

Harmonizing Attention: Training-free Texture-aware Geometry Transfer

Eito Ikuta¹ Yohan Lee¹ Akihiro Iohara¹ Yu Saito¹ Toshiyuki Tanaka²

¹DATAGRID Inc. ²Graduate School of Informatics, Kyoto University



Figure 1. Examples of texture-aware geometry transfer. Our method seamlessly transfers geometric features (holes, cracks, droplets, etc.) from source images (red rectangles in the insets of the top row) to target images with different surface textures (main images in the top row).

Abstract

Creating images where surface patterns of one object - such as cracks, holes, or grooves - are precisely transferred onto objects made of different materials remains a challenging task in computer graphics. For example, recreating the exact pattern of wood grain cracks on a metallic surface, while maintaining the realistic metallic texture, requires sophisticated technical solutions. In this study, we introduce Harmonizing Attention, a new method that can automatically extract these surface patterns from photographs and recreate them with different materials, while preserving natural-looking textures. Our approach achieves this through a novel attention mechanism that can process multiple reference images simultaneously, without requiring additional training. This makes the method both practical and efficient for real-world applications, opening up new possibilities in augmented reality, image editing, and beyond.

1. Introduction

In computer vision, image harmonization is the task of seamlessly integrating a foreground object from one image into the background of another to produce a cohesive composite. A significant challenge is achieving visual harmony by adjusting the foreground’s appearance to align with the background. Traditional methods [1–10] have primarily addressed image harmonization by focusing on color and illumination adjustments.

Recent advances in deep learning and diffusion models [11–21] have enabled sophisticated image composition and harmonization, as exemplified by ObjectStitch [20] and TF-ICON [21]. These methods primarily focus on style transfer, encompassing a broad range of attributes such as color schemes, brushstrokes, textures, and patterns that collectively determine an image’s visual appearance. However, the selective transfer of geometrical features—which we term “geometry” in this paper—such as holes, cracks, droplets, and dents from one material to another, independently of material-specific surface texture (e.g., wood grain or fabric weave), remains a complex challenge. To the best

of our knowledge, existing techniques have not successfully addressed this issue. Such geometry transfer is crucial for creating realistic and aesthetically pleasing composite images, particularly in applications requiring the integration of complex textures from different materials, like transferring wood hole geometry onto metal surfaces. Traditional and recent deep-learning-based image harmonization methods struggle with geometry transfer due to their focus on color, illumination, and style adjustments of the object itself. This limitation highlights the need for a fundamentally different approach that can effectively decouple geometry from material-specific textures in image harmonization tasks.

To address this challenge, we propose Harmonizing Attention, a novel training-free diffusion-model-based approach for texture-aware geometry transfer. The key aspect of this method is a custom attention mechanism comprising Texture-aligning Attention and Geometry-preserving Attention, enabling reference to both source geometry and target texture information. By replacing the self-attention layers in the diffusion model with these custom attention layers during inversion and generation, we successfully incorporate geometry and texture from different images, as demonstrated in Fig. 1. Notably, since the Harmonizing Attention framework leverages pretrained diffusion models to perform geometry transfer, it requires neither model tuning nor prompt exploration for each dataset.

Our contributions are summarized as follows:

1. We introduce texture-aware geometry transfer, a novel approach for integrating geometry details from different materials into cohesive composite images.
2. We develop Harmonizing Attention, a training-free mechanism enabling models to query information from multiple reference images within self-attention blocks, effectively capturing and transferring material-independent geometry while preserving material-specific textures.
3. Our experiments demonstrate that Harmonizing Attention produces more harmonious and realistic composite images compared with existing techniques.

2. Related Work

Image Harmonization. Conventionally, image harmonization [1–10] has focused on color-to-color transformations to match visual appearances. These methods can be further divided into non-linear transformations [7, 9] and linear transformations [3, 6, 10]. Recently, deep-learning-enabled image harmonization emerged [11–19] and more recently diffusion-model-based harmonization techniques [20, 21] have been developed. One notable example is TF-ICON [21], which equips attention-based text-

to-image diffusion models for image harmonization, enabling cross-domain image-guided composition.

Paintely Image Harmonization. In the context of image harmonization, a similar task called painterly image harmonization [22–27] has been studied, and it might be another promising candidate technique for geometry transfer. The painterly image harmonization technique integrates a photographic foreground into an artistic background, resulting in a visually coherent painting. PHDiffusion [28] represents a significant pivot by adapting the GAN framework (PHDNet [22]) to a diffusion model with an adaptive encoder. Furthermore, TF-GPH [29], which utilizes an image-wise attention-sharing mechanism for general painterly harmonization, is proposed for training-free general painterly harmonization and provides flexible options for attention-based image editing methods.

