Satellite Products and Services Review Board

Algorithm Theoretical Basis Document:

Advanced Scatterometer Ultra High-Resolution Autonomous Tropical Cyclone Center Fix via Deep Learning Object Detection

Compiled by:
NOAA/NESDIS/STAR Ocean Surface Winds Team

Version 1.0

April, 2026

AUTHORS:

Paul S. Chang (NOAA/NESDIS/STAR): *paul.s.chang@noaa.gov*

Zorana Jelenak (UCAR): *zorana.jelenak@noaa.gov*

Matthew McKinney (Global Science & Technology): *matthew.mckinney@noaa.gov*

Seubson Soisuvarn (UCAR): s*eubson.soisuvarn@noaa.gov*

Suleiman Alsweiss (Global Science & Technology): *suleiman.alsweiss@noaa.gov*

LIST OF CHANGES

DOCUMENT TITLE: Advanced Scatterometer UHR TC Center Fix			
DOCUMENT CHANGE HISTORY			
Revision No.	**Date**	**Revision Originator Project Group**	**CCR Approval # and Date**
1.0	04/2026	M. McKinney	N/A

LIST OF CHANGES

DOCUMENT TITLE: Advanced Scatterometer UHR TC Center Fix					
LIST OF CHANGE-AFFECTED PAGES/SECTIONS/APPENDICES					
Version Number	**Date**	**Changed By**	**Page**	**Section**	**Description of Change(s)**
1.0	04/2026	M. McKinney		All	Initial documentation

TABLE OF CONTENTS

LIST OF ACRONYMS

AFRC: Air Force Reserve Command

ASCAT: Advanced Scatterometer

ATBD: Algorithm Theoretical Basis Document

ATCF: Automated Tropical Cyclone Forecasting

IBTrACS: International Best Track Archive for Climate Stewardship

KDE: Kernel Density Estimation

NESDIS: National Environmental Satellite, Data, and Information Service

NHC: National Hurricane Center

NOAA: National Oceanic and Atmospheric Administration

NRCS: Normalized Radar Cross Section

NRT: Near Real-Time

STAR: Center for Satellite Applications and Research

TC: Tropical Cyclone

UHR: Ultra High Resolution

VDM: Vortex Data Message

YOLO: You Only Look Once

Section 1

Ultra High-Resolution Autonomous Tropical Cyclone Center Fix Algorithm Description

1.1. Instrument Overview and Technical Specifications

The Advanced Scatterometer (ASCAT) is a vertically polarized active microwave radar scatterometer instrument that operates at C-band, specifically at a frequency of 5.255 GHz. It was designed to measure near-surface wind vectors over the global ocean. The instrument transmits and receives radar pulses using two sets of three fixed fan-beam antennas (fore, mid, and aft beams) located on each side of the satellite ground track. The ASCAT instrument series has been flown on board the European Meteorological Operational (MetOp) polar-orbiting satellites, including MetOp-A, MetOp-B, and MetOp-C. While standard near real-time operational ASCAT ocean surface wind vectors are traditionally produced at 12.5 km and 25 km swath grid resolutions, these coarser resolutions leave large data voids in inner coastal zones. This document describes the theoretical basis and algorithms used to generate specialized Ultra High-Resolution (UHR) products, expanding the data's utility for coastal winds, sea ice detection, and tropical cyclone analysis.

Accurate determination of tropical cyclone (TC) center locations is a fundamental requirement for intensity forecasting, structural analysis, and downstream wind product derivation. The ASCAT UHR product provides ocean surface wind vectors at approximately 1 to 2.5 km resolution (Lindsley, 2010; Soisuvarn, 2023), capable of resolving the inner core structure of tropical cyclones through its measurement of the normalized radar cross section (NRCS, or σ^0).

Current methods for determining TC center locations from scatterometer data rely on either temporal and spatial interpolation of infrequent “best track” positions or manual expert interpretation of convergence patterns in the backscatter fields. The former approach introduces temporal aliasing errors that grow with the time gap between observations, and neither are scalable for real-time operational processing.

This ATBD describes an automated TC center fix algorithm that employs deep learning object detection to identify storm center locations directly from ASCAT UHR σ^0 fore-aft difference ($\Delta\sigma^0$) imagery. The algorithm uses a YOLOv11 Nano model trained on 2,490 hurricane passes spanning 2020–2025, combined with σ^0 normalization and Gaussian kernel density estimation to produce probabilistic center location estimates. The system is designed to operate on resource-constrained hardware (single NVIDIA RTX 3000 Quadro GPU with 6 GB VRAM), enabling deployment at operational centers without dedicated GPU clusters.

