

Deep Learning-based Image and Video Inpainting: A Survey

Weize Quan^{1,2} · Jiayi Chen^{1,2} · Yanli Liu³ · Dong-Ming Yan^{1,2,✉} · Peter Wonka⁴

Received: date / Accepted: date

Abstract Image and video inpainting is a classic problem in computer vision and computer graphics, aiming to fill in the plausible and realistic content in the missing areas of images and videos. With the advance of deep learning, this problem has achieved significant progress recently. The goal of this paper is to comprehensively review the deep learning-based methods for image and video inpainting. Specifically, we sort existing methods into different categories from the perspective of their high-level inpainting pipeline, present different deep learning architectures, including CNN, VAE, GAN, diffusion models, etc., and summarize techniques for module design. We review the training objectives and the common benchmark datasets. We present evaluation metrics for low-level pixel and high-level perceptual similarity, conduct a performance evaluation, and discuss the strengths and weaknesses of representative inpainting methods. We also discuss related real-world applications. Finally, we discuss open challenges and suggest potential future research directions.

Keywords Image inpainting · Video inpainting · Deep learning · Generation

✉ D.-M. Yan
yandongming@gmail.com

¹ MAIS & NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing, China

² School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

³ College of Computer Science, Sichuan University, Chengdu, China

⁴ Computer, Electrical and Mathematical Science and Engineering Division, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia




Fig. 1: Application examples of inpainting techniques: photo restoration (top left: image from (Bertalmio et al., 2000)), text removal (top right: image from (Bertalmio et al., 2000)), undesired target removal (bottom left: image from (Chen, 2018)), and face verification (bottom right: image from (Zhang et al., 2018c)).

1 Introduction

Image and video inpainting (Masnou and Morel, 1998; Bertalmio et al., 2000) refers to the task of restoring missing/occluded regions of a digital image or video with plausible and natural content. Inpainting is an underconstrained problem with multiple plausible solutions, especially if there are large missing regions. Inpainting has many important applications in multiple fields, such as cultural relic restoration, virtual scene editing, digital forensics, and film and television production, etc. Fig. 1 shows some important applications of inpainting techniques. Video is composed of multiple images exhibiting temporal coherence, therefore, video inpainting is closely related to image inpainting, where the former often learns from or extends the latter. For this reason, we simultaneously review image and video




Fig. 2: The rough number of papers on image and video inpainting per year.

inpainting in this survey, and the number of papers is shown in Fig. 2.

Early image inpainting methods mainly depend on low-level features of corrupted images, including PDE-based methods (Bertalmio et al., 2000; Ballester et al., 2001; Tschumperlé and Deriche, 2005) and patch-based methods (Efros and Leung, 1999; Barnes et al., 2009; Darabi et al., 2012; Huang et al., 2014; Herling and Broll, 2014; Guo et al., 2018). PDE-based approaches usually propagate the information from the boundary to create a smooth inpainting. It is possible to propagate edge information, but it is hard to inpaint textures. Instead of only considering the boundary information, patch-based approaches recover the unknown regions by matching and duplicating similar patches of known regions. For smaller areas, patch-based methods can inpaint textures and also inpaint complete objects if similar objects are available in other image regions. However, these traditional methods have limited ability to generate new semantically plausible content, especially for large missing regions and missing regions that are not similar to other image regions. A comprehensive review on classical image inpainting methods is beyond our scope, and we refer readers to the surveys (Guillemot and Meur, 2014; Jam et al., 2021) for more details.

By contrast, deep learning holds the promise to inpaint large regions and also inpaint new plausible content that was learned from a larger set of images. In the beginning convolutional neural networks (CNNs) and generative adversarial networks (GANs) were the most popular choices in the inpainting literature. CNNs are a class of feed-forward neural networks that consist of convolutional, activation, and down-/up-sampling lay-

ers. They learn a highly non-linear mapping from the input image to the output image. GANs are a type of generative model consisting of a generator and a discriminator that estimates the data distribution through an adversarial process. Recently, more attention has been paid to the transformer architecture and generative diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020). Transformers are a prevalent network architecture based on parallel multi-head attention modules. Compared to the locality of CNNs, transformers have a better ability for contextual understanding. Diffusion probabilistic models are a type of latent variable model, which mainly contain the forward process, the reverse process, and the sampling procedure. Diffusion models learn to reverse a stochastic process (*i.e.*, diffusion process) that progressively destroys data via adding noise. These deep learning-based image inpainting methods can achieve attractive results that surpass traditional methods in many aspects. From the perspective of the high-level inpainting pipeline, existing inpainting methods can be classified into three categories: a single-shot framework, a two-stage framework, and a progressive framework. Orthogonal to these main approaches, different technical methods can be observed in their realization, including mask-aware design, attention mechanisms, multi-scale aggregation, transform domain, deep prior guidance, multi-task learning, structure representations, loss functions, etc.

Compared with images, video data has an additional time dimension. Therefore, video inpainting not only fills in reasonable content in the missing regions for each frame but also aims to recover a temporally consistent solution. Because of this close relationship between image inpainting and video inpainting, many technical ideas used in image inpainting are often applied and extended to solve video inpainting tasks. Traditional video inpainting methods are usually based on patch sampling and synthesis (Wexler et al., 2007; Granados et al., 2012; Newson et al., 2014; Huang et al., 2016). These methods have limited ability to synthesize consistent content and capture complex motion and are often computationally expensive. To address these shortcomings, many deep learning-based methods have been proposed and achieved significant progress. There mainly exist four research directions: 3D CNN-based methods, shift-based methods, flow-guided methods, and attention-based methods. The core idea of these methods is to transfer information from neighboring frames to the target frame. 3D CNNs are the direct extension of 2D CNNs and work in an end-to-end manner. However, they often suffer from spatial misalignment and high computational cost. Shift-based methods can address these limitations to some extent, but within a limited

temporal window only. Flow-guided approaches can produce higher resolution and temporally consistent results but are vulnerable to imperfect optical flow completion due to occlusion and complex motion. Attention-based methods fuse known information from short and long distances. Unfortunately, inaccurate attention score estimation often leads to blurry results.

To our knowledge, there are several papers that review the deep learning-based inpainting works in the literature. Elharrouss et al. (2020) categorizes image inpainting methods into sequential-based, CNN-based, and GAN-based methods, and reviews related papers. To improve on their work, we also discuss common methodological approaches, loss functions, and evaluation metrics. We also add more discussion about further research directions and include newer work. Jam et al. (2021) reviews the traditional and deep learning-based image inpainting methods. However, they paid much attention to the traditional methods but have significantly fewer deep learning-based works compared to our survey. Weng et al. (2022) reviews some GAN-based image inpainting methods, but is generally shorter. Moreover, these existing surveys do not review the image and video inpainting simultaneously.

2 Image Inpainting

For the restoration of missing regions in an image, the results sometimes are not unique, especially for large missing areas. Consequently, there mainly exists two lines of research in the literature: (1) deterministic image inpainting and (2) stochastic image inpainting. Given a corrupted image, deterministic image inpainting methods only output an inpainted result while stochastic image inpainting methods can output multiple plausible results with a random sampling process. Inspired by multi-modal learning, some researchers have recently focused on text-guided image inpainting by providing additional information with text prompts.

2.1 Deterministic Image Inpainting

From the perspective of a high-level inpainting pipeline, existing works for deterministic image inpainting usually adopt three types of frameworks: single-shot, two-stage, and progressive. The single-shot framework usually adopts a generator network with a corrupted image as input and an inpainted image as output; The two-stage framework mainly consists of two generators, where the first generator achieves a rough result and then the second generator improves upon it; The pro-




Fig. 3: Representative pipeline of the single-shot inpainting framework. The generator takes as input the concatenation of a binary mask and a corrupted image and outputs the completed image. Training objectives are used for training the generator.

gressive framework applies one or more generators to iteratively recover missing regions along the boundary.

2.1.1 Single-shot framework

Many existing inpainting methods adopt a single-shot framework, as shown in Fig. 3. It essentially learns a mapping from a corrupted image to a complete image. The framework usually consists of generators and corresponding training objectives.

Generators. To improve the inpainting ability of the generator, there exist several lines of research: mask-aware design, attention mechanism, multi-scale aggregation, transform domain, encoder-decoder connection, and deep prior guidance.

(1) Mask-aware design.

The missing regions (indicated with a binary mask) have different shapes and convolutional operations overlapping with these missing regions may be the source of visual artifacts. Therefore, some researchers proposed mask-aware solutions for classical convolutional operation and normalization. Inspired by the inherent spatially varying property of image inpainting, Ren et al. (2015) designed a Shepard interpolation layer where the feature map and mask both conduct the same convolution operation. Its output is the fraction of feature convolution and mask convolution results. Mask convolution can simultaneously update the mask. To better handle various irregular holes and evolve the hole during mask updating, Liu et al. (2018) proposed a mask-guided convolution operation, *i.e.*, partial convolution, which distinguishes between the valid region and hole in a convolutional window. Xie et al. (2019) proposed trainable bidirectional attention maps to extend the partial convolution (Liu et al., 2018), which can adaptively learn the feature re-normalization and mask-updating.

Different from the feature normalization considered by previous methods, Yu et al. (2020) focused on the

mean and variance shift-related normalization and introduced a spatial region-wise normalization into the inpainting network. Wang et al. (2020c) designed a visual consistency network for blind image inpainting. They first predicted the damaged regions yielding a mask, and then applied an inpainting network with the proposed probabilistic context normalization, which transfers the mean and variance from known features to unknown parts building on different layers. Inspired by filling holes with pixel priorities (Criminisi et al., 2004; Zhang et al., 2019b), Wang et al. (2021c) used a structure priority (in low-resolution features) and a texture priority (in high-resolution features) in partial convolution (Liu et al., 2018). Wang et al. (2021a) proposed a dynamic selection network to utilize the valid pixels better. Specifically, they designed a validness migratable convolution to dynamically sample the convolutional locations, and a regional composite normalization module to adaptively composite batch, instance, and layer normalization on mask-based selective feature maps. Zhu et al. (2021) learned to derive the convolutional kernel from the mask for each convolutional window and proposed a point-wise normalization that produces the mask-aware scale and bias for batch normalization.

(2) Attention mechanism.

Attention is a prevalent tool to model correlation in the field of natural language processing Vaswani et al. (2017) and computer vision (Wang et al., 2018b; Fu et al., 2019). Attention is better at accessing features of distant spatial locations than convolution. In the literature, Yu et al. (2018) was the first to introduce a contextual attention mechanism to image inpainting. This pioneering work inspired many following works. To enhance both visual and semantic coherence, Zeng et al. (2019) proposed a pyramid-context encoder network with an attention transfer method, where the attention score computed in a high-level feature is used for low-level feature updating. Instead of using one fixed patch size for attention computation, Wang et al. (2019b) proposed a multi-scale contextual attention model with two different patch sizes followed by a channel attention block (Hu et al., 2018). Wang et al. (2020b) introduced a multistage attention module that performs large patch swapping in the first stage and small patch swapping in the next stage. Qin et al. (2021) combined spatial-channel attention (Chen et al., 2017) and a spatial pyramid structure to construct a multi-scale attention unit (MSAU). This unit separately conducts spatial attention on four feature maps obtained by different dilation convolutions and then applies augmented channel attention on concatenated attentive features. Zhang et al. (2022e) proposed a structure and texture interac-

tion network for image inpainting. They designed a texture spatial attention module to recover texture details with robust attention scores guided by coarse structures and introduced a structure channel excitation module to recalibrate structures according to the difference between coarse and refined structures.

In addition, some recent works proposed image inpainting networks based on vision transformers (Dosovitskiy et al., 2021). Deng et al. (2021) proposed a contextual transformer network to complete the corrupted images. Their network mainly depends on the multi-scale multi-sub-head attention, which is extended from the original multi-head attention proposed by (Vaswani et al., 2017). Cao et al. (2022) incorporated rich prior information from the ViT-based masked autoencoder (MAE) (He et al., 2022) into image inpainting. Specifically, the pre-trained MAE model provides the features prior to the encoder of the inpainting network and the attention prior to making the long-distance relationship modeling easier. Instead of using shallow projections or large receptive field convolutions to sequence the incomplete image, Zheng et al. (2022a) designed a restrictive CNN head with a small and non-overlapping receptive field as token representation. Deng et al. (2022) modified multi-head self-attention by inserting a Laplace distance prior, which computes the similarity considering the features and their locations simultaneously.

(3) Multi-scale aggregation.

In the literature on image processing, multi-scale aggregation is a common method to fuse information from different resolutions. Wang et al. (2018c) designed a generative multi-column inpainting network, consisting of three convolution branches with different filter kernel sizes, to fuse multi-scale feature representations. To create a smooth transition between the inpainted regions with existing content, Hong et al. (2019) proposed a deep fusion network with multiple fusion modules and reconstruction loss applied on multi-scale layers. The fusion module merged predicted content with the input image via a learnable alpha composition. Hui et al. (2020) proposed a dense multi-scale fusion module, which fuses hierarchical features obtained by multiple convolutional branches with different dilation rates. Zheng et al. (2021b) designed a progressive multi-scale fusion module to extract multi-scale features in parallel and progressively fuse these features, yielding more representative local features. Inspired by the high-resolution network (HRNet) for visual recognition (Sun et al., 2019; Wan et al., 2021), Wang et al. (2021c) introduced a parallel multi-resolution fusion network for image inpainting. This network can simultaneously conduct inpainting in multiple resolutions with mask-aware and attention-guided representation fusion methods. Phutke

and Murala (2021) also followed a multi-path design, where they introduce four concurrent branches with different resolutions in the encoder. A residual module with diverse receptive fields is designed as the building block of the encoder. Cao and Fu (2021) proposed a multi-scale sketch tensor network for man-made scene inpainting. This network reconstructs different types of structures by adding constraints on predicted lines, edges, and coarse images with different scales. Different from the mask-blind processing (Li et al., 2020b; Qin et al., 2021) of multi-scale features produced by convolution with different dilation rates, Zeng et al. (2022) carefully designed a gated residual connection, which considers the difference between holes and valid regions. They also proposed a soft mask-guided PatchGAN, where the discriminator is trained to predict the soft mask obtained by Gaussian filtering.

(4) Transform domain.

Instead of conducting image inpainting in the spatial domain, some existing works designed inpainting frameworks in a transformed domain via the DWT (discrete wavelet transform) (Daubechies, 1990) and the FFT (fast Fourier transform). Wang et al. (2020a) recast the image inpainting problem as predicting low-frequency semantic structures and high-frequency texture details. Specifically, they decomposed the corrupted image into different frequency components via the Haar wavelet transform (Mallat, 1989), designed a multi-frequency probabilistic inference model to predict the frequency content in missing regions, and inversely transformed back to image space. Yu et al. (2021a) adopted a similar inpainting pipeline. For the multi-frequency completion, they proposed a frequency region attentive normalization module to align and fuse the features with different frequencies and applied two discriminators to two high-frequency streams. Li et al. (2021) extracted high-frequency subbands as the texture and introduced a DWT loss to constrain the fidelity of low- and high-frequency subbands. LaMa (Suvorov et al., 2022) combined the residual design (He et al., 2016) and fast Fourier convolution (Chi et al., 2020) to construct a fast Fourier convolution residual block, which is integrated into the encoder-decoder network to handle large mask inpainting. Lu et al. (2022) further improved LaMa by introducing various types of masks and adding the focal frequency loss (Jiang et al., 2021) to constrain the spectrum of the images.

(5) Encoder-decoder connection.

Some works modify the basic encoder-decoder architecture by introducing carefully designed feature connections. Shift-Net (Yan et al., 2018) modified the U-Net architecture by introducing a specific shift-connection layer, which shifts the encoder features of the valid re-

gion to the missing regions with a guidance loss. Dolhansky and Ferrer (2018) introduced an eye inpainting network that merges the identifying information of the reference image encoding as a code. Shen et al. (2019) designed a densely connected generative network for semantic image inpainting. They combined four symmetric U-Nets with dense skip connections. Liu et al. (2020) introduced a mutual encoder-decoder CNN, fusing the texture and structure features (from the shallow and deep layers of the encoder), to jointly restore the structure and texture with feature equalization. Similarly, Guo et al. (2021) designed a two-stream image inpainting network, which combines a structure-constrained texture synthesis submodel and a texture-guided structure reconstruction submodel. In addition, they introduced a bi-directional gated feature fusion module and a contextual feature aggregation module to fuse and refine the resulting images. Feng et al. (2022) inserted generative memory into the classical encoder-decoder network to jointly exploit the high-level semantic reasoning and the pixel-level content reasoning. Based on (Liu et al., 2020), Liu et al. (2022) inferred the texture and structure with a content-consistent reference image through a feature alignment module.

(6) Deep prior guidance.

To enhance the performance of the inpainting generator, some works have explored the deep prior from a single image or a large image database. Lempitsky et al. (2018) utilized a randomly-initialized generator network as the prior to completing the corrupted image by only reconstructing the known regions. Gu et al. (2020) proposed mGANprior by incorporating a pre-trained GAN as prior for image inpainting. Specifically, this method reconstructs the incorrupt regions while filling in the missing areas by adaptively merging multiple generative feature maps from different latent codes. Richardson et al. (2021) developed a pixel2style2pixel (pSp) framework for image inpainting. They introduced an encoder consisting of a feature pyramid and multiple mapping networks to encode the damaged image into extended latent space $\mathcal{W}+$ (18 512-dimensional style vectors), which is the extension of latent space \mathcal{W} (Karras et al., 2019), and reused a pre-trained StyleGAN generator as priors to achieve the complete image. To handle the large missing regions and complex semantics, Wang et al. (2022b) designed a dual-path image inpainting framework with GAN inversion (Xia et al., 2022). Given a corrupted image, the inversion path infers the close latent code and extracts the corresponding multi-layer features from the trained GAN model, and the feed-forward path fills the missing regions by merging the above semantic priors with a deformable fusion module. To guarantee the invariance of the valid

area in the corrupted and completed images, Yu et al. (2022b) modified the GAN inversion pipeline (Richardson et al., 2021) by designing the mapping network with a pre-modulation module and introducing $\mathcal{F}\&\mathcal{W}$ + latent space, where \mathcal{F} are the feature maps of the corrupted image.

Training objectives. The training objective is a very important component of deep learning-based image inpainting methods. Pixel-wise reconstruction loss, perceptual loss (Johnson et al., 2016), style loss (Gatys et al., 2016), and adversarial loss (Goodfellow et al., 2014) are the prevalent training objectives. The adversarial loss is obtained by a discriminator network. Pathak et al. (2016) and Li et al. (2019b) adopted the discriminator (stacked convolution and down-sampling) from DCGAN (Radford et al., 2016). Considering Pathak et al. (2016)’s method struggles to maintain local consistency with the surrounding regions, Iizuka et al. (2017) proposed local and global discriminators, which generate more realistic contents. Yu et al. (2018) proposed a patch-based discriminator, which can be regarded as the generalized version of local and global discriminators (Iizuka et al., 2017). This patch-based discriminator is subsequently used in many following works. Liu et al. (2021c) designed two discriminators with small- and large-scale receptive fields to guide the inpainting network for fine-grained image detail generation.

Besides, researchers have also introduced some carefully designed losses. Li et al. (2017) introduced a semantic parsing loss for face completion. Yeh et al. (2017) proposed context and prior losses to search the closest encoding in the latent image manifold for inferring the missing content. Vo et al. (2018) proposed a structural reconstruction loss, which is the combination of reconstruction errors in pixel and feature space. For explicitly exploring the structural and textural coherence between filled contents and their surrounding contexts, Li et al. (2019a) utilized the local intrinsic dimensionality (Houle, 2017a,b) in the image- and patch-level to measure and constrain the alignment between data sub-manifolds of inpainted contents and those of the valid pixels. To stabilize the training process of face inpainting, *i.e.*, weakening gradient vanishing and model collapse, Han and Wang (2021) trained the generator via neuro-evolution and optimized the generator’s parameters by mutation and crossover.

