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

SEA ICE DETECTION USING SATELLITE IMAGES

[1]Ayishath Nafiya B, [2]Ambili M, [3]Gayathri T,[4]Dr. Parvathraj K M M, [5]Prof. Daya Naik
[1,2,3] Student, [4,5]Associate Professor
[1,2,3,4,5]Artificial Intelligence and Machine Learning,
[1,2,3,4,5]Srinivas Institute of Technology, Mangalore, India

Abstract : Monitoring sea ice is vital for understanding climate change, ensuring safe marine navigation, and supporting polar research. This project presents an AI-based system that detects and classifies sea ice concentrations using Sentinel-2 satellite imagery focused on the Hudson Bay region. By employing a U-Net convolutional neural network, the model performs detailed image segmentation to categorize sea ice into meaningful classes. The workflow includes automated data acquisition, preprocessing to manage cloud cover and class imbalance, and the use of augmented data to strengthen model learning. The model achieves reliable performance with 83% accuracy and a mean Intersection over Union (IoU) of 0.44. A web-based interface has been developed to make this system accessible for real-time chart generation. The approach demonstrates a scalable and adaptable framework that enhances existing methods of sea ice analysis and supports environmental monitoring and operational decision-making in polar regions.

***IndexTerms* -** Sea Ice Segmentation, Sentinel-2 Imagery, U-Net Architecture, Satellite Image Processing, Climate Monitoring, Deep Learning, Convolutional Neural Network, Arctic Navigation, Image Classification, Environmental Mapping.

I. INTRODUCTION

Sea ice is crucial to the global climate system, affecting ocean circulation, weather patterns, and ecosystems. Monitoring sea ice concentration accurately is vital for comprehending climate change, ensuring maritime safety, and facilitating environmental oversight. Conventional methods for producing sea ice charts depend on the manual interpretation of geospatial data, which is labor-intensive and limited in scope. With the increasing availability of satellite imagery, there is a rising demand for automated and scalable solutions to analyze and map sea ice concentrations with enhanced detail and efficiency. This project aims to develop a model based on convolutional neural networks (CNN) to generate sea ice concentration charts from Sentinel-2 satellite imagery, focusing specifically on the Hudson Bay region. By employing advanced image segmentation techniques, such as the U-Net architecture, the model provides pixel-level classifications for ice concentration. This method not only simplifies the chart generation process but also enhances spatial resolution and accuracy compared to traditional manual techniques. The project addresses challenges such as class imbalances in datasets and improves data collection processes, striving to deliver a robust and scalable solution. The overarching goal of this initiative extends beyond Hudson Bay, aiming for global applications. By implementing a systematic and automated workflow, this framework could produce detailed and precise ice charts for other Arctic and Antarctic areas. These innovations have the potential to significantly improve environmental research, maritime operations, and policy-making concerning polar regions, thereby deepening the understanding of sea ice dynamics in a changing climate.

A. Background

The polar regions play a crucial role in the Earth's climate system, with sea ice acting as a key component that enables heat exchange between the ocean and the atmosphere. In recent years, the rapid melting of sea ice attributed to climate change has raised concerns about rising sea levels, altered ocean circulation patterns, and disrupted ecosystems. It is essential to monitor sea ice concentration to comprehend these changes and their implications. Accurate and timely sea ice data is also vital for maritime operations, as vessels navigating Arctic waters depend on reliable ice charts for safe passage. Traditional methods for creating sea ice charts involve the manual analysis of geospatial data, such as that provided by the Canadian Ice Service. While these methods have produced valuable insights, they are labor-intensive, prone to errors, and limited in terms of spatial resolution and scalability. The increasing availability of satellite imagery from missions like Sentinel-2, which offers high-resolution optical data on a global scale, presents a significant opportunity to automate this process, thereby improving its efficiency and accuracy. Advances in deep learning, particularly in image segmentation, offer a powerful solution to this challenge. Convolutional neural networks (CNNs) like U-Net have demonstrated exceptional effectiveness in categorizing images, making them particularly well-suited for sea ice detection tasks. This project aims to leverage these technologies to create an automated, scalable system for generating sea ice charts. By combining satellite imagery and geospatial data within a deep learning framework, the project aspires to transform sea ice monitoring, addressing both regional and global needs.

