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

Deep learning for automated surveillance of sea ice dynamics

Matias Uusinoka[1], Arttu Polojärvi[1], Jari Haapala[2]
[1] Aalto University, Espoo, Finland
[2] Finnish Meteorological Institute, Helsinki, Finland

ABSTRACT

Effective monitoring of sea ice conditions is required for safe maritime operations and research on sea ice. Common surveillance systems rely on radar imagery from ships and coastal stations, which are often handled through a time-consuming process. Currently existing automated systems offer continuous monitoring but are often limited to coarse spatial resolutions and with computational efficiency restricting their use for detailed analysis and real-time navigation support. We leverage state-of-the-art deep learning-based optical flow architectures for continuous, high-resolution surveillance of sea ice dynamics with shipborne and coastal radar systems. By employing these architectures, we can accurately capture the relative motion and deformation fields of sea ice with considerably lower computational overhead. We demonstrate the applicability of the methods with operational radar systems to demonstrate their efficiency and accuracy.

KEY WORDS: Deep learning; Optical Flow; Ice Radar; Automation; Ice Dynamics

INTRODUCTION

Maritime operations in ice-covered waters benefit from accurate, near-real-time observations into sea ice dynamics. As Arctic marine transit routes become increasingly viable due to changing climate conditions and, for example, as offshore wind energy harvesting expands into regions with seasonal ice cover, the ability to monitor and predict sea ice dynamics grows in importance. Traditional satellite-based methods (e.g. Ninnis et al., 1986; Thomas et al., 2011 or RADARSAT Geophysical Processor System), in addition to their expenses, often lack the spatial or temporal resolution necessary for local-scale real-time applications. Similarly, digital image correlation approaches (Chu et al., 1985; Peters & Ranson, 1982) used typically in laboratory settings do not readily transfer to challenging field conditions faced with ship or coastal radar systems. This sets a need for robust, high-resolution solutions directly suitable for radar imagery captured from ships or coastal radar stations. Despite the importance of radar-based sea ice monitoring, existing automated methods for full-field motion and deformation detection remain sparse. An approach by Karvonen (2016) employs "virtual radar ice buoys" based on the Lucas-Kanade optical flow method (Lucas & Kanade, 1981) to track drift in automation but is computationally intensive and provides limited spatial coverage of ice motion. Furthermore, standard satellite-derived techniques and digital image correlation

systems are often confounded by speckle noise and the non-stationary vertical plane of typical ship-based radar, rendering them unsuitable for real-time deployment.

In response to these challenges, we recently presented (Uusinoka et al., 2025a) a deep learning tool that harnesses the Recurrent All-Pairs Field Transforms (RAFT) architecture (Teed and Deng, 2020) and extended it with a novel temporal resolution tree (TRT) for improved robustness against noise and for increased accuracy. Unlike prior work, the framework offers full-field motion and deformation estimates, can be operated with minimal hardware, and is generalizable effectively to a variety of radar platforms. By integrating recurrent neural network components and carefully designed multi-scale temporal features, this method provides superior performance even when confronted with the high-frequency noise typical of maritime radars.

Beyond the core deep learning algorithm, this paper emphasizes the possibility for an automation pipeline that supports the method in practical settings. We demonstrate that the system can continuously intake radar data streams to produce full-field motion and deformation maps down to pixel scale in near-real time. This level of automation is of high use in operational environments, where consistent coverage and rapid decision-making are needed. We validate the versatility and efficiency of our method through case studies featuring ship radar data from the RV Polarstern and RV Kronprins Haakon, as well as coastal radar data from the Baltic Sea. In each scenario, the system is shown to accurately track ice motion, detect deformation zones. We also discuss scalability with parallelized implementations and coarse-graining methods for real-time implementations.

