



Advances in diffusion models for image data augmentation: a review of methods, models, evaluation metrics and future research directions

Panagiotis Alimisis¹ · Ioannis Mademlis¹ · Panagiotis Radoglou-Grammatikis^{2,3} · Panagiotis Sarigiannidis² · Georgios Th. Papadopoulos¹

Accepted: 12 January 2025
© The Author(s) 2025

Abstract

Image data augmentation constitutes a critical methodology in modern computer vision tasks, since it can facilitate towards enhancing the diversity and quality of training datasets; thereby, improving the performance and robustness of machine learning models in downstream tasks. In parallel, augmentation approaches can also be used for editing/modifying a given image in a context- and semantics-aware way. Diffusion Models (DMs), which comprise one of the most recent and highly promising classes of methods in the field of generative Artificial Intelligence (AI), have emerged as a powerful tool for image data augmentation, capable of generating realistic and diverse images by learning the underlying data distribution. The current study realizes a systematic, comprehensive and in-depth review of DM-based approaches for image augmentation, covering a wide range of strategies, tasks and applications. In particular, a comprehensive analysis of the fundamental principles, model architectures and training strategies of DMs is initially performed. Subsequently, a taxonomy of the relevant image augmentation methods is introduced, focusing on techniques regarding semantic manipulation, personalization and adaptation, and application-specific augmentation tasks. Then, performance assessment methodologies and respective evaluation metrics are analyzed. Finally, current challenges and future research directions in the field are discussed.

Keywords Image data augmentation · Diffusion models · Generative artificial intelligence · Evaluation metrics

1 Introduction

Modern computer vision has been dominated by the so-called Deep Learning (DL) paradigm, which relies on the use of large-scale on Deep Neural Networks (DNNs) (Chai et al. 2021). DNNs have so far exhibited outstanding performance in a wide set of visual understanding tasks. However, this eminent visual interpretation and reasoning capability

Extended author information available on the last page of the article

is accompanied by the increased need for ever larger and sufficiently diverse training datasets (Shrestha and Mahmood 2019). On the other hand, as image analysis tasks become increasingly intricate and demanding, the ability of DNNs to generalize robustly is hindered by limitations in training data quantity, diversity, and potential bias (Mumuni and Mumuni 2022). As a result, data requirements has emerged as a rather prominent topic, since a sufficient volume of training samples is essential for fully harnessing the capabilities of DNNs (Zhang et al. 2022). On the contrary, real-world image datasets, especially regarding specific-targeted application domains, often suffer in these aspects, even to the point of containing perfectly correlated training images that are proven to be essentially redundant (Yang et al. 2022).

Image augmentation constitutes a common preprocessing step in machine learning workflows, designed to increase the visual diversity of the training dataset, without introducing additional independent images. The goal is to enhance the learning process, while operating within a fixed set of data (Shorten and Khoshgoftaar 2019). Augmentation addresses limitations in dataset size, by automatically creating additional variants of existing training images. These variants are generated by directly modifying the original images, ensuring that the transformed images differ in appearance, but retain their semantic content (Xu et al. 2023a). In contrast, general image generation (or image synthesis) involves sampling from a model that approximates the overall data distribution of the dataset, rather than transforming individual training images. This approach generates entirely new images that are not directly derived from specific originals, but they are instead representative of the underlying characteristics of the whole dataset (Cao et al. 2024). Extending the dataset with such augmented images, either synthesized or modified, increases its diversity and improves, in many cases, the downstream learning and recognition performance of DNNs that are being trained using it Zhou et al. (2023b). This behaviour stems from image augmentation essentially acting as an additional regularizing mechanism while training the DNN and, as a result, helping to prevent overfitting (Perez and Wang 2017; Shorten and Khoshgoftaar 2019).

Traditional approaches for image augmentation, such as geometric transformations (e.g., image rotation, flip, crop, scaling, horizontal/vertical translation, squeezing, etc.) and color space adjustments or photometric transforms (e.g., blurring, sharpening, jittering, etc.) are still very common (Xu et al. 2023; Shorten and Khoshgoftaar 2019; Yang et al. 2022). Multiple transformations of this type can be composed together, so that an even wider set of augmented images can be generated from the original dataset. These methods leverage domain knowledge to produce synthetic examples similar to the initial ones. More recently proposed image augmentation methods in this general vein are a set of strategies for systematically corrupting the original images, in order to generate augmented variants. This category of methods includes, among others: a) ‘Mixup’ (Zhang et al. 2017), which uses convex combinations of pairs of training images and their labels, b) ‘Cutout’ (DeVries and Taylor 2017), which randomly masks square regions of an input image, c) ‘Cutmix’ (Yun et al. 2019), which randomly combines two training images by masking the first with a region of the second (and vice versa), d) ‘Patchshuffle’ (Kang et al. 2017), which uses a kernel filter to randomly swap the pixel values in a sliding window, e) ‘Copy-Paste’ (Dwivedi et al. 2017), which pastes segmented object instances onto random background images, f) ‘Co-Mixup’ (Kim et al. 2021), which optimizes saliency-guided mixing of multiple inputs while encouraging diversity among outputs, g) ‘RandAugment’ (Cubuk et al. 2020), which randomly selects N transformations from a predefined set, each applied with magnitude M ,

eliminating the need for separate policy search, h) ‘GuidedMixup’ (Kang and Kim 2023), which uses saliency maps to guide image mixing while preserving salient regions through pixel-wise mixing ratios, and i) ‘CAL-AUG’ (Rao et al. 2021), which randomly replaces learned attention maps with counterfactual ones.

The effectiveness though of the above-mentioned relatively simple and straightforward augmentation methods is being increasingly challenged by the complexity and variability of contemporary image analysis demands. Although such strategies can be effective in increasing data diversity for simple tasks, they are mostly unable to capture the underlying structure and complex relationships present in high-dimensional image data. Additionally, many of them require domain-specific knowledge and dataset-specific calibration, in order to be applied correctly (Wu et al. 2023d). Moreover, the needs of DNNs for large training datasets and effective regularization are ever-growing, rendering image augmentation a critically important component of modern machine learning (Zhang et al. 2022).

Unlike traditional methods, which directly manipulate existing images to generate variants, Diffusion Models (DMs) can be exploited for image augmentation, by learning to synthesize new, realistically-looking and plausible images, given a training dataset (Zhang et al. 2023f; Xu et al. 2024; Asperti et al. 2023). These are general image generation algorithms, aiming to learn a model of the underlying distribution of the selected training dataset and then allow its sampling, in order to synthesize novel images. These generated images do not have a direct one-to-one correspondence with the original training ones, but instead reflect the dataset’s underlying statistical patterns. Generative methods can be used for various tasks, such as procedural content generation and simulation, but they can also implement dataset expansion, as a specific type of image augmentation at the dataset level. Additionally, when coupled with a conditioning mechanism, they can replace traditional augmentation approaches and synthesize modified variants of specific given images, through conditional generation that allows one-to-one correspondence between an original training image and a synthesized variant. DMs are a sophisticated class of generative DNNs that excel in implicitly modeling the underlying data generating distribution and the structure of complex images. This capability allows them to essentially sample fake novel images from their training dataset’s distribution, which are simultaneously diverse, highly realistic and representative of unseen data scenarios, as they encompass subtle details and preserve the inherent structure of the original dataset (Zhang et al. 2022; Trabucco et al. 2023). Thus, they can be utilized for meaningfully augmenting the latter.

The learning paradigm of DMs, which relies on iteratively applying noise to the training images and subsequently learning to reverse the process, has shown significant promise in image augmentation, when compared against competing generative models (e.g., Generative Adversarial Networks) (Ho et al. 2020). Additionally, recent advancements in DMs enable the conditioning of the image synthesis process via class labels, textual descriptions, or input images (Rombach et al. 2022). This level of user control allows for targeted image augmentation, generating images that fulfill specific requirements based on the task at hand.

The recent advancement in generative image synthesis through DMs and multimodal strategies (e.g., text-conditioned image creation) has been complemented by the use of large-scale pretraining on massive datasets, in the vein of the Foundation Model (FM) trend (Rombach et al. 2022; Podell et al. 2023; Esser et al. 2024; Saharia et al. 2022b). This approach has led to the availability of pretrained DMs that can generate images with natural variations in appearance (e.g., changing the design of the graffiti on a truck, as illustrated in

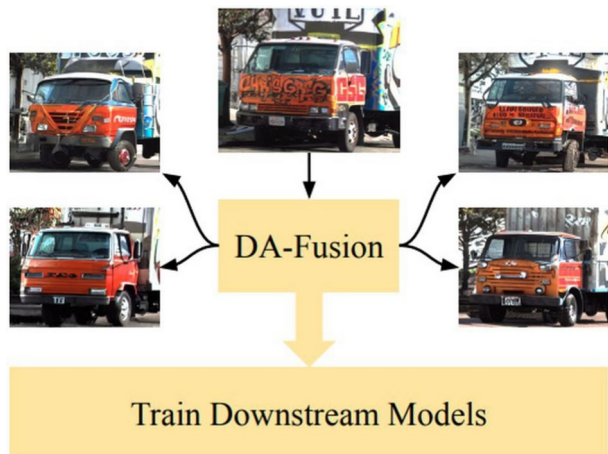
Fig. 1) and, hence, can be directly exploited for sophisticated image augmentation without significant human effort (Dunlap et al. 2023).

Despite the increasing importance of DMs in image augmentation, there is a *gap* in the relevant literature, since no existing review focuses specifically on this application and its particularities. Published surveys on DMs do not focus on image augmentation in a targeted manner, given the recent progress and achievements in the field. Existing surveys either focus on traditional image augmentation (Yang et al. 2022), or provide a general overview of DMs without delving into their specific application for image augmentation (Cao et al. 2024; Yang et al. 2023b; Song et al. 2024). For instance, the study of Mumuni and Mumuni (2022) presents a comprehensive taxonomy of image augmentation approaches, including input space transformations, feature space augmentation, data synthesis and meta-learning based methods. However, it does not cover the latest advancements concerning the use of DMs. In contrast, the work of Croitoru et al. (2023) presents three generic diffusion modeling frameworks, which are based on denoising diffusion probabilistic models, noise conditioned score networks, and stochastic differential equations, but does not explore their potential for image augmentation. Some surveys have touched upon the efficiency aspect of DMs (Ulhaq et al. 2022) or their application in specific domains like medical imaging (Kazerouni et al. 2023; Kebaili et al. 2023). However, these works do not provide a comprehensive overview of DMs for augmentation across various computer vision tasks. Furthermore, a recent survey categorizes augmentation methods based on large learning models, including those based on DMs, but does not focus specifically on them (Zhou et al. 2024).

The above-described situation has *motivated* this study, in an attempt to remedy the identified gap. Its main *contribution* is a systematic, comprehensive and in-depth review of DM-based approaches specifically for image augmentation, covering a wide range of strategies, tasks and applications. In particular, the current work analyzes fundamental principles, introduces a detailed taxonomy of methods, examines evaluation metrics, discusses current challenges and provides perspectives on future research directions.

The remainder of this article is structured in the following way. Section 2 outlines the foundations and underlying principles of DMs. Section 3 describes the different categories of DM-powered image augmentation methods, while Sect. 4 presents and explains these categories in depth. Section 5 surveys the various evaluation metrics used to assess the

Fig. 1 Semantic alteration of truck graffiti design using diffusion models. Image from Trabucco et al. (2023)



performance of DMs for image augmentation. Section 6 discusses the current challenges and limitations associated with the use of DMs for image augmentation. Finally, Sect. 7 draws insights from the preceding discussion, along with suggestions for future research directions.

2 Foundations of diffusion models

Diffusion Models (DMs) are a powerful class of generative models gaining significant traction in image synthesis. Inspired by non-equilibrium thermodynamics (Sohl-Dickstein et al. 2015), they operate by incrementally destroying structure in the data, through an iterative process of adding Gaussian noise (forward diffusion) that progressively transforms the data distribution towards a distribution of pure random noise. Then, a learnable reverse diffusion process that restores structure in the data yields a tractable generative model. Thus, DMs are trained for gradually transforming random noise patterns into samples of the data generating distribution. This section details the principles underlying DMs, clarifying why they are particularly suited for visual data augmentation.

2.1 Forward diffusion process

The Forward Diffusion (FD) process (Ho et al. 2020) is the cornerstone of DMs, as it corrupts the training dataset by sequentially inserting Gaussian noise. Assume that the initial data distribution is $q(x_0)$, where subscript '0' denotes the original/unmodified state of the dataset and $x_0 \sim q(x_0)$ is an image from this dataset. FD proceeds as a sequence q of incrementally noised versions, x_1, x_2, \dots, x_T , which are generated by a Markov chain. The conditional distribution for each step in this sequence, $p(x_t|x_{t-1})$, is modeled as a Gaussian $\mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I})$, where t ranges from 1 to T and it denotes the total noise added to the input image. T corresponds to the total number of diffusion steps, β_1, \dots, β_T is a sequence of variance parameters that define the noise level at each step, \mathbf{I} is the identity matrix matching the dimensionality of input x_0 and $\mathcal{N}(x; \mu, \sigma)$ denotes the normal distribution with mean μ and covariance σ .

A key attribute of FD is that x_t can be sampled at any arbitrary time step t in closed form, using a reparameterization trick:

$$\text{Let } a_t = 1 - \beta_t, \bar{a}_t = \prod_{i=1}^t a_i \quad (1)$$

$$\text{Then } q(x_t|x_0) = \mathcal{N}(x_t|\sqrt{\bar{a}_t}x_0, (1 - \bar{a}_t)\mathbf{I})$$

$$x_t = \sqrt{\bar{a}_t}x_0 + \sqrt{1 - \bar{a}_t}\epsilon,$$

where integer $t \in [1, N]$ and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. Thus, the noisy version x_t can be directly obtained through a cumulative variance adjustment β_t , determined by sequence a_i , where $a_t = 1 - \beta_t$. This allows one to compute any noisy version x_t from the original image x_0 in a single step, without having to iteratively generate the noisy version of all intervening time steps.

2.2 Reverse diffusion process

Following the corruption introduced by FD, the iterative Reverse Diffusion (RD) process aims at recovering the original dataset images from their noisy versions. Instead of directly generating images from noise patterns, a denoising learning model, which can be a DNN, iteratively predicts the noise pattern added to the data at each individual step of the FD process, starting from the final FD output, so that it can be removed. Progressive denoising gradually refines the image across T consecutive steps. This is the so-called Denoising Diffusion Probabilistic Model (DDPM) formulation (Ho et al. 2020). Alternatively, the model can learn the so-called ‘score function’, which is the gradient of the log probability density function of the data with respect to the input. Then, the model’s predictions at each time step can be used to iteratively sample from the distribution, by following the gradient. Such a DM variant is called Score-based Generative Model (SGM) (Song et al. 2020b).

The employed predictive DNN is usually a U-Net CNN (Ronneberger et al. 2015). Regarding the mathematical formulation, the RD process is defined as follows:

$$p_{\theta}(x_{0:T}) = p(x_T) \prod_{t=1}^T p_{\theta}(x_{t-1}|x_t), \quad (2)$$

where $p_{\theta}(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}, \mu_{\theta}(x_t, t), \Sigma(x_t, t))$.

In Ho et al. (2020), the U-Net is trained by the following loss function:

$$L_{simple} = \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_{\theta}(x_t, t)\|_2^2]. \quad (3)$$

where ϵ represents the Gaussian noise added to image x_0 to obtain the noisy version x_t and $\epsilon_{\theta}(x_t, t)$ denotes the noise predicted by the DNN parameterized by θ , given the noisy image x_t and the time step t . In the SGM formulation, $\epsilon_{\theta}(x_t, t)$ is the predicted score and, thus, after training, $\mu_{\theta}(x_t, t)$ can be approximated by a function of $\epsilon_{\theta}(x_t, t)$. Even though L_{simple} (3) does not offer a way to learn $\Sigma_{\theta}(x_t, t)$, it has been shown in Ho et al. (2020) that the best results are obtained by fixing the variance to $\sigma_t^2 \mathbf{I}$, rather than learning it.

RD is an iterative process of T consecutive time steps, starting from noise pattern x_T and gradually recovering the original image x_0 . At each time step t , μ and Σ are computed and a new version of the output image is generated, which subsequently serves as input for the next time step t . This need for sequential generation across T consecutive iterations is a significant limitation of DMs. One simple improvement is to reduce the number of sampling steps, from T to K evenly spaced real numbers between 1 and T (Nichol and Dhariwal 2021). Alternatively, the non-Markovian Denoising Diffusion Implicit Models (DDIMs) (Song et al. 2020a) sample only across S diffusion steps $[t_1, \dots, t_S] \subseteq [1, T]$ during generation:

$$x_{t-1} = \sqrt{\bar{a}_{t-1}} \left(\frac{x_t - \sqrt{1 - a_t} \epsilon_{\theta}(x_t)}{\sqrt{a_t}} \right) + \sqrt{1 - a_{t-1} - \sigma_t^2} \epsilon_{\theta}(x_t) + \sigma_t \epsilon_t. \quad (4)$$

2.3 Guidance

Classifier Guidance (Dhariwal and Nichol 2021) leverages a pretrained closed-set classifier to condition the RD process of a pretrained unconditional DM on a desired class label. The classifier model $p_\phi(y|x_t)$, where ϕ denotes its parameters, supports as many different class labels y as the potential conditioning classes. With this approach and given the SGM formulation, the RD process is adjusted at each time step by the gradient of the log-probability $\nabla_{x_t} \log p_\phi(y|x_t)$ that steers sampling. Thus, $\mu_\theta(x_t, t)$ is approximated by:

$$\epsilon_\theta(x_t, t) + s \times \nabla_{x_t} \log p_\phi(y|x_t), \quad (5)$$

where s is a scaling factor controlling the strength of guidance. This method ensures that the generated samples conform to the target class distribution, without any need to retrain the unconditionally trained DNN.

Classifier-Free Guidance (Ho and Salimans 2022) eliminates the need for an explicit separate classifier model, by conditioning the DM on class labels directly during its training. The DM is trained with both conditional θ_c and unconditional θ_u objectives, alternating between conditioning on labels and generating without labels. At inference time, guidance is implemented by interpolating between the conditional and unconditional scores, so that $\mu_\theta(x_t, t)$ is approximated as follows:

$$\nabla_{x_t} \log p_\theta(x_t|y) = \nabla_{x_t} \log p_{\theta_c}(x_t|y) + w(\nabla_{x_t} \log p_{\theta_c}(x_t|y) - \nabla_{x_t} \log p_{\theta_u}(x_t)), \quad (6)$$

where w is a weight parameter that controls the strength of the guidance.

2.4 Diffusion models in latent space

Despite the faster RD process of DDIMs, image generation in pixel space and in an arbitrary resolution remains a significant bottleneck. To this end, Latent Diffusion Models (LDMs) (Rombach et al. 2022) have been introduced that operate in a latent space, in order to significantly accelerate the generation process. In particular, an LDM relies on an external autoencoder pretrained on a large-scale dataset. Its encoder \mathcal{E} learns to map images $x \in D_x$ into a special latent code $z = \mathcal{E}(x)$ (Van Den Oord et al. 2017; Agustsson et al. 2017). Its decoder \mathcal{D} learns to map such low-dimensional latent representations back to pixel space, so that $\mathcal{D}(\mathcal{E}(x)) \approx x$. Thus, a regular DM or DDIM is trained to generate codes within the latent space. The resulting code can be mapped back to a realistic, high-dimensional image via the pretrained \mathcal{D} .

The LDM can be conditioned on class labels, segmentation masks, or even text, which guide the generation process. Let $c_\theta(y)$ be a model that maps a raw conditioning input y to a conditioning vector.¹ The LDM loss is then formulated as:

$$L_{LDM} = \mathbb{E}_{z \in \mathcal{E}(x), y, \epsilon \in \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_\theta(z_t, t, c_\theta(y))\|_2^2], \quad (7)$$

where t is the time step, z_t is the latent representation noised at step t , ϵ is the unscaled noise sample, and ϵ_θ is the denoising network's prediction. Intuitively, the objective is to correctly

¹In the case of text prompts, this can be any text encoder.

remove the noise added to a latent representation of an image. During training, c_θ and ϵ_θ are jointly optimized to minimize the LDM loss. At inference time, a random noise tensor is sampled and iteratively denoised to produce a new latent image z_0 .

2.5 Types of conditioning

DMs can be conditioned on various types of information to guide the generation process. The most common types of conditioning include:

- **Text Conditioning:** Models like Stable Diffusion (Rombach et al. 2022) use text prompts to guide image generation. The text is typically encoded using CLIP (Radford et al. 2021) or T5 (Raffel et al. 2020) encoders, and the resulting embeddings condition the denoising process through cross-attention mechanisms.
- **Class Label Conditioning:** Models can be conditioned on discrete class labels (Dhariwal and Nichol 2021), enabling class-specific image generation. This is often implemented through class embeddings that are concatenated with or added to the model's intermediate features.
- **Image Conditioning:** Reference images can guide the generation process through various mechanisms: direct concatenation with model features (Saharia et al. 2022a), CLIP image embeddings (Avrahami et al. 2023b), or learned image encoders (Rombach et al. 2022). This enables tasks like image-to-image translation and style transfer.
- **Segmentation Map Conditioning:** Incorporating segmentation maps allows for control over the spatial layout and the structure of the generated images. Methods like (Wang et al. 2022a) incorporate segmentation maps through specialized encoders of feature injection to control the spatial distribution of semantic classes. This enables precise control over object placement and scene composition.
- **Pose and Structure Conditioning:** Human pose maps (Zhang et al. 2023b), edge maps, or semantic segmentation masks can guide the spatial structure of the generated images. These are typically processed through specialized encoders and injected into the model via cross-attention or feature concatenation.
- **Multi-Modal Conditioning:** Certain models combine multiple mechanisms simultaneously. For example, combining text prompts with segmentation masks (Avrahami et al. 2023b) or reference images with class labels (Ruiz et al. 2023) for more precise control over the generation process. The most suitable type of conditioning depends on the specific application and the desired degree of control over the generation procedure. Recent work has focused on developing more sophisticated approaches to conditioning that enable finer-grained control, while maintaining generation quality (Zhang et al. 2023b; Zhao et al. 2024).