B. Problem Statement

The monitoring of sea ice is essential for comprehending the effects of climate change, ensuring maritime safety, and aiding environmental conservation initiatives. Presently, the generation of sea ice concentration charts depends on the manual analysis of geospatial data, such as the Canadian Regional Ice Charts. These techniques are labor-intensive, susceptible to human error, and constrained in their spatial resolution. Furthermore, the ever-changing and seasonal characteristics of sea ice pose difficulties in obtaining precise and timely data, particularly in areas like Hudson Bay, which undergo significant seasonal fluctuations. Given the growing accessibility of high-resolution satellite imagery, there is a pressing need for automated, scalable, and precise tools to evaluate and map sea ice concentrations. Current machi ne learning approaches for analogous tasks frequently encounter issues such as class imbalance (for instance, the overrepresentation of thick ice during winter) and the effective integration of various data sources. This project seeks to address these challenges by utilizing a convolutional neural network to automatically segment satellite images and forecast sea ice concentrations, thereby offering a more efficient and accurate method for sea ice charting.

C. Objectives

This project seeks to create an automated system for producing high-resolution sea ice charts utilizing deep learning methodologies, particularly the U-Net architecture. The system will execute pixel-wise segmentation of Sentinel-2 satellite images to categorize various ice concentration levels, thereby replacing traditional manual charting methods with a more scalable and efficient alternative. The use of high-resolution imagery, along with alignment to ice chart shapefiles, will guarantee both spatial and temporal precision, while also addressing dataset imbalances resulting from seasonal fluctuations through data augmentation and oversampling strategies. The model's performance will be improved by integrating advanced techniques such as dropout layers and incorporating additional data from less represented seasons, aiming for an accuracy exceeding 80% and a mean IoU score above 0.4. This scalable workflow will encompass preprocessing pipelines, the integration of Sentinel-2 spectral bands, and the capability to adapt to both Arctic and Antarctic regions, ensuring versatility and robustness in sea ice detection. A significant outcome of this project will be a user-friendly web application that enables users to upload satellite images and obtain real-time sea ice charts. This prototype will undergo validation through comparisons with existing charts and real-world data, confirming its reliability and accuracy. Additionally, the system is intended to be applicable on a global scale, facilitating long-term monitoring of sea ice trends and aiding climate research initiatives. By automating the creation of high-resolution ice charts, this project plays a vital role in environmental monitoring and offers critical insights into global climate dynamics.

II. LITERATURE SURVEY:

Sea ice monitoring is a crucial aspect of climate research, playing a significant role in regulating global temperatures, ocean circulation, and providing a habitat for marine ecosystems. Its reflective properties influence the Earth's albedo, impacting the global energy balance. Traditional methods of monitoring, such as manual interpretation of geospatial data, have been invaluable but are limited in terms of scope and efficiency. The emergence of satellite remote sensing technologies, including instruments like Sentinel-2 and MODIS, has revolutionized this field by offering detailed, high-resolution datasets. These advancements enable more comprehensive analysis of ice concentration, extent, and thickness, though manual processing remains a bottleneck due to its labor-intensive nature. Machine learning, particularly deep learning, has emerged as a transformative approach to automating tasks in remote sensing. U-Net, a deep learning model specifically designed for image segmentation, has proven highly effective in generating pixel-wise classifications of satellite images. This model's ability to learn spatial and contextual information makes it particularly suitable for detailed sea ice analysis. However, challenges such as cloud interference, atmospheric noise, and seasonal variations in ice coverage pose significant hurdles. Addressing these issues requires robust preprocessing techniques, including cloud masking, data augmentation, and stratified sampling, to enhance model performance and balance datasets. The applications of automated sea ice charting are vast, ranging from safer maritime navigation in ice prone regions to aiding climate research with real time, high-resolution data. Evaluating segmentation models with metrics like accuracy and Intersection over Union (IoU) ensures reliability, with U-Net models often achieving IoU scores of 0.4 to 0.6. Expanding the use of additional satellite bands, such as radar and thermal imaging, along with incorporating underrepresented seasonal data, can further enhance model accuracy and scalability. By integrating these advancements into user-friendly platforms, automated sea ice monitoring systems can provide global applicability, supporting climate research and policy-making efforts.

