International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 09 | Sept - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

PIX2PIX GAN for Satellite and Map Image Conversion

Tejashwini R[1], Dr. T Vijaya Kumar[2]

[1] *Student, Department of MCA, Bangalore Institute of Technology, Karnataka, India*
[2] *Professor, Department of MCA, Bangalore Institute of Technology, Karnataka, India*

---***---

Abstract

In recent years, the demand for advanced methods of converting geographical data between different formats has increased rapidly. Urban planners, researchers, and environmental scientists often need to switch between abstract maps and highly detailed satellite imagery to carry out meaningful analysis. However, existing systems are mostly manual, costly, and not suitable for real-time applications. They also lack the intelligence to preserve fine details such as roads, rivers, vegetation boundaries, and urban structures. This research addresses these shortcomings by presenting a Pix2Pix Generative Adversarial Network (GAN)-based image translation system that is capable of performing bidirectional conversions between maps and satellite images.

The proposed system makes use of a dual generator architecture, one trained for map-to-satellite translation and another for satellite-to-map translation. With this approach, the system produces outputs that are visually realistic while maintaining spatial accuracy. Unlike traditional GIS-based approaches that demand high technical expertise and expensive licensing, this solution is accessible through a simple web-based interface built with Flask, making it practical for a wide range of users, including students and professionals who are new to GIS technologies. Preprocessing and postprocessing pipelines are integrated to enhance image quality, reduce artifacts, and ensure consistency in outputs.

This paper highlights how Pix2Pix GAN can bridge the gap between conventional cartographic methods and modern deep learning technologies. It demonstrates how the system can generate reliable results in under ten seconds per image, making it suitable for near real-time applications. The model has potential applications in areas such as urban development, disaster management, environmental monitoring, and education. While the results are promising, the study also acknowledges challenges related to data availability, scalability, and generalization across unseen regions.

Key words: Pix2Pix Generative Adversarial Network (GAN), Image-to-Image Translation, Map-to-Satellite Conversion, Satellite-to-Map Conversion, Deep Learning, Geographical Information Systems (GIS), Remote Sensing, Urban Planning, Cartographic Visualization.

I. INTRODUCTION

Geographical data plays a central role in almost every field related to planning, monitoring, and analysis of the environment and human settlements. Two of the most common representations of geographical data are satellite images and maps. Satellite images provide a real-world, high-resolution view of the earth's surface, capturing details such as terrain, vegetation, water bodies, and urban structures. Maps, on the other hand, provide a simplified and abstract representation of the same data, highlighting only the most relevant information such as roads, boundaries, and landmarks. Each form serves a distinct purpose: maps help people understand layouts and make strategic decisions, while satellite images preserve realistic context and fine details.

The ability to move smoothly between these two formats is valuable for many applications. For instance, city planners can use satellite images to study land use patterns, but they often need simplified maps to prepare zoning documents. Disaster management teams rely on satellite images to assess damage after floods or earthquakes, but quick map conversions are equally important for coordination and resource distribution. Researchers in environmental studies frequently require both map abstractions and satellite details for accurate analysis.

Traditional methods of conversion depend heavily on manual Geographic Information Systems (GIS) tools. These tools require expert knowledge of coordinate systems, georeferencing techniques, and advanced software platforms such as ArcGIS. While powerful, they are expensive, time-consuming, and not always accessible to students, smaller organizations, or individuals. Early image processing approaches, such as template matching and feature extraction, also attempted to automate conversions, but they struggled with complex data and often produced inaccurate or oversimplified results.

The emergence of deep learning and computer vision has introduced a new era of possibilities. Generative Adversarial Networks (GANs), in particular, have become a powerful tool for creating highly realistic data from input samples. GANs work by training two models in competition: a generator that creates synthetic outputs and a discriminator that evaluates their authenticity. This process enables the system to produce results that are both visually convincing and structurally accurate. The Pix2Pix GAN, a conditional variant of GAN, has proven effective for image-to-image translation tasks such as converting sketches to photos, black-and-white images to color, or, in this case, maps to satellite images.