2.6 Diffusion transformer

The Diffusion Transformer (DiT) (Peebles and Xie 2023) architecture represents a significant departure from U-Net-based designs, replacing convolutional neural layers with Transformer blocks for processing visual data. DiT maintains the same high-level structure as previous DMs, but processes images as sequences of spatial patches in the vein of Vision Transformers (ViT) (Dosovitskiy et al. 2020). The architecture consists of: a) A patchifica-

tion layer that splits images into non-overlapping patches and linearly projects them, b) A sequence of Transformer blocks with self-attention mechanisms, c) Time step and condition embeddings that are added to the patch embeddings, and d) A final decoder layer that reconstructs the spatial input image. The key advantages of DiT over U-Net-LDM include better scaling properties, improved handling of long-range dependencies through self-attention, and more flexible integration of conditioning information (Hatamizadeh et al. 2025; Mo et al. 2023). However, DiT typically requires more computational resources for training, due to the quadratic complexity of the self-attention operations. Recent work has shown that DiT-based models can achieve superior generation quality when scaled to sufficient size, leading to their wide adoption (Peebles and Xie 2023; Fei et al. 2024).

2.7 Foundation diffusion models for image generation

The evolution of DMs for image generation has been marked by several key architectural innovations and training strategies. These developments have resulted in the proliferation of commonly used Foundation Models (FMs) for image synthesis, which are available in a pretrained form and can be employed as a basis for more specialized methods. The architectures can be primarily categorized into *U-Net-based* (Rombach et al. 2022) and *Transformer-based* (Peebles and Xie 2023) ones, with *U-Net-based* methods further divided into those operating directly in *pixel space* versus those performing inference in *latent spaces*. All methods operate in the latent space of a pretrained autoencoder, unless those that perform directly in the pixel space. Moreover, U-Net-based methods employ a conditional LDM as backbone (or DDPM (Ho et al. 2020), if they operate in pixel space), while Transformer-based approaches utilize a conditional DiT as their backbone.

Table 1 illustrates key Foundation Diffusion Models (FDMs) presented in the literature, emphasizing on the following important characteristics: a) Release date, b) License, c) Conditioning mechanisms, d) Text encoder, e) Training dataset, f) Number of parameters, g) Resolution (in pixels), h) Notable features, and i) Indicative/important methods relying on them. As can be seen from Table 1, most FDMs operate in latent space and use text/image conditioning, while U-Net is the dominant neural architecture before the recent emergence of Transformer-based approaches. LAION-5B (Schuhmann et al. 2022) and its variants serve as the predominant training dataset for many open-source FDMs, highlighting its significance in image generation. There is also an evolution in text encoders from CLIP-based models towards more sophisticated language models like T5-XXL, suggesting a trend toward better text understanding.

3 Taxonomy of diffusion models for image augmentation

In this section, an overview of the landscape of DM-based methods for image augmentation is provided. In particular, a taxonomy of the various approaches is defined and graphically illustrated in Fig. 2. More specifically, considering as a criterion the task/goal of each method, DM-powered image augmentation approaches can initially be classified in the following main categories (while each class can be further divided into sub-categories, as will be discussed later in this section):

- **Semantic Manipulation:** The goal is to introduce fine-grained context-aware modifications to an image, while maintaining its main semantic contents (Kawar et al. 2023; Kim et al. 2022; Zhang et al. 2023b).
- **Personalization and Adaptation:** The target is to alter the appearance of the image, so as to better conform to specific datasets, tasks, requirements or user preferences (Ruiz et al. 2023; Gal et al. 2022; Wei et al. 2023).
- **Application-Specific Augmentation:** The goal is to regulate the augmentation process using domain specific knowledge, i.e. introducing modifications that are only meaningful for a given application (e.g., medical imaging, facial recognition, etc.) (Chambon et al. 2022a; Boutros et al. 2023). Complementarily to the categorization described above, Fig. 3 illustrates a timeline representation that contains key recent DM-powered image augmentation methods. Each entry corresponds to a critical milestone work that significantly impacted the research field. Naturally, more recent works are associated with more complex and advanced DM models/architectures, leading also to superior performance. Moreover, Table 2 illustrates key representative methods, belonging to each category defined in the taxonomy of Fig. 2.

3.1 Semantic manipulation

Methods belonging to this category aim to induce subtle and context-aware changes to an image, altering its interpretation or conveying additional contextual information, while preserving the core semantic content and maintaining a coherent and meaningful visual representation (Saharia et al. 2022a; Hertz et al. 2022). Such methods are useful for generating realistic and diverse training samples, while maintaining the original image context. More fined-grained and detailed sub-classes of this category are:

- **Concept Manipulation:** Concept manipulation involves altering the semantic content of an image, by adding, removing or modifying objects, attributes, or even the entire scene (Chen et al. 2024e; Song et al. 2022).
- **Text-Guided Editing:** Text-guided editing leverages natural language descriptions to directly influence the editing process of images, allowing for precise control over the modifications based on textual inputs. This approach combines the strengths of Natural Language Processing (NLP) and computer vision technologies, enabling a nuanced interpretation of text into visual changes (Hertz et al. 2022; Brooks et al. 2023; Kawar et al. 2023).
- **Layout and Region-Based Editing:** Layout and region-based editing involves modifying specific areas or rearranging elements within an image to alter its composition or focus. These methods are crucial for applications that require precise control over spatial arrangements and detailed modifications to image content (Avrahami et al. 2023a, b; Zeng et al. 2023).
- **I2I (Image-to-Image) Translation:** Image-to-Image (I2I) translation methods harness DMs to transform one source image into a different target one, maintaining the core content while altering its style, texture, or modal characteristics. This category is critical for applications ranging from artistic style transfer to functional medical imaging translations (Saharia et al. 2022a; Parmar et al. 2023; Trabucco et al. 2023).
- **Counterfactual Augmentation:** Counterfactual augmentation uses DMs to generate im-

Table 1 Overview of Foundation Diffusion Models (FDMs)

Model	Release Date	License	Conditioning Mechanisms	Text Encoder	Training Data	Params	Res. (pixels)	Notable Features	Indicative Derived Methods
Pixel-Space U-Net FDMs									
Imagen (Saharia et al. 2022b)	May'22	Prop	Text, Img	T5-XXL (Raffel et al. 2020)	LAION-400 M + Prop. (~860 M)	2B	1024 ²	First that explored using frozen LLMs as text encoders	Wang et al. (2023b), Chen et al. (2022a), Hu et al. (2024)
DeepFloyd IF (at StabilityAI 2023)	Apr'24	Open	Text, Img	T5-XXL (Raffel et al. 2020)	LAION-5B (~1B) (Schuhmann et al. 2022)	400 M - 4.3B	1024 ²	Cascaded pixel diffusion	Tang et al. (2024a)
Latent-Space U-Net FDMs									
SD1.x (Rombach et al. 2022)	Dec'21	Open	Text, Img, Mask, Layout, Semantic Map	CLIP ViT - L/14 (Radford et al. 2021)	LAION-5B (Schuhmann et al. 2022)	860M	512 ²	First major open-source LDM	Podell et al. (2023), Cao et al. (2023), Wei et al. (2023)
DALL-E2 (Ramesh et al. 2022)	Apr'22	Prop	Text, CLIP Embeddings	CLIP (Radford et al. 2021)	Prop. (~ 650M)	3.5B	1024 ²	Two-stage generation using CLIP latents	Becker et al. (2023)
SD2.x (Rombach et al. 2022)	Nov'22	Open	Text, Img, Mask, Layout, Semantic Map	OpenCLIP - ViT (Radford et al. 2021)	HLAION-5B (Schuhmann et al. 2022)	860 M	768 ²	Improved text understanding	Zhang et al. (2023b, 2024b)
SD Cascade (Permas et al. 2023)	Feb'23	Open	Text, Img, Mask, Layout, Semantic Map, Edge Map	CLIP-H (Ilharco et al. 2021)	LAION-5B (Schuhmann et al. 2022) (~103 M)	4B	1024 ²	Highly compressed semantic latent space	Phung et al. (2024)
SDXL (Podell et al. 2023)	Jul'23	Open	Text, Img	CLIP ViT-L (Radford et al. 2021) + OpenCLIP ViT-bigG (Cherti et al. 2023)	Prop	2.6B	1024 ²	Dual text encoders with size and crop conditioning	Sauer et al. (2025)
Kandinsky 3 (Arkh- ipkin et al. 2023)	Oct'23	Open	Text, Img, Mask	Fian-UL2 (Tay 2023)	LAION-5B + COYO-700 M + Prop. (stim150 M)	11.9B	1024 ²	Single-stage pipeline with large text encoder for better text understanding	Patel et al. (2024)

Table 1 (continued)

Model	Release Date	License	Conditioning Mechanisms	Text Encoder	Training Data	Params	Res. (pixels)	Notable Features	Indicative Derived Methods
<i>DALL-E</i> 3 (Bettker et al. (2023))	Oct'23	Prop	Text, Img	<i>T5-XXL</i> (Raffel et al. (2020))	Prop. (~ 1 B)	–	1024^2	Training on AI-generated detailed captions improves prompt following ability	–
SDXL Turbo (Sauer et al. (2025))	Feb'24	Open	Text, Img	CLIP-ViT-g-14 (Radford et al. (2021))	Distilling from SDXL	3.5B	512 ²	Single-step generation	Su et al. (2024)
Transformer-Based FDMs									
PixArt- α (Chen et al. (2023c))	Sep'23	Open	Text	Flan-T5-XXL (Raffel et al. (2020))	SAM-LLAVA + JourneyDB + Prop. (~ 25 M)	600 M	1024^2	Efficient training with only 1% cost of SOTA models	Chen et al. (2024a)
PixArt- δ (Chen et al. (2024a))	Jan'24	Open	Text, Img	T5 (Raffel et al. (2020))	Prop (~ 120 K)	600 M	1024^2	Fast 4-step sampling	–
SD 3 (Esser et al. (2024))	Mar'24	Mixed	Text, Img	T5-XXL (Raffel et al. (2020))	Prop. (~ 1 B)	0.8B – 8B	256^2	Rectified flow formulation enables efficient few-step sampling	–

Fig. 2 Taxonomy of DM-based image augmentation methods

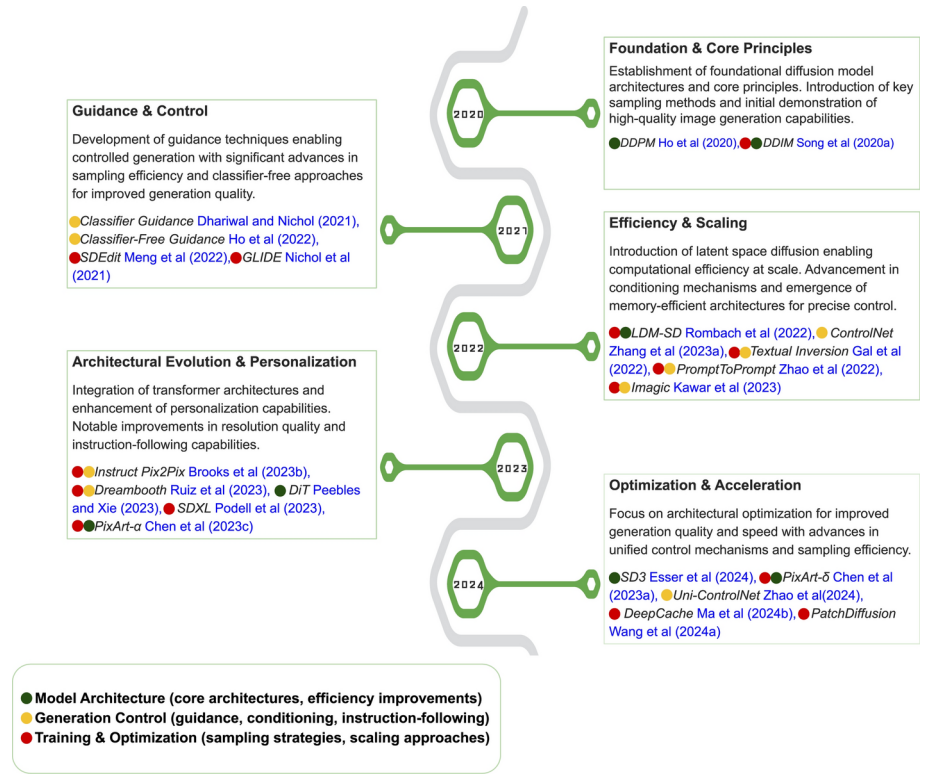
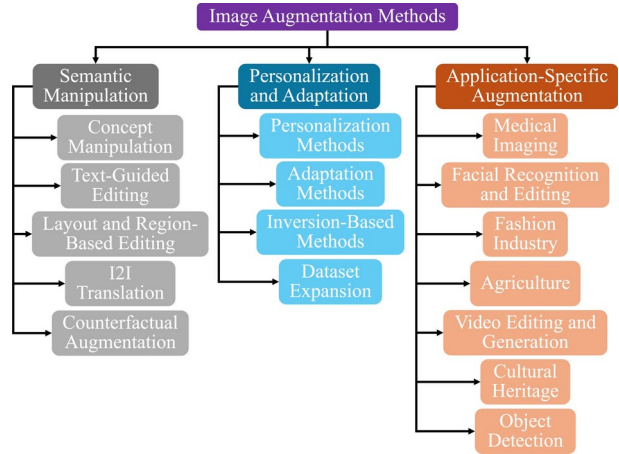


Fig. 3 Timeline representation of key recent DM-powered image augmentation methods

Table 2 DM-powered methods for image augmentation

Category	Subcategory	Methods
Semantic Manipulation	Concept Manipulation	Chen et al. (2024e), Luo et al. (2023), Zhao et al. (2022), Zhang et al. (2024c), Song et al. (2022), Huang et al. (2023a), Brack et al. (2022), Rando et al. (2022), Schramowski et al. (2023), Wasserman et al. (2024), Gandikota et al. (2023), (2024), Heng and Soh (2024), Kim et al. (2023), Kumari et al. (2023a), Ni et al. (2023), Zhang et al. (2024b)
	Text-Guided Editing	Kawar et al. (2023), Yu et al. (2024), Nichol et al. (2021), Hertz et al. (2022), Chen et al. (2024b), Lin et al. (2024), Huang et al. (2023b), Brooks et al. (2023), Wang et al. (2023a), Yang et al. (2024), Jin et al. (2024), Santos et al. (2024), Avrahami et al. (2023b), Yang et al. (2023a), Kirstain et al. (2023), Balaji et al. (2022), Geng et al. (2024)
	Layout and Region-Based Editing	Zeng et al. (2023), Chen et al. (2023d), Xue et al. (2023), Schnell et al. (2024), Lugmayr et al. (2022), Yu et al. (2023a), Ackermann and Li (2022), Couairon et al. (2022), Avrahami et al. (2022a, 2023a), Sarukkai et al. (2024), Xie et al. (2023b), Xiao et al. (2023), Zhang et al. (2023b), Levin and Fried (2023)
	I2I (Image-to-Image) Translation	Meng et al. (2021), Xu et al. (2024), Kim et al. (2022), Kwon and Ye (2022), Parmar et al. (2023), Su et al. (2022), Tumanyan et al. (2023), Wang et al. (2022a), Ma et al. (2023), Trabucco et al. (2023), Cao et al. (2023), Saharia et al. (2022a), Michaeli and Fried (2024), Rahat et al. (2024), Lingenberg et al. (2024)
	Counterfactual Augmentation	Sanchez et al. (2022), Sanchez and Tsafiris (2022), Madaan and Bedathur (2023), Yuan et al. (2022), Parihar et al. (2024), Vendrow et al. (2023)
	Personalization and Adaptation	Personalization Methods
Adaptation Methods		Hemati et al. (2023), Wu et al. (2023b), Dunlap et al. (2022), Zang et al. (2023), Zhu et al. (2023), Qiu et al. (2023), Islam et al. (2024a, b)
Inversion-Based Methods		Zhou et al. (2023a, b), Zhang et al. (2023e), Li et al. (2023b), Wallace et al. (2023), Tang et al. (2024b), Kwon et al. (2022), Mokady et al. (2023)
Dataset Expansion		Zhang et al. (2022), Li et al. (2023a), Ye et al. (2023), Wang et al. (2022b), Bansal and Grover (2023), Yin et al. (2023), Sheynin et al. (2022), Blattmann et al. (2022), Chen et al. (2024c)
Application-Specific Augmentation	Medical Imaging	Akrout et al. (2023), Sagers et al. (2022), Ali et al. (2022), Pinaya et al. (2022), Hu et al. (2022), Rouzrokh et al. (2022), Wolleb et al. (2022), Chambon et al. (2022a, b), Guo et al. (2023), Xia et al. (2022), Packhäuser et al. (2023)
	Facial Recognition and Editing	Boutros et al. (2023), Huang et al. (2024)
	Fashion Industry	Li et al. (2023c), Kong et al. (2023)
	Agriculture	Deng and Lu (2023), Muhammad et al. (2023), Chen et al. (2023a)
	Video Editing and Generation	Shin et al. (2024), Wu et al. (2023a)
	Cultural Heritage	Cioni et al. (2023)
Object Detection	Fang et al. (2024), Zhang et al. (2023c), Li et al. (2025), Tang et al. (2024c), Ma et al. (2024a)	

ages that represent hypothetical scenarios or what-if analyses, often used for enhancing model explainability and robustness. This includes generating counterfactual scenarios, where key elements are altered to assess potential outcomes. Such methods are useful in fields such as medical imaging and policy-making, where understanding the impact of variable changes is crucial (Sanchez and Tsafaris 2022; Sanchez et al. 2022).

3.2 Personalization and adaptation

Personalization and adaptation methods tailor the augmentation process to better suit specific datasets, tasks or user preferences (Gal et al. 2022; Ruiz et al. 2023). These methods enhance the relevance and effectiveness of augmented data, by finetuning models to align with particular requirements.

- **Personalization Methods:** Personalization methods aim at adapting DMs to generate content that meets specific user needs or preferences. These approaches often involve finetuning models, leveraging textual or visual inputs, and optimizing for personalized outputs (Gal et al. 2022; Ruiz et al. 2023; Wei et al. 2023).
- **Adaptation Methods:** Adaptation methods tailor DMs to different domains or tasks, improving their versatility and performance in new or varied contexts. These approaches are crucial for ensuring that models can generalize well across different datasets and applications (Qiu et al. 2023; Hemati et al. 2023).
- **Inversion-Based Methods:** Inversion-based DM methods leverage the ability to invert the diffusion process for enhancing image editing and augmentation capabilities, by obtaining the original image from a noisy or perturbed version (Mokady et al. 2023; Zhang et al. 2023e).
- **Dataset Expansion:** Dataset expansion methods utilize DMs to synthetically generate additional images, given an original input dataset. The goal is to address the limitations of small-scale datasets and to enhance the diversity of training images. This is critical for improving the generalization and robustness of machine learning models, particularly when the acquisition of extensive labeled datasets is impractical (Zhang et al. 2022; Li et al. 2023a).

3.3 Application-specific augmentation

Application-specific augmentation methods tailor the augmentation process so as to meet the unique requirements and properties of a given application domain, i.e. they extensively rely on using detailed and domain specific knowledge (Chambon et al. 2022a; Boutros et al. 2023). Such particular characteristics can be exploited for improving or guiding the augmentation process. Typical domains with unique requirements or properties are medical imaging, facial recognition, fashion industry, agriculture, etc.

- **Medical Imaging:** DMs for image augmentation are extensively employed to generate high-fidelity synthetic medical images, enhance existing datasets, and improve the robustness of diagnostic models. These methods address challenges such as data scarcity, variability in medical conditions, as well as the need for anonymized training data (Chung et al. 2022; Ali et al. 2022).

- **Other Domain-Specific Applications:** DMs have also been effectively applied to a wide range of other domain-specific applications, taking into account the specific requirements of each field. Examples include facial recognition and editing (Boutros et al. 2023), fashion industry (Li et al. 2023c), agriculture (Deng and Lu 2023), etc.

4 DM-powered methods for image augmentation

This section details the fundamental principles and mechanisms of DM-based methods for image augmentation, based on the taxonomy of approaches discussed in Sect. 3. The main benefits of each (sub)category are highlighted; thus, providing key insights regarding their practical usage.

4.1 Semantic manipulation

Semantic manipulation transforms the image appearance, while partially preserving its semantic content (Brooks et al. 2023), or manipulates the depicted semantic concept, potential textual elements, or layout. As discussed in Sect. 3, its subcategories are concept manipulation, text-guided editing, layout and region-based editing, image-to-image (I2I) translation, and counterfactual augmentation. In many cases the goal is to transform a specific input ‘reference’ image to an augmented variant, a process generally called ‘editing’.