3. Methods

Our objective is to synthesize an image I^{out} that seamlessly integrates surface geometry information from a source image I^{src} into a target background image I^{tar} . This integration is guided by a 0-1-valued foreground mask M^{src} , preserving the geometry of $M^{\text{src}} \odot I^{\text{src}}$ while maintaining textural continuity with I^{tar} , independently of the source’s original texture, where \odot denotes the Hadamard (elementwise) product. To achieve this, we propose a novel framework called Harmonizing Attention. As illustrated in Fig. 2, this framework leverages Stable Diffusion (SD) [30] and encompasses both inversion and generation processes. The key aspect of our approach is modifying self-attention computation during both inversion and generation to query additional information from I^{src} and I^{tar} , enabling a more coherent and context-aware transfer process. Furthermore, to enhance textural continuity within the target region, we utilize the SD inpainting model rather than the text-to-image model.

Editing. The initial stage of our framework involves preparation of a geometry image I^{geo} , which serves as a cornerstone for the subsequent image generation process. A geometry image is simply an image patch of a source image cropped by a 0-1 source mask M^{src} , with affine transformations \mathcal{T} (shifting, scaling, and rotation) according to the user’s needs, as well as the color adjustment described below (see the uppermost panel in Fig. 2). The ultimate objective of our framework is to seamlessly synthesize this geometry image with the target image.

The color adjustment of the geometry image plays a crucial role in that it facilitates the inversion process to work properly. By aligning the color profile of the source region with that of the target image, we expect that it will

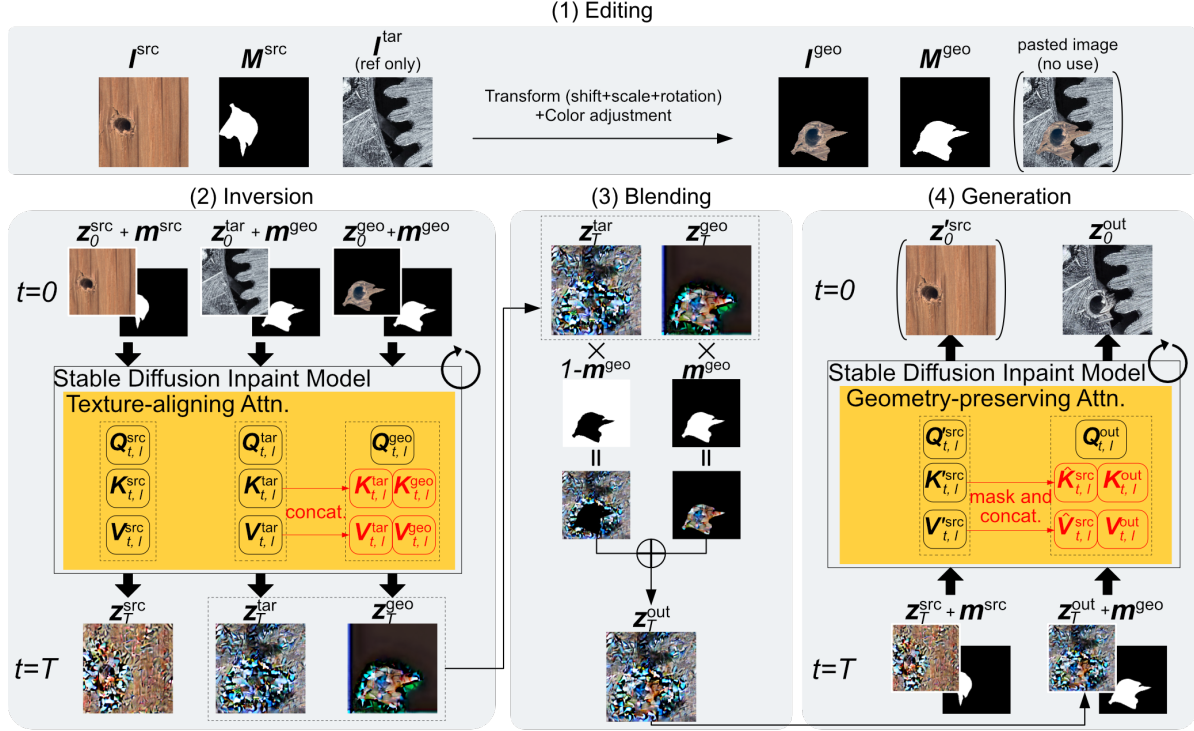


Figure 2. The overview of Harmonizing Attention framework. The process consists of four main stages: (1) Editing: The masked region of the source image $M^{\text{src}} \odot I^{\text{src}}$ is transformed, positioned, and color-adjusted to create a geometry image I^{geo} and a corresponding mask M^{geo} . (2) Inversion: The process commences with the compression of the source, target, and geometry images via VAE, yielding clear latents z_0^{src} , z_0^{tar} , and z_0^{geo} respectively. Subsequently, these latents, along with the source and geometry masks resized to the latent space, are fed into an SD inpainting model. This step results in the generation of inverted noisy latents z_T^{src} , z_T^{tar} , and z_T^{geo} . Here we replace the self-attention computation for the geometry image with Texture-aligning Attention, which incorporates information from the target image to align the geometry image with the target domain. (3) Blending: The latents of the target and geometry images are combined using the geometry mask. (4) Generation: The blended latents undergo denoising, with modified self-attention named Geometry-preserving Attention, to reference the source image, preserving the geometry while ensuring seamless integration. For the sake of enhanced clarity and readability, all latents are presented in pixel space, rather than in the VAE latent space.

