POAC'25
St. John's,
Newfoundland and
Labrador, Canada

Proceedings of the 28th International Conference on
Port and Ocean Engineering under Arctic Conditions
Jul 13-17, 2025
St. John's, Newfoundland and Labrador
Canada

A Sensor Fusion Platform for Semantic Segmentation of Sea Ice Imagery

Richard R.J. Duan[1], K. Andrea Scott[1], Corwin G.J. MacMillan[1],
Robert M. Gash[2], and Zhao Pan[1]
[1] University of Waterloo, Waterloo, ON, Canada
[2] Ocean Coastal and River Engineering Research Centre, National Research Council Canada,
St. John's, Canada

ABSTRACT

Declining sea ice coverage in the Arctic is expected to open up routes such as the Northwest Passage for commercial shipping for greater parts of the year. However, collisions with ice are still a major risk for ships travelling through ice-covered waters. On-board ice experts can help with identifying the nearby ice conditions for navigation, but when presented with a field of different types of ice that are visually similar, it is difficult to accurately assess the ice conditions to plot a safe route. We introduce a multi-sensor system that leverages the strengths of sensor fusion to provide unique information that can highlight the differences in ice beyond what a human expert can see. Unique data of river ice was collected in February 2025 along the shores of the Ottawa and St. Lawrence rivers as an initial test of the system. A small, labelled dataset was then created to test the performance of some basic sensor fusion methods on a pre-trained image segmentation network.

KEY WORDS: Sensor Fusion; Sea Ice Navigation; Image Processing; Semantic Segmentation; Deep Learning.

INTRODUCTION

As global warming progresses, the Arctic sea ice concentration is projected to decrease at an increasing rate. Consequently, previously impassable routes through the Arctic are becoming possible options for commercial shipping, including the Northwest Passage (NWP) through the Canadian Arctic Archipelago. For example, a simulation conducted by Somanathan et al. (2009) shows that a trip between St. John's, Newfoundland and Yokohama, Japan can be shortened by 3500 nautical miles – a 33% decrease – by using the NWP. The significant savings in distance travelled and decreased reliance on the Panama/Suez canals make these routes an attractive choice for ship operators. Additionally, the increased marine traffic through Canadian waters also provides an opportunity for stimulating economic activity in the area.

Accurate information of the upcoming ice conditions is crucial for safe navigation in ice-covered waters. According to Garvin (2020), while modern ice charts and satellite imagery are crucial for general route planning, on-site assessments from specially trained ice experts are still required for real-time navigation. Typically, on-board observations are limited to what can be seen ahead by the observer for near-ship ice conditions. However, such observations are somewhat subjective in nature and can become ineffective in adverse weather. Furthermore, the human eye is limited to seeing only the narrow spectrum of visible light, losing out on important information present in the rest of the spectrum and in the polarization of light.

For a system to outperform human observations, it is clear that no single sensor can return useful information in all environmental conditions. A review by Fung et al. (2017) found that in the field of autonomous vehicles, the state-of-the-art perception systems all leverage the power of fusing together information from a complementary assortment of different sensors (such as cameras, Global Positioning System, radar, and LIDAR for autonomous driving). To effectively use the multimodal data collected by different sensors, an appropriate data fusion strategy must be chosen.

In general, there are three different points along the data processing pipeline where fusion can occur. Data level or early fusion combines the calibrated data from all sensors and is usually applied when the data format is similar across the sensors, such as in the case of multiple cameras. This strategy will be used for the experiments covered in this paper. Feature level or middle fusion first conducts feature extraction on the calibrated data from each sensor separately, before combining the detected features to be further analyzed. Finally, decision level or late fusion combines the analyzed output from each individual sensor to generate a single decision, and is usually applied to extremely different data formats that are difficult to combine.

In this study, we propose a novel sensor system using early fusion that can provide beyond human-visible information of nearby ice conditions to aid in real-time navigation. We introduce a new multimodal dataset of river ice scenes collected by the sensor system in land-based trials, and evaluate the performance of several sensor fusion methodologies using the UPerNet semantic segmentation network.

