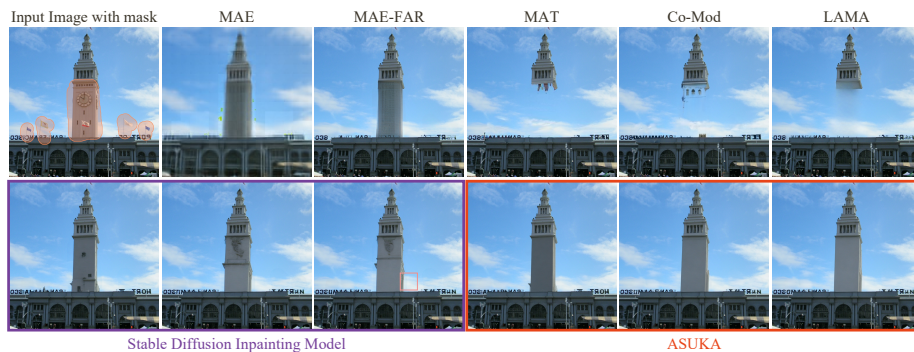


# Towards Context-Stable and Visual-Consistent Image Inpainting

Yikai Wang\*, Chenjie Cao\*, Ke Fan, Xiangyang Xue, and Yanwei Fu†

School of Data Science, Fudan University  
{yikaiwang19, yanweifu}@fudan.edu.cn

Project Page: <https://yikai-wang.github.io/asuka/>  
MISATO Dataset: <https://github.com/Yikai-Wang/asuka-misato>



**Fig. 1:** Comparison on 1024<sup>2</sup> image (depicted from MegaDepth [26]) between our proposed ASUKA and reconstruction-based (MAE [19]) as well as generative-based (MAE-FAR [7], MAT [25], Co-Mod [59], LaMa [45], Stable Diffusion [39]) inpainting models. MAE offers stable estimates, yet **falls short in texture detail**. GAN-based inpainting struggles with **low fidelity**. Powerful SD **hallucinates random elements** and suffers from **mask-unmask visual inconsistency** (color of clouds in the orange square). ASUKA achieves context-stable and visual-consistent inpainting.

**Abstract.** Recent progress in inpainting increasingly relies on generative models, leveraging their strong generation capabilities for addressing large irregular masks. However, this enhanced generation often introduces context-instability, leading to arbitrary object generation within masked regions. This paper proposes a balanced solution, emphasizing the importance of unmasked regions in guiding inpainting while preserving generation capacity. Our approach, Aligned Stable Inpainting with UnKnown Areas Prior (ASUKA), employs a Masked Auto-Encoder (MAE) to produce reconstruction-based prior. Aligned with the powerful Stable Diffusion inpainting model (SD), ASUKA significantly improves context stability. ASUKA further adopts an inpainting-specialized decoder, highly reducing the color inconsistency issue of SD and thus ensuring more visual-consistent inpainting. We validate effectiveness of inpainting algorithms on benchmark dataset Places 2 and a collection of several existing datasets, dubbed MISATO, across diverse domains



**Fig. 2:** Object removal on  $1024^2$  images. ASUKA enjoys context-stable (first to second rows) and visual-consistent (second to third rows) inpainting compared with SD.

and masking scenarios. Results on these benchmark datasets confirm ASUKA’s efficacy in both context-stability and visual-consistency compared to SD and other inpainting algorithms.

**Keywords:** Image Inpainting · Context-Stability · Visual-Consistency

## 1 Introduction

Image inpainting [5] fills masked areas of images while maintaining consistency with the unmasked regions. Traditional inpainting algorithms [5, 13, 18, 24, 40] often result in blurred synthesis when reconstructing masked regions, as indicated by [36]. The advent of Generative Adversarial Networks (GANs) has led to works [8, 15, 17, 21, 27, 35, 45, 46, 59] that utilize GANs to fill complex mask structures, achieving impressive inpainting results. However, they still struggle with general challenging inpainting cases, particularly in filling large holes.

The recent Stable Diffusion Inpainting (SD) [39] model, with its large model capacity and extensive dataset training, acts as a versatile tool for image inpainting. However, we have observed a phenomenon of "*context-instability*" problem, where the model generates random, unreasonable elements to fill masked regions of images, as depicted in first to second rows in Fig. 2, as well as Fig. 1 and Fig. 8. This issue appears to come from the random masking strategy used to train SD. This strategy introduces training cases where foreground objects are completely masked and the SD is forced to fill masked regions with foreground objects. Consequently, SD will hallucinate unreasonable objects devoid

of contextual information. Adjusting prompts may reduce this risk, but the best prompt is image-dependent, making it infeasible for practical applications.

Furthermore, the inpainted results of SD may suffer from "*visual inconsistency*" problem. This problem, less explored in academia but critical for real-world industrial applications, results in significant color discrepancies between inpainted and unmasked regions and exhibits smear-like traces in the image, as shown in the second to third rows in Fig. 2, as well as Fig. 1 and Fig. 8. Essentially, this visual-inconsistency comes from the misalignment between the pixel distributions of filled results and original images due to information loss in the Kullback-Leibler Variational Autoencoder (KL-VAE) used in SD, as illustrated in Fig. 5. To improve visual consistency, an inpainting-specialized VAE decoder is necessary to explicitly mitigate information loss based on available ground-truth information of unmasked images.

To efficiently tackle context-instability and visual inconsistency issues, we enhance the SD [39] with regression-based reconstruction and distribution-aligned generation. This results in an improved image inpainting model that avoids generating unreasonable elements in the masked region and reduces mask-unmask color inconsistency. The vanilla SD [39] adopts a KL-VAE to compress image into latent and perform diffusion in the latent space. We manipulate its generation and decoding processes for better context stability and visual consistency. **i)** We propose using the Masked Auto-Encoder (MAE) [19] as a prior to guide and stabilize the SD’s generation process. As shown in Fig. 2, MAE yields stable yet blurred results, while SD may produce implausible content despite its impressive generation capacity. By aligning MAE prior with SD objectives, we improve context-stability without compromising performance. **ii)** We redesign the KL-VAE decoder in SD to address color inconsistencies between masked and unmasked regions. The new decoder, acting as a local harmonization model conditioned on unmasked image pixels, significantly enhances visual consistency.

