Can Location Embeddings Enhance Super-Resolution of Satellite Imagery?

Daniel Panangian* Ksenia Bittner
The Remote Sensing Technology Institute
German Aerospace Center (DLR), Wessling, Germany

{daniel.panangian, ksenia.bittner}@dlr.de

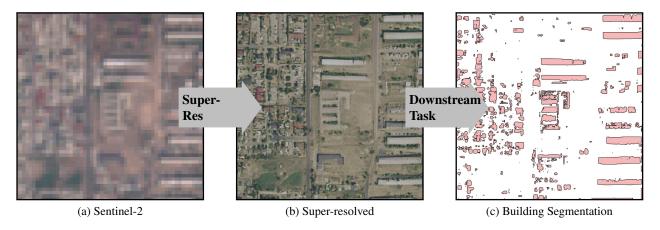


Figure 1. Illustration of the super-resolution pipeline and its application on an area of the city of Tashkent, Uzbekistan. (a) Sentinel-2 input image with low resolution, (b) super-resolved output, and (c) building segmentation result.

Abstract

Publicly available satellite imagery, such as Sentinel-2, often lacks the spatial resolution required for accurate analysis of remote sensing tasks including urban planning and disaster response. Current super-resolution techniques are typically trained on limited datasets, leading to poor generalization across diverse geographic regions. In this work, we propose a novel super-resolution framework that enhances generalization by incorporating geographic context through location embeddings. Our framework employs Generative Adversarial Networks (GANs) and incorporates techniques from diffusion models to enhance image quality. Furthermore, we address tiling artifacts by integrating information from neighboring images, enabling the generation of seamless, high-resolution outputs. We demonstrate the effectiveness of our method on the building segmentation task, showing significant improvements over state-ofthe-art methods and highlighting its potential for real-world applications.

1. Introduction

Recent advancements in remote sensing technologies have provided a wealth of satellite imagery, enabling applications in areas such as urban planning, disaster management, environmental monitoring, and resource management. Among these applications, building segmentation is particularly critical, especially for rapidly urbanizing regions. High-resolution imagery is crucial for accurately mapping and analyzing structures; however, limitations in publicly available satellite imagery, such as Sentinel-2, pose challenges due to their relatively low spatial resolution (10–20 m per pixel). To overcome these limitations, superresolution techniques have emerged as essential tools to enhance image quality and enable finer detail extraction from low resolution images [11].

Generative Adversarial Networks (GANs) [10] have been previously employed for super-resolution of remote sensing images, but their application has largely been confined to limited datasets, for example NAIP imagery, which covers only regions within the United States [30]. Models trained on these datasets often fail to generalize effectively to global regions, producing suboptimal results when applied elsewhere. Moreover when attempting to upscale

^{*}Corresponding author



Figure 2. Visual comparison of satellite images from geographically diverse regions using the proposed super-resolution pipeline. The left image in each pair represents the low-resolution Sentinel-2 input, while the right shows the corresponding super-resolved output, demonstrating the generalization of our method across different global urban landscapes.

large areas using tiling techniques, noticeable patch artifacts often emerge, degrading the overall quality of the output [30].

In recent years, the use of prior information through embeddings has gained popularity, most notably with the success of text embeddings in various tasks. Building on this trend, there has been a growing interest in leveraging geographic context via location embeddings. In this work, we introduce a novel approach that leverages location embedding to enhance super-resolution models. Additionally, we improve the GAN's performance by incorporating techniques used diffusion models. We also conduct experiments to address common patching issues caused by tiling, drawing inspiration from recent advancements in seamless image synthesis.

In summary, our contributions are threefold:

- 1. We develop first location-guided super-resolution model for remote sensing, designed to enhance generalization across diverse geographic regions by integrating location embeddings directly into the model.
- We improve a GAN-based super-resolution model's architecture by integrating attention mechanisms techniques to improve the scalability, and context understanding of the model.