Related Works

Trials have been conducted in the past to collect data aboard ships travelling through ice, both in simulated and true ice conditions. Brown et al. (2023) introduced a suite of cameras and sensors – including stereoscopic and 360° cameras, LIDAR, and an electromagnetic sensor – that were mounted on the Canadian Coast Guard ship *Henry Larsen* in March of 2022. The study was exploratory, evaluating the performance of individual sensors on sea ice, and thus fusion of different sensor outputs was not explored. Zhou et al. (2023) captured optical data by attaching a camera to a model ship in an ice basin, which was then used to train YOLACT, a deep learning network, to perform instance segmentation on the scene. As the goal was to estimate ice concentration and floe size distribution, only three classes (ship, ice, and water) were used for labelling this dataset.

Additionally, there has been increasing interest recently in using computer vision and deep learning models to identify and distinguish different types of sea ice for navigational purposes. Kim et al. (2019) introduced an object classification network to distinguish between nine different categories of ice, trained on publicly available close-range imagery of ice. Zhang et al. (2022) modified DeepLabv3+, a state-of-the-art semantic segmentation model, to conduct pixel-wise ice image segmentation. The dataset used in the study consisted of manually labelled images captured by an optical camera, mounted on an ice-strengthened cargo ship as it travelled through the Arctic.

As the introduced system includes a unique combination of sensors, field trials were conducted in February 2025 to collect first-of-its-kind data on river ice along the Ottawa and St. Lawrence rivers. The data was processed to create a dataset of spatially and temporally registered images from each sensor. A small subset of the dataset was manually labelled for preliminary testing of different early fusion strategies using a pre-trained image segmentation network.

METHODOLOGY

Hardware Description

Due to the flexibility allowed via sensor fusion, a number of different sensor types were considered for this study. Field tests conducted by Perovich (1994) on several types of sea ice surfaces showed large differences in reflectance in the near-infrared wavelength range and in polarized light, revealing a sensitivity of sensors that can enhance the ability to identify and classify different types of sea ice. Additional inspiration for sensor choices were taken from the tangential fields of autonomous driving and robotics. For example, Rankin and Matthies (2010) used several different sensors – including colour cameras, polarization cameras, multispectral cameras, and various infrared wavelength cameras – to detect mud hazards ahead of an unmanned ground vehicle. Emphasis was placed on the ease of integration into a relatively self-contained system that can withstand the harsh environment that it will be exposed to. The following sensors were included, with additional information on each of cameras available in Table 1:

Visual (RGB) Camera – this camera most closely resembles what a person can see and will also be used as a baseline for comparison.

Stereoscopic Cameras – a set of two corresponding visual cameras which, when calibrated, can provide depth information about the scene.

Thermal Infrared Camera – already widely used on marine vessels and autonomous vehicles due to good performance in low-light conditions. The differences in temperature between ice and water can help with identification in daylight conditions as well.

Polarization Camera – provides information on how light is linearly polarized as a result of being reflected off of different surfaces, which can be processed to give additional information about features such as the surface angle and roughness.

Inclinometer – measures the angle of the sensors relative to the horizon. This information can aid in reprojecting the individual camera views to reconstruct the ice scene.

Table 1. Details of the cameras used.

Camera	Field of View (Horizontal/Vertical)	Resolution	Framerate
Visual (RGB)	92°/61°	3840 x 2160	29.97 FPS
Stereoscopic	92°/61°	3840 x 2160	29.97 FPS
Thermal	25°/20°	640 x 480	29.97 FPS
Polarization	19°/16°	1224 x 1024	1 FPS

Three GoPro HERO 12 Black cameras were used to act as visual cameras, with two cameras serving as a pair of stereoscopic cameras, and the remaining one acting as the baseline. The GoPros were chosen due to their built-in weatherproofing and acceptable optical clarity. A FLIR MD625 marine thermal camera and a LUCID Vision Labs Triton polarization camera were also chosen for their weather ratings, at the cost of being more difficult to integrate into the system. All sensors were attached to a common rigid frame made of aluminum extrusions using custom-made mounts. The cameras were carefully placed as to ensure that each camera had a clear, unobstructed view of the scene ahead. The sensor frame with the cameras was then attached to a tripod to allow for portability and ease of adjustment during data collection. The setup as described can be seen in Figure 1a. A digital level was used to measure the angle of inclination of the cameras.

