

Rethinking cluster-conditioned diffusion models for label-free image synthesis

Nikolas Adaloglou

Heinrich Heine University of Dusseldorf
 adaloglo@hhu.de

Felix Michels

Heinrich Heine University of Dusseldorf
 felix.michels@hhu.de

Tim Kaiser

Heinrich Heine University of Dusseldorf
 tikai103@hhu.de

Markus Kollmann

Heinrich Heine University of Dusseldorf
 markus.kollmann@hhu.de

Abstract

Diffusion-based image generation models can enhance image quality when conditioned on ground truth labels. Here, we conduct a comprehensive experimental study on image-level conditioning for diffusion models using cluster assignments. We investigate how individual clustering determinants, such as the number of clusters and the clustering method, impact image synthesis across three different datasets. Given the optimal number of clusters with respect to image synthesis, we show that cluster-conditioning can achieve state-of-the-art performance, with an FID of 1.67 for CIFAR10 and 2.17 for CIFAR100, along with a strong increase in training sample efficiency. We further propose a novel empirical method to estimate an upper bound for the optimal number of clusters. Unlike existing approaches, we find no significant association between clustering performance and the corresponding cluster-conditional FID scores. Code is available at <https://github.com/HHU-MMBS/cedm-official-wavc2025>

1. Introduction

Diffusion models have enabled significant progress in many visual generative tasks, such as image synthesis [26, 36, 53] and manipulation [75]. Conditioning diffusion models on human-annotated data is today’s standard practice as it significantly improves the image fidelity [7, 21, 23]. Image-level conditioning is typically realized by using associated text captions [61, 62] or class labels, if available [26, 49]. Nonetheless, human-annotated labels are costly and often contain inaccuracies [6, 9], while publicly available image-text pairs can be non-descriptive [8].

Large, human-annotated datasets are prohibitively costly [55] and hard to obtain in a plethora of real-life applications such as medical image synthesis [38]. Human annotation is subject to the label collection procedure and can have vary-



Figure 1. An ideal image-level conditioning should group images based on shared patterns, shown in the same row, which do not always align with human labels, indicated above each image (CIFAR100 [42] samples).

ing annotation hierarchies [3, 10, 28]. An ideal conditioning signal would be based solely on shared characteristics corresponding to general visual concepts, as illustrated in Fig. 1. To provide a ground truth (GT) label-free alternative that is applicable to unlabelled datasets, we focus on cluster-based conditioning on diffusion models. Image clustering refers to algorithmically assigning a semantic group to an image, called a cluster, given an *a priori* number of such groups [13, 17, 35].

Maximizing cluster alignment with ground truth (GT) labels does not guarantee optimal cluster-conditional image generation. To create visually distinguishable groups (Fig. 1), a clustering algorithm would sometimes group images into more fine-grained groups than the GT labels, i.e., sub-grouping analog and digital devices that map to the

same label or merge similarly looking images with different labels (merging different types of trees).

Deep image clustering has recently seen significant progress [4, 73, 76], especially when utilizing off-the-shelf features learned from large-scale pre-training [2, 3, 71]. Currently, the relationship between clustering performance and cluster-conditional image synthesis performance remains unclear for several reasons. First, clustering metrics capture the alignment to GT labels, whereas generative models may benefit more from visually distinguishable groups [31]. Second, clustering metrics cannot be computed for unlabeled datasets. Third, existing metrics are not suitable for fairly evaluating cluster assignments with varying numbers of clusters. Therefore, they cannot help to determine the optimal cluster granularity w.r.t. generative performance for an arbitrary dataset. This aspect has not been sufficiently explored to date [23, 46].

Existing cluster-conditioned generative approaches [23, 31, 46] often adopt simple generative baselines [26] and have mainly investigated k -means clustering on balanced classification datasets. The adoption of k -means is justified due to its ability to work with off-the-shelf feature extractors and its non-parametric nature, apart from setting the number of clusters. Conversely, it is well-established that k -means is suboptimal for image clustering because its cluster assignments are highly imbalanced [31, 71]. While cluster-conditioned approaches are sensitive to the number of clusters [23], there is currently no method to pre-determine this number. Previous methods have overlooked this [23, 31, 46], significantly hindering their applicability to unlabeled datasets. Recent feature-based image clustering methods have demonstrated superior performance compared to k -means [2, 71], and their efficacy in conditional image generation is yet to be investigated. For the above reasons, cluster-conditioning methods underperform GT conditional ones to date.

In this paper, we systematically study clustering as a label-free condition for diffusion models. We demonstrate that optimal cluster granularity achieves state-of-the-art FID scores (1.67 for CIFAR10 and 2.17 for CIFAR100) while enhancing training sample efficiency. We propose a computationally efficient method to derive an upper cluster bound using feature-based clustering, which narrows the search space for optimal image synthesis performance. We validate this upper bound across three datasets and two generative metrics. Finally, we find no significant association between clustering performance and cluster-conditional image synthesis performance across various clustering methods, performance metrics, and pre-trained models.

2. Related Work

2.1. Conditional generative models

Early attempts on controllable image synthesis adopted generative adversarial networks (GANs) [48]. By conditioning both the generator and discriminator on human labels, GANs can produce images for a specific GT label, even at the scale of ImageNet [12]. Recently, diffusion models (DMs) have emerged as an expressive and flexible category of generative models [27, 37, 59, 64, 65]. Internal guidance-based methods [22, 27, 29, 30] further improved the flexibility and visual fidelity of DMs during sampling such as classifier and classifier-free guidance. DMs exhibit an enormous number of label-conditional variants [26, 36, 41, 49, 54]. Recently, Stein et al. [66] proposed the adoption of Fréchet DINOv2 [50] distance as it aligns better with human preferences and demonstrated that DMs achieve the highest perceptual realism.