 We adapt a seamless image synthesis method to superresolution, tackling common tiling artifacts problem by incorporating neighboring image data, ensuring the generation of continuous, high-resolution satellite imagery.

As shown in Fig. 1 and Fig. 2, our results demonstrate super-resolution of low-resolution satellite imagery, marking a key step forward in the application of publicly available dataset for global-scale remote sensing tasks.

2. Related Works

2.1. Remote-Sensing Super-Resolution

Super-resolution techniques in remote sensing have evolved from early Convolutional Neural Network (CNN) models like SRCNN [6] to more advanced generative models. CNN-based approaches, such as RCAN [34], use attention mechanisms to enhance important image features, resulting in sharper outputs. However, these models are computationally demanding and struggle to scale efficiently for large global datasets. GANs have further enhanced SR by introducing adversarial loss functions to improve perceptual quality. Models like SRGAN [18] and ESRGAN [29] generate realistic textures and high-frequency details with residual-in-residual dense blocks, making them suitable for

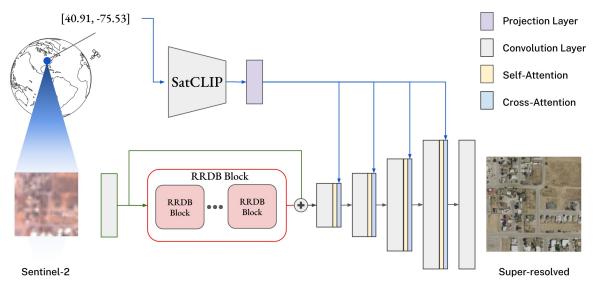


Figure 3. Summary of the proposed architecture. The input image from Sentinel-2 is processed through a shallow feature extraction network (CNN), followed by the RRDB block for deep feature extraction. Location embeddings, extracted using the SatCLIP module, are projected via MLP layers (purple) and injected into the network during the upsampling phase. The self-attention (yellow) and cross-attention (blue) mechanisms facilitate the integration of the image features and location embeddings, enabling the network to produce a super-resolved image that is both high in visual quality and geographically consistent.

large-scale remote sensing tasks due to their faster inference times. However, GAN-based methods can still introduce artifacts in regions with complex textures, affecting reliability. More recently, diffusion models like SR3 [26] have become state-of-the-art for generating high-resolution images, utilizing recursive refinement and attention mechanisms to surpass CNNs and GANs in visual fidelity. However, their high computational demands and slow processing make them less practical for large-scale satellite imagery applications. At the same time, there's a growing trend of integrating priors or different modalities to control and improve image generation and restoration. Models like GLEAN [3] and FeMaSR [4] utilize pretrained large generative models, such as StyleGAN [12] or VQGAN [8] for super-resolution. Some approaches also use modalities like text prompts to guide image generation [9, 25]. In the remote sensing domain, methods like DSen2 [17] and Real-GDSR [22] use multi-modal inputs, while DiffusionSAT [13] integrates metadata, such as time and geographic coordinates, to control image generation via conditional generative trained on big large data [25]. Diffusion-based techniques offer greater control and flexibility but their scalability remains a significant challenge, especially for largescale remote sensing applications.

The datasets available for remote sensing superresolution largely dictate the performance and generalization of the models trained on them. Many datasets use low-resolution Sentinel-2 imagery due to its free availability and balanced temporal and spatial resolution [5,20]. Highresolution data is often sourced from Google Earth [7], PlanetScope [24], or WorldView [16], which are regionally constrained and fail to cover diverse geographic conditions, which limits the practical applicability of models trained on them. To address this, we are exploring the use of available modalities and features to improve model generalization without relying solely on expensive high-resolution data.