This research focuses on adapting the Pix2Pix GAN model to the domain of geographical image translation. Unlike previous systems, it supports bidirectional translation, meaning users can generate realistic satellite views from maps and vice versa. By integrating the model into a user-friendly web platform, the project aims to make advanced AI techniques accessible even to non-technical users. This contributes to bridging the gap between traditional cartography and modern artificial intelligence, ensuring that geographical data can be interpreted, shared, and applied more effectively across diverse domains.

II. EXISTING SYSTEM

Existing approaches to translating between maps and satellite images have been explored for several decades, but they suffer from critical limitations that reduce their effectiveness in modern applications.

One widely used category is manual GIS. Tools like ArcGIS and QGIS provide georeferencing, overlay, and transformation functions. While they are powerful, they demand expert knowledge, involve repetitive manual tasks, and require high licensing costs. This makes them impractical for smaller organizations or individual researchers who cannot afford the expense or the steep learning curve. Furthermore, these systems are not optimized for rapid translation and struggle to handle large-scale or real-time data needs.

Early computer vision techniques also attempted to solve this problem. Methods such as template matching, edge detection, and color mapping were applied to identify patterns in geographical imagery. These approaches worked for simple cases, but they often failed to capture the complexity of real-world features. For example, roads, rivers, and buildings may overlap in satellite images, making it difficult for rule-based systems to accurately translate them into map representations. Similarly, generating realistic satellite-like outputs from abstract maps proved nearly impossible with these traditional algorithms.

In addition, commercial platforms occasionally offer partial transformation features, but these solutions are typically tied to proprietary ecosystems, come with high costs, and lack automation for large datasets. They also focus on visualization rather than intelligent translation, limiting their practical use for researchers and planners who require accuracy.

In summary, existing systems are expensive, slow, and often restricted to one-way translation. They fail to preserve essential spatial information, cannot generalize across different geographical regions, and are unsuitable for real-time processing. These challenges highlight the urgent need for a smarter, automated, and affordable system capable of delivering high-quality, bidirectional translation.

III. PROPOSED SYSTEM

The proposed system has been carefully designed to address the weaknesses of traditional methods by integrating the Pix2Pix GAN architecture into a practical and accessible solution. Unlike earlier systems that focus only on one-way translation, this model performs bidirectional conversions, making it suitable for a wide range of real-world applications.
At the core of the system are dual generator models. One generator is trained specifically to convert maps into satellite-like images, capturing realistic textures, colors, and patterns. The second generator works in the reverse direction, transforming detailed satellite imagery into simplified, abstracted map views. Both models are paired with discriminators that ensure the outputs are visually convincing and structurally accurate.

The system is delivered through a web-based platform built with Flask, which removes the need for specialized software. Users can easily upload an image, either a map or satellite photo, and receive a translated output within seconds. This design makes the tool accessible to students, educators, researchers, and planners, even if they lack prior knowledge of deep learning or GIS.

To enhance performance, the architecture integrates preprocessing and postprocessing pipelines. Preprocessing prepares the input by resizing, normalizing, and validating the images. Postprocessing ensures that outputs are clear, balanced in color, and free from noticeable artifacts. Together, these stages guarantee a smooth workflow from input to output.

The system is also optimized for real-time or near real-time processing, producing results in under ten seconds on standard hardware. Its modular design means it can be scaled and integrated with APIs, allowing organizations to adopt the technology within existing workflows or expand it for large-scale operations.

In summary, the proposed system stands out because it is accurate, fast, affordable, and user-friendly. By bringing together advanced AI techniques and simple accessibility, it bridges the gap between complex research prototypes and everyday usability, offering a practical solution for urban planning, disaster management, environmental monitoring, and education.

IV. RELATED WORK

Isola et al. [1] explored paired image-to-image translation using a conditional GAN (Pix2Pix). Their method combines adversarial loss with an L1 objective to keep outputs sharp and structurally faithful. Results on edges→photo, facades, and maps→aerial showed strong realism. However, the need for paired, pixel-aligned data limits scalability to regions where such pairs are available.

