prediction

Jamie Stirling  
Durham University  
United Kingdom

jamie.s.stirling@durham.ac.uk

Chris G. Willcocks  
Durham University  
United Kingdom

christopher.g.willcocks@durham.ac.uk

Noura Al-Moubayed  
Durham University  
United Kingdom

noura.al-mouyed@durham.ac.uk

Hubert P. H. Shum  
Durham University  
United Kingdom

hubert.shum@durham.ac.uk

## Abstract

Conditional discrete generative models struggle to faithfully compose multiple input conditions. To address this, we derive a theoretically-grounded formulation for composing discrete probabilistic generative processes, with masked generation (absorbing diffusion) as a special case. Our formulation enables precise specification of novel combinations and numbers of input conditions that lie outside the training data, with concept weighting enabling emphasis or negation of individual conditions. In synergy with the richly compositional learned vocabulary of VQ-VAE and VQ-GAN, our method attains a 63.4% relative reduction in error rate compared to the previous state-of-the-art, averaged across 3 datasets (positional CLEVR, relational CLEVR and FFHQ), simultaneously obtaining an average absolute FID improvement of  $-9.58$ . Meanwhile, our method offers a  $2.3\times$  to  $12\times$  real-time speed-up over comparable methods, and is readily applied to an open pre-trained discrete text-to-image model for fine-grained control of text-to-image generation.

## 1. Introduction

Conditional image generation models struggle to faithfully satisfy multiple input conditions simultaneously, especially for novel combinations outside the training data [1] (compositional generalisation [2]).

In the area of continuous iterative image generation approaches, earlier works [4–6] have proposed methods for improving controllability of image generation via composition of energy-based models and diffusion models. These models improve condition-output alignment, significantly exceeding non-composed baselines in terms of accuracy and image quality [6]. However, such approaches do not extend to discrete sample spaces, which offer a number of trade-offs and





Figure 1. Scatter plots of compositional generation error vs FID on 3 datasets (3 input components). Our method lies on the Pareto front of all results (see Appendix for full scatter plots) while achieving the lowest or joint lowest error among the baselines.



"Vibrant, bustling outdoor market with colourful stalls" AND "Vendors and shoppers from various cultures interacting" AND "Hanging lanterns illuminating the scene" AND "Fluffy kitten with big eyes and pink nose" AND "Soft lighting casting gentle shadows" AND "Snow-capped mountains in the distance, reflected in water" AND "Rustic, wooden rowboat floating on misty lake" AND "Morning light filtering through mist, golden glow"

Figure 2. Compositional text-to-image results with captions (zooming recommended). *Top*: single-prompt baseline. *Bottom*: composed multi-prompt (ours). Our method enables the composition of multiple conditions, conferring an advantage over the single-prompt baseline.



Figure 3. Concept negation with text-to-image (left baseline, right ours): Our method allows more precise control over the outputs of an existing pre-trained model (aMUSEd [3]). The baseline handles poorly the negation/removal of characteristics which commonly co-occur with the subject of the image (e.g. a king with *no* crown). Each image is selected from three runs as the most representative of the prompt (ours and baseline).

outright improvements over their continuous counterparts, including data-efficiency and temperature control[7–10].

Discrete image generation methods, including autoregressive [8] and masked sampling [3, 9–11], offer advantages in generation speed, image quality and diversity over continuous analogues [12–14]. Until now, these advantages have remained mutually exclusive from those of compositional generation, which have thus far been limited to *continuous* approaches only [5, 6].

We address this gap by deriving a theoretically-grounded framework for composing discrete generative model outputs, enabling faithful multi-condition generation. Applying this framework to conditional parallel token prediction (absorbing diffusion), we leverage the expressive and richly compositional “visual language” emergent in discrete representation methods like VQ-VAE and VQ-GAN. Concept weighting provides an additional degree of controllability, enabling emphasis or negation of individual conditions.

We train and evaluate our approach on positional CLEVR [6, 15], relational CLEVR [6, 15] and FFHQ [16]. Our method achieves state-of-the-art error rates across all 9 quantitative settings (3 datasets with 1, 2, or 3 input components), as shown in Figure 1, equating to an average relative error rate reduction of 63.4%. We outperform the next-best method (ranked by error rate) in terms of FID for 7 out of 9 settings, meanwhile maintaining the significant speed advantages of parallel token prediction, offering a  $2.3\times$  to  $12\times$  speed-up in real time over comparable continuous approaches. We further show that our method can be applied to an open pre-trained text-to-image parallel token model (aMUSEd [3]) with appealing visual results (Figures ??).

The contributions of this research are summarised as:

- Extending the product-of-experts principle, we derive a theoretically-grounded formulation for composing conditional distributions in the discrete sequential generation setting. This result is generalised to arbitrary discrete iterative processes, encompassing both masked and autoregressive sampling.
- We adapt the formulation to conditional parallel token prediction (absorbing diffusion [9]), enabling efficient, high-

quality, and controllable image synthesis in the discrete latent space of VQ-VAE and VQ-GAN.

- We demonstrate empirically that our method achieves state-of-the-art accuracy and competitive FID scores across 3 datasets, while offering a  $2.3\times$  to  $12\times$  speed-up in real time, over comparable continuous methods.
- We release source code of our method under the MIT license.