Formally, to enhance the context-stability and visual-consistency of image inpainting model, we present the Aligned Stable Inpainting with Unknown Areas prior (ASUKA) framework. **i)** To enhance context-stability, we replace the text-condition part of the SD with the MAE prior. We fine-tune the MAE to ensure adaptability to practical inpainting scenarios, as MAE is trained on uniformly distributed masks but inpainting usually contains large continuous masks. We train an alignment module to align MAE prior with SD using the standard SD objective while freezing the fine-tuned MAE and SD. **ii)** To address visual-consistency issues during decoding, we redesign the decoder of the KL-VAE in SD. We inject additional inputs of unmasked image pixels and 0-1 masks to the decoder. This new decoder is trained to perform local harmonization task, taking as inputs of inharmonious latent and taking as outputs the original harmonious image. We generate inharmonious latent via the proposed color and latent augmentation techniques to incorporate KL-VAE information loss and generation artifacts of SD. Then the decoder is trained to recover the original image from corrupted latent based on the ground-truth information of unmasked

image. These steps collectively enable ASUKA to achieve more context-stable and visual-consistent inpainting results.

To evaluate the effectiveness of inpainting algorithms across various scenarios and mask shapes, in addition to the benchmark dataset Places 2 [61], we further utilize an evaluation dataset named MISATO, which selects representative testing images from Matterport3D [10], Flickr-Landscape [28], MegaDepth [26], and COCO 2014 [29]. This dataset thus covers four distinct domains—landscape, indoor, building, and background—making it diverse enough to serve as a benchmark for evaluation. Experiments on MISATO and Places 2, especially with large irregular masks, validate the efficacy of ASUKA.

**Contributions.** ASUKA enhances image inpainting from context-stability and visual-consistency while leveraging the generation capacity of the frozen SD inpainting model. It achieves this through two main components: *i) Context-Stable Alignment*: ASUKA aligns the stable MAE prior with SD to provide a context-stable estimation of masked regions, replacing the text-condition part of SD with MAE prior. *ii) Visual-Consistent Alignment*: ASUKA re-formulates the decoding from SD latent to image as a local harmonization task, trains an inpainting-specialized decoder to align masked and unmasked regions during decoding and thus mitigates color inconsistencies caused by information loss from the SD.

## 2 Related Works

**Image inpainting** is a longstanding challenge to infill the missing regions with consistent pixels. Classical approaches, relying on patchmatch [4, 14, 56] and partial differential equation [5, 6, 9] are constrained by low-level image features and often fail to produce reasonable results for large holes. GAN-based [17] inpainting methods [7, 25, 36, 53, 59] have evolved with adaptive convolutions [30, 53, 54], attention mechanism [23, 51, 52, 55], and high-resolution generalization in frequency learning [12, 45, 50]. Methods like Co-Mod [59] prioritize adversarial learning and plausible generation to address the challenging ill-posed inpainting issue [25, 60], yielding robust and high-quality generations, even for large masks. However, they suffer from unstable results and uncontrollable hallucinations conditioned on different random latent codes. On the other hand, methods emphasizing higher reconstruction penalties [7, 35, 45] achieve more stable generations, but struggle with large masks, often resulting in blur artifacts. Recent diffusion models [39, 41] have rapidly developed and demonstrated impressive inpainting results. However, these models share similarities with GAN-based methods, learning to align distributions rather than achieving faithful pixel space reconstruction. As a consequence, diffusion-based approaches also suffer from context-instability issue.

**Adapting Stable Diffusion.** SD [39] is one of the most widely used diffusion models. SD incorporates a KL-VAE to learn a latent space preserving image semantics at a lower resolution, along with a U-Net for diffusion within this latent space. Various tuning and aligning methods have been developed to introduce additional conditions to SD, including image-inversion for text-guided image-to-image translation [34], textual-inversion for personalization [16], LoRA

fine-tuning [22], and controlnet [57] to add other types of conditions. Our objective is to enhance context-stability while preserving the generation capacity of SD. Consequently, we refrain from fine-tuning the U-Net backbone of SD. The text-condition is not utilized in inpainting tasks. Textual-inversion [16] introduces learnable tokens in Clip text encoders to improve subject description. ObjectStitch [43] employs the CLIP visual feature to represent the subject. In our approach, we advocate removing the text-condition and instead guiding SD using the Masked Auto-Encoder (MAE) prior for masked region.


**Information loss in SD latent space.** Although claimed only eliminates imperceptible details, the KL-VAE used by SD causes distortion in the reconstruction of images. See Fig. 6a for illustrative examples. OpenAI [2] proposes a larger decoder to improve the decoding quality of SD’s latent. Luo *et al.* [33] propose a frequency-augmented decoder to address the super-resolution case. Zhu *et al.* [62] propose to preserve unmasked regions during decoding. In this paper, we ensure the low-frequency color visual-consistency in the decoding process of SD.

**Masked Image-Modeling [3]** (MIM) is an active research area in self-supervised learning. Typical MIM methods [3, 11, 19, 49] split images into visible and masked patches, learning to estimate masked patches from visible patches. Training targets for visible patches encompass pixel values [19], HOG features [47], and high-level semantic features [48]. While the primary objective of MIM is representation learning, its potential effectiveness in image generation is also noteworthy. Cao *et al.* [7] adopts MAE features and attention scores to assist the convolutional inpainting model better in handling long-distance dependencies. In contrast, this paper uses MAE prior to enhance the context-stability of SD.


### 3 Methodology

**Problem setup.** Inpainting takes as inputs a masked image to complete with a mask to indicate the missing region. The target of inpainting is to fill the missing region based on the information of unmasked regions to generate high-fidelity images. In this paper, we focus on the standard inpainting task without utilizing other conditions. We adopt stable diffusion inpainting model [39] and focus on the general issues of inpainting models, i) **context-instability issue**: unstable and uncontrollable hallucinations, yielding random elements generated in the masked region; ii) **visual-inconsistency issue**: mask-unmask color inconsistency issue, yielding smear-like traces in the masked region.

**Overview.** The framework of the proposed Aligned Stable inpainting with Unknown Areas prior (ASUKA) is illustrated in Fig. 3. ASUKA adopts the pre-trained stable diffusion inpainting model (SD) [39]. Our target is to provide more context-stable and visual-consistent inpainting results while fully exploiting the generation capacity of SD. ASUKA includes **i)** a *context-stable alignment* to align stable Masked Auto-Encoder (MAE) prior for masked region with SD and **ii)** a *visual-consistent alignment* to align ground-truth unmasked region with generated masked region during decoding. To this end, we freeze the U-Net of SD, while replacing the text-condition part with our proposed MAE prior to



**Fig. 3:** ASUKA tackles the context-instability issue by adopting the MAE to provide a stable prior for frozen SD to maintain the generation capacity while increasing context-stability. For the visual-inconsistency issue, ASUKA utilizes an inpainting-specialized decoder to achieve mask-unmask color consistency when decoding SD latent.