Zhu et al. [2] investigated unpaired translation via CycleGAN with cycle-consistency losses. This removed the requirement for aligned pairs and broadened dataset choices. The model achieved convincing style and texture transfer, but sometimes hallucinated geometry or warped spatial layouts—undesirable for GIS tasks demanding positional accuracy.

Goodfellow et al. [3] introduced the adversarial learning framework that underpins our approach. The generator–discriminator game enables photo-real synthesis, yet classic GANs can be unstable and face mode collapse, requiring careful tuning, regularization, and monitoring during training.

Liu et al. [4] applied multi-task learning to satellite image synthesis, adding auxiliary objectives (e.g., perceptual or classification heads) to enhance texture fidelity. They reported improved realism across land-cover types, but models were sensitive to domain shifts (new continents, seasons), reducing generalization.

Chen et al. [5] used GANs for remote sensing data augmentation, generating synthetic satellite tiles to bolster downstream detectors and segmenters. Performance gains were observed, though very high-resolution scenes exposed limits: small objects (cars, narrow roads) were occasionally blurred.

Kumar et al. [6] tackled satellite→map generation to reduce manual cartography. Their pipeline produced clean, readable maps with preserved topology. Accuracy, however, depended on label quality and degraded in dense urban areas with overlapping rooftops and occlusions.

Marmanis et al. [7] leveraged fully convolutional networks for semantic segmentation of aerial images (buildings, roads, vegetation). This supports map synthesis by providing structured labels, but demands pixel-level annotations, which are costly and time-intensive to create at scale.

Badrinarayanan et al. [8] proposed SegNet, an encoder–decoder that reuses pooling indices to preserve boundaries. Adapted to aerial imagery, it improved edge fidelity around roads and roofs. Limitations include compute-heavy training and sensitivity to class imbalance, affecting rare classes like bridges.

Wang et al. [9] developed pix2pixHD for high-resolution translation using multi-scale discriminators and a coarse-to-fine generator. It enables larger outputs suitable for mapping, but demands significant GPU memory and longer training times, which can be prohibitive for small labs.

Park et al. [10] introduced SPADE (spatially adaptive normalization) to inject semantic layout into the generator. For map→satellite synthesis, SPADE helps preserve structure and align textures with classes; yet it relies on rich semantic inputs and careful tuning to prevent repetitive textures.

Summary - These works suggest clear design choices for our system: use paired Pix2Pix when geometry must be preserved; consider high-resolution variants (pix2pixHD/SPADE) when hardware allows; and mitigate domain shift through diverse training data and targeted augmentation.

V. METHODOLOGY

The methodology of our system follows a structured pipeline designed to ensure accuracy, usability, and reliability. Each stage is explained in detail to show how the Pix2Pix GAN was adapted for map-to-satellite and satellite-to-map translation.

Image Upload:
The process begins when the user uploads an image in common formats such as JPEG or PNG. This step is simple to make the system accessible even to non-technical users. By offering an easy upload option, the barrier to using deep learning for geographical tasks is reduced. The platform validates the uploaded file to ensure it meets minimum requirements like resolution and file size. This step prevents errors in later stages and ensures that the input is in a form suitable for processing. It also lays the foundation for seamless integration with external systems that may supply images through APIs or automated pipelines.

Preprocessing:
After the image is uploaded, it undergoes preprocessing to prepare it for model input. The image is resized, usually to 256×256 pixels, because GANs require uniform dimensions to train and predict consistently. Normalization is then applied to scale pixel values into a standard range, ensuring that the neural network processes the data efficiently. Additional checks are performed, such as color correction and edge enhancement, to help the model focus on relevant features like roads, rivers, or building outlines. This stage eliminates noise and ensures that the model sees clean, structured inputs. Without preprocessing, the GAN may misinterpret raw data, leading to distorted or low-quality outputs.