III. METHODOLOGY:

The methodology for this project involves a systematic approach to data collection, preprocessing, model development, evaluation, and deployment. Each step ensures that the system accurately detects and maps sea ice concentrations from satellite imagery.

A. Data Collection

The data collection process forms the foundation of this project, focusing on acquiring and preparing satellite imagery and geospatial ice concentration data for sea ice segmentation. Sentinel-2 satellite images were chosen as the primary data source due to their high spatial resolution (10 meters) and availability of multispectral bands. Specifically, bands 3 (green), 4 (red), and 8 (near-infrared) were used, as they effectively capture the spectral properties of sea ice, distinguishing it from open water and land. The images were obtained using the Sentinelhub API, which provided a streamlined approach to accessing this data. The Hudson Bay region was selected as the area of interest due to its significant seasonal variability in sea ice coverage and the availability of corresponding labeled datasets. In addition to satellite imagery, labeled sea ice concentration data was sourced from the Canadian Regional Ice Charts provided by the National Snow and Ice Data Center (NSIDC). These charts, which detail ice concentrations based on the SIGRID-3 classification, were used to create pixel-wise masks 3 | P a g e for training and validation. The dataset spanned from January 2016 to July 2018, capturing multiple seasonal cycles of ice formation and melting. By aligning the temporal data of the satellite images with the weekly updates of the ice charts, the project ensured consistency between observed satellite features and their labelled counterparts. This temporal synchronization was critical for generating reliable training data.

The EO-Learn Python library was employed to automate data collection and preprocessing workflows. The Hudson Bay region was divided into smaller geographical tiles to facilitate easier processing. Satellite images were filtered based on cloud cover, and a custom workflow was developed to exclude pixels obstructed by clouds or atmospheric interferences. The final dataset consisted of image-mask pairs, with each pixel categorized into one of eight classes: open water, intermediate ice concentrations, thick ice, fast ice, or land. Challenges such as class imbalances, resulting from seasonal ice coverage variations, were addressed using data augmentation techniques like flipping, rotation, and over-sampling underrepresented classes. This comprehensive and systematic data collection process laid a robust groundwork for training the sea ice segmentation model.

B. Data Preprocessing

3.2. Data Preprocessing In order to simplify the analysis of sea ice concentrations, the original 31 SIGRID-3 classes were consolidated into 8 broader categories. These 8 classes represent varying levels of ice concentration and specific geographic features, including open water and land. The class definitions are as follows: 0 for less than 10% ice, 1 for 10-30% ice, 2 for 30 50% ice, 3 for 50-70% ice, 4 for 70-90% ice, 5 for 90-100% ice, 6 for fast ice (thick ice that is attached to the coastline), and 7 for land. The class distribution across the dataset revealed a strong class imbalance, with the majority of the images being dominated by open water, 90-100% ice, and land. This imbalance is visually represented in the class distribution chart, highlighting the challenge of training a model to accurately predict less common classes, such as intermediate ice concentrations.

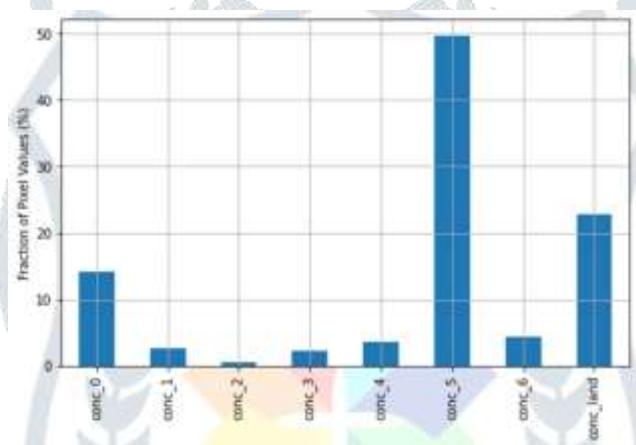

Fig:1 Pixel-wise class distribution over all images

To address this imbalance and optimize model performance, a structured data input pipeline was implemented. The dataset was first split into training and validation sets, with 80% of the images allocated for training and 20% for validation. The split was stratified to ensure that the most common class in each image was well-represented in both datasets. Additionally, within the training set, images that contained under-represented classes were over sampled to give these classes greater weight during training. This was achieved by duplicating images with high amounts of under-represented pixels, helping the model learn to classify these less common classes more effectively. To further enhance the robustness of the model and reduce overfitting, random image augmentation techniques were applied. These included random left-right and up down flips, as well as random rotations of up to ±5 degrees. In cases where the image was rotated, the corners were mapped to black, and the corresponding land areas in the masks were adjusted accordingly. This augmentation approach helped improve the generalization of the model by presenting it with a more diverse range of training images, as illustrated by the sample image-mask pair shown below.