## 2. Related Work

We examine the literature from the perspectives of *compositional methods* and *discrete generative methods*, with our work lying in the intersection of the two.

**Continuous compositional generation:** Earlier works proposed methods for conditioning continuous generative models [4–6] based on the conjunction and negation of input attributes, drawing on formal analogues to product-of-experts [17]. To the best of our knowledge, no analogous method exists for composing flow matching models [14]. [6] introduces an approach that enhances the capabilities of text-conditioned diffusion models in generating complex and photo-realistic images based on textual descriptions.

Many compositional approaches apply a weighting factor to each component, which resembles the mathematical form of **classifier-free guidance (CFG)** which is readily applied to both continuous and discrete image generation [3, 18].

Applications-wise, [19] proposes applying composed diffusion to a multi-constraint design task, which is shown to generalise to designs which are more complex than those in the training data (compositional generalization). Similarly, [20] composes multiple energy terms to constrain the generation of motion from text. These methods demonstrate the practical utility and contemporary relevance of compositional approaches.

Beyond the literature, our work diverges significantly from existing work with continuous methods, in providing the first (to the best of our knowledge) formulation for composing arbitrary *discrete* generative processes.

**Composition in sequential tasks:** Earlier work in reinforcement learning [21] has sought to use ideas relevant to



Figure 4. Conceptual product space: Example of composing two concept spaces using our framework: {"a cat", "a dog", "an apple", "a cherry"}  $\times_{AND}$  {"a red object", "an orange object", "a yellow object", "a green object", "a blue object", "an indigo object", "a violet object"}.

our own for composing *policies* for the purposes of compositional generalisation in multi-timestep environments. The main idea shared with our work is that of *multiplying* constituent “primitives” (as opposed to *additive* composition, as in mixture-of-experts [22]). This shares some superficial similarities with our work on compositional discrete image generation, where multiple attributes are required to be expressed in the same output image; however, it differs significantly in both the application and the formalism.

**Discrete representation learning and sampling:** Discrete representation learning has emerged as the discrete counterpart to continuous VAE approaches [7, 23]. Discrete representation learning is based on the concept of vector-quantization (VQ) [24], whereby features from a continuous vector space are mapped to an element of a finite set of learned codebook vectors. This VQ family of models includes VQ-VAE [7] and VQ-GAN [8]. VQ-family approaches require a secondary prior model to be trained to sample from the discrete latent space, which can be computationally expensive at both train- and inference- time [7–9]. An alternative approach is to predict multiple tokens in parallel transformer encoder [3, 9, 10] (akin to masked language modelling [25]), which introduces a controllable trade-off between sample speed and generation quality. The per-image generation time is still linear in the size ( $W \times H$ ) of the image, albeit with a smaller constant than autoregressive models [9]. Masked prediction can be further extended to unaligned conditional 2D-3D image synthesis [26].

Beyond autoregressive and masked generation approaches, there exists a handful of alternative discrete generative methods. *Discrete flow matching* [27] and bitwise autoregressive modelling [28] are notable recent examples.

### 3. Method

In this section, we present our method for composing generative models for controllable sampling from discrete representation spaces of images.

The overview of our approach is shown in Figure 5. We first derive a novel formulation that directly informs our method of composing conditional distributions over discrete spaces. Next, we show how this extends generally to all discrete sequential generation tasks, in which categorical variables are sampled iteratively to produce a complete sample (e.g. via autoregressive or masked models). We show how this result can be specifically adapted for conditional parallel token prediction [9, 10] to achieve high-quality and accurate controllable image synthesis. This is further enhanced by concept weighting, which allows the relative importance of input conditions to be increased, decreased or negated to the desired effect. We note that similar results can be obtained for other types of generative models (provided they are iterative and discrete, see Appendix for examples). This section lays the groundwork for our later experiments with compositional sampling from the latent space of VQ-VAE and VQ-GAN.

#### 3.1. Composing Conditional Categorical Distributions

Our framework is derived from the simplifying assumption employed by the product of experts [17], in which the input conditions  $c_1, \dots, c_n$  are independent of each other conditional on the output variable  $x$ , i.e.  $P(c_1, c_2|x) \propto P(c_1|x)P(c_2|x)$  for any pair of distinct attributes  $(c_1, c_2)$ . A consequence of this is that the probability of an outcome  $x$  given two or more conditions  $c_1, c_2, \dots$  is proportional to