1.2. Data Preparation and Image Generation

1.2.1. 2D Binned Image Generation

The primary input to the center detection algorithm is the σ^0 fore-aft difference field ($\Delta\sigma^0$) from ASCAT UHR NetCDF files. The data processing pipeline loads longitude, latitude, and $\Delta\sigma^0$ arrays, applies a validity mask where the number of ambiguities equals zero, and normalizes longitudes to the [–180°, 180°] range.

The $\Delta\sigma^0$ field is converted to RGB images through an 2D binning algorithm. The global swath data is binned onto a regular grid using SciPy's binned_statistic_2d function with a mean aggregation statistic. The grid resolution is configured at 3.2 degrees per chunk, with each chunk producing a 220 × 220 pixel image (110 pixels per grid cell with 50% overlap between adjacent chunks). The $\Delta\sigma^0$ values are color-mapped using a divergent blue-white-red colormap with fixed limits of [–15, +15] dB.

Optional Gaussian noise is added to the binned values for data augmentation during training: val_matrix += ε × N(0, σ(val_matrix)), where ε is a configurable noise parameter (default: 0.05). The chunks retain their geographic metadata (longitude and latitude bin edges) for subsequent coordinate projection of YOLO detections back to geographic space.

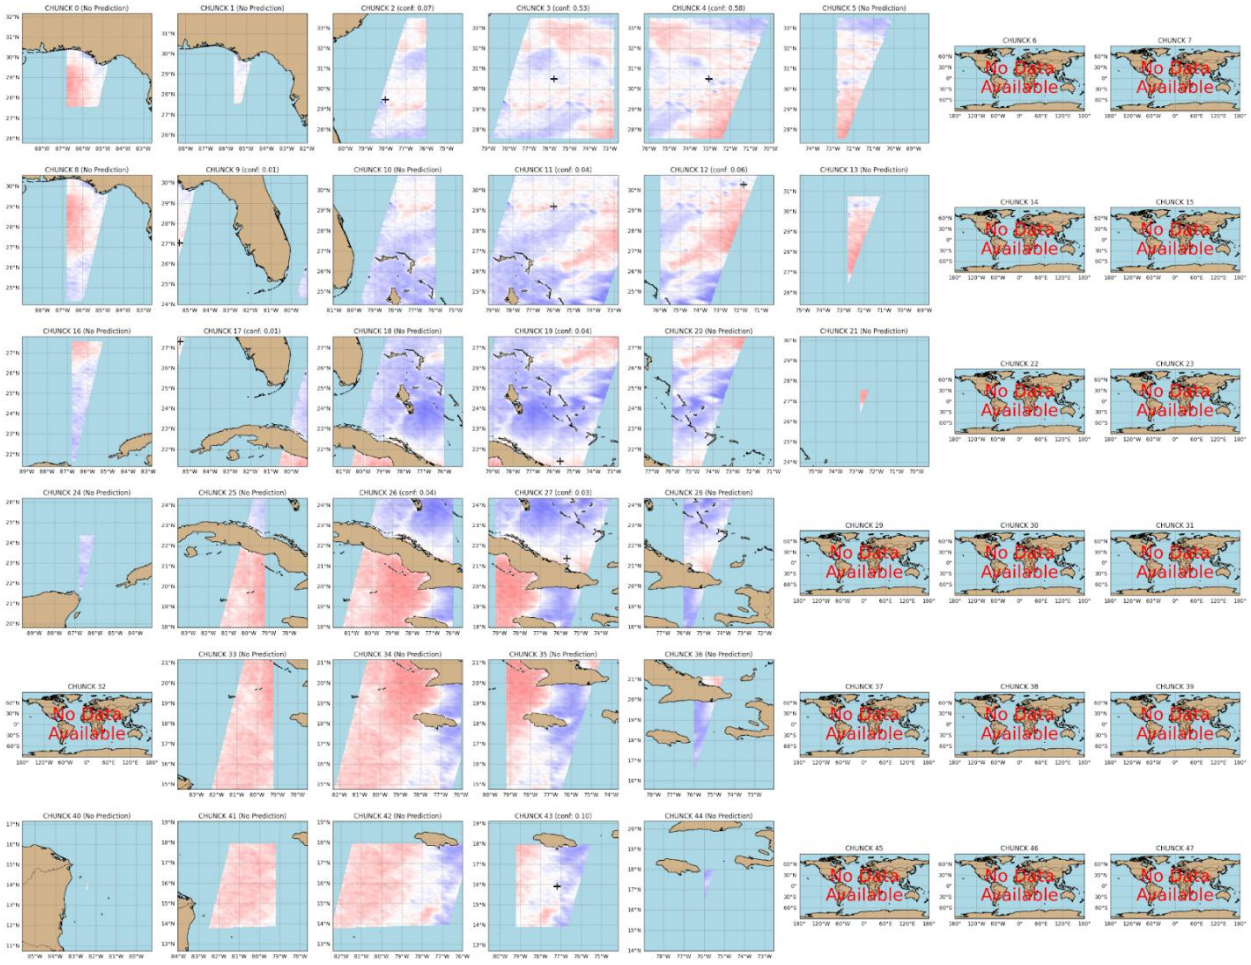

Figure 1-1: Example of input images collected from an ASCAT-B UHR $\Delta\sigma^0$ pass over Melissa on October 29th, 2025. The pass is split into many chunks that can each be individually processed by the AI model. Each chunk is 220 x 220 pixels with 110 pixels overlapping with each neighboring image.

1.2.2. Forecast-Constrained Image Generation

For operational deployment, a forecast-constrained image generator variant produces images only within a configurable radius (default: 100 km) of the NHC or ATCF forecast track position. The algorithm computes the pass median time, interpolates the forecast track to that time, and applies a Haversine distance filter to skip chunks whose centers exceed the search radius. This spatial pruning dramatically reduces the number of images processed, improving throughput for near real-time applications. If no forecast track is available, the algorithm falls back to full-swath image generation.

1.3. YOLO Object Detection Model

The object detection model is based on the YOLOv11 Nano (YOLOv11n) architecture from Ultralytics, a single-stage detector optimized for edge deployment. The model file size is approximately 5.5 MB, enabling inference on resource-constrained GPUs (Jocher, 2024). The model performs single-class object detection where the target class is “center” (the TC eye or circulation center). The detection output for each image is a set of bounding boxes with associated confidence scores, from which the box center coordinates are extracted as candidate TC center positions.