METHODS

Radar data

We perform case studies with three different radar datasets (Fig. 1a-c): ship radar data from RV Polarstern during the Multidisciplinary drifting Observatory for the Study of the Arctic Climate (MOSAiC) expedition (Nicolaus et al., 2022), ship radar data from RV Kronprins Haakon (KPH) gathered during the NPI Arctic Ocean Cruise II (Hop et al., 2024), and coastal radar (CR) data gathered at the Baltic Sea during March 2011. With the MOSAiC and CR datasets, a radar signal digitizing system was deployed as described in Karvonen (2016) and Oikkonen et al. (2017). The KPH dataset employs an onboard digitizing system used in navigation. With the ship-based radar systems a continuous scans were used to capture local sea ice conditions under dynamic navigation, while with the coastal radar a stationary installation along the Baltic Sea coastline was used to provide repeated sweeps of nearshore ice fields. The radar datasets differ in terms of noise profiles (Fig. 1d-f) and scene complexity, and thus provide excellent cases for exemplifying the generalization of the method.

The first set, MOSAiC, consisted of 144 ship radar frames at 10-minute resolution capturing pack ice conditions around the RV Polarstern in the high Arctic over a 24-hour period. The ice conditions ranged from stable pack ice to leads and ridges forming along fracture lines. The second dataset, KPH comprised 288 frames at 5-minute resolution, and presented frequent interactions between discrete ice floes moving, shearing, and colliding over short timescales. In nearshore setting, the CR data from the Baltic represents highly compact and deformed ice field in the coastal boundary zone between drift ice and fast ice regions (Oikkonen et al., 2016).

An overview of the image dimensions, pixel resolution, and temporal sampling for these datasets is summarized in Table 1. The MOSAiC dataset, originally captured at a radar range

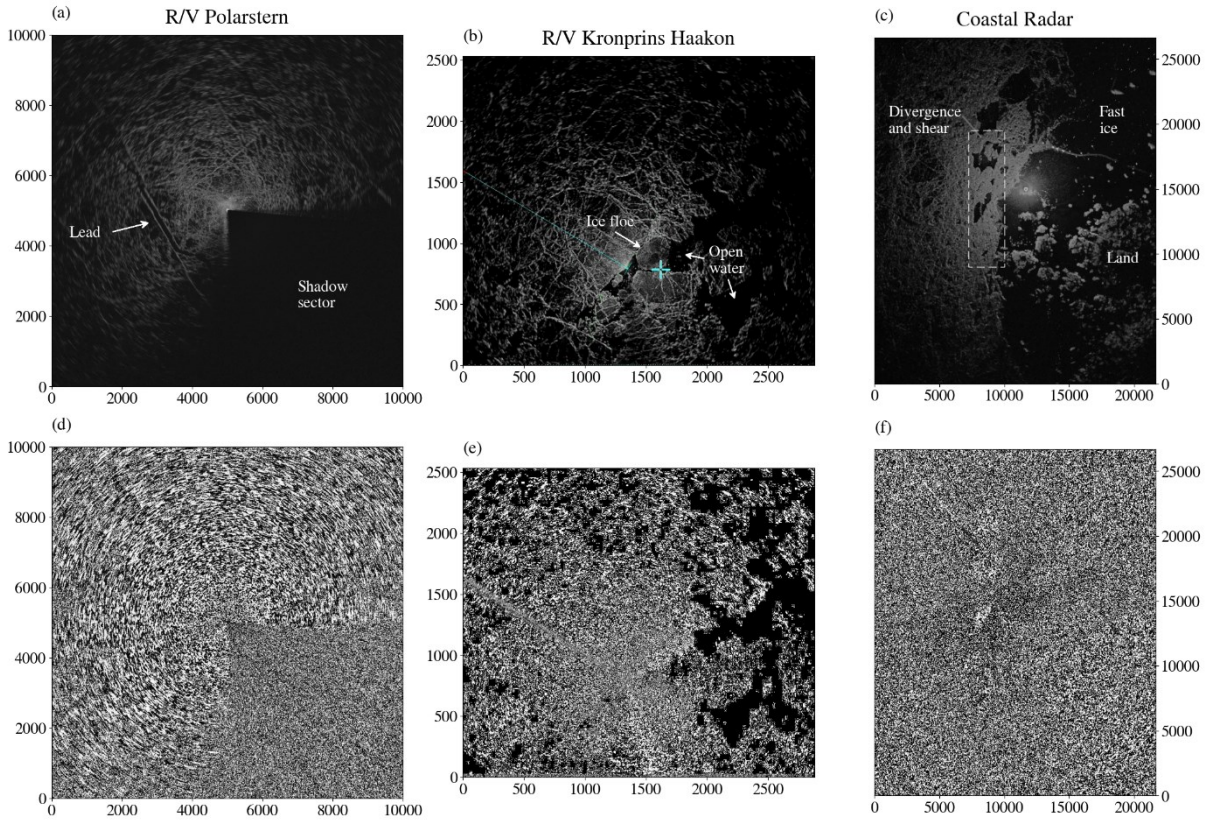

Figure 1. (a) to (c) Example frames from each dataset, (d) to (f) corresponding noise patterns.