4.1.1 Concept manipulation

Several methods for concept manipulation specialize in placing objects within images using the pretrained SD model, often being themselves training-free (Chen et al. 2024e; Luo et al. 2023; Zhao et al. 2022; Zhang et al. 2024c). In particular, they condition SD on new objects specified by text prompts, Web-retrieved images or a high-frequency map of the target object, stitched with the scene at the desired location (generated from Contrastive Language-Image Pre-training (CLIP) (Radford et al. 2021)) for placing them at the background of an input image. The new objects might be checked for semantic consistency via CLIP embedding similarity. For example, the method in Chen et al. (2024e) employs identity feature extraction (using a self-supervised DINOv2 model (Oquab et al. 2023)), detail feature extraction and feature injection to seamlessly integrate the target object into the scene, by feeding the ID tokens and the detail maps into the pretrained SD as guidance to generate the final composition. It also supports additional controls, like user-drawn masks to indicate the desired shape of the object during inference.

The method in Song et al. (2022) adapts the pretrained SD model for realistic object integration into scenes. It uses a content adaptor module that maps visual features from the input image object to a text embedding space to condition the SD. It is first pretrained on image-text pairs to learn semantics and then finetuned with SD to preserve object appearance. The SD model takes in the background image and the object embedding to generate a composite image. An example is depicted in Fig. 4, where various methods are compared on object adding at a specific location in the target image. The first column displays the desired object, while the second one the target image and the desired location of the object.



Fig. 4 Comparison of various semantic manipulation methods: **a** Desired object (col. 1), **b** Target image and desired object location (col. 2), **c** Copy-And-Paste (col. 3), **d** BLIP (Li et al. 2022) (col. 4), **e** SDEdit (Meng et al. 2021) (col. 5), **f** ObjectStitch (Song et al. 2022) (col. 6). Image from Song et al. (2022)

The methods ‘Composer’ (Huang et al. 2023a) and ‘Stable Artist’ (Brack et al. 2022) provide fine-grained control of the image generation process, by leveraging operations in the latent space. The ‘Composer’ decomposes an image into representative factors, such as text description, depth map, sketch, color histogram, etc., and trains a DM (Guided Language to Image Diffusion for Generation and Editing (GLIDE) (Nichol et al. 2021)) conditioned on these factors, allowing customizable content creation by recombining the aforementioned factors. ‘Stable Artist’ (Brack et al. 2022) uses Semantic Guidance (SEGA) to steer the diffusion process along multiple semantic directions corresponding to editing prompts, enabling subtle edits as well as changes in composition and style without masks or finetuning. SEGA allows the user to control the latent space representation, by calculating guidance vectors between the noise estimates of the original prompt and the editing prompts, and applying these to shift the unconditional noise estimate.

Large T2I models like SD can also replicate undesired behavior (Birhane and Prabhu 2021) or generate inappropriate content, such as copyrighted artworks (Jiang et al. 2023) or explicit images (Schramowski et al. 2023). To address these issues, several methods have been proposed, which can be organized in four main groups:

Image Post-Processing (Rando et al. 2022): These methods filter out inappropriate content from the generated images after the generation process.

Inference Guidance (Schramowski et al. 2023): These methods guide the diffusion process during inference to avoid generating undesired concepts. For example, Safe Latent Diffusion (Schramowski et al. 2023) defines an ‘unsafe’ textual concept and uses it to guide the diffusion process away from generating inappropriate content.

Image Inpainting (Yu et al. 2023a; Wasserman et al. 2024): These methods remove undesired objects or regions from images and inpaint the missing parts with appropriate content. They often use a mask to specify the region to be inpainted.

Model Finetuning (Gandikota et al. 2023, 2024; Heng and Soh 2024; Kim et al. 2023; Kumari et al. 2023a; Ni et al. 2023; Zhang et al. 2024b): These methods finetune the pre-trained DM to prevent it from generating undesired concepts. They often use approaches like concept erasure (Gandikota et al. 2023), self-distillation (Kim et al. 2023), or degeneration-tuning (Ni et al. 2023) to remove the unwanted concepts from the model’s learned representations.

4.1.2 Text-guided editing

Unlike text-prompt-based concept manipulation, which directly uses the text prompt to guide the editing process, text-guided editing methods first optimize a text embedding to reconstruct the input reference image (Kawar et al. 2023; Yu et al. 2024; Nichol et al. 2021). The reconstruction is the output of a pretrained conditional LDM, i.e., the outcome of the RD process. Additionally, a target text embedding is the CLIP representation of a textual prompt, which is given as a condition to the LDM. In the end, these methods interpolate between the optimized embedding and the target text embedding, so that conditioning on this interpolated vector eventually generates the desired edited/augmented image. For example, Imagic (Kawar et al. 2023) optimizes a text embedding, finetunes the DM to better reconstruct the input image, and then linearly interpolates between the optimized embedding and the target text embedding. Subsequently, this interpolated embedding is provided as a condition to the finetuned DM, so that it generates the desired edited augmented image.

A different approach to influence how the image is built word-by-word is by incorporating cross-attention maps within the diffusion process. These maps determine which parts of the image ('pixels') focus on specific elements ('tokens') of the text prompt at different stages of image creation, as illustrated in Fig. 5. Methods like (Hertz et al. 2022) and Chen et al. (2024b) use cross-attention maps to align textual descriptions with image regions, enabling nuanced image alterations. Some methods focus on segmenting images into learnable regions that can be individually manipulated, based on text commands. For example, the methods of Lin et al. (2024) and Huang et al. (2023b) divide the image into regions, by using pretrained models for feature extraction, and then apply text-guided editing to each region separately. This enhances the granularity of edits and allows for more precise control over specific parts of the image.

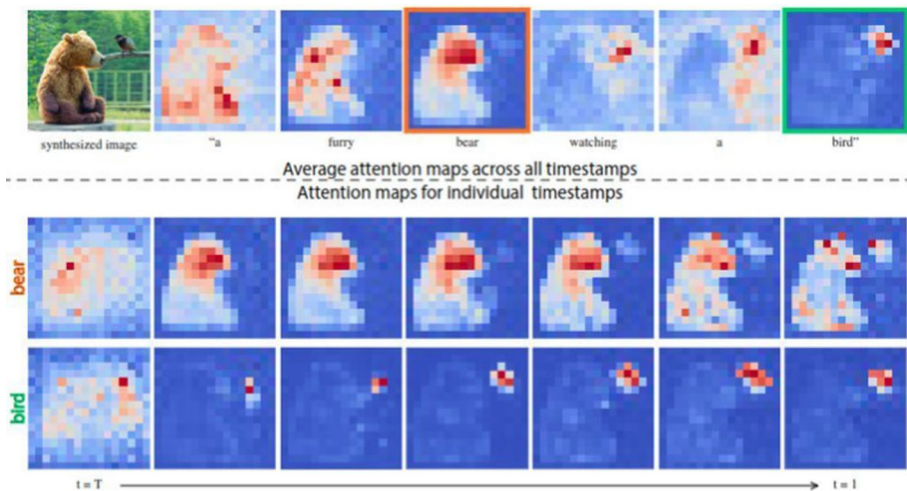


Fig. 5 Cross-attention maps of a text-conditioned diffusion image generator. The top row illustrates the average attention masks for each word in the prompt that synthesize the image on the left. The bottom rows depict the attention maps from different diffusion steps with respect to the words 'bear' and 'bird'. Image from Hertz et al. (2022)

In general, editing images based on human instructions allows a higher level of control over the DM's actions. These methods typically utilize a Large Language Model (LLM) to guide the editing process via human-like instructions (Brooks et al. 2023; Wang et al. 2023a; Yang et al. 2024; Jin et al. 2024; Santos et al. 2024; Geng et al. 2024). For example, InstructPix2Pix (Brooks et al. 2023) makes use of GPT-3 to generate a synthetic dataset of editing instructions (e.g., input caption: 'photograph of a girl riding a horse', edited caption: 'photograph of a girl riding a dragon') and subsequently a pretrained SD model is utilized to generate synthetic images matching the captions before/after. Subsequently, an SD model is trained on this generated dataset (by optimizing the LDM loss), so as to learn to perform image edits. Classifier-free guidance can also be used along with a text prompt to guide the diffusion process and to generate new images that comply better with their conditioning (Bansal et al. 2023). This improves the alignment between the generated images and the text prompt, without requiring a separate classifier model.

Moving towards a different direction, the method of Avrahami et al. (2023b) introduces a 'spatio-textual representation' that combines CLIP image embeddings of object segments during training and CLIP textual embeddings of local prompts during inference. This representation is incorporated into the DM, by concatenating it with the noisy image or latent code, and the DM is finetuned accordingly (by minimizing the LDM loss). Multi-conditional classifier-free guidance is employed to control the relative importance of each condition, involving a fine-grained variant using separate guidance scales and a fast variant using a single scale for the joint probability of all conditions.

Paint-by-Example (Yang et al. 2023a) introduces an exemplar-based editing approach, where it automatically merges a reference image into a source one. Self-supervised training utilizes the object's bounding box as a binary mask and the image patch inside as the reference image to reconstruct the source image. An information bottleneck compresses the reference image with a CLIP class token and injects the decoded feature into the diffusion process for better content understanding. Strong augmentations reduce train-test mismatches and handle irregular mask shapes. Controllability is achieved by training with irregular masks and using a classifier-free sampling strategy to adjust similarity between the edited and reference images.

X&Fuse (Kirstain et al. 2023) presents a general approach for conditioning on visual information in T2I generation: processing a reference and the input image separately, using shared ResBlocks, and then concatenating them (before attention blocks) to allow interaction between the two images. eDiff-I (Balaji et al. 2022) introduces an ensemble of expert denoisers, each specialized for a particular stage of the generation process, to capture distinct behaviors and to enhance model capacity without increasing computational cost during inference. The study compares T5 textual embeddings (Raffel et al. 2020), CLIP textual embeddings and optionally CLIP image embeddings during training. Subsequently, an inference approach (Paint-with-words) is presented, where the user can select phrases from the text prompt and can create binary masks of objects, which are provided as inputs to the model in order to control the spatial location of the objects. This is performed by modifying the cross-attention matrix between the image and the text features.

4.1.3 Layout and region-based editing

Several methods use text prompts to guide the generation and manipulation of images, based on layout and region information. For example, the approaches of Zeng et al. (2023) and Chen et al. (2023d) enable semantic image synthesis and geometric control using text prompts. These methods often utilize a separate layout encoder to model spatial and semantic information into a format suitable for image generation. Indicatively, the latter consists of a precision-based mask pyramid that represents region shapes at multiple resolutions combined with text embeddings in Zeng et al. (2023). An example of such an approach is depicted in Fig. 6. In Chen et al. (2023d), the layout encoder translates geometric layouts into text prompts by mapping locations, classes, and conditions into text tokens, in order to extract spatial information from the input image's layout, and then conditions the DM on this layout encoding.

ControlNet (Zhang et al. 2023b) is another powerful approach for layout and region-based editing. It introduces a DNN architecture that adds spatial conditioning controls to large, pretrained T2I DMs, like SD. ControlNet creates a trainable copy of the model's encoding layers and connects it to the original model using 'zero convolutions'. This allows for efficient finetuning on small datasets for various conditioning tasks, such as edge detection, pose estimation, and depth mapping. The architecture can process both text prompts and conditioning images (e.g., edge maps, pose maps, depth maps) as inputs, making it highly versatile for different types of spatial control in image generation and editing.

Certain methods focus on generating images from coarse layouts or scribbles. For instance, the method of Xue et al. (2023) utilizes SD to generate freestyle images, by integrating semantic text embeddings with spatial layouts into SD. It represents each semantic class in the input layout using a text concept, which are then encoded into text embeddings. A Rectified Cross-Attention (RCA) module is subsequently introduced to inject these text semantics into the corresponding layout regions within the diffusion model's U-Net cross-attention layers. By finetuning just the U-Net with the integrated RCA on layout-image




Fig. 6 Indicative layout and region-based editing results with different precision levels. For each input layout, the images are progressively modified (starting with same noise), so that the generated images at different precision levels can have similar styles. Image from Zeng et al. (2023)

pairs, the pretrained model can generate images conditioned on both user-specified layouts and free-form text prompts, enabling capabilities like binding new attributes to objects and generating unseen object classes. The method of Schnell et al. (2024) utilizes ControlNet (Zhang et al. 2023b) to generate synthetic images conditioned on scribble labels and text prompts. It employs classifier-free guidance (10% of the conditioning scribble inputs are randomly dropped and replaced with a learnable embedding) and introduces an encoding ratio to adjust the diversity and photorealism of the generated images (by performing fewer FD steps), allowing for a trade-off between mode coverage and sample fidelity.

Inpainting methods (Lugmayr et al. 2022; Yu et al. 2023a) make use of DMs to fill-in missing regions of an image based on a given mask. They typically condition the DM on the masked image and the mask itself, and then generate content to fill-in the masked region. A subset of methods (Ackermann and Li 2022; Couairon et al. 2022; Avrahami et al. 2022a, 2023a) employs a multi-stage diffusion process with mask guidance, in order to achieve high-resolution image editing and inpainting.

Collage Diffusion (Sarukkai et al. 2024) is a method that takes a user-defined sequence of layers as input, called Collage. It consists of a full-image text string describing the entire image to be generated, along with a sequence of layers ordered from back to front; each layer consists of an RGBA image (alpha-masked input image) and a text describing it. The method modifies the text-image cross-attention in the DM to achieve spatial fidelity, while extending ControlNet to preserve appearance fidelity on a per-layer basis. Collage Diffusion also allows users to control the harmonization-fidelity trade-off for each layer by specifying desired noise levels, enabling layer-by-layer image editing. Similarly, SmartBrush (Xie et al. 2023b) utilizes text and shape guidance for object inpainting with DMs, enabling users to control the inpainted content based on both text descriptions and object masks. Fast-Composer (Xiao et al. 2023) is a tuning-free multi-subject image generation method that augments text prompts with visual features extracted from input images using an image encoder. It models text prompts and input images as embeddings, using pretrained CLIP encoders, and then uses an Multilayer Perceptron (MLP) to augment the text embeddings with the visual features. The method trains the image encoder, MLP module, and U-Net with a denoising loss using a subject-augmented image-text paired dataset, while localizing cross-attention maps using the reference subject's segmentation mask to prevent identity blending in multi-subject generation. It also employs delayed subject conditioning in iterative denoising to balance identity preservation and editability.


The method of Levin and Fried (2023) is designed for allowing editing with pixel-level control over the amount of applied modification, using a 'change map'. The latter is a matrix of dimensions identical to the spatial resolution of the original input image, describing the strength of the edit to be applied at each location. Different approaches can be used to generate it, such as Segment-Anything (Kirillov et al. 2023), MiDas (Ranftl et al. 2020) or even manually drawn change maps. The method can operate in latent space and has been evaluated on SD, SDXL, Arkhipkin et al. (2023) and DeepFloyd IF (at StabilityAI 2023) pretrained LDMs, using text prompts (either manually created, or by reversing the input image into CLIP and BLIP (Li et al. 2022)) along with change maps to guide the inference process.

4.1.4 I2I (image-to-image) translation

Several methods make use of conditional DMs for I2I translation (Michaeli and Fried 2024; Rahat et al. 2024). For example, SDEdit (Meng et al. 2021) utilizes a conditional DM, guided by stochastic differential equations, to perform image editing. CycleNet (Xu et al. 2024) and DiffusionCLIP (Kim et al. 2022) finetune a conditional DM with a CLIP-based loss to ensure that the generated image matches the target image’s text description. Saharia et al. (2022a) employ conditional DMs to learn the distribution $p(y|x)$ for various tasks, like colorization, inpainting, JPEG restoration, and uncropping. The model incorporates a U-Net architecture with self-attention layers, conditioning on the input image through concatenation. During training, it predicts the noise added to the original image, minimizing L2 or L1 loss between the predicted and the actual noise. Then, the inference process involves iterative denoising over 1000 timesteps, starting from Gaussian noise and progressively refining the output. The method’s strength lies in its ability to handle multiple tasks with a single architecture, eliminating the need for task-specific customizations (Fig. 7).

A different group of methods focus on disentangling style and content for more targeted I2I translation (Kwon and Ye 2022; Lingenberg et al. 2024). For instance, the approach of Kwon and Ye (2022) makes use of separate encoders for style and content, and then injects these representations into the DM to generate the translated image. Other methods explore few-shot or zero-shot I2I translation. For example, pix2pix-zero (Parmar et al. 2023) first inverts the input image to obtain a noise map, using DDIM and BLIP. These noise maps

Fig. 7 Indicative examples of an Image-to-Image DM that can generate high-fidelity outputs across a variety of tasks, without incorporating task-specific customizations. Image from Saharia et al. (2022a)



are then regularized to improve editability, using an autocorrelation objective. Edit directions are automatically discovered by generating diverse sentences for the source and target domains, and computing the mean difference between their CLIP embeddings. In order to preserve image structure during editing, a novel cross-attention guidance technique is employed, which matches the edited cross-attention maps with reference maps corresponding to the original structure. The method builds upon the SD model, utilizing its CLIP conditioning and cross-attention layers. On the other hand, Dual Diffusion Implicit Bridges (DDIBs) (Su et al. 2022) encompass two independently trained DMs (one for the source domain and one for the target one) to perform I2I translation, without requiring joint training on paired data.

The method of Tumanyan et al. (2023) injects spatial features and self-attention maps extracted from the diffusion process of a source image into the generation process of the translated image, enabling fine-grained control over the generated structure, while complying with the target text prompt. Additionally, the method of Wang et al. (2022a) leverages a pretrained T2I DM (GLIDE) as a generative prior, consisting of a base model at 64×64 resolution and an upsampling model to 256×256 . The framework adopts an encoder-decoder architecture, where a task-specific encoder maps input conditions (e.g., segmentation masks, sketches) to the semantic latent space of the pretrained diffusion decoder. The DM is pretrained on diverse text-image pairs, enabling the latent space to be conditioned on highly semantic text embeddings. The pipeline involves pretraining the DM, training a task-specific encoder using a two-stage finetuning scheme, finetuning the diffusion upsampler with adversarial training, and sampling the base model and upsampler during inference, using the encoded input and normalized classifier-free guidance. Moreover, another relevant framework is presented in Ma et al. (2023), aiming at joint subject- and text-conditional image generation. The DM (SD) generates high-quality images that align semantically with input texts, while preserving input image subjects. It leverages CLIP encoders to map texts and images into a unified multi-modal latent space. A fusing sampling strategy balances noise predictions between the unified condition and pure text condition to avoid overfitting to background information, i.e., the final noise estimate fuses the two predictions using a given ratio.

DA-Fusion (Trabucco et al. 2023) is another method that aims to alter the semantic contents of an image. It relies on the pretrained SD and modifies the input reference image based on class label guidance. Novel visual concepts outside the DM's training dataset are handled by inserting and finetuning new token embeddings in the text encoder, using the textual inversion approach (see Sect. 4.2). Input images are spliced into the diffusion process at a random timestep to guide the generation, rather than generating from scratch. Randomizing the splice timestep provides diverse augmentation intensities.

MasaCtrl (Cao et al. 2023) introduces a mutual self-attention mechanism within U-Net, allowing the model to query correlated local structures and textures from a source image, while maintaining consistency with a target editing prompt. The architecture modifies the standard U-Net blocks by converting self-attention into mutual self-attention, where query features originate from the current denoising process, while key and value features are derived from the input image's diffusion process. In order to prevent confusion between foreground and background elements, MasaCtrl employs a mask-guided mutual self-attention strategy, utilizing masks extracted from cross-attention maps.

4.1.5 Counterfactual augmentation

A significant number of methods make use of conditional DMs to generate counterfactual images. For instance, the approach of Sanchez et al. (2022) utilizes a conditional DM to generate ‘healthy’ counterfactuals of brain MRI images with tumors, by conditioning the model on the tumor mask. Diff-SCM (Sanchez and Tsafaris 2022) employs a conditional DM to estimate the effects of interventions in a causal framework, by conditioning the model on the intervention variable $do(class)$, corresponding to ‘how the image should change in order to be classified as another class’. For instance, for a given photo of a dog in a park, the dog should be replaced by a cat, while leaving the rest of the image unchanged if $do(cat)$ holds (Fig. 8).

A great portion of methods focus on generating counterfactuals for specific applications. For example, the methods of Madaan and Bedathur (2023) and Yuan et al. (2022) generate counterfactual images to explain the model’s behavior and to improve its robustness. On the other hand, a different set of approaches utilize counterfactual augmentation to address bias and fairness issues in datasets and models. For example, the work of Parihar et al. (2024) utilizes a conditional DM to generate counterfactual images that balance the distribution of sensitive attributes (e.g., gender, race) in the dataset. The model is conditioned on the desired attribute distribution and generates images that match it, while preserving the remaining aspects of the image.

Another set of methods focus on expanding a given dataset with generated counterfactual or Out-of-Distribution (OOD) examples, in order to improve the robustness of a model (e.g., classifier) trained on the expanded dataset. For example, the method of Vendrow et al. (2023) utilizes a conditional DM to generate counterfactual examples that change specific attributes of the input image (e.g., object position, color, texture), while preserving the overall scene structure. These counterfactual examples are used to diagnose and mitigate model failures on OOD data.

4.2 Personalization and adaptation

Personalization and adaptation constitute very common methodologies in image augmentation tasks. In the remainder of this subsection, respective DM-based approaches are detailed, grouped according to the taxonomy presented in Sect. 3.