2.2. Alternative conditioning of generative models

Unlike internal guidance [27, 29, 30], external conditioning signals are computed from the training set using additional models, which is commonly referred to as *self-conditioning*. Self-conditioning signals can be roughly divided into image-level and sub-image-level. Image-level conditioning refers to a single condition for all pixels in an image, such as cluster assignments or text captions [59]. Sub-image-level conditions refer to specific parts or regions of an image [5, 31]. For instance, Hu et al. [31] extend DMs to incorporate conditioning of bounding boxes or segmentation masks.

Feature-based conditioning. Using image or text features can also provide an informative conditioning signal [11, 45, 57, 63]. Bordes et al. [11] condition DMs directly on image representations (*oracle features from real images*) and reveal their impact on visual fidelity. Instance-conditioned GAN combines GT labels with features from each image’s NN set [15]. Ramesh et al. [57] leverage the zero-shot capability of CLIP to condition DMs on language-guided image representations to improve diversity. Zhou et al. provide a language-free framework [79] that directly generates text features for images using CLIP. Still, CLIP models require an *a priori* set of descriptive candidate captions, such as label names.

Cluster-based conditioning. Similar to label conditioning, conditioning on cluster assignments facilitates DMs by allowing them to specialize on a distinct set of shared visual features [23, 31, 46]. In [23], Bao et al. computed the k -means clusters offline using contrastive learned features on the training data. However, their approach could not outperform the label-conditioned models. The first cluster-based approach that achieves competitive performance compared to GT labels leverages pre-trained feature extractors [31],

which is the closest to our work. In [31], the authors attempt to weakly establish a correlation between clustering and cluster-conditional generative performance for different self-supervised models. The observed correlation cannot be tested on unlabeled datasets, and its sensitivity to the number of clusters is unknown. While it was demonstrated that DINO [14] provides the most informative k -means clusters [31], their method lacks a strategy for choosing the cluster granularity for unlabeled datasets.

3. Method

3.1. Notations and prerequisites

We consider the frequent scenario where we have access to an **unlabeled** dataset D and a pre-trained feature extractor $g(\cdot)$. We denote the number of ground truth (GT) labels as C_{GT} , which we assume to be unknown. To distinguish cluster-based conditioning from the image clustering task, we denote the number of *visual groups* C_V of D as the optimal number of clusters w.r.t. image synthesis (e.g. as measured by FID [25]). We adopt the diffusion model approach of Karras et al. (EDM [36]) as a baseline, which introduced various improvements to the standard DDPM [26]. We provide an overview of EDM in the supplementary material. Although we focus on EDM throughout this work, our method can be applied to any conditional generative model.

TEMI clustering. Given an *a priori* determined number of clusters C and a feature extractor $g(\cdot)$, TEMI [2] first mines the m nearest neighbors (NN) of all $\mathbf{x} \in D$ in the feature space of $g(\cdot)$ based on their cosine similarity. We denote the set of NN for \mathbf{x} by $S_{\mathbf{x}}$. During training, TEMI randomly samples \mathbf{x} from D and \mathbf{x}' from $S_{\mathbf{x}}$ to generate image pairs with similar (visual) features. A self-distillation framework is introduced to learn the cluster assignments with a teacher and student head $h_t(\cdot)$ and $h_s(\cdot)$ that share the same architecture (i.e. 3-layer MLP) but differ w.r.t. their parameters. The features of the image pair are fed to the student and teacher heads $h_s^i(\mathbf{z})$, $h_t^i(\mathbf{z}')$, where $i \in \{1, \dots, H\}$ is the head index and $\mathbf{z} = g(\mathbf{x})$, $\mathbf{z}' = g(\mathbf{x}')$. The outputs of the heads are converted to probabilities $q_s(c|\mathbf{x})$ and $q_t(c|\mathbf{x}')$ using a softmax function. The TEMI objective,

$$\mathcal{L}_{\text{TEMI}}^i(\mathbf{x}, \mathbf{x}') := -\frac{1}{H} \sum_{j=1}^H \sum_{c'=1}^C q_t^j(c'|\mathbf{x}) q_t^j(c'|\mathbf{x}') \log \sum_{c=1}^C \frac{(q_s^i(c|\mathbf{x}) q_t^i(c|\mathbf{x}'))^\gamma}{\tilde{q}_t^i(c)}, \quad (1)$$

maximizes the pointwise mutual information between images \mathbf{x} and \mathbf{x}' , using the clustering index c as information bottleneck [35]. Here, $\tilde{q}_t^i(c)$ is an estimate of the teacher cluster distribution $\mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [q_t^i(c|\mathbf{x})]$, which can be com-

puted by an exponential moving average over mini-batches B defined as $\tilde{q}_t^i(c) \leftarrow \lambda \tilde{q}_t^i(c) + (1 - \lambda) \frac{1}{|B|} \sum_{\mathbf{x} \in B} q_t^i(c|\mathbf{x})$.

