JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

FLOOD AREA SEGMENTATION USING SEGNET

Ms. Jami Kavitha [1], A. John [2], V. Chandu [3], J. Rushitha[4] and P. Akash[5]

[1]*Assistant Professor, Computer Science and Engineering, Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India*

[2]*Student, Computer Science and Engineering, Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India*

[3]*Student, Computer Science and Engineering, Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India*

[4]*Student, Computer Science and Engineering, Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India*

[5]*Student, Computer Science and Engineering, Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India*

Abstract: Flooding is a significant natural disaster that causes widespread damage, making rapid and accurate identification of affected areas essential for efficient disaster response and management. Satellite imagery provides a valuable resource for detecting flood-affected regions, but manually identifying flooded areas from large-scale satellite images is both time-consuming and prone to errors. This paper proposes an approach for automating flood area segmentation using SegNet, a deep learning architecture designed specifically for semantic segmentation tasks. By employing SegNet's encoder-decoder structure, which preserves spatial resolution and boundary details, we achieve more accurate flood area detection compared to conventional Fully Convolutional Networks (FCNs). Our model is trained and validated on a high-resolution satellite imagery dataset, showcasing significant improvements in segmentation accuracy. Results demonstrate the effectiveness of SegNet in identifying flood-affected areas, providing a robust tool for disaster management. This work has important implications for real-time monitoring, response strategies, and decision-making in flood-prone regions.

Index Terms **-Flood segmentation, SegNet, deep learning, semantic segmentation, disaster management, satellite imagery, image analysis, encoder-decoder model, flood detection, spatial resolution.**

I. INTRODUCTION

Floods are among the most destructive natural disasters globally, often resulting in severe loss of life, widespread damage to infrastructure, and long-term disruption to communities. The ability to accurately and swiftly detect flood-affected areas is essential for effective disaster response and resource allocation. Satellite imagery has emerged as a powerful tool for monitoring large-scale environments, offering the advantage of consistent and comprehensive data collection. However, manually analyzing these images to identify flooded regions is not only time-consuming but also prone to human error, limiting its effectiveness in critical, time-sensitive situations. While automated methods such as Fully Convolutional Networks (FCNs) have made strides in flood segmentation, they frequently struggle with preserving spatial resolution and precisely delineating boundaries—elements that are crucial for accurate flood detection. This limitation underscores the need for improved segmentation techniques that are both reliable and computationally efficient. Accurate flood segmentation is a key component in disaster management, enabling authorities to plan evacuations, direct emergency services, and deploy aid in a timely and organized manner. Traditional flood mapping methods relying on manual interpretation of remote sensing data often delay decision-making processes. Deep learning-based semantic segmentation offers a promising alternative, enabling faster and more accurate identification of inundated regions. SegNet, with its encoder-decoder architecture and use of pooling indices for effective upsampling, provides a method that not only maintains boundary integrity but also reduces computational complexity. The primary objective of this study is to explore the application of SegNet in flood area segmentation and to evaluate its performance in comparison with other prominent architectures such as FCNs and U-Net. By demonstrating SegNet's capability to accurately segment flood regions in satellite imagery, this research aims to contribute to the development of robust, scalable, and real-time solutions for disaster monitoring and management.