1.3.1. Training Dataset

The training dataset comprises 2,490 ASCAT UHR hurricane passes from 2020 through March 2025, spanning multiple ocean basins. Each pass is processed into one or more 220 × 220 pixel $\Delta\sigma^0$ images with corresponding bounding box labels. Ground truth labels are generated by linearly interpolating IBTrACS best track center positions to the satellite pass time and projecting the interpolated center onto the image coordinate system (Knapp, 2010).

The dataset is partitioned chronologically: passes from 2020–2023 (1,917 images) serve as the training set, while passes from 2024–March 2025 (573 images) serve as the validation set. This temporal boundary simulates near real-time deployment conditions where the model has never seen storms from the evaluation period. The dataset includes all tropical depressions, tropical storm, and hurricanes with available best tracks, with Saffir-Simpson thresholds applied to the IBTrACS maximum sustained wind speed.

Table 1-1: Saffir-Simpson Category Wind Speed Thresholds

Category	Minimum (kt)	Maximum (kt)
Tropical Storm (TS)	34	63
Category 1	64	82
Category 2	83	95
Category 3	96	112
Category 4	113	136
Category 5	137	—

1.3.2. Deep Ensemble Training

The training runs 1,000 epochs, 220 × 220 pixel image size, single-class detection mode, and AdamW optimizer with cosine learning rate scheduling. To reduce epistemic uncertainty in center predictions, a deep ensemble of N = 13 models is trained (Lakshminarayanan, 2017). Each ensemble member is trained on a random 50% subsample of the training data (fraction = 0.5), with diversity further encouraged through: randomized seeds (uniform integers in [0, 10^8]), dropout rate of 0.1, multi-scale augmentation, and non-deterministic operations. Each ensemble member produces an independent set of bounding box detections, and the aggregated detections across all 13 models provide a distribution of center estimates from which the model's uncertainty can be quantified (Lakshminarayanan, 2017).

Table 1-2: Deep Ensemble Training Configuration

Parameter	Value
Ensemble Size	10
Training Fraction	0.5 (50% per model)
Epochs	1,000
Image Size	220 × 220 pixels
Optimizer	AdamW
Learning Rate Schedule	Cosine annealing
Dropout	0.1
Label Smoothing	0.1
Multi-Scale	Enabled
Deterministic	Disabled
Detection Mode	Single class (center)

1.4. Likelihood Estimation

The center location algorithm combines two independent likelihood fields: (1) a KDE likelihood derived from YOLO detection coordinates across the ensemble members, and (2) a σ^0 likelihood derived from the normalized backscatter fields. The two likelihoods are fused through element-wise multiplication to produce a combined likelihood field that leverages both the deep learning-based pattern recognition and the physically anchor to backscatter.