Some researchers introduced additional training objectives via multi-task learning. Liao et al. (2018a) presented a novel collaborative framework by training a generator simultaneously on multiple tasks, *i.e.*, face completion, landmark detection, and semantic parsing. To enhance the inpainting capability of the network for image structure, Yang et al. (2020) designed a struc-

ture restoration branch in the decoder and explicitly inserted the structure features into the primary inpainting process. Appropriate semantic guidance is a suitable tool for image inpainting (Song et al., 2018b), inspired by this, Liao et al. (2020, 2021b,a) proposed a unified framework to jointly predict the segmentation maps and recover the corrupted images. Specifically, Liao et al. (2020, 2021b) designed a semantic guidance and evaluation network that iteratively updates and evaluates a semantic map and infers the missing contents in multiple scales. However, this method may create implausible textures and blurry boundaries, especially on mixed semantic regions. To solve this problem, Liao et al. (2021a) devised a semantic-wise attention propagation module to apply the attention operation on the same semantic regions. They also introduced two coherence losses to constrain the consistency between the semantic map and the structure and texture of the inpainted image. Zhang et al. (2020b) studied how to improve the visual quality of inpainted images and proposed a pixel-wise dense detector for image inpainting. This detection-based framework can localize the artifacts of completed images, and the corresponding position information is combined with the reconstruction loss to better guide the training of the inpainting network. Zhang et al. (2021) introduced the semantic prior estimation as a pretext task with a pre-trained multi-label classification model, and then utilized the learned semantic priors to guide the inpainting process through a spatially-adaptive normalization module (Park et al., 2019). Yu et al. (2022a) jointly solved image reconstruction, semantic segmentation, and edge texture generation. Each branch is implemented with a transformer network, and a multi-scale spatial-aware attention block is developed to guide the main image inpainting branch from the other two branches. Similar to (Zhang et al., 2020b), Zhang et al. (2022d) first localized the perceptual artifacts from the completed image, and then used this information to guide the iterative refinement process. They also manually annotated an inpainting artifact dataset.

2.1.2 Two-stage framework

Coarse-to-fine methods. This kind of method first applies a generator to fill the holes with coarse contents, and then refine them via the second generator, as shown in Fig. 4(a). Yu et al. (2018) modified the generative inpainting framework with cascaded coarse and refinement networks. In the refinement stage, they designed a contextual attention module modeling the long-term correlation to facilitate the inpainting process.

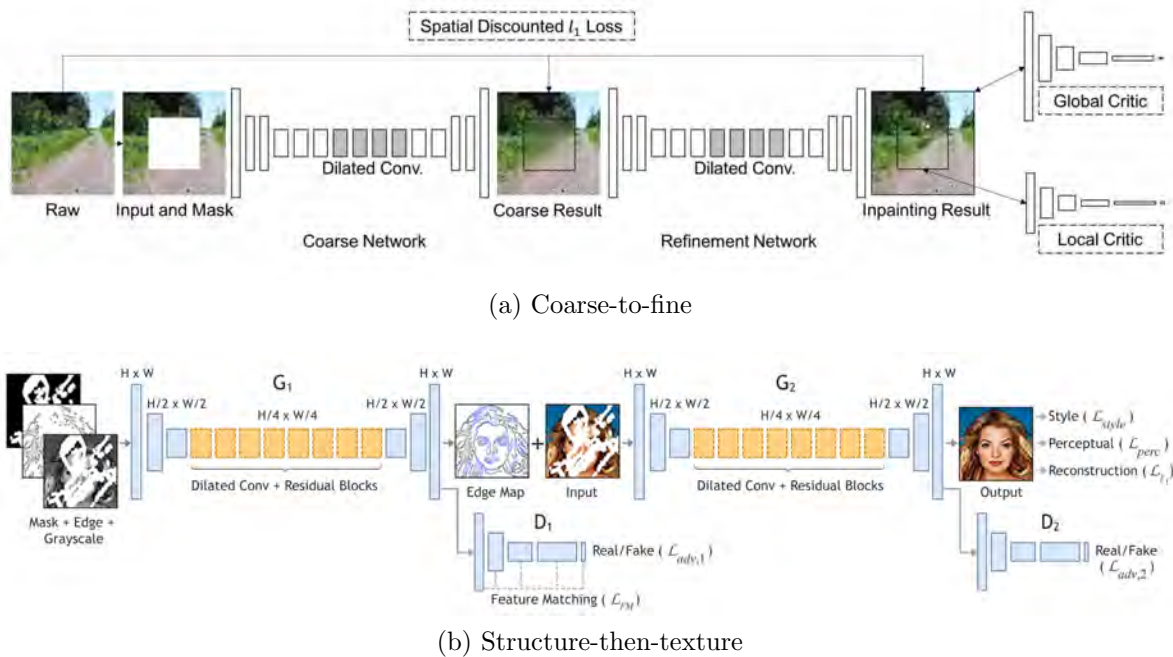


Fig. 4: Two types of the two-stage inpainting framework: (a) coarse-to-fine (Yu et al., 2018) where the first network predicts an initial coarse result and the second network predicts a refined result; (b) structure-then-texture (Nazeri et al., 2019) where the first network predicts a structure map and the second network predicts a complete image. An apparent difference between these two types is that the structure-then-texture methods explicitly predict the structure map in the first stage.

Many later works refined different aspects of this classical coarse-to-fine framework. Inspired by mask-aware convolution (Liu et al., 2018) for irregular holes, Yu et al. (2019) improved the previous network (Yu et al., 2018) by introducing gated convolution that adaptively perceives the mask location. In the coarse stage, Ma et al. (2019) proposed region-wise convolutions and a non-local operation to process the discrepancy and correlation between intact and damaged areas. PEPSI Sagong et al. (2019) modified the two-stage feature encoding processes in (Yu et al., 2018) by sharing the encoding network and organizing the coarse and fine inpainting network in a parallel manner. PEPSI can enhance the inpainting capability while reducing the number of convolution operations and computational resources. To further reduce the network parameters, Shin et al. (2021) extended PEPSI by replacing the original dilated convolutional layers (Yu and Koltun, 2016) with a so-called rate-adaptive version, which shares the weights for each layer but produces dynamic features via dilation rates-related scaling and shifting operations. The contextual attention proposed by (Yu et al., 2018) has a limited ability to model the relationships between patches inside the holes, therefore, Liu et al. (2019) introduced a coherent semantic attention layer,

which can enhance the semantic relevance and feature continuity in the attention computation of hole regions. In (Yu et al., 2018), several dilated convolutions are applied to enlarge the receptive field. Li et al. (2020b) replaced the dilated convolution with a spatial pyramid dilation ResNet block with eight different dilation rates to extract multi-scale features. Navasardyan and Ohanyan (2020) designed a patch-based onion convolution mechanism to continuously propagate information from known regions to the missing ones. This convolution mechanism can capture long-range pixel dependencies and achieve high efficiency and low latency. Wadhwa et al. (2021) proposed a hypergraph convolution with a trainable incidence matrix to generate globally semantic completed images and replaced the regular convolutions with gated convolution in the discriminator to enhance the local consistency of inpainted images.

Due to the computational overhead and the lack of supervision for the contextual attention in (Yu et al., 2018), Zeng et al. (2021b) removed this attention block and learned its patch-borrowing behavior with a so-called contextual reconstruction loss. Based on the insight that recovering different types of missing areas need a different scope of neighboring areas, Quan et al.

(2022) designed a local and global refinement network with small and large receptive fields, which can be directly applied to the end of existing networks to further enhance their inpainting capability. Kim et al. (2022) developed a coarse-super-resolution-refine pipeline, where they add a super-resolution network to reconstruct finer details after the coarse network and introduce a progressive learning mechanism to repair larger holes.

Some works adopt a coarse-to-fine framework to obtain high-resolution inpainting. Yang et al. (2017) designed a two-stage inpainting framework consisting of a content network and a texture network. The former predicts the holistic content in the low resolution (128×128) and the latter iteratively optimizes the texture details of missing regions from low to high resolution (512×512). Song et al. (2018a) developed an image-to-feature network to infer coarse results, and then designed a patch-swap method to refine the coarse features. The swapped feature map is translated to a complete image via a Feature2Image network. In addition, this framework can be directly used for high-resolution inpainting by upsampling the complete image as the input of refine stage with a multi-scale inference. Yi et al. (2020) proposed a contextual residual aggregation mechanism for ultra high-resolution image inpainting (up to 8K). Specifically, a low-resolution inpainting result was first predicted via a two-stage coarse-to-fine network and then the high-resolution result was generated by adding the large blurry image with the aggregated residuals, which are obtained by aggregating weighted high-frequency residuals from contextual patches. Zhang et al. (2022c) focused on image inpainting for 4K or more resolution. They first fill the hole via LaMa (Suvorov et al., 2022), predict depth, structure, and segmentation map from the initially completed image, then generate multiple candidates with a multiply-guided PatchMatch (Barnes et al., 2009), and finally choose a good output using the proposed auto-curation network. To complete the high-resolution image with limited resources, these methods first predicted the coarse content at the low-resolution level and then refine the texture details at the high-resolution level (sometimes with multi-scale inferences).

Other works also follow the basic coarse-to-fine strategy, but they are clearly different from the framework proposed by (Yu et al., 2018). After obtaining the coarse result with an initial prediction network, Li et al. (2019d) applied a super-resolution network as the refinement stage to produce high-frequency details. Roy et al. (2021) predicted the coarse results in the frequency domain by learning the mapping of the DFT of the corrupted image and its ground truth. Based on the insight that patch-based methods (Barnes et al., 2009; He and Sun,

2012) fill the missing regions with high-quality texture details, Xu et al. (2021) proposed a texture memory-augmented patch synthesis network with a patch distribution loss after the coarse inpainting network.

Structure-then-texture methods. Structure and texture are two important components of the image, therefore, some works decompose the image inpainting as the structure inference and the texture restoration, as shown in Fig. 4(b). Sun et al. (2018b) designed a two-stage head inpainting obfuscation network. The first stage generates facial landmarks and the second stage recovers the head image guided by the landmarks. Song et al. (2019) first estimated the facial geometry including landmark heatmaps and parsing maps, and then concatenated these results with a corrupted face image as the input of the complete network to recover face images and disentangle masks. Liao et al. (2018b) and Nazeri et al. (2019) both proposed an edge-guided image inpainting method, which first estimates the edge map for the missing regions, and then utilizes this edge map prior to predicting the texture details. Similarly, Xiong et al. (2019) explicitly disentangled the image inpainting problem into two sub-tasks of foreground contour prediction and content completion. To improve the structural guidance of coarse edge maps, Ren et al. (2019) introduced another representation of the structure, *i.e.*, the edge-preserving smoothing via filtering operation. Based on the structure reconstruction of the first network, they inpainted missing regions using appearance flow. Shao et al. (2020) combined the edge map and color aware map as the representation of the structure, where the former is captured via the Canny operator (Canny, 1986) and the latter is obtained through Gaussian blur with a large kernel. For the specific Manga inpainting, Xie et al. (2021) first completed a semantic structure map, including the structural lines and the ScreenVAE map (a point-wise representation of screen-tones) (Xie et al., 2020), using a semantic inpainting network. Then, the completed semantic map is used for guiding the appearance synthesis. Wang et al. (2021b) designed an external-internal learning inpainting framework. It first reconstructs the structures in the monochromatic space using the knowledge externally learned from large datasets. Based on internal learning, then, it applies a multi-stage network to recover the color information via iterative optimization. Besides the edge map used in (Nazeri et al., 2019), Yamashita et al. (2022) incorporated the depth image to provide the boundaries between different objects. Their method first completed the masked edge and depth images separately and then recovered the missing regions via an RGB image inpainting network taking as input the concatenation of masked images, inpainted edges, and depth images. To

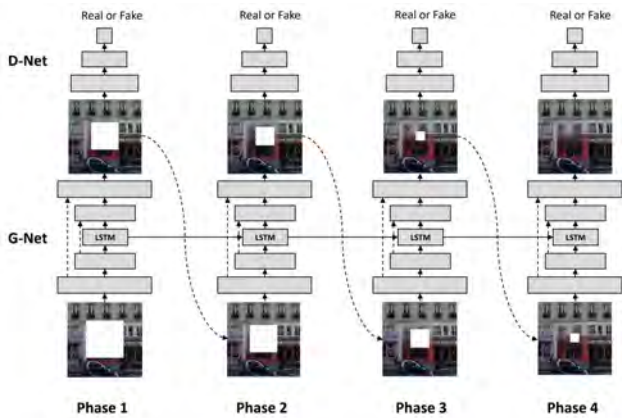


Fig. 5: Progressive image inpainting. The image comes from (Zhang et al., 2018a).

contain richer structural information, Wu et al. (2022) choose the local binary pattern (LBP) (Ojala et al., 1996, 2002), which describes the distribution information of edges, speckles, and other local features (Zhang et al., 2010). In (Wu et al., 2022), the first network infers the LBP information of the holes, and the second network with spatial attention conducts the actual image inpainting. Dong et al. (2022) utilized a transformer to complete the holistic structure in a grayscale space and proposed a masking positional encoding for large irregular masks.

In addition, semantic segmentation maps are also used as the proxy of structure (Song et al., 2018b; Qiu et al., 2021; Zhou et al., 2021). Song et al. (2018b) introduced the semantic segmentation information into the image inpainting process to improve the recovered boundary between different class regions. They first predict the segmentation map of missing regions via a U-Net and then recover the missing contents with the guidance of the above inpainted semantic map using the second generator network. Song et al. (2018b) utilized the pre-classification algorithm (Felzenszwalb and Huttenlocher, 2004) to extract a semantic structure map. After the completion of the semantic map, they employed a spatial-channel attention module to generate the texture information. Zhou et al. (2021) first predicted the complete segmentation map via a segmentation reconstructor, and then recovered fine-grained texture details with an image generator based on a relation network. The relation network is an extension of SPADE (Park et al., 2019) to better modulate features via spatially-adaptive normalization with the relation graph.

2.1.3 Progressive frameworks

Following the basic idea of traditional inpainting methods, some works have been proposed to exploit progressive inpainting with deep models. As shown in Fig. 5, the progressive methods iteratively fill in the holes from the boundary to the center of the holes, and the missing area gradually becomes smaller until it disappears. Zhang et al. (2018a) formulated image inpainting as a sequential problem, where the missing regions are filled in four inpainting phases. They designed an LSTM (long short-term memory) (Hochreiter and Schmidhuber, 1997)-based framework to string these four inpainting phases together. However, this method cannot handle irregular holes common in real-world applications. Guo et al. (2019) devised a residual architecture to progressively update irregular masks and introduced a full-resolution network to facilitate feature integration and texture reconstruction. Inspired by structure-guided inpainting methods (Nazeri et al., 2019; Xiong et al., 2019), Li et al. (2019c) proposed a progressive reconstruction with a visual structure network to incorporate structure information into the visual features step by step, which can generate a more structured image. Progressive inpainting methods have the potential to fill in large holes, however, it is still difficult due to the lack of constraints on the hole center. To handle this drawback, Li et al. (2020c) designed a recurrence feature reasoning network with consistent attention and weighted feature fusion. This network recurrently infers and gathers the hole boundaries of the feature map so as to progressively strengthen the constraints for estimating internal contents. Zeng et al. (2020b) proposed an iterative inpainting method with confidence feedback for high-resolution images. SRInpaintor (Li et al., 2022a) combined super-resolution and the transformer in a progressive pipeline. It reasons about the global structure in low resolution, and progressively refines the texture details in high resolution.

To this end, we organize the important and prevalent technical aspects for the network design, as shown in Table 1.

2.2 Stochastic Image Inpainting

Image inpainting is an underdetermined inverse problem. Therefore, multiple plausible solutions exist. We use the term stochastic image inpainting to refer to methods capable of producing multiple solutions with a random sampling process.

VAE-based methods. A variational autoencoder (VAE) (Kingma and Welling, 2014) is a generative model that combines an encoder and a decoder. The encoder

Table 1: The summary of important techniques for deep learning-based image inpainting.

Aspects	Blocks	Core idea
mask-aware convolution	Shepard interpolation (Ren et al., 2015) partial convolution (Liu et al., 2018) gated convolution (Yu et al., 2019) priority-guided partial convolution (Wang et al., 2021c) contextual attention (Yu et al., 2018)	translation variant interpolation convolution on valid regions adaptive gating structure and texture priority
Attention	coherent semantic attention (Liu et al., 2019) multi-scale attention module (Wang et al., 2019b) multi-scale attention unit (Qin et al., 2021)	background patches with high similarity to the coarse prediction correlation between patches within the hole attention with two patch sizes attention with four different dilation rates
Normalization	region normalization (Yu et al., 2020) probabilistic context normalization (Wang et al., 2020c) regional composite normalization (Wang et al., 2021a) point-wise normalization (Zhu et al., 2021) frequency region attentive normalization (Zhu et al., 2021)	spatial and region-wise transfer mean and variance batch, instance, and layer normalization mask-aware batch normalization align low- and high-frequency features
Discriminator	global discriminator (Pathak et al., 2016) local discriminator (Iizuka et al., 2017) patch-based discriminator (PatchDis) (Yu et al., 2019) conditional multi-scale discriminator (Li et al., 2020b) soft mask-guided PatchDis (Zeng et al., 2022)	entire image corrupted region dense local patches PatchDis with two different scales central parts of the missing regions

learns an appropriated latent space and the decoder transforms sampled latent representations back into new data. Zheng et al. (2019) proposed a two-branch completion network, where the reconstructive branch models the prior distribution of missing parts and reconstructs the original complete image from this distribution. The generative branch infers the latent conditional prior distribution for the missing areas. This framework is optimized by balancing the variance of the conditional distribution and the reconstruction of the original training data. Zheng et al. (2021a) extended this work by estimating the distributions in a separate training stage and introducing the patch-level short-long term attention module. For stochastic fashion image inpainting, Han et al. (2019) decomposed the inpainting process as the shape and appearance generation. The network design for these two generation tasks mainly adopts the VAE architecture. Based on a pre-trained VAE on facial images, Tu and Chen (2019) first searched for the possible set of solutions in the coding vector space for the corrupted image, and then recovers possible face images with the decoder of the VAE. Zhao et al. (2020) proposed an instance-guided conditional image-to-image translation network to learn conditional completion distribution. Specifically, they first encode the instance and masked images into two probability feature spaces, and then design a cross-semantic attention layer to fuse two feature maps. A decoder is finally used to generate the inpainted image. However, Han et al. (2019) and Zhao et al. (2020) often suffer from distorted structures and blurry textures due to the joint optimization of structure and appearance. Peng et al. (2021) designed a two-stage pipeline, where the

first stage produces multiple coarse results with different structures based on a hierarchical vector quantized variational auto-encoder, and the second stage synthesizes the texture under the guidance of the discrete structural features.

GAN-based methods. GAN (Goodfellow et al., 2014) learns the data distribution via an adversarial process. A generator is applied to transform sampled Gaussian random noise into image space and a discriminator is used to differentiate the real sample and fake sample. Based on the premise that the degree of freedom increases from the hole boundary to the hole center, Liu et al. (2021a) introduced a spatially probabilistic diversity normalization to modulate the pixel generation with diversity maps. Considering that minimizing the classical reconstruction loss hampers the diversity of results, they also proposed a perceptual diversity loss that maximizes the distance of two generated images in the feature space. By combining the image-conditional and unconditional generative architectures, Zhao et al. (2021) proposed a co-modulated GAN for large-scale image inpainting. Technically, they encode the incomplete input image into a conditional latent vector, which is then concatenated with the original style vector of StyleGAN2 (Karras et al., 2020). To enhance the diversity and control of image inpainting, Zeng et al. (2021a) applied the patch matching from the training samples on the basis of coarse inpainted results. In particular, they designed the nearest neighbor-based pixel-wise global matching (from a single image) and compositional matching (from multiple images). Inspired by CoModGAN (Zhao et al., 2021), Zheng et al. (2022b) proposed a cascaded mod-

ulation GAN, which combines the global modulation and the spatially-adaptive modulation in each scale of the decoder, and replaces the common convolution with fast Fourier convolution (Chi et al., 2020) in the encoder. To directly complete the high-resolution image, Li et al. (2022b) proposed a mask-aware transformer module with a dynamic mask updating as (Liu et al., 2018). This module conducts non-local interactions only using partially valid tokens in a shifted-window manner Liu et al. (2021d). Following (Chen et al., 2019; Karras et al., 2019), they developed a style manipulation module for stochastic generations.

Flow-based methods. Normalizing Flows (Tabak and Vandeen-Eijnden, 2010; Dinh et al., 2014; Rezende and Mohamed, 2015) are a generative method that constructs a complex probability distribution by assembling a sequence of invertible mappings. Inspired by Glow (Kingma and Dhariwal, 2018) and its conditional extension (Lugmayr et al., 2020), Wang et al. (2022a) proposed a conditional normalizing flow network to learn the probability distribution of structure priors. Then, another generator is applied to produce the final complete image with rich texture.