Fig. 8 Counterfactuals on ImageNet generated by Diff-SCM. From left to right: a random image sampled from the data distribution and its counterfactuals $do(class)$, corresponding to ‘how the image should change in order to be classified as another class?’. Image from Sanchez and Tsafaris (2022)

4.2.1 Personalization methods

Personalization implies that the DM must be adapted to generate content that meets specific user needs or preferences. These methods often involve finetuning pretrained models, leveraging textual or visual inputs, and optimizing for personalized outputs.

A significant number of methods make use of finetuning to personalize pretrained T2I DMs for subject-driven generation. For example, DreamBooth (Ruiz et al. 2023) finetunes the DM within a few-shot learning setting, using a small set (3–5) of images of a single subject paired with a text prompt containing a unique identifier and the name of the class the subject belongs to (e.g., ‘A [V] dog’). These images contain a specific subject and serve as the training dataset that the DM needs to learn, while the text prompt simply describes the new subject. In particular, the method applies a class preservation loss to ensure that the generated images maintain the identity of the subject. Similarly, HyperDreamBooth (Ruiz et al. 2024) is a method for efficient personalized DMs that introduces three key components, namely Lightweight DreamBooth (LiDB), a HyperNetwork for fast personalization, and rank-relaxed fast finetuning. LiDB decomposes the rank-1 LoRA (Hu et al. 2021) weight space using a random orthogonal incomplete basis, resulting in smaller personalized models. Additionally, the HyperNetwork, consisting of a ViT encoder and transformer decoder, predicts LiDB residuals from input images, using diffusion denoising and weight-space losses. Eventually, rank-relaxed fast finetuning captures fine-level details by relaxing the LoRA rank and finetuning with the predicted HyperNetwork weights, improving subject fidelity while maintaining fast personalization (Fig. 9).

A different set of methods make use of the so-called ‘textual inversion’ (Gal et al. 2022) to personalize T2I generation. This process learns a new text embedding (termed ‘pseudo-word’) that represents a specific visual concept, using a small set of images depicting that concept. This allows the model to generate images of the concept, using natural language descriptions that include the pseudo-word. ProSpect (Zhang et al. 2023d) extends this idea by learning a collection of pseudo-words (called a ‘prompt spectrum’) that capture different visual attributes (e.g., material, style, layout) of the concept.

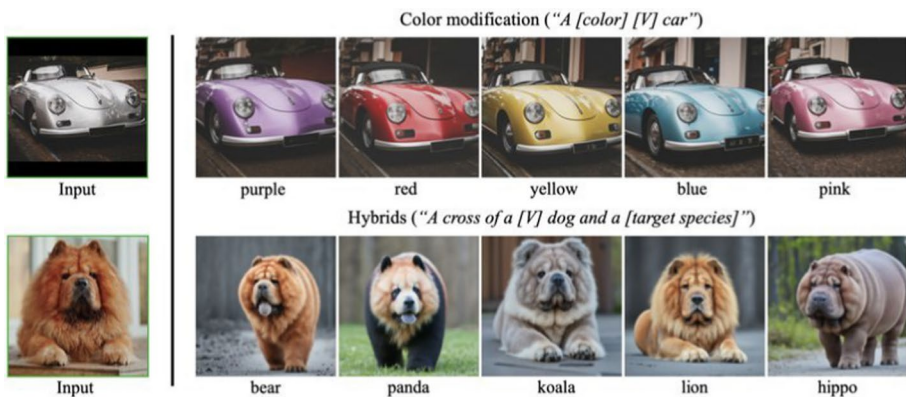


Fig. 9 Example of personalization-oriented image augmentation, where the given subject’s properties need to be modified, while preserving its key visual features that forge its identity. Image from Ruiz et al. (2023)

On a different basis, another group of methods focus on combining multiple concepts for personalized generation. For example, the approaches of Kumari et al. (2023b) and Vinker et al. (2023) utilize a compositional mechanism, where different concepts are represented by separate text embeddings that can be combined to generate novel images. StyleDrop (Sohn et al. 2023) and DreamArtist (Dong et al. 2022) adopt a style-based approach, where the style of the generated image is controlled by a learned embedding.

Another conceptualization comprises methods that make use of encoder-based approaches to personalize T2I generation. For example, ELITE (Wei et al. 2023) utilizes an image encoder to map user-provided input images to text embeddings, which are then used to condition the pretrained SD. This allows the model to generate personalized images that match the style and content of the input. Similarly, the methods of Gal et al. (2023) employs a cross-attention mechanism to inject visual features from an image encoder into the DM at different layers.

SuTI Chen et al. (2024d) is an apprenticeship learning framework that trains expert DMs (Imagen (Saharia et al. 2022b)) for each subject on image-text clusters, where each cluster contains 3–10 images-prompts. Then, K expert DMs are fine tuned on K clusters in order to learn particular subjects belonging to the clusters. Subsequently, a single DM is trained on the outputs of these expert DMs. This approach enables SuTI to compose specialized knowledge from numerous expert models into a single, versatile model. During inference, it can generate customized images for new subjects, without some kind of optimization.

InstantBooth (Shi et al. 2024) is an alternative approach that enables personalized generation without requiring deployment-time finetuning of the DM itself. It introduces a learnable image encoder to convert input images to a textual embedding and adapter layers to inject rich visual features for better identity preservation. The original model weights are frozen and only the new components are trained. A unique identifier token V represents the input concept in the text prompt, while the input images are cropped and background-masked.

Perfusion (Tewel et al. 2023) is a compact and efficient architecture based on key-locked rank-1 editing, which addresses key challenges in personalized generation, such as overfitting and maintaining high visual fidelity, while allowing creative control. Key-locked rank-1 editing prevents overfitting by locking the keys of a learned concept to those of its supercategory in the cross-attention layers, ensuring the concept inherits the generative prior without deviating too much. The method maintains high visual fidelity, by learning concept-specific values in an extended latent space to capture the concept's unique appearance. Moreover, the rank-1 update is gated based on the similarity between the current encoding and the target concept, allowing fine-grained control over the influence of each concept at inference time.

Taming Encoder (Jia et al. 2023) is a method for generating images of customized objects specified by users, without the need for lengthy per-object optimization, as required by previous approaches. In particular, given an input image and a text description of the desired output, the image object encoder (CLIP) computes the object embedding, while the text encoder computes the text embedding. These two are then passed to Imagen (Saharia et al. 2022b) for generating the final output in a single forward pass. Training data is prepared using a captioning model (PaLI (Chen et al. 2022b)) and binary masks to isolate the object. The framework is jointly trained on domain-specific and general-domain datasets, using cross-reference regularization and object-embedding dropping. The entire network, except from the object encoder, is tuned using the same objective as the pretrained DM.

DisenBooth (Chen et al. 2023b) is a framework for subject-driven T2I generation that utilizes a textual identity-preserving embedding and a visual identity-irrelevant embedding. In particular, an identity-preserving branch maps the subject's identity to a special text token using a CLIP text encoder, while an identity-irrelevant branch employs a pretrained CLIP image encoder and a learnable mask to extract identity-irrelevant features, which are then aligned with the text feature space using an MLP with skip connection (Adapter). During finetuning, DisenBooth utilizes denoising, weak denoising, and contrastive embedding objectives, and employs parameter-efficient finetuning using LoRA. After finetuning, DisenBooth enables flexible and controllable generation, by combining the subject's identity token with other text descriptions and optionally inheriting characteristics from a reference image.

HiPer (Han et al. 2023) personalizes pretrained SD models for text-driven image manipulation, using a single input image. It decomposes the CLIP text embedding space into initial tokens and end tokens for preserving subject identity. In particular, given an input image, the input image's prompt, and the target desired image's prompt, HiPer optimizes the embedding in the DM's latent space. During inference, the target prompt's embedding is concatenated with the optimized HiPer embedding to condition the pretrained SD model, generating a manipulated image that preserves the subject's identity, while incorporating the target prompt's semantics. HiPer's main contribution lies on its decomposition of the text embedding space for separate control of semantics and identity, requiring only a single source image without finetuning the DM.

4.2.2 Adaptation methods

A significant portion of adaptation methods rely on finetuning to adjust pretrained DMs to new domains. For example, the method of (Hemati et al. 2023) employs the pretrained SD model to generate synthetic images that bridge gaps between domains and reduce non-iidness in training data. It receives an image as input from one domain and a guidance attribute (either a text prompt or an image) from another one of the same class, using the LDM to create interpolated synthetic images. The approach of Wu et al. (2023b) introduces an optimization method to find the optimal soft combination weights of the input image description and the target image description text embeddings at each denoising step. The weights are optimized to match the target attribute using a CLIP-based loss, while preserving other content using a perceptual loss. The method of Dunlap et al. (2022) makes use of a domain-specific CLIP model to guide the diffusion process towards generating images that match the target domain (Fig. 10). Moreover, the method of Zang et al. (2023) utilizes a domain-specific discriminator to filter out low-quality or out-of-domain samples during training.

Another group of approaches focuses on adapting pretrained DMs to specific tasks or applications. Indicative examples are the cases of nuclei segmentation (Yu et al. 2023b) or knee osteoarthritis severity classification (Chowdary et al. 2023), by finetuning on datasets of nuclei images or X-ray images, respectively.

DomainStudio (Zhu et al. 2023) comprises a recent method that adapts pretrained DMs to target domains using limited data, by introducing a pairwise similarity loss to maintain the relative distances between the generated samples during domain adaptation and designing a high-frequency detail enhancement approach to improve generation quality. Additionally, Orthogonal Finetuning (OFT) (Qiu et al. 2023), which is a method for adapting

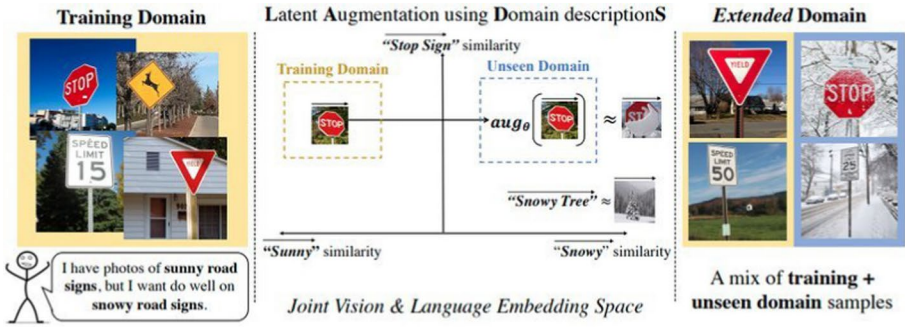


Fig. 10 Example of DM-based domain adaptation for image augmentation, where a model trained on images from a source domain can be applied to images from a target one. Image from Dunlap et al. (2022)

DMs to new tasks, injects trainable orthogonal matrices into the model’s attention layers, transforming neuron weights multiplicatively. These matrices, initialized as identity and kept orthogonal, enable stable finetuning that preserves the model’s semantic knowledge. OFT receives subject images or control signals with text prompts as input and generates images following the subject identity or control while matching the prompt. The orthogonal transformation of neurons maintains their pairwise angles, leading to faster convergence, better quality and controllability.

The methods of Islam et al. (2024a) and Islam et al. (2024b) introduce a prompt-guided approach that enhances both in- and cross-domain image classification. Specifically, they leverage image editing with carefully designed conditional prompts to generate augmented images based on concatenation of original and edited content, followed by blending with fractal patterns, which reduces overfitting and ensures robust training.

4.2.3 Inversion-based methods

Inversion-based methods typically optimize a latent code, i.e., the latent-space representation of a user-provided input image that contains an object to be learnt by the DM, so that it can reconstruct the image using the DM. Subsequently, the latent code is manipulated so as to steer the DM to perform the desired editing or augmentation. Various types of inversion-based image augmentation approaches are summarized in Fig. 11.

A broad set of methods employ inversion to generate high-quality synthetic data for training other DNNs. For example, the methods of Zhou et al. (2023b) and Zhou et al. (2023a) generate synthetic images, by inverting real images to the latent space, and then sampling new latent codes around the inverted ones. This facilitates the production of diverse and realistic augmentations of the original dataset, which can significantly improve the performance of downstream classification models.

Another group of methods adopt inversion to enable precise image editing and style transfer. For example, Null-Text Inversion (Mokady et al. 2023) consists of two main components, namely pivotal inversion, which estimates an initial diffusion trajectory using DDIM inversion, and ‘null-text’ optimization, which optimizes only the unconditional ‘null-text’ embeddings for each diffusion timestep. Given an input image and its associated text caption, the method inverts the image to obtain a latent code and optimized null-text

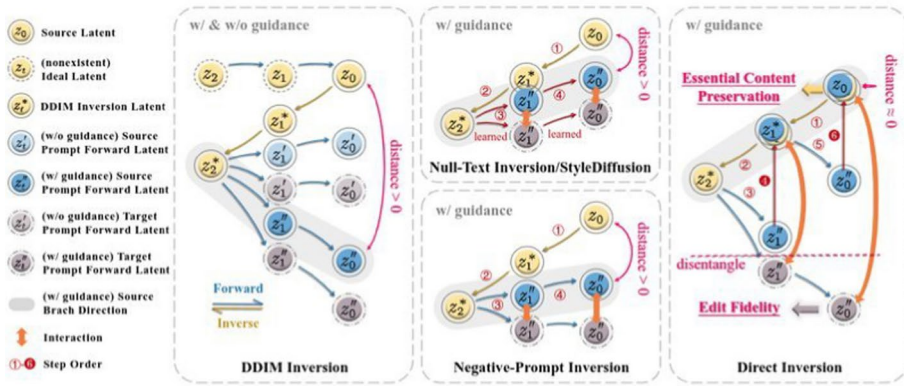


Fig. 11 Various types of inversion-based approaches for diffusion-powered image editing. Image from Ju et al. (2023)

embeddings. These can then be used to edit the image intuitively using approaches such as P2P (Hertz et al. 2022), by modifying the text prompt, while preserving the model weights and conditional embeddings.

Following a similar conceptualization, the method of Zhang et al. (2023e) enables artistic style transfer from a single reference painting to an input image, using a pretrained SD. In particular, it introduces an attention-based textual inversion module that learns a style representation ("[C]") from the CLIP image embedding of the reference painting, through multi-layer cross-attention. This learnt textual embedding is then encoded into a DM caption conditioning format to guide the generation of new images that combine the content of the input image with the artistic style of the reference painting. Specifically, the method involves extracting the image embedding of the reference painting, learning the corresponding textual embedding via the inversion module, encoding it into DM conditioning format, applying stochastic inversion to the content image to obtain an initial latent noise map, and generating the output image conditioned on the textual embedding and an inverted noise map (noisy version of the input image).

In a closely related way, the method of Li et al. (2023b) learns to invert the input image to value embeddings. while preserving the keys and attention maps from the original model. The respective architecture consists of a frozen CLIP image encoder, a learnable mapping network, and the SD model. The keys control the output image structure, while the values determine the object style. The editing process involves DDIM inversion to generate latent codes and attention maps, training the mapping network to reconstruct these while preserving object-like attention, and using the trained network with target prompt embeddings for editing.

Following a different line of research, a particular group of methods focus on improving the inversion process itself. For example, EDICT (Wallace et al. 2023) utilizes a coupled DM to perform inversion, where one model learns to map the image to the latent space and the other learns to map the latent code back to the original image space. This allows for more accurate and efficient inversion, compared to optimizing the latent code directly. LocInv Tang et al. (2024b) utilizes a localization-aware inversion process that optimizes the latent code to match the spatial attention maps of the input image, which facilitates in preserv-

ing the local structure and the details of the image during editing. Moreover, the method of Kwon et al. (2022) introduces an asymmetric reverse process (Asypr) that discovers a semantic latent space, called ‘h-space’, in frozen pretrained DMs. It modifies the RD process by shifting the predicted noise in the bottleneck feature maps of the U-Net architecture, while preserving the direction pointing to the current timestep. This breaks destructive interference and enables semantic manipulation of the generated images. The h-space exhibits desirable properties for editing, such as homogeneity, linearity, composability, robustness, and consistency across timesteps.

4.2.4 Dataset expansion

Conditional DMs are often used for generating synthetic images sampled from the underlying distribution of a given dataset. For example, the method of Zhang et al. (2022) trains a class-conditional DM on a small dataset and then uses it to generate additional images per class. The generated examples are filtered using a classifier to ensure that they are realistic and match the target class distribution. Similarly, the approach of Li et al. (2023a) utilizes a semantics-guided DM to generate synthetic images that match the semantic layout and object categories of the original dataset. The method of Wang et al. (2022b) comprises a single-scale pixel-level DDPM that learns the internal patch distribution of a single natural image. It employs a U-Net denoising network with a restricted patch-level receptive field, enabling it to capture the image patch statistics without memorizing the entire image. The model is trained at a single scale, avoiding error accumulation as observed in progressive growing approaches. The method of Ye et al. (2023) leverages a time-conditioned, U-Net-powered LDM and a two-stage training process to generate realistic synthetic images and to improve classification performance. The first stage involves large-scale pretraining on unlabeled data to learn common features for unconditional image synthesis. The second stage finetunes the model on a small labeled dataset, enabling conditional synthesis guided by a latent classifier. Similarly, the method of Bansal and Grover (2023) generates synthetic examples that are close to the training data in feature space but far away in the image space, which facilitates towards improving the model’s ability to generalize to novel visual concepts (Fig. 12).

A different group of methods leverage the ability of DMs to generate high-quality images from text descriptions. For example, TTIDA Yin et al. (2023) finetunes a pretrained text-to-text (T2T) model (GLIDE (Nichol et al. 2021)) to generate diverse text descriptions of




Fig. 12 Examples of DM-based dataset expansion image augmentation. Image from Bansal and Grover (2023)

objects and scenes, and then uses a T2I DM to generate corresponding images. This allows the generation of large-scale synthetic datasets with rich annotations, which can be used to train more accurate and robust vision models.

KNN-Diffusion (Sheynin et al. 2022) consists of a multi-modal CLIP encoder, a non-trainable retrieval index, and a trainable DM conditioned on the retrieved embeddings. During training, the model receives an image as input, while the CLIP image embedding and its k nearest neighbors are used to condition the generation. During inference, the model receives a text prompt as input, while the CLIP text embedding and its k nearest neighbors are used for conditioning. The DM iteratively denoises a noise vector, guided by the input embedding and the retrieved neighbors, in order to generate the output image. The retrieved neighbors bridge the distribution gap between the image and the text embeddings.

The Retrieval-augmented Diffusion Model (RDM) (Blattmann et al. 2022) combines a trainable conditional DM with a fixed external database of diverse visual examples and a non-trainable retrieval function. During training, the method retrieves from the database the k nearest neighbors for each image, using CLIP embeddings, and encodes them with a pretrained encoder. Subsequently, it conditions the generative decoding head on these encoded representations to generate the target image. During inference, the database and the retrieval function can be flexibly swapped to enable unconditional, class-conditional, text-conditional sampling, or style transfer, by retrieving neighbors based on different criteria or using a database with a different visual style. This approach augments a comparatively small generative model with a large external memory, allowing it to compose novel images based on relevant retrieved information rather than memorizing the full training data; thereby, reducing model size, while enhancing performance and flexibility.

4.3 Application-specific augmentation

This subsection discusses application-specific DM-powered image augmentation approaches, i.e., methods that take into account particular facts and characteristics that are only present in the examined application domain. Common application fields that have been studied so far are medical imaging, facial recognition, object detection and agriculture, to name a few.

4.3.1 Medical imaging

In the field of medical imaging, image augmentation methods typically employ conditional DMs to generate synthetic medical samples. For example, the methods of Akrouf et al. (2023) and Sagers et al. (2022) train class-conditional DMs on datasets of skin lesion and use them to generate additional examples of each lesion type. The generated images are used to augment the training data and to improve the performance of skin disease classification models. Similarly, the methods of Ali et al. (2022) and Packhäuser et al. (2023) make use of DMs to generate synthetic chest X-ray images, which are exploited to train more robust models for detecting thoracic abnormalities.

A particular group of methods focus on generating specific types of medical images, such as brain Magnetic Resonance Imaging (MRI) or retinal Optical Coherence Tomography (OCT). For example, the approach of Pinaya et al. (2022) trains a DM on a dataset of brain MRI images and uses it to generate synthetic images with different neurological conditions.

The method of Hu et al. (2022) trains a DM on retinal OCT images and uses it to generate denoised and super-resolved versions of the images.

A different set of methods employs DMs for image inpainting and anomaly detection. For example, the method of Rouzrokh et al. (2022) trains a DM to inpaint brain MRI images with tumors, by conditioning the model on the tumor mask. This allows for the generation of complete and realistic brain images from partial or corrupted scans. In contrast, the approach of Wolleb et al. (2022) employs DDIMs to generate synthetic healthy images from diseased individuals. In particular, a U-Net architecture is trained to iteratively denoise noisy representations, while a binary classifier guides the denoising process towards the healthy class. The method assumes unpaired training data of healthy and diseased images with only image-level labels. For an unseen test image, the DDIM noising process encodes its anatomical information into a noisy representation, which is then iteratively denoised using classifier guidance to generate a corresponding healthy synthetic image. The anomaly map is computed as the pixel-wise difference between the input and the synthetic images, highlighting the diseased regions. The noise level and classifier gradient scale control the trade-off between preserving input details and translating to the healthy class.