be easier to bring the geometry image closer to the visual domain of the target through the inversion procedure. Because the geometry information is expressed through localized color differences rather than absolute values, we employ a simple uniform color shift, although more sophisticated color adjustment methods would also be feasible. Specifically, we compute this shift by calculating the difference between two color averages: (1) $c^{\text{src}} \in \mathbb{R}^3$, which is the mean color of the source image I^{src} within the region formed by subtracting the source mask M^{src} from its dilated version (i.e., the boundary of the complementary mask $M_{\text{comp}}^{\text{src}}$ in mathematical morphology, where the complementary mask M_{comp} of a mask M is defined by replacing 0 with 1 and vice versa in M); (2) $c^{\text{tar}} \in \mathbb{R}^3$, which is the mean color of the target image I^{tar} within an analogous region, defined by subtracting the transplantation-area mask $M^{\text{geo}} := \mathcal{T}(M^{\text{src}})$ from its dilated version, where the transplantation-area mask M^{geo} defines the transplantation region in the target-image coordinate. This process is

represented by the following equation:

$$p_{x,y}^{\text{geo}} = \tilde{p}_{x,y}^{\text{geo}} + aM_{x,y}^{\text{geo}}(c^{\text{tar}} - c^{\text{src}}), \quad (1)$$

where $\tilde{p}_{x,y}^{\text{geo}}, p_{x,y}^{\text{geo}} \in \mathbb{R}^3$ represent the pixel values at coordinates (x, y) in the image before and after color adjustment, respectively. $M_{x,y}^{\text{geo}} \in \{0, 1\}$, value of M^{geo} at coordinates (x, y) , becomes 1 only if the pixel (x, y) in the target image is within the transplantation region. Equation (1) then shows that the color adjustment is done only within the transplantation region. The scalar parameter $a \in [0, 1]$ controls the strength of the color adjustment. It is anticipated that employing a small value of a would result in inadequate integration of the transplantation region with the target image, while a large value of a is expected to compromise the integrity of the source geometry.

Inversion. The geometry image I^{geo} inherently retains substantial textural information from the source image. Consequently, a direct pasting of the geometry image onto

the target image I^{tar} (“pasted image” in Fig. 2) may result in textural discontinuities at the boundaries of M^{geo} in the final output. To address this issue, our objective here is to derive a latent representation that effectively translates the geometry image into the visual domain of the target image through inversion. This domain-shifted latent representation facilitates seamless blending within the latent space, thereby enabling a more naturalistic geometry transplantation. It should be noted that since we leverage SD, the diffusion process operates in the latent space encoded by the variational autoencoder (VAE) instead of the pixel space.

To achieve this, we simultaneously invert the geometry and target images by Denoising Diffusion Implicit Models (DDIM) inversion [31], with replacing self-attention for the geometry image with a novel attention computation which we call Texture-aligning Attention, which additionally incorporates target image information. Let z denote the latent representations in the VAE-encoded space, where specific instances will be distinguished by appropriate subscripts in the subsequent discussion. The standard self-attention computation is formulated as:

$$A(z) = \text{Softmax} \left(\frac{QK^\top}{\sqrt{d}} \right) V, \quad (2)$$

where $A(\cdot)$ denotes the functional representation of self-attention computation, where d and \top denote the dimension of the latent space and the transposition of a matrix, respectively, and where Q , K , and V represent the query, key, and value, respectively, all computed from the latents z corresponding to the same image. Note that the diffusion step t and the layer index l are described in Fig. 2, but omitted here for readability. To align the source image domain with the target image domain, we replace the self-attention computation in the inversion of the geometry image with the following formula:

$$A_{\text{TA}}(z^{\text{geo}}; z^{\text{tar}}) = \text{Softmax} \left(\frac{Q^{\text{geo}} \begin{bmatrix} K^{\text{geo}} \\ K^{\text{tar}} \end{bmatrix}^\top}{\sqrt{d}} \right) \begin{bmatrix} V^{\text{geo}} \\ V^{\text{tar}} \end{bmatrix}, \quad (3)$$

where $A_{\text{TA}}(\cdot)$ denotes Texture-aligning Attention, and where the superscripts geo and tar denote whether the respective Q , K , and V are computed from the geometry image feature map z_{geo} or the target image feature map z_{tar} . As previously elucidated, the geometry image preserves the relative color variations associated with surface geometry through the color adjustment, while simultaneously approximating the visual characteristics of the target image. Consequently, by retaining both the conventional self-attention key (K^{geo}) and value (V^{geo}) alongside their counterparts derived from the target image (K^{tar} and V^{tar}), we anticipate a synergistic effect. This dual-source attention mechanism is expected to efficiently facilitate both

geometry preservation and material transformation during the inversion process. A noisy latent representation z_T^{geo} of the geometry image I^{geo} is then obtained as a result of T iterations of the inversion step using this custom attention.