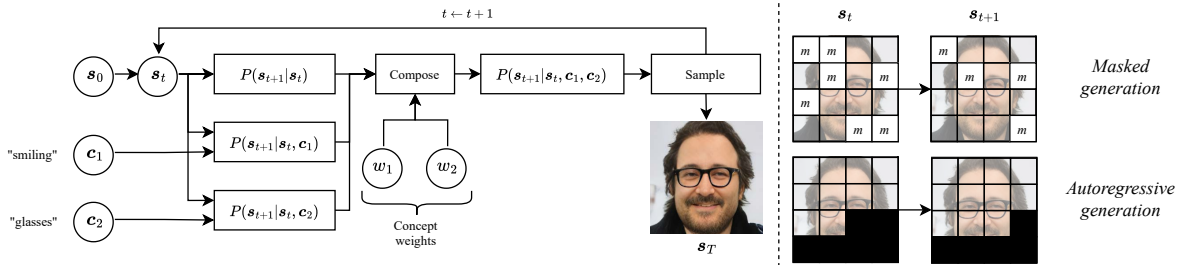


Figure 5. Overview of our approach. *Left*: We start with a generic “empty” state  $s_0$  or the preceding state  $s_t$  and a set of input concepts  $c_1, c_2, \dots$ . These are used to compute the unconditional distribution over  $s_{t+1}$  and the conditional distributions given each of  $c_1, c_2, \dots$ . These are *composed*, optionally incorporating *concept weights*, to produce an accurate estimate of the distribution conditioned on *all* inputs. The process is repeated until a “full” state  $s_T$  is obtained, e.g. a grid of VQ-VAE latents. *Right*: Our general formulation encompasses any *additive* discrete generative process, i.e. where the representation of  $s_t$  is fully determined by  $s_{t+1}$  (masked and autoregressive generation shown as examples).

the product of the probabilities of each condition given  $\mathbf{x}$ . Intuitively, taking the product of different categorical distributions in this way is analogous to taking the intersection of the sample spaces of two or more conditional distributions, thus (ideally) resulting in samples which embody all of the specified conditions. We discuss the benefits and limitations of this in Section 5. Following previous work with composition of continuous models [6, 29], we factorize the distribution of a  $k$ -way categorical variable  $\mathbf{x}$  conditioned on  $n$  variables as

$$P(\mathbf{x}|\mathbf{c}_1, \dots, \mathbf{c}_n) \propto P(\mathbf{x}) \prod_{i=1}^n P(\mathbf{c}_i|\mathbf{x}). \quad (1)$$

Applying Bayes’ theorem [30], this can be re-written as

$$\begin{aligned} P(\mathbf{x}|\mathbf{c}_1, \dots, \mathbf{c}_n) &\propto P(\mathbf{x}) \prod_{i=1}^n \frac{P(\mathbf{x}|\mathbf{c}_i)P(\mathbf{c}_i)}{P(\mathbf{x})}, \\ &\propto P(\mathbf{x}) \prod_{i=1}^n \frac{P(\mathbf{x}|\mathbf{c}_i)}{P(\mathbf{x})}. \end{aligned} \quad (2)$$

We are able to eliminate the  $P(\mathbf{c}_i)$  term in (2) as a consequence of normalising the values of  $P(\mathbf{x} = x_i|\dots)$  for all  $x_i$ , such that they sum to 1. Please refer to Appendix for the full derivation.

### 3.2. Composition for Sequential Generative Tasks

So far, we have shown how our approach applies to conditional generation with a single categorical output  $\mathbf{x}$ . In practice, many generative tasks involve sampling multiple categorical variables (tokens or latent codes) over a number of time steps [31, 32] where the sampling of a new state  $s_{t+1}$  at each successive step  $t$  is conditioned on the previous state (in addition to the specified conditions  $c_1, \dots, c_n$ ). Formulating this alongside the result in (2) gives the following general expression for composing discrete sequential

generation tasks (see Appendix for further explanation):

$$P(s_{t+1}|s_t, \mathbf{c}_1, \dots, \mathbf{c}_n) \propto P(s_{t+1}|s_t) \prod_{i=1}^n \frac{P(s_{t+1}|s_t, \mathbf{c}_i)}{P(s_{t+1}|s_t)}. \quad (3)$$

We observe that this applies generally to any generative process in which each successive step is conditioned on the output of previous steps, including autoregressive language modelling [32], masked language modelling [33], as well as autoregressive [8] and non-autoregressive [9] approaches for image generation. In the remainder of this paper, we maintain a particular focus on conditional parallel (masked) token prediction [9], which we use for composed sampling from the latent space of VQ-VAE [7] and VQ-GAN [8] for high-fidelity image synthesis.

A key practical consideration concerning the result in (3) is that estimates must be obtained for each conditional probability distribution  $P(s_{t+1}|s_t, \mathbf{c}_i)$  in addition to  $P(s_{t+1}|s_t)$ . In each of our experiments (Section 4), we ensure that, during training, conditional information is zero-masked with a set probability (0.1) per sample, thus allowing us to obtain  $P(s_{t+1}|s_t)$  at inference time by supplying zeros in place of the condition encoding.

### 3.3. Composed Parallel Token Prediction

Parallel token prediction [9–11] is a non-autoregressive alternative to next-token prediction [8] for generative sampling from a discrete latent space. This allows a direct trade-off between sampling speed and image generation quality [9] by controlling the rate at which tokens are sampled.