Focusing on a different aspect, another group of methods focus on adapting pretrained DMs to the medical domain. For example, the method of Chambon et al. (2022b) finetunes a pretrained T2I DM on a dataset of chest X-ray images and radiology reports, using textual inversion and cross-attention control. This allows for the generation of synthetic chest X-ray images with specific abnormalities and attributes, based on natural language descriptions. Similarly, the approach of Chambon et al. (2022a) finetunes the pretrained SD on a dataset of chest X-rays and radiology reports to adapt it to generate synthetic medical images from text prompts. The key components are a frozen Variational Auto-Encoder (VAE) encoder for compressing inputs, a U-Net denoiser conditioned on text embeddings from a CLIP encoder and a VAE decoder. During inference, random noise is encoded and progressively denoised by the U-Net guided by the text prompt, ultimately yielding a synthetic chest X-ray reflecting the described imaging findings (Fig. 13).

More recently, several methods have aimed at further advancing the augmentation performance and quality of the generated images, making use of more recent and sophisticated DM architectures. The method of Guo et al. (2023) introduces PD-DDPM, an accelerated DM for medical image segmentation that uses a pre-segmentation network to generate noisy segmentation predictions, which are then denoised in fewer steps compared to a vanilla DM. The method of Xia et al. (2022) employs a DDPM that is trained on paired low-dose Computed Tomography (CT) and normal-dose CT images. During the FD process, Gaussian noise is gradually added to the normal-dose image based on a time varying schedule, resulting in a noisy image following a standard normal distribution. During the RD process, the U-Net learns to recover the clean normal-dose image from the noisy one, conditioned on




Fig. 13 Examples of application-specific DM-based image augmentation (medical imaging), given a specific textual prompt. Image from Chambon et al. (2022a)

the corresponding low-dose image, by predicting and removing the noise at each timestep. To improve sampling efficiency, a fast ordinary differential equation (ODE) solver, termed DPM-Solver (Lu et al. 2022a), is integrated into the RD process, allowing for faster sampling with fewer steps, while maintaining or improving denoising performance compared to the original DDPM.

4.3.2 Other domain-specific applications

Apart from medical imaging, DMs have also been successfully applied to a wide range of domain-specific applications, as detailed in the followings.

Facial Recognition and Editing: The methods of Boutros et al. (2023) and Huang et al. (2024) focus on generating high-fidelity synthetic faces to improve recognition accuracy. These approaches typically use conditional DMs trained on large-scale facial datasets and generate new faces, by sampling from the learned distribution. Some methods also utilize approaches like multi-level text-related augmentation (Wu et al. 2023c) or disentangled representations (Ding et al. 2023) to enable more fine-grained control over the generated faces, so as to change specific attributes or expressions (Fig. 14).

Fashion Industry: The methods of Li et al. (2023c) and Kong et al. (2023) enable realistic virtual try-on experiences, allowing users to visualize clothing items on different body types and poses. These approaches typically use conditional DMs trained on datasets of clothing images and corresponding body poses, and generate new try-on images, by conditioning on the desired clothing item and pose (Fig. 15).

Agriculture: The methods of Deng and Lu (2023), Muhammad et al. (2023) and Chen et al. (2023a) focus on enhancing datasets for plant disease detection and weed recognition; thereby, improving the accuracy of agricultural models. Such approaches typically employ conditional DMs trained on datasets of plant images with different diseases or weed types, and generate new examples, by sampling from the learned distribution.




Fig. 14 Examples of application-specific DM-based image augmentation (facial recognition and editing), involving different expressions, lighting and poses. Image from Ding et al. (2023)




Fig. 15 Examples of application-specific DM-based image augmentation (fashion industry), using different input images. Image from Li et al. (2023c)

Video Editing and Generation: The methods of Shin et al. (2024) and Wu et al. (2023a) enable high-quality video manipulation and generation from text prompts. These approaches typically rely on the use of 3D or video diffusion models that can generate coherent sequences of video frames, based on textual descriptions or input videos.

Cultural Heritage: The method of Cioni et al. (2023) adopts an LDM to generate diverse, semantically consistent variations of artworks, conditioned on textual descriptions. In particular, the LDM performs a conditional denoising diffusion process in the latent space learned by a convolutional autoencoder, taking the original artwork image and its encoded caption as inputs. By varying the generation seed, the LDM outputs multiple variations of the artwork that preserve the content described in the caption. These synthetic image-caption pairs augment the original art dataset, bridging the domain gap between natural images and artworks, while improving visual grounding of artistic concepts. The augmented dataset, containing both real and synthetic data, facilitates more effective training of downstream vision-language models, like image captioning and cross-modal retrieval ones.

Object Detection: The approaches of Fang et al. (2024), Zhang et al. (2023c), Li et al. (2025), Tang et al. (2024c) and Ma et al. (2024a) leverage DMs to generate synthetic data, significantly enhancing model training and performance. These methods make use of conditional DMs trained on object detection datasets and generate new samples, by conditioning on the desired object categories and bounding boxes.

4.4 Comparison of DM-based methods for image augmentation

Having presented the various methodologies for DM-based image augmentation, this section realizes a detailed and in-depth comparison among the most important methods of each category, in order to provide further critical insights. In particular, Table 3 analyzes several key/important methods, focusing on the following aspects: a) The particular task to be performed, b) If/which FDM is used, c) Input requirements, d) Key exhibited innovation, e) Notable limitations, and f) Fine-tuning requirements. The considered methods, which cover all main categories of the proposed taxonomy (presented in Fig. 2 and Table 2) were selected based on the following criteria: a) High citation impact and community adoption, b) Introduction of novel technical approaches or architectures, and c) Demonstration of significant practical capabilities or limitations that influenced subsequent research works.

From the observation of Table 3, several key insights can be derived. In particular, SD is shown to be the most widely used FDM across multiple tasks, mainly due to its open-source

Table 3 In-depth comparison of key DM-based augmentation methods

Method	Task	FDM	Input requirements	Key innovation	Notable limitations	Fine-tuning requirements
Semantic Manipulation						
Yang et al. (2023a)	I2I Translation	SD	Source image + reference image + mask	Content bottleneck prevents trivial copying	May not preserve fine details	Self-supervised training on OpenImages
Kim et al. (2022)	I2I Translation	–	Image + text prompt	Multi-attribute control via noise combination	Limited to pretrained domains	Reference images for domain finetuning
Lugmayr et al. (2022)	Inpainting	–	Image + binary mask	Resampling strategy for mask-agnostic inpainting	Significantly slower than GAN-based methods	–
Meng et al. (2021)	Editing	–	User strokes or patches	No task-specific training needed	Quality depends on initial noise level	–
Kawar et al. (2023)	Editing	ImageGen SD	Image + text prompt	Linear semantic embedding interpolation	Slow optimization (8min/image)	Target text embedding optimization
Brooks et al. (2023)	Instruction-based Editing	SD v1.5	Image + text instruction	Direct instruction followed without optimization	Cannot handle spatial reasoning	Image-instruction-output triplets training
Avrahami et al. (2022b)	Local Editing	–	Image + mask + text prompt	Spatially blended noise guidance for local editing	30 s generation time per image	–
Chen et al. (2024e)	Object Teleportation	SD 2.1	Target object + scene image + location box	Zero-shot object teleportation with identity preservation	Small detail preservation issues	Light UNet decoder fine-tuning
Huang et al. (2023a)	Multi-condition Generation	–	Multiple optional conditions (text/sketch/depth/etc.)	Composable conditions with exponential control space	Performance tradeoff in single-condition scenarios	60 M image multi-task training
Zhang et al. (2023b)	Image Control	SD 1.5/2.1	Image + text prompt	Zero convolution prevents catastrophic forgetting	Higher GPU memory and training time	Trainable copy with zero convolution
Personalization and Adaptation						
Ruiz et al. (2023)	Personalized Generation	SD	3–5 concept images + text prompt	Concept-specific generation with few shots	Context-appearance entanglement issues	Class-specific prior preservation loss

Table 3 (continued)

Method	Task	FDM	Input requirements	Key innovation	Notable limitations	Fine-tuning requirements
Gal et al. (2022)	Personalized Generation	SD v1.4	3–5 concept images	Single word embedding for concept representation	Lengthy optimization (2 h per concept)	5000 optimization steps with reconstruction loss
Qiu et al. (2023)	Personalized Generation	SD 1.5	5–50 concept images + text prompt	Orthogonal transformation preserving hyperspherical energy	Color and detail distortions	Orthogonal weight transformation with energy preservation
Mokady et al. (2023)	Editing	SD v1	Image + source caption	Null-text optimization without model tuning	Requires editable parts in source caption	–
Zhang et al. (2023e)	Style Transfer	SD	Content image + style image	Single-image style learning via attention	Long training time (20min/image)	–
Wu et al. (2023b)	Style Transfer	SD v1.4	Image + text prompt	Text embedding optimization for attribute editing	Struggles with fine-grained edits	No fine-tuning, only embedding optimization
Bansal and Grover (2023)	Classification	SD 1.5	Class labels or images	Zero-shot data generation for robustness	Limited to ImageNet-like datasets	–
Zhang et al. (2022)	Classification	SD 1.4, DALL-E2	Small-scale seed dataset	Guided imagination for class-consistent generation	Generated samples less informative than real ones	Domain adaptation fine-tuning
Application Specific Augmentation						
Sanchez et al. (2022)	Brain Lesion Detection	–	Multi-channel MRI scans	Counterfactual healthy image generation	Limited to brain tumor scenarios	Healthy/unhealthy classification training
Pinaya et al. (2022)	Brain MRI Generation	–	Age, sex, brain metrics	Anatomically-controlled brain MRI synthesis	Limited to training distribution range	Conditional training on demographic variables
Fang et al. (2024)	Object Detection	SD	Visual prior + text prompt	Category-calibrated CLIP filtering	Performance drops with excessive synthetic data	–

characteristic. Additionally, an important trade-off between processing speed and output quality is observed, where methods requiring optimization or fine-tuning typically produce better results, but at the cost of longer processing times (often ranging from minutes to hours per image). Another notable trend is the evolution from the usage of single- to multi-condition mechanisms, where methods increasingly support various input modalities (i.e., text, images, masks, sketches) for more precise control; however, this increased flexibility is often accompanied with higher computational requirements and more complex training procedures. Investigating category-specific patterns, semantic manipulation methods typically focus on architectural innovations and generally require less fine-tuning, often implementing zero- or few-shot techniques; these methods also emphasize user control through various input modalities. On the other hand, personalization and adaptation methods demonstrate a clear trend towards optimization-heavy approaches, exhibiting notably longer processing times (often at the scale of hours) and more complex fine-tuning requirements, in order to preserve identity and style characteristics. Moreover, application-specific methods are characterized by their specialized nature, requiring domain-specific training data and often implementing custom architectures or loss functions tailored to their selected application fields; these methods typically prioritize task-specific performance metrics over general-purpose applicability.

5 Evaluation metrics

Evaluating the performance of DM-powered image augmentation methods is crucial to understanding their impact on visual image analysis tasks. This section outlines the methodologies used to assess both the efficacy and efficiency of such frameworks, while a detailed quantitative comparison of DM-based and traditional image augmentation methods is also provided.

5.1 Quantitative evaluation

Quantitative evaluation involves the measurement of the numerical improvements in model performance, as well as the perceptual quality and diversity of the augmented images. This type of evaluation provides objective metrics that can be directly compared and analyzed.

5.1.1 Improvement in model performance

The primary quantitative evaluation involves the measurement of the (potential) improvement in downstream task performance metrics of learned models (e.g., classifiers), trained while incorporating augmented images that have been generated by the DM. The more the downstream model's performance increase, compared to the baseline of realizing training using only the original non-augmented dataset, the better the DM that conducted the augmentation is considered to be. The following cases are most commonly met:

- **Classification Tasks:** Metrics such as accuracy, precision, recall, and F1-score are used to evaluate the performance of classification models. These metrics provide a comprehensive overview regarding how well the model can distinguish between different

classes (Trabucco et al. 2023; Chowdary et al. 2023; Packhäuser et al. 2023; Sagers et al. 2022; Bansal and Grover 2023).

- Segmentation Tasks: Metrics such as Intersection over Union (IoU) and Dice coefficient are used for evaluating segmentation models. IoU measures the overlap between the predicted segmentation and the ground truth one, while the Dice coefficient assesses the similarity between these two (Xie et al. 2023a; Schnell et al. 2024; Sanchez et al. 2022; Valvano et al. 2024).
- Object Detection Tasks: Metrics such as mean Average Precision (mAP), precision and recall are crucial for evaluating object detection models. mAP calculates the average precision across different recall levels and is a standard metric for assessing the accuracy of object detectors. On the other hand, precision and recall provide insights into the detector's ability to correctly identify objects and to minimize false positives (Fang et al. 2024; Zhang et al. 2023c). Analyzing these metrics makes it easier to determine the degree of improvement in the downstream DNN's predictive capabilities, which is introduced by training set augmentation.

5.1.2 Perceptual quality and diversity

To evaluate the quality and diversity of augmented images, the following metrics are most commonly used:

- Fréchet Inception Distance (FID): FID (Heusel et al. 2017) measures the distance between the distribution of the generated images and the real ones in the feature space of a pretrained Inception network (Szegedy et al. 2016). It captures both the quality and the diversity of the generated images. Lower FID scores indicate higher quality and more diverse generated images (Esser et al. 2024; Ho et al. 2020; Rombach et al. 2022; Xie et al. 2023b; Avrahami et al. 2023b; Gandikota et al. 2024; Zhang et al. 2023b; Pinaya et al. 2022; Couairon et al. 2022; Fu et al. 2024).
- Inception Score (IS): IS (Salimans et al. 2016) evaluates the quality of the generated images based on the confidence of the class predictions of a pretrained Inception model. High IS values imply that the generated images are diverse and that each image is recognized with high confidence as belonging to a specific class (Blattmann et al. 2022; Gafni and Wolf 2020; Luo et al. 2023; Chen et al. 2023a).
- Kernel Inception Distance (KID): KID (Bińkowski et al. 2018) comprises an alternative to the case of FID, while it estimates the maximum mean discrepancy between the feature representations of the generated and the real image samples. Lower KID scores indicate better quality and similarity to real images (Muhammad et al. 2023; Li et al. 2023c; Kumari et al. 2023b; Meng et al. 2021).
- Perceptual Metrics: Metrics such as Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al. 2018) and Structural Similarity Index Measure (SSIM) (Wang et al. 2004) are used to assess the perceptual similarity between generated and real samples, or between the input and the output of image editing tasks (Kulikov et al. 2023; Wang et al. 2022b; Li et al. 2023b; Zhang et al. 2024d; Qiu et al. 2023; Chambon et al. 2022b; Xu et al. 2024). Often estimating and combining all above metrics together provides a comprehensive overview of the augmented images' quality and variability.

5.2 Qualitative evaluation

Qualitative evaluation involves subjective assessments of visual realism and relevance of the augmented images. Such types of evaluations are crucial for ensuring that the synthetic data generated by DMs is not only technically sound, but also perceptually convincing and contextually appropriate.

5.2.1 Visual quality of augmented images

Assessing the visual quality of augmented images involves inspection by experts, aiming to determine how closely the synthetic images resemble real ones. Expert evaluators are requested to analyze the following main, among others, aspects factors:

- **Realism:** Evaluation of whether the augmented images are indistinguishable from real ones, focusing on characteristics such as texture, lighting, and color consistency (Sanchez and Tsaftaris 2022; Kwon and Ye 2022; Dong et al. 2022).
- **Detail Preservation:** It is essential to ensure that the augmented images retain critical details necessary for the task, such as fine-grained textures and structural integrity (Ruiz et al. 2023; Kwar et al. 2023; Dong et al. 2022).
- **Editing Consistency:** Experts assess how closely the augmented images maintain semantic coherence with the original input and the editing instructions. This involves the evaluation of whether the edits are applied accurately to the intended regions, while preserving the overall context and structure of the image (Wang et al. 2023a; Kwar et al. 2023; Tang et al. 2024b; Li et al. 2023b). Overall, the above described subjective assessment metrics facilitate towards validating the visual authenticity and usability of the generated images.

5.2.2 Relevance and contextual appropriateness

Evaluating the relevance and contextual appropriateness of augmented images ensures that they maintain semantic coherence and they are suitable for the specific application at hand. This involves the following main, among others, aspects:

- **Contextual Consistency:** Experts review whether the augmented images fit well within the expected context of the task. For example, in medical imaging, synthetic images should accurately reflect the characteristics of the examined disease or the condition being modeled (Gandikota et al. 2023; Han et al. 2023; Ruiz et al. 2024; Wang et al. 2023a; Zhang et al. 2023e; Tewel et al. 2023).
- **Semantic Accuracy:** The augmented images should convey the correct semantic information, avoiding any misleading or nonsensical variations. This is critical for applications where accurate representation of objects and scenes is of paramount importance (Huang et al. 2023b; Li et al. 2023a).
- **Task-Specific Features:** The relevance of augmented images is also evaluated based on their utility for the specific task at hand. For instance, in object detection, the images should contain objects that are correctly annotated and positioned, in order to enhance efficient model training (Chambon et al. 2022a; Valvano et al. 2024). The above de-

scribed qualitative assessment ensures that the augmented images not only look realistic, but also serve their intended purpose in a satisfactory way.

5.3 Quantitative evaluation of DM-based and traditional image augmentation methods

Having presented the various quantitative and qualitative evaluation metrics, this subsection realizes a comprehensive quantitative comparison among DM-based and traditional image augmentation methods. In particular, Table 4 reports the performance merits of traditional (namely the approaches of Kang and Kim (2023); Yun et al. (2019); Rao et al. (2021); Cubuk et al. (2020); Kim et al. (2021); Dwibedi et al. (2017); Zhang et al. (2017); DeVries and Taylor (2017)) and key DM-based augmentation methods, applied separately to common downstream tasks, namely image classification, object detection and semantic segmentation. The table illustrates the reported improvement that each DM-based augmentation method (denoted DM-AUG) induces, compared to the considered traditional one (denoted BASIC-AUG).

From the observation of Table 4, it can be seen that almost all DM-based augmentation methods demonstrate performance improvements across a wide range of tasks and datasets, over conventional augmentation approaches. The improvement level varies across different tasks, exhibiting substantial (or outstanding) gains in some cases and more modest enhancements in others. It needs to be highlighted though that this increase in performance is also associated with correspondingly significant computational overhead for DM-based methods. In particular, DM-powered generation time may vary considerably across methods and hardware setups, ranging between 0.43 and 6.6 s per image using high-performing GPUs (e.g., H800, V100, RTX3090, etc.), while traditional methods (e.g., Cutout, GridMask, etc.) operate in approximately 0.008 s per image (Zhang et al. 2022; Fu et al. 2024). Indicatively, the total generation time of DM-based methods for whole datasets may vary for relatively small datasets from 2 h (e.g., Waterbirds (Sagawa et al. 2019)) up to 7 h (e.g., iWildCam (Koh et al. 2021)), while certain methods may require up to 10 h to generate 100K images (Dunlap et al. 2024; Du et al. 2024). Apart from computational cost, DM-powered methods also often require additional storage space for maintaining the generated images, unlike traditional augmentation approaches that can be performed on-the-fly (Islam et al. 2024b). Moreover, prompt quality dependency is usually a common bottleneck that requires proper construction and careful selection; poor prompt selection could lead to unrealistic or unsuitable generated images that may degrade performance (Islam et al. 2024a, b).

6 Challenges and future research directions

Image augmentation using DM-powered methodologies constitutes a promising approach for enhancing the diversity and quality of training datasets. However, despite the rapid progress and significant advances that have been observed in the field in the recent period, several open challenges, which at the same time comprise future research direction, remain to be addressed. These can broadly be divided into general ones, which concern the application of DMs in general and regardless of the particular application case, and challenges that are particularly important for image augmentation.