2.2. Location Encoders

Location encoders transform geographic coordinates into rich, contextual embeddings that capture environmental and spatial dependencies, such as terrain, climate, and human land use. These models turned out to be effective in tasks like geolocation, spatial prediction, and environmental monitoring by embedding geographic context into deep learning systems. Models like GPS2Vec [31] and CSP [19] have applied location-based embeddings to tasks such as image retrieval and species classification, but they faced limitations in generalization due to uneven data distribution across geographic regions. GeoCLIP [2] further refined geolocation tasks using contrastive learning but still encountered similar biases. A more comprehensive approach emerged with SatCLIP [15], which uses satellite data to generate embeddings that generalize across global landscapes, allowing it to perform well in tasks like population density estimation and environmental prediction.

While recent developments like GeoSynth [27] have used these embeddings for satellite image generation, our work focuses on using geographic location encoders for image super-resolution. By incorporating location-based features from models like SatCLIP, we aim to enhance satellite image resolution while accounting for diverse geographic

and environmental conditions. This approach allows our super-resolution models to perform more effectively in remote sensing tasks by embedding spatial context directly into the resolution enhancement process, improving image clarity across varying terrains and climates.

3. Method

3.1. Problem Formulation

Given a low-resolution satellite image and its corresponding geographic location, our goal is to reconstruct a high-resolution image that reflects both fine visual details and the correct geographic characteristics. Unlike traditional super-resolution, where reconstruction relies solely on pixel data, our approach integrates location-based context to ensure the image aligns with real-world geography.

To achieve this, we address two main questions:

- How can we integrate location embeddings into the super-resolution network? These embeddings, can be integrated by concatenating or modulating them with image feature maps at various stages of the network. This allows the network to incorporate spatial context, enhancing the super-resolved image's geographic and environmental consistency.
- How do we design better conditioning to control the generated image? The location data must be integrated effectively with image features to ensure the reconstructed image is geographically consistent.

We introduce a cross-attention module to fuse the features from the low-resolution image with the location embeddings. The cross-attention mechanism allows the model to dynamically focus on regions of the image that require more refinement based on the geographic context. This helps the network associate specific textures or colors, such as vegetation or urban areas, with the location, enhancing the image's overall realism and detail. To further enhance the geographic consistency of the generated high-resolution images, we design a location-matching discriminator. This discriminator ensures that the textures and visual patterns in the high-resolution image align with the geographic context provided by the location embeddings.

3.2. Model Overview

As illustrated in Fig. 3, our proposed super-resolution model integrates satellite imagery and location embeddings for high-resolution image reconstruction. The architecture begins with the extraction of shallow features via a convolutional layer, which capture basic image attributes like edges and textures. The Residual-in-Residual Dense Blocks (RRDBs) then processes these shallow features to capture high-level, complex details essential for high-fidelity super-resolution.

A key element of our model is the integration of SatCLIP embeddings, which capture geographical context and urban patterns based on the image's coordinates. This step plays a vital role in incorporating location-specific information, helping the model generalize effectively across diverse regions. The SatCLIP embeddings are projected using Multi-Layer Perceptrons (MLPs) and are introduced into the network during the upsampling phase. We chose this, as the upsampling stage focuses on refining elements like texture and color without modifying the spatial features that were already captured during earlier feature extraction. By injecting the embeddings at this phase, the location-specific details such as textures and colors are adjusted in a way that enhances realism, while preserving the spatial structure.

3.3. Super-Resolution Network

Our model builds upon the foundation of ESRGAN [29], a state-of-the-art framework for image super-resolution. ESRGAN improves over traditional SRGAN by using RRDB blocks, which enhance the network's capacity and ease of training without the need for batch normalization. These blocks are designed to improve the recovery of complex textures and reduce artifacts, a challenge seen in traditional SRGAN. In this work we try to further improve it by utilizing self-attention mechanisms in the upsampling step to allow the network to focus on different parts of the image, capturing long-range dependencies and spatial relationships between regions. Applying self-attention at each stage helps the model maintain spatial consistency across different resolutions and overall coherence during the super-resolution process.