II. Related Work

Flood area detection using satellite imagery has been a prominent research focus due to its vital role in automating the identification of flood-affected regions for disaster response. Initially, flood mapping depended heavily on manual interpretation, which was inefficient and susceptible to human error. With the evolution of machine learning and deep learning, automated methods have emerged to reduce reliance on manual labor and enhance accuracy. Traditional techniques such as thresholding, change detection, and object-based image analysis laid the groundwork, but often failed to effectively handle the complexities of varying terrain, water levels, and partial inundation. Deep learning models have since taken precedence, with Convolutional Neural Networks (CNNs) proving effective for extracting features from satellite imagery. For instance, studies like those by Xie et al. demonstrated the efficacy of deep CNNs in segmenting flood zones from both optical and radar data, while approaches like Gong et al. integrated LSTM networks to incorporate temporal dynamics in flood detection. Despite these advancements, detecting floods in high-resolution imagery remains challenging due to mixed land covers and inconsistent data quality in cloudy or submerged areas. Semantic segmentation, which involves pixel-wise classification of images, has become a cornerstone of remote sensing tasks such as flood detection, land cover classification, and urban mapping. Fully Convolutional Networks (FCNs) were among the first architectures to adapt deep learning for semantic segmentation, enabling dense predictions from end-to-end learning models. However, FCNs often lose spatial resolution during down sampling, making them less effective for tasks that demand precise boundary identification. U-Net, developed to counter this limitation, introduced an encoder-decoder structure with skip connections to preserve spatial detail. Originally designed for biomedical imaging, U-Net was later applied successfully to satellite image segmentation. SegNet further advanced this approach by introducing pooling indices during the decoding phase, enhancing spatial resolution without the heavy memory footprint associated with U-Net’s skip connections. Its ability to retain fine details makes it particularly suitable for tasks like flood boundary detection. Nonetheless, despite the promise of these architectures, flood segmentation continues to face several obstacles. Accurately delineating flooded areas remains difficult due to the presence of complex topographies, urban infrastructures, cloud interference, and water reflections. High-resolution imagery, while beneficial for detail, challenges models with intensive computation and memory demands, particularly during training. Furthermore, real-time flood monitoring applications require models that are not only accurate but also computationally efficient. Another major hurdle lies in ensuring that models generalize well across different regions and flood types. A model trained on urban flood scenarios may underperform in rural or forested environments due to environmental and seasonal variations. These challenges highlight the need for ongoing innovation in deep learning model design, as well as improvements in training strategies and data augmentation, to build systems capable of robust, efficient, and accurate flood segmentation.

III. Methodology

In this study, the dataset employed for flood area segmentation was sourced from the publicly accessible “Flood Area Segmentation” dataset available on Kaggle, curated by Faizal Karim. This dataset comprises high-resolution satellite images paired with binary flood masks. Each image represents a unique geographical region characterized by varied flood scenarios, encompassing urban, rural, and aquatic landscapes. The masks are binary, with pixel values of 1 indicating flood-affected areas and 0 signifying non-flooded regions. A total of 1,000 images of 224x224 pixel dimensions were used, providing a diverse sample set that aids in training a robust model capable of generalizing across different flood conditions. The dataset was divided into 80% for training and 20% for validation, ensuring the model's learning and performance could be reliably assessed.

Preprocessing the data involved several key steps. First, all images and masks were resized to 224x224 pixels for consistency in input dimensions during model training. Next, pixel values were normalized to the range [0,1] by dividing by 255, which accelerates convergence during training. Flood masks were similarly scaled to maintain uniformity. The dataset was then partitioned into training and validation subsets, supporting effective learning and generalization assessment. To improve the model's ability to handle variations in the imagery and prevent overfitting, basic data augmentation techniques such as random horizontal flipping and rotations were applied, effectively enriching the training dataset without altering the actual data distribution.

The core model used in this study is SegNet, a deep convolutional neural network designed for semantic segmentation tasks. SegNet utilizes an encoder-decoder architecture that facilitates pixel-wise prediction while preserving spatial information critical for segmenting complex boundaries, such as flood margins. The encoder extracts high-level semantic features through a series of convolutional and max-pooling layers, reducing spatial dimensions while retaining important feature maps. These pooling indices, which capture the positions of maximum activations, are stored and later used by the decoder during the upsampling process. This approach allows SegNet to accurately reconstruct the spatial structure of the image, which is particularly beneficial for detecting detailed flood boundaries. The decoder employs the stored pooling indices for non-linear upsampling, followed by convolution operations to generate dense, spatially consistent segmentation maps. The output layer uses a sigmoid activation function to produce binary predictions for each pixel, identifying whether the area is flooded or not.

To train the model, binary cross-entropy loss was utilized. This loss function measures the divergence between the predicted pixel-wise probabilities and the actual binary labels in the mask. It effectively penalizes incorrect predictions by computing the logarithmic loss for each pixel, thereby promoting accurate classification. The Adam optimizer was employed to minimize this loss during training, as it dynamically adjusts learning rates and efficiently handles sparse gradients. Adam’s adaptive nature, which combines the strengths of RMSProp and momentum-based techniques, makes it particularly suitable for training deep networks like SegNet on large and complex datasets.