The hyperparameter $\gamma \in (0.5, 1]$ avoids the collapse of all sample pairs into a single cluster, and $\lambda \in (0, 1)$ is a momentum hyperparameter. The loss function Eq. (1) is further symmetrized and averaged over heads to obtain the final training loss. After training, the cluster assignments $c^*(\mathbf{x}) = \arg \max_c q_t(c|\mathbf{x})$ and the empirical cluster distribution $q(c)$ on the training set is computed from the TEMI head with the lowest loss. We use $c^*(\mathbf{x})$ as a condition of the generative model during training and sample from $q(c)$ to generate images. Adaloglou et al. [2] show that the value $\gamma = 0.6$ enforces a close to uniform cluster utilization and achieves state-of-the-art clustering accuracy for C_{GT} [2].

3.2. Cluster-conditional EDM (C-EDM)

We consider k -means and TEMI clusters using off-the-shelf feature extractors with C_V clusters to compute cluster assignments. We experimentally observed that $C_V > C_{GT}$, even though this is not a requirement of our approach. We highlight that clustering using $C > C_{GT}$ while enforcing a uniform cluster utilization (such as TEMI with $\gamma = 0.6$) has not been previously explored [2, 32, 44, 71] as it reduces the cluster alignment with the GT labels. We denote the cluster-conditioned EDM model as *C-EDM*.

A direct estimation of C_V is hard to obtain for unlabeled datasets. In principle, C_V can be found using a hyperparameter search with FID as an evaluation metric, which is computationally expensive. In contrast to existing approaches that perform a restricted hyperparameter search around C_{GT} , we propose a new metric that allows us to derive an upper bound for C_V (Sec. 3.3), which requires no prior knowledge about the dataset.

3.3. Estimating the upper cluster bound

We aim to find an upper cluster bound C_{max} using TEMI ($\gamma = 0.6$), such that $C_V < C_{max}$. During this computation, we neither use the generative model nor any additional information. We highlight that previous works [23, 31, 46] have overlooked this design choice and iterate within a small range around C_{GT} . We denote the number of utilized clusters (one training sample is assigned to it) as C^u after TEMI clustering. The TEMI cluster utilization ratio is defined as $r_C := \frac{C^u}{C} \leq 1$. Importantly, there is no guarantee that the full cluster spectrum will be utilized. Unlike k -means that always has $r_C = 1$, we observe that as the number of clusters C increases, the TEMI cluster utilization ratio r_C typically decreases, which provides dataset-specific information. Intuitively, TEMI clustering with $\gamma = 0.6$ is enforcing $r_C \rightarrow 1$, and the observation of $r_C < 1$ is an indication that the maximum number of TEMI clusters is reached for $\gamma = 0.6$.

The search for C_{max} involves doubling the number of

Generative methods	CIFAR10		CIFAR100		FFHQ-64	FFHQ-128
	Unlabelled	GT	Unlabelled	GT	Unlabelled	Unlabelled
DDPM [26, 31]	3.17	-	10.7	9.7	-	-
PFGM++ [74]	-	1.74	-	-	2.43	-
EDM [36]	1.97	1.79	-	-	2.39	-
EDM (our reproduction)	2.07	1.81	3.41	2.21	2.53	5.93
Self-conditional methods						
IDDPM+k-means ($k=10$) [23]	2.23	-	-	-	-	-
DDPM+k-means [†] ($k=400$) [31]	-	-	9.6	-	-	-
C-EDM+k-means [†] (Ours)	1.69	-	2.21	-	1.99	-
C-EDM+TEMI [†] (Ours)	1.67	-	2.17	-	2.09	4.40
EDM+oracle features [†] (Ours)	2.21	-	2.25	-	1.77	-

Table 1. State-of-the-art generative model comparison: FID (\downarrow) for various GT labeled and unlabelled benchmarks. We use $C_V = 100, 200, 400$ clusters for CIFAR10, CIFAR100, and FFHQ-64, respectively. Self-conditional methods with [†] use the pre-trained DINO ViT-B feature extractor. Ground truth and oracle feature conditioning results are marked in gray color as they are **not** a fair comparison.

clusters C at each iteration until the r_C falls below a threshold, $r_C \leq \alpha$, followed by a more fine-grained grid search. C_{max} is then defined as the highest C for which $r_C > \alpha$. This results in a worst-case time complexity of $O(\log_2(\frac{C_{max}}{C_{start}}))$. This is conceptually similar to the elbow method [39, 70]. In this sense, the proposed heuristic has the same limitations as the elbow method. Yet, it is currently the only method that provides a practitioner with a starting point for an unlabelled dataset. After detecting C_{max} , we perform a bounded grid search to find $C_V \in [2, C_{max}]$ using the generative model. To obtain a simple cross-dataset estimate, we empirically find a cutoff threshold $r_C \leq 0.96$ to work well across three datasets.

4. Experimental evaluation

4.1. Datasets, models, and metrics

Following prior works [2, 31, 78], we use DINO ViT-B pre-trained on ImageNet [20] for image clustering. We report the Fréchet inception distance (FID) [25] to quantify the image generation performance as it simultaneously captures visual fidelity and diversity. To facilitate future comparisons, we also compute the Fréchet DINOv2 distance (FDD) [66] that replaces the features of InceptionNet-v3 [68, 69] with DINOv2 ViT-L [50]. FID and FDD are averaged over three independently generated sample sets of 50K images each. To measure the cluster alignment w.r.t. GT labels, we use the adjusted normalized mutual information (ANMI) as in [32].