Inversion of the source image I^{src} and the target image I^{tar} is performed independently, employing the standard self-attention mechanism without any modifications. The resulting noisy latent representations z_T^{src} and z_T^{tar} are then stored for subsequent use in the generation stage and in the blending stage, respectively.

Blending. Utilizing the latent representations obtained in the inversion stage, in the blending stage (see the lower-middle panel in Fig. 2) we perform a latent-space blending between the geometry and target images. Let m^{geo} be the target mask M^{geo} resized to the latent space. The blended latent representation z_T^{out} to be used for generating a synthesized image is computed from the latents z_T^{geo} and z_T^{tar} for the geometry and target images at diffusion step T as follows:

$$z_T^{\text{out}} = z_T^{\text{geo}} \odot m^{\text{geo}} + z_T^{\text{tar}} \odot m_{\text{comp}}^{\text{geo}}, \quad (4)$$

where $m_{\text{comp}}^{\text{geo}}$ denotes the complementary mask of m^{geo} , defined by replacing 0 with 1 and vice versa in m^{geo} .

Generation. In the generation stage (see the lower-right panel in Fig. 2), we obtain the final output image by denoising the blended latent code z_T^{out} synthesized in the blending stage. Although we initialize the generation process with the latents that explicitly incorporate a portion of the source image, it turns out that applying standard denoising procedures may potentially result in the undesirable removal of the transplanted content along with the noise, leading to the loss of the geometry.

To address this challenge, we introduce a novel attention computation named Geometry-preserving Attention. Specifically, we concurrently perform denoising on the source image latent representation z_T^{src} obtained during the inversion stage and replace the self-attention computation for the blended latent representation z_T^{out} with Geometry-preserving Attention by augmenting the key and value components with information derived from the source image. The process is formally expressed by the following equation:

$$A_{\text{GP}}(z^{\text{geo}}; z^{\text{tar}}) = \text{Softmax} \left(\frac{Q^{\text{geo}} \begin{bmatrix} K^{\text{out}} \\ \hat{K}^{\text{src}} \end{bmatrix}^\top}{\sqrt{d}} \right) \begin{bmatrix} V^{\text{out}} \\ \hat{V}^{\text{src}} \end{bmatrix}, \quad (5)$$

$$\hat{K}^{\text{src}} = \{K_i^{\text{src}} \mid m_i^{\text{src}} \neq 0, i \in \{1, \dots, H \times W\}\}, \quad (6)$$

$$\hat{V}^{\text{src}} = \{V_i^{\text{src}} \mid m_i^{\text{src}} \neq 0, i \in \{1, \dots, H \times W\}\}. \quad (7)$$




Figure 3. Qualitative comparison of image generation results for geometry transfer. The methods to be compared are set to maximize the quality of the generated images, and the highest quality results are selected and included. Each column shows the output result by the compared method for harmonizing and the output result by our method.

Here, A_{GP} represents the functional form of the Geometry-preserving Attention. The superscripts out and src denote that the corresponding variables are associated with the z^{out} or z^{src} , respectively. It is important to note that K_i^{src} and V_i^{src} are used instead of K_i^{src} and V_i^{src} to distinguish them from the keys and values used during the inversion process. H and W denote the height and width of the latent feature map z^{src} , $m^{src} \in \{0, 1\}^{H \times W}$ represents the source mask M^{src} resized to the latent space, and $m_i^{src} \in \{0, 1\}$ represents i -th element of it.

4. Experiments

4.1. Setup

We use the publicly available SD inpainting model checkpoint on HuggingFace¹ as the backbone network. For our experiments, both source image I^{src} and target image I^{tar} are cropped to a uniform size of 512×512 pixels. When performing color adjustment as expressed in Eq. (1), we set the scalar parameter a to 0.5. We set the number of diffusion steps T to 25 and employ the DDIM sampler for both the inversion and generation processes, with no prompt input (i.e., an empty string is used) in either phase.

During the inversion phase, we leverage the insights from Garibi *et al.* [32], which demonstrates that iterative inversion sampling at each diffusion step enhances inversion accuracy. Specifically, we perform 5 sampling iterations per

¹<https://huggingface.co/runwayml/stable-diffusion-inpainting> (no longer available)

diffusion step in the inversion procedure.