1.4.1. KDE Likelihood from YOLO Detections

YOLO bounding box center coordinates are projected from image pixel space back to geographic coordinates using the stored chunk metadata:

$$lon = \frac{x}{W}(lon_{max} - lon_{min}) + lon_{min}$$

$$lat = lat_{max} - \frac{y}{H}(lat_{max} - lat_{min})$$

where (x, y) are the pixel coordinates of the bounding box center, W and H are the image dimensions (220 pixels), and the lon/lat bounds come from the chunk's binning grid. Detections from all ensemble members are merged, with their associated confidence scores serving as KDE weights. A Gaussian kernel density estimate is constructed over the geographic detection coordinates using SciPy's gaussian_kde function with a fixed bandwidth of 0.1. The KDE is evaluated on the full swath coordinate grid and normalized to sum to unity.

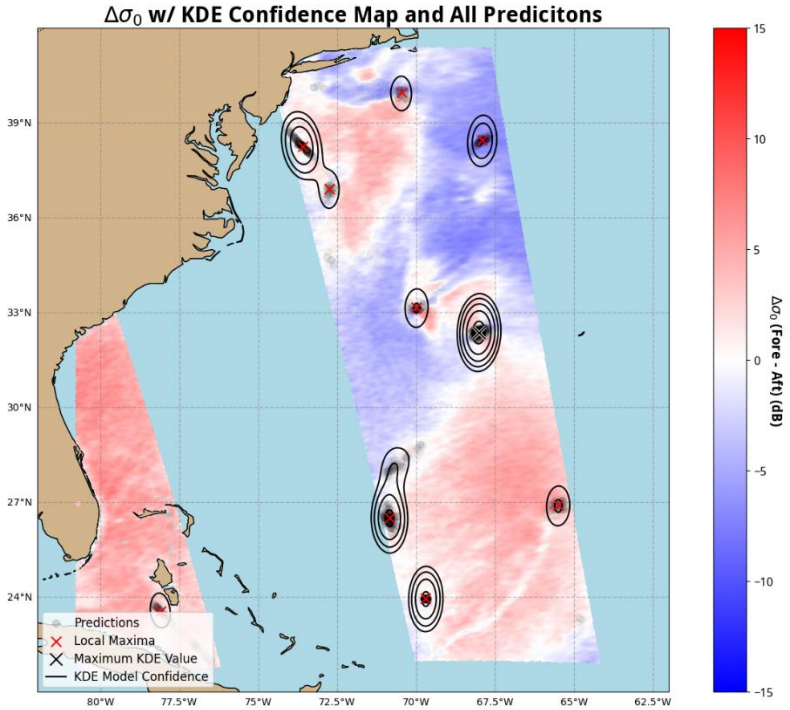

Figure 1-2: ASCAT-B ascending Δσ0 pass captured on October 31st, 2025. Solid black contours delineate the KDE model probabilities. The KDE maxima are marked by red "X". The highest KDE probability prediction is indicated by a black and white "X".

1.4.2. Normalized σ^0 Likelihood

An independent likelihood field is derived from normalizing ASCAT σ^0. The processing pipeline proceeds as follows:

(1) σ^0 normalization: dB values are converted to linear units via $\sigma_{linear} = 10^{\sigma^0/10}$. An incidence angle correction is applied: $\sigma_{corrected} = \sigma^0 - \text{slope} \times (\text{inc} - \theta_{ref})$, where slope = −0.23 dB/degree and $\theta_{ref} = 55°$.
(2) Multi-beam summation: The corrected σ^0-0 values from the fore, mid, and aft beams are summed: $\sigma_{corrected\_sum} = \sigma_{fore\_corrected} + \sigma_{mid\_corrected} + \sigma_{aft\_corrected}$.
(3) Thresholding and normalization: Values below the 90th percentile and below 1.0×10^{-5} are set to zero. A final Gaussian filter (variance = 50) is applied for spatial smoothing, followed by normalization to sum to unity.