resolution of 8.33 m, was bilinearly interpolated and averaged to 10 m resolution following Uusinoka et al. (2025b). Additionally, distant regions with weak backscatter were cropped out in the MOSAiC and CR datasets to avoid spurious estimations in areas of minimal pixel intensities. To highlight the deployable nature of our automated system, the KPH and CR datasets were used without any preprocessing. The temporal resolutions were chosen based on data availability and ice motion intensity to allow the optical flow system to capture the displacement fields more accurately. Overall, these heterogeneous data sources allow us to demonstrate that the proposed system yields consistent and robust ice motion estimates under varying conditions and noise types.

Table 1. Radar data properties for case studies. The MOSAiC dataset was used with spatial averaging resulting in the reduced image resolution.

Dataset	Cropped extent (km^2)	Image dimensions	Pixel resolution (m)	Temporal resolution (min)	Preprocessing
MOSAiC	100	(1000, 1000)	10	10	Spatial averaging based on bilinear interpolation
KPH	7	(864, 760)	3.33	5	No preprocessing
CR	578	(650, 800)	33.33	5	No preprocessing

Motion estimation with RAFT

Our motion-detection framework is based on the deep learning architecture Recurrent All-Pairs Field Transforms (RAFT), originally introduced by Teed and Deng (2020). RAFT builds dense correlation volumes between two frames and uses a recurrent update operator to iteratively refine flow estimates. RAFT has shown strong generalization and state-of-the-art accuracy on diverse optical flow benchmarks. Nonetheless, directly applying RAFT to sea ice radar imagery requires addressing the challenges posed by low signal-to-noise ratios, irregular intensities, and small displacements typical for ice movement.

To overcome these challenges, we modify RAFT by introducing a temporal resolution tree (TRT). Instead of estimat4ing the optical flow at only one temporal interval, we use coarser time steps first (30 min or longer), where larger pixel displacements are more easily identified, and then progressively refine the solution at finer time steps (10 min or shorter). This hierarchical process relies on forward warping the coarser-resolution flow fields onto the finer time-step frames, accounting for the differences in displacement magnitude and image coordinates between intervals (Gehrig et al., 2021) and is formalized in Uusinoka et al. (2025a) as

$$F_{r\to r+1}(x) = \frac{\sum_{\forall x_{r_i}} k_b\left(x - g_{r_i}(x_{r_i})\right) F_{r_i}(x_{r_i})}{\sum_{\forall x_{r_i}} k_b\left(x - g_{r_i}(x_{r_i})\right)},$$

where $g_{r_i}(x) = x + \alpha_{r_i} F_{r_i}(x)$ estimates the source pixel location adjusted by the flow at resolution r_i, $k_b(a) = \max\{0, 1 - |a_x|\} \cdot \max\{0, 1 - |a_y|\}$ is the bilinear interpolation kernel, and $\alpha_{r_i} = \Delta t_{r_{i+1}} / \Delta\, t_{r_i}$ is the temporal scaling factor. The final high-temporal-resolution estimates serve as initial estimate for RAFT's update operator, which further refines the estimated Eulerian displacement field. This ensures that even very small displacements remain detectable, despite noise or rapidly varying radar returns. Compared to other methods of initializing RAFT, TRT consistently yields higher accuracy with minimal extra computation time.

