Optimizing Meteorological Station Placement for High-Resolution Field Reconstruction in Mountainous Terrain

Anna Poltronieri [1] Nikolas Olson Aksamit [1]
[1]Department of Mathematics and Statistics, UiT, Norway

UiT The Arctic University of Norway

Overview

Goal: Optimally place sensors to capture independent information streams, and use them to reconstruct the full wind field over complex mountainous terrain from sparse observations.

What we do: We generalize a compressed sensing framework [1], using PALM [2] simulations as a surrogate training library, to identify optimal station locations under real-world constraints (terrain accessibility, fixed sensors). We validate that independence is preserved in practice and demonstrate full wind field reconstruction from four stations at Ullstinden, Norway.

Mountain stations: Gill WindSonic4 2D ultrasonic anemometer, 40 Hz sampling, 5-minute averages, Tilsig IoT 5G data logger, 2.5 m mast.

Valley station: propeller anemometer, 1-minute averages, full meteorological station for snow science, 3 m tripod.

Base framework

1. **Training library:** PALM simulations for 8 inflow directions × 3 snapshots → $X \in \mathbb{R}^{n\times 24}$
2. **Spatial modes:** SVD of $X = \Psi\Sigma V^\top$ extracts the dominant flow patterns from the library, retaining r modes (one per sensor) → Ψ_r
3. **Sensor placement:** QR pivoting on $\Psi_r^\top$ selects r locations, encoded in the selection matrix $C \in \mathbb{R}^{r\times n}$, that maximize the information content of $A = C\Psi_r$
4. **Reconstruction:** given station measurements y, solve $\hat{a} = (C\Psi_r)^\dagger y$ (Moore-Penrose pseudo-inverse) and recover the full field as $\hat{x} = \Psi_r\hat{a}$

Spatial constraint

Sensor candidates are restricted to physically accessible terrain by weighting each row of Ψ_r before QR pivoting:

$$\tilde{\Psi}_r = W\Psi_r, \quad \text{with } W = \mathrm{diag}(w_i) \text{ and } w_i \in [0,1]$$

In the binary case ($w_i \in \{0,1\}$), inaccessible locations are excluded entirely. The framework generalizes to soft constraints via a Gaussian decay $w_i = \exp(-d_i^2/\sigma^2)$, where d_i is the distance from the center of the accessible region and σ controls the decay rate.