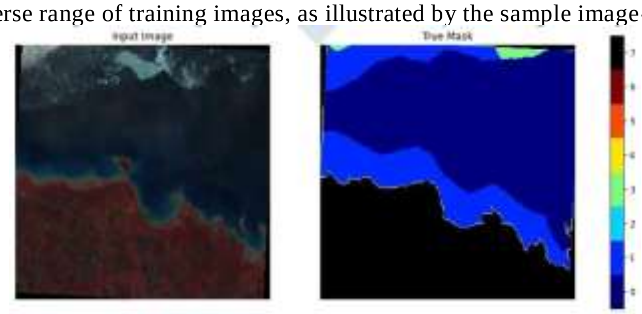

Fig:2 Sample image/mask pair training data

C. Model Development

- **Model Architecture**: A U-Net convolutional neural network was chosen for its proven ability in image segmentation tasks. The architecture was adapted with dropout layers to prevent overfitting.
- **Training Process**: The model was trained on 80% of the dataset, with 20% reserved for validation. The training employed a stratified sampling strategy to balance class representation.
- **Evaluation Metrics**: Performance was assessed using accuracy and mean Intersection over Union (IoU), with a target of achieving at least 83% accuracy and a mean IoU score of 0.44.

IV. ALGORITHM DESIGN:

The algorithm developed for this project utilizes a Convolutional Neural Network (CNN), specifically employing a U-Net architecture, to classify satellite images of Hudson Bay into eight distinct categories that reflect different levels of sea ice concentration. The preprocessing workflow includes dividing the geographical area into manageable tiles through EO-Learn and aligning satellite images with relevant ice charts to generate accurate segmentation masks. The dataset comprises 3,392 images, organized into eight categories, which were balanced using oversampling methods and stratified splitting to mitigate class imbalances. To improve the model's robustness, data augmentation techniques such as random flips and rotations were implemented. and rotations, was applied to enhance model robustness. The U-Net model features a contraction-expansion architecture, enabling detailed pixel-level classification by combining convolutional layers, max-pooling, and upsampling. The performance of the segmentation model was evaluated using Intersection over Union (IoU), calculated as:

$$[illegible] = \frac{[illegible]}{[illegible]}$$

A. Data Preprocessing and Augmentation:

The preprocessing phase involved careful handling of satellite imagery to ensure quality data for training the model. EO-Learn was used to slice large geographical areas into smaller, manageable tiles (EOPatches), ensuring the satellite images were processed in a spatially efficient manner. Each image was synchronized with the corresponding ice chart, which provided precise sea ice concentration data, enabling the creation of accurate segmentation masks for training. Data augmentation techniques, such as random flips and rotations, were applied to enhance the robustness of the model and prevent overfitting, particularly in regions where certain classes (like intermediate ice concentrations) were underrepresented.These preprocessing and augmentation steps significantly contributed to improving model generalization, especially for handling the class imbalance observed in the dataset.

B. Future Opportunities and Improvements:

While the current U-Net model demonstrates promising results, there is still room for further enhancement. The notable class imbalance, especially in areas with dense ice and open water, could be addressed by collecting more data during transitional seasons, such as spring or early summer, when a broader range of ice concentrations is present. Additionally, incorporating extra satellite data from various wavelength bands, including thermal infrared, may improve the model's ability to detect subtle changes in ice concentration. Further refinements to the U-Net architecture, such as the incorporation of attention mechanisms or hybrid models that combine deep learning with traditional image processing methods, could lead to even more accurate and detailed segmentation of sea ice. These advancements could enable the development of automated, high-resolution ice charts that exceed the precision of existing manual methods, ultimately aiding climate research and enhancing maritime safety.