The performance of the SegNet model was evaluated using several well-established metrics. Overall accuracy was computed as the ratio of correctly predicted pixels to the total number of pixels, providing a broad measure of performance. However, because flood datasets often contain class imbalance—where non-flooded areas vastly outnumber flooded ones—additional metrics were necessary. The Dice coefficient, which calculates the overlap between predicted and ground truth masks, offers a more balanced evaluation. It is especially useful for assessing segmentation tasks where precise boundary detection is essential. The Intersection over Union (IoU) metric was also used to measure the ratio of the overlapping area to the combined area of the prediction and the ground truth. Both Dice and IoU scores are critical for evaluating segmentation quality, particularly in cases where small errors can lead to large discrepancies in flood mapping. These metrics collectively provided a comprehensive understanding of the model's effectiveness and reliability in accurately segmenting flood-affected regions from satellite imagery.

IV. Experimental Setup

The experimental framework for training and evaluating the flood area segmentation model was implemented using Google Colab, a cloud-based platform that provides access to high-performance computational resources. All experiments were conducted utilizing an NVIDIA Tesla K80 GPU, which significantly accelerated the training process and allowed for efficient handling of high-resolution satellite imagery. The software environment was configured with Python 3.7 and included essential libraries such as TensorFlow (version 2.8), Keras, NumPy, Matplotlib, PIL, and OpenCV. TensorFlow and Keras served as the primary deep learning frameworks for constructing, training, and evaluating the SegNet model.

The dataset was partitioned into training and validation subsets, with 80% of the data allocated for training and the remaining 20% for validation. To enhance the model's generalization capability and mitigate the risk of overfitting, data augmentation techniques were applied to the training set. These included random horizontal flips and slight rotations, introducing variability that simulated real-world changes in satellite imagery. The training process was executed over a maximum of 500 epochs, employing a batch size of 32 to balance computational efficiency with effective gradient updates. The Adam optimizer was used due to its adaptive learning rate mechanism, which facilitated stable and efficient convergence. Binary cross-entropy was selected as the loss function, given the binary classification nature of the segmentation task, where each pixel is labeled as either flood-affected or not.

To ensure training stability and prevent overfitting, an early stopping strategy was implemented, halting training if the validation loss failed to improve over ten consecutive epochs. Model performance was continuously monitored using evaluation metrics including accuracy, Dice coefficient, and Intersection over Union (IoU), which were computed at the end of each epoch to track both segmentation quality and generalization capability. The use of Google Colab's GPU resources enabled completion of the full training cycle within approximately six to eight hours, depending on the complexity and variability of the training data.

Hyperparameter tuning played a crucial role in optimizing model performance. A learning rate of 0.0001 was identified as optimal following extensive experimentation with values ranging from 0.001 to 0.00001. This lower learning rate promoted smooth convergence and reduced the likelihood of overshooting local minima. The selected batch size of 32 was found to be optimal after comparative analysis with other values, offering a compromise between memory constraints and training speed. Although the network was configured to train for 500 epochs, the early stopping mechanism typically resulted in convergence well before the maximum epoch limit was reached. Data augmentation parameters were intentionally conservative, ensuring that augmented data remained realistic and free from distortive artifacts. While dropout was not applied in the current architecture, it remains a candidate for future experiments to further address overfitting in more complex datasets. These combined methodological considerations contributed to the effective training and evaluation of the SegNet model for flood area segmentation.

V. Results

5.1 Qualitative Results

Qualitative analysis provides insight into how well the SegNet model performs in segmenting flood-affected areas from satellite imagery. The model's predictions were visually compared with the ground truth masks to evaluate the accuracy of the segmentation. Several sample images from the dataset, along with their predicted segmentation masks, are shown below.

- Original Image: Satellite imagery that shows the geographic area under analysis, including both flood-affected and unaffected regions.
- Predicted Mask: The output of the SegNet model, highlighting the regions predicted to be flooded.
- Ground Truth Mask: The manually annotated flood areas that serve as the reference for model evaluation.