(a) Use MAE prior for image-to-image translation (start from 80% noise rate) via SD achieves poor inpainting results.

(b) Alignment module.

**Fig. 4:** (a) Use MAE as initial value v.s. as condition. (b) Alignment module.

increase context-stability. To align the MAE prior to SD, we introduce an alignment module, which is trained via the training objective of SD. Additionally, to align masked and unmasked regions during decoding and resolve the problem of information loss from the KL-VAE decoder which causes mask-unmask color inconsistency, we train an inpainting-specialized decoder to decode the latent back to the image space for seamless visual-consistency. Combined together, ASUKA achieves more context-stable and visual-consistent inpainting.

### 3.1 Masked Auto-Encoder Prior

**MAE as context-stable prior.** While recent generative models rely on random noise to provide more diverse generation results, it leads to the generation of random objects unexpectedly. In contrast, MAE is known to provide a context-stable estimation of masked regions based purely on the unmasked regions. As MAE is trained on the L2 reconstruction loss, we can regard the estimation of MAE as a mean estimation, which can be utilized to provide a context-stable prior for generative models to not generate new concepts. However, MAE itself results in average and blur generations and cannot reconstruct detailed textures of the masked region, and works poorly if we use MAE prior as the initial values for the SD to perform inpainting in image-to-image style, as illustrated in Fig. 4a. To this end, we adopt the MAE to provide prior to stabilizing SD.

**Training MAE.** The original MAE is trained to estimate random masks uniformly distributed in the image, while inpainting task usually contains large continuous masks. Inspired by Cao *et al.* [7], we fine-tune the MAE to inpainting masks. To adapt MAE for more practical inpainting scenarios, we construct a systematic masking strategy. The mask basis contains: object-shape mask, irregular mask, and regular mask. We collect object-shape masks from COCO [29] object segments. We use irregular masks from previous studies, including Co-Mod mask [59] and LaMa [45] mask. The regular masks contains rectangle and complement rectangle mask. To ensure generalization and coverage, for each mask we generate from mask basis with the probability of 50% object-shape, 40% irregular, and 10% regular. For object-shape mask basis, we combine it with irregular mask with the chance of 50%. This construction of mask style estimates the masks occurs in inpainting tasks, especially for the object removal and user-specified irregular masks. We control the mask ratio in the range of  $[0.1, 0.75]$  to follow the training scenario of MAE. For masks smaller than the ratio of 75%, we enlarge the mask ratio to 75% with randomly selected mask regions. This benefits ASUKA to tackle the large hole inpainting task.

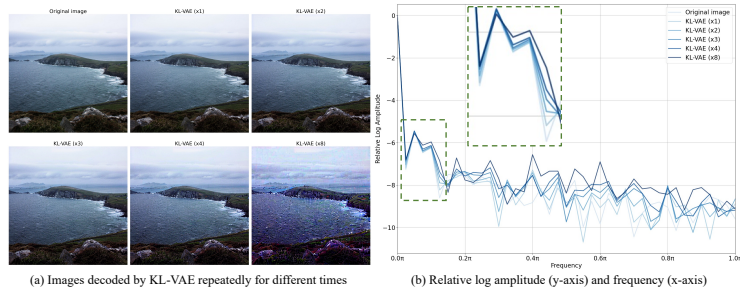
### 3.2 Align MAE Prior with SD for Context-Stability

**Replace text-condition with MAE prior.** SD is not trained on MAE priors. Fortunately, the text-condition part of SD is not utilized in inpainting, as we do not assume a text condition for inpainting task. Hence, we propose to replace the text-condition part of SD with our proposed MAE prior condition. However, as we do not fine-tune the SD, it cannot aligns well with the MAE prior. To achieve a proper guidance for SD, we introduce the alignment module to align MAE with SD in both dimension and distribution perspective, as shown in 4b.

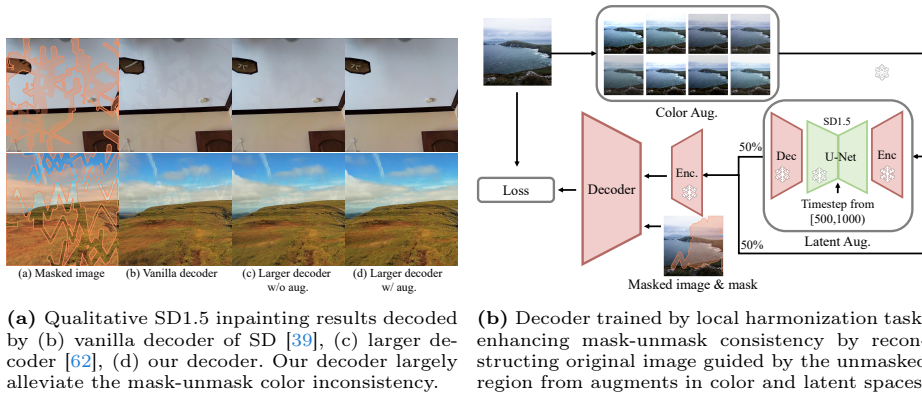
**Dimension alignment.** Particularly, the MAE prior  $F_{\text{MAE}}$  is of size  $N_m \times M_m$ , where  $N_m$  is the sequence length and  $M_m$  is the feature dimension. To align it with the SD condition of size  $N_s \times M_s$ , we adopt a linear layer to map the feature dimension from  $M_m$  to  $M_s$  and set  $N_s = N_m$  to preserve the local MAE prior.

**Distribution alignment.** After aligning the dimension, we use self-attention blocks to learn to better guiding SD, leading to the condition  $C_{\text{MAE}}$  for SD. We train our alignment module using the standard diffusion objective with the same masking strategy used to train the MAE, keeping other modules frozen.

**Misalignment.** When training the alignment module with the set (input image, MAE prior, SD result), misalignment issues may arise. For example, if an object is completely masked, the MAE will predict the masked area with background, whereas the SD is trained to recreate the object. This difference can lead the alignment module to mistakenly disregard the MAE prior. To address this, we improve the SD’s adherence to the MAE prior by substituting the MAE predicted prior with the MAE reconstructed prior at a probability of  $p$ . The MAE reconstructed prior involves using MAE to recreate the image without masking any area, ensuring MAE has access to all information needed for reconstruction. This approach helps train the alignment module to better guide the SD.