V. MODEL BUILDING:

The model developed for this project is built upon the architecture, a well-established convolutional neural network (CNN) design specifically suited for image segmentation tasks. U Net consists of two primary components: a contraction path and an expansion path. The contraction path uses successive convolutional layers, activation functions, and max-pooling operations to capture increasingly abstract features of the image. This is followed by the expansion path, where up-sampling and concatenation with high resolution features from the contraction path help the network reconstruct the spatial dimensions of the image, pixel by pixel. This structure allows U-Net to effectively capture fine details and accurately classify each pixel, making it highly effective for tasks like semantic segmentation. In the case of sea ice detection, U-Net can identify the boundaries and varying concentrations of ice in satellite imagery, providing detailed and accurate predictions. For this specific project, the base U-Net model was adapted from the Dstl Satellite Imagery Feature Detection Kaggle competition, which similarly focused on pixel-wise classification of satellite images. To enhance performance and reduce overfitting, dropout layers were incorporated into the model. Dropout helps the network generalize better by randomly omitting certain neurons during training, thus preventing the model from over relying on specific features. The network was trained over 100 epochs, with the goal of minimizing the loss function and maximizing the Intersection over Union (IoU) metric, which is commonly used to evaluate segmentation accuracy. The model's performance was closely monitored through training and validation loss curves, as well as a confusion matrix, highlighting its strength in predicting open water and land, though it faced challenges with intermediate ice concentrations.

A. U-Net

U-Net is a widely used convolutional neural network (CNN) architecture for image segmentation tasks, particularly in medical imaging and remote sensing applications. The architecture consists of two main paths: a contraction path and an expansion path. The contraction path is responsible for extracting features from the input image through a series of convolutional layers, ReLU activations, and maxpooling operations. This part of the network progressively reduces the spatial dimensions of the image, while increasing the depth of feature maps, capturing increasingly abstract features. The expansion path, on the other hand, employs upsampling and concatenation with high-resolution features from the contraction path to gradually reconstruct the spatial dimensions of the output. The final output is a pixel-wise classification of the input image. This design ensures that U-Net performs well in both localization (determining the spatial location of objects in the image) and classification tasks, making it particularly effective for problems like semantic segmentation. In the context of this project, U-Net's architecture is well-suited to the task of classifying sea ice, as it allows for fine-grained predictions of ice concentration at the pixel level.

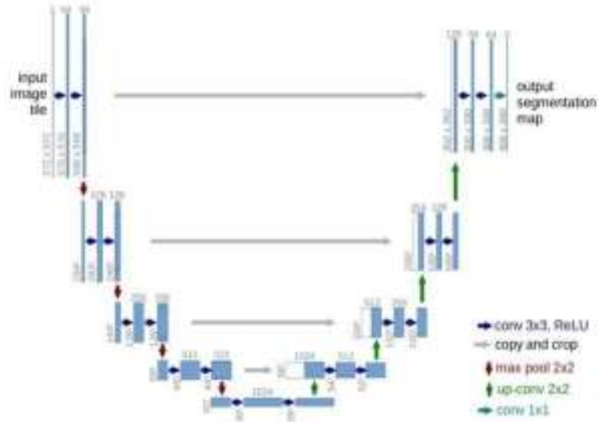

=Fig:3 Base U-Net Architecture

B. Model definition

The model created for this initiative is founded on an adapted version of the U-Net architecture, originally utilized in the Dstl Satellite Imagery Feature Detection competition on Kaggle. This competition aimed to classify pixels within satellite images, making the U-Net architecture an appropriate selection for the task. The competition model was modified to meet the specific requirements of sea ice detection by integrating dropout layers to reduce overfitting. Dropout serves to prevent the network from becoming overly dependent on any single feature during training, thereby enhancing its ability to generalize to new, unseen data. Additionally, the architecture was further refined for the task through parameter fine-tuning and the application of techniques such as data augmentation and regularization to boost performance. The model was crafted to predict the classification of each pixel in a satellite image, categorizing it into various ice concentrations or other pertinent classes, including land or open water. This methodology guarantees that the model can accurately capture the complex details of the sea ice environment, resulting in precise and detailed ice charts.