4.2. Datasets and Metrics

We use images from MVTec AD [33,34] and Pixabay² in our experiments. For the quantitative evaluation, we employ the following three metrics: LPIPS [35], CLIP [36], and DISTS [37], since metrics specifically tailored for geometry transfer do not currently exist. These established metrics allow us to assess our method’s performance in terms of image harmonization, and compare ours with existing methods. Each metric is evaluated against both the composite image background (denoted $LPIPS_{(bg)}$, $CLIP_{(bg)}$, and $DISTS_{(bg)}$) and the local region of the stretched and harmonized image (denoted $LPIPS_{(fg)}$, $CLIP_{(fg)}$, and $DISTS_{(fg)}$). Here, (bg) and (fg) refer to background and foreground, respectively. It is noteworthy that LPIPS and CLIP have previously been used as evaluation metrics in the existing literature [21, 29]. A total of 150 generated images are used for the quantitative performance evaluation, consisting of 10 composite foreground images and 15 background images, resulting in 150 combinations. Scores are calculated for each method, and the averages of the scores are presented.

4.3. Baselines

For comparison, we select four diffusion-model-based methods, including two image harmonization methods

²<https://pixabay.com/>

PHDiffusion [28] and TF-GPH [29], one image composition method TF-ICON [21], and one image editing method Paint by Example [38]. Since our Harmonizing Attention does not require prompts, we manually set suitable prompts for generating each sample with TF-ICON. For all methods, we use the default or recommended hyperparameters.

4.4. Qualitative Comparison

We present a qualitative comparison of our proposed method, Harmonizing Attention, with four existing techniques: Paint by Example [38], TF-ICON [21], PHDiffusion [28], and TF-GPH [29]. Paint by Example and TF-ICON struggled with consistent texturing and maintaining original geometries. PHDiffusion showed better texture preservation but struggled with accurate geometry transfer. TF-GPH achieved a balance between texture and geometry alignment compared with previous methods, but still showed subtle texture inconsistencies and degradation of geometric properties upon closer inspection.

In contrast, our proposed method excelled in several critical aspects of image harmonization. It preserved the geometry and structural details of foreground objects with high fidelity. Simultaneously, it successfully adapted the texture and material properties to seamlessly match the background. This dual capability resulted in the most natural and realistic geometry integration across various scenarios, from textured surfaces such as wood to complex materials like metal parts.

User Study To further compare the apparent quality of each technique, we conducted a user study with 105 participants from CrowdWorks³. Participants rated 15 images, shown in Fig. 3, on a scale from 1 to 5 (1 being the worst and 5 being the best) on the basis of the following criteria: background preservation (referred to as Quality Of Background, QOB), foreground preservation (referred to as Quality Of Foreground, QOF), and seamless composition (referred to as Quality Of Composition, QOC). Each qualitative rating score was calculated by applying weights to each response choice (5 for the best, 4, 3, 2, and 1 for the worst), and averaging all the responses. For more details, please refer to supplementary materials.

The results are listed in Table 1. Our method consistently outperformed existing methods across all evaluated metrics. The consistent superiority across these different metrics demonstrates the ability of our method to produce high-quality harmonized images with improved boundary definition, improved feature representation, and better overall content coherence compared with the existing techniques.

³<https://crowdworks.jp>

Table 1. Qualitative evaluation results from user study. QOB, QOF, and QOC represent scores for background preservation, foreground preservation, and seamless composition, respectively. Higher scores indicate better performance across all criteria.

Method	QOB \uparrow	QOF \uparrow	QOC \uparrow
Paint by Example	3.11	1.69	2.44
TF-ICON	2.41	2.40	2.38
PHDiffusion	3.10	2.64	3.26
TF-GPH	3.12	2.79	3.40
Ours	3.44	2.96	3.56

4.5. Quantitative Comparison

We conducted a comprehensive quantitative evaluation to assess our method’s performance against state-of-the-art approaches for geometry composition in target background images. Table 2 presents results across multiple metrics, including LPIPS, CLIP, and DISTs, for both background and foreground components. Our method achieved the lowest LPIPS_(bg) score of 0.266 and DISTs_(bg) score of 0.179, indicating superior preservation of background structural integrity. Furthermore, ours excelled in foreground semantic consistency, evidenced by the highest CLIP_(fg) score of 91.352. While PHDiffusion and TF-GPH marginally outperformed ours on CLIP_(bg) and LPIPS_(fg) respectively, our method maintained competitive scores across all metrics. This balanced performance, particularly our excellence in semantic preservation coupled with strong perceptual and structural similarity scores, suggests the effectiveness of our approach in achieving harmonious image composites. These results collectively demonstrate our method’s advancement of the state-of-the-art in image harmonization across multiple dimensions of visual quality.

4.6. Ablation Studies

4.6.1 Effectiveness of Color Shift

We assessed our color shift method’s efficacy by varying the parameter a in Eq. (1) and comparing it with histogram matching, a method that fully adjusts the source image’s color to match the target. The experiment was conducted under four conditions: (i) no color adjustment (i.e., (1) with $a = 0.0$), (ii) color adjustment via the proposed color shift method (1) with $a = 0.5$, (iii) color adjustment via the proposed color shift method (1) with $a = 1.0$, and (iv) color adjustment via histogram matching. It should be noted that, except for color adjustment, all other settings were consistent across all three conditions.