3.4. Location Conditioning

We then introduce location conditioning through crossattention after self-attention mechanisms. With the image features serving as the query, and the location embeddings as key-value pairs. This enables the model to fuse visual data from the image with the corresponding geographic context. This approach allows the model to first establish the internal consistency before applying more complex, external guidance in each upsampling stage.

However, despite these advantages, we add more explicit mechanism to evaluate and guide the generation process. We introduce a location-matching discriminator into the model. This discriminator not only evaluates the visual similarity between super-resolved and real images but also ensures geographic accuracy by comparing generated images with their true locations. Additionally, during training, we introduce random false locations, forcing the discriminator to distinguish between correct and incorrect geographic contexts. This encourages the model to not only integrate location information effectively but also ensures that the generated image consistently aligns with their respec-

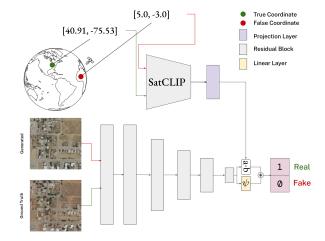


Figure 4. Overview of the location-matching Discriminator. The input consists of an image paired with real and false geographic coordinates, processed through the SatCLIP module. The image and the location embeddings are then passed through the discriminator network. The output of the discriminator is a binary classification (real or fake) based on whether the location matches the input image. The SatCLIP location embeddings guide the discriminator, enforcing geographic consistency in the image generation process.

tive location embeddings. As seen in Figure 4 the architecture of the discriminator is inspired by projection discriminator [21], a method designed to incorporate conditional information (such as location embeddings in our case) into the discriminator in a more effective way than traditional methods like concatenation. The projection discriminator applies a projection layer that takes the inner product between the condition and the intermediate features of the image within the discriminator. This leads to more robust conditioning than concatenation, as it directly influences the discriminator's decision-making by embedding the condition deeper into its structure.

3.5. Training Constraints

We define the *location-matching loss* as:

$$\mathcal{L}_{\text{loc_match}} = \mathbb{E}_{x_{\text{HR}}, x_{\text{LR}}, c, \hat{c}} \Big[\log D(x_{\text{HR}}, c) \\ + \log \left(1 - D(x_{\text{HR}}, \hat{c}) \right) \\ + \log \left(1 - D(G(x_{\text{LR}}, c), \hat{c}) \right) \Big]$$
(1)

where $x_{\rm HR}$ represents the high-resolution image, $x_{\rm LR}$ is the corresponding low-resolution image, c denotes the true location embedding, and \hat{c} is a false location embedding. This loss encourages the model to correctly associate geographic information by ensuring that images are matched with their true locations while distinguishing them from false ones.

The total objective function is:

$$L(G, D) = L_{\text{pix}} + L_{\text{perceptual}} + L_{\text{CLIP}} + L_{\text{OSM}} + L_{\text{loc match}}$$
(2)

While pixel loss and perceptual loss handle pixel-level and perceptual fidelity, we also incorporate CLIP loss and OpenStreetMap (OSM) discriminator loss from Satlas-SR [30]. CLIP loss ensures that the generated images align with high-level semantic content, while OSM loss ensures accurate reconstruction of geographic entities like buildings and roads based on OSM data. We enhance the OSM loss by preserving spatial relationships between these objects, ensuring geographic coherence in the generated images. Together, these losses balance visual quality and geographic accuracy. Our addition of the location-matching loss further enforces spatial correctness, ensuring that the generated images are not only perceptually accurate but also location-guided.

3.6. Local Padding

At inference, we replace all padding with adjacent neighboring patches, following the technique proposed by Latif et al. [1]. Instead of using conventional zero-padding, which can introduce artificial boundaries, this approach leverages the information from surrounding patches. By incorporating neighboring patches into the padding, we ensure smoother transitions and eliminate visible seams that can occur when concatenating patches. This method results in more consistent and seamless outputs during inference.