C. Training and Predictions

The model was trained over 100 epochs, where an approach refers to a full pass through the training dataset. The goal during training was to minimize the loss function, which quantifies the difference between the predicted and actual pixel classifications. In addition to the loss, the model's performance was evaluated using the Intersection over Union (IoU) metric, also known as the Jaccard index. IoU is a popular metric for evaluating segmentation tasks, as it measures the overlap between the predicted and true pixel classes. The training and validation loss, as well as the mean IoU, were tracked over the course of the training process. The results showed that the model achieved the lowest validation loss around the 50th epoch, after which it started to overfit the training data. The model's weights were saved at this point to ensure the best performance on unseen validation data. A confusion matrix was also generated to assess the model's ability to classify different ice concentrations. The confusion matrix highlighted the model's strong performance in predicting open water and land, but also indicated that it struggled with intermediate ice concentrations (10-90%). This suggests that the model would benefit from additional data specifically focused on these intermediate ice concentrations, which are more prevalent during transitional periods such as spring and early summer when the ice is melting. The validation results, including true and predicted masks for several test images, further illustrated the model's ability to capture the general structure of sea ice, even though there was a class imbalance in the dataset, with land and solid ice being more frequent. Despite this imbalance, the model was able to generate relatively accurate predictions of localized ice concentrations in many cases. Interestingly, the model was also able to provide finer details compared to published ice charts, suggesting that with additional training data, this approach could potentially lead to more detailed and accurate ice charts than are currently available. This could have significant implications for improving the understanding of sea ice dynamics and for better supporting operational applications such as navigation and climate modeling.

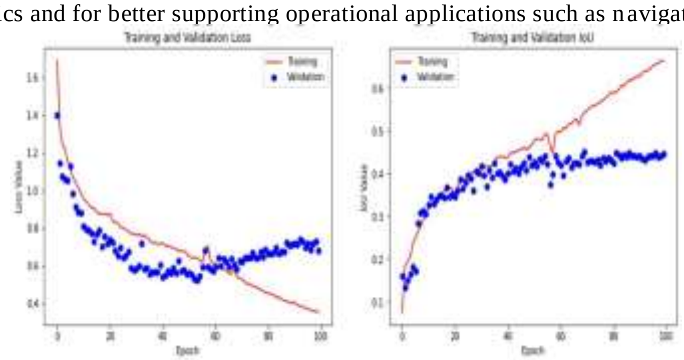

Fig:4 Model performance over training epochs

VI RESULT:

The U-Net-based convolutional neural network successfully generated detailed sea ice concentration charts from Sentinel-2 satellite imagery, focusing on the Hudson Bay region. By employing advanced segmentation techniques, the model accurately mapped sea ice into eight predefined classes, achieving a classification accuracy of 83% and a mean IoU score of 0.44. This highlights its capability to transform raw satellite data into actionable insights for monitoring sea ice dynamics and supporting environmental and navigational efforts. A detailed analysis revealed that the model performed well in predicting dominant classes such as open water, solid ice, and land, which are abundant in the dataset. However, the intermediate ice concentration classes (90% coverage) were less accurately predicted, highlighting a data imbalance issue. A confusion matrix analysis underscored this limitation, suggesting that additional springtime data could enhance the model's performance in these categories. The visualization of predictions against the true masks demonstrated the model's ability to generate finer ice concentration details compared to traditional ice charts. This indicates the potential for such models to produce higher-resolution and more detailed ice charts, which could benefit navigation and environmental monitoring efforts. Examples of predicted outputs show promising alignment with actual data, though further refinement through data augmentation and extended training datasets remains necessary. Overall, the findings underscore the feasibility of leveraging convolutional neural networks for automated sea ice mapping, with future work focusing on addressing class imbalances and exploring additional spectral bands to enhance accuracy.