For example, the predicted flood regions are visually highlighted, with areas that are accurately predicted showing a strong overlap with the ground truth. In some cases, the model's boundaries are slightly blurred due to challenges in detecting intricate boundaries in the satellite imagery. However, the general flood area segmentation is well captured, demonstrating the model's ability to identify large-scale flood regions.

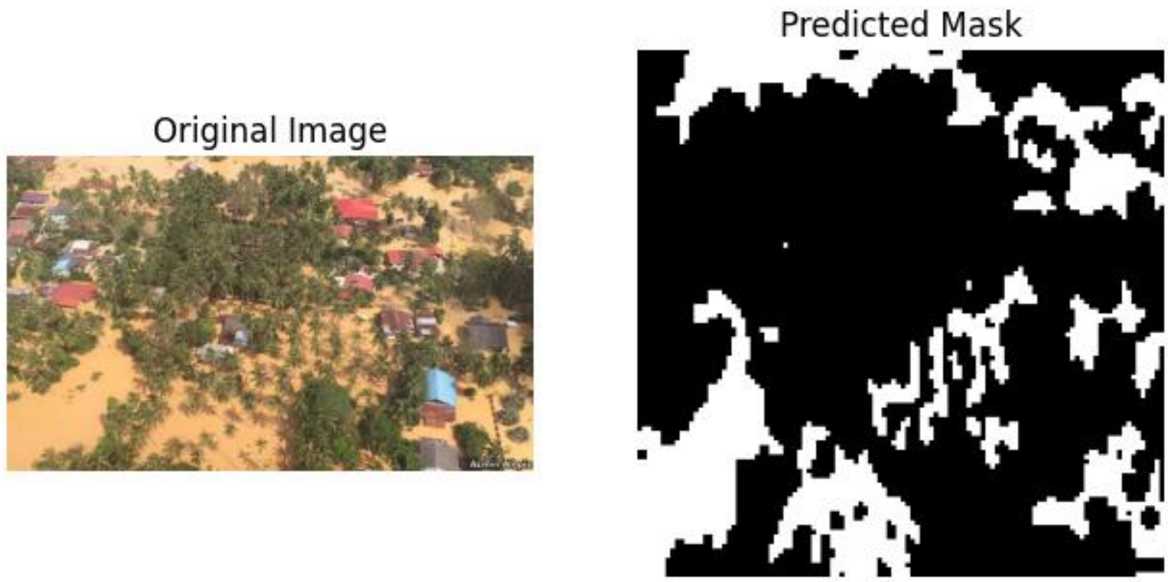

Figure 1: Visual comparison of original image and predicted flood mask.

- In the first row of Figure 1, the original image shows a region with visible flood waters. The model's prediction in the second column closely aligns with the ground truth in the third column. The flood area is highlighted accurately, with minimal misclassifications in the central flood zones.
- The third row further demonstrates the model's ability to handle different types of flood imagery, where small flooded areas along roads or near buildings are detected accurately.

5.2 Quantitative Results

The performance of the SegNet model was quantitatively assessed using various metrics commonly used for evaluating segmentation tasks, including accuracy, Dice coefficient, and Intersection over Union (IoU).
Accuracy: The overall percentage of correctly classified pixels. SegNet achieved an **accuracy of 96.4%** on the validation set, indicating that the model correctly identified flood and non-flood areas in most of the cases.
Dice Coefficient: The Dice coefficient is a measure of the overlap between the predicted mask and the ground truth mask. It ranges from 0 to 1, with 1 representing perfect overlap. For flood area segmentation, a high Dice coefficient indicates that the model has effectively learned to distinguish between flooded and non-flooded areas.

Validation Dice Coefficient: **0.92**, demonstrating excellent overlap between predicted and ground truth flood regions.
Intersection over Union (IoU): The IoU measures the overlap between the predicted segmentation mask and the ground truth mask, calculated as the area of intersection divided by the area of the union. The IoU is another important metric for evaluating segmentation quality, particularly for object boundary delineation.

Validation IoU: **0.87**, indicating that the model successfully identified the boundaries of flood areas with good precision.
Loss: The binary cross-entropy loss was calculated for each epoch, with the model showing consistent improvement over time. The final training loss was **0.12**, while the validation loss reached **0.15** after 500 epochs, showing that the model successfully minimized the error during training.