We follow the default hyperparameter setup for EDM [36] and TEMI [2]. We denote the number of samples (in millions) seen during training as M_{img} . We use $M_{img} = 200$ when comparing C-EDM with other baselines and state-of-the-art methods (Tab. 1), and $M_{img} = 100$ when comparing across different number of clusters. We train

the TEMI clustering heads for 200 epochs per dataset. All the experiments were conducted on 4 NVIDIA A100 GPUs with 40GB VRAM each. On this hardware, TEMI clustering was more than $50\times$ faster than training EDM on FFHQ-64, which requires more than 2 days for 200 M_{img} with a batch size of 512. Additional implementation details and hyperparameters can be found in the supplementary material. We verify our approach across CIFAR10, CIFAR100 [42] and FFHQ. CIFAR10 and CIFAR100 have 50K samples and 32^2 resolution images, while we use the 64x64 and 128x128 versions of FFHQ (FFHQ-64, FFHQ-128) consisting of 70K samples. Finally, diffusion sampling methods are not included, as they can be incorporated into any diffusion model [27, 40, 60].

4.2. State-of-the-art comparison for image synthesis

Comparison with state-of-the-art unconditional generative models. With respect to FID, we found $C_V = 100, 200, 400$ to be close to optimal cluster granularities for CIFAR10, CIFAR100, and FFHQ-64, respectively. We use these values to report FIDs compared to GT conditional and unconditional generative models in Tab. 1. Compared to unconditional methods, C-EDM achieves an average relative FID improvement of 24.4% and 24.9% using TEMI and k-means clusters, respectively. More importantly, previous cluster-conditional approaches did not achieve near-state-of-the-art FIDs because they: a) adopted non-competitive diffusion baselines such as DDPM [26], b) did not consider pre-trained feature extractors for clustering [23], c) did not use the optimal cluster granularity [31]. For instance, Hu et al. [31] used 400 clusters on CIFAR100 using DINO ViT-B for clustering, while Fig. 3 shows that 200 clusters lead to a superior FID using C-EDM (11.2% relative improvement).

Comparison with state-of-the-art GT conditional




Figure 2. FID (y-axis) versus seen samples during training in millions (x-axis). TEMI and k-means clusters are computed using the representations of DINO ViT-B [14]. We used $C_V = 100, 200, 400$ for CIFAR10, CIFAR100 and FFHQ-64 respectively. The training sample efficiency compared to the unconditional baseline is indicated by the arrow. Best viewed in color.

generative models. Intriguingly, using C-EDM with C_V clusters, we report small improvements compared to GT label conditioning on CIFAR10 and CIFAR100 in Tab. 1. For instance, 4.77% mean relative improvement in FID using TEMI. Even though cluster-conditioning is primarily designed for unconditional generation, such as FFHQ, we demonstrate that it can match or outperform GT labels, which showcases the effectiveness of C-EDM.

Sample efficiency. In addition to the reported gains in FID, we study the sample efficiency of C-EDM during training across datasets in Fig. 2. On CIFAR10 and CIFAR100, the training sample efficiency for $C = C_V$ compared to the unconditional model ($C = 1$) peaks at $5\times$, where C-EDM with $M_{img} = 40$ outperforms the unconditional model at $M_{img} = 200$. On FFHQ-64, which is not a classification dataset, we report a sample efficiency of $3.3\times$. TEMI and k -means demonstrate identical sample efficiency compared to the unconditional model. Intuitively, a more informative conditioning signal enables learning the data distribution faster. Upon visual inspection, we could identify FFHQ clusters for $C = C_V$ with easily distinguishable visual patterns, such as groups of images with beanies, smiling faces, glasses, sunglasses, hats, and kids (see supplementary). Finally, we find that the sample efficiency of C-EDM against GT conditional EDM heavily depends on the quality and granularity of the GT labels.

4.3. Cluster utilization ratio and discovered upper bounds

We depict how the TEMI cluster utilization r_C changes in tandem with FID for different numbers of clusters in Fig. 3. Crucially, this approach also applies to FFHQ, where no GT labels exist. Moreover, we can discern certain patterns across datasets: (i) C_V always has a high utilization ratio, (ii) the majority of experiments with $C < C_{max}$ outperform the unconditional model at $M_{img} = 100$, and (iii)

$C_V > C_{GT}$ for CIFAR10 and CIFAR100, which is in line with [31]. Even though the choice of $\alpha = 0.96$ is not guaranteed to be generally applicable and is based on empirical evidence, a more or less strict choice can be used based on the practitioner’s computational budget. The introduced upper bound can be computed with negligible computational overhead and without access to GT labels or training the generative model, which on large scales (ImageNet) can require up to 4MWh per experiment [37].

4.4. Investigating the connection between clustering and cluster-conditional image synthesis

FID and ANMI. On image clustering benchmarks, TEMI ($\gamma = 0.6$) outperforms k -means, where we respectively measure an ANMI of 60.6% versus 59.2% on CIFAR10 and 72.3% versus 67.6% on CIFAR100. However, their generative performance displays negligible differences in terms of FID for $C = C_V$. We emphasize that the imbalanced clusters of k -means are penalized when computing clustering metrics such as ANMI on balanced classification datasets, while during image synthesis, this is naturally mitigated by sampling from $q(c)$. Tab. 1 suggests that imbalanced cluster assignments are beneficial for unlabelled datasets such as FFHQ, where we report a relative gain of 4.8% using k -means clusters compared to TEMI.