Parallel token prediction can be thought of as the gradual un-masking of a collection of discrete latent codes  $z_0$  given the partial reconstruction from a previous time step (as well as additional conditioning information). This corresponds directly to the next-state prediction formulation in (3), but with the time labels  $t$  reversed in order to reflect the “reverse process” which characterizes diffusion models (both

continuous [34] and discrete [9]):

$$P(\mathbf{z}_{t-1}|\mathbf{z}_t, \mathbf{c}_1, \dots, \mathbf{c}_n) \propto P(\mathbf{z}_{t-1}|\mathbf{z}_t) \prod_{i=1}^n \frac{P(\mathbf{z}_{t-1}|\mathbf{z}_t, \mathbf{c}_i)}{P(\mathbf{z}_{t-1}|\mathbf{z}_t)}. \quad (4)$$

In (4),  $\mathbf{z}_t$  is an intermediate, partially unmasked representation at each time step, and  $\mathbf{z}_{t-1}$  represents the distribution over image representations with fewer masked tokens. In practise, following earlier work with parallel token prediction, the model is trained to directly predict the fully unmasked representation  $\mathbf{z}_0$  (as opposed to intermediate states) so as to maximise training stability [9]. At inference time, we compute the composed unmasking probabilities as:

$$P(\mathbf{z}_0|\mathbf{z}_t, \mathbf{c}_1, \dots, \mathbf{c}_n) \propto P(\mathbf{z}_0|\mathbf{z}_t) \prod_{i=1}^n \frac{P(\mathbf{z}_0|\mathbf{z}_t, \mathbf{c}_i)}{P(\mathbf{z}_0|\mathbf{z}_t)}. \quad (5)$$

Image representations are then unmasked one or more tokens at a time, corresponding to a trade-off between sample speed (more tokens per iteration) and image generation quality (fewer tokens per iteration) [9].

### 3.4. Concept Weighting to Improve Controllability

We introduce an additional set of hyperparameters  $w_1, \dots, w_n$  to specify the relative importance of each condition  $\mathbf{c}_1, \dots, \mathbf{c}_n$  respectively, motivated by similar research employing diffusion models for image generation [6]. Restating (3) in terms of log-probabilities and introducing these weighting terms gives:

$$\begin{aligned} & \log P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{c}_1, \dots, \mathbf{c}_n) \\ &= \log P(\mathbf{s}_{t+1}|\mathbf{s}_t) \\ & \quad + \sum_{i=1}^n w_i [\log P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{c}_i) - \log P(\mathbf{s}_{t+1}|\mathbf{s}_t)]. \end{aligned} \quad (6)$$

This expression (6) can be readily reformulated into the corresponding form for parallel token prediction (4). A key observation here is that setting a concept weight  $w_i$  to *negative* (e.g.  $-1$ ) has the intuitive effect of *negating* the corresponding condition  $\mathbf{c}_i$  by excluding image representations which correspond to  $\mathbf{c}_i$  from the sample space. Altogether, the prompt-weighting approach provides an additional degree of controllability over model outputs by enabling conditions to be emphasized ( $w_i > 1$ ), de-emphasized ( $w_i < 1$ ) or even negated ( $w_i < 0$ ) as required. We demonstrate the practical utility of this feature in our qualitative experiments (Section 4). We do not include *disjunction* in our evaluation for reasons explained in Appendix.

In practice, we use our compositional framework to sample from the discrete latent space of VQ-VAE [7] and VQ-GAN [8], which are powerful and practical approaches for

encoding images and other high-dimensional modalities as collections of discrete latent codes (visual tokens) while producing high-fidelity reconstructions and generated samples.

### 3.5. Discrete Encoding and Decoding of Images

In order to compose categorical distributions for generating images, we must also define an invertible mapping between RGB images and discrete latent representations. We utilise convolutional down-sampling and up-sampling (autoencoder) to map between RGB image space and latent embedding space. Following the original VQ-VAE [7] formulation, we employ nearest-neighbour vector quantization, in which encoder outputs are mapped to their nearest neighbour in a learned codebook. Specifically, for each encoder output vector  $z_e$ , the corresponding quantized vector is computed as the nearest codebook entry  $e_c$ , where

$$c = \arg \min_j \|z_e - e_j\|_2 \quad (7)$$

and  $e_0, e_1, \dots, e_{K-1}$  are entries in a learned vector codebook of length  $K$ .

Since the quantization step is non-differentiable, it is necessary to estimate the gradients during backpropagation. For this purpose, we apply straight-through gradient estimation [35], whereby during backpropagation, the gradients are copied directly from the decoder input  $z_c$  to the encoder output  $z_e$ . We adopt this vector quantization approach for all experiments, including the embedding and commitment loss terms from the original VQ-VAE formulation [7].

## 4. Experiments

### 4.1. Datasets

Following earlier work [6] in evaluating compositional generalisation for image generation, we employ three datasets for training and evaluation: Positional CLEVR [6, 15], Relational CLEVR [6, 15] and FFHQ [16] (full description of datasets in Appendix). These three datasets are chosen to represent a range of unique and challenging compositional tasks (conditioned on object position, object relations, and facial attributes respectively). For each of the three datasets, we train a VQ-VAE or VQ-GAN model to enable encoding and decoding between the image space and the discrete latent representation space, in addition to a conditional parallel token prediction model (encoder-only transformer) which learns to unmask discrete latent representations, optionally conditioned on an encoded input annotation.