1.4.3. Combined Likelihood

The combined likelihood is computed as the element-wise product of the normalized σ^0 likelihood and the KDE likelihood:

$$L_{combined} = L_{\sigma_0} \times L_{kde}$$

This multiplicative fusion ensures that the combined likelihood is high only where both the deep learning detections and the physics-based backscatter structure agree on the center location. The product is smoothed with a Gaussian filter ($\sigma = 10$) and normalized to sum to unity.

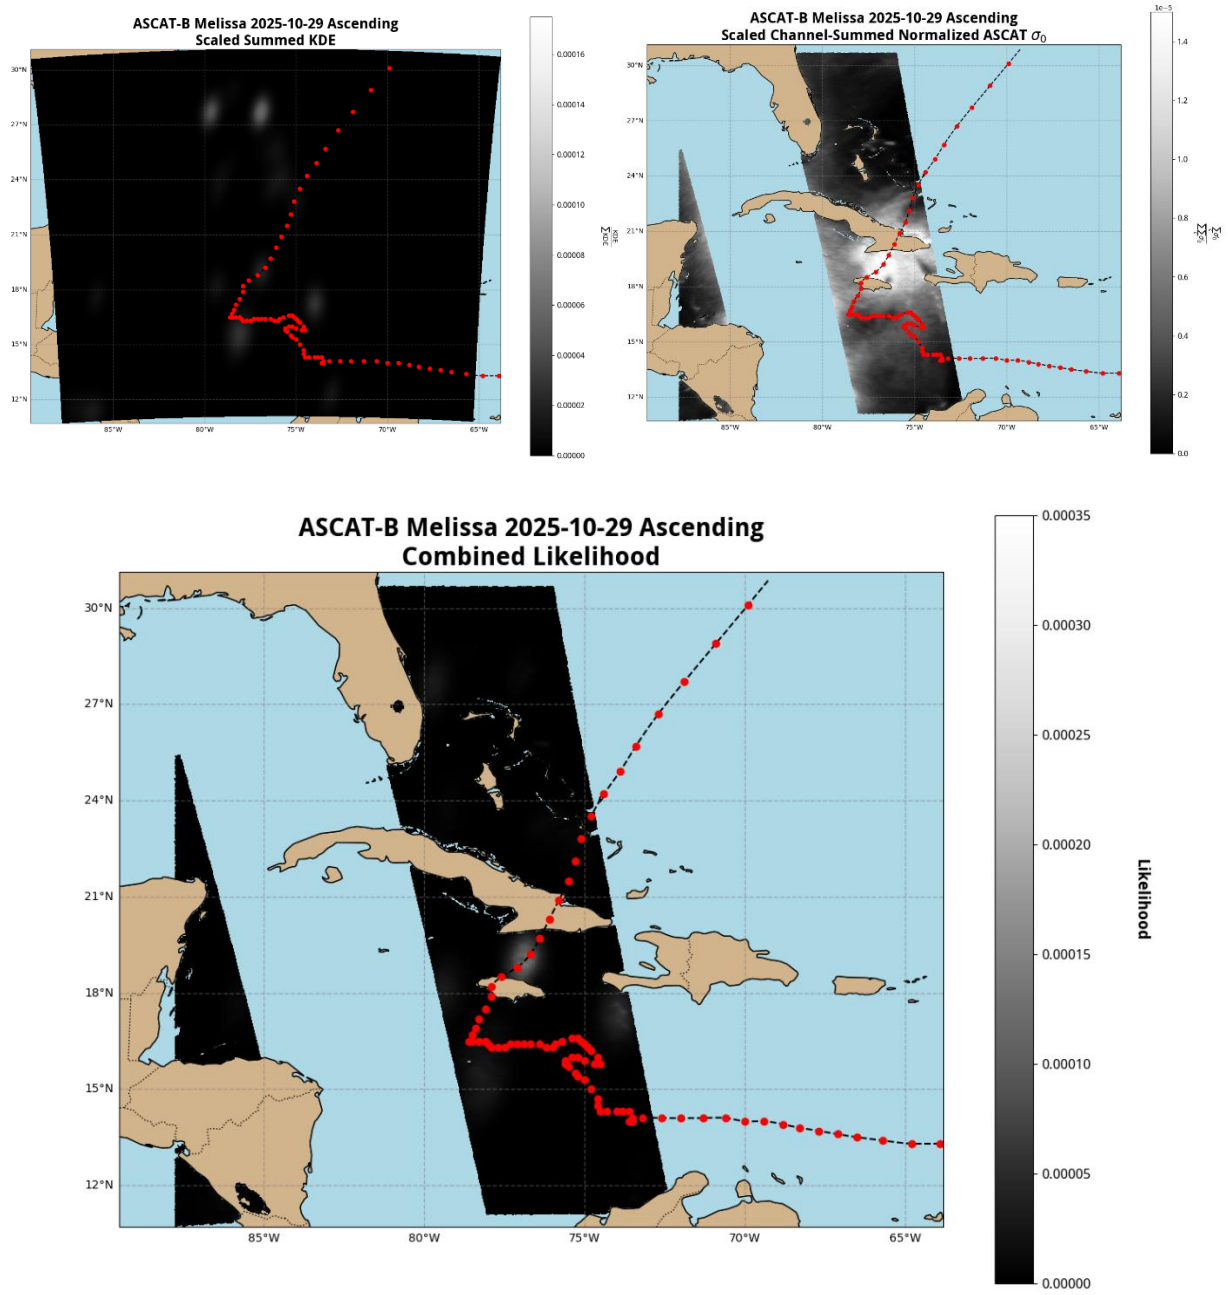

Figure 1-3: ASCAT-B ascending pass captured on October 29th, 2025. Top Left: KDE model output. Top Right: Summed normalized ⟂0. Bottom: Combined likelihood.

1.5. Center Location Algorithms

1.5.1. KDE Maxima Detection

Local maxima in the KDE likelihood field are identified using SciPy's maximum_filter with a window size of 13 pixels. A cell is classified as a local maximum if its value equals the local maximum within the window and the combined likelihood at that cell exceeds 1.0×10^{-4}. The geographic coordinates of all qualifying maxima are recorded as candidate center positions.

1.5.2. Combined Likelihood Maximum

The candidate center position with the highest combined likelihood is identified as the primary center estimate. This represents the location where the deep learning-based KDE and the normalized σ^0 most strongly agree.

1.5.3. Forecast-Bounded Solutions

Operationally, center estimates are constrained to lie within a forecast uncertainty envelope. For each candidate center from the KDE, the Haversine distance to the most recent ATCF or NHC forecast position is computed. The forecast-bounded solution selects the candidate with the minimum distance to the forecast track, providing a constrained center estimate that mitigates false detections far from the expected storm location. By comparing the forecast bounded solution and the primary center estimate with the highest combined likelihood, there can be greater awareness of instability in either the ATCF forecast or the model prediction.

1.5.4. σ^0 Minima Detection