To ensure proper analysis of the collected data, the cameras needed to be synchronized with each other. A GoPro The Remote was used to connect to the three GoPro cameras via Bluetooth for software synchronization. However, due to the incompatibilities between the different cameras and the analog nature of the thermal camera, manual synchronization through clapping was used. Considering the 29.97 frames per second rate at which the GoPro and analog capture device recorded at, this sets an upper limit on the temporal difference within a set of image captures to be about 17 milliseconds. As ships and ice are relatively slow-moving objects, the temporal difference between synchronized views was considered acceptable.

Camera Calibration

As with any application involving computer vision, camera calibration is important to ensure that the objects in view are not distorted by the optical properties of the camera and lens. To conduct camera calibration, a routine incorporating the method introduced by Zhang (2000) was used. The use of multiple cameras introduces the additional problems of registering the different views to a single, common view for analysis. As such, a careful calibration scheme was required for this sensor suite. To find the intrinsic parameters of each individual camera, a planar pattern was shown in multiple different positions and orientations throughout the field of view of the camera. Usually, a simple printed pattern is sufficient for the routine. However, the inclusion of a thermal camera required a different method, as it is unable to distinguish between printed colours. Thus, a custom calibration board had to be designed that can be visible to all cameras at the same time.

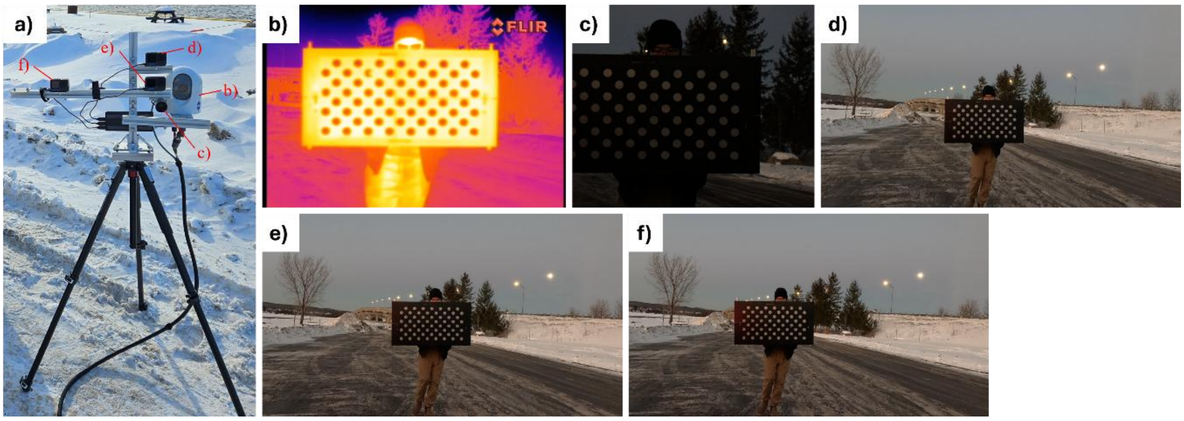

Figure 1. a) Full sensor system with each camera labelled. Camera calibration was conducted using a rigid asymmetric circle grid board, with an example frame shown from each of the synchronized views of b) the thermal camera, c) the polarization camera, d) the baseline RGB camera, e) and f) the left and right pair of RGB stereoscopic cameras.

A review by ElSheikh et al. (2023) identified various designs for thermal camera calibration boards that have been previously implemented, and these generally fall within two categories – active and passive. Active calibration boards include a powered heating or cooling system to create a temperature difference that is visible on the thermal camera. Passive boards rely on an external heating source, such as the sun shining on dark areas, to provide the temperature difference. Both types of boards can also incorporate materials with different thermal emissivities to enhance the thermal contrast.

Although active boards tend to result in sharper details as a result of the relatively instant temperature change, passive boards are far simpler and more portable, especially considering weatherproofing for cold, outdoor conditions. Thus, a passive board made up of a composite of black acrylic sheets bonded to unfinished aluminum sheets was made. An asymmetrical circle grid pattern, required to remove ambiguity in stereoscopic calibration, was laser cut into the acrylic as the calibration pattern. Circles were used instead of the classic square checkerboard due to their good performance even with blurry features. The calibration process as seen from each camera is shown in Figure 1b-f.