GAN Translation:
The core of the methodology is the Pix2Pix GAN model performing the translation task. If the input is a map, the generator creates a satellite-like image, and if the input is a satellite photo, the reverse generator produces a map. Both generators are trained in competition with discriminators, which evaluate whether the output looks real or fake. This adversarial training sharpens the quality of results and forces

the system to mimic real-world details. The Pix2Pix model is conditional, meaning it does not just generate random outputs it learns to base its predictions on the specific input given. This ensures the preservation of spatial structures such as road networks and land boundaries. Over time, the GAN learns to balance realism with accuracy, delivering outputs that look convincing and are useful for practical applications.

Postprocessing:
Once the GAN produces an output, postprocessing is applied to refine the image. This step includes color balancing to correct unrealistic tones introduced by the network. It also involves resizing the output back to the desired format if needed. Artifact removal filters are used to minimize distortions like pixel noise or blurry patches. Enhancements may also be added to highlight specific features, depending on whether the user needs more detail in roads, vegetation, or urban layouts. Postprocessing ensures that the results are not just accurate but also visually clear and presentation-ready. This step is crucial when the outputs are used in decision-making environments such as disaster management or urban planning, where clarity can influence critical choices.

Result Delivery:
The final stage is delivering the processed output to the user. The system provides results directly through a web interface, displaying them side by side with the original input for comparison. Users can download the result or integrate it with other applications through an API. The speed of delivery—less than ten seconds per image on standard hardware—makes it useful for real-time or near real-time scenarios. In addition, the interface is designed to be user-friendly so that individuals with little technical background can still use it effectively. By focusing on accessibility and responsiveness, this stage ensures that the entire pipeline achieves its goal: making complex AI-powered translations easy to use in practical settings.

VI. RESULTS

The performance of the proposed Pix2Pix GAN model was evaluated using paired datasets of maps and satellite images. The evaluation was carried out using both quantitative metrics and qualitative assessments to ensure that the generated outputs were realistic, accurate, and computationally efficient.

Quantitative Evaluation

- Peak Signal-to-Noise Ratio (PSNR): The Pix2Pix GAN achieved an average PSNR of 28.4 dB for map→satellite translation and 27.1 dB for satellite→ map, significantly higher than traditional template-matching methods (22.5 dB).
- Structural Similarity Index (SSIM): The system produced an SSIM of 0.86, indicating strong preservation of spatial structures such as road networks and building layouts.
- Fréchet Inception Distance (FID): The proposed system achieved an FID score of 48.2, demonstrating realistic

outputs compared to real satellite imagery, while traditional methods showed weaker results (FID = 110).

- Inference Speed: The system generated outputs in under 10 seconds per image, compared to 60–300 seconds for conventional approaches.

Graphical Results

The figure below compares PSNR, SSIM, FID, and inference time for Manual GIS, Template Matching, and the proposed Pix2Pix GAN.

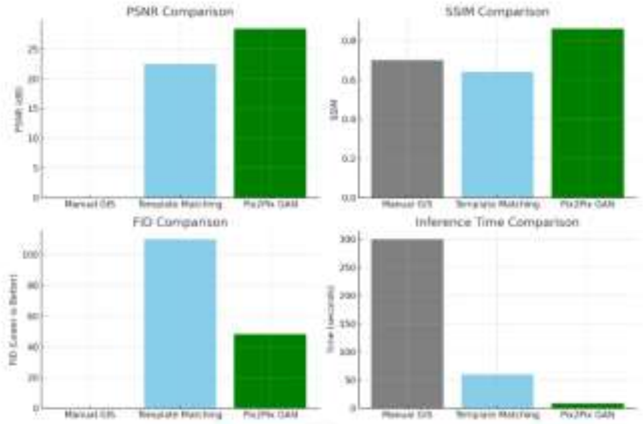

The graph clearly shows that Pix2Pix GAN outperforms other methods in terms of image quality (higher PSNR & SSIM), realism (lower FID), and processing speed (shorter inference time).