MLM-based methods. To produce a stochastic structure in the missing region, Yu et al. (2021b) and Wan et al. (2021) adopted a sequence prediction pipeline based on a masked language model (MLM). Yu et al. (2021b) proposed a bidirectional and auto-regressive transformer as the low-resolution stochastic-structure generator, which predicts masked token (missing regions) via a top- \mathcal{K} sampling strategy during inference. Then, a texture generator was applied to generate multiple inpainted results. Similarly, Wan et al. (2021) proposed a Transformer-CNN framework. They first apply a transformer training with MLM objective to produce a low-resolution image with pluralistic structures and some coarse textures, and then utilize an encoder-decoder network to enhance the local texture details of the high-resolution complete image.

Diffusion model-based methods. Diffusion models (DM) are emerging generative models for image synthesis. Here, we only review diffusion model-based inpainting methods, and we refer readers to the surveys (Yang et al., 2023; Croitoru et al., 2023) about a comprehensive introduction to diffusion models. Generally, diffusion-based inpainting models employ a U-Net architecture. The training objectives are usually based on $\mathcal{L}_{DM} = \mathbb{E}_{x, \epsilon \in \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_{\theta}(x_t, t)\|_2^2]$, where $t = 1 \dots T$, x_t is a noised version of x , and $\epsilon_{\theta}(\cdot, t)$ is a neural network. In the literature, existing works mainly focused on the sampling strategy design and the computational cost reduction.

(1) Sampling strategy design.

Based on an unconditionally pre-trained denoising diffusion probabilistic model (DDPM) (Ho et al., 2020), Lugmayr et al. (2022) modified the standard denoising strategy by sampling the masked regions from the diffusion model and sampling the unmasked areas from the given image. To preserve the background and improve the consistency, Xie et al. (2023) added an extra mask prediction to the diffusion model. In the inference stage, the predicted mask is used to guide the sampling process.

(2) Computational cost reduction.

Instead of applying the diffusion process in pixel space, Esser et al. (2021) utilized a multinomial diffusion process (Hoogetboom et al., 2021; Austin et al., 2021) on a discrete latent space and autoregressively factorized models for the reverse process. These designs enable ImageBART to generate high-resolution images, *e.g.*, 300×1800 . Similarly, Rombach et al. (2022) proposed a latent diffusion model (LDM) to reduce the training cost of DMs while boosting visual quality, which can be applied to the image inpainting task at a high resolution of 1024^2 pixels. To overcome the limitation of massive iterations in the diffusion model, Li et al. (2022c) proposed a spatial diffusion model (SDM) with decoupled probabilistic modeling, where the mean term refers to the inpainted result and the variance term measures the uncertainty. Instead of starting with random Gaussian noise in the reverse conditional diffusion, Chung et al. (2022) remarkably reduced the number of sampling steps with a better initialization by starting from forward-diffused data.

2.3 Text-guided Image Inpainting

Text-guided image inpainting takes an incomplete image and text description as input and generates text-aligned inpainting results. The main challenge lies in how to fuse the text and image semantic features, and how to focus on effective information in the text. Zhang et al. (2020a) proposed a dual attention mechanism to obtain the semantic feature of the masked region by finding unmatched words compared to the image and applied DAMSM loss (Xu et al., 2018b) to measure the similarity of text and image. Lin et al. (2020) introduced an image-adaptive word demand module that removes redundant information and aggregates text features in the coarse stage. They also proposed a text-guided attention loss that pays more attention to the reconstruction of the region affected by the text. Zhang et al. (2020c) encoded text and image to sequential data and exploited the transformer architecture to let cross-modal features interact. To ensure that the inpainted

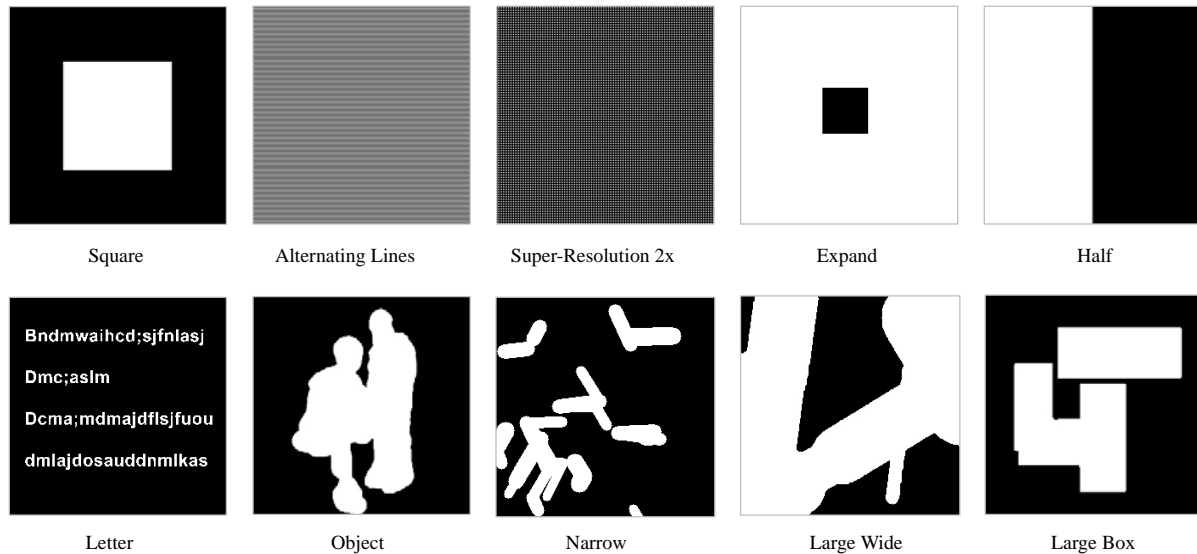


Fig. 6: Representative examples of masks.

image matches the text, they took the masked text and inpainted image as input to restore the text prompt. [Wu et al. \(2021\)](#) incorporated word-level and sentence-level textual features into a two-stage generator by introducing a dual-attention module. To eliminate the effect of the background, the mask reconstruction module was devised to recover the corrupted object mask. [Xie et al. \(2022\)](#) applied multi-head self-attention as text-image interactive encoder. They created a semantic relation graph to compute non-Euclidean semantic relations between text and image, and used graph convolution to aggregate node features. [Li et al. \(2023\)](#) followed a coarse-to-fine image inpainting framework. They first employed a visual-aware textual filtering mechanism to adaptively concentrate on required words and then inserted filtered text features into the coarse network. Unlike [\(Zhang et al., 2020c\)](#), they directly reconstructed text descriptions from inpainted images to guarantee multi-modal semantic alignment. To better preserve the non-defective regions during the text guidance, [Ni et al. \(2023\)](#) proposed a defect-free VQGAN to control receptive spreading and a sequence-to-sequence module to enable visual-language learning from multiple different perspectives, including text descriptions, low-level pixels, and high-level tokens. Recent methods are based on diffusion models. [Shukla et al. \(2023\)](#) focused on how to generate a high-quality text prompt to guide a text-to-image model-based inpainting network by analyzing inter-object relationships. They first constructed a scene graph based on object detector outputs and ex-

panded it via a graph convolution network to obtain the features of the corrupted node. Finally, the generated text prompt and masked image were fed to the diffusion model to obtain the inpainted result. [Wang et al. \(2023\)](#) found that object masks would force the inpainted images to rely more on text descriptions instead of the random mask. Then, they proposed Imagen Editor fine-tuned from Imagen ([Saharia et al., 2022b](#)) with a new convolutional layer and designed an object masking strategy for better training. To facilitate the systematic evaluation of text-guided image inpainting, they established a benchmark called EditBench.

2.4 Inpainting Mask

In the development of image inpainting techniques, various artificial masks have been introduced. These masks can be roughly divided into two categories: regular masks and irregular masks. Fig. 6 summarizes these masks, where white pixels indicate missing regions.

Regular masks. A *square hole* that blocks the center area or random location are generally easier to construct. [Lugmayr et al. \(2022\)](#) introduced more regular masks, including *Super-Resolution 2 \times* (reserving pixels with a stride of 2), *Alternating lines* (removing every second row), *Expand* (leaving a small center crop of the input image), and *Half* (masking the half of the input image).

Irregular masks. *Letter masks* ((Bertalmio et al., 2000; Bian et al., 2022)) and *object-shaped masks* ((Criminisi et al., 2004; Yi et al., 2020)) are particularly designed for specific tasks, for example, caption removal and object removal. Liu et al. (2018) introduced free-form masks, where the former collected random streaks and arbitrary holes from the results of the occlusion/disocclusion mask estimation method. The irregular masks shared by (Liu et al., 2018) are very common in the existing inpainting methods. Suvorov et al. (2022) further split free-form masks into *narrow masks*, *large wide masks*, and *large box masks*, where two types of large masks are generated via an aggressive mask method sampling polygonal chains with a high random width and rectangles of random aspect ratios, respectively.

2.5 Loss Functions

For image inpainting, the loss functions affect features of different sizes. At the lowest level, a pixel reconstruction loss aims to recover the exact pixel values. We further discuss the total-variational (TV) loss (Rudin et al., 1992), feature consistency loss, the perceptual loss (Johnson et al., 2016), style loss (Gatys et al., 2016), and adversarial loss (Goodfellow et al., 2014).

As input, an inpainting network accepts an input image \mathbf{I}_{in} and a binary mask \mathbf{M} describing the missing regions (where 0 means the valid pixel and 1 means the missing pixel). The output of the network is a complete image \mathbf{I}_{out} . The loss functions are formulated as follows.

Pixel-wise reconstruction loss. In the literature, the pixel-wise reconstruction loss often has two types: ℓ_1 loss (Eq. (1)) and weighted ℓ_1 loss (Eq. (3)). The key point is how the valid and unknown regions differ in the loss function. The detailed formulations are as follows,

$$\mathcal{L}_{wpr} = \|(\mathbf{I}_{gt} - \mathbf{I}_{out})\|_1. \quad (1)$$

where \mathbf{I}_{gt} is the ground-truth complete image.

$$\begin{aligned} \mathcal{L}_{valid} &= \frac{1}{\text{sum}(\mathbb{1} - \mathbf{M})} \|(\mathbf{I}_{gt} - \mathbf{I}_{out}) \odot (\mathbb{1} - \mathbf{M})\|_1, \\ \mathcal{L}_{hole} &= \frac{1}{\text{sum}(\mathbf{M})} \|(\mathbf{I}_{gt} - \mathbf{I}_{out}) \odot \mathbf{M}\|_1, \end{aligned} \quad (2)$$

where \odot is the element-wise product operation, and $\text{sum}(\mathbf{M})$ is the number of non-zero elements in \mathbf{M} . Then the weighted ℓ_1 loss is formulated as

$$\mathcal{L}_{pr} = \mathcal{L}_{valid} + \alpha \cdot \mathcal{L}_{hole}, \quad (3)$$

where α is the balancing factor. It is well known that the ℓ_1 loss can capture the low-frequency components, whereas it struggles to restore the high-frequency components (Isola et al., 2017; Ledig et al., 2017).

Total-variation loss. Total-variation loss can be applied to ameliorate the potential checkerboard artifacts introduced by the perceptual loss. The formulation is:

$$\begin{aligned} \mathcal{L}_{tv} &= \|\mathbf{I}_{mer}(i, j + 1) - \mathbf{I}_{mer}(i, j)\|_1 \\ &\quad + \|\mathbf{I}_{mer}(i + 1, j) - \mathbf{I}_{mer}(i, j)\|_1. \end{aligned} \quad (4)$$

where $\mathbf{I}_{mer} = \mathbf{I}_{out} \odot \mathbf{M} + \mathbf{I}_{gt} \odot (\mathbb{1} - \mathbf{M})$ is the merged (completed) image.

Feature consistency loss. This loss constrains extracted feature maps of the prediction with guidance from ground truth images:

$$\mathcal{L}_{fc} = \sum_{y \in \Omega} \|\Phi_m(\mathbf{I}_{in})_y - \Phi_n(\mathbf{I}_{gt})_y\|_2^2. \quad (5)$$

where Ω is the missing regions, $\Phi_m(\cdot)$ is the feature map of the selected layer in the inpainting network, and $\Phi_n(\cdot)$ is the feature map of the corresponding layer in the inpainting network or pre-trained VGG models. $\Phi_m(\cdot)$ and $\Phi_n(\cdot)$ must have the same shape.

Perceptual loss. The perceptual loss is first proposed in style transfer and super-resolution tasks. This loss measures the semantic/content difference between inpainted and ground-truth images, and thus encourages the inpainting generator to restore the semantics of missing regions. The perceptual loss is computed in high-level feature representations and is formulated as:

$$\mathcal{L}_{per} = \sum_i \|\Psi_i(\mathbf{I}_{out}) - \Psi_i(\mathbf{I}_{gt})\|_1 + \|\Psi_i(\mathbf{I}_{mer}) - \Psi_i(\mathbf{I}_{gt})\|_1, \quad (6)$$

where $\Psi_i(*)$ is the feature map of i -th layer in the VGG-16/19 network (Simonyan and Zisserman, 2014) pre-trained on ImageNet (Deng et al., 2009). Instead of using the common VGG network, Suvorov et al. (2022) suggested using a base network with a fast-growing receptive field for large-mask inpainting and utilized the pre-trained segmentation network (ResNet50 with dilated convolutions (Zhou et al., 2018)) to compute the so-called high receptive field perceptual loss. Note that, some works used 2-norm in Eq. (6) to compute perceptual loss.

Style loss. Similar to the perceptual loss, the style loss also depends on higher-level features extracted from a pre-trained network. This loss is applied to penalize the style difference between inpainted and ground-truth images, *e.g.*, texture details and common patterns. Mathematically, the style loss measures the similarities of Gram matrices of image features, instead of



Fig. 7: Some examples of image inpainting datasets.

the feature reconstruction in the perceptual loss. The detailed formulation can be written,

$$\mathcal{L}_{sty} = \sum_i \|\Phi_i(\mathbf{I}_{out}) - \Phi_i(\mathbf{I}_{gt})\|_1 + \|\Phi_i(\mathbf{I}_{mer}) - \Phi_i(\mathbf{I}_{gt})\|_1, \quad (7)$$

where $\Phi_i(\cdot) = \Psi_i(\cdot)\Psi_i(\cdot)^T$ is the Gram matrix (Gatys et al., 2016).

Besides using Gram matrices to model the style information, the mean and standard deviation of image features are commonly used in style transfer (Huang and Belongie, 2017; Deng et al., 2020). The formulation is written as,

$$\begin{aligned} \mathcal{L}_{sty_mean} &= \sum_i \|\mu(\Psi_i(\mathbf{I}_{out})) - \mu(\Psi_i(\mathbf{I}_{gt}))\|_2 \\ &\quad + \|\mu(\Psi_i(\mathbf{I}_{mer})) - \mu(\Psi_i(\mathbf{I}_{gt}))\|_2, \\ \mathcal{L}_{sty_std} &= \sum_i \|\sigma(\Psi_i(\mathbf{I}_{out})) - \sigma(\Psi_i(\mathbf{I}_{gt}))\|_2 \\ &\quad + \|\sigma(\Psi_i(\mathbf{I}_{mer})) - \sigma(\Psi_i(\mathbf{I}_{gt}))\|_2, \\ \mathcal{L}_{sty_meanstd} &= \mathcal{L}_{sty_mean} + \mathcal{L}_{sty_std}, \end{aligned} \quad (8)$$

where $\mu(*)$, $\sigma(*)$ are the mean and standard deviation, computed over spatial dimensions independently for each sample and each channel.

Adversarial loss. GANs (Goodfellow et al., 2014) are widely used in many image generation tasks. They employ an adversarial loss to force the output distribution to be close to the “real” distribution. The adversarial loss can counteract blurry results and enhance the visual realism of the output image. Therefore, it is often applied in GAN-based inpainting networks. To compute the adversarial loss, a discriminator network (D) is necessary, which interacts with the generator network (G). The hinge version (Lim and Ye, 2017) of the

adversarial loss can be formulated as:

$$\begin{aligned} \mathcal{L}_D &= \mathbb{E}_{\mathbf{I} \sim p_{data}(\mathbf{I})} [\max(0, 1 - D(\mathbf{I}_{gt}))] \\ &\quad + \mathbb{E}_{\mathbf{I}_{mer} \sim p_{\mathbf{I}_{mer}}(\mathbf{I}_{mer})} [\max(0, 1 + D(\mathbf{I}_{mer}))], \end{aligned} \quad (9)$$

where $D(\mathbf{I}_{gt})$ and $D(\mathbf{I}_{mer})$ are the logits output from discriminator D . The objective function for generator G can be denoted as:

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{I}_{mer} \sim p_{\mathbf{I}_{mer}}(\mathbf{I}_{mer})} [D(\mathbf{I}_{mer})]. \quad (10)$$

Except for the above hinge version, other types of adversarial losses are also adopted: GAN (Goodfellow et al., 2014), WGAN (Arjovsky et al., 2017), LSGAN (Mao et al., 2017), etc.

2.6 Datasets

In the literature, there are six prevalent and public datasets for evaluating image inpainting. These datasets cover various types of images, including faces (CelebA and CelebA-HQ), real-world encountered scenes (Places2), street scenes (Paris), texture (DTD), and objects (ImageNet). Several examples are shown in Fig. 7. The details of the datasets are described below.

- CelebA dataset (Liu et al., 2015): A large-scale face attribute dataset that contains 10,177 identities, each of which has about 20 images. In total, CelebA has 202,599 face images, each with 40 attribute annotations.
- CelebA-HQ dataset (Karras et al., 2018): The high-quality version of CelebA (Liu et al., 2015) with JPEG artifacts removal, super-resolution operation, and cropping, etc. This dataset consists of 30,000 face images.

- Places2 dataset (Zhou et al., 2017): A large-scale scene recognition dataset. Places365-Standard has 365 scene categories. The training set has 1,803,460 images and the validation set contains 18,250 images.
- Paris StreetView dataset (Doersch et al., 2012): This dataset consists of street-level imagery. It contains 14,900 images for training and 100 images for testing.
- DTD dataset (Cimpoi et al., 2014): A describable texture dataset consisting of 5,640 images. According to human perception, these images are divided into 47 categories with 120 images per category.
- ImageNet dataset (Deng et al., 2009): A large-scale benchmark for object category classification. There are about 1.2 million training images and 50 thousand validation images.

2.7 Evaluation Protocol

The evaluation metrics can be classified into two categories: pixel-aware metrics and (human) perception-aware metrics. The former focus on the precision of reconstructed pixels, including ℓ_1 error, ℓ_2 error, and PSNR (peak signal-to-noise ratio), SSIM (the structural similarity index) (Wang et al., 2004), and MS-SSIM (multi-scale SSIM) (Wang et al., 2003). The latter pay more attention to the visual perception quality, including FID (Fréchet inception distance) (Heusel et al., 2017), LPIPS (learned perceptual image patch similarity) (Zhang et al., 2018b), P/U-IDS (paired/unpaired inception discriminative score) (Zhao et al., 2021), and user study results. The detailed descriptions are given in the following.