Dataset Acquisition

A preliminary set of land-based data of river ice was collected in February 2025 at various locations along the Ottawa and St. Lawrence rivers, as shown in Figure 2. Although river ice formation and dynamics are different compared to sea ice – especially since sections of the St. Lawrence Seaway are managed waterways as opposed to free-flowing water – it nonetheless provided an opportunity to test out the sensor suite on similar-looking data in a variety of environmental conditions. Time spent recording at each location ranged from 5-30 minutes, providing hundreds of gigabytes of raw video footage. Data was collected across multiple days at a total of six locations, referred to as scenes; Table 2 provides details noted during the collection process. Camera calibration was conducted at the start and end of the trip to ensure consistency in calibration properties throughout the acquisition process.

Table 2. Description of data acquisition conditions and details.

Scene	Date	Location	Length of Capture	Lighting/Weather Conditions	Notes
1	2025-02-11	Hawkesbury, ON (45°37'05.1"N 74°35'54.1"W)	20 min.	Good/Clear	Stationary ice, tested different camera angles looking down from bridge
2	2025-02-12	Berthier-sur-Mer, QC (46°56'06.8"N 70°44'08.8"W)	30 min.	Good/Clear	Relatively fast-flowing brash ice
3	2025-02-12	Saint-Jean-de-l'Île-d'Orléans, QC (46°54'56.0"N 70°53'46.8"W)	5 min.	Poor/Clear	After sunset; fast-flowing brash ice difficult to see in most views
4	2025-02-13	Montréal, QC (45°30'46.7"N 73°31'43.0"W)	5 min.	Medium/Snow	At dawn with overcast sky and heavy snowfall; mostly small clumps of thin ice
5	2025-02-15	Québec, QC (46°46'13.7"N 71°14'54.3"W)	25 min.	Good/Clear	At sunrise with prominent reflection of the sun; relatively fast-moving brash ice
6	2025-02-15	Rivière-du-Loup, QC (47°50'46.4"N 69°34'18.0"W)	10 min.	Good/Clear	Large swaths of stationary ice in foreground with fast-moving ice ~400 m from shore

The GoPro cameras were set to record using the *Linear* lens mode with auto exposure turned on. The exposure on the polarization camera was manually set at each location according to the lighting conditions, as auto exposure in the provided interfacing software tended to overexpose the image. The thermal camera did not have options to adjust video settings other than the video colour palette, which was set to red-hot.

Figure 2. Map of the data collection locations in red for each scene as specified in Table 2.

Data Pre-Processing

To prepare the dataset, the videos recorded from each camera were initially manually synchronized by temporally aligning the synchronization action across each view. Given the known framerate of each video, the synchronized frame times can then be propagated throughout the entire video for that scene. The synchronized frames were then extracted at 1 frame per second, matching the framerate of the polarization camera. The dataset size was further reduced by taking every 5th or 10th frame depending on the scene, as there was little change from frame-to-frame due to the relatively slow-moving ice. Each frame was then undistorted according to the calibration parameters found using the Computer Vision Toolbox in Matlab (2025). As there is a small shift in perspective between each view due to the physical distances between the cameras, each view was warped to that of the baseline camera using manually selected control points for each scene. Finally, the frames were cropped and scaled to 512×512, containing only those pixels which are present across all cameras when warped to the baseline view.

The polarization camera required some additional processing to obtain linear polarization information. Using images captured at each of the four polarization angles (0°, 45°, 90°, 135°), three Stokes parameters can be calculated as follows:

$$S_0 = I_0 + I_{90} \tag{1}$$
$$S_1 = I_0 - I_{90} \tag{2}$$
$$S_2 = I_{45} - I_{135} \tag{3}$$

where I_0, I_{45}, I_{90}, I_{135} are the pixel intensities at each of the four polarization angles, and S_0, S_1, S_2 are the three Stokes parameters that describe the linear polarization state of that pixel. Thus, the degree of linear polarization (DoLP) and angle of linear polarization (AoLP) can be calculated as follows:

$$p = \frac{\sqrt{S_1^2 + S_2^2}}{S_0} \tag{4}$$

$$2\psi = \tan^{-1}\left(\frac{S_2}{S_1}\right) \tag{5}$$

where p is the DoLP and ψ is the AoLP in radians. The resultant polarization images contain a significant amount of salt-and-pepper noise due to the sensitivity in measuring polarizing light, so a median filter was applied to remove some of the noise. Additionally, histogram equalization was applied to each frame to enhance contrast and make details more visible.