Deformation calculation

Once pixel-scale Eulerian displacements are obtained between consecutive radar frames, we compute Lagrangian measures of deformation to characterize sea ice evolution of deformation processes, including ridging, rafting, and fracturing. We form Lagrangian trajectories by bilinearly interpolating the displacement fields over time, effectively tracking each ice parcel. We define the trajectories in the simple form of

$$\mathbf{X}(i, j, t_{n+1}) = \mathbf{X}(i, j, t_n) + \tilde{\mathbf{F}}_{\boldsymbol{t_n}}(\mathbf{X}(i, j, t_n),$$

where $\mathbf{X}(i, j, t)$ is a material point for which we define the trajectory, and $\tilde{\mathbf{F}}_{\boldsymbol{t}}$ is the bilinear interpolation map over the flow field $\mathbf{F}_t$. Subsequent integration yields a material path for each pixel, and the deformation tensor is determined by mapping these trajectories back to their reference coordinates (Gurtin, 1982). The Lagrangian strain measures thus retain a history of the motion.

In this work, when referring to deformation, we use the measure of total deformation rate that is estimated based on velocity gradients associated with each deformation cell. As described in

detail in Uusinoka et al. (2025b), divergence rate, $\dot{E}_d$ (ice field opening or converging), shear rate, $\dot{E}_s$ (sliding between ice parcels), and total deformation rate, $\dot{E}_t$, are respectively estimated from the principal values, E_1 and E_2, of the Green-Lagrange strain tensor, **E**, as

$$E_d = E_1 + E_2, \quad E_s = \frac{1}{2}(E_1 - E_2), \quad \text{and} \quad E_t = \sqrt{E_d^2 + E_s^2}$$

Automated system architecture

We encapsulate our motion estimation method within a workflow (Fig. 2) designed for continuous data ingestion and minimal operator intervention. The system accepts a live feed of radar images and processes each new frame as soon as it is available. The incoming frames are buffered when needed to construct TRT at multiple time resolutions. A monitoring routine checks for missing or corrupted frames, and if errors arise, the pipeline reverts to lower temporal resolutions. This error-handling mechanism prevents abrupt disruptions in flow estimation and mitigates the impact of data dropouts. Our framework also logs intermediate output (correlation volumes, motion fields, trajectories, and deformation fields) in an archive, enabling retrospective quality assessments without halting the real-time operations. The final output can be visualized onboard for navigation decisions, integrated into a central forecasting center for ice management, or stored for subsequent scientific analysis.

The entire pipeline runs on a single CPU and GPU in standard configurations. In this work, all experiments were performed with a single NVIDIA V100 GPU and one CPU, with ICE-RAFT using a previously trained deep learning model described in Uusinoka et al. (2025a). No specialized cluster computing is required. Implementation leverages open-source libraries for GPU-accelerated convolutional neural networks and image manipulation. If processing demands increase, users may scale vertically by assigning each layer of TRT to a separate GPU. We use minimal parameter tuning within the neural network as our experiments indicate that