4. Experiments

4.1. Dataset

We utilize the S2-NAIP dataset [30], which pairs low-resolution Sentinel-2 images (64×64 pixels) with high-resolution National Agriculture Imagery Program (NAIP) imagery (512×512 pixels). This dataset is geographically situated within the United States. We apply random cropping to the Sentinel-2 images, reducing them to 32×32 pixels, while concurrently cropping the corresponding NAIP images to 256×256 pixels.

The dataset is geographically divided by Universal Transverse Mercator (UTM) zones. We chose UTM zones 13N, 18N, and 19N as the validation and test set for our experiments, as illustrated in Figure 5. These zones were selected based on their spatial positioning relative to the training zones: zone 13N is located in the middle of the training zones, while 18N and 19N are located at the far east. This geographic separation allows us to rigorously test the generalization ability of the model across different landscapes. We select central and edge UTM zones to assess the efficacy of embedding mechanisms within our model. Central UTM zones represent regions centrally located within the

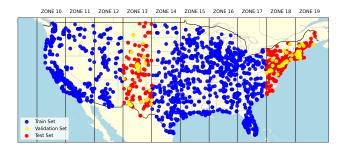


Figure 5. Visualization of the data split by UTM zones. The train set includes samples from UTM zones 10-17, while the validation and test set consists of samples from UTM zones 13N, 18N, and 19N. These zones were selected to evaluate how well the model generalizes across different geographic regions.

dataset's spatial distribution, providing a baseline for typical geographic variability. In contrast, edge UTM zones are chosen from the peripheries to introduce challenges related to spatial extrapolation and embedding generalization.

4.2. Implementation Details

Our model is implemented using the BasicSR [28] framework, a widely-used platform for super-resolution tasks. To benchmark our approach, we compare it against Satlas-SR [30], which serves as our baseline. Satlas-SR primarily contributes by integrating CLIP loss for faster convergence [23] and an object-aware discriminator that ensures spatial consistency. These components are incorporated into our evaluation to provide a fair and comprehensive comparison. Additionally, we compare our results with Stable Diffusion [25] with ControlNet [32], which focuses on conditional image generation rather than traditional super-resolution. ControlNet enhances Stable Diffusion by allowing the network to condition on additional control signals, such as edge maps or segmentation maps, to generate images that adhere more strictly to the provided conditions. A similar approach is explored in Diffusion-SAT [13], which uses diffusion models for remote-sensing image generation and other tasks such as SR and inpainting by leveraging metadata like geolocation and temporal information. Due to the unavailability of its conditional generation code at the time of our experiments, we did not include DiffusionSAT in our evaluation.

For quantitative evaluation, we use four key metrics: Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) for pixel-level accuracy, Learned Perceptual Image Patch Similarity (LPIPS) [33] for perceptual quality, and the CLIP score [23] to evaluate imagetext alignment. In addition to these metrics, we evaluate the performance of the super-resolved images on semantic segmentation task. Specifically, we fine-tune the Segment

Anything Model (SAM) [14] on NAIP imagery and building segmentation labels derived from OSM. We then assess the quality of the segmentation using metrics such as mean Intersection over Union (mIoU), F1 score for building detection, and mean F1 score (mF1) for buildings. This provides additional insights into the effectiveness of our superresolution model in solving downstream tasks in remote sensing. More detailed descriptions of baselines, hyperparameters, and training details can be found in the supplementary Sec. S2.

5. Results

In this section, we present both qualitative and quantitative comparisons of our model against the baselines, including Satlas-SR and Stable Diffusion with ControlNet (SD + ControlNet). We evaluate the models on super-resolution and building segmentation tasks, showcasing the effectiveness of our approach. Additionally, we provide results on the impact of local padding in the supplementary Sec. S3.