### 4.2. Model Training

We train a VQ-GAN model to reconstruct FFHQ samples at  $256 \times 256$  resolution, as well as VQ-VAE models for each of CLEVR and Relational CLEVR at  $128 \times 128$  resolution. These choices of resolution follow earlier work in




Figure 6. Effect of varying the  $w_{\text{smile}}$  concept weight from  $-3.0$  to  $3.0$  while keeping  $w_{\text{male}} = w_{\text{no\_glasses}} = 3.0$ .

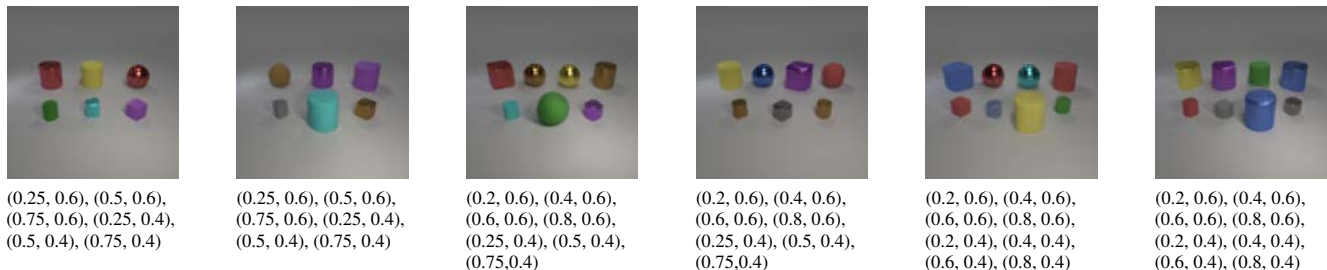


Figure 7. Compositional out-of-distribution generation: Positional CLEVR training images contain no more than 5 objects per image, but our compositional method allows 6 or more objects to appear in the same image via compositional sampling.

compositional generation with these 3 datasets [6]. We find in practice that VQ-VAE (without the adversarial loss) is sufficient for high-fidelity reconstruction of the two CLEVR datasets due to the smaller resolution and visual simplicity, while VQ-GAN is required for realistic reconstructions of FFHQ. Unlike [6], our choice of training regime produces FFHQ images directly at  $256 \times 256$ , so a post-upsampling step is not required during evaluation. We train with a perceptual loss [36] in addition to the MSE loss for all datasets (and the learned adversarial loss for FFHQ). Full details of model training are in Appendix.

### 4.3. Quantitative Evaluation of Compositional Generation

For each dataset, following [6] we evaluate (compositionally) generated image samples according to both FID (Fréchet Inception Distance [37]) and binary classification accuracy (defined as whether a specified attribute, or collection of attributes, is present or not in the corresponding generated output image according to a pre-trained classifier). These two metrics are chosen in order to assess each model’s ability to match the target distribution from the perspective of both perceptual image quality and visual accuracy. We recognise that FID has limitations as a measure of quality/diversity [38], however it remains a standard metric for evaluation of image generation methods [6, 12, 39, 40], and serves as a highly informative tool for our purposes.

We conduct all quantitative experiments for 1, 2 and 3 attributes per image for all 3 datasets, totalling 9 quantitative experiments. We use a temperature of 0.9 when generating samples for our quantitative experiments. Details of how accuracy scores are obtained are in Appendix. Excluded

from our experiments are unconditional methods such as R3GAN [39] and those with a primary application to text-to-image (MM-DiT [40]).

In comparison to the 6 baseline compositional results reported in [6], our method *exceeds or matches the accuracy of the previous state-of-the-art in all nine settings*, while attaining highly competitive FID scores across the three datasets. Particularly noteworthy are our accuracy results for Positional CLEVR, for which our method attains an error rate of 0.70%, 1.82% and 4.96% on 1, 2 and 3 input components respectively, where the previous state-of-the-art attained error rates of 13.58%, 59.20% and 68.64% respectively (Table 1). We see similarly dramatic improvements much harder Relational CLEVR dataset (Table 2) and significant improvements on the FFHQ dataset (Table 3).

We speculate that the dramatic accuracy improvements offered by our method can be attributed to the fact that the introduction of the discrete representation learning step (VQ-VAE or VQ-GAN) facilitates the emergence of an expressive and richly compositional “visual language” to represent images, while the conditional parallel token model offers a strongly regularised and highly calibrated model for the visual language. We conjecture that these effects synergise to produce an efficient, accurate and robust compositional method.

While the FID scores of our method are not the best in all instances (Tables 1, 2, 3), our method obtains lower FID than the next-best competitor (ranked by error rate) in 7 out of 9 experiments. We recognise that there will always be a trade-off between accuracy and FID; by design, higher concept weights result in stronger/more biased expression of the desired features (Fig. 6). Nonetheless, scatter plots of

Table 1. Quantitative results (error rate and FID score) on the Positional CLEVR dataset

Method	1 Component		2 Components		3 Components	
	Err (%) ↓	FID ↓	Err (%) ↓	FID ↓	Err (%) ↓	FID ↓
StyleGAN2-ADA [41]	62.72±1.37	57.41	-	-	-	-
StyleGAN2 [42]	98.96±0.29	51.37	99.96±0.04	23.29	100.00±0.00	19.01
LACE [4]	99.30±0.24	50.92	100.00±0.00	22.83	100.00±0.00	19.62
GLIDE [43]	99.14±0.26	61.68	99.94±0.06	38.26	100.00±0.00	37.18
EBM [5]	29.46±1.29	78.63	71.78±1.27	65.45	92.66±0.74	58.33
Composed GLIDE [6]	<u>13.58±0.97</u>	<u>29.29</u>	<u>40.80±1.39</u>	<u>15.94</u>	<u>68.64±1.31</u>	<u>10.51</u>
Ours	<b>0.70±0.24</b>	<b>13.76</b>	<b>1.82±0.38</b>	<b>15.30</b>	<b>4.96±0.61</b>	<b>16.23</b>

error rate against FID indicate that our method empirically lies on the Pareto front in all 9 settings (please see Fig. 1 and Appendix for full scatter plots).