The qualitative results under these conditions are shown in Fig. 4. For reference, we also include pasted images, which are geometry images obtained under each condition directly pasted onto the target image. Comparing (i) and (ii), (iii) in Fig. 4, we can see that without color adjust-

Table 2. Quantitative evaluation results for geometry composition in a given target background image. Arrows next to each score indicate score interpretation: \downarrow lower is better, \uparrow higher is better. The minimum value for LPIPS and DISTS is 0, and the maximum value for CLIP is 100. A total of 150 images were used for the evaluation.

Method	LPIPS _(bg) \downarrow	LPIPS _(fg) \downarrow	CLIP _(bg) \uparrow	CLIP _(fg) \uparrow	DISTS _(bg) \downarrow	DISTS _(fg) \downarrow
Paint By Example [38]	0.424	0.273	70.694	84.348	0.256	0.241
TF-ICON [21]	0.434	0.392	69.470	81.422	0.309	0.315
PHDiffusion [28]	0.408	0.255	73.624	84.879	0.276	0.196
TF-GPH [29]	0.324	0.252	71.234	86.936	0.223	0.237
Ours	0.266	0.255	68.180	91.352	0.179	0.250

ment, the texture remained unchanged, while with color shift, the geometry was reconstructed with a texture that was blended into the target image. Also, comparing (ii), (iii) and (iv), we find that when histogram matching was used, the relative color changes within the source image area to be transplanted were not preserved, resulting in loss of three-dimensionality. Finally, comparing (ii) and (iii), a stronger color shift tends to improve texture blending, but tends to reduce the three-dimensionality of the object, making it difficult to maintain the original geometry. These results suggest that a modest shift in the color of the source image to that of the target image is effective in transferring geometry.

4.6.2 Effectiveness of Texture-aligning Attention

We examined the effectiveness of Texture-aligning Attention. As previously explained, the primary objective of the inversion phase is to align the geometry image, created from the source image, with the visual domain of the target image. Intuitively, one might expect that using only the keys K^{tar} and values V^{tar} computed from the target image in Eq. (3), while eliminating K^{geo} and V^{geo} derived from the geometry image, would yield results more harmoniously integrated with the target image.

To evaluate the impact of each K , V on this alignment, we conducted an ablation study comparing three inversion configurations: (i) employing only target-image-derived attention (i.e., eliminating K^{geo} and V^{geo}), and (ii) utilizing only geometry-image-derived attention (i.e., removing K^{tar} and V^{tar}). (iii) using both geometry-image-derived and target-image-derived attention components. It is noted that we used the same settings for editing, blending, and generation.

Figure 5 presents a comparative analysis of generation results under the three conditions. Contrary to the initial intuition, the samples for condition (i) demonstrate that using only K^{tar} and V^{tar} does not lead to better integration with the target image. Instead, as shown in condition (iii), using geometry-derived attention components yields results that are more harmoniously blended with the target image. Even when comparing conditions with geometry-derived attention, the presence of target-derived attention tends to accentuate the source geometry. A comparative




Figure 4. Qualitative comparison of images generated with different types of color adjustment. Here are examples of two paired images. The first columns are the background and foreground images to be input to our method. For each pair, the top line shows the pasted image with the foreground target attached to the background, and the bottom line shows the output from our method for each pasted image. Each line is the result of (i) using the original color image, (ii) and (iii) are color-shifted images of the target image using the color shift parameter, and (iv) using the color matched image with histogram matching method.

analysis of columns (ii) and (iii) in Fig. 5 revealed notable distinctions. In the first row, the absence of target-derived attention in column (ii) results in the disappearance of geometry. Conversely, in the second row, column (ii) exhibits partial omission of geometry. These observations suggest the effectiveness of employing both geometry-derived and target-derived attention in geometry transfer.

4.6.3 Effectiveness of Geometry-preserving Attention

We investigated the significance of Geometry-preserving Attention. Specifically, we compared generation results under three conditions: (i) employing only self-attention (i.e., eliminating K^{src} and V^{src}), and (ii) utilizing only source-image-derived attention (i.e., removing K^{out} and V^{out}). (iii) using both self-attention and source-image-derived attention components (i.e., using Geometry-preserving Atten-




Figure 5. Qualitative comparison of images generated with different types of attention during the inversion phase. The first column shows pasted images, which we refer to as geometry images overlaid on target images, as shown in Fig. 2. The second column shows results for replacing K^{geo} and V^{geo} with K^{tar} and V^{tar} . The third column shows results with the ordinary self-attention computation. The fourth column shows results with Texture-aligning Attention, where target- and geometry-derived keys and values are concatenated as shown in Eq. (3).

tion as defined in Eq. (5)). Figure 6 provides a qualitative analysis of the generation results. Comparing (i) and (ii) or (iii), it was shown that the absence of source image attention leads to a significant degradation in geometry preservation, particularly evident in the transplantation of smaller regions. As for the comparison of (ii) and (iii), samples in the third and fourth rows illustrate that only using source-image-derived attention results in overemphasized geometry when transplantation regions become larger. These results indicate that although source-image-derived attention plays a key role in maintaining geometry information, the influence needs to be balanced by using self-attention simultaneously.