Table 4 Quantitative comparison of DM-based and traditional image augmentation methods. ('Std aug' refers to standard simple geometric transformations, such as rotation, crop, etc)

Dataset	BASIC-AUG method	DM method	BA-SIC-AUG (%)	DM-AUG (%)	Gain (%)
Classification-Top-1 Accuracy					
<i>Aircraft</i> (Maji et al. 2013)	Guided-SR (Kang and Kim 2023)	Chen et al. (2024c)	77.38	84.79	+ 9.57
	Cutmix (Yun et al. 2019)	Wang et al. (2024b)	89.44	90.25	+ 0.90
	CAL-AUG (Rao et al. 2021)	Michaeli and (Fried 2024)	81.90	87.40	+ 6.71
<i>Caltech101</i> (Fei-Fei et al. 2004)	GuidedAP (Kang and Kim 2023)	Islam et al. (2024b)	84.32	85.84	+ 1.80
	GuidedAP (Kang and Kim 2023)	Islam et al. (2024a)	84.32	85.76	+ 1.70
	RandAugment (Cubuk et al. 2020)	Li et al. (2023a)	58.55	59.17	+ 1.05
	RandAugment (Cubuk et al. 2020)	Fu et al. (2024)	60.10	85.30	+ 41.93
	RandAugment (Cubuk et al. 2020)	Zhang et al. (2022)	57.80	65.10	+ 12.61
<i>Cars</i> (Krause et al. 2013)	RandAugment (Cubuk et al. 2020)	Li et al. (2023a)	86.55	88.53	+ 2.28
	RandAugment (Cubuk et al. 2020)	Fu et al. (2024)	65.80	87.10	+ 32.37
	RandAugment (Cubuk et al. 2020)	Zhang et al. (2022)	43.20	75.70	+ 75.23
	Guided-SR (Kang and Kim 2023)	Chen et al. (2024c)	91.01	93.04	+ 2.23
	Cutmix (Yun et al. 2019)	Wang et al. (2024b)	94.73	95.21	+ 0.50
	RandAugment (Cubuk et al. 2020) + Cutmix (Yun et al. 2019)	Michaeli and (Fried 2024)	92.70	93.80	+ 1.18
	GuidedAP (Kang and Kim 2023)	Islam et al. (2024b)	90.27	91.30	+ 1.14
	GuidedAP (Kang and Kim 2023)	Islam et al. (2024a)	90.27	91.26	+ 1.09
<i>CIFAR-100</i> (Krizhevsky et al. 2009)	GuidedAP (Kang and Kim 2023)	Islam et al. (2024a)	81.20	82.50	+ 1.70
	Cutmix (Yun et al. 2019)	Li et al. (2023a)	77.56	75.75	- 2.37
	GuidedMixup (Kang and Kim 2023)	Islam et al. (2024b)	81.20	82.58	+ 1.69
<i>Flowers</i> (Nilsback and Zisserman 2008)	RandAugment (Cubuk et al. 2020)	Li et al. (2023a)	41.97	45.61	+ 8.67
	Cutmix (Yun et al. 2019)	Zhang et al. (2022)	83.80	88.30	+ 5.36
	GuidedMixup (Kang and Kim 2023)	Wang et al. (2024b)	99.40	99.54	+ 0.14
<i>iWildCam</i> (Koh et al. 2021)	RandAugment (Cubuk et al. 2020)	Rahat et al. (2024)	76.78	85.37	+ 11.18
	Cutmix (Yun et al. 2019)	Dunlap et al. (2023)	77.56	84.87	+ 9.42
<i>CUB</i> (Wah et al. 2011)	Cutmix (Yun et al. 2019)	Rahat et al. (2024)	70.48	73.16	+ 3.80
	Co-Mixup (Kim et al. 2021)	Chen et al. (2024c)	79.41	80.82	+ 1.77
	Cutmix (Yun et al. 2019)	Wang et al. (2024b)	87.23	87.56	+ 0.37
<i>Pets</i> (Parkhi et al. 2012)	RandAugment (Cubuk et al. 2020)	Li et al. (2023a)	57.45	73.71	+ 28.39
	RandAugment (Cubuk et al. 2020)	Fu et al. (2024)	61.50	86.50	+ 40.65
	RandAugment (Cubuk et al. 2020)	Zhang et al. (2022)	48.00	73.40	+ 52.91
Few-shot Classification (8 examples per class)-Top-1 Accuracy					
<i>MS COCO</i> (Lin et al. 2014)	Std aug	Trabucco et al. (2023)	42.00	47.00	+ 11.90
	Std aug	Lingenberg et al. (2024)	47.00	57.00	+ 21.67
Object Detection - mAP					
<i>PASCAL VOC</i> (Everingham et al. 2010)	RandAugment (Cubuk et al. 2020)	Li et al. (2025)	78.20	79.10	+ 1.15
<i>DIOR-R</i> (Cheng et al. 2022)	CopyPaste + Flip (Dwibedi et al. 2017)	Tang et al. (2024c)	38.75	41.69	+ 7.58

Table 4 (continued)

Dataset	BASIC-AUG method	DM method	BA-SIC-AUG (%)	DM-AUG (%)	Gain (%)
Out-of-Distribution Classification - OOD Accuracy					
(Sagawa et al. 2019)	RandAugment (Cubuk et al. 2020)	Dunlap et al. (2023)	30.32	46.63	+ 53.79
	Mixup (Zhang et al. 2017)	Chen et al. (2024c)	72.52	76.16	+ 5.03
	Cutmix (Yun et al. 2019)	Wang et al. (2024b)	71.23	72.47	+ 1.74
Segmentation - mIoU					
<i>KITTI</i> (Fritsch et al. 2013)	Cutout (DeVries and Taylor 2017)	Ma et al. (2024a)	78.65	82.20	+ 4.51
	Cutout (DeVries and Taylor 2017)	Ma et al. (2024a)	77.65	78.34	+ 0.88
	Cutout (DeVries and Taylor 2017)	Ma et al. (2024a)	92.35	92.72	+ 0.40
<i>PASCAL VOC</i> (Everingham et al. 2010)	Std aug	Schnell et al. (2024)	78.10	78.90	+ 1.02
Medical Image Classification-Top-1 Accuracy					
<i>OrganMNIST</i> (Yang et al. 2021)	RandAugment (Cubuk et al. 2020)	Zhang et al. (2022)	79.60	80.70	+ 1.38
<i>PathMNIST</i> (Yang et al. 2021)	RandAugment (Cubuk et al. 2020)	Zhang et al. (2022)	79.20	86.90	+ 9.72
<i>BreastMNIST</i> (Yang et al. 2021)	RandAugment (Cubuk et al. 2020)	Zhang et al. (2022)	68.70	77.40	+ 12.66
<i>Shenzhen TB</i> (Jaeger et al. 2014)	RandAugment (Cubuk et al. 2020)	Fu et al. (2024)	75.50	83.50	+ 10.59

6.1 Computational cost and efficiency

A significant factor associated with the application of any DM-based architecture concerns its very high computational cost. In particular, DM require substantial computational resources and can be time-consuming for both training and inference, which may delay their development and deployment. On the contrary, a trained Generative Adversarial Network (GAN) Generator synthesizes an output image directly, i.e., with just one forward pass. The iterative denoising process of DMs renders it relatively difficult to scale such methods to real-world applications, large datasets and complex tasks. For example, SD requires more than 200,000 GPU hours (Rombach et al. 2022) to be trained on the LAION-5B dataset with one A100 40GB GPU. Approaches such as DDIM and DPM-Solver++ (Lu et al. 2022b) target to speed up the sampling process and to improve efficiency.

Given this context, a significant avenue for future research comprises the development of more efficient DM architectures and optimization methods. Promising directions include knowledge distillation approaches (Luo 2023) to create lightweight models (Song et al. 2024), quantization (Shang et al. 2023) or advanced parallelization strategies (Shih et al. 2024). Recent work on progressive distillation (Salimans and Ho 2022) and one-step generation (Yin et al. 2024) shows potentials for dramatically reducing inference time. Moreover, research into specialized acceleration (Ma et al. 2024b; Wang et al. 2024a) could facilitate towards bridging the efficiency gap between DMs and other generative approaches.

6.2 Lack of fine-grained control and interpretability

Interpretability and controllability of DMs is generally problematic, making it challenging to understand how they generate their outputs (e.g., given specific conditions). The users have limited ability to precisely control specific attributes, objects, or regions in the generated images, which can hinder the usefulness of these methods for certain applications.

As a result, it is important for future research to explore advanced conditioning mechanisms and interpretability tools. Research directions may include developing more sophisticated attention-based control mechanisms (Hertz et al. 2022) and creating interactive interfaces for fine-grained manipulation (Lee et al. 2024).

6.3 Limited diversity and realism of generated data

One of the most prevalent challenges is the (relatively) constrained diversity and realism of the synthetic images generated by DMs. While DMs have shown impressive capabilities in generating high-quality images, they often struggle to capture the full diversity and complexity of real-world data distributions. This limitation can lead to a domain gap between the synthetic and real data, which can hinder the effectiveness of using the generated data for training downstream tasks, like image classification and object detection.

Several methods have been introduced so far that aim to address this issue by incorporating approaches that improve the diversity and realism of the generated samples, such as using language enhancement (Dunlap et al. 2022) and post-filtering (e.g., CLIP-based filtering (Kim et al. 2022)). Further relevant research towards the direction of domain adaptation (Chopra et al. 2024) or style transfer (Kwon and Ye 2022) may facilitate towards bridging the gap between real and synthetic image distributions.

6.4 Model overfitting and catastrophic forgetting

One of the most common challenges faced by many of the methods covered in this study comprises the issue of model overfitting and catastrophic forgetting (Cywiński et al. 2024). In particular, overfitting occurs when a model learns to fit the training data too closely, at the expense of generalization to new, unseen data. This can be especially problematic when working with limited training data, as is often the case in domains like medical imaging. On the other hand, catastrophic forgetting refers to the tendency of trainable models to forget previously learned representations, when being finetuned on new data or tasks. This is a major hurdle for approaches that aim to adapt pretrained models to specific domains or styles. Directly finetuning large models on limited datasets can lead to a rapid loss of the originally learned knowledge structures.

Recent approaches that target to mitigate this Zeng et al. (2024); Zhong et al. (2024); Zając et al. (2023) leverage various methodologies, including selective parameter updates, dual-stream architectures, and careful management of the denoising process, to maintain model generalization while adapting to new tasks or concepts. Future research along the promising directions of relevant continual learning (Smith et al. 2023) or metalearning (Zhang et al. 2024a) approaches may bear the potentials for further alleviating any remaining issues.

6.5 Evaluation metrics and benchmarks

Evaluating the effectiveness of DM-powered image augmentation remains a challenge itself. Traditional metrics like FID and KID may not fully assess/capture the quality and diversity of the generated samples, especially for complex or specialized domains. Additionally, the lack of standardized benchmarks and evaluation protocols makes it difficult to compare different methods and to assess their generalization and augmentation abilities.

To this end, developing more comprehensive and both general as well as domain-specific evaluation metrics, while also establishing common benchmarks and datasets, comprises an important area for future research for enabling more rigorous and consistent evaluation (Betzalet et al. 2024).

6.6 Ethical considerations and bias

Large-scale T2I DMs are often trained on Web-scraped datasets, which can contain harmful stereotypes, offensive content, and biases related to gender, race, age, and other sensitive attributes. These biases can get amplified in the generated images, leading to unfair representations and perpetuating societal stereotypes. An additional ethical consideration is the use of copyrighted or private data in the training datasets without proper consent or attribution. This raises significant concerns regarding the ownership and fair use of the generated images, as well as the potential for privacy violations.

Such concerns may partially be mitigated by technical measures, which, however, would require research on detecting and mitigating biases in both training and generated images. Similarly, creating fairness-aware training procedures (Friedrich et al. 2023) and establishing ethical guidelines for data collection and model deployment (Zhang et al. 2023a) are potentially worthy avenues for investigation, towards more transparent and trustworthy DMs.

7 Conclusion

Image augmentation comprises a fundamental task in modern computer vision, since it allows the enhancement of training datasets with realistic synthetic samples, permitting the automatic context- and semantics-aware editing of given reference images, etc. Diffusion Models (DMs) have shown significant promise in generating realistic and diverse images, capturing complex relationships and structures in high-dimensional image data. Moreover, the ability to condition the generation process using class labels, textual descriptions, or visual prompts allows for targeted augmentation, generating images that fulfill specific requirements based on the task at hand.

This study provided a comprehensive overview of the recent advancements in DM-powered image augmentation, a taxonomy of the main categories of existing methods, insights regarding the practical usage of DM-powered techniques for semantic manipulation, personalization and adaptation, and application-specific image augmentation, and a review of the relevant performance evaluation metrics. Promising and significant avenues for future research include the improvement of the efficiency of DMs, i.e., reducing computational costs and improving scalability, enhancing interpretability and control over generated

images, as well as boosting diversity and realism of synthetic data. Finally, developing new, robust evaluation metrics and addressing ethical considerations emerges as a critical key to the progress and responsible deployment of DMs in image augmentation.

Acknowledgements The research leading to the results of this paper has received funding from the European Union's Horizon Europe research and development programme under grant agreement No. 101073876 (Ceasefire).

Author contributions Panagiotis Alimisis, Ioannis Mademlis and Georgios Th. Papadopoulos performed the literature review and prepared a draft of the manuscript. Panagiotis Radoglou-Grammatikis and Panagiotis Sargiannidis performed reviewing and editing of the manuscript. Georgios Th. Papadopoulos was responsible for receiving the funding to implement the study.

Funding The research leading to these results received funding from the European Commission under Grant Agreement No. 101073876 (Ceasefire).

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors have no Conflict of interest to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ackermann J, Li M (2022) High-resolution image editing via multi-stage blended diffusion. arXiv preprint [arXiv:2210.12965](https://arxiv.org/abs/2210.12965)
- Agustsson E, Mentzer F, Tschannen M, et al. (2017) Soft-to-hard vector quantization for end-to-end learning compressible representations. *Advances in neural information processing systems* 30
- Akrouf M, Gyepesi B, Holló P, et al. (2023) Diffusion-based data augmentation for skin disease classification: impact across original medical datasets to fully synthetic images. In: *International conference on medical image computing and computer-assisted intervention*, Springer, pp 99–109
- Ali H, Murad S, Shah Z (2022) Spot the fake lungs: generating synthetic medical images using neural diffusion models. In: *Irish conference on artificial intelligence and cognitive science*, Springer, pp 32–39
- Arkipkin V, Filatov A, Vasilev V, et al. (2023) Kandinsky 3.0 technical report. [arXiv:2312.03511](https://arxiv.org/abs/2312.03511)
- Asperti A, Evangelista D, Marro S et al. (2023) Image embedding for denoising generative models. *Artif Intell Rev* 56(12):14511–14533
- Avrahami O, Lischinski D, Fried O (2022a) Blended diffusion for text-driven editing of natural images. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 18208–18218
- Avrahami O, Lischinski D, Fried O (2022b) Blended diffusion for text-driven editing of natural images. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pp 18208–18218
- Avrahami O, Fried O, Lischinski D (2023a) Blended latent diffusion. *ACM Trans Graph (TOG)* 42(4):1–11

- Avrahami O, Hayes T, Gafni O, et al. (2023b) Spatext: spatio-textual representation for controllable image generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 18370–18380
- Balaji Y, Nah S, Huang X, et al. (2022) ediff-i: text-to-image diffusion models with an ensemble of expert denoisers. arXiv preprint [arXiv:2211.01324](https://arxiv.org/abs/2211.01324)
- Bansal H, Grover A (2023) Leaving reality to imagination: Robust classification via generated datasets. arXiv preprint [arXiv:2302.02503](https://arxiv.org/abs/2302.02503)
- Bansal A, Chu HM, Schwarzschild A, et al. (2023) Universal guidance for diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 843–852
- Betker J, Goh G, Jing L et al. (2023) Improving image generation with better captions. *Comput Sci* 2(3):8
- Betzalel E, Penso C, Fetaya E (2024) Evaluation metrics for generative models: an empirical study. *Mach Learn Knowl Extract* 6(3):1531–1544
- Bińkowski M, Sutherland DJ, Arbel M, et al. (2018) Demystifying mmd gans. arXiv preprint [arXiv:1801.01401](https://arxiv.org/abs/1801.01401)
- Birhane A, Prabhu VU (2021) Large image datasets: a pyrrhic win for computer vision? In: 2021 IEEE winter conference on applications of computer vision (WACV), IEEE, pp 1536–1546
- Blattmann A, Rombach R, Oktay K, et al. (2022) Semi-parametric neural image synthesis. arXiv preprint [arXiv:2204.11824](https://arxiv.org/abs/2204.11824)
- Boutros F, Grebe JH, Kuijper A, et al. (2023) Idiff-face: synthetic-based face recognition through fizzy identity-conditioned diffusion model. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 19650–19661
- Brack M, Schramowski P, Friedrich F, et al. (2022) The stable artist: steering semantics in diffusion latent space. arXiv preprint [arXiv:2212.06013](https://arxiv.org/abs/2212.06013)
- Brooks T, Holynski A, Efros AA (2023) Instructpix2pix: learning to follow image editing instructions. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 18392–18402
- Cao M, Wang X, Qi Z, et al. (2023) Masactrl: tuning-free mutual self-attention control for consistent image synthesis and editing. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 22560–22570
- Cao H, Tan C, Gao Z, et al. (2024) A survey on generative diffusion models. *IEEE transactions on knowledge and data engineering*
- Chai J, Zeng H, Li A et al. (2021) Deep learning in computer vision: a critical review of emerging techniques and application scenarios. *Mach Learn Appl* 6:100134
- Chambon P, Bluethgen C, Delbrouck JB, et al. (2022a) Roentgen: vision-language foundation model for chest x-ray generation. arXiv preprint [arXiv:2211.12737](https://arxiv.org/abs/2211.12737)
- Chambon P, Bluethgen C, Langlotz CP, et al. (2022b) Adapting pretrained vision-language foundational models to medical imaging domains. arXiv preprint [arXiv:2210.04133](https://arxiv.org/abs/2210.04133)
- Chen W, Hu H, Saharia C, et al. (2022a) Re-Imagen: retrieval-augmented text-to-image generator. arXiv preprint [arXiv:2209.14491](https://arxiv.org/abs/2209.14491)
- Chen X, Wang X, Changpinyo S, et al. (2022b) Pali: a jointly-scaled multilingual language-image model. arXiv preprint [arXiv:2209.06794](https://arxiv.org/abs/2209.06794)
- Chen D, Qi X, Zheng Y, et al. (2023a) Deep data augmentation for weed recognition enhancement: a diffusion probabilistic model and transfer learning based approach. In: 2023 ASABE annual international meeting, American society of agricultural and biological engineers, p 1
- Chen H, Zhang Y, Wu S, et al. (2023b) Disenbooth: identity-preserving disentangled tuning for subject-driven text-to-image generation. arXiv preprint [arXiv:2305.03374](https://arxiv.org/abs/2305.03374)
- Chen J, Yu J, Ge C, et al. (2023c) Pixart- α : fast training of diffusion transformer for photorealistic text-to-image synthesis. arXiv preprint [arXiv:2310.00426](https://arxiv.org/abs/2310.00426)
- Chen K, Xie E, Chen Z, et al. (2023d) Geodiffusion: text-prompted geometric control for object detection data generation. arXiv preprint [arXiv:2306.04607](https://arxiv.org/abs/2306.04607)
- Chen J, Wu Y, Luo S, et al. (2024a) Pixart- $\{\delta\}$: fast and controllable image generation with latent consistency models. arXiv preprint [arXiv:2401.05252](https://arxiv.org/abs/2401.05252)
- Chen M, Laina I, Vedaldi A (2024b) Training-free layout control with cross-attention guidance. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision, pp 5343–5353
- Chen R, Wang Z, Zhang KY, et al. (2024c) Decoupled data augmentation for improving image classification. arXiv preprint [arXiv:2411.02592](https://arxiv.org/abs/2411.02592)<https://api.semanticscholar.org/CorpusID:273821653>
- Chen W, Hu H, Li Y, et al. (2024d) Subject-driven text-to-image generation via apprenticeship learning. *Adv Neural Inform Proc Syst* 36
- Chen X, Huang L, Liu Y, et al. (2024e) Anydoor: zero-shot object-level image customization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 6593–6602
- Cheng G, Wang J, Li K et al. (2022) Anchor-free oriented proposal generator for object detection. *IEEE Trans Geosci Remote Sens* 60:1–11

- Cherti M, Beaumont R, Wightman R, et al. (2023) Reproducible scaling laws for contrastive language-image learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 2818–2829
- Chopra S, Kothawade S, Aynaou H, et al. (2024) Source-free domain adaptation with diffusion-guided source data generation. arXiv preprint [arXiv:2402.04929](https://arxiv.org/abs/2402.04929)
- Chowdary PN, Vardhan GV, Akshay MS, et al. (2023) Enhancing knee osteoarthritis severity level classification using diffusion augmented images. arXiv preprint [arXiv:2309.09328](https://arxiv.org/abs/2309.09328)
- Chung H, Lee ES, Ye JC (2022) Mr image denoising and super-resolution using regularized reverse diffusion. IEEE Trans Med Imaging 42(4):922–934
- Cioni D, Berlincioni L, Becattini F, et al. (2023) Diffusion based augmentation for captioning and retrieval in cultural heritage. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 1707–1716
- Couairon G, Verbeek J, Schwenk H, et al. (2022) Diffedit: diffusion-based semantic image editing with mask guidance. arXiv preprint [arXiv:2210.11427](https://arxiv.org/abs/2210.11427)
- Croitoru FA, Hondru V, Ionescu RT, et al. (2023) Diffusion models in vision: a survey. IEEE Transactions on pattern analysis and machine intelligence
- Cubuk ED, Zoph B, Shlens J, et al. (2020) Randaugment: practical automated data augmentation with a reduced search space. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops, pp 702–703
- Cywiński B, Deja K, Trzeciński T, et al. (2024) Guide: guidance-based incremental learning with diffusion models. ArXiv abs/2403.03938. <https://api.semanticscholar.org/CorpusID:268253681>
- Deng B, Lu Y (2023) Stable diffusion for data augmentation in coco and weed datasets. arXiv preprint [arXiv:2312.03996](https://arxiv.org/abs/2312.03996)
- DeVries T, Taylor GW (2017) Improved regularization of convolutional neural networks with cutout. arXiv preprint [arXiv:1708.04552](https://arxiv.org/abs/1708.04552)
- Dhariwal P, Nichol A (2021) Diffusion models beat gans on image synthesis. Adv Neural Inform Proc Systms 34:8780–8794
- Ding Z, Zhang X, Xia Z, et al. (2023) Diffusionrig: learning personalized priors for facial appearance editing. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 12736–12746
- Dong Z, Wei P, Lin L (2022) Dreamartist: towards controllable one-shot text-to-image generation via positive-negative prompt-tuning. arXiv preprint [arXiv:2211.11337](https://arxiv.org/abs/2211.11337)
- Dosovitskiy A, Beyer L, Kolesnikov A, et al. (2020) An image is worth 16x16 words: transformers for image recognition at scale. arXiv preprint [arXiv:2010.11929](https://arxiv.org/abs/2010.11929)
- Du X, Sun Y, Zhu J, et al. (2024) Dream the impossible: outlier imagination with diffusion models. Adv Neural Inform Proc Syst 36
- Dunlap L, Mohri C, Guillory D, et al. (2022) Using language to extend to unseen domains. In: The eleventh international conference on learning representations
- Dunlap L, Umino A, Zhang H, et al. (2023) Diversify your vision datasets with automatic diffusion-based augmentation. Conference on neural information processing systems
- Dunlap L, Umino A, Zhang H, et al. (2024) Diversify your vision datasets with automatic diffusion-based augmentation. Adv Neural Inform Proc Syst 36
- Dwibedi D, Misra I, Hebert M (2017) Cut, paste and learn: surprisingly easy synthesis for instance detection. In: Proceedings of the IEEE international conference on computer vision, pp 1301–1310
- Esser P, Kulal S, Blattmann A, et al. (2024) Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first international conference on machine learning
- Everingham M, Van Gool L, Williams CK et al. (2010) The pascal visual object classes (voc) challenge. Int J Compu Vision 88:303–338
- Fang H, Han B, Zhang S, et al. (2024) Data augmentation for object detection via controllable diffusion models. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision, pp 1257–1266
- Fei Z, Fan M, Yu C, et al. (2024) Scaling diffusion transformers to 16 billion parameters. arXiv preprint [arXiv:2407.11633](https://arxiv.org/abs/2407.11633)
- Fei-Fei L, Fergus R, Perona P (2004) Learning generative visual models from few training examples: an incremental Bayesian approach tested on 101 object categories. In: 2004 Conference on computer vision and pattern recognition workshop, IEEE, pp 178–178
- Friedrich F, Brack M, Struppek L, et al. (2023) Fair diffusion: instructing text-to-image generation models on fairness. arXiv preprint [arXiv:2302.10893](https://arxiv.org/abs/2302.10893)
- Fritsch J, Kuehnl T, Geiger A (2013) A new performance measure and evaluation benchmark for road detection algorithms. In: 16th International IEEE conference on intelligent transportation systems (ITSC 2013), IEEE, pp 1693–1700