C-EDM matches EDM when $C = C_{GT}$. In Fig. 3, we observe that the GT label conditioning closely follows the FID of TEMI clusters for $C = C_{GT}$. In Tab. 2, we leverage the two annotation levels of CIFAR100 [42], specifically the 20 GT superclasses and the 100 GT labels to benchmark how different grouping methods perform in image synthesis. Apart from image clustering, we create text-based pseudo-labels with CLIP (OpenCLIP ViT-G/14 [33]) similar to [47]. Then, CLIP pseudo-labels are derived from zero-shot classification [55] using the arg max of the text-image similarity after softmax [1]. For reference, we pro-




Figure 3. FID (left y-axis) and TEMI cluster utilization ratio r_C (right y-axis) across different numbers of clusters C (x-axis) using C-EDM, evaluated at $M_{img} = 100$. The green area indicates the discovered cluster range $[2, C_{max})$ for $r_C \leq \alpha = 0.96$.

	C		
	20	100	200
CIFAR100			
Human annotation (GT)	3.10	2.41	-
CLIP pseudo-labels	3.38	2.42	-
k -means clusters	3.09	2.41	2.36
TEMI clusters	2.93	2.37	2.31

Table 2. FID for different grouping methods as a condition to EDM ($M_{img} = 100$). We use the 20 GT superclasses and 100 GT labels of CIFAR100. CLIP pseudo-labels are computed using zero-shot classification pseudo-labels based on the GT label names.

vide the FID with the number of visual groups $C_V = 200$ for k -means and TEMI clusters. Interestingly, all methods attain a similar FID for $C_{GT} = 100$, while TEMI achieves the best FID only for $C = 20$. This suggests that generative performance is mostly invariant to the grouping method, including human annotation, given a feature extractor that captures general visual concepts. Next, we investigate the impact of the pre-trained feature extractor.

The number of visual groups is insensitive to the chosen feature extractor. The broader adoption of the discovered C_V for each dataset requires our analysis to be insensitive to the chosen feature extractor. In Fig. 4, we show that the pre-trained models considered in this work achieve similar FIDs across different numbers of clusters even though the ANMI varies up to 5.4% on CIFAR10. Feature extractors differ in terms of the number of parameters (86M up to 1.8 billion), pre-training dataset size (1.2M up to 2B), and objective (MoCO [18], DINO [14], iBOT [77], CLIP [33,55]). Our finding is consistent with concurrent work on the scale of ImageNet [45] that shows that linear probing accuracy is not indicative of generative performance for feature-based

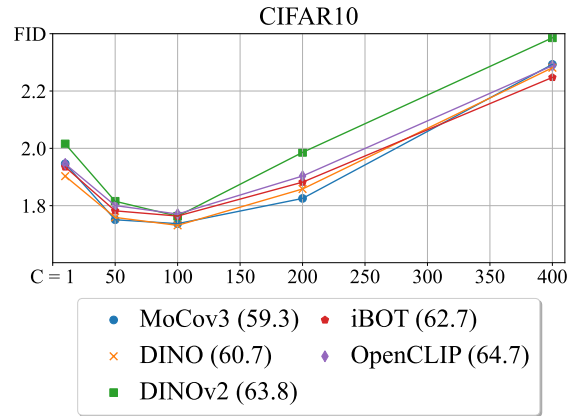


Figure 4. FID (y-axis) across different numbers of clusters C (x-axis) using C-EDM with TEMI with different feature extractors. The ANMI is shown in parentheses for $C_V=100$.

conditioning. The fact that differences in discriminative performance do not translate to improvements in FID is in contrast with Hu et al. [31].

Dataset-specific features do not improve generative performance. To realize dataset-dependent adaptation of the feature extractor, we perform self-supervised fine-tuning on CIFAR10 and CIFAR100, starting from the ImageNet weights in Tab. 3. We train all layers for 15 epochs similar to [58] using the DINO framework, resulting in “DINO adapted weights”. We find large improvements in classification accuracy and ANMI, which is in line with the fact that the learned features are more dataset-specific [56,58,72]. However, the gains in classification and clustering do not translate into conditional image generation, which suggests that image synthesis requires features that capture general visual concepts.

C-EDM+TEMI	CIFAR10 ($C = 100$)				CIFAR100 ($C = 200$)			
	ACC \uparrow	ANMI \uparrow	FID \downarrow	FDD \downarrow	ACC \uparrow	ANMI \uparrow	FID \downarrow	FDD \downarrow
DINO ImageNet weights	96.4	60.6	1.73	153.5	82.3	72.3	2.31	250.7
DINO adapted weights	98.5	63.41	1.74	152.8	87.8	78.2	2.39	250.8

Table 3. Effect of self-supervised fine-tuning (adapted weights) on the pre-trained DINO ViT-B. ACC stands for the 20-NN classification accuracy using the DINO features. We measure FID and FDD for both models at $M_{img} = 100$.




Figure 5. Top 1-NN cosine similarity AUROC (left y-axis) and Frchet distance between the C-EDM and unconditional samples (uFID) for different cluster sizes C (x-axis). For the computation of AUROC, we use the official test splits.

5. Discussion

Highly fine-grained clusters lead to out-of-distribution samples. To quantify the degree of out-of-distribution for the generated C-EDM samples, we measure the AUROC using the top-1 NN cosine similarity (1-NN) using DINO ViT-B [67] and compare with the test split of CIFAR10 and CIFAR100 (Fig. 5). We note that 1-NN is independent of the data distribution of the generated samples and thus does not account for diversity. The highest C considered for CIFAR10 and CIFAR100 produces the highest AUROC, suggesting that the features of the generated samples are pushed away from the training data. We hypothesize that this behavior originates from highly specialized clusters that cannot always be generated from the initial noise, resulting in visual artifacts. In parallel, we measure how similar the C-EDM samples are compared to the unconditional ones using FID. We call this metric *unconditional FID* (uFID). We observe that uFID increases as C increases, suggesting that the generated samples for larger C are farther away in feature space than the unconditional samples.