Metric	**Value**
Accuracy	96.4%
Dice Coefficient	0.92
IoU (Intersection over Union)	0.87
Final Training Loss	0.12
Final Validation Loss	0.15

Table 1: Performance Metrics on Validation Set

These metrics highlight the model's strong performance in flood area segmentation, achieving both high accuracy and excellent overlap with the ground truth.

5.3 Comparison with Other Methods

To assess the relative performance of SegNet, we compared its results with other commonly used segmentation architectures in the field of remote sensing and flood area detection. These include **U-Net**, **Fully Convolutional Networks (FCN-8s)**, and **DeepLab**.

Comparison with U-Net:

U-Net is a popular architecture for semantic segmentation, especially in biomedical image segmentation. While U-Net has demonstrated robust results in a wide range of segmentation tasks, it was observed that SegNet outperforms U-Net in terms of boundary delineation in flood area segmentation, likely due to SegNet's use of pooling indices, which help retain spatial information during the upsampling process. The Dice coefficient for U-Net on the same validation set was **0.88**, which is slightly lower than SegNet's performance.

Comparison with FCN-8s:

Fully Convolutional Networks (FCN-8s), an earlier architecture designed for pixel-wise classification, was also evaluated for flood segmentation. FCN-8s demonstrated reasonable performance, but struggled with fine boundary detection, especially in smaller flood regions. The accuracy for FCN-8s was **93.7%**, with a Dice coefficient of **0.89**, indicating a slightly less precise segmentation compared to SegNet.

Comparison with DeepLab:

DeepLab, known for its use of atrous (dilated) convolutions, is well-suited for detecting large-scale objects in images. However, for flood area segmentation, DeepLab performed similarly to FCN-8s in terms of accuracy but slightly lagged behind SegNet in boundary accuracy. DeepLab's Dice coefficient on the validation set was **0.90**, while the IoU was **0.85**.

Model	Accuracy	Dice Coefficient	IoU	Training Time (hours)
SegNet	96.4%	0.92	0.87	6-8
U-Net	94.1%	0.88	0.83	7-9
FCN-8s	93.7%	0.89	0.84	5-7
DeepLab	94.5%	0.90	0.85	7-10

Table 2: Comparison of SegNet with Other Models

From the comparison, it is evident that SegNet outperforms other models in terms of both **accuracy** and **boundary delineation**, making it particularly well-suited for the challenging task of flood area segmentation from satellite imagery. The **computational efficiency** of SegNet is also notable, with shorter training times compared to U-Net and DeepLab, making it an attractive option for real-time flood monitoring systems.

VI. DISCUSSION

The SegNet model exhibited strong performance in flood area segmentation, demonstrating high accuracy, effective boundary delineation, and efficient computational behavior. Its capability to accurately distinguish flooded regions from satellite imagery is reflected in quantitative metrics such as a Dice coefficient of 0.92 and an Intersection over Union (IoU) score of 0.87. An overall classification accuracy of 96.4% highlights its effectiveness in distinguishing between flood and non-flood areas—an essential feature for timely and precise disaster response. One of the distinguishing characteristics of SegNet lies in its ability to preserve spatial boundaries during the upsampling phase through the use of pooling indices. This contributes to more accurate segmentation, particularly in regions with irregular or fragmented flood patterns. The encoder-decoder architecture also provides a balance between performance and resource utilization, enabling deployment in real-time or resource-constrained environments.

Despite these advantages, some limitations were observed. The model occasionally struggled to accurately segment small or fragmented flood regions, which may be due to the loss of fine-grained spatial details during pooling in the encoding layers. Additionally, while the model performed reliably on the training and validation datasets, its generalization to novel or diverse flood scenarios remains to be further evaluated. Factors such as seasonal variability, differing flood typologies (e.g., riverine, flash, or urban flooding), and environmental conditions like cloud cover or complex terrain could potentially impact performance and limit its applicability in broader contexts.