- Fu Y, Chen C, Qiao Y, et al. (2024) Dreamda: generative data augmentation with diffusion models. arXiv preprint [arXiv:2403.12803](https://arxiv.org/abs/2403.12803)
- Gafni O, Wolf L (2020) Wish you were here: context-aware human generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 7840–7849
- Gal R, Alaluf Y, Atzmon Y, et al. (2022) An image is worth one word: personalizing text-to-image generation using textual inversion. arXiv preprint [arXiv:2208.01618](https://arxiv.org/abs/2208.01618)
- Gal R, Arar M, Atzmon Y et al. (2023) Encoder-based domain tuning for fast personalization of text-to-image models. *ACM Trans Graphics (TOG)* 42(4):1–13
- Gandikota R, Materzynska J, Fiotto-Kaufman J, et al. (2023) Erasing concepts from diffusion models. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 2426–2436
- Gandikota R, Orgad H, Belinkov Y, et al. (2024) Unified concept editing in diffusion models. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision, pp 5111–5120
- Geng Z, Yang B, Hang T, et al. (2024) Instructdiffusion: a generalist modeling interface for vision tasks. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition, pp 12709–12720
- Guo X, Yang Y, Ye C, et al. (2023) Accelerating diffusion models via pre-segmentation diffusion sampling for medical image segmentation. In: 2023 IEEE 20th international symposium on biomedical imaging (ISBI), IEEE, pp 1–5
- Han I, Yang S, Kwon T, et al. (2023) Highly personalized text embedding for image manipulation by stable diffusion. arXiv preprint [arXiv:2303.08767](https://arxiv.org/abs/2303.08767)
- Hatamizadeh A, Song J, Liu G, et al. (2025) Diffit: diffusion vision transformers for image generation. In: European conference on computer vision, Springer, pp 37–55
- Hemati S, Beitollahi M, Estiri AH, et al. (2023) Cross domain generative augmentation: domain generalization with latent diffusion models. arXiv preprint [arXiv:2312.05387](https://arxiv.org/abs/2312.05387)
- Heng A, Soh H (2024) Selective amnesia: a continual learning approach to forgetting in deep generative models. *Adv Neural Inform Proc Syst* 36
- Hertz A, Mokady R, Tenenbaum J, et al. (2022) Prompt-to-prompt image editing with cross attention control. arXiv preprint [arXiv:2208.01626](https://arxiv.org/abs/2208.01626)
- Heusel M, Ramsauer H, Unterthiner T, et al. (2017) Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Adv Neural Inform Proc Syst* 30
- Ho J, Salimans T (2022) Classifier-free diffusion guidance. arXiv preprint [arXiv:2207.12578](https://arxiv.org/abs/2207.12578)
- Ho J, Jain A, Abbeel P (2020) Denoising diffusion probabilistic models. *Adv Neural Inform Proc Syst* 33:6840–6851
- Hu EJ, Shen Y, Wallis P, et al. (2021) Lora: low-rank adaptation of large language models. arXiv preprint [arXiv:2106.09685](https://arxiv.org/abs/2106.09685)
- Hu D, Tao YK, Oguz I (2022) Unsupervised denoising of retinal oct with diffusion probabilistic model. In: Medical imaging 2022: Image processing, SPIE, pp 25–34
- Hu H, Chan KC, Su YC, et al. (2024) Instruct-imagen: image generation with multi-modal instruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 4754–4763
- Huang L, Chen D, Liu Y, et al. (2023a) Composer: creative and controllable image synthesis with composable conditions. arXiv preprint [arXiv:2302.09778](https://arxiv.org/abs/2302.09778)
- Huang N, Tang F, Dong W, et al. (2023b) Region-aware diffusion for zero-shot text-driven image editing. arXiv preprint [arXiv:2302.11797](https://arxiv.org/abs/2302.11797)
- Huang Z, Geng H, Wang H, et al. (2024) Data augmentation for facial recognition with diffusion model. In: CVPR 2024 workshop SyntaGen: harnessing generative models for synthetic visual datasets, <https://openreview.net/forum?id=GXMlanJ6rC>
- Ilharco G, Wortsman M, Wightman R et al. (2021). Openclip. <https://doi.org/10.5281/zenodo.5143773>
- Islam K, Zaheer MZ, Mahmood A, et al. (2024a) Diffusemix: label-preserving data augmentation with diffusion models. 2024 IEEE/CVF conference on computer vision and pattern recognition (CVPR) pp 27611–27620. <https://api.semanticscholar.org/CorpusID:270045834>
- Islam K, Zaheer MZ, Mahmood A, et al. (2024b) Genmix: effective data augmentation with generative diffusion model image editing. arXiv preprint [arXiv:2412.02366](https://arxiv.org/abs/2412.02366)<https://api.semanticscholar.org/CorpusID:274446078>
- Jaeger S, Candemir S, Antani S et al. (2014) Two public chest X-ray datasets for computer-aided screening of pulmonary diseases. *Quant Imaging Med Surg* 4(6):475
- Jia X, Zhao Y, Chan KC, et al. (2023) Taming encoder for zero fine-tuning image customization with text-to-image diffusion models. arXiv preprint [arXiv:2304.02642](https://arxiv.org/abs/2304.02642)
- Jiang HH, Brown L, Cheng J, et al. (2023) Ai art and its impact on artists. In: Proceedings of the 2023 AAAI/ACM conference on AI, ethics, and society. Association for computing machinery, New York, NY, USA, AIES '23, p 363–374, <https://doi.org/10.1145/3600211.3604681>
- Jin Y, Ling P, Dong X, et al. (2024) Reasonpix2pix: instruction reasoning dataset for advanced image editing. arXiv preprint [arXiv:2405.11190](https://arxiv.org/abs/2405.11190)

- Ju X, Zeng A, Bian Y, et al. (2023) Direct inversion: boosting diffusion-based editing with 3 lines of code. arXiv preprint [arXiv:2310.01506](https://arxiv.org/abs/2310.01506)
- Kang M, Kim S (2023) Guidedmixup: an efficient mixup strategy guided by saliency maps. In: Proceedings of the AAAI conference on artificial intelligence, pp 1096–1104
- Kang G, Dong X, Zheng L, et al. (2017) Patchshuffle regularization. arXiv preprint [arXiv:1707.07103](https://arxiv.org/abs/1707.07103)
- Kawar B, Zada S, Lang O, et al. (2023) Imagic: text-based real image editing with diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 6007–6017
- Kazerouni A, Aghdam EK, Heidari M, et al. (2023) Diffusion models in medical imaging: a comprehensive survey. *Med Image Anal* p 102846
- Kebaili A, Lapuyade-Lahorgue J, Ruan S (2023) Deep learning approaches for data augmentation in medical imaging: a review. *J Imaging* 9(4):81
- Kim JH, Choo W, Jeong H, et al. (2021) Co-mixup: saliency guided joint mixup with supermodular diversity. arXiv preprint [arXiv:2102.03065](https://arxiv.org/abs/2102.03065)
- Kim G, Kwon T, Ye JC (2022) Diffusionclip: text-guided diffusion models for robust image manipulation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 2426–2435
- Kim S, Jung S, Kim B, et al. (2023) Towards safe self-distillation of internet-scale text-to-image diffusion models. arXiv preprint [arXiv:2307.05977](https://arxiv.org/abs/2307.05977)
- Kirillov A, Mintun E, Ravi N, et al. (2023) Segment anything. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 4015–4026
- Kirstain Y, Levy O, Polyak A (2023) X & fuse: fusing visual information in text-to-image generation. arXiv preprint [arXiv:2303.01000](https://arxiv.org/abs/2303.01000)
- Koh PW, Sagawa S, Marklund H, et al. (2021) Wilds: a benchmark of in-the-wild distribution shifts. In: International conference on machine learning, PMLR, pp 5637–5664
- Kong C, Jeon D, Kwon O, et al. (2023) Leveraging off-the-shelf diffusion model for multi-attribute fashion image manipulation. In: Proceedings of the IEEE/cvf winter conference on applications of computer vision, pp 848–857
- Krause J, Stark M, Deng J, et al. (2013) 3D object representations for fine-grained categorization. In: Proceedings of the IEEE international conference on computer vision workshops, pp 554–561
- Krizhevsky A, Hinton G, et al. (2009) Learning multiple layers of features from tiny images.
- Kulikov V, Yadin S, Kleiner M, et al. (2023) Sinddm: a single image denoising diffusion model. In: International conference on machine learning, PMLR, pp 17920–17930
- Kumari N, Zhang B, Wang SY, et al. (2023a) Ablating concepts in text-to-image diffusion models. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 22691–22702
- Kumari N, Zhang B, Zhang R, et al. (2023b) Multi-concept customization of text-to-image diffusion. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 1931–1941
- Kwon G, Ye JC (2022) Diffusion-based image translation using disentangled style and content representation. arXiv preprint [arXiv:2209.15264](https://arxiv.org/abs/2209.15264)
- Kwon M, Jeong J, Uh Y (2022) Diffusion models already have a semantic latent space. arXiv preprint [arXiv:2210.10960](https://arxiv.org/abs/2210.10960)
- Lee J, Jung DS, Lee K, et al. (2024) Semanticdraw: towards real-time interactive content creation from image diffusion models. <https://arxiv.org/abs/2403.09055>, [arXiv:2403.09055](https://arxiv.org/abs/2403.09055)
- Levin E, Fried O (2023) Differential diffusion: giving each pixel its strength. arXiv preprint [arXiv:2306.00950](https://arxiv.org/abs/2306.00950)
- Li J, Li D, Xiong C, et al. (2022) Blip: bootstrapping language-image pre-training for unified vision-language understanding and generation. In: International conference on machine learning, PMLR, pp 12888–12900
- Li B, Wang X, Xu X, et al. (2023a) Semantic-guided generative image augmentation method with diffusion models for image classification. In: AAAI conference on artificial intelligence, <https://api.semanticscholar.org/CorpusID:256615374>
- Li S, van de Weijer J, Hu T, et al. (2023b) Stylediffusion: prompt-embedding inversion for text-based editing. arXiv preprint [arXiv:2303.15649](https://arxiv.org/abs/2303.15649)
- Li Z, Wei P, Yin X, et al. (2023c) Virtual try-on with pose-garment keypoints guided inpainting. In: Proceedings of the IEEE/CVF international conference on computer vision (ICCV), pp 22788–22797
- Li Y, Dong X, Chen C, et al. (2025) A simple background augmentation method for object detection with diffusion model. In: European conference on computer vision, Springer, pp 462–479
- Lin TY, Maire M, Belongie S, et al. (2014) Microsoft coco: common objects in context. In: Computer vision—ECCV 2014: 13th European conference, Zurich, Switzerland, september 6–12, 2014, proceedings, Part V 13, Springer, pp 740–755
- Lin Y, Chen YW, Tsai YH, et al. (2024) Text-driven image editing via learnable regions. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 7059–7068
- Lingenberg T, Reuter M, Sudhakaran G, et al. (2024) Diagen: diverse image augmentation with generative models. arXiv preprint [arXiv:2408.14584](https://arxiv.org/abs/2408.14584)

- Lu C, Zhou Y, Bao F et al. (2022a) Dpm-solver: a fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Adv Neural Inform Proc Syst* 35:5775–5787
- Lu C, Zhou Y, Bao F, et al. (2022b) Dpm-solver++: fast solver for guided sampling of diffusion probabilistic models. arXiv preprint [arXiv:2211.01095](https://arxiv.org/abs/2211.01095)
- Lugmayr A, Danelljan M, Romero A, et al. (2022) Repaint: inpainting using denoising diffusion probabilistic models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 11461–11471
- Luo W (2023) A comprehensive survey on knowledge distillation of diffusion models. arXiv preprint [arXiv:2304.04262](https://arxiv.org/abs/2304.04262)
- Luo XJ, Wang S, Wu Z, et al. (2023) Camdiff: camouflage image augmentation via diffusion model. arXiv preprint [arXiv:2304.05469](https://arxiv.org/abs/2304.05469)
- Ma Y, Yang H, Wang W, et al. (2023) Unified multi-modal latent diffusion for joint subject and text conditional image generation. arXiv preprint [arXiv:2303.09319](https://arxiv.org/abs/2303.09319)
- Ma F, Qi W, Zhao G, et al. (2024a) Erase, then redraw: a novel data augmentation approach for free space detection using diffusion model. arXiv preprint [arXiv:2409.20164](https://arxiv.org/abs/2409.20164)
- Ma X, Fang G, Wang X (2024b) Deepcache: accelerating diffusion models for free. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 15762–15772
- Madaan N, Bedathur S (2023) Diffusion-guided counterfactual generation for model explainability. In: *XAI in action: past, present, and future applications*
- Maji S, Rahtu E, Kannala J, et al. (2013) Fine-grained visual classification of aircraft. arXiv preprint [arXiv:1306.5151](https://arxiv.org/abs/1306.5151)
- Meng C, He Y, Song Y, et al. (2021) Sdedit: guided image synthesis and editing with stochastic differential equations. arXiv preprint [arXiv:2108.01073](https://arxiv.org/abs/2108.01073)
- Michaeli E, Fried O (2024) Advancing fine-grained classification by structure and subject preserving augmentation. ArXiv abs/2406.14551. <https://api.semanticscholar.org/CorpusID:270620582>
- Mo S, Xie E, Chu R et al. (2023) Dit-3d: exploring plain diffusion transformers for 3d shape generation. *Adv Neural Inform Proc Syst* 36:67960–67971
- Mokady R, Hertz A, Aberman K, et al. (2023) Null-text inversion for editing real images using guided diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 6038–6047
- Muhammad A, Salman Z, Lee K et al. (2023) Harnessing the power of diffusion models for plant disease image augmentation. *Front Plant Sci*. <https://doi.org/10.3389/fpls.2023.1280496>
- Mumuni A, Mumuni F (2022) Data augmentation: a comprehensive survey of modern approaches. *Array* 16:100258. <https://doi.org/10.1016/j.array.2022.100258>
- Ni Z, Wei L, Li J, et al. (2023) Degeneration-tuning: using scrambled grid shield unwanted concepts from stable diffusion. In: *Proceedings of the 31st ACM international conference on multimedia*, pp 8900–8909
- Nichol AQ, Dhariwal P (2021) Improved denoising diffusion probabilistic models. In: *International conference on machine learning*, PMLR, pp 8162–8171
- Nichol A, Dhariwal P, Ramesh A, et al. (2021) Glide: towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint [arXiv:2112.10741](https://arxiv.org/abs/2112.10741)
- Nilsback ME, Zisserman A (2008) Automated flower classification over a large number of classes. In: *2008 Sixth Indian conference on computer vision, graphics & image processing*, IEEE, pp 722–729
- Oquab M, Darcet T, Moutakanni T, et al. (2023) Dinov2: learning robust visual features without supervision. arXiv preprint [arXiv:2304.07193](https://arxiv.org/abs/2304.07193)
- Packhäuser K, Follé L, Thamm F, et al. (2023) Generation of anonymous chest radiographs using latent diffusion models for training thoracic abnormality classification systems. In: *2023 IEEE 20th international symposium on biomedical imaging (ISBI)*, IEEE, pp 1–5
- Parihar R, Bhat A, Basu A, et al. (2024) Balancing act: distribution-guided debiasing in diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 6668–6678
- Parkhi OM, Vedaldi A, Zisserman A, et al. (2012) Cats and dogs. In: *2012 IEEE conference on computer vision and pattern recognition*, IEEE, pp 3498–3505
- Parmar G, Kumar Singh K, Zhang R, et al. (2023) Zero-shot image-to-image translation. In: *ACM SIG-GRAPH 2023 conference proceedings*, pp 1–11
- Patel M, Jung S, Baral C, et al. (2024) λ -eclipse: multi-concept personalized text-to-image diffusion models by leveraging clip latent space. arXiv preprint [arXiv:2402.05195](https://arxiv.org/abs/2402.05195)
- Peebles W, Xie S (2023) Scalable diffusion models with transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*, pp 4195–4205
- Perez L, Wang J (2017) The effectiveness of data augmentation in image classification using deep learning. arXiv preprint [arXiv:1712.04621](https://arxiv.org/abs/1712.04621)
- Pernias P, Rampas D, Richter ML, et al. (2023) Wuerstchen: an efficient architecture for large-scale text-to-image diffusion models. [arXiv:2306.00637](https://arxiv.org/abs/2306.00637)

- Phung Q, Ge S, Huang JB (2024) Coherent zero-shot visual instruction generation. arXiv preprint [arXiv:2406.04337](https://arxiv.org/abs/2406.04337)
- Pinaya WH, Tudosiu PD, Dafflon J, et al. (2022) Brain imaging generation with latent diffusion models. In: MICCAI workshop on deep generative models, Springer, pp 117–126
- Podell D, English Z, Lacey K, et al. (2023) Sdxl: improving latent diffusion models for high-resolution image synthesis. arXiv preprint [arXiv:2307.01952](https://arxiv.org/abs/2307.01952)
- Qiu Z, Liu W, Feng H et al. (2023) Controlling text-to-image diffusion by orthogonal finetuning. *Adv Neural Inform Proc Syst* 36:79320–79362
- Radford A, Kim JW, Hallacy C, et al. (2021) Learning transferable visual models from natural language supervision. In: International conference on machine learning, PMLR, pp 8748–8763
- Raffel C, Shazeer N, Roberts A et al. (2020) Exploring the limits of transfer learning with a unified text-to-text transformer. *J Mach Learn Res* 21(140):1–67
- Rahat F, Hossain MS, Ahmed MR, et al. (2024) Data augmentation for image classification using generative ai. arXiv preprint [arXiv:2409.00547](https://arxiv.org/abs/2409.00547)
- Ramesh A, Dhariwal P, Nichol A, et al. (2022) Hierarchical text-conditional image generation with clip latents. arXiv preprint [arXiv:2204.06125](https://arxiv.org/abs/2204.06125) 1(2):3
- Rando J, Paleka D, Lindner D, et al. (2022) Red-teaming the stable diffusion safety filter. arXiv preprint [arXiv:2210.04610](https://arxiv.org/abs/2210.04610)
- Ranftl R, Lasinger K, Hafner D et al. (2020) Towards robust monocular depth estimation: mixing datasets for zero-shot cross-dataset transfer. *IEEE Trans Pattern Anal Mach Intell* 44(3):1623–1637
- Rao Y, Chen G, Lu J, et al. (2021) Counterfactual attention learning for fine-grained visual categorization and re-identification. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 1025–1034
- Rombach R, Blattmann A, Lorenz D, et al. (2022) High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 10684–10695
- Ronneberger O, Fischer P, Brox T (2015) U-net: convolutional networks for biomedical image segmentation. In: Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18, Springer, pp 234–241
- Rouzrokh P, Khosravi B, Faghani S, et al. (2022) Multitask brain tumor inpainting with diffusion models: a methodological report. arXiv preprint [arXiv:2210.12113](https://arxiv.org/abs/2210.12113)
- Ruiz N, Li Y, Jampani V, et al. (2023) Dreambooth: fine tuning text-to-image diffusion models for subject-driven generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 22500–22510
- Ruiz N, Li Y, Jampani V, et al. (2024) Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 6527–6536
- Sagawa S, Koh PW, Hashimoto TB, et al. (2019) Distributionally robust neural networks for group shifts: on the importance of regularization for worst-case generalization. arXiv preprint [arXiv:1911.08731](https://arxiv.org/abs/1911.08731)
- Sagers LW, Diao JA, Groh M, et al. (2022) Improving dermatology classifiers across populations using images generated by large diffusion models. arXiv preprint [arXiv:2211.13352](https://arxiv.org/abs/2211.13352)
- Saharia C, Chan W, Chang H, et al. (2022a) Palette: image-to-image diffusion models. In: ACM SIGGRAPH 2022 conference proceedings, pp 1–10
- Saharia C, Chan W, Saxena S et al. (2022b) Photorealistic text-to-image diffusion models with deep language understanding. *Adv Neural Inform Proc Syst* 35:36479–36494
- Salimans T, Ho J (2022) Progressive distillation for fast sampling of diffusion models. arXiv preprint [arXiv:2202.00512](https://arxiv.org/abs/2202.00512)
- Salimans T, Goodfellow I, Zaremba W, et al. (2016) Improved techniques for training gans. *Adv Neural Inform Proc Syst* 29
- Sanchez P, Tsafaris SA (2022) Diffusion causal models for counterfactual estimation. arXiv preprint [arXiv:2202.10166](https://arxiv.org/abs/2202.10166)
- Sanchez P, Kascenas A, Liu X, et al. (2022) What is healthy? Generative counterfactual diffusion for lesion localization. In: MICCAI workshop on deep generative models, Springer, pp 34–44
- Santos R, Silva J, Branco A (2024) Pix2pix-on-the-fly: leveraging llms for instruction-guided image editing. arXiv preprint [arXiv:2403.08004](https://arxiv.org/abs/2403.08004)
- Sarukkai V, Li L, Ma A, et al. (2024) Collage diffusion. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision, pp 4208–4217
- Sauer A, Lorenz D, Blattmann A, et al. (2025) Adversarial diffusion distillation. In: European conference on computer vision, Springer, pp 87–103
- Schnell J, Wang J, Qi L, et al. (2024) Scribblegen: generative data augmentation improves scribble-supervised semantic segmentation. [arXiv:2311.17121](https://arxiv.org/abs/2311.17121)