- ℓ_1 error: The mean absolute differences between the complete image (\mathbf{I}_c) and the ground-truth image (\mathbf{I}_g).
- ℓ_2 error: The mean squared differences between the complete image and the ground-truth image.
- PSNR: It is mainly used to measure the quality of reconstruction of the complete image. Its formulation is $\text{PSNR} = 20 \cdot \log_{10}(255) - 10 \cdot \log_{10}(\text{MSE})$, where MSE is the mean squared error between the complete image and the ground-truth image.
- SSIM: Instead of estimating absolute errors, SSIM measures the similarity in structural information by incorporating luminance masking and contrast masking. It is written as $\text{SSIM} = \frac{(2\mu_{\mathbf{I}_c}\mu_{\mathbf{I}_g} + c_1)(2\sigma_{\mathbf{I}_c\mathbf{I}_g} + c_2)}{(\mu_{\mathbf{I}_c}^2 + \mu_{\mathbf{I}_g}^2 + c_1)(\sigma_{\mathbf{I}_c}^2 + \sigma_{\mathbf{I}_g}^2 + c_2)}$, where μ and σ refer to the average and the variance, respectively; and $c_1 = 0.01^2$ and $c_2 = 0.03^2$ are two variables to stabilize the division.
- MS-SSIM: Dosselmann and Yang (2011) illustrated that SSIM is very close to the windowed mean squared error and Wang et al. (2003) highlighted the single-scale nature of SSIM as a drawback. As an alternative, MS-SSIM embraces more flexibility for image quality assessment. To compute the MS-SSIM, two input images are iteratively processed with low-pass filters and downsampled with a stride of 2 (in total, five scales). Then, the contrast comparison and structure comparison are computed at each scale and the luminance comparison is calculated at the last scale. These measurements are combined with appropriate weights (Wang et al., 2003).
- FID: The Fréchet inception distance compares two sets of images. It computes a Gaussian with mean and covariance (\mathbf{m}, \mathbf{C}) and a Gaussian ($\mathbf{m}_g, \mathbf{C}_g$) from deep features of the set of completed images and the set of ground-truth images. Specifically, FID is defined as $\text{FID} = \|\mathbf{m} - \mathbf{m}_g\|_2^2 + \text{Tr}(\mathbf{C} + \mathbf{C}_g - 2(\mathbf{C}\mathbf{C}_g)^{\frac{1}{2}})$.
- LPIPS: The distance of multi-layer deep features of complete and ground-truth images. Let $\mathbf{F}_c, \mathbf{F}_g \in \mathbb{R}^{H_l \times W_l \times C_l}$ denote the channel-wise normalized features in the l -th layer, the LPIPS is given by $\text{LPIPS} = \sum_l \frac{1}{H_l W_l} \sum_{h,w} \|\mathbf{W}_l \odot (\mathbf{F}_{chw}^l - \mathbf{F}_{ghw}^l)\|_2^2$, where $\mathbf{W}_l \in \mathbb{R}^{C_l}$ is the channel weight vector.
- P/U-IDS: The linear separability of complete and ground-truth images in a pre-trained feature space. Let $\phi(\cdot)$ denote the Inception v3 model mapping the image to the 2048D feature space, $f(\cdot)$ be the decision function of the SVM, the P-IDS is formulated as $\text{P-IDS} = \Pr\{f(\phi(\mathbf{I}_c)) > f(\phi(\mathbf{I}_g))\}$. Due to the unpaired nature, U-IDS is obtained by directly calculating the misclassification rate.
- User Study: FID, LPIPS, and P/U-IDS cannot be able to comprehensively evaluate the visual quality of complete images, therefore, a user study is often conducted to complement the above metrics. User studies typically let a human chooses a preferred image among two (or multiple) images generated from two (or multiple) competitors. Based on the collected votes, the preference ratio is calculated for comparison.

2.8 Performance Evaluation

2.8.1 Representative Image Inpainting Methods

We qualitatively and quantitatively compare some representative image inpainting methods: RFR (Li et al., 2020c), MADF (Zhu et al., 2021), DSI (Peng et al., 2021), CR-Fill (Zeng et al., 2021b), CoModGAN (Zhao

Table 2: Quantitative comparison of several representative image inpainting methods on CelebA-HQ and Places2. ‡ Higher is better. † Lower is better. From M1 to M6, the mask ratios are 1%-10%, 10%-20%, 20%-30%, 30%-40%, 40%-50%, and 50%-60%, respectively. Because of the heavy inference time, we do not show the results of RePaint for M1, M2, M4, and M6.

	Dataset	CelebA-HQ						Places2					
	Mask	M1	M2	M3	M4	M5	M6	M1	M2	M3	M4	M5	M6
ℓ_1 (%) †	RFR	1.59	2.47	3.58	4.90	6.44	9.47	0.83	2.20	3.93	5.83	7.96	11.37
	MADF	0.47	1.30	2.40	3.72	5.26	8.43	0.80	2.18	3.96	5.91	8.10	11.68
	DSI	0.60	1.65	3.08	4.80	6.83	11.11	0.88	2.42	4.48	6.75	9.32	13.82
	CR-Fill	0.79	2.15	3.95	6.01	8.33	13.18	0.78	2.17	4.02	6.11	8.46	12.43
	CoModGAN	0.48	1.38	2.66	4.28	6.20	10.53	0.72	2.05	3.83	5.89	8.27	12.58
	LGNet	0.46	1.28	2.38	3.72	5.27	8.38	0.68	1.89	3.51	5.33	7.41	10.86
	MAT	0.83	1.74	3.00	4.52	6.30	9.98	1.07	2.53	4.48	6.69	9.20	13.70
	RePaint	-	-	3.37	-	7.47	-	-	-	4.96	-	10.01	15.27
PSNR †	RFR	36.39	31.87	29.07	26.87	25.09	22.51	35.74	30.24	27.24	25.13	23.48	21.33
	MADF	39.68	33.77	30.42	27.95	25.99	23.07	36.17	30.37	27.17	25.00	23.31	21.10
	DSI	37.68	31.74	28.39	25.88	23.91	20.87	35.40	29.47	26.15	23.91	22.19	19.75
	CR-Fill	35.67	29.87	26.60	24.29	22.53	19.70	36.35	30.32	26.96	24.63	22.85	20.50
	CoModGAN	39.56	33.15	29.41	26.62	24.49	21.16	37.00	30.82	27.35	24.92	23.05	20.43
	LGNet	40.04	33.99	30.54	27.99	26.01	23.12	37.62	31.61	28.18	25.84	24.05	21.69
	MAT	38.44	32.62	29.21	26.70	24.72	21.78	35.66	29.76	26.41	24.09	22.30	19.81
	RePaint	-	-	28.38	-	23.16	-	-	-	26.04	-	21.72	18.99
SSIM †	RFR	0.991	0.976	0.957	0.932	0.902	0.834	0.983	0.952	0.911	0.862	0.805	0.699
	MADF	0.995	0.984	0.967	0.945	0.917	0.848	0.984	0.953	0.910	0.859	0.800	0.690
	DSI	0.992	0.976	0.951	0.918	0.877	0.778	0.982	0.945	0.892	0.832	0.763	0.636
	CR-Fill	0.988	0.965	0.931	0.890	0.842	0.729	0.985	0.954	0.909	0.855	0.794	0.675
	CoModGAN	0.994	0.981	0.960	0.929	0.891	0.792	0.987	0.957	0.914	0.860	0.796	0.671
	LGNet	0.995	0.985	0.968	0.945	0.917	0.849	0.988	0.963	0.925	0.878	0.823	0.714
	MAT	0.993	0.980	0.959	0.931	0.897	0.814	0.983	0.948	0.898	0.839	0.772	0.645
	RePaint	-	-	0.952	-	0.867	-	-	-	0.892	-	0.750	0.606
MS-SSIM †	RFR	0.992	0.976	0.956	0.933	0.900	0.830	0.986	0.960	0.924	0.880	0.828	0.731
	MADF	0.994	0.983	0.966	0.942	0.913	0.846	0.987	0.961	0.923	0.877	0.824	0.722
	DSI	0.992	0.976	0.952	0.919	0.878	0.784	0.984	0.952	0.905	0.850	0.785	0.664
	CR-Fill	0.987	0.963	0.928	0.887	0.839	0.732	0.987	0.960	0.920	0.872	0.814	0.704
	CoModGAN	0.994	0.980	0.958	0.926	0.888	0.793	0.988	0.961	0.921	0.870	0.810	0.692
	LGNet	0.995	0.984	0.968	0.945	0.917	0.851	0.990	0.968	0.935	0.894	0.844	0.744
	MAT	0.994	0.980	0.960	0.932	0.898	0.818	0.986	0.957	0.913	0.859	0.796	0.676
	RePaint	-	-	0.953	-	0.870	-	-	-	0.903	-	0.771	0.633
FID †	RFR	0.86	1.68	2.67	3.77	5.21	7.60	2.62	5.99	9.47	12.90	16.62	22.13
	MADF	0.52	1.55	3.28	5.43	8.35	13.54	2.15	5.58	9.20	13.08	17.36	24.42
	DSI	0.59	1.58	3.01	4.50	6.51	9.76	2.51	6.52	11.35	15.99	21.75	29.38
	CR-Fill	1.06	2.86	5.26	7.79	11.23	19.52	2.37	6.24	10.54	15.17	20.36	26.43
	CoModGAN	0.44	1.25	2.45	3.65	5.03	6.89	2.11	5.63	9.58	13.65	17.68	22.58
	LGNet	0.39	1.06	2.08	3.16	4.61	7.07	1.97	5.25	8.90	13.02	17.60	25.99
	MAT	0.41	1.13	2.05	2.96	4.05	5.43	2.13	5.47	9.26	13.00	16.62	21.88
	RePaint	-	-	2.14	-	4.24	-	-	-	8.85	-	15.90	21.58
LPIPS †	RFR	0.015	0.028	0.042	0.060	0.081	0.118	0.021	0.047	0.074	0.106	0.142	0.201
	MADF	0.009	0.025	0.048	0.077	0.109	0.168	0.014	0.038	0.068	0.102	0.141	0.209
	DSI	0.010	0.026	0.048	0.074	0.104	0.160	0.018	0.047	0.085	0.125	0.169	0.242
	CR-Fill	0.017	0.043	0.074	0.107	0.143	0.212	0.016	0.042	0.076	0.114	0.156	0.226
	CoModGAN	0.008	0.022	0.041	0.065	0.092	0.143	0.016	0.044	0.080	0.121	0.164	0.236
	LGNet	0.006	0.017	0.031	0.048	0.069	0.108	0.014	0.035	0.064	0.096	0.132	0.198
	MAT	0.007	0.019	0.035	0.054	0.077	0.120	0.014	0.040	0.073	0.111	0.152	0.224
	RePaint	-	-	0.038	-	0.093	-	-	-	0.077	-	0.167	0.259



Fig. 8: Qualitative comparison of representative image inpainting methods on CelebA-HQ (the first three rows) and Places2 (the last four rows).

Table 3: Model computational complexity statistics.

Model	#Parameter	GPU Memory	Infer. time
RFR	30.59 M	1.23 G	28.95 ms
MADF	85.14 M	2.42 G	15.59 ms
DSI	70.32 M	6.54 G	40.20 s
CR-Fill	4.10 M	0.96 G	9.18 ms
CoModGAN	79.80 M	1.71 G	42.24 ms
LGNet	115.00 M	1.52 G	13.59 ms
MAT	59.78 M	1.69 G	78.35 ms
RePaint	552.81 M	4.14 G	6 min 30 s

et al., 2021), LGNet (Quan et al., 2022), MAT (Li et al., 2022b), RePaint (Lugmayr et al., 2022). The test mask is from (Liu et al., 2018). Specifically, RFR follows a progressive inpainting strategy, MADF adopts a mask-aware design, DSI generates stochastic structures

with hierarchical vq-vae, CR-Fill designs an attention-free generator, CoModGAN embeds the known content of corrupted images into style vectors of styleGAN2, LGNet introduces local and global refinement networks with different receptive fields, MAT designs a mask-aware transformer architecture, and RePaint utilizes a pre-trained unconditional diffusion model.

Table 2 reports the quantitative results of these advanced image inpainting methods on CelebA-HQ and Places2 datasets. In this experiment, we use the irregular masks shared by (Liu et al., 2018) for the evaluation. From this table, we can find that MS-SSIM is very close to SSIM in the CelebA-HQ dataset; MS-SSIM is consistently higher than SSIM in the Places2 dataset and this phenomenon is more apparent for large masks. The reason may be that face images are rel-

Table 4: Quantitative comparison of different loss functions on CelebA-HQ (“C”) and Paris StreetView (“P”). † Higher is better. ‡ Lower is better. “16” refers to Eq. (3) with $\alpha = 6$, and the remaining loss settings both include “16” (We omit it for simplicity). “percept” refers to Eq. (6) based on pretrained VGG16; “resnetpl” refers to Eq. (6) based on the pre-trained segmentation network ResNet50, which is proposed by (Suvorov et al., 2022).; “style” refers to Eq. (7); “stylemeanstd” refers to Eq. (8); “percept_style” refers to Eq. (6) plus Eq. (7); “lsgan” refers to Eq. (9) and (10). Different percentage numbers in the first row refer to the hole ratios, where a large number implies large missing regions. Following the common setting, the test mask is from (Liu et al., 2018).

	Mask	1%-10%		10%-20%		20%-30%		30%-40%		40%-50%		50%-60%	
	Dataset	C	P	C	P	C	P	C	P	C	P	C	P
ℓ_1 (%) †	16	0.46	0.57	1.26	1.53	2.34	2.81	3.63	4.25	5.14	5.93	8.37	9.07
	percept	0.45	0.57	1.24	1.53	2.30	2.81	3.58	4.26	5.08	5.95	8.33	9.11
	resnetpl	0.45	0.59	1.25	1.58	2.32	2.88	3.60	4.35	5.10	6.06	8.35	9.21
	style	0.48	0.59	1.33	1.60	2.47	2.97	3.85	4.53	5.45	6.37	8.86	9.78
	stylemeanstd	0.46	0.59	1.27	1.60	2.39	2.96	3.75	4.51	5.35	6.34	8.80	9.75
	percept_style	0.47	0.60	1.30	1.63	2.43	3.00	3.79	4.58	5.40	6.42	8.83	9.84
	lsgan	0.47	0.59	1.30	1.61	2.42	2.97	3.76	4.52	5.33	6.34	8.67	9.72
PSNR †	16	40.03	38.74	34.13	33.17	30.76	29.92	28.27	27.67	26.29	25.88	23.23	23.28
	percept	40.14	38.77	34.19	33.17	30.81	29.91	28.31	27.65	26.33	25.85	23.23	23.25
	resnetpl	40.12	38.56	34.18	33.03	30.80	29.83	28.29	27.61	26.31	25.83	23.23	23.26
	style	39.63	38.36	33.65	32.71	30.24	29.37	27.72	27.07	25.74	25.22	22.71	22.62
	stylemeanstd	39.91	38.38	33.89	32.81	30.42	29.49	27.85	27.20	25.83	25.35	22.72	22.70
	percept_style	39.78	38.20	33.74	32.60	30.31	29.30	27.76	27.01	25.76	25.20	22.68	22.58
	lsgan	39.71	38.40	33.73	32.78	30.39	29.50	27.91	27.18	25.96	25.35	22.95	22.72
SSIM †	16	0.995	0.991	0.985	0.973	0.969	0.946	0.948	0.911	0.921	0.867	0.847	0.767
	percept	0.995	0.991	0.985	0.973	0.970	0.946	0.949	0.911	0.921	0.866	0.847	0.765
	resnetpl	0.995	0.991	0.985	0.972	0.970	0.945	0.948	0.909	0.921	0.865	0.848	0.764
	style	0.995	0.991	0.983	0.971	0.966	0.940	0.943	0.902	0.913	0.853	0.834	0.746
	stylemeanstd	0.995	0.991	0.984	0.971	0.968	0.941	0.944	0.903	0.914	0.854	0.835	0.747
	percept_style	0.995	0.990	0.984	0.970	0.967	0.940	0.943	0.901	0.913	0.852	0.834	0.745
	lsgan	0.995	0.991	0.984	0.971	0.967	0.941	0.944	0.903	0.915	0.854	0.839	0.746
FID †	16	0.56	4.74	1.57	13.74	3.31	26.55	5.38	40.79	8.37	57.49	15.18	86.51
	percept	0.53	4.64	1.51	13.41	3.20	26.13	5.22	40.35	8.18	57.22	14.63	88.10
	resnetpl	0.52	4.62	1.47	13.23	3.11	25.60	5.13	39.08	7.99	54.75	13.81	83.93
	style	0.42	3.91	1.13	10.67	2.25	19.65	3.38	28.87	5.00	39.09	7.90	57.00
	stylemeanstd	0.44	4.21	1.21	11.42	2.38	20.54	3.65	29.68	5.36	39.59	8.55	56.38
	percept_style	0.40	3.98	1.13	10.91	2.26	19.76	3.42	29.12	5.07	39.28	7.87	57.07
	lsgan	0.54	4.26	1.57	11.72	3.34	21.54	5.57	31.21	8.85	41.80	16.03	60.01
LPIPS †	16	0.011	0.016	0.032	0.048	0.063	0.091	0.102	0.142	0.144	0.197	0.222	0.298
	percept	0.010	0.015	0.029	0.045	0.057	0.086	0.092	0.134	0.129	0.185	0.200	0.279
	resnetpl	0.009	0.015	0.027	0.044	0.052	0.083	0.083	0.126	0.116	0.174	0.174	0.259
	style	0.007	0.011	0.018	0.031	0.033	0.056	0.052	0.086	0.073	0.120	0.117	0.186
	stylemeanstd	0.008	0.013	0.021	0.035	0.037	0.062	0.055	0.092	0.076	0.125	0.119	0.189
	percept_style	0.006	0.011	0.018	0.032	0.033	0.057	0.052	0.087	0.074	0.122	0.118	0.188
	lsgan	0.009	0.013	0.028	0.036	0.054	0.066	0.086	0.098	0.123	0.135	0.196	0.208

atively regular and uniform compared to the natural scene images in Places2, and thus the latter is more sensitive to structural similarity with different scales. Among these methods, MAT and RePaint have relatively superior FID, especially for large masks ($> 30\%$), while CoModGAN and LGNet perform better in PSNR. For DSI, the inpainting performance on CelebA-HQ is slightly better than that on Places2, and the possible reason is that the structure of face images is easier to model than diverse natural scene images. CR-Fill has limited inpainting performance.

Fig. 8 shows the visual results of some representative image inpainting methods on CelebA-HQ and Places2 datasets. MADF adopts a mask-aware design, which can predict reasonable structures (the second and third rows), but has limited ability for detail restoration. By contrast, MAT has better inpainting performance with mask-aware transformer blocks. Through introducing local and global refinement with different receptive fields, LGNet can perceive local details (the black stroke in the first row) and global structure (the second row). For large missing regions, RFR can recover the helicopter rotor blade (the fourth row) and

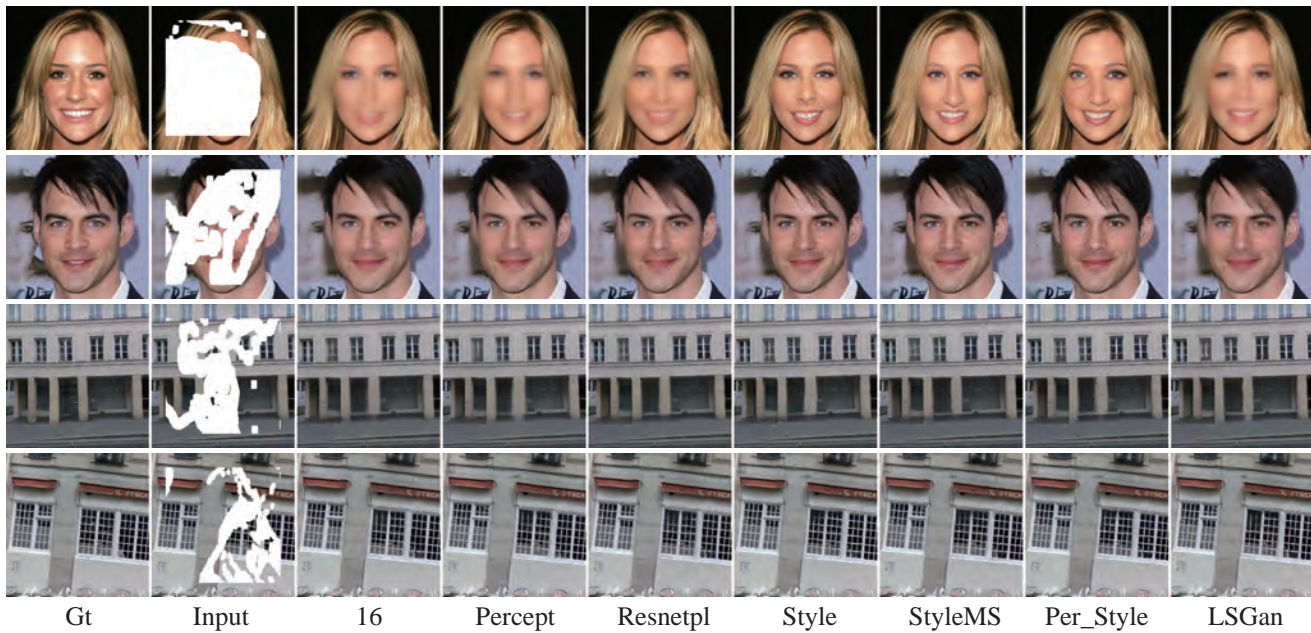


Fig. 9: Qualitative comparison of different loss functions on CelebA-HQ (the first two rows) and Paris StreetView (the last two rows). “StyleMS” refers to “stylemeanstd”; “Per_Style” refers to “Percept_style”.

waterfall (the fifth row) with progressive inpainting. With the help of the generative capability of the unconditional modulated model (StyleGAN2), CoModGAN demonstrates relatively good inpainting performance (the fourth and sixth rows). DSI can perceive the structure with hierarchical VQ-VAE (the third and fourth rows). Based on the powerful generation ability of the diffusion model, RePaint can correctly infer the missing background (the sixth row) and the human body (the seventh row). Interestingly, it may have an incorrect semantic prediction (a woman’s head in the waterfall of the fifth row). Due to the implicit attention mechanism and simple network, CR-Fill achieves comparatively inferior inpainted results, which is also consistent with the quantitative comparisons as shown in Table 2.