**Fig. 5:** The illustration of the information loss of KL-VAE. (a) The hue of the image is shifted. (b) KL-VAE still suffers from non-ignorable shifts in low-frequency fields.



(a) Qualitative SD1.5 inpainting results decoded by (b) vanilla decoder of SD [39], (c) larger decoder [62], (d) our decoder. Our decoder largely alleviate the mask-unmask color inconsistency. (b) Decoder trained by local harmonization task, enhancing mask-unmask consistency by reconstructing original image guided by the unmasked region from augments in color and latent spaces.

**Fig. 6:** ASUKA trains a specialized-decoder to solve the color-inconsistency issue from the information loss of KL-VAE and domain gap between generated and real latent.

### 3.3 Mask-Unmask Align during Decoding for Visual-Consistency

**Information loss of KL-VAE.** SD performs all its generative processes in the latent space and subsequently decodes these latent codes back to image space using KL-VAE. Despite the decoder being trained to reconstruct the image, it encounters challenges associated with information loss. Particularly in tasks like inpainting, we have ground-truth values for the unmasked region. Though Rombach *et al.* [39] claimed that the diffusion model should prioritize the informative semantic compression, while the VAE and GAN are used to tackle perceptual compression with high-frequency details, we argue that *low-frequency semantic loss in KL-VAE could not be neglected*, as verified in Fig. 5 (b). The KL-VAE will not only noticeably degrades high-frequency details but also shifts in semantic hue. This shift can be verified by repeated reconstruction with KL-VAE, as shown in Fig. 5 (a) where larger shift is observed during repeated reconstruction. As human is sensitive to low-frequency information changes in the image, even subtle color shifts can induce significant visual-inconsistencies. This visual-inconsistency issue is more severe under the irregular or large mask structures.



**Unmask-region conditioned decoder.** The aforementioned information loss is prohibitive to be addressed by the SD U-Net alone, especially given the already degraded input latent codes. Zhu *et al.* [62] adopts a larger decoder, with additional inputs of masked images. However, it still fails to handle the incompatible color and texture between the original images and compressed ones in challenging scenes as verified in Fig. 6a (c). The low-frequency gap between degraded and original images makes it challenging to explicitly address this issue. To solve the mask-unmask visual-inconsistency issue, we reformulate the decoding as a local harmonization task to preserve the visual consistency. As shown in Fig. 3, compared with SD decoder, our decoder involves additional inputs of masked images in the pixel-wise color space and the 0-1 mask.


**Mask-unmask visual-consistent decoder.** To train the decoder to ensure visual-consistency between generated latent and unmasked pixels, we re-formulate the decoding as a local harmonization task. Specifically, we propose the color and latent augmentation as shown in Fig. 6b to estimate and enlarge the visual-inconsistency during decoding. We follow the standard KL-VAE training pipeline of SD, but replacing the inputs with augmented ones. Particularly, we use the original image as the reconstruction target and use color and latent augmentation to corrupt input image, simulating the information loss of KL-VAE and domain gap between generated and real latent, respectively. This forces the decoder to reconstruct the un-corrupted image based on the ground-truth unmasked region.

**i) Color augmentation:** The first information loss comes from the imperfect VAE reconstruction. VAE is trained to encode and decode the input image for reconstruction, but suffers from color shift as shown in Fig. 6a (b). Empirically, further conditioned on unmasked image alleviate but not solve the color inconsistency issue, as shown in Fig. 6a (c). Hence, we need to explicitly train the decoder to ensure color consistency. To this end, we augment all training images in brightness, contrast, saturation, and hue, and requires the decoder to reconstruct original image conditioned on the unaugmented unmasked image. This encourages the decoder to fully utilize the information from unmasked regions.

**ii) Latent augmentation:** The other information loss comes from the distribution gap between SD generated latent and real latent. This gap also causes color inconsistency even if we alleviate the VAE reconstruction loss, as shown in Fig. 7a. To simulate this loss, we incorporate the artifacts generated from the SD to train the decoder. However, denoising to real images iteratively is notably time-consuming, even with DDIM [42]. To balance the efficiency and efficacy, we design a one-step estimation of SD. As our target is to capture the generation gap, we use the clean latent  $\mathbf{z}_0$  and all-zero mask  $\mathbf{M}$  as conditions. This tells the SD all the needed information to generate the clean latent, ensuring the generated latent preserves content and only shift from the distribution gap. We follow the standard pipeline to estimate  $\mathbf{z}_0$  with modified conditions as:

$$\hat{\mathbf{z}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{z}_t - \sqrt{1 - \bar{\alpha}_t}\varepsilon_\theta([\mathbf{z}_t; \mathbf{z}_0; \mathbf{M}], t)), \quad (1)$$

where the timestep  $t$  is randomly sampled from [500, 1000);  $\bar{\alpha}_t$  indicates the prescribed variance schedule in SD;  $\varepsilon_\theta(\cdot)$  is the frozen SD take as inputs noised



**Fig. 7:** (a) Ablation of latent augmentation; (b) MISATO evaluation examples.

$z_t$ , unmasked  $z_0$ , and all-zero masking  $M$ . The large step denoising is chosen to increase the distribution gap, as empirically SD could produce reliable results in small  $t$  given the unmasked latent condition  $z_0$ . Then we decode  $z_0$  to image as the latent augmented inputs. This makes latent augmentation an off-line strategy. We apply latent augmentation to 50% training images. We fine-tune the decoder based on [62], showcasing superior consistency as compared in Fig. 6a.