The clusters’ granularity level determines conditional generative performance. The quality of feature representations is typically determined by linear separability w.r.t. GT labels [51]. However, our experimental analysis shows that in image synthesis, the quality of image-level conditioning primarily depends on the granularity of the cluster assignments (Tab. 2 and Fig. 3). Overall, we find no

significant connection between the task of feature-based clustering and cluster-conditional image generation (Fig. 4 and Tab. 3). The pre-trained model and the categorization method do not severely affect generative performance for the considered datasets.

Generative metric and optimal number of clusters. Here, we investigate the impact of the choice of generative metric by comparing FDD and FID. We highlight that the only difference is that FID and FDD use InceptionV3 and DINOv2 features, respectively. The cross-dataset evaluation in Fig. 6 shows that FID and FDD do not always agree with respect to the number of visual groups. The largest disagreement is observed on CIFAR10. Still, the discovered upper bound C_{max} always includes the number of visual groups. Moreover, we evaluate both metrics across training iterations and found that while FID fluctuates after $M_{img} = 120$, FDD decreases monotonically (supplementary material). Finally, we measured additional feature-based metrics such as precision, recall, diversity, coverage, MSS, and inception score without success. Similar to [19, 34, 43, 52, 66], we argue that new generative metrics must be developed.

Visual comparison of C-EDM to EDM. In Fig. 7, we map clusters to their respective CIFAR100 classes and produce samples with C-EDM and GT conditional EDM using the same noise. Samples exhibit high visual similarity, corroborating with Tab. 2 and Fig. 2. More visualizations are provided in the supplementary material.




Figure 6. FID (left y-axis) and FDD (right y-axis) versus number of clusters (x-axis) using C-EDM at $M_{img} = 100$. The non-filled markers indicate the unconditional EDM.




Figure 7. Comparing C-EDM samples generated with learned cluster assignments (top row, $C=100$) to EDM samples generated with GT labels, on CIFAR100. The clusters are mapped to GT classes using the Hungarian algorithm.

Measuring the confidence of the generated samples using TEMI. In Fig. 8, we show the generated examples with the lowest and highest maximum softmax probability [24] of the TEMI head as a measure of confidence. For comparison, we show unconditional samples generated using the same initial noise in the denoising process. Visual inspection shows that low-confidence C-EDM samples do not have coherent semantics compared to the unconditional ones, leading to inferior image quality. We hypothesize that the sampled condition for the low-confidence samples is in conflict with the existing patterns in the initial noise. Increasing the number of conditions likely leads to worse image-condition alignment [27, 29, 61]. We argue that internal guidance methods could be employed to increase the image quality in such cases, which is left for future work.

By contrast, highly confident C-EDM samples show more clearly defined semantics than unconditional ones. When the low frequencies, such as the object’s shape, remain intact, cluster conditioning aids in refining local pixel patterns. Confident C-EDM samples consist of simple pixel

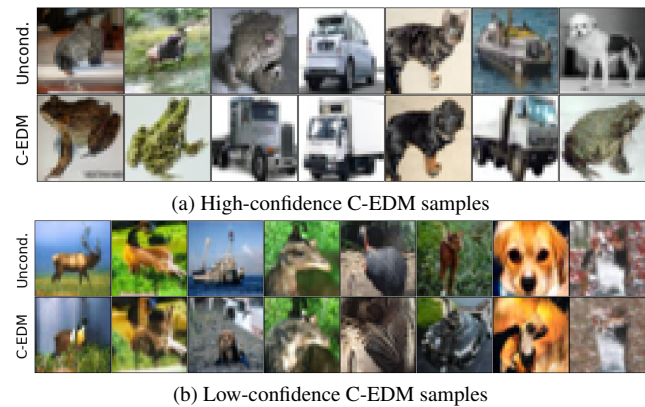


Figure 8. Generated high (a) and low (b) confidence CIFAR100 samples. The top row depicts the unconditional (*Uncond.*) samples, while the bottom row shows the generated samples using C-EDM with TEMI ($C = 100$). Images on the same column are produced with the same initial noise.

patterns that are easy to generate, such as a white background. Confidence can also be leveraged in future works in rejection sampling schemes [16]. More samples are provided in the supplementary material.

6. Conclusion

In this paper, a systematic empirical study was conducted focusing on conditioning diffusion models with cluster assignments. It was demonstrated that cluster conditioning achieves state-of-the-art FID on three generative benchmarks while attaining strong sample efficiency. To reduce the search space for estimating the visual groups of the dataset, a novel method that computes an upper cluster bound based solely on clustering was proposed. Finally, our experimental study indicates that generative performance using cluster assignments depends primarily on the granularity of the assignments.