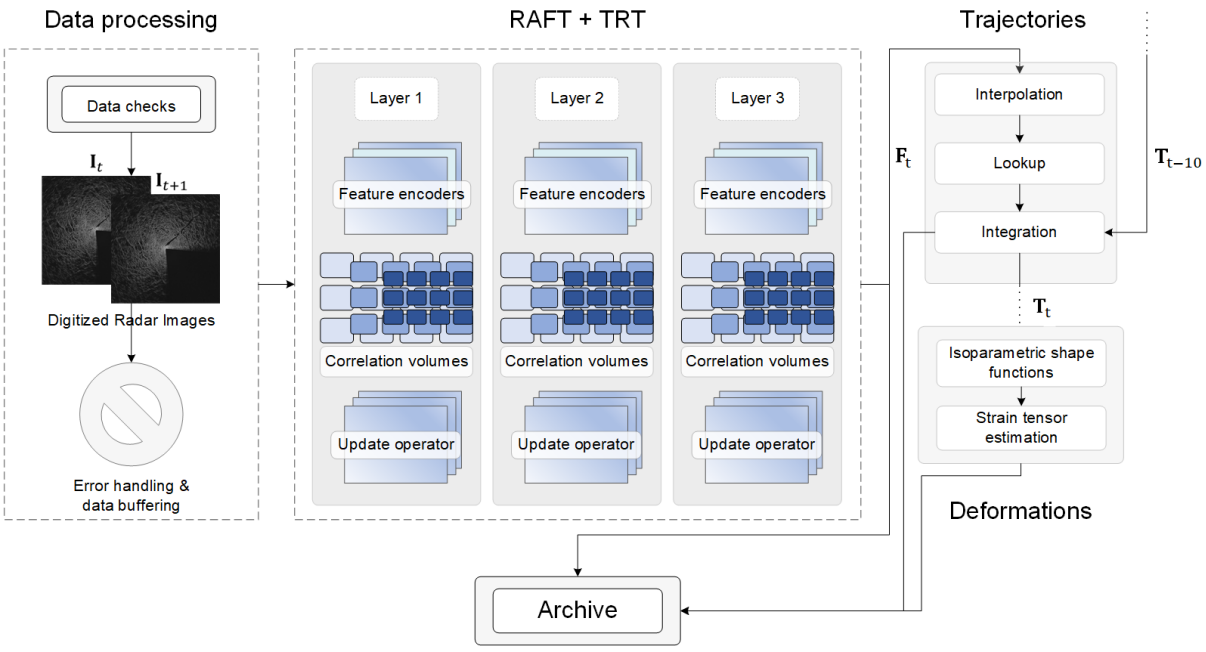

Figure 2. Visualization of the data pipeline from radar data checks to Lagrangian deformation estimates as an extension of the framework in (Uusinoka et al., 2025a).

the default model, trained according to Teed and Deng (2020), performs robustly on sea ice data without any further finetuning of the model weight. However, domain-specific finetuning is straightforward if particularly noisy or specialized radar setups are encountered.

RESULTS

Displacement fields

Fig. 2 presents the spatial distribution of cumulative Eulerian displacement estimates over 24-hour period for each of the three radar datasets. For the MOSAiC data (Fig. 2a), the proposed method captures the motion of large rigid ice bodies with minimal noise, while the discontinuities in the displacement field align well with visually identified leads, ridges, and fractures (see Fig. 1a). Despite the spatially and temporally inconsistent noise signature in this dataset (Fig. 1d), the displacement estimates accurately track ice movement, indicating that the neural network remains robust against such variability. Notably, the method detects even very slow displacements, down to the order of ten meters over the 24-hour period, which surpasses earlier conservative lower bound estimate for the method reported by Uusinoka et al. (2025a). This improved accuracy reflects the relatively simple rigid-body motion in winter pack ice conditions, where distinct fracture lines delineate clearly separable ice floes.

In contrast, the KPH dataset displays a more complex displacement field comprising multiple floes with varying velocities and directions, as well as a region of anchored ice that appears stationary relative to the ship's position. Although such complexity involves frequent collisions and shear zones, the method consistently reproduces the observed displacement gradients, confirming its ability to handle challenging these types of highly dynamic scenarios. The main artifact seems to arise in open water regions, which appear almost free of any pixel intensities in the radar imagery (Fig 1b and 1e) and consequently yield artificially high displacement values. This effect stems from a lack of consistent speckle or backscatter that the gradient-based neural network can match. Additionally, the RAFT+TRT combination is trained with the inclusion of high-frequency noise over the training set for which the stationary features are expected to include some noise patterns. This issue is partially mitigated in the MOSAiC and