## 4 Evaluation Dataset

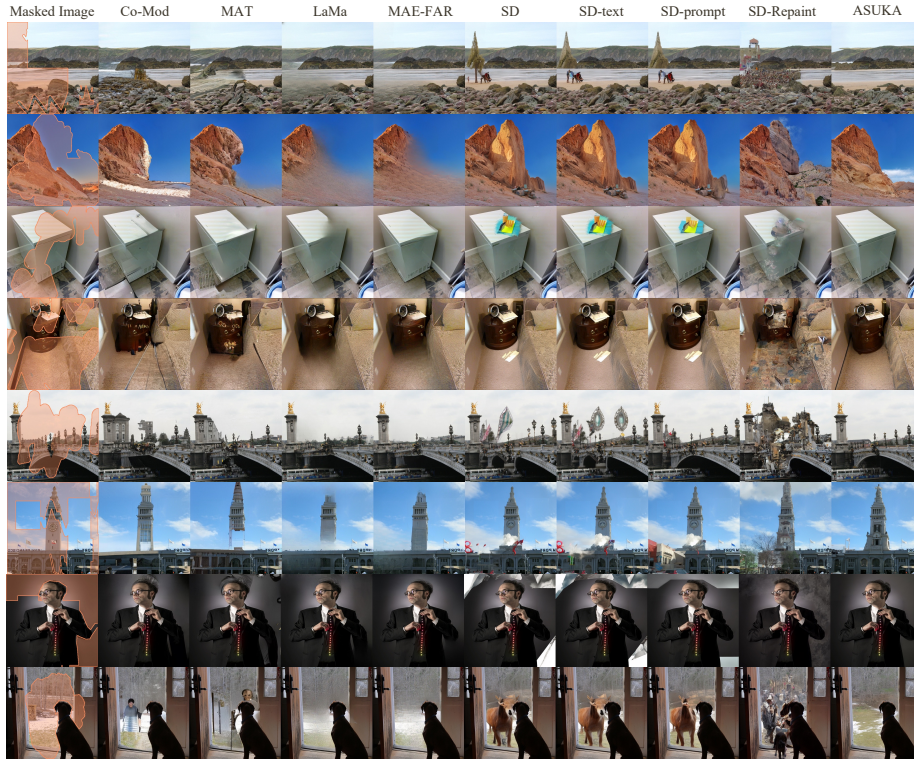
In addition to the standard benchmark Places 2 [61], to validate the inpainting performance across different domains and mask styles, we construct a evaluation dataset, dubbed as MISATO, from Matterport3D [10], Flickr-Landscape [28], MegaDepth [26], and COCO 2014 [29]. The first three datasets contains images from indoor, outdoor landscape and building, respectively. COCO has object segmentation. We select 500 representative examples of size  $512^2$  from each dataset. The MISATO dataset is available at <https://github.com/Yikai-Wang/asuka-misato>.

The principle of constructing MISATO is to select the most representative and diverse examples. To this end, for first three datasets, we use CLIP visual model [37] to extract semantic visual features. Then we use BisectingKMeans [44] to cluster each dataset into 500 clusters, and select the cluster centers as the evaluation data. The selected data are center cropped and then resized to  $512^2$ . For COCO, we focus on the background inpainting. To this end, for each data we identify the foreground with provided segmentation and remove it from the generated masks, yielding a dataset specified for purely background inpainting.

Combined together, MISATO contains 2000 examples from four inpainting domains, indoor, outdoor landscape, building, and background. As shown in Fig. 7b, we adopt the masking strategy as in Sec. 3.1 but excluding the rectangle and complement rectangle masks. The masking ratio is set as  $[0.2, 0.8]$ .

## 5 Experiments

**Implementation details.** ASUKA utilizes the stable diffusion text-guided inpainting model v1.5 as the generative backbone. We train ASUKA on Places2 [61].



**Fig. 8:** ASUKA v.s. Co-Mod [59], MAT [25], LaMa [45], MAE-FAR [7], SD [39], and SD-Repaint [32] in  $512^2$  landscape, indoor, building and background inpaintings. SD runs unconditional generation. SD-text uses “background” prompt as we have no caption in inpainting. SD-prompt uses learned prompt, under the same training setting of ASUKA. SR-Repaint uses Repaint [32] to modify SD outputs. GAN-based methods generate blurred results; SD-based methods hallucinate unreasonable objects and suffer from color shift. ASUKA achieves context-stable and visual-consistent inpainting.

For the MAE [19] used in ASUKA, we train on images of size  $256^2$ , which is efficient and produce context-stable guidance for SD to generate high-resolution images. We fine-tune from the pre-trained MAE with a batch size of 1024. We use four self-attention blocks for the alignment module, training from scratch with AdamW [31] of learning rate  $5e-2$  with the standard diffusion objective. We set  $p$  as 100% and linearly decay it to 10% in the first 2K training steps and then freeze. We fine-tune the decoder based on [62] for 50K steps with a batch size of 40 and learning rate of  $8e-5$  with cosine decay. We use ColorJitter for color augmentation, with brightness 0.15, contrast 0.2, saturation 0.1, and hue 0.03.

**Evaluation metrics.** We use the Learned Perceptual Image Patch Similarity (LPIPS) [58] to calculate the patch-level image distances, Fréchet Inception Distance (FID) [20] to compare the distribution distance between generated images and real images, and Paired/Unpaired Inception Discriminative Score (P-IDS/U-IDS) [59] to measure the human-inspired linear separability.

**Table 1:** Quantitative comparison on MISATO and Places 2. Co-Mod [59], MAT [25], LaMa [45], MAE-FAR [7] and SD-Repaint [32] are state-of-the-art inpainting methods. SD [39] performs unconditional generation. SD-text uses “background” text prompt to guide generation. SD-prompt uses learnable prompts trained specifically for inpainting, using the same training setting as ASUKA, performing prompt-guided generation. ASUKA and SD variants use the stable diffusion text-guided inpainting model v1.5.


Dataset Method	MISATO (2k images)				Places 2 (36.5k images)			
	LPIPS↓	FID↓	U-IDS↑	P-IDS↑	LPIPS↓	FID↓	U-IDS↑	P-IDS↑
Co-Mod [59]	0.179	17.421	0.243	0.109	0.267	5.794	0.274	0.096
MAT [25]	0.176	17.261	0.255	0.122	0.202	3.765	0.348	0.195
LaMa [45]	0.155	15.436	0.260	0.135	0.202	6.693	0.247	0.050
MAE-FAR [7]	<b>0.142</b>	13.283	0.282	0.153	<b>0.174</b>	3.559	0.307	0.105
SD [39]	0.168	12.812	0.345	0.211	0.193	1.514	0.375	0.207
SD-text	0.164	12.603	0.337	0.207	0.191	1.506	0.373	0.202
SD-prompt	0.160	12.517	0.331	0.204	0.189	1.477	0.390	0.234
SD-Repaint [32]	0.227	27.861	0.016	0.007	0.251	12.466	0.217	0.045
ASUKA	0.150	<b>11.460</b>	<b>0.368</b>	<b>0.256</b>	0.183	<b>1.230</b>	<b>0.413</b>	<b>0.287</b>

**Competitors.** We consider three variants of SD [39]: i) SD, prompted by "", performing unconditional generation; ii) SD-text, prompted by "background" as we have no captions in inpainting; iii) SD-prompt, prompted by learnable tokens which are trained using the same training pipeline of ASUKA. We also compare with state-of-the-art inpainting algorithms Co-Mod [59], MAT [25], LaMa [45], MAE-FAR [7], and SD-Repaint [32].