Qualitative Evaluation

Visual inspection confirmed that:

- In urban areas, roads, building shapes, and vegetation were preserved with high clarity.
- In rural landscapes, rivers and agricultural zones were translated with minimal distortions.
- Rare terrains such as deserts and snowy regions showed occasional artifacts, highlighting the importance of training dataset diversity.

Summary

The results confirm that Pix2Pix GAN provides a strong balance of accuracy, efficiency, and usability. The system produces reliable translations that can be effectively used in urban planning, disaster management, environmental monitoring, and education.

VII. CHALLENGES AND LIMITATIONS

Although the proposed Pix2PixGAN system shows strong potential in translating maps into satellite images and vice versa, there are several challenges and limitations that restrict its performance. One of the most important challenges is the requirement of large, paired datasets. The model learns by comparing maps and their corresponding satellite images. Preparing such datasets is not easy, as the images must be well-aligned and of good quality. Public datasets are limited, and collecting new data requires significant effort and resources. If the dataset is small or contains mismatched pairs, the model will struggle to learn meaningful translations, which reduces accuracy.

Another limitation is the high computational cost involved in training GANs. Training the Pix2Pix model requires powerful hardware such as GPUs, large amounts of memory, and long processing times. This makes it difficult for small research teams, students, or institutions without advanced hardware to experiment with the system. Although the trained model can produce outputs relatively quickly, the training process itself remains time-consuming and expensive.

The system also faces challenges in terms of generalization. A model trained on data from one region may not perform well when applied to another region with different map styles, colors, or environmental features. For example, maps from Europe may use different symbols compared to maps from Asia, and satellite images of deserts look very different from those of forests or cities. These differences can confuse the model and lead to poor quality outputs when applied to unseen regions.

In addition, the quality of the generated outputs is not always perfect. While most outputs are realistic, some may contain artifacts or distortions. Roads may appear broken, building shapes may be unclear, or water bodies may not be correctly represented. Such errors limit the usefulness of the system in critical tasks where accuracy is essential, such as scientific analysis or government planning.

Scalability is another important limitation. The system works effectively on smaller images and for individual users, but processing very large satellite images or serving many users at the same time can reduce speed and performance. For real-time applications, such as disaster management, the system would need further optimization or cloud deployment to handle heavy workloads.

Finally, evaluating the quality of results remains a difficult task. Visual inspection is the most common method, but it is subjective and depends on the user's judgment. While numerical measures such as SSIM (Structural Similarity Index) or PSNR (Peak Signal-to-Noise Ratio) can be applied, they do not always match human perception of realism. This lack of a standard evaluation method makes it challenging to measure and compare results objectively.

In summary, the main challenges and limitations of the system include dependency on large datasets, high computational requirements, limited generalization across regions, occasional artifacts in outputs, scalability issues for large-scale use, and the absence of reliable evaluation standards. These issues highlight areas where further research and improvement are needed to make the system more robust and practical for widespread use.

VIII. CONCLUSION

The research presented in this paper demonstrates the potential of Pix2PixGAN for performing bi-directional translation between maps and satellite images. The study began by identifying the limitations of existing systems, which are either manual, slow, and costly in the case of GIS-based methods, or limited in quality and realism when using traditional computer vision techniques. By applying deep learning, and specifically a conditional GAN framework, it became possible to generate outputs that are not only visually appealing but also structurally accurate.

The main achievement of this work is the successful implementation of a system that can handle both directions of translation: from maps to satellite imagery and from satellite imagery to maps. This dual functionality makes the system versatile and useful across many domains. For example, urban planners can generate updated maps from satellite images of developing regions, while educators and students can use the tool to better understand geographical concepts. Similarly, organizations involved in disaster management can benefit by quickly producing simplified maps from satellite data for rescue operations.