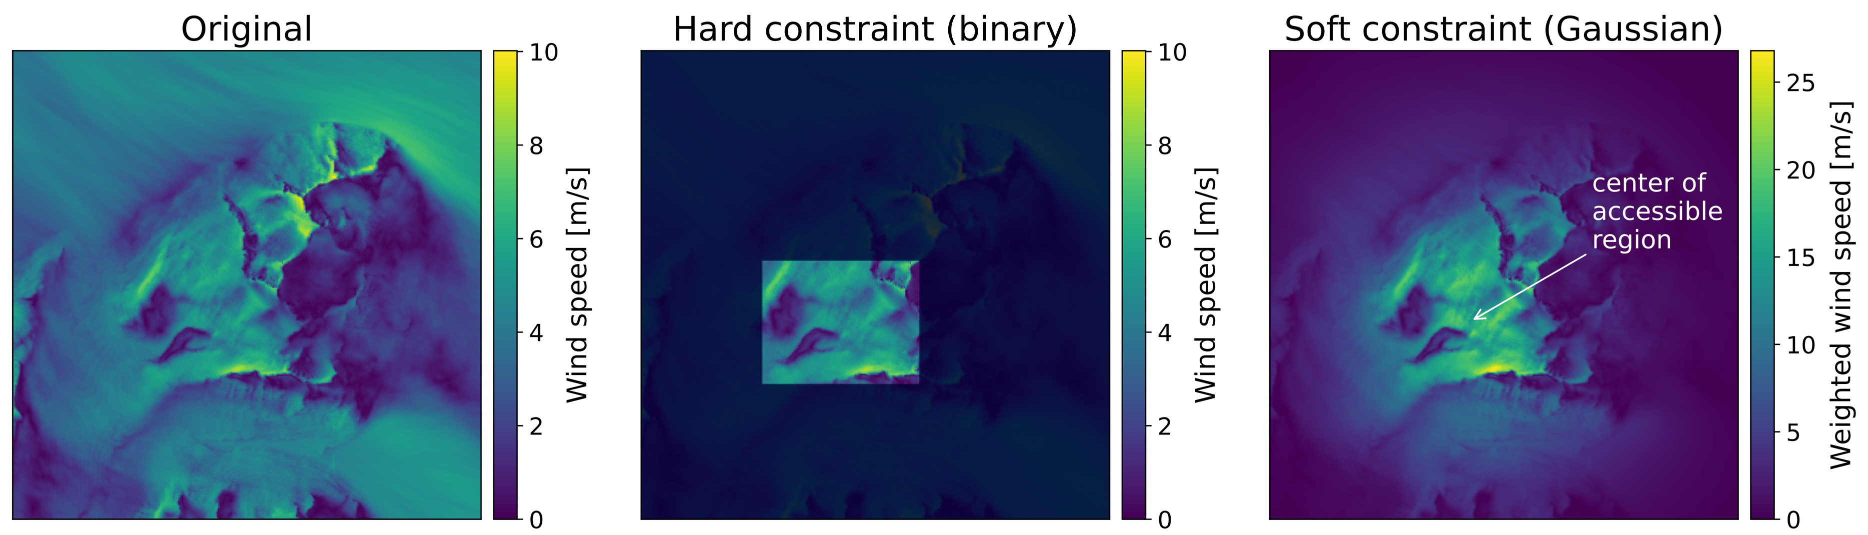

Fixed sensor

One sensor is prescribed *a priori* due to logistical constraints. Its contribution to the flow basis is projected out before QR pivoting, ensuring the remaining $r-1$ sensors cover the orthogonal complement:

$$\tilde{\Psi}_r^\perp = \tilde{\Psi}_r - \mathbf{e}_1(\mathbf{e}_1^\top\tilde{\Psi}_r), \quad \text{with } \mathbf{e}_1 = \frac{\tilde{\psi}_{\text{fixed}}}{\|\tilde{\psi}_{\text{fixed}}\|}$$

where $\tilde{\psi}_{\text{fixed}}$ is the row of $\tilde{\Psi}_r$ corresponding to the fixed sensor location.

Possible extension: weighted directions

We generalize the base framework to weight each column of X by the local wind climatology, forming:

$$X_w = XD, \quad \text{with } D = \mathrm{diag}(w_k) \text{ and } w_k = 1 + \beta(\bar{f}_k - 1)$$

where $\bar{f}_k$ is the normalized frequency of direction k and β is a softening parameter. The SVD of X_w then prioritizes flow patterns that occur most frequently at the site.