- Schramowski P, Brack M, Deiseroth B, et al. (2023) Safe latent diffusion: mitigating inappropriate degeneration in diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 22522–22531
- Schuhmann C, Beaumont R, Vencu R et al. (2022) Laion-5b: an open large-scale dataset for training next generation image-text models. *Adv Neural Inform Proc Syst* 35:25278–25294
- Shang Y, Yuan Z, Xie B, et al. (2023) Post-training quantization on diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 1972–1981
- Sheynin S, Ashual O, Polyak A, et al. (2022) Knn-diffusion: image generation via large-scale retrieval. *arXiv preprint arXiv:2204.02849*
- Shi J, Xiong W, Lin Z, et al. (2024) Instantbooth: personalized text-to-image generation without test-time finetuning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 8543–8552
- Shih A, Belkhal S, Ermon S, et al. (2024) Parallel sampling of diffusion models. *Adv Neural Inform Proc Syst* 36
- Shin C, Kim H, Lee CH, et al. (2024) Edit-a-video: single video editing with object-aware consistency. In: Asian conference on machine learning, PMLR, pp 1215–1230
- Shorten C, Khoshgoftaar TM (2019) A survey on image data augmentation for deep learning. *J Big Data* 6:1–48
- Shrestha A, Mahmood A (2019) Review of deep learning algorithms and architectures. *IEEE Access* 7:53040–53065
- Smith JS, Hsu YC, Zhang L, et al. (2023) Continual diffusion: continual customization of text-to-image diffusion with c-lora. *arXiv preprint arXiv:2304.06027*
- Sohl-Dickstein J, Weiss E, Maheswaranathan N, et al. (2015) Deep unsupervised learning using nonequilibrium thermodynamics. In: International conference on machine learning, PMLR, pp 2256–2265
- Sohn K, Ruiz N, Lee K, et al. (2023) Styledrop: text-to-image generation in any style. *arXiv preprint arXiv:2306.00983*
- Song J, Meng C, Ermon S (2020a) Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*
- Song Y, Sohl-Dickstein J, Kingma DP, et al. (2020b) Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*
- Song Y, Zhang Z, Lin Z, et al. (2022) Objectstitch: generative object compositing. *arXiv preprint arXiv:2212.00932*
- Song W, Ma W, Zhang M et al. (2024) Lightweight diffusion models: a survey. *Artif Intell Rev* 57(6):161
- at StabilityAI DL (2023) DeepFloyd IF: a novel state-of-the-art open-source text-to-image model with a high degree of photorealism and language understanding. <https://www.deepfloyd.ai/deepfloyd-if>, retrieved on 2023-11-08
- Su X, Song J, Meng C, et al. (2022) Dual diffusion implicit bridges for image-to-image translation. *arXiv preprint arXiv:2203.08382*
- Su D, Hou J, Li G, et al. (2024) Generative dataset distillation based on diffusion model. *arXiv preprint arXiv:2408.08610*
- Szegedy C, Vanhoucke V, Ioffe S, et al. (2016) Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 2818–2826
- Tang C, Wang K, van de Weijer J (2024a) Iterinv: iterative inversion for pixel-level t2i models. In: 2024 IEEE international conference on multimedia and expo (ICME), IEEE, pp 1–6
- Tang C, Wang K, Yang F, et al. (2024b) Locinv: localization-aware inversion for text-guided image editing. *arXiv preprint arXiv:2405.01496*
- Tang D, Cao X, Wu X, et al. (2024c) Aerogen: enhancing remote sensing object detection with diffusion-driven data generation. *arXiv preprint arXiv:2411.15497*
- Tay Y (2023) A new open source flan 20b with ul2. Yi Tay URL: <https://www.yitay.net/blog/flan-ul2-20b> [accessed 2023-04-03]
- Tewel Y, Gal R, Chechik G, et al. (2023) Key-locked rank one editing for text-to-image personalization. In: ACM SIGGRAPH 2023 conference proceedings, pp 1–11
- Trabucco B, Doherty K, Gurinas M, et al. (2023) Effective data augmentation with diffusion models. *arXiv preprint arXiv:2302.07944*
- Tumanyan N, Geyer M, Bagon S, et al. (2023) Plug-and-play diffusion features for text-driven image-to-image translation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 1921–1930
- Ulhaq A, Akhtar N, Pogrebna G (2022) Efficient diffusion models for vision: a survey. *arXiv preprint arXiv:2210.09292*
- Valvano G, Agostino A, De Magistris G, et al. (2024) Controllable image synthesis of industrial data using stable diffusion. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision, pp 5354–5363

- Van Den Oord A, Vinyals O, et al. (2017) Neural discrete representation learning. *Adv Neural Inform Proc Syst* 30
- Vendrow J, Jain S, Engstrom L, et al. (2023) Dataset interfaces: diagnosing model failures using controllable counterfactual generation. arXiv preprint [arXiv:2302.07865](https://arxiv.org/abs/2302.07865)
- Vinker Y, Vovnov A, Cohen-Or D et al. (2023) Concept decomposition for visual exploration and inspiration. *ACM Trans Graphics (TOG)* 42(6):1–13
- Wah C, Branson S, Welinder P, et al. (2011) The caltech-ucsd birds-200-2011 dataset.
- Wallace B, Gokul A, Naik N (2023) Edict: exact diffusion inversion via coupled transformations. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 22532–22541
- Wang Z, Bovik AC, Sheikh HR et al. (2004) Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Proc* 13(4):600–612
- Wang T, Zhang T, Zhang B, et al. (2022a) Pretraining is all you need for image-to-image translation. arXiv preprint [arXiv:2205.12952](https://arxiv.org/abs/2205.12952)
- Wang W, Bao J, Zhou W, et al. (2022b) Sindiffusion: learning a diffusion model from a single natural image. arXiv preprint [arXiv:2211.12445](https://arxiv.org/abs/2211.12445)
- Wang Q, Zhang B, Birsak M, et al. (2023a) Instructedit: improving automatic masks for diffusion-based image editing with user instructions. arXiv preprint [arXiv:2305.18047](https://arxiv.org/abs/2305.18047)
- Wang S, Saharia C, Montgomery C, et al. (2023b) Imagen editor and editbench: advancing and evaluating text-guided image inpainting. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 18359–18369
- Wang Z, Jiang Y, Zheng H, et al. (2024a) Patch diffusion: faster and more data-efficient training of diffusion models. *Adv Neural Inform Proc Syst* 36
- Wang Z, Wei L, Wang T, et al. (2024b) Enhance image classification via inter-class image mixup with diffusion model. 2024 IEEE/CVF conference on computer vision and pattern recognition (CVPR) pp 17223–17233. <https://api.semanticscholar.org/CorpusID:268733015>
- Wasserman N, Rotstein N, Ganz R, et al. (2024) Paint by inpaint: learning to add image objects by removing them first. arXiv preprint [arXiv:2404.18212](https://arxiv.org/abs/2404.18212)
- Wei Y, Zhang Y, Ji Z, et al. (2023) Elite: encoding visual concepts into textual embeddings for customized text-to-image generation. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 15943–15953
- Wolleb J, Bieder F, Sandkühler R, et al. (2022) Diffusion models for medical anomaly detection. In: International conference on medical image computing and computer-assisted intervention, Springer, pp 35–45
- Wu JZ, Ge Y, Wang X, et al. (2023a) Tune-a-video: one-shot tuning of image diffusion models for text-to-video generation. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 7623–7633
- Wu Q, Liu Y, Zhao H, et al. (2023b) Uncovering the disentanglement capability in text-to-image diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 1900–1910
- Wu R, Wang Y, Shi H, et al. (2023c) Towards prompt-robust face privacy protection via adversarial decoupling augmentation framework. arXiv preprint [arXiv:2305.03980](https://arxiv.org/abs/2305.03980)
- Wu T, Ye S, Chen S, et al. (2023d) Detail reinforcement diffusion model: augmentation fine-grained visual categorization in few-shot conditions. arXiv preprint [arXiv:2309.08097](https://arxiv.org/abs/2309.08097)
- Xia W, Lyu Q, Wang G (2022) Low-dose ct using denoising diffusion probabilistic model for $20 \times$ speedup. arXiv preprint [arXiv:2209.15136](https://arxiv.org/abs/2209.15136)
- Xiao G, Yin T, Freeman WT, et al. (2023) Fastcomposer: Tuning-free multi-subject image generation with localized attention. arXiv preprint [arXiv:2305.10431](https://arxiv.org/abs/2305.10431)
- Xie J, Li W, Li X, et al. (2023a) Mosaicfusion: diffusion models as data augmenters for large vocabulary instance segmentation. arXiv preprint [arXiv:2309.13042](https://arxiv.org/abs/2309.13042)
- Xie S, Zhang Z, Lin Z, et al. (2023b) Smartbrush: text and shape guided object inpainting with diffusion model. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 22428–22437
- Xu M, Yoon S, Fuentes A et al. (2023) A comprehensive survey of image augmentation techniques for deep learning. *Pattern Recognit* 137:109347
- Xu S, Ma Z, Huang Y, et al. (2024) Cyclenet: rethinking cycle consistency in text-guided diffusion for image manipulation. *Adv Neural Inform Proc Syst* 36
- Xue H, Huang Z, Sun Q, et al. (2023) Freestyle layout-to-image synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 14256–14266
- Yang J, Shi R, Ni B (2021) Medmnist classification decathlon: a lightweight autml benchmark for medical image analysis. In: 2021 IEEE 18th international symposium on biomedical imaging (ISBI), IEEE, pp 191–195

- Yang S, Xiao W, Zhang M, et al. (2022) Image data augmentation for deep learning: a survey. arXiv preprint [arXiv:2204.08610](https://arxiv.org/abs/2204.08610)
- Yang B, Gu S, Zhang B, et al. (2023a) Paint by example: exemplar-based image editing with diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 18381–18391
- Yang L, Zhang Z, Song Y et al. (2023) Diffusion models: a comprehensive survey of methods and applications. *ACM Comput Surv* 56(4):1–39
- Yang L, Zeng B, Liu J, et al. (2024) Editworld: simulating world dynamics for instruction-following image editing. arXiv preprint [arXiv:2405.14785](https://arxiv.org/abs/2405.14785)
- Ye J, Ni H, Jin P, et al. (2023) Synthetic augmentation with large-scale unconditional pre-training. In: International conference on medical image computing and computer-assisted intervention, Springer, pp 754–764
- Yin Y, Kaddour J, Zhang X, et al. (2023) Ttida: controllable generative data augmentation via text-to-text and text-to-image models. arXiv preprint [arXiv:2304.08821](https://arxiv.org/abs/2304.08821)
- Yin T, Gharbi M, Zhang R, et al. (2024) One-step diffusion with distribution matching distillation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 6613–6623
- Yu T, Feng R, Feng R, et al. (2023a) Inpaint anything: segment anything meets image inpainting. arXiv preprint [arXiv:2304.06790](https://arxiv.org/abs/2304.06790)
- Yu X, Li G, Lou W, et al. (2023b) Diffusion-based data augmentation for nuclei image segmentation. In: International conference on medical image computing and computer-assisted intervention, Springer, pp 592–602
- Yu H, Luo H, Wang F, et al. (2024) Uncovering the text embedding in text-to-image diffusion models. arXiv preprint [arXiv:2404.01154](https://arxiv.org/abs/2404.01154)
- Yuan J, Pinto F, Davies A, et al. (2022) Not just pretty pictures: toward interventional data augmentation using text-to-image generators. arXiv preprint [arXiv:2212.11237](https://arxiv.org/abs/2212.11237)
- Yun S, Han D, Oh SJ, et al. (2019) Cutmix: regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 6023–6032
- Zajac M, Deja K, Kuzina A, et al. (2023) Exploring continual learning of diffusion models. arXiv preprint [arXiv:2303.15342](https://arxiv.org/abs/2303.15342)
- Zang Z, Luo H, Wang K, et al. (2023) Boosting unsupervised contrastive learning using diffusion-based data augmentation from scratch. arXiv preprint [arXiv:2309.07909](https://arxiv.org/abs/2309.07909)
- Zeng Y, Lin Z, Zhang J, et al. (2023) Scenecomposer: any-level semantic image synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 22468–22478
- Zeng W, Yan Y, Zhu Q, et al. (2024) Infusion: preventing customized text-to-image diffusion from overfitting. arXiv preprint [arXiv:2404.14007](https://arxiv.org/abs/2404.14007)
- Zhang H, Cisse M, Dauphin YN, et al. (2017) mixup: beyond empirical risk minimization. arXiv preprint [arXiv:1710.09412](https://arxiv.org/abs/1710.09412)
- Zhang R, Isola P, Efros AA, et al. (2018) The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 586–595
- Zhang Y, Zhou D, Hooi B, et al. (2022) Expanding small-scale datasets with guided imagination. arXiv 2022. arXiv preprint [arXiv:2211.13976](https://arxiv.org/abs/2211.13976)
- Zhang C, Zhang C, Zhang M, et al. (2023a) Text-to-image diffusion models in generative AI: a survey. arXiv preprint [arXiv:2303.07909](https://arxiv.org/abs/2303.07909)
- Zhang L, Rao A, Agrawala M (2023b) Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 3836–3847
- Zhang M, Wu J, Ren Y, et al. (2023c) Diffusionengine: diffusion model is scalable data engine for object detection. arXiv preprint [arXiv:2309.03893](https://arxiv.org/abs/2309.03893)
- Zhang Y, Dong W, Tang F et al. (2023d) Prospect: prompt spectrum for attribute-aware personalization of diffusion models. *ACM Trans Graphics (TOG)* 42(6):1–14
- Zhang Y, Huang N, Tang F, et al. (2023e) Inversion-based style transfer with diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 10146–10156
- Zhang Z, Han L, Ghosh A, et al. (2023f) Sine: single image editing with text-to-image diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 6027–6037
- Zhang B, Luo C, Yu D, et al. (2024a) Metadiff: meta-learning with conditional diffusion for few-shot learning. In: Proceedings of the AAAI conference on artificial intelligence, pp 16687–16695
- Zhang G, Wang K, Xu X, et al. (2024b) Forget-me-not: learning to forget in text-to-image diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 1755–1764
- Zhang Z, Lin M, Ji R (2024c) Objectadd: adding objects into image via a training-free diffusion modification fashion. arXiv preprint [arXiv:2404.17230](https://arxiv.org/abs/2404.17230)
- Zhang Z, Zheng J, Fang Z, et al. (2024d) Text-to-image editing by image information removal. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision, pp 5232–5241

- Zhao H, Sheng D, Bao J, et al. (2022) X-paste: revisiting scalable copy-paste for instance segmentation using clip and stablediffusion. arXiv preprint [arXiv:2212.03863](https://arxiv.org/abs/2212.03863)
- Zhao S, Chen D, Chen YC, et al. (2024) Uni-controlnet: all-in-one control to text-to-image diffusion models. *Adv Neural Inform Proc Syst* 36
- Zhong J, Guo X, Dong J, et al. (2024) Diffusion tuning: transferring diffusion models via chain of forgetting. arXiv preprint [arXiv:2406.00773](https://arxiv.org/abs/2406.00773)
- Zhou Y, Sahak H, Ba J (2023a) Training on thin air: improve image classification with generated data. arXiv preprint [arXiv:2305.15316](https://arxiv.org/abs/2305.15316)
- Zhou Y, Sahak H, Ba J (2023b) Using synthetic data for data augmentation to improve classification accuracy. In: *Proceedings of the international conference on machine learning (ICML) Workshops*
- Zhou Y, Guo C, Wang X, et al. (2024) A survey on data augmentation in large model era. arXiv preprint [arXiv:2401.15422](https://arxiv.org/abs/2401.15422)
- Zhu J, Ma H, Chen J, et al. (2023) Domainstudio: fine-tuning diffusion models for domain-driven image generation using limited data. arXiv preprint [arXiv:2306.14153](https://arxiv.org/abs/2306.14153)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Panagiotis Alimisis¹ · Ioannis Mademlis¹ · Panagiotis Radoglou-Grammatikis^{2,3} · Panagiotis Sarigiannidis² · Georgios Th. Papadopoulos¹

✉ Georgios Th. Papadopoulos
g.th.papadopoulos@hua.gr

Panagiotis Alimisis
csi23301@hua.gr

Ioannis Mademlis
imademlis@hua.gr

Panagiotis Radoglou-Grammatikis
pradoglou@uowm.gr; pradoglou@k3y.bg

Panagiotis Sarigiannidis
psarigiannidis@uowm.gr

¹ Department of Informatics and Telematics, Harokopio University of Athens, Thiseos 70, Athens 17676, Attiki, Greece

² Department of Electrical and Computer Engineering, University of Western Macedonia, Active Urban Planning Zone, Kozani 50150, Kozani, Greece

³ K3Y, Studentski district, Vitosha quarter, bl. 9, Sofia BG 1700, Sofia, Bulgaria

Terms and Conditions

Springer Nature journal content, brought to you courtesy of Springer Nature Customer Service Center GmbH (“Springer Nature”).

Springer Nature supports a reasonable amount of sharing of research papers by authors, subscribers and authorised users (“Users”), for small-scale personal, non-commercial use provided that all copyright, trade and service marks and other proprietary notices are maintained. By accessing, sharing, receiving or otherwise using the Springer Nature journal content you agree to these terms of use (“Terms”). For these purposes, Springer Nature considers academic use (by researchers and students) to be non-commercial.

These Terms are supplementary and will apply in addition to any applicable website terms and conditions, a relevant site licence or a personal subscription. These Terms will prevail over any conflict or ambiguity with regards to the relevant terms, a site licence or a personal subscription (to the extent of the conflict or ambiguity only). For Creative Commons-licensed articles, the terms of the Creative Commons license used will apply.

We collect and use personal data to provide access to the Springer Nature journal content. We may also use these personal data internally within ResearchGate and Springer Nature and as agreed share it, in an anonymised way, for purposes of tracking, analysis and reporting. We will not otherwise disclose your personal data outside the ResearchGate or the Springer Nature group of companies unless we have your permission as detailed in the Privacy Policy.

While Users may use the Springer Nature journal content for small scale, personal non-commercial use, it is important to note that Users may not:

1. use such content for the purpose of providing other users with access on a regular or large scale basis or as a means to circumvent access control;
2. use such content where to do so would be considered a criminal or statutory offence in any jurisdiction, or gives rise to civil liability, or is otherwise unlawful;
3. falsely or misleadingly imply or suggest endorsement, approval, sponsorship, or association unless explicitly agreed to by Springer Nature in writing;
4. use bots or other automated methods to access the content or redirect messages
5. override any security feature or exclusionary protocol; or
6. share the content in order to create substitute for Springer Nature products or services or a systematic database of Springer Nature journal content.

In line with the restriction against commercial use, Springer Nature does not permit the creation of a product or service that creates revenue, royalties, rent or income from our content or its inclusion as part of a paid for service or for other commercial gain. Springer Nature journal content cannot be used for inter-library loans and librarians may not upload Springer Nature journal content on a large scale into their, or any other, institutional repository.

These terms of use are reviewed regularly and may be amended at any time. Springer Nature is not obligated to publish any information or content on this website and may remove it or features or functionality at our sole discretion, at any time with or without notice. Springer Nature may revoke this licence to you at any time and remove access to any copies of the Springer Nature journal content which have been saved.

To the fullest extent permitted by law, Springer Nature makes no warranties, representations or guarantees to Users, either express or implied with respect to the Springer nature journal content and all parties disclaim and waive any implied warranties or warranties imposed by law, including merchantability or fitness for any particular purpose.

Please note that these rights do not automatically extend to content, data or other material published by Springer Nature that may be licensed from third parties.

If you would like to use or distribute our Springer Nature journal content to a wider audience or on a regular basis or in any other manner not expressly permitted by these Terms, please contact Springer Nature at

onlineservice@springernature.com