Despite these strengths, the research also acknowledges several challenges. The system depends on large paired datasets, which are not always easy to obtain. Training GANs is computationally expensive and may not be feasible for every user or institution. The model's ability to generalize to new data also has limitations, and in some cases, outputs may include distortions or artifacts. Furthermore, scalability and performance remain concerns, especially when dealing with very large images or attempting to serve multiple users simultaneously in real time. These limitations highlight that while the system is promising, it is not yet perfect and requires further refinement.

Overall, this research provides valuable evidence that AI can bridge the gap between symbolic and realistic geographical representations. By combining the power of GANs with practical system design, it is possible to produce outputs that are useful, accurate, and visually realistic. The work represents a step toward more intelligent and automated geospatial tools that can reduce human effort, lower costs, and increase accessibility. With further improvements and enhancements, systems like this could become standard tools for planners, researchers, educators, and even government agencies.

In conclusion, the study proves that Pix2PixGAN is not only capable of handling the task of map-to-satellite and satellite-to-map translation but also offers a practical framework that can be built upon in future research. It highlights the importance of deep learning in transforming the way geographical data is processed and used, paving the way for more advanced, scalable, and user-friendly applications in the future.

IX. FUTURE ENHANCEMENTS

Although the current system provides encouraging results and demonstrates the practical potential of Pix2PixGAN for translating maps into satellite images and vice versa, there are several directions in which it can be improved. These future enhancements will not only make the system more robust and reliable but will also broaden its usability across different domains and user groups.

One of the most significant improvements would be the integration of CycleGAN into the workflow. At present, the system depends on paired datasets, meaning that each map must have an exact satellite image counterpart. Collecting such datasets is time-consuming and sometimes impractical. CycleGAN can work with unpaired data, which would make the system more flexible and allow it to learn from larger and more diverse datasets. This enhancement would also help the system adapt to new regions where paired datasets are unavailable.

Another possible direction is the use of more advanced GAN architectures such as StyleGAN or SPADE. These models are known for producing highly realistic and detailed outputs in other computer vision applications. By integrating such models into the system, it may be possible to further improve the sharpness, color accuracy, and natural appearance of the generated images. This could make the outputs more reliable for professional use in fields such as urban planning, scientific research, and environmental monitoring.

A key area for improvement lies in image resolution. Currently, the system processes images at a fixed size (256×256 pixels) for efficiency in training and testing. While this is sufficient for research demonstrations, real-world applications often require high-resolution images. Future versions of the system should be able to handle larger and more detailed images without losing quality. Supporting higher resolutions would make the system more suitable for professional cartography, satellite analysis, and government-level planning.

Another enhancement would be to extend the system to incorporate three-dimensional terrain data. At present, the model focuses only on two-dimensional images. However, terrain information such as elevation and slope is critical in many real applications, including disaster management, construction planning, and environmental studies. Integrating 3D translation capabilities would allow the system to generate terrain-aware maps and satellite views, making it far more versatile.

In addition to improving technical features, future work should also focus on deployment and scalability. Running the system on cloud platforms would enable it to process large datasets more efficiently and serve multiple users simultaneously. This would be particularly valuable in real-time scenarios such as emergency response or live monitoring of urban development. Similarly, developing a mobile application would bring the system directly to end-users in the field, allowing planners, researchers, and students to access the functionality on handheld devices without needing advanced hardware.

Finally, another important enhancement is the introduction of explainable AI features. One of the common criticisms of GANs is that they work like a "black box," making it difficult for users to understand why a certain output was produced. By adding interpretability tools that explain how the model processes data and makes decisions, the system could increase user trust and become more acceptable in sensitive applications, such as government planning or academic research.

In summary, the future of this system lies in expanding its flexibility through unpaired learning, improving realism through advanced GAN architectures, supporting high-resolution and 3D data, scaling through cloud deployment and mobile apps, and increasing trust through explainable AI. Together, these enhancements would transform the system from a research prototype into a robust, scalable, and widely adopted solution for geospatial image translation.

X. REFERENCES

[1] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). *Generative Adversarial Nets*. Advances in Neural Information Processing Systems (NeurIPS), 27, 2672–2680.