4.7. Limitation

The primary limitation of our work lies in its difficulty in transferring extremely large or small geometries. As evident from Eqs. (3) and (5), the ratio of texture-related to geometry-related information in our customized attention calculation is strongly dependent on the size of the geometry transfer area. In addition, the generated output occasionally exhibits significant geometric deviations from intended results or produces unrelated content. This limitation might be attributed to our method’s reliance on a pretrained SD model without additional training, potentially constraining its ability to fully capture texture and geometry. Future improvements could focus on dynamically adjusting attention based on transfer area size, developing more robust atten-




Figure 6. Qualitative comparison of images generated with different attention computation during the generation phase. The first column shows pasted images. The second column shows results for using only self-attention key K^{out} and value V^{out} during generation. The third column shows results for utilizing only source-image-derived key \hat{K}^{src} and value \hat{V}^{src} . The fourth column shows results with Geometry-preserving Attention, where source- and self-derived keys and values are concatenated as shown in Eq. (5).

tion mechanisms for diverse image combinations, and expanding generatable geometries through advanced training or architectural improvements. These refinements aim to address the current limitations and further improve the versatility and reliability of our approach across a wider range of geometry and image combinations.

5. Conclusion

In this work, we introduce Harmonizing Attention, a novel approach that facilitates the effective capture and transfer of material-independent geometry while preserving material-specific textural continuity. Our method uses custom Texture-aligning and Geometry-preserving Attention during inversion and generation processes, respectively, enabling the simultaneous referencing of source geometry and target texture information. Our approach achieves effective geometry transfer without requiring additional training or prompt engineering. Our method presented herein not only improves the creation of photorealistic composites but also expands the horizons of computer vision applications, ranging from augmented reality to advanced image editing.

References

- [1] D. Cohen-Or, O. Sorkine, R. Gal, T. Leyvand, and Y.-Q. Xu, “Color harmonization,” *ACM Transactions on Graphics*, vol. 25, no. 3, pp. 624–630, 2006. 1, 2
- [2] J. Jia, J. Sun, C. Tang, and H. Shum, “Drag-and-drop pasting,” *ACM Transactions on Graphics*, vol. 25, no. 3, pp. 631–637, 2006. 1, 2
- [3] J.-F. Lalonde and A. A. Efros, “Using color compatibility for assessing image realism,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2007. 1, 2
- [4] P. Perez, M. Gangnet, and A. Blake, “Poisson image editing,” *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 313–318, 2003. 1, 2
- [5] F. Pitie, A. C. Kokaram, and R. Dahyot, “N-dimensional probability density function transfer and its application to color transfer,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2005. 1, 2
- [6] E. Reinhard, M. Ashikhmin, B. Gooch, and P. Shirley, “Color transfer between images,” *IEEE Computer Graphics and Applications*, vol. 21, no. 5, pp. 34–41, 2001. 1, 2
- [7] K. Sunkavalli, M. K. Johnson, W. Matusik, and H. Pfister, “Multi-scale image harmonization,” *ACM Transactions on Graphics*, vol. 29, no. 4, pp. 125:1–125:10, 2010. 1, 2
- [8] M. W. Tao, M. K. Johnson, and S. Paris, “Error-tolerant image compositing,” in *Computer Vision – ECCV 2010* (K. Daniilidis, P. Maragos, and N. Paragios, eds.), vol. 6311 of *Lecture Notes in Computer Science*, (Berlin, Heidelberg), pp. 31–44, Springer, 2010. 1, 2
- [9] S. Xue, A. Agarwala, J. Dorsey, and H. E. Rushmeier, “Understanding and improving the realism of image composites,” *ACM Transactions on Graphics*, vol. 31, no. 4, pp. 84:1–84:10, 2012. 1, 2
- [10] J.-Y. Zhu, P. Krähenbühl, E. Shechtman, and A. A. Efros, “Learning a discriminative model for the perception of realism in composite images,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 1, 2
- [11] W. Cong, J. Zhang, L. Niu, L. Liu, Z. Ling, W. Li, and L. Zhang, “Dovenet: Deep image harmonization via domain verification,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8394–8403, 2020. 1, 2
- [12] X. Cun and C.-M. Pun, “Improving the harmony of the composite image by spatial-separated attention module,” *IEEE Transactions on Image Processing*, vol. 29, pp. 4759–4771, 2020. 1, 2
- [13] Z. Guo, D. Guo, H. Zheng, Z. Gu, B. Zheng, and J. Dong, “Image harmonization with transformer,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 14870–14879, October 2021. 1, 2
- [14] Y. Jiang, H. Zhang, J. Zhang, Y. Wang, Z. Lin, K. Sunkavalli, S. Chen, S. Amirghodsi, S. Kong, and Z. Wang, “SSH: A self-supervised framework for image harmonization,” in *Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4812–4821, 2021. 1, 2
- [15] J. Ling, H. Xue, L. Song, R. Xie, and X. Gu, “Region-aware adaptive instance normalization for image harmonization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9357–9366, 2021. 1, 2
- [16] W. Cong, L. Niu, J. Zhang, J. Liang, and L. Zhang, “Bargainnet: Background-guided domain translation for image harmonization,” in *Proceedings of the 2021 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1–6, 2021. 1, 2
- [17] K. Sofiuk, P. Popenova, and A. Konushin, “Foreground-aware semantic representations for image harmonization,” in *WACV*, pp. 1620–1629, 2021. 1, 2
- [18] Z. Guo, Z. Gu, B. Zheng, J. Dong, and H. Zheng, “Transformer for image harmonization and beyond,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 11, pp. 12960–12977, 2022. 1, 2
- [19] Y. Hang, B. Xia, W. Yang, and Q. Liao, “Scs-co: Self-consistent style contrastive learning for image harmonization,” in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19678–19687, 2022. 1, 2
- [20] Y. Song, Z. Zhang, Z. Lin, S. Cohen, B. Price, J. Zhang, S. Y. Kim, and D. Aliaga, “Objectstitch: Object compositing with diffusion model,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 18310–18319, June 2023. 1, 2
- [21] S. Lu, Y. Liu, and A. W.-K. Kong, “TF-ICON: Diffusion-based training-free cross-domain image