References

- [1] Nikolas Adaloglou, Felix Michels, Tim Kaiser, and Markus Kollmann. Adapting contrastive language-image pretrained (clip) models for out-of-distribution detection. *arXiv e-prints*, pages arXiv-2303, 2023. [5](#)
- [2] Nikolas Adaloglou, Felix Michels, Hamza Kalisch, and Markus Kollmann. Exploring the limits of deep image clustering using pretrained models. In *34th British Machine Vision Conference 2023, BMVC 2023, Aberdeen, UK, November 20-24, 2023*. BMVA, 2023. [2](#), [3](#), [4](#)
- [3] Nikolas Adaloglou, Felix Michels, Kaspar Senft, Diana Petrusheva, and Markus Kollmann. Scaling up deep clustering methods beyond imagenet-1k. *arXiv preprint arXiv:2406.01203*, 2024. [1](#), [2](#)
- [4] Elad Amrani, Leonid Karlinsky, and Alex Bronstein. Self-supervised classification network. In *European Conference on Computer Vision*, pages 116–132. Springer, 2022. [2](#)
- [5] Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 843–852, 2023. [2](#)
- [6] Björn Barz and Joachim Denzler. Do we train on test data? purging cifar of near-duplicates. *Journal of Imaging*, 6(6):41, 2020. [1](#)
- [7] Georgios Batzolis, Jan Stanczuk, Carola-Bibiane Schönlieb, and Christian Etmann. Conditional image generation with score-based diffusion models. *arXiv preprint arXiv:2111.13606*, 2021. [1](#)
- [8] James Betker, Gabriel Goh, Li Jing, TimBrooks, Jianfeng Wang, Linjie Li, LongOuyang, JuntangZhuang, JoyceLee, YufeiGuo, WesamManassra, PrafullaDhariwal, CaseyChu, YunxinJiao, and Aditya Ramesh. Improving image generation with better captions. [1](#)
- [9] Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xi-aohua Zhai, and Aäron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020. [1](#)
- [10] Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xi-aohua Zhai, and Aäron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020. [1](#)
- [11] Florian Bordes, Randall Balestriero, and Pascal Vincent. High fidelity visualization of what your self-supervised representation knows about, 2022. [2](#)
- [12] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis, 2019. [2](#)
- [13] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 132–149, 2018. [1](#)
- [14] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. [3](#), [5](#), [6](#)
- [15] Arantxa Casanova, Marlene Careil, Jakob Verbeek, Michal Drozdal, and Adriana Romero Soriano. Instance-conditioned gan. *Advances in Neural Information Processing Systems*, 34:27517–27529, 2021. [2](#)
- [16] George Casella, Christian P Robert, and Martin T Wells. Generalized accept-reject sampling schemes. *Lecture Notes-Monograph Series*, pages 342–347, 2004. [8](#)
- [17] Jianlong Chang, Lingfeng Wang, Gaofeng Meng, Shiming Xiang, and Chunhong Pan. Deep adaptive image clustering. In *Proceedings of the IEEE international conference on computer vision*, pages 5879–5887, 2017. [1](#)
- [18] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9640–9649, 2021. [6](#)
- [19] Min Jin Chong and David Forsyth. Effectively unbiased fid and inception score and where to find them. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6070–6079, 2020. [7](#)
- [20] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [4](#)
- [21] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. [1](#)
- [22] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. [2](#)
- [23] Bao et al. Why are conditional generative models better than unconditional ones? In *NeurIPS Workshop on Score-Based Methods*, 2022. [1](#), [2](#), [3](#), [4](#)
- [24] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations*, 2017. [8](#)
- [25] 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. [3](#), [4](#)
- [26] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. [1](#), [2](#), [3](#), [4](#)
- [27] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. [2](#), [4](#), [8](#)
- [28] Laura Hollink, Aysenur Bilgin, and Jacco Van Ossenbruggen. Is it a fruit, an apple or a granny smith? predicting the basic level in a concept hierarchy. *arXiv preprint arXiv:1910.12619*, 2019. [1](#)
- [29] Susung Hong, Gyuseong Lee, Wooseok Jang, and Seungryong Kim. Improving sample quality of diffusion models using self-attention guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7462–7471, October 2023. [2](#), [8](#)
- [30] Vincent Tao Hu, Yunlu Chen, Mathilde Caron, Yuki M Asano, Cees GM Snoek, and Bjorn Ommer. Guided diffu-