5.1. Super-Resolution

In Tab. 1, we observe that Satlas-SR delivers the highest PSNR and LPIPS value and of 15.6481, indicating superior pixel-level reconstruction accuracy. However, our model slightly outperforms all others in CLIP score, achieving 0.9261, which highlights its strong alignment with reference data, particularly in terms of semantic relevance. This demonstrates competitive performance of our model in both pixel-level fidelity and perceptual accuracy with other methods

Figure 6 visually compares the super-resolved images generated by our method and baselines. In the first row, depicting an industrial area, the super-resolved image produced by our model preserves not only finer building edges and textures but also better captures the road structures compared to Satlas-SR and SD + ControlNet. Satlas-SR fails to produce clear outlines and shapes particularly for the buildings and roads. SD + ControlNet, on the other hand, hallucinates vegetation in areas where there is none, leading to visually unrealistic elements in the scene. Our model achieves a better balance between sharpness and naturalness, recovering both roads and building structures better than other baselines.

In the second row, depicting a suburban area, our method offers better contrast and more details, making it more visually aligned with the ground truth. The roads and building edges are more clearly defined, enhancing the overall realism. In comparison, SD + ControlNet does a reasonable job at reconstructing some buildings but introduces repetitive patterns that give the image an unnatural, synthetic appearance. Additionally, the variation in vegetation color creates an inconsistency, making it seem like the image was captured in a different season. Satlas-SR struggles



Figure 6. Performance comparison of super-resolution methods on satellite imagery from the test set. Each row shows a distinct scene with the Sentinel-2 input, super-resolved outputs, and NAIP image for reference.

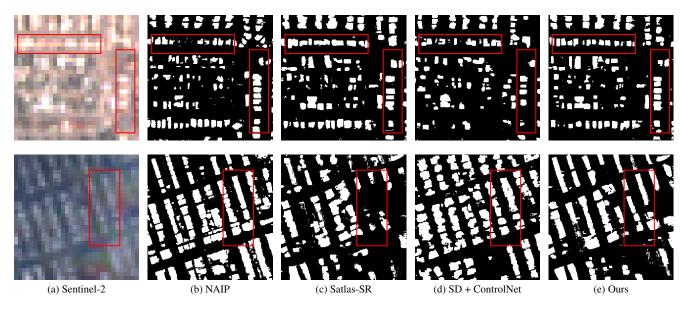


Figure 7. Building segmentation results of different super-resolution methods on the test set. The NAIP column represents the results of the segmentation model using NAIP images as input, serving as the reference for evaluating the accuracy of the other approaches.

Table 1. Quantitative comparison of super-resolution methods on the test set

Model	PSNR	SSIM	LPIPS ↓	CLIP
Satlas-SR	15.6481	0.2621	0.4786	0.9227
SD + ControlNet	13.4392	0.1594	0.5640	0.8980
Ours	14.3973	0.2221	0.4855	0.9261

with over-reconstructing buildings, adding structures where there should be none, leading to a cluttered and less accurate representation.

In the third row, all methods struggle to accurately reconstruct the complex roof structures. SD + ControlNet generates more details in the roofs, giving them a more realistic texture. However, it fails to properly reconstruct the roads, with noticeable distortions and misaligned intersections. Satlas-SR and our model handle the roads better, maintaining clearer paths and intersections. Satlas-SR, however, suffers from a lack of contrast and struggles with the reconstruction of small houses, resulting in a more blurred appearance overall. Our model stands out by producing higher contrast and better-defined small structures compared to Satlas-SR, especially in the residential areas with smaller houses. While SD + ControlNet adds realistic texture to some features, the overall balance in road and building reconstruction, combined with superior contrast and clarity, makes our method the closest to the ground truth, especially when considering the finer details in the smaller residential buildings and road layouts.

Table 2. Segmentation performance of our method and baselines. Our method outperforms others.