In addition, we evaluate the computational complexity of the representative inpainting methods in terms of the number of parameters, GPU memory of single image inference, and inference time on a GPU (the time of a forward pass through the networks). The statistical results are shown in Table 3. CR-Fill implicitly learns the patch-borrowing behavior without an attention layer, its model is the smallest and thus needs less GPU memory and running time. Because RePaint is based on a diffusion model, it has the largest number of parameters and a very long inference time. The GPU memory and inference time of DSI are also very high. LGNet follows a coarse-to-fine framework with local and global refinement, therefore, the number of parameters is high.

The running time of MAT and CoModGAN is relatively high because the former conducts many attention computations and the latter has multiple style modulations with progressive growing. RFR and MADF are in the middle.

2.8.2 Loss Functions

As summarized in Sec. 2.5, many loss functions have been proposed for image inpainting. In this part, we evaluate the effect of each loss term. We train an inpainting network with different loss settings on the CelebA-HQ and Paris StreetView datasets. This network consists of two downsampling layers, 11 ResNet residual blocks with dilation, and two upsampling layers. The corresponding numerical results are reported in Table 4. In the case of masks at 1%-10%, the SSIM values of different loss settings are (almost) the same for CelebA-HQ and Paris StreetView datasets. The reason is that different loss settings only have a slight impact on the inpainting of very small missing regions. We can see that pixel-wise reconstruction loss (“16”) provides the baseline performance. After adding the perceptual loss (“percept”), FID and LPIPS are improved. Compared with “percept”, “resnetpl” achieves slightly better results, especially for the large mask. The style loss can remarkably decrease the FID and LPIPS at the expense of PSNR and SSIM. In other words, there exists a trade-off between pixel-wise reconstruction loss




Fig. 10: Two representative examples of object removal.

and style loss, where the former focuses on low-level pixel recovery, and the latter emphasizes visual quality. A similar finding is reported and studied in (Blau and Michaeli, 2018). In addition, combining the perceptual loss with style loss (“percept_style”) has a very slight effect on the results compared to only style loss (“style”). The style loss based on Gram matrix (“style”) and style loss based on mean and standard deviation (“stylemeanstd”) have comparable results. Comparing with adversarial loss (“lsgan”), style loss (“style”) obtain significantly lower FID and LPIPS.

Fig. 9 illustrates the corresponding qualitative comparison. “16” fills the missing regions with smooth structures and textures. After introducing the content loss, this phenomenon is slightly improved, for example, the nose and mouth of the first row are better recovered in column “Percept”. Compared with “Percept”, the inpainted results of “Resnetpl” have slightly improved visual quality, which is attributed to the perceptual loss computation with higher receptive field (Suvorov et al., 2022). We can find that the results of “Style” are significantly superior to the previous three columns, especially for the restoration of texture details. This is consistent with the numerical results in Table 4. For three settings with style loss, *i.e.*, “Style”, “StyleMS”, and “Per_Style”, “Style” and “Per_Style” are on par, “StyleMS” is slightly worse. The performance of “LS-Gan” is in between “Percept” and “Style”.

2.9 Inpainting-based Applications

Image inpainting can be used in many real-world applications, such as object removal, text editing, old photo restoration, image compression, text-guided image editing, etc.




Fig. 11: Representative samples of text editing.

2.9.1 Object Removal

Almost all image editing tools include the function of object removal, which is directly accomplished with image inpainting. To illustrate the capability of several current inpainting methods on the object removal application, we apply the respective trained models to remove objects from selected real-world images with different scenes, and the corresponding results are shown in Fig. 10. The first row is generated by CNN-based method (Suvorov et al., 2022) and the second one is inpainted by a transformer-based method (Zheng et al., 2022a). These two methods can achieve visually realistic results, successfully removing the objects highlighted with shadow markers.

2.9.2 Text Editing

On social media sites, users often share their pictures and also want to hide their personal information for privacy. For real-time text translation applications in smartphones, the original content needs to be replaced with the translated version. These text editing-related tasks can be solved via inpainting techniques. Fig. 11(a) shows the results of text removal with the method proposed by (Quan et al., 2022), and Fig. 11(b) illustrates two samples of text replacement from (Wu et al., 2019). These results have a pleasing visual quality.

2.9.3 Old Photo Restoration

Photos are helpful to record important moments. Unfortunately, some photos are damaged over time, resulting in various missing regions. Image inpainting can be used to recover these incomplete photos automatically. It is difficult to collect paired training data for this task, therefore, we synthesize old photos using the Pascal VOC dataset Everingham et al. (2015) inspired by Wan et al. (2020). Specifically, we collect some paper and



Fig. 12: Several representative examples of old photo restoration.

scratch texture images to simulate the realistic defects in the old photos. To blend the above texture images with the VOC images, we randomly choose a mode from three candidates (screen, lighten-only, and layer addition) with a random opacity. In addition, some operations, *e.g.*, random flipping, random position, rescaling, cropping, etc. are also used for augmenting the diversity of texture images. To this end, the paired samples of the original VOC images and the corresponding blended results are used for training the inpainting network (Quan et al., 2022). Fig. 12 shows several examples of old photo restoration, where the inpainting method restores the original appearance of old photos.

2.9.4 Image Compression

Image compression is a fundamental image processing technique to reduce the cost of storage or transmission of digital images. This technique mainly consists of two stages: compression and reconstruction. The former reduces the data size to obtain a sparse image representation and the latter reconstructs the original image. Different from waveform-based methods, Carlsson (1988) proposed a sketch-based method to obtain a sparse representation and reconstructed images via an interpolation process. Galić et al. (2008) introduced partial differential equation (PDE)-based inpainting to image compression, where image coding and decoding both are based on edge-enhancing anisotropic diffusion. Recently, some researchers (Baluja et al., 2019; Dai et al., 2020; Schrader et al., 2023) applied deep learning methods to generate the sampling mask and reconstruct the image with an inpainting network. Fig. 13 shows several examples of image compression with inpainting, where the reconstructed images have good quality based on adaptive sparse sampling with inpainting.




Fig. 13: Two representative examples of image compression with inpainting. From left to right: input image, sampling mask, sampled image, and reconstructed image. Images come from (Dai et al., 2020).




Fig. 14: Selected examples of text-based image editing. Each group includes the input image with mask (red transparent) and text prompts and edited results. Images come from (Xie et al., 2023).

2.9.5 Text-guided image editing

Image inpainting is a basic processing tool for image editing. Recent generative models based on probabilistic diffusion have the powerful capability of text-to-image generation, which provides the potential for text-guided image editing with diffusion model-based image inpainting approaches. For example, diffusion-based SmartBrush (Xie et al., 2023) edited images with the guidance of text and shape. Fig. 14 illustrates several samples generated by SmartBrush. The first row adds new objects and the second row replaces original objects with new contents. We can see that the edited results have high visual realism and are consistent with text prompts.

3 Video Inpainting

3.1 Method

Unlike images, videos have an additional temporal dimension which provides extra information about objects or camera movement. This information helps networks to obtain a better understanding of the context of the video. Therefore, the video inpainting task aims to ensure both spatial consistency and temporal coherence. Existing deep learning-based video inpainting methods can be roughly divided into four categories: 3D CNN-based approaches, shift-based approaches, flow-guided approaches, and attention-based approaches. We refer the readers to more conventional methods in (Ilan and Shamir, 2015).

3.1.1 3D CNN-based Approaches

To deal with the temporal dimension, researchers proposed 3D CNN-based approaches, which often combine temporal restoration and image inpainting. Wang et al. (2019a) proposed a two-stage pipeline to jointly infer temporal structure and spatial texture details. The first sub-network processes the low-resolution videos with a 3D CNN, and the second sub-network completes the original-resolution video frames with an extended 2D inpainting network (Iizuka et al., 2017). Inspired by the gated convolution in image inpainting (Yu et al., 2019), Chang et al. (2019a) proposed a 3D gated convolution and a temporal SN-PatchGAN for free-form video inpainting. They also integrated the perceptual loss (Johnson et al., 2016) and style loss (Gatys et al., 2016) into the training objective. Hu et al. (2020) proposed a two-stage video inpainting network, where they obtain a coarse inpainting result with a 3D CNN and then fuse inpainting proposals generated by matching valid pixels and pixels in coarse inpainting results.

3.1.2 Shift-based Approaches

Considering the high computational cost of 3D convolution, Lin et al. (2019) proposed a generic temporal shift module (TSM) to capture temporal relationships with high efficiency. This technique is extended for video inpainting. Chang et al. (2019b) developed a learnable gated TSM, which combines a TSM with learnable shifting kernels and gated convolution (Yu et al., 2019). They also equipped the 2D convolution layers in SN-PatchGAN (Yu et al., 2019) with gated TSM. However, TSM often leads to blurry content due to misaligned features. To solve this, Zou et al. (2021) proposed a spatially-aligned TSM (TSAM), aligning

features to the current frame after shifting features. The alignment process is based on estimated flow with a validity mask. Ouyang et al. (2021) applied an internal learning strategy for video inpainting, which implicitly learns the information shift from valid regions to unknown parts in a single video sample. They also designed the gradient regularization term and the anti-ambiguity loss term for temporal consistency reconstruction and realistic detail generation. Ke et al. (2021) presented an occlusion-aware video object inpainting method. Specifically, they completed the object shape with a transformer-based network, recovered the flow within the completed object region under the guidance of the object contour, and filled missing content with an occlusion-aware TSM after the flow-guided pixel propagation.

3.1.3 Flow-guided Approaches

Optical flow is a common tool to model the temporal information in videos, which is also applied to solve video inpainting. Based on the completed flow, the missing pixels in the current frame can be filled by propagating pixels from neighboring frames. Kim et al. (2019b) modeled the video inpainting task as a multi-to-single frame inpainting problem and proposed a 3D-2D encoder-decoder network VINet. This network includes several flow and mask sub-networks in a progressive manner. They also introduced the flow and warp loss to further enforce temporal consistency. Chang et al. (2019c) proposed a three-stage video inpainting framework consisting of a warping network, an inpainting network, and a refinement network. In the warping network, bilinear interpolation is used to recover background flow without learning. Then the refinement network selected the best candidate from two frames completed by warping and inpainting network to generate the final output. Zhang et al. (2019a) applied internal learning to infer both frames and flow from input random noise and used flow generation loss to enhance temporal coherence. Xu et al. (2019) proposed a flow-guided completion framework consisting of three steps. It first fills the incomplete optical flow with stacked CNN networks, then propagates pixels from known regions to holes with inpainted flow guidance, and finally completes unseen regions with an image inpainting network (Yu et al., 2019). To reduce the over-smoothing in the boundary regions during flow completion, they leveraged hard flow example mining to encourage the network to produce sharp edges. To solve the same problem, Gao et al. (2020) explicitly completed motion edges and used them to guide flow completion. In addition, they introduced

a non-local flow connection to enable content propagation from distant frames.

These previous methods cannot guarantee the consistency of flow, and even small errors in the flow may lead to geometric distortion in the video. Inspired by this, [Lao et al. \(2021\)](#) transformed the background of a 3D scene to a 2D scene template and learned the mapping of the template to the mask in the image. Given that the complex motion of objects between consecutive frames will increase the difficulty to recover flow, [Zhang et al. \(2022b\)](#) introduced an inertia prior in flow completion to align and aggregate flow features. To alleviate the spatial incoherence problem, they proposed an adaptive style fusion network to correct the distribution in the warped regions with the guidance of feature distribution in valid regions. [Kang et al. \(2022\)](#) offset the weaknesses of the error accumulation of a multi-stage pipeline in flow-based methods by introducing an error compensation strategy, which iteratively detects and corrects the inconsistency errors during the flow-guided pixel propagation.

The above hand-crafted flow-based methods restored videos with high computation and memory consumption because these processes cannot be accelerated by GPU. To speed up training and inference, [Li et al. \(2022d\)](#) proposed an end-to-end framework. They propagated features based on completed flow in low resolution and used deformable convolution to decrease the distortion caused by errors in flow. The temporal focal transformer blocks were stacked to aggregate local and non-local features.

3.1.4 Attention-based Approaches

The attention mechanism is often applied to model the contextual information and enlarge the spatial-temporal window. [Oh et al. \(2019\)](#) recurrently calculated the attention scores between the target and reference frames, and progressively filled holes of the target frame from the boundary. [Lee et al. \(2019\)](#) firstly aligned frames by an affine transformation, and then copied pixels based on the similarity between the target frame and aligned reference frames. [Woo et al. \(2020\)](#) proposed a coarse-to-fine framework for video inpainting. The first stage roughly recovers the target holes based on the computed homography between the target and reference frames, and the second stage refines the filled contents with non-local attention. They also introduced an optical flow estimator to enhance temporal consistency. Considering the motion of the foreground objects is diverse, the choice of reference frames becomes more important. While other methods take neighboring frames or frames in a specific distance as reference frames, [Li](#)

[et al. \(2020a\)](#) dynamically updated long-term reference frames after aggregating short-term aligned features.

Instead of a frame-by-frame inpainting strategy, [Zeng et al. \(2020a\)](#) adopted a “multi-to-multi” mechanism to fill in the holes in all input frames. Specifically, they proposed a spatial-temporal transformer network (STTN) to compute attention in both spatial and temporal dimensions. Based on STTN ([Zeng et al., 2020a](#)), [Liu et al. \(2021b\)](#) separated feature maps into overlapping patches, enabling more interactions between neighboring patches. In addition, they modified the common transformer block by inserting soft split and soft composition modules into the feed-forward network. [Chen et al. \(2021\)](#) proposed an interactive video inpainting method to jointly perform object segmentation and video inpainting with user guidance. For network design, they introduce a spatial time attention block to update the target frames’ features with the reference frames’ features. [Zhang et al. \(2022a\)](#) designed a flow-guided transformer to combine the flow and the attention. They first utilized the completed flow to propagate pixels from neighboring frames, and then synthesized the remaining missing regions with a flow-guided spatial transformer and a temporal transformer.

These attention-based methods still suffer from blurry content in high frequency due to mapping videos into a continuous feature space. By learning a specific codebook for each video and using subscripts of code to represent images, [Ren et al. \(2022\)](#) transformed videos to a discrete latent space. Then a discrete latent transformer was applied to infer content in masked regions.

Table 5 summarizes the technical details of existing video inpainting methods.

3.2 Loss Functions

Video inpainting is very close to image inpainting. Therefore, many loss functions for training image inpainting networks are also applied to train video inpainting models, including reconstruction loss, GAN loss, perceptual loss, and style loss. To complete the corrupted flow, two losses are often used:

Flow loss. Similar to the image reconstruction loss, the flow loss measures the difference between inpainted flow and its ground-truth version, which is defined as:

$$\mathcal{L}_{flow} = \|O_{i,j} \odot (F_{i,j} - \hat{F}_{i,j})\|_1, \quad (11)$$

where $\hat{F}_{i,j}$ is the inpainted optical flow from frame i to frame j , $F_{i,j}$ is the ground-truth flow estimated by pre-trained flow estimation networks, *e.g.*, FlowNet2 ([Ilg et al., 2017](#)) and PWC-Net ([Sun et al., 2018a](#)), and $O_{i,j}$ denotes the occlusion map obtained by the forward-backward consistency check.

Table 5: Summary of video inpainting methods. Like image inpainting, we also split existing video inpainting approaches into three types according to the number of stages: 1) one-stage framework (①) usually designs a generator to recover the missing contents for each frame; 2) two-stage framework (②) often consists of two networks for different purposes; and 3) multi-stage framework (Ⓜ) splits video inpainting into multiple steps.

Category	Method	Stage	Loss details						
			L1 loss	GAN loss	Perceptual loss	Style loss	TV loss	Flow loss	Warp loss
3D CNN	Wang et al. (2019a)	②	✓						
	Chang et al. (2019a)	①	✓	✓	✓	✓			
	Hu et al. (2020)	②	✓	✓					
Shift	Chang et al. (2019b)	①	✓	✓	✓	✓			
	Zou et al. (2021)	①	✓	✓	✓	✓			
	Ouyang et al. (2021)	①	✓						
	Ke et al. (2021)	Ⓜ	✓	✓				✓	✓
Flow	Kim et al. (2019b)	①	✓					✓	✓
	Chang et al. (2019c)	Ⓜ	✓	✓	✓				
	Zhang et al. (2019a)	①	✓		✓				✓
	Xu et al. (2019)	Ⓜ	✓					✓	
	Gao et al. (2020)	Ⓜ	✓					✓	
	Lao et al. (2021)	②	✓						✓
	Zhang et al. (2022b)	Ⓜ	✓	✓				✓	✓
	Li et al. (2022d)	①	✓	✓				✓	
	Kang et al. (2022)	Ⓜ	✓	✓				✓	
Attention	Oh et al. (2019)	Ⓜ	✓		✓	✓	✓		
	Lee et al. (2019)	②	✓		✓	✓	✓		
	Woo et al. (2020)	②	✓	✓				✓	✓
	Li et al. (2020a)	①	✓		✓	✓			
	Zeng et al. (2020a)	①	✓	✓					
	Liu et al. (2021b)	①	✓	✓					
	Chen et al. (2021)	②	✓		✓	✓			
	Zhang et al. (2022a)	②	✓	✓				✓	✓
Ren et al. (2022)	②	✓					✓		

Table 6: Quantitative comparisons of representative video inpainting methods on YouTube-VOS and DAVIS dataset. ‡ Higher is better. † Lower is better. *: our results using the method described in STTN (Zeng et al., 2020a), and numerical differences may be due to different optical flow models during evaluation.

Methods	YouTube-VOS				DAVIS			
	PSNR‡	SSIM‡	VFID†	FWE($\times 10^{-2}$)†	PSNR‡	SSIM‡	VFID†	FWE($\times 10^{-2}$)†
VINet (Kim et al., 2019b)	29.20	0.9434	0.072	0.1490 / -	28.96	0.9411	0.199	0.1785 / -
DFVI (Xu et al., 2019)	29.16	0.9429	0.066	0.1509 / -	28.81	0.9404	0.187	0.1880 / 0.1608*
LGTSM (Chang et al., 2019b)	29.74	0.9504	0.070	0.1859 / -	28.57	0.9409	0.170	0.2566 / 0.1640*
CAP (Lee et al., 2019)	31.58	0.9607	0.071	0.1470 / -	30.28	0.9521	0.182	0.1824 / 0.1533*
FGVC (Gao et al., 2020)	29.68	0.9396	0.064	- / 0.0858*	30.24	0.9444	0.143	- / 0.1530*
STTN (Zeng et al., 2020a)	32.34	0.9655	0.053	0.1451 / 0.0884*	30.67	0.9560	0.149	0.1779 / 0.1449*
FuseFormer (Liu et al., 2021b)	33.16	0.9673	0.051	- / 0.0875*	32.54	0.9700	0.138	- / 0.1336*
FGT (Zhang et al., 2022a)	32.11	0.9598	0.054	- / 0.0860*	32.39	0.9633	0.1095	- / 0.1517*
ISVI (Zhang et al., 2022b)	32.80	0.9611	0.048	- / 0.0856*	33.70	0.967	0.1028	- / 0.1509*
E2FGVI (Li et al., 2022d)	33.50	0.9692	0.046	- / 0.0864*	32.71	0.9700	0.096	- / 0.1383*

Warp loss. This loss encourages image-flow consistency:

$$\mathcal{L}_{warp} = \|I_i - I_j(\hat{F}_{i,j})\|_1, \quad (12)$$

where $I_j(\hat{F}_{i,j})$ refers to the warped result of the frame I_j using the generated flow $\hat{F}_{i,j}$ through backward warping.

3.3 Datasets

For video inpainting, three common video datasets, *i.e.*, FaceForensics (Rössler et al., 2018), DAVIS (Perazzi et al., 2016) and YouTube-VOS (Xu et al., 2018a), are used for training and evaluation.