Fig:5 SeaIceMaskComparison-1

Fig:6 SeaIceMaskComparison-2

VII. FUTURE WORK:

This study highlights the potential benefits and challenges of using convolutional neural networks for the automated mapping of sea ice. While the model demonstrated moderate accuracy and provided valuable insights, there is significant room for improvement. A key area of focus should be addressing the class imbalance, particularly regarding the underrepresented concentrations of intermediate ice within the dataset. Collecting additional data during transitional periods, such as spring and early summer when ice variability is at its peak, could improve the model's performance and generalization capabilities. Moreover, an important direction for future research involves leveraging additional spectral bands from Sentinel-2 imagery, including shortwave infrared, to better distinguish between different ice concentrations and surface types. Incorporating temporal data, such as time-series analyses of ice conditions, may also enhance predictive accuracy by capturing the dynamic processes of ice formation and melting. Beyond improving accuracy, scaling the model for global applications is a crucial next step. This necessitates adjustments to the workflow to accommodate larger, geographically diverse datasets and applying the model to regions with varying ice characteristics. Furthermore, integrating the model into real-time operational systems, complete with automated data collection and processing pipelines, could enable faster and more efficient production of sea ice maps for practical uses in navigation and environmental monitoring. Future research should also explore the integration of this model with other data sources, such as atmospheric and oceanographic variables, to provide a more comprehensive understanding of sea ice dynamics. By addressing these challenges, this research can progress towards developing a more accurate, scalable, and practical solution for global sea ice monitoring, thereby contributing to climate research and informed decision-making.

VIII. CONCLUSION:

This research investigated the application of convolutional neural networks, particularly the U-Net architecture, for the automation of sea ice concentration mapping utilizing Sentinel-2 satellite imagery, with a specific emphasis on the Hudson Bay area. The findings indicate that the model is exceptionally proficient in segmenting satellite images into distinct sea ice categories, achieving an accuracy rate of 83% and a mean Intersection over Union (IoU) of 0.44. The resulting ice charts correspond closely with conventional datasets while presenting the opportunity for enhanced spatial resolution. This capability holds significant importance for uses in environmental monitoring, maritime navigation, and climate studies. The model's efficacy highlights the advantages of AI-based methodologies in handling extensive and intricate geospatial datasets. Nonetheless, issues such as class imbalances within the training data particularly concerning intermediate ice concentration's underscore the necessity for additional data gathering and refinement. Broadening the dataset to encompass a wider variety of seasonal imagery, especially during transitional phases like spring and early summer, could enhance the model's accuracy across all categories. Furthermore, utilizing the complete spectrum of Sentinel-2's spectral bands beyond just the visible and near-infrared could further improve the model's predictive performance. Prospective avenues for this research involve scaling the model for global applications by integrating it with larger datasets and investigating real-time implementation in operational frameworks. The potential to merge such models with other remote sensing technologies and machine learning techniques could transform our approach to monitoring and addressing changes in polar and subpolar ice systems. Additionally, the integration of domain-specific insights, such as oceanic and atmospheric conditions, may offer further context to refine the predictions.

In conclusion, this project showcases the transformative power of deep learning in geospatial analysis and paves the way for innovative, scalable solutions to pressing environmental challenges. By bridging the gap between satellite imagery and actionable insights, this research contributes to advancing sustainable and data-driven approaches to Arctic and global ice monitoring.

IX. REFERENCES:

[1]. Canadian Ice Service. (2009). Canadian Ice Service Arctic Regional Sea Ice Charts in SIGRID-3 Format, Version 1. Boulder, Colorado, USA: National Snow and Ice Data Center. DOI: https://doi.org/10.7265/N51V5BW9.

[2]. Copernicus Sentinel Data (2021). Sentinel-2 Satellite Imagery. Accessed via Sentinel Hub API. https://eolearn.readthedocs.io.

[3]. Ronneberger, O., Fischer, P., & Brox, T. (2015). UNet: Convolutional Networks for Biomedical Image Segmentation. Retrieved from: https://arxiv.org/abs/1505.04597.

[4]. Sylvester, A. (2021). Arctic Sea Ice Image Masking Dataset. Retrieved from: https://www.kaggle.com/alexandersylvester/arcticsea-ice-image-masking.

[5].WMO. (2004). SIGRID-3: A Vector Archive Format for Sea Ice Charts. World Meteorological Organization. https://library.wmo.int/doc_num.php?explnum_id=9270.

[6]. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press. https://www.deeplearningbook.org.

[7]. Chollet, F. (2017). Deep Learning with Python. Manning Publications. ISBN: 9781617294433.

[8]. Kingma, D. P., & Ba, J. (2014). Adam: A Method for Stochastic Optimization. Retrieved from: https://arxiv.org/abs/1412.6980.