#### 4.4. Qualitative Experiments

Here we also qualitatively investigate the usefulness of our approach outside of the rigorous quantitative experimental settings. In particular, we investigate the controllability offered by logical conjunction and negation of prompts, as well as the qualitative effect of concept weighting. In addition to the models trained for our quantitative experiments, some of the experiments below apply our method to a pre-trained text-to-image parallel token prediction model ([10]). We choose the aMUSEd [3] implementation of MUSE because it is publicly accessible (as of the time of writing), in addition to being trained on a large, open dataset [44], which facilitates the open-ended generation which we aim to explore here. Additional qualitative results are in Appendix.

**Concept negation:** Fig. 3 demonstrates the application of concept negation using aMUSEd text-to-image parallel token prediction [3]. We compare each example against a single-prompt CFG baseline using the same model. We focus on problematic cases where the underlying text-image model is incapable of properly interpreting negation in the linguistic sense, which is especially pertinent when the concept being negated may be considered an essential characteristic of the concept from which it is being negated (e.g. “a king **not** wearing a crown”). Our method allows us to achieve more specific outputs without changing or fine-tuning the underlying model, even in cases where the underlying model fails to comprehend the original negated prompt.

**Out-of-distribution generation:** We demonstrate our model’s ability to generalise to compositions of conditions that are not seen in training. We focus on the (positional) CLEVR dataset, in which individual training samples have at most 5 objects per image. Fig. 7 contains generated samples for input conditions specifying between 6 and 8 objects per image. We make two key observations of Fig. 7: (1) that our method successfully generalises outside the distribution of the training data and (2) that re-running the same input gives varied outputs, i.e. the model has not over-fit to always generate the same objects in the same position.

**Varying the concept weight:** Fig. 6 illustrates the effect of varying the concept weighting parameter  $w$  for a specific input condition (in this case, the weighting of the “no glasses” attribute of FFHQ). Keeping other concept weights the same ( $w_{\text{no\_glasses}} = w_{\text{male}} = 3.0$ ), we vary  $w_{\text{smile}}$  from  $-3.0$  to  $3.0$ . The outputs in Fig. 6 are consistent with the expectation that the concept weighting method should allow for an interpretable degree of controllability over the generated outputs. Similar visualisations for the other two FFHQ attributes are in Appendix.

#### 4.5. Parameter Count and Sampling Time

Table 4 contains a comparison of our method to the most similar methods in the literature (composed EBM [5] and composed GLIDE [6]). We compare our method against these methods in particular for 2 reasons: (1) they are compositional and iterative like our own method, and (2) they are generally closest to ours (lowest) in terms of error rate on the three datasets studied. We compare in terms of total parameters and the time taken to generate both a single image and a batch of 25 images on our hardware (NVIDIA GeForce RTX 3090) with 3 input conditions (Positional CLEVR dataset). Runs of baseline methods use the official PyTorch implementations from [6] with default settings (corresponding to the baseline results in Tables 1, 2 and 3). The results in Table 4 show that our method runs in a fraction of the time of existing approaches while having a comparable number of parameters (and smaller error rate: see Tables 1, 2, 3). Altogether, we see a  $2.3\times$  to  $12.0\times$  speedup across our speed experiments compared with the baselines.

### 5. Discussion and Limitations

Through varied quantitative and qualitative experiments, we have demonstrated that our formulation for compositional generation with iterative sampling methods is readily applicable to a range of tasks for both newly trained and out-of-the-box pre-trained models. We demonstrated state-of-the-art performance in terms of the error rate of the generated results, in addition to obtaining competitive sample quality as measured by FID scores. This is achieved with minimal extra cost in terms of memory, since only the accumulated log-probability outputs need to be retained at inference time. The

Table 2. Quantitative results (error rate and FID score) on the Relational CLEVR dataset

Method	1 Component		2 Components		3 Components	
	Err (%) ↓	FID ↓	Err (%) ↓	FID ↓	Err (%) ↓	FID ↓
StyleGAN2-ADA [41]	<u>32.29</u> ±1.32	<u>20.55</u>	-	-	-	-
StyleGAN2 [42]	79.82±1.14	22.29	98.34±0.36	30.58	99.84±0.11	31.30
LACE [4]	98.90±0.30	40.54	99.90±0.09	40.61	99.96±0.04	40.60
GLIDE [43]	53.80±1.41	<b>17.61</b>	91.14±0.80	<b>28.56</b>	98.64±0.33	40.02
EBM [5]	<b>21.86</b> ±1.17	44.41	<u>75.84</u> ±1.21	55.89	<u>95.74</u> ±0.57	58.66
Composed GLIDE [6]	39.60±1.38	29.06	78.16±1.17	29.82	97.20±0.47	<b>26.11</b>
Ours	<b>21.84</b> ±1.17	30.00	<b>56.94</b> ±1.40	<u>28.87</u>	<b>85.70</b> ±0.99	<u>30.34</u>