In additional method is being explored to retrieve an even more refined center fix via the normalized σ^0 minima. Local maxima in the normalized field σ^0 are identified using SciPy's minimum_filter with a window size of 13 pixels. A cell is classified as a local minimum if its value equals the local minimum within the window and the combined likelihood at that cell exceeds 1.0×10^{-4}. These points center fixes act very similar to KDE maxima candidate centers, but are more vulnerable to noise in the σ^0 field.

1.6. Reference Track Systems

1.6.1. IBTrACS Best Track

The International Best Track Archive for Climate Stewardship (IBTrACS) v04r01 provides post-season reanalyzed TC center positions at 3- to 6-hourly intervals (Knapp, 2010). The implementation uses lazy loading to handle the large NetCDF dataset, filtering by storm name and season year. B-spline interpolation ($k = 3$, requiring ≥ 4 data points) is used to generate smooth continuous tracks between reporting times.

1.6.2. NHC/ATCF Forecast Tracks

Two sources of real-time forecast tracks are supported. NHC forecast advisory files are parsed using regex-based extraction of initialization time, lead time, latitude, longitude, and maximum wind speed. The most recent initialization before the target prediction time is selected. ATCF A-Deck files provide multi-model objective aid forecasts, from which the best available technique is selected according to a hierarchical priority: OFCI/OFCL (official), HCCA, TVCN, EAII, AVNO, and others. The selected technique's forecast track is interpolated to the prediction time using B-spline interpolation.

1.6.3. Vortex Data Messages (VDM)

Vortex Data Messages from NOAA and Air Force Reserve Command (AFRC) Hurricane Hunter reconnaissance flights provide in-situ TC center fixes. VDM tracks are filtered by storm name and date and interpolated using linear interpolation. VDM data serves as an independent validation reference with higher accuracy than IBTrACS, although it is scarce in comparison to IBTrACS.

1.6.4. Temporal Interpolation Methods

Two interpolation methods are available for all track sources. Linear interpolation (numpy interp1d) requires a minimum of 2 data points and is used for VDM tracks where data is inconsistent and limited. B-spline interpolation (scipy splprep/splev, k = 3) requires a minimum of 4 data points and provides smoother continuous tracks, used for IBTrACS and ATCF forecast tracks.