Data Labelling

In total, there were over 600 synchronized and cropped frames across all scenes. Of those, every 5th frame was picked out to be manually labelled, giving a total of 133 fully labelled images across the six scenes. The dataset was manually labelled from the cropped view of the baseline camera using the Computer Vision Annotation Tool (CVAT) online (2025). Large, relatively simple features such as the sky and background were manually labelled and propagated with small changes throughout the appropriate frames. More complex features, such as small pieces of brash ice, were labelled semi-automatically using the built-in interface to the Segment Anything Model (SAM) introduced by Kirillov et al. (2023) with given prompts.

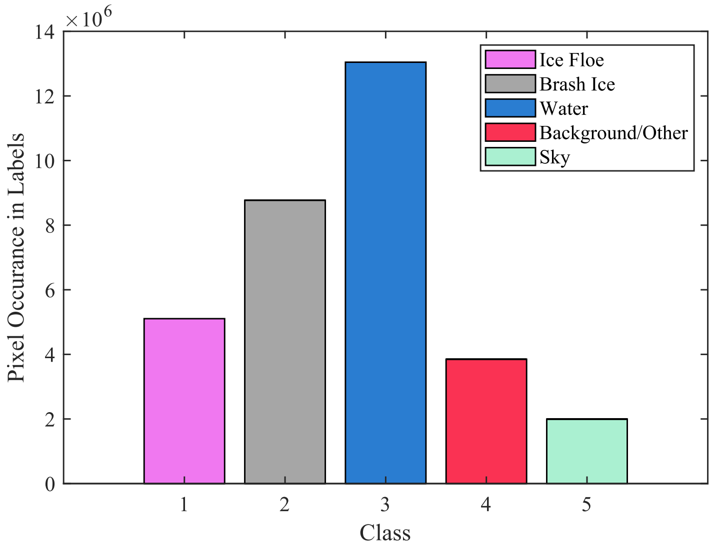

Figure 3. Histogram of class distribution across the labelled images. The bar colours correspond to the label colours in the dataset. There is a large bias towards the water class and a lack of the sky class in the dataset due to the slight downwards angle during data collection.

The label classes were based on features described by the World Meteorological Organization (2014). In total, 5 classes were used in labelling, with the class distribution across the labelled images is shown in Figure 3:

- **Ice Floe** – Large, flat pieces of floating ice with relatively smooth texture
- **Brash Ice** – Congregations of many different small fragments of floating ice, with rough and jagged texture
- **Water** – Open, ice-free water
- **Background/Other** – Items in frame that are not ice or water, such as trees or hills in the background

- **Sky** – Areas above the horizon that are not obstructed by other objects or the background