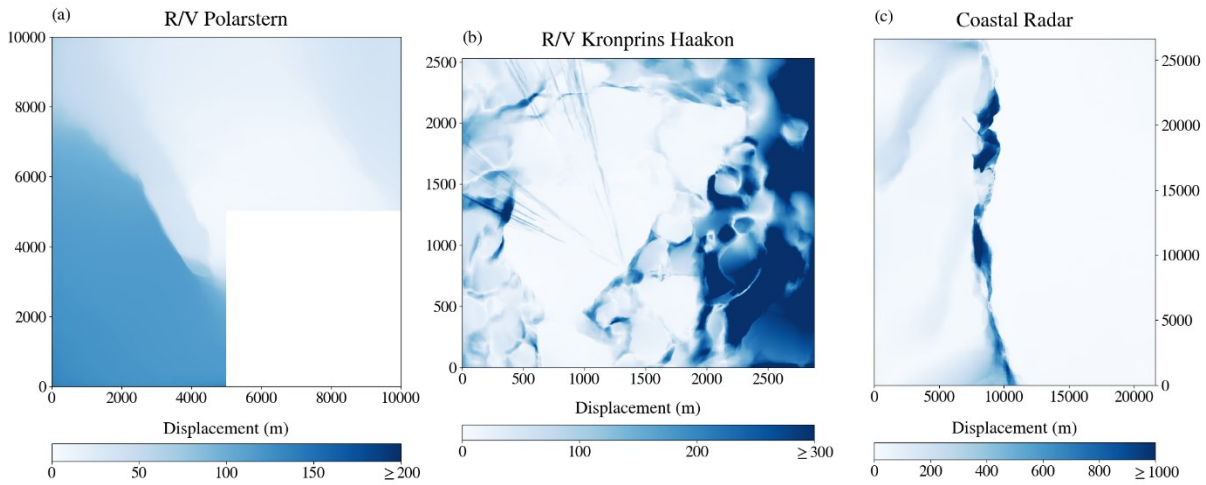

Figure 3. The Eulerian displacement fields of the three radar datasets summed over each pixel during a 24-hour period.

CR datasets by the presence of at least some uniform background noise that correlates with the icescape. Nonetheless, in ice-covered areas with less open water, the predicted motion fields exhibit close agreement with manual interpretations.

The CR dataset, again, demonstrates that the method accurately captures diverging and shearing parts of the ice motion, most prominently along a quasi-linear deformation zone. On top of this complex motion pattern, signs of compression and granular-like interactions are detected away from the main deformation zone. Regions of stationary backscatter, mainly coastal land or immobile fast ice, remain classified with effectively zero displacement, confirming that the system handles static features appropriately. Although each radar dataset differs in terms of noise profiles and ice conditions, the automated network detects ice motion accurately in all three cases with few artifacts, pointing to strong generalization of the approach.

Deformation fields

Fig. 3 shows Lagrangian total deformation rates averaged over each deformation cell during the corresponding 24-hour periods. These fields appear largely free of background noise, while localizing deformation zones with considerable precision. For the MOSAiC dataset, the major fracture line evident in Fig. 1a is clearly highlighted, accompanied by smaller ridges and leads in the vicinity. The radar's coverage also reveals ridging near the edges of the scanned region even though the quality of the radar signal is already considerably weaker at these areas.

In the KPH dataset, the evolution of ice floes become clear when viewed through deformation rates. Collisions between the rigid floes in the upper-left corner, which was initially interpreted as loosely consolidated pack ice, become evident from the velocity gradients. While open water regions produce inflated displacement estimates that result in slightly elevated deformation rates, the method is still able to distinguish discrete floes and their interaction zones with high spatial accuracy. Similarly, in the CR data, the prominent fracture zone appears as a region of intense deformation, and smaller clusters of mix-mode deformation also emerge in the vicinity of the main zone. Although these estimates match the manually observed ice behavior

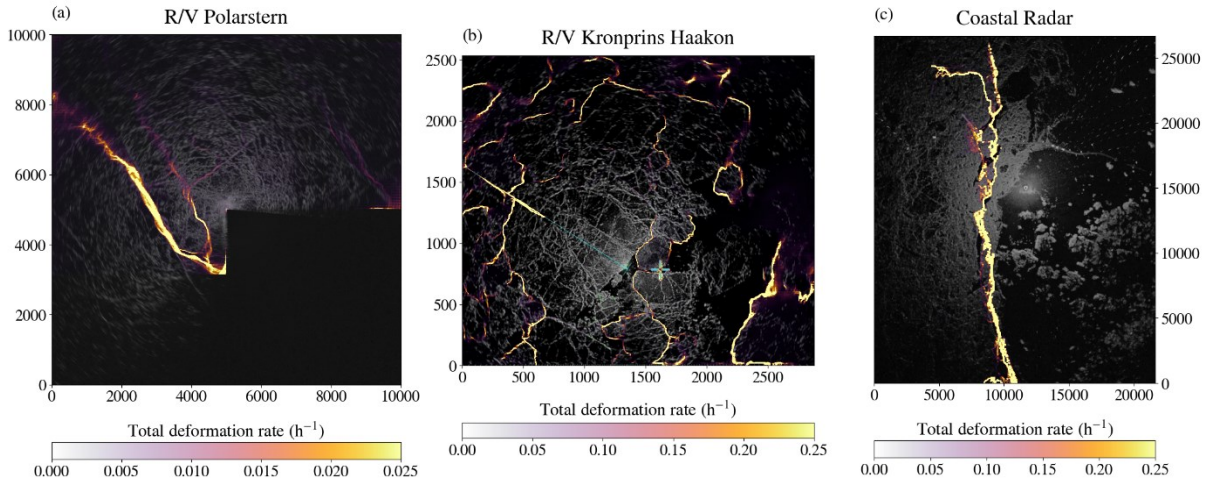

Figure 4. Lagrangian deformation rates summed over each deformation cell during a 24-hour period.