## 5.1 Comparison on Benchmarks

**Quantitative comparison** results are reported in Tab. 1. ASUKA, although based on frozen SD, achieves significant and consistent improvements over SD across all the evaluation metrics, and enjoys the state-of-the-art results on FID, U-IDS and P-IDS. Particularly, the U-IDS and P-IDS coincide with human preference [59], and has a potential upper bound of 0.5, indicating the superiority of ASUKA by a large margin. The inconsistency results between LPIPS and FID, U-IDS, P-IDS, partially comes from the frozen SD, where ASUKA achieves consistent improvements but is bounded to the frozen U-Net. In summary, the results validate the contribution of ASUKA that yields context-stable and visual-consistent inpainting through frozen SD. This improvement is both validated in in-distribution dataset Places 2 and out-of-distribution dataset MISATO.

**Qualitative comparison** examples are shown in Fig. 8. **i)** The state-of-the-art inpainting algorithms usually suffer from unnatural generation, for example the unnatural boundaries in the third and fourth rows, and failed inpainting of tower in the third-to-last row. LaMa and MAE-FAR sometimes lead to blurred inpainting results, especially in the scenario of large continuous masks. **ii)** The SD variants usually suffer from the context-stability issue and hallucinate unreasonable objects, in almost all the illustrated images. **iii)** Particularly, for the



**Fig. 9:** The curse of self-attention, causing the MAE falsely estimate the masked region and powerful text-guided diffusion models fail to generation content based on text prompts. ASUKA potential circumvents this issue by using a blank paper image as the input to the MAE to provide correct prior.

background mask shown in the second-to-last row, state-of-the-art inpainting algorithms cannot identify the object precisely, and extend the human in hairs and elbows. SD hallucinates random elements in the masked region. **iv)** In contrast, ASUKA enjoys context-stable and visual-consistent inpainting.

## 5.2 Further Analysis of ASUKA

**Ablation of MAE prior.** We compare our fine-tuned MAE with directly adopting the MAE trained in [7]. To this end, we train ASUKA with the MAE in [7] using the same training strategy and compare the results in Tab. 2a. Results suggest the improvements of fine-tuning MAE, especially on FID and U-IDS. This improvement comes from the better adaptation on the real-world masks.

**Ablation of alignment module.** Our alignment module includes a dimension and a distribution alignment. We validate the effectiveness of each module by comparing with the following variants step by step: i) *linear*: Use linear layer to align feature dimension only; ii) *attn*: Based on i, further use a single self-attention block to align the distribution; iii) *cross  $\mathcal{A}$* : we instead use learnable query and 4 cross-attention layers to learn the MAE prior. ASUKA adopts 4 self-attention blocks. Results are shown in Tab. 2b. The self-attention block shows improved results compared with only align dimension and cross-attention block. Using 4 self-attention blocks improves the capacity of the alignment module.

**Ablation of decoder.** For the decoder, we compare ASUKA with i) *VAE*: the decoder used in SD; ii) *+ cond.*: the decoder conditioned on unmasked image [62]; iii) *+ color*: only use color augmentation to train the decoder; Results are in Tab. 2d, showing the superiority of our augmented decoder.

**User-Study.** To evaluate the user preference on inpainting algorithms, we conduct an user-study. Specifically, we randomly select 40 testing images. We ask the user to select the best inpainting results from the following perspectives respectively: i) Context-stability (C.S.): the generated region should be context-stable with surrounding unmasked region, with a preference of not generating new elements; ii) Visual-consistency (V.C.): the visual consistency between masked and unmasked regions. We collect 100 valid anonymous questionnaire results, and report the average selection ratio among all the inpainting algorithms in Tab. 2c. This result validate the efficacy of ASUKA on alignment with human preference.

**Table 2:** Ablation studies of ASUKA.

(a) Comparison of ASUKA using pre-trained (p.t.) MAE v.s. fine-tuned (f.t.) MAE.

	MAE	LPIPS↓	FID↓	U-IDS↑	P-IDS↑
p.t.	0.151	11.513	0.354	<b>0.258</b>	
f.t.	<b>0.150</b>	<b>11.460</b>	<b>0.368</b>	0.256	

(c) User-study of top-1 ratio among all the inpainting results. Context-stability (C.S.) measures the coherence between masked region and unmasked surroundings, with a preference of not generating new elements; Visual-consistency (V.C.) measures the color consistency.

Model	C.S.(%)	V.C.(%)
Co-Mod [59]	3.98	4.98
MAT [25]	7.40	3.20
LaMa [45]	8.18	8.28
MAE-FAR [7]	4.88	5.60
SD [39]	10.58	5.75
SD-text	7.70	15.83
SD-prompt	16.18	15.78
SD-Repaint [32]	1.60	0.55
ASUKA	<b>39.43</b>	<b>40.05</b>

(b) Ablation of different alignment modules. *Linear* adopts linear layer; *attn* adopts a self-attention layer; *cross x4* adopts 4 cross-attention layers; ASUKA adopts 4 self-attention layers.

Align	LPIPS↓	FID↓	U-IDS↑	P-IDS↑
linear	0.155	11.934	0.361	0.227
attn	0.152	11.613	0.362	0.234
cross x4	0.152	11.762	<b>0.368</b>	0.238
ASUKA	<b>0.150</b>	<b>11.460</b>	<b>0.368</b>	<b>0.256</b>

(d) Comparison of different decoders for SD. VAE [39] is the original decoder used by SD; + cond. [62] is the decoder conditioned on unmasked image; + color uses the color augmentation strategy to perform local harmonization task; Ours further combines latent augmentation strategy to handle the gap between generated latent and real latent.