EXPERIMENTS

Sensor Fusion Methodology

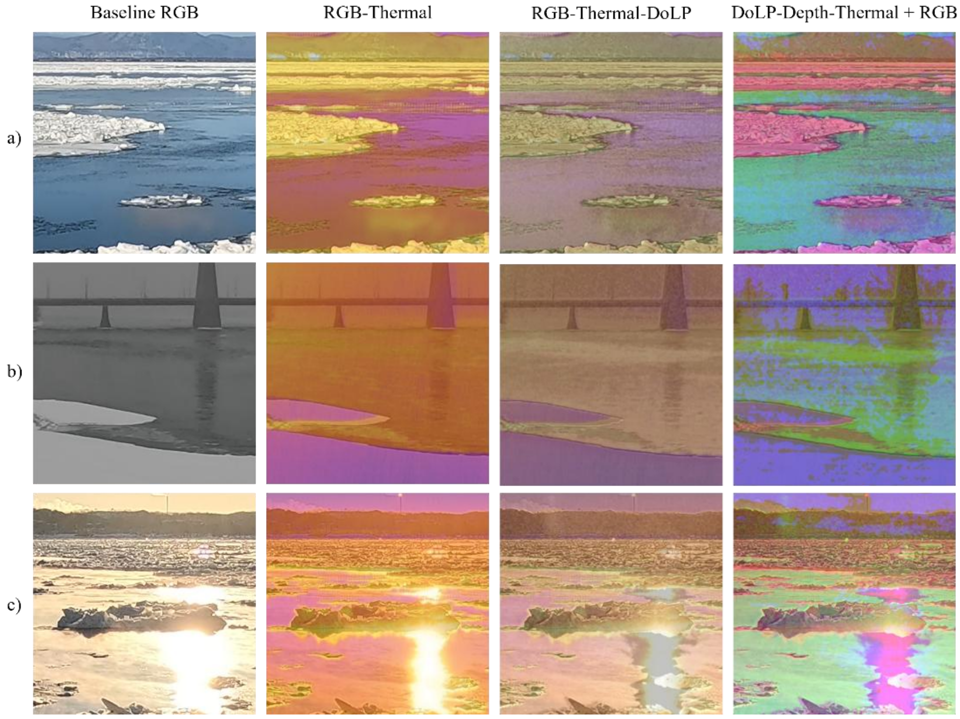

Figure 4. Example images from the (columns from left to right) baseline RGB, RGB-Thermal, RGB-Thermal-DoLP, and DoLP-Depth-Thermal + RGB datasets in three different settings: a) daylight, b) snow, and c) sunrise.

An important step in multimodal data fusion is to select an appropriate fusion strategy to use. From reviewing the different fusion strategies covered by Brenner et al. (2023), early fusion was chosen as it is simple to implement with similar data formats across all sensors, being in the form of captured images. Due to the inclusion of depth and thermal data as an addition to RGB optical data being extremely prevalent in autonomous driving and surveillance, inspiration was drawn from some of the fusion strategies used in those fields.

The main methods used to fuse the different images together were concatenation and blending to create a new 3-channel RGB image. Three new sets of artificial images were created using different combinations of sensor outputs and fusion methods. One set was generated by simply blending the baseline and thermal images at an equal 1:1 opacity ratio (dubbed as RGB-Thermal). Another set added polarization data, blending the baseline, thermal, and DoLP images at an opacity ratio of 2:1:1 respectively (dubbed as RGB-Thermal-DoLP). A final third set used concatenation to include depth data by first creating a false-colour image with DoLP, depth, and thermal acting as the "red", "green", and "blue" channels respectively. The false-colour image was then blended with the baseline view at an equal 1:1 opacity ratio (dubbed as DoLP-Depth-Thermal + RGB). A few examples of the fused images are shown in Figure 4.

Image Segmentation Network

The Unified Perceptual Parsing Network (UPerNet) proposed by Xiao et al. (2018) was used to benchmark the performance of each set of fused data compared to the baseline. UPerNet is based on the Feature Pyramid Network (FPN), which itself fuses together feature information extracted by a backbone network at multiple levels. Introduced by He et al. (2016), a ResNet encoder with a depth of 50 layers was used as the backbone network due to its good general performance in image segmentation tasks.

The low-light scene 3 in Table 2 was not used in the experiments, as the lack of information in most of the images had a detrimental effect on the accuracy of the model, resulting in a final set of 125 images. To offset the effects of a small dataset, transformations were applied to the images for training. The images were augmented through a combination of random cropping, random rotation, adding random Gaussian blur, and random jittering of the contrast and colour saturation. To quantify the performance in each experiment, standard semantic segmentation metrics of mean intersection over union (mIoU) and mean pixel-wise accuracy (mAcc) were used. An 80:20 training-test split was used, giving a total test size of 25 images chosen randomly while making sure that the test dataset had at least one image included from each of the five tested scenes, as per Table 2.

RESULTS

The outputs of the network on the different datasets are shown in a few example images in Figure 5. The performance metrics on each dataset in terms of mIoU and mAcc are included in Table 3. The network performed well at separating between water and ice in most cases, even with the reflection of the sun being present. This is perhaps due to the fact that the reflections are only visible on the surface of the water in the dataset. Although water and reflections are very visibly different, the network is able to associate reflections with the presence of water. There are, however, small inconsistencies mostly with sky being confused with water and the border between the background and sky. Additionally, the network struggled when faced with background objects beyond the scenery, such as the bridge with a number of fine pillars. Both of these results are likely due to their infrequent occurrence during training, as the sky class has the lowest distribution and there are only a few frames with non-scenery background labels present.