Table 3. Quantitative results (error rate and FID score) on the FFHQ dataset

Method	1 Component		2 Components		3 Components	
	Err (%) ↓	FID ↓	Err (%) ↓	FID ↓	Err (%) ↓	FID ↓
StyleGAN2-ADA [41]	8.94±0.81	<b>10.75</b>	-	-	-	-
StyleGAN2 [42]	41.10±1.39	<u>18.04</u>	69.32±1.30	<u>18.06</u>	83.04±1.06	<u>18.06</u>
LACE [4]	2.40±0.43	28.21	<u>4.34</u> ±0.58	36.23	<u>19.12</u> ±1.11	34.64
GLIDE [43]	1.34±0.33	20.30	51.32±1.41	22.69	<u>72.76</u> ±1.26	21.98
EBM [5]	1.26±0.32	89.95	6.90±0.72	99.64	69.99±1.30	335.70
Composed GLIDE [6]	<u>0.74</u> ±0.24	18.72	7.32±0.74	<b>17.22</b>	31.14±1.31	<b>16.95</b>
Ours	<b>0.22</b> ±0.13	21.52	<b>0.62</b> ±0.22	28.25	<b>0.82</b> ±0.26	33.80

Table 4. Parameter counts and sample times for 3 input components

Method	Parameters	Sample Time Per Image	Sample Time Per Batch of 25
EBM	33 mil	5.99s±0.17	108.57s±0.93
Composed GLIDE	385 mil	4.92s±0.17	73.92s±0.70
Ours	108 mil	<b>2.11s</b> ±0.39	<b>9.08s</b> ±0.39

simplicity of our method offers further advantages, including ease of implementation (facilitating integration with existing discrete generation pipelines) in addition to improved interpretability, since the composition operator can be thought of as directly taking the “intersection” between two discrete distributions.

The strong quantitative metrics of our method are complemented by its *out-of-distribution* generation capability and *controllability*. The significance of the results of our quantitative experiments is further reinforced by the fact that we used the same experimental settings for all three of the datasets studied, without extensive fine-tuning of hyperparameters, training runs or model architecture. Our method further provides a 2.3× to 12× speedup over comparable approaches on our hardware.

Similarly to compositional methods for continuous processes [5, 6], our method requires  $(n + 1)$  times the number of feed-forward operations compared to standard iterative approaches of the same architecture, where  $n$  is the number of conditions imposed on the output. This is a direct consequence of the mathematical formulation of the approach, however this is largely mitigated by the fact that our method can produce accurate and high-quality outputs in only a small number of iterations [5, 6].

Our method makes the necessary assumption that in-

put conditions are conditionally independent given  $\mathbf{x}$ , i.e.  $P(\mathbf{c}_1, \mathbf{c}_2|\mathbf{x}) \propto P(\mathbf{c}_1|\mathbf{x})P(\mathbf{c}_2|\mathbf{x})$  for all condition pairs  $\mathbf{c}_1, \mathbf{c}_2$ . It is possible in practical scenarios that this underlying assumption of our approach does not hold, for example due to biases in the training data. The importance-weighting capability of our method can mitigate this in part by allowing the user to compensate for potential biases. However, we speculate that greater robustness would be better achieved through an unbiased backbone model. Training unbiased models for image generation is beyond the scope of this work, and remains an open challenge, especially in the context of text-to-image generation [45]. Further to this, we have not yet explored principled methods for choosing the condition weighting coefficients  $w_i$ , which may be an interesting direction for future work (e.g. producing a learned concept-weighting policy).

## 6. Conclusion

We have proposed a novel method for enabling precisely controllable conditional image generation by composing discrete iterative generative processes. The empirical success of our method across the axes of sampling speed, error rate and FID demonstrates a conceptual step beyond the previous state-of-the-art for compositional generation. We further show that our approach can be applied to an out-of-the-box pre-trained text-to-image model to allow for principled and controllable generation without any fine-tuning. The prospect of applying our method outside of image generation (such as multi-prompt text generation) remains an intriguing possibility for future work. Altogether, we believe our work provides a strong foundation for future work in the direction of controllable generation in discrete spaces.