[9]. Zhu, X., Vondrick, C., Fowlkes, C., & Ramanan, D. (2016). Do We Need More Training Data or Better Models for Object Detection? Retrieved from: https://arxiv.org/abs/1511.06930.

[10]. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770-778. DOI: https://doi.org/10.1109/CVPR.2016.90.

[11]. Simonyan, K., & Zisserman, A. (2015). Very Deep Convolutional Networks for Large-Scale Image Recognition. Retrieved from: https://arxiv.org/abs/1409.1556.

[12]. 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: https://doi.org/10.1109/CVPR.2015.7298965.

[13]. Zhang, Z., Liu, Q., & Wang, Y. (2018). Road Extraction by Deep Residual U-Net. IEEE Geoscience and Remote Sensing Letters, 15(5), 749-753. DOI: https://doi.org/10.1109/LGRS.2018.2802944.

[14].Girshick, R. (2015). Fast R-CNN. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), 1440-1448. DOI: https://doi.org/10.1109/ICCV.2015.169.

[15]. Redmon, J., & Farhadi, A. (2018). YOLOv3: An Incremental Improvement. Retrieved from: https://arxiv.org/abs/1804.02767.

[16]. Springenberg, J. T., Dosovitskiy, A., Brox, T., & Riedmiller, M. (2015). Striving for Simplicity: The All Convolutional Net. Retrieved from: https://arxiv.org/abs/1412.6806.

[17]. Paszke, A., Gross, S., Chintala, S., et al. (2017). Automatic Differentiation in PyTorch. In Advances in Neural Information Processing Systems (NeurIPS). Retrieved from: https://pytorch.org.

[18]. LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-Based Learning Applied to Document Recognition. Proceedings of the IEEE, 86(11), 22782324. DOI: https://doi.org/10.1109/5.726791.

[19]. Szegedy, C., Liu, W., Jia, Y., et al. (2015). Going Deeper with Convolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 1-9. DOI: https://doi.org/10.1109/CVPR.2015.7298594.

[20]. Isola, P., Zhu, J. Y., Zhou, T., & Efros, A. A. (2017). Image-to-Image Translation with Conditional Adversarial Networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 5967-5976. DOI: https://doi.org/10.1109/CVPR.2017.632.

[21]. 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: https://doi.org/10.1109/TPAMI.2016.2644615.

[22]. Li, W., Wang, F., Li, M., & Gao, J. (2018). DeepUNet: A Deep Fully Convolutional Network for Pixel-Level Sea-Land Segmentation. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 11(11), 3954-3962. DOI: https://doi.org/10.1109/JSTARS.2018.2869579.

[23].Sutton, R. S., & Barto, A. G. (2018). Reinforcement Learning: An Introduction. Second Edition. MIT Press. http://incompleteideas.net/book/the-book.html.

[24]. Bishop, C. M. (2006). Pattern Recognition and Machine Learning. Springer. ISBN: 978-0387310732.

[25]. Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely Connected Convolutional Networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 4700-4708. DOI: https://doi.org/10.1109/CVPR.2017.243.

[26]. Abadi, M., Barham, P., Chen, J., et al. (2016). TensorFlow: A System for Large-Scale Machine | Page Learning. In Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI), 265-283. https://www.tensorflow.org.

[27]. Zhang, H., Dana, K., Shi, J., et al. (2018). Context Encoding for Semantic Segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 7151-7160. DOI: https://doi.org/10.1109/CVPR.2018.00747.

[28]. Lin, T. Y., Dollár, P., Girshick, R., et al. (2017). Feature Pyramid Networks for Object Detection. In Proceedings of the IEEE Conference

on Computer Vision and Pattern Recognition (CVPR), 2117-2125. DOI: https://doi.org/10.1109/CVPR.2017.106.
[29].Zhou, Z., Siddiquee, M. M. R., Tajbakhsh, N., & Liang, J. (2018). UNet++: A Nested U-Net Architecture for Medical Image Segmentation. Retrieved from: https://arxiv.org/abs/1807.10165.
[30]. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In Proceedings of the 25th International Conference on Neural Information Processing Systems (NeurIPS), 1097-1105. DOI: https://doi.org/10.1145/3065386.