Quantitatively, the network performed best on the baseline dataset, followed by the DoLP-depth-thermal + RGB dataset at a significant drop in both mIoU and mAcc. This indicates that the current method of early fusion is causing confusion somewhere, making the network less capable of accurately making predictions. It is likely that the added noise from the different sources have not been effectively filtered out prior to the fusion process, and as a result the additional noise is overwhelming any potential information that can be gained from the other sensors. The effect of noise on the data can be clearly seen in Figure 4b, where relatively flat areas of water and sky in the baseline dataset appear significantly grainier with randomly added artifacts in the fused datasets. Overall, early fusion through simple image blending and concatenation alone was not effective at extracting additional sources of useful information compared to the baseline, based on our preliminary results using the pre-trained UPerNet network.

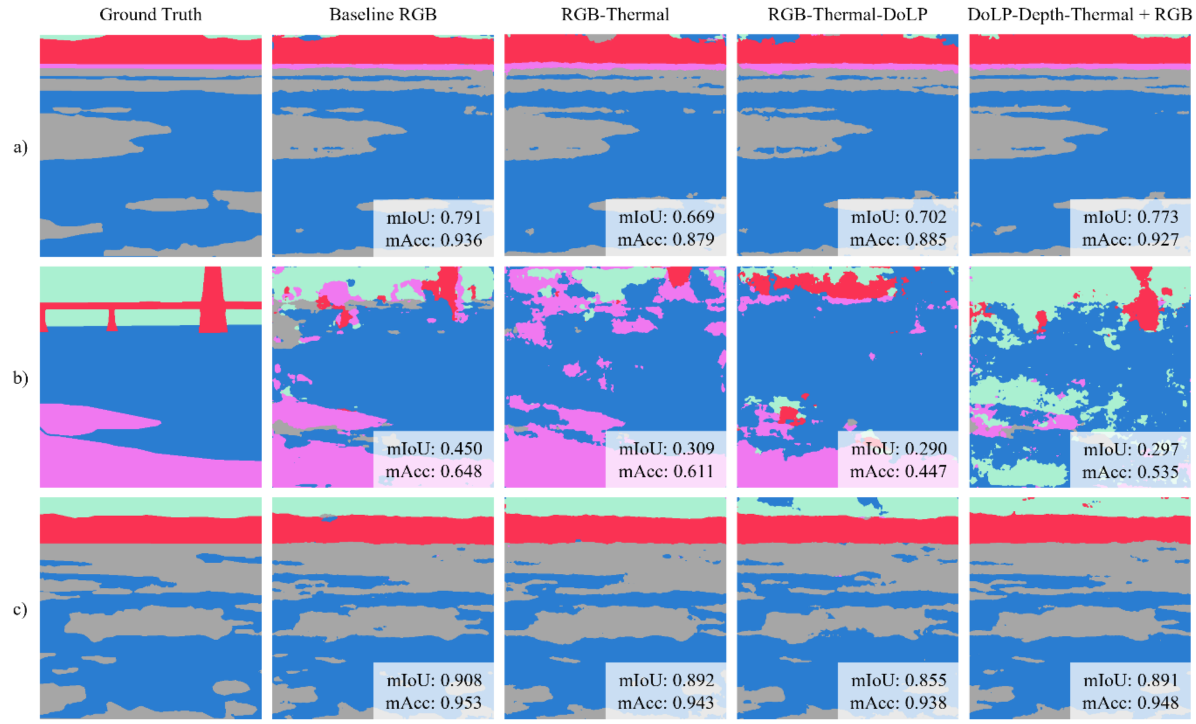

Figure 5. Examples of predictions from UPerNet with mIoU and mAcc metrics compared to the labelled ground truth for the three settings shown in Figure 4. The label colours correspond to the classes in Figure 3. The network performs well across the board, with some inconsistency in the sky. However, the network struggled to process the bridge in b) due to the lack of man-made structures in the dataset.

Table 3. UPerNet performance on the different datasets.