1.7. Evaluation and Validation

The evaluation pipeline processes ASCAT UHR passes from 2025 across both ASCAT-B and ASCAT-C instruments. Passes are discovered using regex-based filename parsing with deduplication logic to retain only the latest version of each unique pass. INVEST-designated systems are excluded. A parallel processing framework (ProcessPoolExecutor with 15 workers) enables batch evaluation of the full archive.

Quality control filters are applied during evaluation: passes where the interpolated center falls outside the swath data or in the extended swath region (where L2B model wind speeds are unavailable) are flagged. Wind speeds exceeding 132.2 knots are treated as anomalous and set to NaN due to an error that exists in ASCAT UHR processing.

All distance errors are computed using the Haversine formula on a spherical Earth (R = 6,371 km):

$$d = 2R \cdot asin(\sqrt{a})$$

$$a = sin^2(\Delta lat/2) + cos(lat_1) \times cos(lat_2) \times sin^2(\Delta lon/2)$$

Errors are computed between each center estimate method (KDE maxima with highest combined likelihood, KDE maxima closest to ATCF forecast) and each reference track (IBTrACS, VDM, ATCF forecast).

A VDM validation engine provides baseline error statistics by comparing VDM flight-observed center positions against IBTrACS best track interpolations at the same times. This baseline quantifies the inherent uncertainty in the IBTrACS validation, which represents a lower bound on the achievable prediction accuracy when using IBTrACS as ground truth. The need to measure error in IBTrACS is necessary since it is the baseline metric used to evaluate our center fixes. An example of error in IBTrACS is shown in Figure 1-4.