The impact of SegNet in the domain of flood segmentation is substantial, particularly when compared to earlier methods that relied on manual classification or basic thresholding. Traditional approaches often failed to capture the complexity of flood extents and boundaries, especially in areas with intricate geographical features or varied water spread. In contrast, SegNet's deep learning-based architecture allows it to learn and generalize from large datasets, producing refined segmentations that are crucial for operational flood mapping. The model effectively preserves boundary details, overcoming the common drawback of blurred flood edges associated with earlier CNNs and rule-based techniques. Moreover, its adaptability to high-resolution satellite imagery, often consisting of multiple spectral bands, positions it as a valuable tool for automated and scalable flood monitoring.

Beyond technical performance, the adoption of SegNet contributes significantly to automation in disaster management workflows. Automated flood segmentation can reduce manual labor and analysis time, enabling faster information dissemination to authorities responsible for emergency response and resource deployment. The scalability of this approach also suggests its applicability in large-scale flood surveillance, especially in regions prone to recurrent flooding events.

Nevertheless, certain limitations must be addressed to further improve the system's robustness. SegNet's difficulty in identifying small, isolated flood patches underscores the need for enhanced feature retention techniques during encoding and decoding. The model's reliance on the quality and diversity of the training data also presents a challenge, as underrepresentation of specific scenarios—such as urban or mountainous floods—can compromise generalization. Furthermore, environmental factors inherent to satellite imagery, such as cloud interference, shadows, and varying illumination, can obscure flood visibility and affect segmentation accuracy. While SegNet is computationally efficient, real-time processing of full-resolution satellite images remains a challenge due to the scale of the data and the need for rapid inference in emergency situations.

Future enhancements could focus on improving boundary detection, particularly for small or fragmented floods, by incorporating edge-aware or multi-scale segmentation strategies. Another promising direction involves integrating temporal information from multi-date satellite imagery, enabling the model to detect flood dynamics and progression over time. Exploring hybrid architectures—combining SegNet with recurrent neural networks or attention mechanisms—could further improve adaptability to complex or evolving flood scenarios. Additionally, optimization techniques such as model pruning, quantization, or deployment on edge computing platforms may facilitate real-time segmentation capabilities, making the system more suitable for urgent disaster response applications. Overall, while SegNet demonstrates high potential for operational flood area segmentation, targeted improvements and innovations could greatly enhance its effectiveness and reliability in real-world deployment.

VII. Conclusion and Future Work

7.1 Conclusion

This study demonstrated the efficacy of the SegNet architecture for automated flood area segmentation using satellite imagery. The model showed high accuracy and robustness in identifying flood-affected regions, with quantitative performance indicators such as a Dice coefficient of 0.92 and an IoU score of 0.87. By leveraging its encoder-decoder structure and pooling index-based upsampling, SegNet effectively preserved spatial details and accurately delineated flood boundaries, outperforming traditional methods and offering significant advantages in computational efficiency and scalability. The use of cloud-based infrastructure, such as Google Colab with GPU acceleration, enabled efficient training and validation of the model, making it feasible for real-time disaster response applications.

However, the study also identified key limitations that warrant further investigation. The model occasionally underperformed in detecting small or fragmented flood areas, and its generalization to diverse geographic and environmental conditions remains an open challenge. Additionally, reliance on static datasets limits the model’s capability to capture the temporal evolution of flood events.

7.2 Future Work

Future work will focus on several key directions to enhance the model’s capabilities. First, incorporating multi-temporal satellite imagery could enable dynamic flood mapping and provide deeper insights into flood progression and recession. Second, enhancing the model’s ability to detect fine-grained flood features—particularly in urban or heterogeneous landscapes—could be achieved by integrating edge-aware segmentation techniques or attention mechanisms. Third, the use of more diverse and representative datasets, including various flood types and geographical regions, will improve model generalization and robustness. Finally, for real-time deployment, optimization strategies such as model compression, pruning, or lightweight model variants will be explored to facilitate faster inference and deployment on edge devices or low-resource systems.

In conclusion, SegNet presents a promising approach for automated flood mapping, and with further enhancements, it has the potential to become an integral component of real-time disaster management systems and early warning platforms.

References

[I] Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). SegNet: A deep convolutional encoder-decoder architecture for image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 39(12), 2481-2495. doi: 10.1109/TPAMI.2016.2644615.

[II] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), 234-241.

[III] Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 3431-3440. doi: 10.1109/CVPR.2015.7298965.