Dataset	Baseline RGB	RGB-Thermal	RGB-Thermal-DoLP	DoLP-Depth-Thermal + RGB
mIoU[*]	**0.863**	0.781	0.782	*0.817*
mAcc[*]	**0.931**	0.887	0.882	*0.904*

[*] Best performance is **bolded**, second best is *italicized*.

CONCLUSIONS

A novel sensor suite consisting of a baseline RGB camera, a pair of stereoscopic RGB cameras, a thermal camera, and a polarization camera was constructed and tested on river ice data. In the process, a first-of-its-kind multimodal dataset with fully calibrated and registered images was created, along with a small subset of labelled images for neural network training. A number of different early fusion methods were tested using the pre-trained UPerNet network, although none managed to outperform the baseline case, likely due to noise from the different cameras. Further image processing, different training and fine-tuning methods, and a more sophisticated fusion strategy will likely be needed to make full use of the available sources of information, especially during non-ideal conditions for RGB optical data.

ACKNOWLEDGEMENTS

We would like to acknowledge the National Research Council of Canada for financially supporting the project and providing access to the indoor ice tank facility for initial small-scale experiments. We would also like to acknowledge the partial support from the Natural Sciences and Engineering Research Council of Canada Undergraduate Student Research Award (USRA) during the course of this project.

REFERENCES

Brenner, M., Reyes, N. H., Susnjak, T., & Barczak, A. L. C., 2023. RGB-D and Thermal Sensor Fusion: A Systematic Literature Review. *IEEE Access*, 11, pp. 82410-82442

Brown, J., King, M., & Briggs, R., 2023. Evaluation of Near-Ship Ice Conditions from Ship Bourne Sensors. *The 33rd International Ocean and Polar Engineering Conference.*

CVAT.ai Corporation, 2025. *About CVAT*. [Online] Available at: https://www.cvat.ai/about [Accessed 11 March 2025].

ElSheikh, A., Abu-Nabah, B. A., Hamdan, M. O., & Tian, G.-Y., 2023. Infrared Camera Geometric Calibration: A Review and a Precise Thermal Radiation Checkerboard Target. *Sensors*, 23(7), 3479.

Fung, M.L., Chen, M.Z.Q., & Chen, Y. H., 2017. Sensor fusion: A review of methods and applications. *29th Chinese Control and Decision Conference (CCDC)*, pp. 3853-3860.

Garvin, M., 2020. Review of Technologies for Real-Time Shipboard Ice Severity Sensing. *Smart Ship Technology*, pp. 53-62.

He, K., Zhang, X., Ren, S., & Sun, J., 2016. Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770-778.

Kim, E.K., Dahiya, G.S., Løset, S., Skjetne, R., 2019. Can a computer see what an ice expert sees? Multilabel ice objects classification with convolutional neural networks. *Results in Engineering*, 4, 100036.

Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., & Girshick, R., 2023. *Segment Anything*. [Online] (Updated 5 April 2023) Available at: https://doi.org/10.48550/arXiv.2304.02643 [Accessed 20 March 2025].

Perovich, D.K., 1994. Light reflection from sea ice during the onset of melt. *Journal of Geophysical Research*, 99(C2), pp. 3351-3359.

Rankin, A.L. & Matthies, L.H., 2010. Passive sensor evaluation for unmanned ground vehicle mud detection. *Journal of Field Robotics*, 27(4), pp. 473-490.

Somanathan, S., Flynn, P., & Szymanski, J., 2009. The Northwest Passage: A Simulation. *Transportation Research Part A: Policy and Practice*, 43(2), pp. 127-135.

The Mathworks, Inc., 2025. *MATLAB*. [Online] Available at: https://www.mathworks.com/products/matlab.html [Accessed 18 March 2025].

World Meteorological Organization (WMO), 2014. *WMO Sea-Ice Nomenclature*.

Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J., 2018. Unified Perceptual Parsing for Scene Understanding. *Computer Vision – ECCV 2018*, 5, pp. 432-448.

Zhang, Z., 2000. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11), pp. 1330-1334.

Zhou, L., Cai, J., & Ding, S., 2023. The Identification of Ice Floes and Calculation of Sea Ice Concentration Based on a Deep Learning Method. *Remote Sensing*, 15(10), 2663.