- FaceForensics: A face forgery detection video dataset consisting of 1,004 videos. Among them, 854 videos



Fig. 15: Qualitative comparisons of representative video inpainting methods on YouTube-VOS and DAVIS dataset. The light blue mask highlights the corrupted regions. The first three columns are random masks and the remaining two columns are object masks.

are used for training and the rest are used for evaluation.

- DAVIS dataset: A densely annotated video segmentation dataset contains 150 videos with challenging motion-blur and appearance motions. For the data split, 60 videos are used for training and 90 videos for testing.
- YouTube-VOS dataset: A large-scale video object segmentation dataset containing 4,453 video clips and 94 object categories. The video clips have on average 150 frames and show various scenes. The original data split, *i.e.*, 3,471/474/508, is adopted for experimental comparisons.

3.4 Evaluation Protocol

Video contains many image frames, therefore, the two most widely-used metrics in image inpainting (*i.e.*, PSNR and SSIM) are also used for video quality assessment. In addition, there are two other video-specific metrics (considering the temporal aspect), *i.e.*, flow warping error (FWE) [Lai et al. \(2018\)](#) and video-based Fréchet inception distance (VFID) [Wang et al. \(2018a\)](#). The former evaluates the temporal stability of inpainted videos and the latter measures the perceptual realism in the video setting.

- FWE: The flow warping error between two consecutive video frames is calculated as $\mathcal{E}(\mathbf{I}_t, \mathbf{I}_{t+1}) = \frac{1}{\sum_{n=1}^N \mathbf{M}_t^n} \sum_{n=1}^N \mathbf{M}_t^n \|\mathbf{I}_t^n - \hat{\mathbf{I}}_{t+1}^n\|_2^2$, where \mathbf{M}_t is a binary mask indicating non-occluded areas and $\hat{\mathbf{I}}_{t+1}$ is

the warped frame of \mathbf{I}_{t+1} . The non-occlusion mask can be estimated by using the method Ruder et al. (2016). Then, the warping error of a video is defined as the average error over the entire frames, and the formulation is $\mathcal{E} = \frac{1}{T-1} \sum_{t=1}^{T-1} \mathcal{E}(\mathbf{I}_t, \mathbf{I}_{t+1})$.

- VFID: A variant of FID for video evaluation. Instead of using a pre-trained image recognition network, the spatiotemporal feature map of each video is extracted via a pre-trained video recognition network, *e.g.*, I3D Carreira and Zisserman (2017). Then, the VFID is calculated following the same procedure as the FID.

3.5 Performance Evaluation

In this section, we report the performance evaluation of representative video inpainting methods.

Table 6 shows the numerical results on YouTube-VOS and DAVIS datasets. We use the evaluated masks shared by (Liu et al., 2021b). Early video inpainting methods based on 3D convolution (*e.g.*, VNet (Kim et al., 2019b)) and shift (*e.g.*, LGTSM (Chang et al., 2019b)) have relatively limited inpainting performance. After introducing optical flow and attention mechanisms, the quality of video inpainting is remarkably improved. DFVI (Xu et al., 2019) generates the baseline result with flow guidance, and FGVC (Gao et al., 2020) achieves better performance by completing flow with sharp edges and propagating information from distant frames. ISVI (Zhang et al., 2022b) obtains more exact flow completion under the inertia prior, and thus enhances the inpainting quality. STTN (Zeng et al., 2020a) and FuseFormer (Liu et al., 2021b) both design video inpainting frameworks through stacking multiple transformer blocks with multi-scale attention and dense patch-wise attention, respectively. FGT (Zhang et al., 2022a) and E2FGVI (Li et al., 2022d) combine the flow completion and transformer as a whole, and the end-to-end pipeline as adopted by E2FGVI is slightly better.

Fig. 15 illustrates some inpainted results with different types of scenes and masks. From the first and second columns, we find that the flow-based pixel propagation methods, including FGVC, FGT, and ISVI, have a good ability to recover the texture details and objects with the guidance of neighboring frames. Through contextual correlation modeling, transformer-based video inpainting methods, such as STTN, FuseFormer, and E2FGVI, can complete the structure of objects, *e.g.*, the window of a bus in the third column. Compared to STTN, FuseFormer introduces more dense attention computation (with overlapping), which can help the global structure recovery, *e.g.*, the trunk in the fourth column and the post in the last column. In the



Fig. 16: Several representative examples of blind video decaptioning produced by (Chu et al., 2021).

fourth column, the coverage area is better filled with the realistic grass texture by the ISVI method, which is attributed to the more accurate flow completion compared to FGVC and FGT.

3.6 Applications

3.6.1 Blind Video Decaptioning

Blind video decaptioning aims to automatically remove subtitles and recover the occluded regions in videos without mask information. Kim et al. (2019a) designed an encoder-decoder framework based on 3D convolution. They applied residual learning to directly touch the corrupted regions and leveraged feedback connections to enforce temporal coherence with the warping loss. However, this method often suffers from the problem of incomplete subtitle removal. Chu et al. (2021) proposed a two-stage video decaptioning network including a mask extraction module and a frame attention-based decaptioning module. Several examples produced by (Chu et al., 2021) are shown in Fig. 16. The regions originally covered by subtitles are filled with plausible content.

3.6.2 Dynamic Object Removal

A common practical application of video inpainting technology is to automatically remove undesired objects, which are static or dynamic at the time of recording. In this part, we show two examples of dynamic object removal with the recent video inpainting methods (Liu et al., 2021b; Ren et al., 2022; Kang et al., 2022). As shown in Fig. 17, the regions covered by dynamic objects can be filled with plausible content.



Fig. 17: Three examples of dynamic object removal produced by FuseFormer (Liu et al., 2021b), DLFormer (Ren et al., 2022), and ECFVI (Kang et al., 2022).

4 Future Work

Image and video inpainting essentially is a conditional generative task, therefore, the common generative models, such as VAE and GAN, are often adopted by the existing inpainting methods. Currently, diffusion models have become the most popular generative models with powerful capability of content synthesis. DMs would have the potential to improve the performance of image and video inpainting and may attract a lot of research effort in the future. For this promising direction, several challenging problems need to be solved.

How to use large pre-trained diffusion models (e.g., denoising diffusion) for image inpainting?

DMs synthesize an image by a sequential application of denoising steps, which are conducted in pixel or latent space. For the inpainting task, the core idea is to fill in the missing regions while preserving the originally valid content. Some researchers have made preliminary attempts, such as Palette (Saharia et al., 2022a), Blended Diffusion (Avrahami et al., 2022), and ControlNet (Zhang and Agrawala, 2023), etc. One research challenge is how to inject conditioning information into the denoising processes of large pre-trained diffusion models. Following the pipeline of diffusion models, they need many iterations to generate the final image and thus require a longer inference time compared to existing VAE- and GAN-based approaches. Another research challenge is to implement fast inpainting methods based on diffusion. Also, while video-based generative diffusion models are still in their infancy, it is expected that large pre-trained video generation models will become available in the near future. Leveraging these models for video inpainting will be an interesting task once these models become available.

How to use large pre-trained models for joint text and image embedding (e.g., the latest CLIP style architecture) for image inpainting?

Mainstream inpainting methods are uncontrollable, where the inpainted content is unknown in advance and sometimes this is undesired for users. Reference-based inpainting cannot fully satisfy this requirement. On the other hand, recent studies (Rombach et al., 2022; Hertz et al., 2022; Parmar et al., 2023) have shown that large pre-trained diffusion models with massive text-image pairs can synthesize high-quality images with rich low-level attributes and details. In addition, Zhao et al. (2023) implied that such pre-trained DMs also contain high-level visual concepts. As a result, text-guided inpainting based on the large pre-trained text-to-image diffusion models would be able to fill the content under the control of users. The first problem is to design the appropriate prompt exactly indicating the user’s intention. It is also challenging to merge the image embedding from the user prompt with the corresponding embedding of the input corrupted image. In addition, text-based video inpainting will be a great avenue for future work.

How to scale up training to datasets of 5B images (e.g. LAION)?

Deep learning models are hungry for training datasets. Currently, advanced diffusion models are pre-trained on large-scale datasets containing millions or even billions of text-image data pairs. However, these models are mainly dominated by several industrial research labs, where the datasets and training processes are not transparent to the research community. Very recently, the largest text-image dataset LAION-5B (Schuhmann et al., 2022) containing 5.8 billion samples is publicly available. In future work, it is worth designing efficient methods for image and video inpainting that are trained on such very large datasets directly.

How to utilize image data and pre-trained image inpainting models to improve the models of video inpainting?

In addition to considering the spatial aspect as in image inpainting, video inpainting also needs to consider the temporal aspect. Therefore, it is important and beneficial to transfer the inpainting ability from image to video. A simple and direct solution is to take the result of image inpainting on each frame as the initialization and then revise the spatial and temporal consistency via carefully designed deep models. Another possible research line is to take the well-trained image inpainting models as the backbone and aggregate the multiple frames in the feature space with appropriate modules, such as deformable convolution or attention. It’s still worth exploring combining the pre-trained image inpainting model with deep video prior.

How to create a large video dataset of 5B videos and leverage it for video inpainting?

Like image inpainting, taking advantage of large pre-trained text-video diffusion models may be a new research direction for video inpainting. However, current text-video DMs are trained on datasets with 10 million captioned videos, which inevitably limits the generation and generalization ability of DMs. One potential direction of future research is to collect large-scale text-video datasets (*e.g.*, 5B pairs) and design the pre-training methods scaling up to this amount. As we all know, video inpainting is more difficult compared to its image counterpart. Therefore, it is valuable to spend time on all aspects of large video datasets: building large publicly available video datasets, generating large diffusion methods for video synthesis and using these pre-trained methods for video inpainting, and separately designing and training large-scale video architectures directly.

5 Concluding Remarks

The prevalence of visual data, including images and video, promotes the development of related processing technologies, *e.g.*, image and video inpainting. Due to their practical applications in many fields, these techniques have attracted great attention from both the industrial and research communities over the past decade. We presented a review of deep learning-based methods for image and video inpainting. Specifically, we outline different aspects of the research, including a taxonomy of existing methods, training objectives, benchmark datasets, evaluation protocols, performance evaluation, and real-world applications. Future research directions are also discussed.

Although current deep learning-based inpainting approaches have achieved remarkable performance improvement, there are still several limitations: (1) Uncertainty of artifacts. The results generated by inpainting methods often exhibit visual artifacts, which are difficult to predict and prevent. There is almost no research work to systematically and comprehensively study these artifacts. (2) Specificity. Current inpainting models are usually trained on specific datasets, for example, face images or natural scene images. In other words, models trained on face images have bad predictions on natural scene images, and vice versa. Not enough models are trained on large scale datasets such as LIAON. (3) Large-scale inpainting. Current advanced inpainting methods still have limited performance on large-scale missing regions. Many methods are based on attention mechanisms, which are more fragile in large-scale scenarios. (4) High training costs. Current deep learning-based inpainting methods often need one or more weeks on multiple GPUs, which places very high demands on

resource consumption. (5) Long inference time. Diffusion model-based methods can achieve better inpainting performance, however, they need a very long running time, which limits the application scope of inpainting techniques.

Deep image/video inpainting techniques have a wide range of real-world applications, however, they also raise potential ethical issues that need to be carefully considered and addressed: (1) Security risks. Inpainting-based visual data editing, *e.g.*, object removal, may maliciously be exploited, such as tampering with visual data or altering evidence. (2) Ownership and copyright. When there is no appropriate authorization, deep inpainting techniques used to manipulate and enhance images/videos could raise questions about ownership and copyright. The inpainting result may strongly resemble or be strongly inspired by copyrighted material. (3) Historical accuracy. Inpainting methods can be used for the restoration of old photos/films or artworks. This process could raise risks of inadvertently changing the initial creative intention or historical accuracy of the content, which requires careful verification by domain experts. (4) Bias. If not properly trained, an inpainting model may introduce bias or unfairness, especially when the training data is biased or unrepresentative. This has the potential to perpetuate social prejudices or inaccurately portray certain groups.

References

- Arjovsky M, Chintala S, Bottou L (2017) Wasserstein Generative Adversarial Networks. *In: Int. Conf. Mach. Learn.*, vol 70, pp 214–223
- Austin J, Johnson DD, Ho J, Tarlow D, van den Berg R (2021) Structured Denoising Diffusion Models in Discrete State-Spaces. *In: Adv. Neural Inform. Process. Syst.*, vol 34, pp 17981–17993
- Avrahami O, Lischinski D, Fried O (2022) Blended Diffusion for Text-Driven Editing of Natural Images. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 18208–18218
- Ballester C, Bertalmio M, Caselles V, Sapiro G, Verdera J (2001) Filling-in by joint interpolation of vector fields and gray levels. *IEEE Trans Image Process* 10(8):1200–1211
- Baluja S, Marwood D, Johnston N, Covell M (2019) Learning to Render Better Image Previews. *In: IEEE Int. Conf. Image Process.*, pp 1700–1704
- Barnes C, Shechtman E, Finkelstein A, Goldman DB (2009) PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Trans Graph* 28(3):24

- Bertalmio M, Sapiro G, Caselles V, Ballester C (2000) Image inpainting. *In: Proc. ACM SIGGRAPH*, pp 417–424
- Bian X, Wang C, Quan W, Ye J, Zhang X, Yan DM (2022) Scene text removal via cascaded text stroke detection and erasing. *Computational Visual Media* 8:273–287
- Blau Y, Michaeli T (2018) The Perception-Distortion Tradeoff. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 6228–6237
- Canny J (1986) A computational approach to edge detection. *IEEE Trans Pattern Anal Mach Intell* (6):679–698
- Cao C, Fu Y (2021) Learning a Sketch Tensor Space for Image Inpainting of Man-Made Scenes. *In: Int. Conf. Comput. Vis.*, pp 14509–14518
- Cao C, Dong Q, Fu Y (2022) Learning Prior Feature and Attention Enhanced Image Inpainting. *In: Eur. Conf. Comput. Vis.*
- Carlsson S (1988) Sketch based coding of grey level images. *Sign Process* 15(1):57–83
- Carreira J, Zisserman A (2017) Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 4724–4733
- Chang YL, Liu ZY, Lee KY, Hsu W (2019a) Free-form video inpainting with 3d gated convolution and temporal patchgan. *In: Int. Conf. Comput. Vis.*, pp 9066–9075
- Chang YL, Liu ZY, Lee KY, Hsu W (2019b) Learnable gated temporal shift module for deep video inpainting. *In: Brit. Mach. Vis. Conf.*
- Chang YL, Yu Liu Z, Hsu W (2019c) Vornet: Spatio-temporally consistent video inpainting for object removal. *In: IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pp 1785–1794
- Chen C, Cai J, Hu Y, Tang X, Wang X, Yuan C, Bai X, Bai S (2021) Deep Interactive Video Inpainting: An Invisibility Cloak for Harry Potter. *In: ACM Int. Conf. Multimedia*, p 862–870
- Chen L, Zhang H, Xiao J, Nie L, Shao J, Liu W, Chua TS (2017) SCA-CNN: Spatial and Channel-Wise Attention in Convolutional Networks for Image Captioning. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 6298–6306
- Chen P (2018) Video retouch: Object removal. <http://www.12371.cn/2021/02/08/ARTI1612745858192472.shtml>
- Chen T, Lucic M, Houlsby N, Gelly S (2019) On Self Modulation for Generative Adversarial Networks. *In: Int. Conf. Learn. Represent.*
- Chi L, Jiang B, Mu Y (2020) Fast Fourier Convolution. *In: Adv. Neural Inform. Process. Syst.*, vol 33, pp 4479–4488
- Chu P, Quan W, Wang T, Wang P, Ren P, Yan DM (2021) Deep video decaptioning. *In: Brit. Mach. Vis. Conf.*
- Chung H, Sim B, Ye JC (2022) Come-Closer-Diffuse-Faster: Accelerating Conditional Diffusion Models for Inverse Problems through Stochastic Contraction. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 12403–12412
- Cimpoi M, Maji S, Kokkinos I, Mohamed S, Vedaldi A (2014) Describing Textures in the Wild. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 3606–3613
- Criminisi A, Perez P, Toyama K (2004) Region filling and object removal by exemplar-based image inpainting. *IEEE Trans Image Process* 13(9):1200–1212
- Croitoru FA, Hondru V, Ionescu RT, Shah M (2023) Diffusion models in vision: A survey. *IEEE Trans Pattern Anal Mach Intell* 45(9):10850–10869
- Dai Q, Chopp H, Pouyet E, Cossairt O, Walton M, Katsaggelos AK (2020) Adaptive image sampling using deep learning and its application on x-ray fluorescence image reconstruction. *IEEE Trans Multimedia* 22(10):2564–2578
- Darabi S, Shechtman E, Barnes C, Goldman DB, Sen P (2012) Image Melding: combining inconsistent images using patch-based synthesis. *ACM Trans Graph (Proc SIGGRAPH)* 31(4):1–10
- Daubechies I (1990) The wavelet transform, time-frequency localization and signal analysis. *IEEE Trans Inf Theory* 36(5):961–1005
- Deng J, Dong W, Socher R, Li LJ, Li K, Fei-Fei L (2009) Imagenet: A large-scale hierarchical image database. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 248–255
- Deng Y, Tang F, Dong W, Sun W, Huang F, Xu C (2020) Arbitrary Style Transfer via Multi-Adaptation Network. *In: ACM Int. Conf. Multimedia*, p 2719–2727
- Deng Y, Hui S, Zhou S, Meng D, Wang J (2021) Learning Contextual Transformer Network for Image Inpainting. *In: ACM Int. Conf. Multimedia*, p 2529–2538
- Deng Y, Hui S, Meng R, Zhou S, Wang J (2022) Hour-glass Attention Network for Image Inpainting. *In: Eur. Conf. Comput. Vis.*, pp 483–501
- Dinh L, Krueger D, Bengio Y (2014) Nice: Non-linear independent components estimation. *Int Conf Learn Represent Worksh*
- Doersch C, Singh S, Gupta A, Sivic J, Efros AA (2012) What makes paris look like paris? *ACM Trans Graph* 31(4):101:1–101:9
- Dolhansky B, Ferrer CC (2018) Eye In-painting with Exemplar Generative Adversarial Networks. *In:*