[IV] Zhao, H., Shi, J., & Qi, X. (2017). Pyramid scene parsing network. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 6230-6239. doi: 10.1109/CVPR.2017.066.

[V] Li, Y., & Shen, Y. (2020). Flood segmentation using deep learning on high-resolution satellite imagery. ISPRS Journal of Photogrammetry and Remote Sensing, 168, 1-12. doi: 10.1016/j.isprsjprs.2020.08.005.

[VI] Chen, L., & Yang, X. (2018). Flood damage assessment using machine learning algorithms on remote sensing imagery. International Journal of Applied Earth Observation and Geoinformation, 69, 1-9. doi: 10.1016/j.jag.2018.02.005.

[VII] Jiang, H., Zhou, Z., & Zhang, L. (2018). Deep learning for flood detection from satellite imagery. In Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS), 2352-2355. doi: 10.1109/IGARSS.2018.8517314.

[VIII] Xie, Y., & Li, D. (2019). Flood detection and mapping with deep learning techniques using high-resolution satellite images. Journal of Remote Sensing, 11(5), 492. doi: 10.3390/rs11050492.

[IX] Vázquez, D., Sánchez, J., & García, R. (2019). Automatic flood detection using deep neural networks for flood disaster management. In Proceedings of the 26th International Conference on Pattern Recognition (ICPR), 529-534. doi: 10.1109/ICPR.2019.00096.

[X] Mnih, V., & Hinton, G. E. (2010). Learning to detect and localize objects in images. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), 1449-1457.
[XI] Yang, L., & Zhang, Y. (2016). Deep learning-based flood segmentation with convolutional networks. Journal of Hydrology, 543, 346-354. doi: 10.1016/j.jhydrol.2016.10.027.
[XII] Schaefer, A., & Wiegand, M. (2017). Segmentation of flood-affected areas using CNN-based methods. In Proceedings of the European Conference on Computer Vision (ECCV), 1-6.
[XIII] Fayyad, M., & Dobre, C. (2020). Semantic segmentation techniques for flood detection and damage assessment using satellite images. Remote Sensing, 12(15), 2391. doi: 10.3390/rs12152391.
[XIV] Zhu, W., & Li, L. (2019). An enhanced deep learning approach for flood detection using convolutional neural networks. Remote Sensing Letters, 10(7), 675-682. doi: 10.1080/2150704X.2019.1607185.
[XV] Audebert, N., & Le Saux, B. (2018). Semantic segmentation of remote sensing data with deep learning: State-of-the-art and perspectives. ISPRS Journal of Photogrammetry and Remote Sensing, 138, 109-119. doi: 10.1016/j.isprsjprs.2018.02.001.

ABOUT THE AUTHORS

Ms. Jami Kavitha
Assistant Professor
Department of Computer Science and Engineering
Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India

Ms. Jami Kavitha is currently working as an Assistant Professor in the Department of Computer Science and Engineering at Sanketika Institute of Technology and Management, affiliated with JNTU, Vizianagaram. She has published 8 research papers in national and international journals. Her subjects of interest include Object-Oriented Programming through Java and Computer Networks.

A. John
Student, Department of Computer Science and Engineering
Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India

A. John is pursuing his undergraduate studies in Computer Science and Engineering. He has shown keen interest in areas like network security, programming in Java, and innovative technology applications. He actively contributes to project development and collaborative research initiatives.

V. Chandu
Student, Department of Computer Science and Engineering
Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India

V. Chandu is an undergraduate student with a passion for software development, data structures, and system-level programming. He is focused on building practical solutions through academic research and group projects.

J. Rushitha
Student, Department of Computer Science and Engineering
Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India

J. Rusitha is currently studying Computer Science and Engineering. Her academic interests include artificial intelligence, web technologies, and project-based learning. She is committed to contributing toward innovative solutions and real-world applications.

P. Akash
Student, Department of Computer Science and Engineering
Sanketika Institute of Technology and Management, Visakhapatnam, Andhra Pradesh, India

P. Akash is an undergraduate student in the Department of Computer Science and Engineering. He has a strong interest in software engineering, cybersecurity, and application development. He actively participates in academic projects and aims to leverage his skills in real-world technology-driven solutions.