## References

- [1] Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation. In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 78723–78747. Curran Associates, Inc., 2023. 1
- [2] Baihan Lin, Djallel Bouneffouf, and Irina Rish. A survey on compositional generalization in applications. *arXiv preprint arXiv:2302.01067*, 2023. 1
- [3] Suraj Patil, William Berman, Robin Rombach, and Patrick von Platen. amused: An open muse reproduction. *arXiv preprint arXiv:2401.01808*, 2024. 2, 3, 7
- [4] Weili Nie, Arash Vahdat, and Anima Anandkumar. Controllable and compositional generation with latent-space energy-based models. *Advances in Neural Information Processing Systems*, 34:13497–13510, 2021. 1, 2, 7, 8
- [5] Yilun Du, Shuang Li, and Igor Mordatch. Compositional visual generation with energy based models. *Advances in Neural Information Processing Systems*, 33:6637–6647, 2020. 2, 7, 8
- [6] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B Tenenbaum. Compositional visual generation with composable diffusion models. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVII*, pages 423–439. Springer, 2022. 1, 2, 4, 5, 6, 7, 8
- [7] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2, 3, 4, 5
- [8] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis, 2020. 2, 3, 4, 5
- [9] Sam Bond-Taylor, Peter Hesse, Hiroshi Sasaki, Toby P. Breckon, and Chris G. Willcocks. Unleashing transformers: Parallel token prediction with discrete absorbing diffusion for fast high-resolution image generation from vector-quantized codes. In *European Conference on Computer Vision (ECCV)*, 2022. 2, 3, 4, 5
- [10] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023. 2, 3, 7
- [11] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11315–11325, 2022. 2, 4
- [12] Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. *Advances in Neural Information Processing Systems*, 32, 2019. 2, 6
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [14] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. 2
- [15] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2901–2910, 2017. 2, 5
- [16] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 2, 5
- [17] G.E. Hinton. Products of experts. *9th International Conference on Artificial Neural Networks: ICANN '99*, 1999:1–6, 1999. doi: 10.1049/CP:1999107510.1049/CP:19991075. 2, 3
- [18] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 2
- [19] Tailin Wu, Takashi Maruyama, Long Wei, Tao Zhang, Yilun Du, Gianluca Iaccarino, and Jure Leskovec. Compositional generative inverse design. *arXiv preprint arXiv:2401.13171*, 2024. 2
- [20] Jianrong Zhang, Hehe Fan, and Yi Yang. Energymogen: Compositional human motion generation with energy-based diffusion model in latent space. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17592–17602, 2025. 2
- [21] Xue Bin Peng, Michael Chang, Grace Zhang, Pieter Abbeel, and Sergey Levine. Mcp: Learning composable hierarchical control with multiplicative compositional policies. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/95192c98732387165bf8e396c0f2dad2-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/95192c98732387165bf8e396c0f2dad2-Paper.pdf). 2
- [22] Saeed Masoudnia and Reza Ebrahimpour. Mixture of experts: a literature survey. *Artificial Intelligence Review*, 42:275–293, 2014. 3
- [23] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019. 3
- [24] Robert Gray. Vector quantization. *IEEE Assp Magazine*, 1(2):4–29, 1984. 3
- [25] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 3
- [26] Abril Corona-Figueroa, Sam Bond-Taylor, Neelanjan Bhowmik, Yona Falinie A Gaus, Toby P Breckon, Hubert PH Shum, and Chris G Willcocks. Unaligned 2d to 3d translation with conditional vector-quantized code diffusion using transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14585–14594, 2023. 3
- [27] Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete

- flow matching. *Advances in Neural Information Processing Systems*, 37:133345–133385, 2024. 3
- [28] Jian Han, Jinlai Liu, Yi Jiang, Bin Yan, Yuqi Zhang, Zehuan Yuan, Bingyue Peng, and Xiaobing Liu. Infinity: Scaling bitwise autoregressive modeling for high-resolution image synthesis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 15733–15744, 2025. 3
- [29] Zhenlin Xu, Marc Niethammer, and Colin A Raffel. Compositional generalization in unsupervised compositional representation learning: A study on disentanglement and emergent language. *Advances in Neural Information Processing Systems*, 35:25074–25087, 2022. 4
- [30] James Joyce. Bayes’ theorem. *The Stanford Encyclopedia of Philosophy*, 2003. 4
- [31] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*, 2020–December, 5 2020. ISSN 10495258. URL <https://arxiv.org/abs/2005.14165v4>. 4
- [32] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. *Advances in Neural Information Processing Systems*, 2017–December:5999–6009, 6 2017. ISSN 10495258. URL <https://arxiv.org/abs/1706.03762v7>. 4
- [33] Jacob Devlin, Ming Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 1:4171–4186, 10 2018. URL <https://arxiv.org/abs/1810.04805v2>. 4
- [34] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 5
- [35] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 5
- [36] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6
- [37] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 6
- [38] Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and Sanjiv Kumar. Re-thinking fid: Towards a better evaluation metric for image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9307–9315, 2024. 6
- [39] Nick Huang, Aaron Gokaslan, Volodymyr Kuleshov, and James Tompkin. The gan is dead; long live the gan! a modern gan baseline. *Advances in Neural Information Processing Systems*, 37:44177–44215, 2024. 6
- [40] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. 6
- [41] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33:12104–12114, 2020. 7, 8
- [42] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 7, 8
- [43] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 7, 8
- [44] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 7
- [45] Yixin Wan, Arjun Subramonian, Anaëlia Ovalle, Zongyu Lin, Ashima Suvarna, Christina Chance, Hritik Bansal, Rebecca Pattichis, and Kai-Wei Chang. Survey of bias in text-to-image generation: Definition, evaluation, and mitigation. *arXiv preprint arXiv:2404.01030*, 2024. 8