- sion from self-supervised diffusion features. *arXiv preprint arXiv:2312.08825*, 2023. [2](#)
- [31] Vincent Tao Hu, David W Zhang, Yuki M Asano, Gertjan J Burghouts, and Cees GM Snoek. Self-guided diffusion models. In *CVPR*, 2023. [2](#), [3](#), [4](#), [5](#), [6](#)
- [32] Zhizhong Huang, Jie Chen, Junping Zhang, and Hongming Shan. Learning representation for clustering via prototype scattering and positive sampling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [3](#), [4](#)
- [33] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, July 2021. If you use this software, please cite it as below. [5](#), [6](#)
- [34] Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and Sanjiv Kumar. Rethinking fid: Towards a better evaluation metric for image generation. *arXiv preprint arXiv:2401.09603*, 2023. [7](#)
- [35] Xu Ji, Joao F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9865–9874, 2019. [1](#), [3](#)
- [36] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022. [1](#), [2](#), [3](#), [4](#)
- [37] Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models, 2023. [2](#), [5](#)
- [38] Amirhossein Kazerooni, Ehsan Khodapanah Aghdam, Moein Heidari, Reza Azad, Mohsen Fayyaz, Ilker Hacıhaliloglu, and Dorit Merhof. Diffusion models for medical image analysis: A comprehensive survey. *arXiv preprint arXiv:2211.07804*, 2022. [1](#)
- [39] David J Ketchen and Christopher L Shook. The application of cluster analysis in strategic management research: an analysis and critique. *Strategic management journal*, 17(6):441–458, 1996. [4](#)
- [40] Dongjun Kim, Yeongmin Kim, Se Jung Kwon, Wanmo Kang, and Il-Chul Moon. Refining generative process with discriminator guidance in score-based diffusion models. *arXiv preprint arXiv:2211.17091*, 2022. [4](#)
- [41] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021. [2](#)
- [42] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [1](#), [4](#), [5](#)
- [43] Tuomas Kynkäänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. The role of imagenet classes in fr \backslash echet inception distance. *arXiv preprint arXiv:2203.06026*, 2022. [7](#)
- [44] Junnan Li, Pan Zhou, Caiming Xiong, and Steven CH Hoi. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*, 2020. [3](#)
- [45] Tianhong Li, Dina Katabi, and Kaiming He. Self-conditioned image generation via generating representations. *arXiv preprint*, 2023. [2](#), [6](#)
- [46] Steven Liu, Tongzhou Wang, David Bau, Jun-Yan Zhu, and Antonio Torralba. Diverse image generation via self-conditioned gans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14286–14295, 2020. [2](#), [3](#)
- [47] Yifei Ming, Ziyang Cai, Jiuxiang Gu, Yiyu Sun, Wei Li, and Yixuan Li. Delving into out-of-distribution detection with vision-language representations. *Advances in Neural Information Processing Systems*, 35:35087–35102, 2022. [5](#)
- [48] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. [2](#)
- [49] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. [1](#), [2](#)
- [50] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. [2](#), [4](#)
- [51] Namuk Park, Wonjae Kim, Byeongho Heo, Taekyung Kim, and Sangdoon Yun. What do self-supervised vision transformers learn? In *The Eleventh International Conference on Learning Representations*, 2023. [7](#)
- [52] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11410–11420, 2022. [7](#)
- [53] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4195–4205, October 2023. [1](#)
- [54] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. [2](#)
- [55] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [1](#), [5](#), [6](#)
- [56] Nima Rafiee, Rahil Gholamipoor, Nikolas Adaloglou, Simon Jaxy, Julius Ramakers, and Markus Kollmann. Self-supervised anomaly detection by self-distillation and negative sampling. In *International Conference on Artificial Neural Networks*, pages 459–470. Springer, 2022. [6](#)
- [57] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. [2](#)
- [58] Tal Reiss, Niv Cohen, Eliahu Horwitz, Ron Abutbul, and Yedid Hoshen. Anomaly detection requires better represen-

- tations. In *European Conference on Computer Vision*, pages 56–68. Springer, 2022. 6
- [59] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. 2022 ieee. In *CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, 2021. 2
- [60] Seyedmorteza Sadat, Jakob Buhmann, Derek Bradely, Otmar Hilliges, and Romann M Weber. Cads: Unleashing the diversity of diffusion models through condition-annealed sampling. *arXiv preprint arXiv:2310.17347*, 2023. 4
- [61] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 1, 8
- [62] Shelly Sheynin, Oron Ashual, Adam Polyak, Uriel Singer, Oran Gafni, Eliya Nachmani, and Yaniv Taigman. Knn-diffusion: Image generation via large-scale retrieval, 2022. 1
- [63] Shelly Sheynin, Oron Ashual, Adam Polyak, Uriel Singer, Oran Gafni, Eliya Nachmani, and Yaniv Taigman. Knn-diffusion: Image generation via large-scale retrieval. *arXiv preprint arXiv:2204.02849*, 2022. 2
- [64] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015. 2
- [65] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
- [66] George Stein, Jesse C. Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Leigh Ross, Valentin Vilecroze, Zhaoyan Liu, Anthony L. Caterini, Eric Taylor, and Gabriel Loaiza-Ganem. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 2, 4, 7
- [67] Yiyu Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. Out-of-distribution detection with deep nearest neighbors. In *International Conference on Machine Learning*, pages 20827–20840. PMLR, 2022. 7
- [68] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 4
- [69] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 4
- [70] Robert L Thorndike. Who belongs in the family? *Psychometrika*, 18(4):267–276, 1953. 4
- [71] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. Scan: Learning to classify images without labels. In *European conference on computer vision*, pages 268–285. Springer, 2020. 2, 3
- [72] Hualiang Wang, Yi Li, Huifeng Yao, and Xiaomeng Li. Clipn for zero-shot ood detection: Teaching clip to say no. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1802–1812, 2023. 6
- [73] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*, pages 478–487. PMLR, 2016. 2
- [74] Yilun Xu, Ziming Liu, Yonglong Tian, Shangyuan Tong, Max Tegmark, and Tommi Jaakkola. Pfgm++: Unlocking the potential of physics-inspired generative models. *arXiv preprint arXiv:2302.04265*, 2023. 4
- [75] Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and Fang Wen. Paint by example: Exemplar-based image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18381–18391, 2023. 1
- [76] Asano YM., Rupprecht C., and Vedaldi A. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations*, 2020. 2
- [77] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021. 6
- [78] Xingzhi Zhou and Nevin L Zhang. Deep clustering with features from self-supervised pretraining. *arXiv preprint arXiv:2207.13364*, 2022. 4
- [79] Yufan Zhou, Ruiyi Zhang, Changyou Chen, Chunyuan Li, Chris Tensmeyer, Tong Yu, Jiuxiang Gu, Jinhui Xu, and Tong Sun. Towards language-free training for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17907–17917, 2022. 2