Model	mIoU	mF1-score	Building	
			IoU	F1-Score
Satlas-SR	51.34	59.96	20.04	29.79
SD + ControlNet	46.55	55.25	14.25	22.60
Ours	55.43	64.04	26.40	36.79
NAIP	59.35	67.88	32.66	43.55

5.2. Segmentation

For the segmentation task, Tab. 2 reports the mIoU and mF1-score metrics for building segmentation on the super-resolved images. Our model outperforms both Satlas-SR and SD + ControlNet across all segmentation metrics. Specifically, we achieve the highest mIoU of 55.43 and mF1-score of 64.04, significantly surpassing the other models. In terms of building segmentation, our model also performs best with IoU and F1-score values of 26.40 and 36.79, respectively, demonstrating that our super-resolved

images provide better foundations for subsequent segmentation tasks. Figure 7 provides a qualitative comparison for building segmentation. Our model's segmentation masks are more accurate and closely resemble the ground truth compared to those generated by Satlas-SR and SD + ControlNet. In both the examples provided, Satlas-SR and SD + ControlNet exhibit issues with boundary detection and misclassification of small buildings, whereas our model provides cleaner and more precise segmentations, especially in densely built environments.

6. Conclusion

In this work, we introduced a novel approach to satellite image super-resolution by implementing location guidance, attention mechanisms, and seamless image synthesis techniques. Our method demonstrates state-of-the-art performance, particularly in preserving structural details of things like buildings and roads compared to other models.

However, we observe the discrepancy between the superresolution and segmentation metrics. Metrics like PSNR, SSIM, and LPIPS are designed to measure pixel-level fidelity and perceptual quality and do not always correlate with the utility of the super-resolved images in downstream tasks such as segmentation. This suggests that traditional super-resolution metrics may not fully capture the aspects of image quality that are most beneficial for subsequent analytical tasks. These findings underscore the need for a more holistic approach to evaluating super-resolution models in remote-sensing images.

Other than that, challenges related to generalization persist. Although our model excels in standard benchmarks, its ability to generalize to regions significantly different from the training data remains limited. Experiments conducted on areas outside the training set reveal that while our model accurately reconstructs building structures, it struggles to generate new textures or styles, often defaulting to hue changes without adapting underlying textures. This limitation is especially evident when coordinates are shifted to test the model's adaptability in unfamiliar areas. Full details of this experiment is provided in the supplementary Sec. S5. These findings indicate that while our method is highly effective in controlled scenarios, improving its ability to adapt to diverse, unseen regions will be crucial for broader real-world applications. Future work will focus on enhancing the generalization capabilities of our model, potentially by incorporating more diverse training datasets, leveraging large generative vision models, or exploring domain adaptation methods to better handle texture variability.

References

 Alhasan Abdellatif, Ahmed H. Elsheikh, and Hannah Menke. Local padding in patch-based gans for seamless infinite-sized texture synthesis, 2024.

- [2] Javier Cepeda et al. Geoclip: Learning multimodal embeddings for geolocation from flickr images. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023. 3
- [3] Kelvin CK Chan, Xintao Wang, Xiangyu Xu, Jinwei Gu, and Chen Change Loy. Glean: Generative latent bank for large-factor image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [4] Chaofeng Chen, Xinyu Shi, Yipeng Qin, Xiaoming Li, Xiaoguang Han, Tao Yang, and Shihui Guo. Real-world blind super-resolution via feature matching with implicit high-resolution priors. In *Proceedings of the 30th ACM International Conference on Multimedia*, MM '22, page 1329–1338, New York, NY, USA, 2022. Association for Computing Machinery. 3
- [5] Julien Cornebise, Ivan Oršolić, and Freddie Kalaitzis. Open high-resolution satellite imagery: The worldstrat dataset – with application to super-resolution. 07 2022. 3
- [6] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(2):295– 307, 2016. 2
- [7] Runmin Dong, Lixian Zhang, and Haohuan Fu. Rrsgan: Reference-based super-resolution for remote sensing image. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–17, 2021. 3
- [8] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis, 2020. 3
- [9] Kanchana Vaishnavi Gandikota and Paramanand Chandramouli. Text-guided explorable image super-resolution. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 25900–25911, June 2024. 3
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances* in neural information processing systems, pages 2672–2680, 2014.
- [11] Svetlana Illarionova, Dmitrii Shadrin, Islomjon Shukhratov, Ksenia Evteeva, Georgii Popandopulo, Nazar Sotiriadi, Ivan Oseledets, and Evgeny Burnaev. Benchmark for building segmentation on up-scaled sentinel-2 imagery. *Remote Sens*ing, 15(9), 2023. 1
- [12] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020. 3
- [13] Samar Khanna, Patrick Liu, Linqi Zhou, Chenlin Meng, Robin Rombach, Marshall Burke, David B. Lobell, and Stefano Ermon. Diffusionsat: A generative foundation model for satellite imagery. In *International Conference on Learn*ing Representations, 2024. 3, 6
- [14] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. arXiv:2304.02643, 2023.