reasonably well, the temporal sampling and the corresponding large per-pixel displacements in this nearshore setting occasionally reduce spatial sharpness. This indicates the possibility of estimating deformation even with smaller temporal resolutions. Overall, however, the Lagrangian deformation fields affirm that the method delivers an accurate description of both large-scale and localized deformation processes in varying ice conditions, data types, and noise patterns.

Data pipeline performance

Table 2 summarizes the runtimes of each data pipeline component for the three datasets. Estimating motion with the RAFT+TRT combination takes between 1-2 seconds per frame, even with three temporal-resolution layers in the TRT, indicating the feasibility of real-time processing even when considering the rate of the radar cycle. Compiling these flow fields into a single, continuous displacement record consumes an additional 0.2-0.4 seconds. The largest computational cost arises from the trajectory formulation component, which can reach 16 seconds per frame with the MOSAiC dataset at the highest spatial resolution, due to the one million trajectories requiring bilinear interpolation. As the bilinear interpolation is performed with a factor of 10 in the 2D spatial domain, we interpolated displacement fields consist of 100 000 000 datapoints. Coarse graining the displacement fields, applying parallelized implementations, or optimizing the implemented code could substantially lower this overhead without compromising accuracy. Deformation calculations, which involve mapping velocity gradients to Lagrangian strain measures, add 3-5.5 seconds. These findings indicate that total runtimes range from approximately 12 seconds per frame for CR to around 23.5 seconds per frame for MOSAiC. Although radar cycles take only a few seconds, the digitizing system is typically run to generate higher quality radar data every few minutes. For the digitized radar data, the processing speed with even the MOSAiC data is sufficient for continuous or near-continuous deployment on a single CPU and GPU. Real-time operational use is thus deemed viable in practical situations.

Table 2. Runtimes for each component in the data pipeline.

Dataset	Motion estimate (sec per frame)	Number of TRT layers	Displacement maps (sec per frame)	Trajectory maps (sec per frame)	Number of trajectories	Deformation maps (sec per frame)	Total runtime (sec per frame)
MOSAiC	1.86 ± 0.92	3	0.43 ± 0.20	15.67 ± 0.70	1000 000	5.45 ± 0.25	23.41
KPH	0.97±0.35	3	0.22 0 ± 0.14	10.0 ± 0.6	656 640	3.78 ± 0.11	14.97
CR	1.65 ± 0.72	3	0.29 ± 0.15	7.0 ± 0.56	520 000	2.92 ± 0.11	11.86

CONCLUSIONS

In this work, we demonstrate that a deep learning-based optical flow system, extended by a temporal resolution tree (TRT) and coupled with an automated data pipeline, effectively captures detailed sea ice dynamics suitable for operational surveillance across diverse radar imaging scenarios. By applying the extended form of the RAFT architecture to shipborne and coastal radar datasets, we capture coherent displacement fields with minimal noise even in the presence of complex ice motions or varying radar backscatter. The results confirm that important spatial discontinuities in ice motion for navigation or surveillance purposes, such as leads, ridges, and fractures, are recognized reliably. While our results confirm method accuracy

qualitatively, future work should include quantitative validation regarding ice convergence and divergence characterization.

The proposed framework handles a range of radar image qualities and temporal sampling intervals, while retaining near-real-time performance. Motion estimation can be carried out in under two seconds per frame on a single GPU with one million datapoints, showing promise for integration into operational systems. The largest computational costs occur in trajectory formulation and subsequent deformation calculations, but these can be substantially reduced through optimization techniques. As a result, the method can scale to larger spatial domains or denser grids without compromising the resolution or accuracy of the estimated flow fields. This approach opens new avenues for better situational awareness and safety in ice-covered waters, supporting navigation, surveillance, and scientific research. Future research will focus on refining these computational steps, optimizing the pipeline efficiency, and exploring network training strategies to enhance robustness against the non-traditional radar noise patterns.