- IEEE Conf. Comput. Vis. Pattern Recog.*, pp 7902–7911
- Dong Q, Cao C, Fu Y (2022) Incremental Transformer Structure Enhanced Image Inpainting With Masking Positional Encoding. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 11358–11368
- Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, Dehghani M, Minderer M, Heigold G, Gelly S, Uszkoreit J, Houlsby N (2021) An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *In: Int. Conf. Learn. Represent.*
- Dosselmann R, Yang XD (2011) A comprehensive assessment of the structural similarity index. *Sign Image and Video Process* 5:81–91
- Efros A, Leung T (1999) Texture synthesis by non-parametric sampling. *In: Int. Conf. Comput. Vis.*, vol 2, pp 1033–1038
- Elharrouss O, Almaadeed N, Al-Maadeed S, Akbari Y (2020) Image Inpainting: A Review. *Neural Process Letters* 51(2):2007–2028
- Esser P, Rombach R, Blattmann A, Ommer B (2021) ImageBART: Bidirectional Context with Multinomial Diffusion for Autoregressive Image Synthesis. *In: Adv. Neural Inform. Process. Syst.*, vol 34, pp 3518–3532
- Everingham M, Eslami SMA, Gool LV, Williams CKI, Winn J, Zisserman A (2015) The pascal visual object classes challenge: A retrospective. *Int J Comput Vis* 111:98–136
- Felzenszwalb PF, Huttenlocher DP (2004) Efficient graph-based image segmentation. *Int J Comput Vis* (59):167–181
- Feng X, Pei W, Li F, Chen F, Zhang D, Lu G (2022) Generative memory-guided semantic reasoning model for image inpainting. *IEEE Trans Circuit Syst Video Technol* 32(11):7432–7447
- Fu J, Liu J, Tian H, Li Y, Bao Y, Fang Z, Lu H (2019) Dual Attention Network for Scene Segmentation. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 3141–3149
- Galić I, Weickert J, Welk M, Bruhn A, Belyaev A, Seidel HP (2008) Image compression with anisotropic diffusion. *J Math Imaging Vis* 31:255–269
- Gao C, Saraf A, Huang JB, Kopf J (2020) Flow-edge guided video completion. *In: Eur. Conf. Comput. Vis.*, pp 713–729
- Gatys LA, Ecker AS, Bethge M (2016) Image Style Transfer Using Convolutional Neural Networks. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 2414–2423
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. *In: Adv. Neural Inform. Process. Syst.*, pp 2672–2680
- Granados M, Kim KI, Tompkin J, Kautz J, Theobalt C (2012) Background Inpainting for Videos with Dynamic Objects and a Free-Moving Camera. *In: Eur. Conf. Comput. Vis.*, pp 682–695
- Gu J, Shen Y, Zhou B (2020) Image Processing Using Multi-Code GAN Prior. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 3009–3018
- Guillemot C, Meur OL (2014) Image inpainting: Overview and recent advances. *IEEE Sign Process Magazine* 31(1):127–144
- Guo Q, Gao S, Zhang X, Yin Y, Zhang C (2018) Patch-based image inpainting via two-stage low rank approximation. *IEEE Trans Vis Comput Graph* 24(6):2023–2036
- Guo X, Yang H, Huang D (2021) Image Inpainting via Conditional Texture and Structure Dual Generation. *In: Int. Conf. Comput. Vis.*, pp 14134–14143
- Guo Z, Chen Z, Yu T, Chen J, Liu S (2019) Progressive Image Inpainting with Full-Resolution Residual Network. *In: ACM Int. Conf. Multimedia*, p 2496–2504
- Han C, Wang J (2021) Face image inpainting with evolutionary generators. *IEEE Sign Process Letters* 28:190–193
- Han X, Wu Z, Huang W, Scott MR, Davis L (2019) FiNet: Compatible and Diverse Fashion Image Inpainting. *In: Int. Conf. Comput. Vis.*, pp 4480–4490
- He K, Sun J (2012) Statistics of Patch Offsets for Image Completion. *In: Eur. Conf. Comput. Vis.*, pp 16–29
- He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 770–778
- He K, Chen X, Xie S, Li Y, Dollár P, Girshick R (2022) Masked Autoencoders Are Scalable Vision Learners. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 16000–16009
- Herling J, Broll W (2014) High-quality real-time video inpainting with pixmix. *IEEE Trans Vis Comput Graph* 20(6):866–879
- Hertz A, Mokady R, Tenenbaum J, Aberman K, Pritch Y, Cohen-Or D (2022) Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:220801626*
- Heusel M, Ramsauer H, Unterthiner T, Nessler B, Hochreiter S (2017) Gans trained by a two time-scale update rule converge to a local nash equilibrium. *In: Adv. Neural Inform. Process. Syst.*, pp 6626–6637
- Ho J, Jain A, Abbeel P (2020) Denoising Diffusion Probabilistic Models. *In: Adv. Neural Inform. Process. Syst.*, vol 33, pp 6840–6851
- Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Comput* 9(8):1735–1780

- Hong X, Xiong P, Ji R, Fan H (2019) Deep Fusion Network for Image Completion. *In: ACM Int. Conf. Multimedia*, pp 2033–2042
- Hoogeboom E, Nielsen D, Jaini P, Forré P, Welling M (2021) Argmax Flows and Multinomial Diffusion: Learning Categorical Distributions. *In: Adv. Neural Inform. Process. Syst.*, vol 34, pp 12454–12465
- Houle ME (2017a) Local Intrinsic Dimensionality I: An Extreme-Value-Theoretic Foundation for Similarity Applications. *In: Int. Conf. Similarity Search App.*, pp 64–79
- Houle ME (2017b) Local Intrinsic Dimensionality II: Multivariate Analysis and Distributional Support. *In: Int. Conf. Similarity Search App.*
- Hu J, Shen L, Sun G (2018) Squeeze-and-excitation networks. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 7132–7141
- Hu YT, Wang H, Ballas N, Grauman K, Schwing AG (2020) Proposal-based video completion. *In: Eur. Conf. Comput. Vis.*, pp 38–54
- Huang JB, Kang SB, Ahuja N, Kopf J (2014) Image completion using planar structure guidance. *ACM Trans Graph (Proc SIGGRAPH)* 33(4):1–10
- Huang JB, Kang SB, Ahuja N, Kopf J (2016) Temporally coherent completion of dynamic video. *ACM Trans Graph* 35(6):1–11
- Huang X, Belongie S (2017) Arbitrary Style Transfer in Real-Time with Adaptive Instance Normalization. *In: Int. Conf. Comput. Vis.*, pp 1510–1519
- Hui Z, Li J, Wang X, Gao X (2020) Image fine-grained inpainting. *arXiv preprint arXiv:200202609*
- Iizuka S, Simo-Serra E, Ishikawa H (2017) Globally and locally consistent image completion. *ACM Trans Graph (Proc SIGGRAPH)* 36(4):1–14
- Ilan S, Shamir A (2015) A survey on data-driven video completion. *Comput Graph Forum* 34(6):60–85
- Ilg E, Mayer N, Saikia T, Keuper M, Dosovitskiy A, Brox T (2017) FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 1647–1655
- Isola P, Zhu JY, Zhou T, Efros AA (2017) Image-to-image translation with conditional adversarial networks. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 1125–1134
- Jam J, Kendrick C, Walker K, Drouard V, Hsu JGS, Yap MH (2021) A comprehensive review of past and present image inpainting methods. *Comput Vis Image Understand* 203:103147
- Jiang L, Dai B, Wu W, Loy CC (2021) Focal Frequency Loss for Image Reconstruction and Synthesis. *In: Int. Conf. Comput. Vis.*, pp 13899–13909
- Johnson J, Alahi A, Fei-Fei L (2016) Perceptual losses for real-time style transfer and super-resolution. *In: Eur. Conf. Comput. Vis.*, pp 694–711
- Kang J, Oh SW, Kim SJ (2022) Error Compensation Framework for Flow-Guided Video Inpainting. *In: Eur. Conf. Comput. Vis.*, pp 375–390
- Karras T, Aila T, Laine S, Lehtinen J (2018) Progressive Growing of GANs for Improved Quality, Stability, and Variation. *In: Int. Conf. Learn. Represent.*
- Karras T, Laine S, Aila T (2019) A Style-Based Generator Architecture for Generative Adversarial Networks. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 4396–4405
- Karras T, Laine S, Aittala M, Hellsten J, Lehtinen J, Aila T (2020) Analyzing and Improving the Image Quality of StyleGAN. *In: IEEE Conf. Comput. Vis. Pattern Recog.*
- Ke L, Tai YW, Tang CK (2021) Occlusion-aware video object inpainting. *In: Int. Conf. Comput. Vis.*, pp 14468–14478
- Kim D, Woo S, Lee JY, Kweon IS (2019a) Deep blind video decaptioning by temporal aggregation and recurrence. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 4263–4272
- Kim D, Woo S, Lee JY, Kweon IS (2019b) Deep video inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 5792–5801
- Kim SY, Aberman K, Kanazawa N, Garg R, Wadhwa N, Chang H, Karnad N, Kim M, Liba O (2022) Zoom-to-Inpaint: Image Inpainting With High-Frequency Details. *In: IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pp 477–487
- Kingma DP, Dhariwal P (2018) Glow: Generative Flow with Invertible 1x1 Convolutions. *In: Adv. Neural Inform. Process. Syst.*, vol 31
- Kingma DP, Welling M (2014) Auto-Encoding Variational Bayes. *In: Int. Conf. Learn. Represent.*
- Lai WS, Huang JB, Wang O, Shechtman E, Yumer E, Yang MH (2018) Learning Blind Video Temporal Consistency. *In: Eur. Conf. Comput. Vis.*, pp 179–195
- Lao D, Zhu P, Wonka P, Sundaramoorthi G (2021) Flow-Guided Video Inpainting with Scene Templates. *In: Int. Conf. Comput. Vis.*, pp 14599–14608
- Ledig C, Theis L, Huszár F, Caballero J, Cunningham A, Acosta A, Aitken A, Tejani A, Totz J, Wang Z, Shi W (2017) Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 105–114
- Lee S, Oh SW, Won D, Kim SJ (2019) Copy-and-paste networks for deep video inpainting. *In: Int. Conf. Comput. Vis.*, pp 4413–4421
- Lempitsky V, Vedaldi A, Ulyanov D (2018) Deep Image Prior. *In: IEEE Conf. Comput. Vis. Pattern Recog.*

- pp 9446–9454
- Li A, Qi J, Zhang R, Ma X, Ramamohanarao K (2019a) Generative Image Inpainting with Submanifold Alignment. *In: Int. Joint Conf. Artificial Intell.*, pp 811–817
- Li A, Zhao S, Ma X, Gong M, Qi J, Zhang R, Tao D, Kotagiri R (2020a) Short-term and long-term context aggregation network for video inpainting. *In: Eur. Conf. Comput. Vis.*, pp 728–743
- Li A, Zhao L, Zuo Z, Wang Z, Xing W, Lu D (2023) MigT: Multi-modal image inpainting guided with text. *Neurocomputing* 520:376–385
- Li B, Zheng B, Li H, Li Y (2021) Detail-enhanced image inpainting based on discrete wavelet transforms. *Sign Process* 189:108278
- Li CT, Siu WC, Liu ZS, Wang LW, Lun DPK (2020b) DeepGIN: Deep Generative Inpainting Network for Extreme Image Inpainting. *In: Eur. Conf. Comput. Vis. Worksh.*, pp 5–22
- Li F, Li A, Qin J, Bai H, Lin W, Cong R, Zhao Y (2022a) Srinpaintor: When super-resolution meets transformer for image inpainting. *IEEE Trans Computational Imaging* 8:743–758
- Li H, Li G, Lin L, Yu H, Yu Y (2019b) Context-aware semantic inpainting. *IEEE Trans Cybern* 49(12):4398–4411
- Li J, He F, Zhang L, Du B, Tao D (2019c) Progressive Reconstruction of Visual Structure for Image Inpainting. *In: Int. Conf. Comput. Vis.*, pp 5961–5970
- Li J, Wang N, Zhang L, Du B, Tao D (2020c) Recurrent Feature Reasoning for Image Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 7757–7765
- Li W, Lin Z, Zhou K, Qi L, Wang Y, Jia J (2022b) MAT: Mask-Aware Transformer for Large Hole Image Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*
- Li W, Yu X, Zhou K, Song Y, Lin Z, Jia J (2022c) Sdm: Spatial diffusion model for large hole image inpainting. *arXiv preprint arXiv:221202963*
- Li Y, Liu S, Yang J, Yang MH (2017) Generative Face Completion. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 5892–5900
- Li Y, Jiang B, Lu Y, Shen L (2019d) Fine-grained Adversarial Image Inpainting with Super Resolution. *In: Int. Joint Conf. Neural Networks*, pp 1–8
- Li Z, Lu CZ, Qin J, Guo CL, Cheng MM (2022d) Towards an end-to-end framework for flow-guided video inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 17562–17571
- Liao H, Funke-Lea G, Zheng Y, Luo J, Zhou SK (2018a) Face Completion with Semantic Knowledge and Collaborative Adversarial Learning. *In: Asian Conf. Comput. Vis.*, vol 11361, pp 382–397
- Liao L, Hu R, Xiao J, Wang Z (2018b) Edge-Aware Context Encoder for Image Inpainting. *In: Int. Conf. Acou. Speech Sign. Process.*, pp 3156–3160
- Liao L, Xiao J, Wang Z, Lin CW, Satoh S (2020) Guidance and Evaluation: Semantic-Aware Image Inpainting for Mixed Scenes. *In: Eur. Conf. Comput. Vis.*
- Liao L, Xiao J, Wang Z, Lin CW, Satoh S (2021a) Image Inpainting Guided by Coherence Priors of Semantics and Textures. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 6539–6548
- Liao L, Xiao J, Wang Z, Lin CW, Satoh S (2021b) Uncertainty-aware semantic guidance and estimation for image inpainting. *IEEE J Selected Topics Sign Process* 15(2):310–323
- Lim JH, Ye JC (2017) Geometric gan. *arXiv preprint arXiv:170502894*
- Lin J, Gan C, Han S (2019) Tsm: Temporal shift module for efficient video understanding. *In: Int. Conf. Comput. Vis.*, pp 7083–7093
- Lin Q, Yan B, Li J, Tan W (2020) Mmfl: Multimodal fusion learning for text-guided image inpainting. *In: ACM Int. Conf. Multimedia*, pp 1094–1102
- Liu G, Reda FA, Shih KJ, Wang TC, Tao A, Catanzaro B (2018) Image inpainting for irregular holes using partial convolutions. *In: Eur. Conf. Comput. Vis.*, pp 85–100
- Liu H, Jiang B, Xiao Y, Yang C (2019) Coherent semantic attention for image inpainting. *In: Int. Conf. Comput. Vis.*, pp 4170–4179
- Liu H, Jiang B, Song Y, Huang W, Yang C (2020) Rethinking Image Inpainting via a Mutual Encoder-Decoder with Feature Equalizations. *In: Eur. Conf. Comput. Vis.*
- Liu H, Wan Z, Huang W, Song Y, Han X, Liao J (2021a) PD-GAN: Probabilistic Diverse GAN for Image Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 9367–9376
- Liu R, Deng H, Huang Y, Shi X, Lu L, Sun W, Wang X, Dai J, Li H (2021b) FuseFormer: Fusing Fine-Grained Information in Transformers for Video Inpainting. *In: Int. Conf. Comput. Vis.*, pp 14040–14049
- Liu T, Liao L, Wang Z, Satoh S (2022) Reference-Guided Texture and Structure Inference for Image Inpainting. *In: IEEE Int. Conf. Image Process.*, pp 1996–2000
- Liu W, Cao C, Liu J, Ren C, Wei Y, Guo H (2021c) Fine-grained image inpainting with scale-enhanced generative adversarial network. *Pattern Recog Letters* 143:81–87
- Liu Z, Luo P, Wang X, Tang X (2015) Deep learning face attributes in the wild. *In: Int. Conf. Comput. Vis.*, pp 3730–3738

- Liu Z, Lin Y, Cao Y, Hu H, Wei Y, Zhang Z, Lin S, Guo B (2021d) Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. *In: Int. Conf. Comput. Vis.*, pp 9992–10002
- Lu Z, Jiang J, Huang J, Wu G, Liu X (2022) GLaMa: Joint Spatial and Frequency Loss for General Image Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pp 1301–1310
- Lugmayr A, Danelljan M, Van Gool L, Timofte R (2020) SRFlow: Learning the Super-Resolution Space with Normalizing Flow. *In: Eur. Conf. Comput. Vis.*, p 715–732
- Lugmayr A, Danelljan M, Romero A, Yu F, Timofte R, Van Gool L (2022) RePaint: Inpainting using Denoising Diffusion Probabilistic Models. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 11451–11461
- Ma Y, Liu X, Bai S, Wang L, He D, Liu A (2019) Coarse-to-Fine Image Inpainting via Region-wise Convolutions and Non-Local Correlation. *In: Int. Joint Conf. Artificial Intell.*, pp 3123–3129
- Mallat SG (1989) A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans Pattern Anal Mach Intell* 11(7):674–693
- Mao X, Li Q, Xie H, Lau RY, Wang Z, Paul Smolley S (2017) Least Squares Generative Adversarial Networks. *In: Int. Conf. Comput. Vis.*
- Masnou S, Morel JM (1998) Level lines based disocclusion. *In: IEEE Int. Conf. Image Process.*, vol 3, pp 259–263
- Navasardyan S, Ohanyan M (2020) Image Inpainting with Onion Convolutions. *In: Asian Conf. Comput. Vis.*
- Nazeri K, Ng E, Joseph T, Qureshi F, Ebrahimi M (2019) EdgeConnect: Structure Guided Image Inpainting using Edge Prediction. *In: Int. Conf. Comput. Vis. Worksh.*
- Newson A, Almansa A, Fradet M, Gousseau Y, Pérez P (2014) Video inpainting of complex scenes. *SIAM J Imaging Sciences* 7(4):1993–2019
- Ni M, Li X, Zuo W (2023) NÜWA-LIP: Language-guided Image Inpainting with Defect-free VQGAN. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 14183–14192
- Oh SW, Lee S, Lee JY, Kim SJ (2019) Onion-peel networks for deep video completion. *In: Int. Conf. Comput. Vis.*, pp 4403–4412
- Ojala T, Pietikäinen M, Harwood D (1996) A comparative study of texture measures with classification based on featured distributions. *Pattern Recog* 29(1):51–59
- Ojala T, Pietikäinen M, Maenpää T (2002) Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans Pattern Anal Mach Intell* 24(7):971–987
- Ouyang H, Wang T, Chen Q (2021) Internal video inpainting by implicit long-range propagation. *In: Int. Conf. Comput. Vis.*, pp 14579–14588
- Park T, Liu MY, Wang TC, Zhu JY (2019) Semantic Image Synthesis With Spatially-Adaptive Normalization. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 2332–2341
- Parmar G, Singh KK, Zhang R, Li Y, Lu J, Zhu JY (2023) Zero-shot image-to-image translation. *arXiv preprint arxiv:2302.03027*
- Pathak D, Krahenbuhl P, Donahue J, Darrell T, Efros AA (2016) Context encoders: Feature learning by inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 2536–2544
- Peng J, Liu D, Xu S, Li H (2021) Generating Diverse Structure for Image Inpainting With Hierarchical VQ-VAE. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 10770–10779
- Perazzi F, Pont-Tuset J, McWilliams B, Van Gool L, Gross M, Sorkine-Hornung A (2016) A Benchmark Dataset and Evaluation Methodology for Video Object Segmentation. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 724–732
- Phutke SS, Murala S (2021) Diverse receptive field based adversarial concurrent encoder network for image inpainting. *IEEE Sign Process Letters* 28:1873–1877
- Qin J, Bai H, Zhao Y (2021) Multi-scale attention network for image inpainting. *Comput Vis Image Understand* 204:103155
- Qiu J, Gao Y, Shen M (2021) Semantic-sca: Semantic structure image inpainting with the spatial-channel attention. *IEEE Access* 9:12997–13008
- Quan W, Zhang R, Zhang Y, Li Z, Wang J, Yan DM (2022) Image inpainting with local and global refinement. *IEEE Trans Image Process* 31:2405–2420
- Radford A, Metz L, Chintala S (2016) Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *In: Int. Conf. Learn. Represent.*
- Ren J, Zheng Q, Zhao Y, Xu X, Li C (2022) DLFormer: Discrete Latent Transformer for Video Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 3511–3520
- Ren JS, Xu L, Yan Q, Sun W (2015) Shepard Convolutional Neural Networks. *In: Adv. Neural Inform. Process. Syst.*, vol 28, p 901–909
- Ren Y, Yu X, Zhang R, Li TH, Liu S, Li G (2019) StructureFlow: Image Inpainting via Structure-aware Appearance Flow. *In: Int. Conf. Comput. Vis.*, pp 181–190