Decoder	LPIPS↓	FID↓	U-IDS↑	P-IDS↑
VAE	0.156	11.949	0.343	0.208
+ cond.	0.151	11.634	0.361	0.231
+ color	0.152	11.603	0.357	0.229
Ours	<b>0.150</b>	<b>11.460</b>	<b>0.368</b>	<b>0.256</b>

**Limitation: The "curse" of self-attention.** The primary limitation of ASUKA stems from the inefficacy of the MAE prior, mainly due to issues within the self-attention module. Specifically, as shown in Fig. 9, the presence of multiple similar objects in an image may lead the MAE to incorrectly predict a similar object in the masked region, conflicting with the goal of object removal. Notably, this curse of self-attention significantly impacts diffusion-based generative models. It results in the inability to accurately follow "blank paper" text prompts, even when employing a substantial classifier-free guidance scale of 9. This issue is not unique to SD but is also a common problem in other advanced text-guided diffusion models, such as OpenAI’s DALL-E 2 [38] and Adobe’s FileFly [1]. Nevertheless, ASUKA has the potential to circumvent this issue by modifying the MAE prior, for instance, by instead using a blank paper image as the input to MAE prior. A more comprehensive solution would involve extra control on self-attention layers in diffusion models, which we leave as future work.

## 6 Conclusion

In this paper, we proposed Aligned Stable inpainting with Unknown Areas prior (ASUKA) to achieve context-stable and visual-consistent inpainting via frozen

Stable Diffusion inpainting model (SD). To achieve context-stability, we adopt a reconstruction-based masked auto-encoder (MAE) as the context-stable prior for masked region purely from unmasked region. Then we align the context-stable prior to frozen SD with the proposed alignment module. To achieve visual-consistency, we resolve the mask-unmask color inconsistency of SD in the decoding process. We train an unmask-region conditioned KL-VAE decoder to perform local harmonization during the decoding process. To validate the efficacy of inpainting algorithms in different image domains and mask types, we introduce an evaluation dataset, named as MISATO, from existing datasets. ASUKA enjoys context-stable and visual-consistent inpainting results and superior than state-of-the-art inpainting algorithms and SD-based variants.

## References

1. Adobe firefly. <https://www.adobe.com/sensei/generative-ai/firefly.html> 14
2. Openai’s consistency decoder. <https://github.com/openai/consistencydecoder> 5
3. Bao, H., Dong, L., Wei, F.: Beit: Bert pre-training of image transformers. arXiv preprint arXiv:2106.08254 (2021) 5
4. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.* **28**(3), 24 (2009) 4
5. Bertalmio, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. In: Proceedings of the 27th annual conference on Computer graphics and interactive techniques. pp. 417–424 (2000) 2, 4
6. Bertalmio, M., Vese, L., Sapiro, G., Osher, S.: Simultaneous structure and texture image inpainting. *IEEE transactions on image processing* **12**(8), 882–889 (2003) 4
7. Cao, C., Dong, Q., Fu, Y.: Learning prior feature and attention enhanced image inpainting. In: European Conference on Computer Vision. pp. 306–322. Springer (2022) 1, 4, 5, 7, 11, 12, 13, 14, 19
8. Cao, C., Fu, Y.: Learning a sketch tensor space for image inpainting of man-made scenes. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14509–14518 (2021) 2
9. Chan, T.F., Shen, J.: Nontexture inpainting by curvature-driven diffusions. *Journal of visual communication and image representation* **12**(4), 436–449 (2001) 4
10. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3D: Learning from RGB-D data in indoor environments. *International Conference on 3D Vision (3DV)* (2017) 4, 10
11. Chen, X., Ding, M., Wang, X., Xin, Y., Mo, S., Wang, Y., Han, S., Luo, P., Zeng, G., Wang, J.: Context autoencoder for self-supervised representation learning. *International Journal of Computer Vision* pp. 1–16 (2023) 5
12. Chu, T., Chen, J., Sun, J., Lian, S., Wang, Z., Zuo, Z., Zhao, L., Xing, W., Lu, D.: Rethinking fast fourier convolution in image inpainting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 23195–23205 (2023) 4
13. Criminisi, A., Pérez, P., Toyama, K.: Object removal by exemplar-based inpainting. 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings. **2**, II–II (2003) 2

14. Criminisi, A., Pérez, P., Toyama, K.: Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on image processing* **13**(9), 1200–1212 (2004) [4](#)
15. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12873–12883 (2021) [2](#)
16. Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion. In: *The Eleventh International Conference on Learning Representations* (2022) [4](#), [5](#)
17. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in neural information processing systems* **27** (2014) [2](#), [4](#)
18. Hays, J., Efros, A.A.: Scene completion using millions of photographs. *ACM Transactions on Graphics (ToG)* **26**(3), 4–es (2007) [2](#)
19. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16000–16009 (2022) [1](#), [3](#), [5](#), [11](#)
20. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017) [11](#)
21. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020) [2](#)
22. Hu, E.J., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. In: *International Conference on Learning Representations* (2021) [5](#)
23. Ko, K., Kim, C.S.: Continuously masked transformer for image inpainting. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 13169–13178 (2023) [4](#)
24. Levin, A., Zomet, A., Weiss, Y.: Learning how to inpaint from global image statistics. *Proceedings Ninth IEEE International Conference on Computer Vision* pp. 305–312 vol.1 (2003) [2](#)
25. Li, W., Lin, Z., Zhou, K., Qi, L., Wang, Y., Jia, J.: Mat: Mask-aware transformer for large hole image inpainting. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022) [1](#), [4](#), [11](#), [12](#), [14](#), [19](#)
26. Li, Z., Snavely, N.: Megadepth: Learning single-view depth prediction from internet photos. In: *Computer Vision and Pattern Recognition (CVPR)* (2018) [1](#), [4](#), [10](#)
27. Liao, L., Xiao, J., Wang, Z., Lin, C.W., Satoh, S.: Guidance and evaluation: Semantic-aware image inpainting for mixed scenes. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVII* 16. pp. 683–700. Springer (2020) [2](#)
28. Lin, C.H., Cheng, Y.C., Lee, H.Y., Tulyakov, S., Yang, M.H.: InfinityGAN: Towards infinite-pixel image synthesis. In: *International Conference on Learning Representations* (2022), <https://openreview.net/forum?id=ufGMqIM0a4b> [4](#), [10](#)
29. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *European conference on computer vision*. pp. 740–755. Springer (2014) [4](#), [7](#), [10](#)
30. Liu, G., Reda, F.A., Shih, K.J., Wang, T.C., Tao, A., Catanzaro, B.: Image inpainting for irregular holes using partial convolutions. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 85–100 (2018) [4](#)



31. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (2018) [11](#)
32. Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Van Gool, L.: Repaint: Inpainting using denoising diffusion probabilistic models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11461–11471 (2022) [11](#), [12](#), [14](#)
33. Luo, F., Xiang, J., Zhang, J., Han, X., Yang, W.: Image super-resolution via latent diffusion: A sampling-space mixture of experts and frequency-augmented decoder approach. arXiv preprint arXiv:2310.12004 (2023) [5](#)
34. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: Sdedit: Guided image synthesis and editing with stochastic differential equations. In: International Conference on Learning Representations (2021) [4](#)
35. Nazeri, K., Ng, E., Joseph, T., Qureshi, F., Ebrahimi, M.: Edgeconnect: Structure guided image inpainting using edge prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (2019) [2](#), [4](#)
36. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2536–2544 (2016) [2](#), [4](#)
37. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021) [10](#)
38. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents [14](#)
39. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10684–10695 (June 2022) [1](#), [2](#), [3](#), [4](#), [5](#), [8](#), [11](#), [12](#), [14](#), [19](#)
40. Roth, S., Black, M.J.: Fields of experts: a framework for learning image priors. 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) [2](#), 860–867 vol. 2 (2005) [2](#)
41. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH 2022 Conference Proceedings. pp. 1–10 (2022) [4](#)
42. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (ICLR) (2021) [9](#)
43. Song, Y., Zhang, Z., Lin, Z., Cohen, S., Price, B., Zhang, J., Kim, S.Y., Aliaga, D.: Objectstitch: Object compositing with diffusion model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18310–18319 (2023) [5](#)
44. Steinbach, M., Karypis, G., Kumar, V.: A comparison of document clustering techniques (2000) [10](#)
45. Suvorov, R., Logacheva, E., Mashikhin, A., Remizova, A., Ashukha, A., Silvestrov, A., Kong, N., Goka, H., Park, K., Lempitsky, V.: Resolution-robust large mask inpainting with fourier convolutions. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 2149–2159 (2022) [1](#), [2](#), [4](#), [7](#), [11](#), [12](#), [14](#), [19](#)
46. Wan, Z., Zhang, J., Chen, D., Liao, J.: High-fidelity pluralistic image completion with transformers. arXiv preprint arXiv:2103.14031 (2021) [2](#)

47. Wei, C., Fan, H., Xie, S., Wu, C.Y., Yuille, A., Feichtenhofer, C.: Masked feature prediction for self-supervised visual pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14668–14678 (2022) [5](#)
48. Wei, L., Xie, L., Zhou, W., Li, H., Tian, Q.: Mvp: Multimodality-guided visual pre-training. In: European Conference on Computer Vision. pp. 337–353. Springer (2022) [5](#)
49. Xie, Z., Zhang, Z., Cao, Y., Lin, Y., Bao, J., Yao, Z., Dai, Q., Hu, H.: Simmim: A simple framework for masked image modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9653–9663 (2022) [5](#)
50. Xu, X., Navasardyan, S., Tadevosyan, V., Sargsyan, A., Mu, Y., Shi, H.: Image completion with heterogeneously filtered spectral hints. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 4591–4601 (2023) [4](#)
51. Yi, Z., Tang, Q., Azizi, S., Jang, D., Xu, Z.: Contextual residual aggregation for ultra high-resolution image inpainting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7508–7517 (2020) [4](#)
52. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Generative image inpainting with contextual attention. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5505–5514 (2018) [4](#)
53. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Free-form image inpainting with gated convolution. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4471–4480 (2019) [4](#)
54. Zeng, Y., Fu, J., Chao, H., Guo, B.: Aggregated contextual transformations for high-resolution image inpainting. *IEEE Transactions on Visualization and Computer Graphics* (2022) [4](#)
55. Zeng, Y., Lin, Z., Yang, J., Zhang, J., Shechtman, E., Lu, H.: High-resolution image inpainting with iterative confidence feedback and guided upsampling. In: European Conference on Computer Vision. pp. 1–17. Springer (2020) [4](#)
56. Zhang, D., Liang, Z., Yang, G., Li, Q., Li, L., Sun, X.: A robust forgery detection algorithm for object removal by exemplar-based image inpainting. *Multimedia Tools and Applications* **77**, 11823–11842 (2018) [4](#)
57. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models (2023) [5](#)
58. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018) [11](#)
59. Zhao, S., Cui, J., Sheng, Y., Dong, Y., Liang, X., Eric, I., Chang, C., Xu, Y.: Large scale image completion via co-modulated generative adversarial networks. In: International Conference on Learning Representations (2020) [1](#), [2](#), [4](#), [7](#), [11](#), [12](#), [14](#), [19](#)
60. Zheng, H., Lin, Z., Lu, J., Cohen, S., Shechtman, E., Barnes, C., Zhang, J., Xu, N., Amirghodsi, S., Luo, J.: Image inpainting with cascaded modulation gan and object-aware training. In: European Conference on Computer Vision. pp. 277–296. Springer (2022) [4](#)
61. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence* **40**(6), 1452–1464 (2017) [4](#), [10](#)
62. Zhu, Z., Feng, X., Chen, D., Bao, J., Wang, L., Chen, Y., Yuan, L., Hua, G.: Designing a better asymmetric vqgan for stablediffusion. arXiv preprint arXiv:2306.04632 (2023) [5](#), [8](#), [9](#), [10](#), [11](#), [13](#), [14](#)

## 7 More Qualitative Examples

Here we provide more qualitative examples on MISATO in Figure 10, Figure 11, Figure 12, Figure 13, and Figure 14. We compare ASUKA with Co-Mod [59], MAT [25], LaMa [45], MAE-FAR [7], and SD [39]. SD performs unconditional generation. SD-text utilizes text prompt of “background”. SD-prompt utilizes trained prompt for inpainting task using the same training setting of ASUKA.

### 7.1 Potential negative impact

As an image editing tool, our proposed ASUKA will generate images based on user intentions for masking specific parts of the image, potentially resulting in unrealistic renderings and posing a risk of misuse. To mitigate this, we plan to prominently watermark the images produced by ASUKA, clearly indicating their synthetic nature.




Fig. 10: More qualitative comparison on MISATO.



Fig. 11: More qualitative comparison on MISATO.



Fig. 12: More qualitative comparison on MISATO.



Fig. 13: More qualitative comparison on MISATO.




Fig. 14: More qualitative comparison on MISATO.