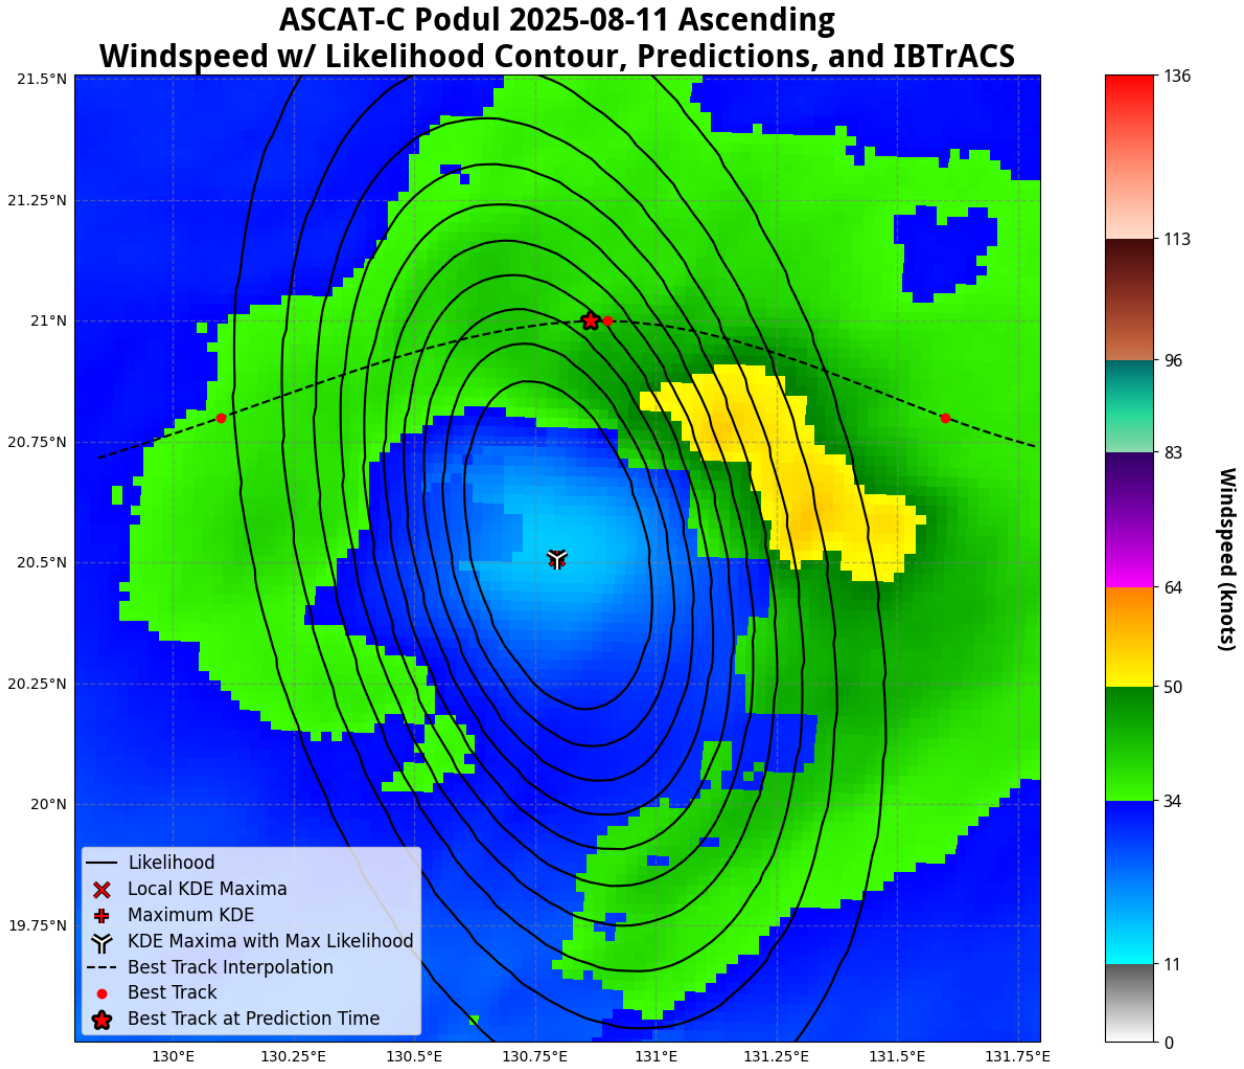

Figure 1-4: Zoomed in ASCAT-C ascending windspeed pass captured on August 11th, 2025. Solid black contours delineate the spatial likelihood selection metric. The modal cluster of the model's center predictions, derived via Kernel Density Estimation (KDE), is marked by a red "X", with the highest-confidence prediction denoted by a red "+". The highest likelihood prediction is indicated by a white "Y". The IBTrACS centers are overlaid as a dashed black line; red circles represent the tracking estimates, and the red star marks the interpolated center at the time of the ASCAT overpass. This image emphasizes the underlying error with an IBTrACS validation metric, since the interpolated best track is 55 km off from the model prediction.

For each evaluated pass, the pipeline generates: (1) nine visualization panels showing the $\Delta\sigma^0$ field, KDE likelihood, σ^0-0 likelihood, combined likelihood, and wind speed/direction fields at both standard and zoomed scales with overlaid center estimates and reference tracks; (2) a metrics CSV file recording Haversine distances between all center estimates and reference tracks; and (3) an interactive web-based viewer for navigating storm visualization data with persistent zoom and pan capabilities.

1.8. References

EUMETSAT. (2014). *Metop-A ASCAT L1 Data Record Release 2 (CF-003): User Guide.* Darmstadt, Germany: European Organisation for the Exploitation of Meteorological Satellites. Retrieved from https://user.eumetsat.int/s3/eup-strapi-media/pdf_ma_ascat_l1_cdr_003_user_guide_a6c37bce29.pdf

Jocher, G. (2024). Ultralytics YOLO11. GitHub. Retrieved from https://github.com/ultralytics/ultralytics

Knapp, K. (2010). The International Best Track Archive for Climate Stewardship (IBTrACS): Unifying Tropical Cyclone Data. *91*(3), 363 - 376. doi:10.1175/2009BAMS2755.1

Lakshminarayanan, B. (2017). Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. Retrieved from https://arxiv.org/abs/1612.01474

Lindsley, R. (2010). Adapting the sir algorithm to ASCAT. 3402-3405. doi:10.1109/IGARSS.2010.5650207

Lindsley, R. D., & Long, D. G. (2016). Enhanced-Resolution Reconstruction of ASCAT Backscatter Measurements. *IEEE Transactions on Geoscience and Remote Sensing, 54*(5), 2589-2601.

Soisuvarn, S. (2013). CMOD5.H—A High Wind Geophysical Model Function for C-Band Vertically Polarized Satellite Scatterometer Measurements. *51*(6), 3744-3760. doi:10.1109/TGRS.2012.2219871

Soisuvarn, S. (2023). High-Resolution Coastal Winds From the NOAA Near Real-Time ASCAT Processor. *61*, 1-12. doi:10.1109/TGRS.2023.3279764