[2] Mirza, M., & Osindero, S. (2014). *Conditional Generative Adversarial Nets*. arXiv preprint arXiv:1411.1784.

[3] Radford, A., Metz, L., & Chintala, S. (2016). *Unsupervised Representation Learning with Deep Convolutional GANs (DCGAN)*. International Conference on Learning Representations (ICLR).

[4] Isola, P., Zhu, J. Y., Zhou, T., & Efros, A. A. (2017). *Image-to-Image Translation with Conditional Adversarial Networks (Pix2Pix)*. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 5967–5976.

[5] Zhu, J. Y., Park, T., Isola, P., & Efros, A. A. (2017). *Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks (CycleGAN)*. IEEE International Conference on Computer Vision (ICCV), 2242–2251.

[6] Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., & Shi, W. (2017). *Photo-Realistic Single Image Super-Resolution Using a GAN (SRGAN)*. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 4681–4690.

[7] Ronneberger, O., Fischer, P., & Brox, T. (2015). *U-Net: Convolutional Networks for Biomedical Image Segmentation*. International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), 234–241.

[8] Park, T., Liu, M. Y., Wang, T. C., & Zhu, J. Y. (2019). *Semantic Image Synthesis with Spatially Adaptive Normalization (SPADE)*. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2337–2346.

[9] Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2018). *Progressive Growing of GANs for Improved Quality, Stability, and Variation*. International Conference on Learning Representations (ICLR).

[10] Karras, T., Laine, S., & Aila, T. (2019). *A Style-Based Generator Architecture for Generative Adversarial Networks (StyleGAN)*. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 4401–4410.

[11] Chen, Y., Wang, H., & Zhang, M. (2020). *Generative Adversarial Networks for Satellite Image Generation*. Remote Sensing, 12(1), 12–29.

[12] Ma, L., Liu, Y., Zhang, X., Ye, Y., Yin, G., & Johnson, B. (2019). *Deep Learning in Remote Sensing Applications: A Meta-Analysis and Review*. ISPRS Journal of Photogrammetry and Remote Sensing, 152, 166–177.

[13] Zhu, X., Tuia, D., Mou, L., Xia, G. S., Zhang, L., Xu, F., & Fraundorfer, F. (2017). *Deep Learning in Remote Sensing: A Comprehensive Review and List of Resources*. IEEE Geoscience and Remote Sensing Magazine, 5(4), 8–36.

[14] Marmanis, D., Schindler, K., Wegner, J. D., Galliani, S., Datcu, M., & Stilla, U. (2016). *Classification with an Edge: Improving Semantic Image Segmentation with Boundary Detection*. ISPRS Journal of Photogrammetry and Remote Sensing, 135, 158–172.

[15] Shermeyer, J., & Van Etten, A. (2019). *The Effects of Super-Resolution on Object Detection Performance in Satellite Imagery*. IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 1482–1491.

[16] Zhu, L., Sun, W., Huang, H., Deng, H., & Zhang, Z. (2020). *High-Resolution Remote Sensing Image Classification Using Transfer Learning and Deep CNNs*. Remote Sensing, 12(2), 298–313.

[17] Liu, Y., Yu, J., Han, L., & Wang, Z. (2018). *Automatic Road Extraction from Satellite Imagery Using Deep Learning and Conditional GANs*. IEEE Transactions on Geoscience and Remote Sensing, 56(6), 3431–3442.

[18] Xu, Y., Wu, L., Xie, Z., & Chen, Z. (2019). *Building Extraction in Very High-Resolution Remote Sensing Imagery Using Deep Learning and Conditional GANs*. Remote Sensing, 11(10), 1168.

[19] Kumar, A., Singh, P., & Gupta, R. (2021). *Map Generation from Satellite Images Using Deep Learning*. International Journal of Applied Earth Observation and Geoinformation, 103, 102445.

[20] Weng, Q., & Zhu, X. X. (2018). *Urban Remote Sensing: Monitoring, Synthesis, and Modeling in the Urban Environment*. Remote Sensing, 10(2), 234–249.