- composition,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2294–2305, 2023. 1, 2, 5, 6, 7
- [22] Y.-H. Tsai, X. Shen, Z. Lin, K. Sunkavalli, X. Lu, and M.-H. Yang, “Deep image harmonization,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3789–3797, 2017. 2
- [23] J.-F. Lalonde and A. A. Efros, “Using color compatibility for assessing image realism,” in *Proceedings of the 2007 IEEE 11th International Conference on Computer Vision (ICCV)*, pp. 1–8, 2007. 2
- [24] H. Wu, S. Zheng, J. Zhang, and K. Huang, “GP-GAN: Towards realistic high-resolution image blending,” in *Proceedings of the 27th ACM International Conference on Multimedia*, pp. 2487–2495, 2019. 2
- [25] L. Tan, J. Li, L. Niu, and L. Zhang, “Deep image harmonization in dual color spaces,” in *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 2159–2167, 2023. 2
- [26] H. Chen, Z. Gu, Y. Li, J. Lan, C. Meng, W. Wang, and H. Li, “Hierarchical dynamic image harmonization,” in *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 1422–1430, 2023. 2
- [27] Y. Xing, Y. Li, X. Wang, Y. Zhu, and Q. Chen, “Composite photograph harmonization with complete background cues,” in *Proceedings of the 30th ACM International Conference on Multimedia*, pp. 2296–2304, 2022. 2
- [28] L. Lu, J. Li, J. Cao, L. Niu, and L. Zhang, “Painterly image harmonization using diffusion model,” in *Proceedings of the 31st ACM International Conference on Multimedia*, p. 233–241, Association for Computing Machinery, 2023. 2, 6, 7
- [29] T.-F. Hsiao, B.-K. Ruan, and H.-H. Shuai, “Training-and-prompt-free general painterly harmonization using image-wise attention sharing,” 2024. arXiv:2404.12900v1 [cs.CV]. 2, 5, 6, 7
- [30] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “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. 2
- [31] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” in *Proceedings of the International Conference on Learning Representation (ICLR)*, 2021. 4
- [32] D. Garibi, O. Patashnik, A. Voynov, H. Averbuch-Elor, and D. Cohen-Or, “ReNoise: Real image inversion through iterative noising,” 2024. arXiv:2403.14602v1 [cs.CV]. 5
- [33] P. Bergmann, K. Batzner, M. Fauser, D. Sattlegger, and C. Steger, “The MVTEC anomaly detection dataset: A comprehensive real-world dataset for unsupervised anomaly detection,” *International Journal of Computer Vision*, vol. 129, pp. 1038–1059, April 2021. 5
- [34] P. Bergmann, M. Fauser, D. Sattlegger, and C. Steger, “MVTEC AD—A comprehensive real-world dataset for unsupervised anomaly detection,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9584–9592, 2019. 5
- [35] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595, 2018. 5
- [36] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, “Learning transferable visual models from natural language supervision,” in *Proceedings of the 38th International Conference on Machine Learning* (M. Meila and T. Zhang, eds.), vol. 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763, PMLR, 18–24 Jul 2021. 5
- [37] K. Ding, K. Ma, S. Wang, and E. P. Simoncelli, “Image quality assessment: Unifying structure and texture similarity,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 5, pp. 2567–2581, 2020. 5
- [38] B. Yang, S. Gu, B. Zhang, T. Zhang, X. Chen, X. Sun, D. Chen, and F. Wen, “Paint by example: exemplar-based image editing with diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 18381–18391, 2023. 6, 7