ACKNOWLEDGEMENTS

MU and AP are grateful for financial support from the Research Council of Finland through the project (347802) DEMFLO: Discrete Element Modeling of Continuous Ice Floes and Their Interaction. Contribution of JH was covered by the the European Union's Horizon 2020 research and innovation programme under grant agreement No 101003826 via project CRiceS (Climate Relevant interactions and feedbacks: the key role of sea ice and Snow in the polar and global climate system). All authors wish to acknowledge CSC – IT Center for Science, Finland, for computational resources under the project (2006428) DEMFLO, the international Multidisciplinary drifting Observatory for the Study of the Arctic Climate (MOSAiC) project with the tag MOSAiC20192020 and the Project_ID:AWI_PS122_00 for providing the ship radar data, and Norwegian Polar Institute for providing the ship radar data from the NPI Arctic Ocean Cruise II.

REFERENCES

Chu, T. C., Ranson, W. F., & Sutton, M. A. (1985). Applications of digital-image-correlation techniques to experimental mechanics. *Experimental mechanics*, 25, pp. 232-244.

Gehrig, M., Millhäusler, M., Gehrig, D., & Scaramuzza, D. 2021. E-raft: Dense optical flow from event cameras. *International Conference on 3D Vision (3DV)* (pp. 197-206). IEEE.

Gurtin, M. E. 1982. An introduction to continuum mechanics. Academic press.

Hop, H., Wold, A., & Misund, O. A. 2024. NPI Arctic Ocean Cruise II, 10–29 August 2023: IMR cruise ID 2023007014. *Norwegian Polar Institute*.

Karvonen, J., 2016. Virtual radar ice buoys–a method for measuring fine-scale sea ice drift. *The Cryosphere*, 10.1, pp.29-42.

Lucas, B. D., & Kanade, T. (1981, August). An iterative image registration technique with an application to stereo vision. In *IJCAI'81: 7th international joint conference on Artificial intelligence*, 2, pp. 674-679.

Nicolaus, M., Perovich, D. K., Spreen, G., Granskog, M. A., von Albedyll, L., Angelopoulos, M., ... & Wendisch, M. 2022. Overview of the MOSAiC expedition: Snow and sea ice. *Elem Sci Anth*, 10(1), 000046.

Ninnis, R. M., Emery, W. J., & Collins, M. J. (1986). Automated extraction of pack ice motion from advanced very high resolution radiometer imagery. *Journal of Geophysical Research: Oceans*, *91*(C9), pp. 10725-10734.

Oikkonen, A., Haapala, J., Lensu, M., & Karvonen, J. (2016). Sea ice drift and deformation in the coastal boundary zone. *Geophysical Research Letters*, *43*(19), pp.10-303.

Oikkonen, A., Haapala, J., Lensu, M., Karvonen, J., & Itkin, P. 2017. Small-scale sea ice deformation during N-ICE 2015: From compact pack ice to marginal ice zone. *Journal of Geophysical Research: Oceans*, 122(6), pp. 5105-5120.

Peters, W. H., & Ranson, W. F. (1982). Digital imaging techniques in experimental stress analysis. *Optical engineering*, 21(3), pp. 427-431.

Teed, Z., & Jia D. 2020. Raft: Recurrent all-pairs field transforms for optical flow. *Computer Vision–ECCV 2020: 16th European Conference*. pp. 402-419. Springer International Publishing.

Thomas, M., Kambhamettu, C., & Geiger, C. A. (2011). Motion tracking of discontinuous sea ice. *IEEE transactions on geoscience and remote sensing*, 49(12), pp. 5064-5079.

Uusinoka, M., Haapala, J., & Polojärvi, A., 2025. Deep learning-based optical flow in fine-scale deformation mapping of sea ice dynamics. *Geophysical Research Letters*, 52.2. https://doi.org/10.1029/2024GL112000.

Uusinoka, M., Haapala, J., Åström, J., Lensu, M., & Polojärvi, A., 2025. Scale invariance in kilometer-scale sea ice deformation. *The Cryosphere*. https://doi.org/10.5194/egusphere-2025-311.