- Rezende DJ, Mohamed S (2015) Variational Inference with Normalizing Flows. *In: Int. Conf. Mach. Learn.*, p 1530–1538
- Richardson E, Alaluf Y, Patashnik O, Nitzan Y, Azar Y, Shapiro S, Cohen-Or D (2021) Encoding in Style: a StyleGAN Encoder for Image-to-Image Translation. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 2287–2296
- Rombach R, Blattmann A, Lorenz D, Esser P, Ommer B (2022) High-Resolution Image Synthesis with Latent Diffusion Models. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 10674–10685
- Rössler A, Cozzolino D, Verdoliva L, Riess C, Thies J, Nießner M (2018) Faceforensics: A large-scale video dataset for forgery detection in human faces. *arXiv preprint arXiv:180309179*
- Roy H, Chaudhury S, Yamasaki T, Hashimoto T (2021) Image inpainting using frequency-domain priors. *J Electronic Imaging* 30(2):023016
- Ruder M, Dosovitskiy A, Brox T (2016) Artistic Style Transfer for Videos. *In: German Conf. Pattern Recog.*, pp 26–36
- Rudin LI, Osher S, Fatemi E (1992) Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* 60(1):259–268
- Sagong Mc, Shin Yg, Kim Sw, Park S, Ko Sj (2019) PEPSI : Fast Image Inpainting With Parallel Decoding Network. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 11352–11360
- Saharia C, Chan W, Chang H, Lee C, Ho J, Salimans T, Fleet D, Norouzi M (2022a) Palette: Image-to-Image Diffusion Models. *In: ACM SIGGRAPH Conf.*
- Saharia C, Chan W, Saxena S, Li L, Whang J, Denton E, Ghasemipour SKS, Gontijo-Lopes R, Ayan BK, Salimans T, Ho J, Fleet DJ, Norouzi M (2022b) Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *In: Adv. Neural Inform. Process. Syst.*
- Schrader K, Peter P, Kämper N, Weickert J (2023) Efficient Neural Generation of 4K Masks for Homogeneous Diffusion Inpainting. *In: Int. Conf. Scale Space Variational Methods Comput. Vis.*, pp 16–28
- Schuhmann C, Beaumont R, Vencu R, Gordon CW, Wightman R, Cherti M, Coombes T, Katta A, Mullis C, Wortsman M, Schramowski P, Kundurthy SR, Crowson K, Schmidt L, Kaczmarczyk R, Jitsev J (2022) LAION-5B: An open large-scale dataset for training next generation image-text models. *In: Adv. Neural Inform. Process. Syst.*
- Shao H, Wang Y, Fu Y, Yin Z (2020) Generative image inpainting via edge structure and color aware fusion. *Sign Process: Image Communication* 87:115929
- Shen L, Hong R, Zhang H, Zhang H, Wang M (2019) Single-Shot Semantic Image Inpainting with Densely Connected Generative Networks. *In: ACM Int. Conf. Multimedia*, p 1861–1869
- Shin YG, Sagong MC, Yeo YJ, Kim SW, Ko SJ (2021) Pepsi++: Fast and lightweight network for image inpainting. *IEEE Trans Neural Networks Learn Syst* 32(1):252–265
- Shukla T, Maheshwari P, Singh R, Shukla A, Kulkarni K, Turaga P (2023) Scene Graph Driven Text-Prompt Generation for Image Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pp 759–768
- Simonyan K, Zisserman A (2014) Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:14091556*
- Sohl-Dickstein J, Weiss E, Maheswaranathan N, Ganguli S (2015) Deep Unsupervised Learning using Nonequilibrium Thermodynamics. *In: Int. Conf. Mach. Learn.*, vol 37, pp 2256–2265
- Song L, Cao J, Song L, Hu Y, He R (2019) Geometry-Aware Face Completion and Editing. *In: AAAI Conf. Artificial Intell.*, pp 2506–2513
- Song Y, Yang C, Lin Z, Liu X, Huang Q, Li H, Kuo CCJ (2018a) Contextual-Based Image Inpainting: Infer, Match, and Translate. *In: Eur. Conf. Comput. Vis.*, pp 3–18
- Song Y, Yang C, Shen Y, Wang P, Huang Q, Kuo CCJ (2018b) SPG-Net: Segmentation Prediction and Guidance Network for Image Inpainting. *In: Brit. Mach. Vis. Conf.*
- Sun D, Yang X, Liu MY, Kautz J (2018a) PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 8934–8943
- Sun K, Xiao B, Liu D, Wang J (2019) Deep High-Resolution Representation Learning for Human Pose Estimation. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 5686–5696
- Sun Q, Ma L, Joon Oh S, Gool LV, Schiele B, Fritz M (2018b) Natural and Effective Obfuscation by Head Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 5050–5059
- Suvorov R, Logacheva E, Mashikhin A, Remizova A, Ashukha A, Silvestrov A, Kong N, Goka H, Park K, Lempitsky V (2022) Resolution-robust Large Mask Inpainting with Fourier Convolutions. *In: Winter Conf. App. Comput. Vis.*, pp 3172–3182
- Tabak EG, Vandenberg E (2010) Density estimation by dual ascent of the log-likelihood. *Commun Math Sci* 8(1):217 – 233
- Tschumperlé D, Deriche R (2005) Vector-valued image regularization with pdes: a common framework

- for different applications. *IEEE Trans Pattern Anal Mach Intell* 27(4):506–517
- Tu CT, Chen YF (2019) Facial Image Inpainting with Variational Autoencoder. In: *Inter. Conf. Intell. Robot. Control Engin.*, pp 119–122
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser u, Polosukhin I (2017) Attention is All You Need. In: *Adv. Neural Inform. Process. Syst.*, p 6000–6010
- Vo HV, Duong NQK, Pérez P (2018) Structural Inpainting. In: *ACM Int. Conf. Multimedia*, p 1948–1956
- Wadhwa G, Dhall A, Murala S, Tariq U (2021) Hyperrealistic Image Inpainting with Hypergraphs. In: *Winter Conf. App. Comput. Vis.*, pp 3911–3920
- Wan Z, Zhang B, Chen D, Zhang P, Chen D, Liao J, Wen F (2020) Bringing Old Photos Back to Life. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp 2747–2757
- Wan Z, Zhang J, Chen D, Liao J (2021) High-Fidelity Pluralistic Image Completion With Transformers. In: *Int. Conf. Comput. Vis.*, pp 4692–4701
- Wang C, Huang H, Han X, Wang J (2019a) Video Inpainting by Jointly Learning Temporal Structure and Spatial Details. In: *AAAI Conf. Artificial Intell.*, pp 5232–5239
- Wang C, Zhu Y, Yuan C (2022a) Diverse Image Inpainting with Normalizing Flow. In: *Eur. Conf. Comput. Vis.*, pp 53–69
- Wang J, Wang C, Huang Q, Shi Y, Cai JF, Zhu Q, Yin B (2020a) Image Inpainting Based on Multi-Frequency Probabilistic Inference Model. In: *ACM Int. Conf. Multimedia*, p 1–9
- Wang N, Li J, Zhang L, Du B (2019b) MUSICAL: Multi-Scale Image Contextual Attention Learning for Inpainting. In: *Int. Joint Conf. Artificial Intell.*, pp 3748–3754
- Wang N, Ma S, Li J, Zhang Y, Zhang L (2020b) Multi-stage attention network for image inpainting. *Pattern Recog* 106:107448
- Wang N, Zhang Y, Zhang L (2021a) Dynamic selection network for image inpainting. *IEEE Trans Image Process* 30:1784–1798
- Wang S, Saharia C, Montgomery C, Pont-Tuset J, Noy S, Pellegrini S, Onoe Y, Laszlo S, Fleet DJ, Soricut R, et al. (2023) Imagen editor and editbench: Advancing and evaluating text-guided image inpainting. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp 18359–18369
- Wang T, Ouyang H, Chen Q (2021b) Image Inpainting with External-internal Learning and Monochromic Bottleneck. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp 5116–5125
- Wang TC, Liu MY, Zhu JY, Liu G, Tao A, Kautz J, Catanzaro B (2018a) Video-to-Video Synthesis. In: *Adv. Neural Inform. Process. Syst.*, vol 31
- Wang W, Zhang J, Niu L, Ling H, Yang X, Zhang L (2021c) Parallel Multi-Resolution Fusion Network for Image Inpainting. In: *Int. Conf. Comput. Vis.*, pp 14559–14568
- Wang W, Niu L, Zhang J, Yang X, Zhang L (2022b) Dual-path Image Inpainting with Auxiliary GAN Inversion. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp 11411–11420
- Wang X, Girshick RB, Gupta A, He K (2018b) Non-local Neural Networks. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp 7794–7803
- Wang Y, Tao X, Qi X, Shen X, Jia J (2018c) Image Inpainting via Generative Multi-column Convolutional Neural Networks. In: *Adv. Neural Inform. Process. Syst.*, pp 331–340
- Wang Y, Chen YC, Tao X, Jia J (2020c) VNet: A Robust Approach to Blind Image Inpainting. In: *Eur. Conf. Comput. Vis.*
- Wang Z, Simoncelli E, Bovik A (2003) Multiscale structural similarity for image quality assessment. In: *Asilomar Conf. Sign. Syst. Comput.*, vol 2, pp 1398–1402
- Wang Z, Bovik AC, Sheikh HR, Simoncelli EP (2004) Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Process* 13(4):600–612
- Weng Y, Ding S, Zhou T (2022) A Survey on Improved GAN based Image Inpainting. In: *Inter. Conf. Consumer Electronics and Comput. Engin.*, pp 319–322
- Wexler Y, Shechtman E, Irani M (2007) Space-time completion of video. *IEEE Trans Pattern Anal Mach Intell* 29(3):463–476
- Woo S, Kim D, Park K, Lee JY, Kweon IS (2020) Align-and-Attend Network for Globally and Locally Coherent Video Inpainting. In: *Brit. Mach. Vis. Conf.*, pp 1–13
- Wu H, Zhou J, Li Y (2022) Deep generative model for image inpainting with local binary pattern learning and spatial attention. *IEEE Trans Multimedia* 24:4016–4027
- Wu L, Zhang C, Liu J, Han J, Liu J, Ding E, Bai X (2019) Editing Text in the Wild. In: *ACM Int. Conf. Multimedia*, p 1500–1508
- Wu X, Xie Y, Zeng J, Yang Z, Yu Y, Li Q, Liu W (2021) Adversarial learning with mask reconstruction for text-guided image inpainting. In: *ACM Int. Conf. Multimedia*, pp 3464–3472
- Xia W, Zhang Y, Yang Y, Xue JH, Zhou B, Yang MH (2022) Gan inversion: A survey. *IEEE Trans Pattern Anal Mach Intell* pp 1–17

- Xie C, Liu S, Li C, Cheng MM, Zuo W, Liu X, Wen S, Ding E (2019) Image Inpainting with Learnable Bidirectional Attention Maps. *In: Int. Conf. Comput. Vis.*, pp 8858–8867
- Xie M, Li C, Liu X, Wong TT (2020) Manga filling style conversion with screentone variational autoencoder. *ACM Trans Graph* 39(6)
- Xie M, Xia M, Liu X, Li C, Wong TT (2021) Seamless manga inpainting with semantics awareness. *ACM Trans Graph* 40(4)
- Xie S, Zhang Z, Lin Z, Hinz T, Zhang K (2023) Smart-Brush: Text and Shape Guided Object Inpainting With Diffusion Model. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 22428–22437
- Xie Y, Lin Z, Yang Z, Deng H, Wu X, Mao X, Li Q, Liu W (2022) Learning semantic alignment from image for text-guided image inpainting. *The Visual Computer* 38(9-10):3149–3161
- Xiong W, Yu J, Lin Z, Yang J, Lu X, Barnes C, Luo J (2019) Foreground-Aware Image Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 5833–5841
- Xu N, Yang L, Fan Y, Yang J, Yue D, Liang Y, Price B, Cohen S, Huang T (2018a) YouTube-VOS: Sequence-to-Sequence Video Object Segmentation. *In: Eur. Conf. Comput. Vis.*, pp 585–601
- Xu R, Li X, Zhou B, Loy CC (2019) Deep flow-guided video inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 3723–3732
- Xu R, Guo M, Wang J, Li X, Zhou B, Loy CC (2021) Texture memory-augmented deep patch-based image inpainting. *IEEE Trans Image Process* 30:9112–9124
- Xu T, Zhang P, Huang Q, Zhang H, Gan Z, Huang X, He X (2018b) AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 1316–1324
- Yamashita Y, Shimosato K, Ukita N (2022) Boundary-Aware Image Inpainting With Multiple Auxiliary Cues. *In: IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pp 619–629
- Yan Z, Li X, Li M, Zuo W, Shan S (2018) Shift-Net: Image Inpainting via Deep Feature Rearrangement. *In: Eur. Conf. Comput. Vis.*, pp 3–19
- Yang C, Lu X, Lin Z, Shechtman E, Wang O, Li H (2017) High-resolution image inpainting using multi-scale neural patch synthesis. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 6721–6729
- Yang J, Qi Z, Shi Y (2020) Learning to Incorporate Structure Knowledge for Image Inpainting. *In: AAAI Conf. Artificial Intell.*, vol 34, pp 12605–12612
- Yang L, Zhang Z, Song Y, Hong S, Xu R, Zhao Y, Zhang W, Cui B, Yang MH (2023) Diffusion models: A comprehensive survey of methods and applications. [2209.00796](https://arxiv.org/abs/2209.00796)
- Yeh RA, Chen C, Lim TY, Schwing AG, Hasegawa-Johnson M, Do MN (2017) Semantic Image Inpainting with Deep Generative Models. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 6882–6890
- Yi Z, Tang Q, Azizi S, Jang D, Xu Z (2020) Contextual Residual Aggregation for Ultra High-Resolution Image Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 7508–7517
- Yu F, Koltun V (2016) Multi-Scale Context Aggregation by Dilated Convolutions. *In: Int. Conf. Learn. Represent.*
- Yu J, Lin Z, Yang J, Shen X, Lu X, Huang TS (2018) Generative image inpainting with contextual attention. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 5505–5514
- Yu J, Lin Z, Yang J, Shen X, Lu X, Huang TS (2019) Free-form image inpainting with gated convolution. *In: Int. Conf. Comput. Vis.*, pp 4471–4480
- Yu T, Guo Z, Jin X, Wu S, Chen Z, Li W, Zhang Z, Liu S (2020) Region Normalization for Image Inpainting. *In: AAAI Conf. Artificial Intell.*, pp 12733–12740
- Yu Y, Zhan F, Lu S, Pan J, Ma F, Xie X, Miao C (2021a) WaveFill: A Wavelet-Based Generation Network for Image Inpainting. *In: Int. Conf. Comput. Vis.*, pp 14114–14123
- Yu Y, Zhan F, WU R, Pan J, Cui K, Lu S, Ma F, Xie X, Miao C (2021b) Diverse Image Inpainting with Bidirectional and Autoregressive Transformers. *In: ACM Int. Conf. Multimedia*, p 69–78
- Yu Y, Du D, Zhang L, Luo T (2022a) Unbiased Multimodality Guidance for Image Inpainting. *In: Eur. Conf. Comput. Vis.*, pp 668–684
- Yu Y, Zhang L, Fan H, Luo T (2022b) High-Fidelity Image Inpainting with GAN Inversion. *In: Eur. Conf. Comput. Vis.*, pp 242–258
- Zeng Y, Fu J, Chao H, Guo B (2019) Learning pyramid-context encoder network for high-quality image inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 1486–1494
- Zeng Y, Fu J, Chao H (2020a) Learning Joint Spatial-Temporal Transformations for Video Inpainting. *In: Eur. Conf. Comput. Vis.*, Springer, pp 528–543
- Zeng Y, Lin Z, Yang J, Zhang J, Shechtman E, Lu H (2020b) High-Resolution Image Inpainting with Iterative Confidence Feedback and Guided Upsampling. *In: Eur. Conf. Comput. Vis.*
- Zeng Y, Gong Y, Zhang J (2021a) Feature learning and patch matching for diverse image inpainting. *Pattern Recog* 119:108036
- Zeng Y, Lin Z, Lu H, Patel VM (2021b) CR-Fill: Generative Image Inpainting With Auxiliary Context

- tual Reconstruction. *In: Int. Conf. Comput. Vis.*, pp 14164–14173
- Zeng Y, Fu J, Chao H, Guo B (2022) Aggregated contextual transformations for high-resolution image inpainting. *IEEE Trans Vis Comput Graph* pp 1–1
- Zhang B, Gao Y, Zhao S, Liu J (2010) Local derivative pattern versus local binary pattern: Face recognition with high-order local pattern descriptor. *IEEE Trans Image Process* 19(2):533–544
- Zhang H, Hu Z, Luo C, Zuo W, Wang M (2018a) Semantic Image Inpainting with Progressive Generative Networks. *In: ACM Int. Conf. Multimedia*, p 1939–1947
- Zhang H, Mai L, Xu N, Wang Z, Collomosse J, Jin H (2019a) An internal learning approach to video inpainting. *In: Int. Conf. Comput. Vis.*, pp 2720–2729
- Zhang J, Niu L, Yang D, Kang L, Li Y, Zhao W, Zhang L (2019b) GAIN: Gradient Augmented Inpainting Network for Irregular Holes. *In: ACM Int. Conf. Multimedia*, p 1870–1878
- Zhang K, Fu J, Liu D (2022a) Flow-Guided Transformer for Video Inpainting. *In: Eur. Conf. Comput. Vis.*, pp 74–90
- Zhang K, Fu J, Liu D (2022b) Inertia-Guided Flow Completion and Style Fusion for Video Inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 5982–5991
- Zhang L, Agrawala M (2023) Adding conditional control to text-to-image diffusion models. [2302.05543](#)
- Zhang L, Chen Q, Hu B, Jiang S (2020a) Text-Guided Neural Image Inpainting. *In: ACM Int. Conf. Multimedia*, p 1302–1310
- Zhang L, Barnes C, Wampller K, Amirghodsi S, Shechtman E, Lin Z, Shi J (2022c) Inpainting at Modern Camera Resolution by Guided PatchMatch with Auto-curation. *In: Eur. Conf. Comput. Vis.*, pp 51–67
- Zhang L, Zhou Y, Barnes C, Amirghodsi S, Lin Z, Shechtman E, Shi J (2022d) Perceptual Artifacts Localization for Inpainting. *In: Computer Vision – ECCV 2022*, pp 146–164
- Zhang R, Isola P, Efros AA, Shechtman E, Wang O (2018b) The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 586–595
- Zhang R, Quan W, Wu B, Li Z, Yan DM (2020b) Pixel-wise dense detector for image inpainting. *Comput Graph Forum* 39(7)
- Zhang R, Quan W, Zhang Y, Wang J, Yan DM (2022e) W-net: Structure and texture interaction for image inpainting. *IEEE Trans Multimedia* pp 1–12
- Zhang S, He R, Sun Z, Tan T (2018c) Demeshnet: Blind face inpainting for deep meshface verification. *IEEE Trans Inf Forensics Secur* 13(3):637–647
- Zhang W, Zhu J, Tai Y, Wang Y, Chu W, Ni B, Wang C, Yang X (2021) Context-Aware Image Inpainting with Learned Semantic Priors. *In: Int. Joint Conf. Artificial Intell.*, pp 1323–1329
- Zhang Z, Zhao Z, Zhang Z, Huai B, Yuan J (2020c) Text-guided image inpainting. *In: ACM Int. Conf. Multimedia*, pp 4079–4087
- Zhao L, Mo Q, Lin S, Wang Z, Zuo Z, Chen H, Xing W, Lu D (2020) UCTGAN: Diverse Image Inpainting Based on Unsupervised Cross-Space Translation. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 5740–5749
- Zhao S, Cui J, Sheng Y, Dong Y, Liang X, Chang EI, Xu Y (2021) Large Scale Image Completion via Co-Modulated Generative Adversarial Networks. *In: Int. Conf. Learn. Represent.*
- Zhao W, Rao Y, Liu Z, Liu B, Zhou J, Lu J (2023) Unleashing text-to-image diffusion models for visual perception. *arXiv preprint arXiv:230302153*
- Zheng C, Cham TJ, Cai J (2019) Pluralistic Image Completion. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 1438–1447
- Zheng C, Cham TJ, Cai J (2021a) Pluralistic free-form image completion. *Int J Comput Vis* 129:2786–2805
- Zheng C, Cham TJ, Cai J, Phung D (2022a) Bridging Global Context Interactions for High-Fidelity Image Completion. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 11512–11522
- Zheng H, Zhang Z, Wang Y, Zhang Z, Xu M, Yang Y, Wang M (2021b) GCM-Net: Towards Effective Global Context Modeling for Image Inpainting. *In: ACM Int. Conf. Multimedia*, p 2586–2594
- Zheng H, Lin Z, Lu J, Cohen S, Shechtman E, Barnes C, Zhang J, Xu N, Amirghodsi S, Luo J (2022b) Image Inpainting with Cascaded Modulation GAN and Object-Aware Training. *In: Eur. Conf. Comput. Vis.*, pp 277–296
- Zhou B, Lapedriza A, Khosla A, Oliva A, Torralba A (2017) Places: A 10 million image database for scene recognition. *IEEE Trans Pattern Anal Mach Intell* 40(6):1452–1464
- Zhou B, Zhao H, Puig X, Xiao T, Fidler S, Barriuso A, Torralba A (2018) Semantic understanding of scenes through the ade20k dataset. *Int J Comput Vis* 127:302–321
- Zhou X, Li J, Wang Z, He R, Tan T (2021) Image Inpainting with Contrastive Relation Network. *In: Int. Conf. Pattern Recog.*, pp 4420–4427
- Zhu M, He D, Li X, Li C, Li F, Liu X, Ding E, Zhang Z (2021) Image inpainting by end-to-end cascaded refinement with mask awareness. *IEEE Trans Image Process* 30:4855–4866

Zou X, Yang L, Liu D, Lee YJ (2021) Progressive temporal feature alignment network for video inpainting. *In: IEEE Conf. Comput. Vis. Pattern Recog.*, pp 16448–16457