- [15] Konstantin Klemmer et al. Satclip: Global, general-purpose location embeddings with satellite imagery. *arXiv preprint arXiv:2311.17179*, 2023. 3
- [16] Paweł Kowaleczko, Tomasz Tarasiewicz, Maciej Ziaja, Daniel Kostrzewa, Jakub Nalepa, Przemysław Rokita, and Michal Kawulok. A real-world benchmark for sentinel-2 multi-image super-resolution. *Scientific Data*, 10, 09 2023.
- [17] Charis Lanaras, José Bioucas-Dias, Silvano Galliani, and Emmanuel Baltsavias. Super-resolution of sentinel-2 images: Learning a globally applicable deep neural network. ISPRS Journal of Photogrammetry and Remote Sensing, 146, 03 2018. 3
- [18] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 4681–4690, 2017.
- [19] Gengchen Mai et al. Contrastive spatial pre-training for geographic data. *Journal of Geographic Information Systems*, 2023. 3
- [20] Julien Michel, Juan Vinasco-Salinas, Jordi Inglada, and Olivier Hagolle. Sen2venμs, a dataset for the training of sentinel-2 super-resolution algorithms. *Data*, 7:96, 07 2022.
- [21] Takeru Miyato and Masanori Koyama. cgans with projection discriminator. ArXiv, abs/1802.05637, 2018. 5
- [22] Daniel Panangian and Ksenia Bittner. Real-gdsr: Real-world guided dsm super-resolution via edge-enhancing residual network. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, 10:185–192, 2024. 3
- [23] Alec Radford, Jong Wook Kim, Jack Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Saurabh Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2021. 6
- [24] Muhammed Razzak, Gonzalo Mateo-García, Luis G'omez-Chova, Yarin Gal, and Freddie Kalaitzis. Multi-spectral multi-image super-resolution of sentinel-2 with radiometric consistency losses and its effect on building delineation. *ArXiv*, abs/2111.03231, 2021. 3
- [25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. 2021. 3, 6
- [26] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. arXiv preprint arXiv:2104.07636, 2021. 3
- [27] Srikumar Sastry, Subash Khanal, Aayush Dhakal, and Nathan Jacobs. Geosynth: Contextually-aware highresolution satellite image synthesis. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern

- Recognition (CVPR) Workshops, pages 460–470, June 2024.
- [28] Xintao Wang, Liangbin Xie, Ke Yu, Kelvin C.K. Chan, Chen Change Loy, and Chao Dong. BasicSR: Open source image and video restoration toolbox. https://github.com/XPixelGroup/BasicSR, 2022. 6
- [29] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Chen Change Loy, Yu Qiao, and Xiaoou Tang. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pages 0–0, 2018. 2, 4
- [30] Piper Wolters, Favyen Bastani, and Aniruddha Kembhavi. Zooming out on zooming in: Advancing super-resolution for remote sensing, 2023. 1, 2, 5, 6
- [31] Yifang Yin, Zhenguang Liu, Ying Zhang, Sheng Wang, Rajiv Ratn Shah, and Roger Zimmermann. Gps2vec: Towards generating worldwide gps embeddings. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 416—419, 2019. 3
- [32] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023. 6
- [33] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [34] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 286–301, 2018. 2