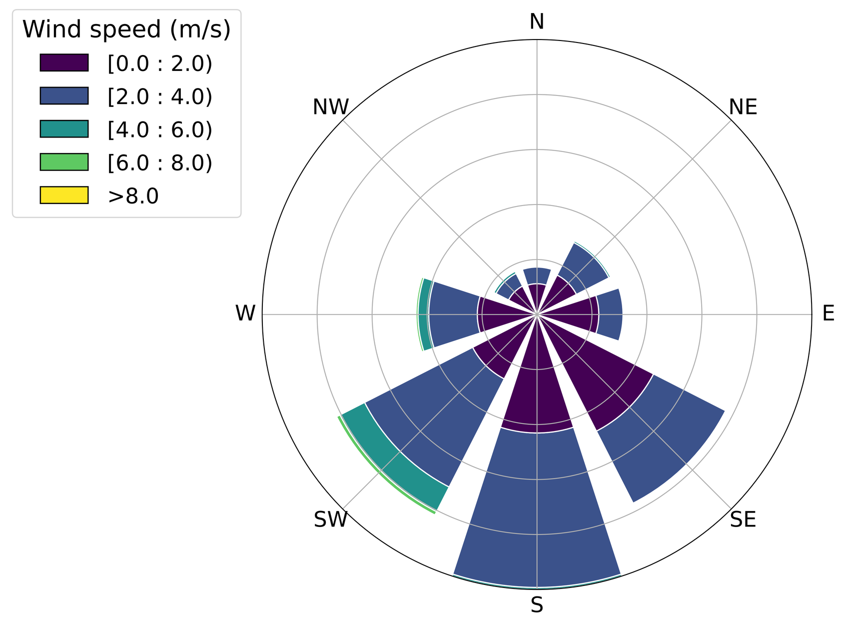

Wind climatology at Ullstinden (ERA5 [3], 5-year average). Southerly and south-westerly winds dominate.

RMSE on held-out snapshots

	Mean	Std
Unweighted	8.26%	0.18%
Weighted	7.93%	0.21%

Independence is preserved

The QR-selected sensor locations are validated against real station measurements using pairwise Normalized Mutual Information (NMI). All values remain below 0.12, confirming that the stations capture largely independent information streams in practice.

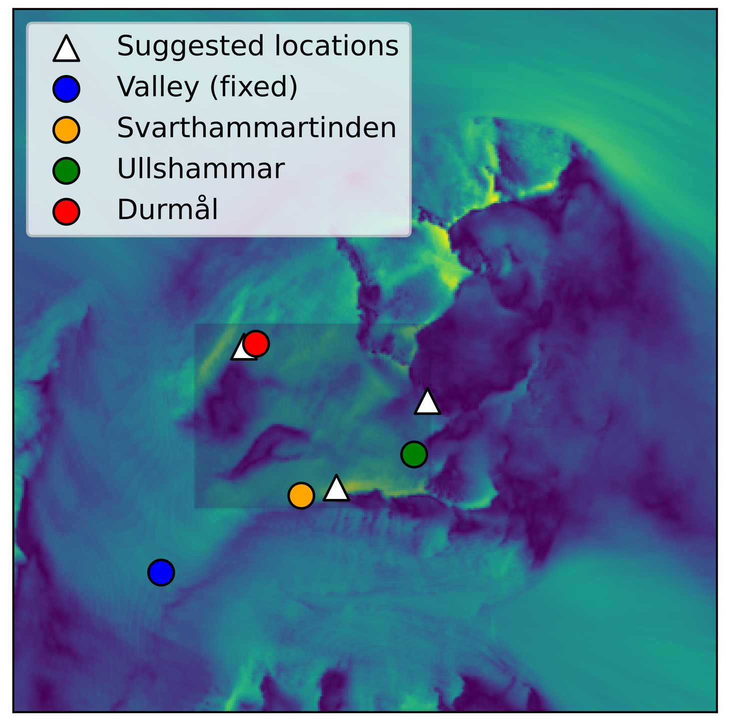

Pairwise NMI

	V	S	U	D
Valley (V)	–	0.097	0.044	0.088
Svarthammartinden (S)	–	–	0.109	0.117
Ullshammar (U)	–	–	–	0.095
Durmål (D)	–	–	–	–

Reconstruction test

Wind components u and v are reconstructed jointly using a stacked basis $[X_u; X_v]$, rather than reconstructing total wind speed directly. This eliminates non-physical negative values and improves accuracy:

$$A = \begin{bmatrix} C\Psi_r^u \\ C\Psi_r^v \end{bmatrix}, \quad \hat{a} = A^\dagger y_{u,v}, \quad |\hat{x}| = \sqrt{(\Psi_r^u\hat{a})^2 + (\Psi_r^v\hat{a})^2}$$

where $y_{u,v} = [u_1,\ldots,u_r,v_1,\ldots,v_r]^\top$ are the observed wind components at the r station locations.

Note: The operator $A = C\Psi_r$ is well-conditioned ($\mathrm{cond}(A) = 3.48$), ensuring that measurement noise is not amplified during the inversion.

Quality gate: Each timestamp is assessed against the library via a normalized vector distance between observed and PALM wind components (u,v) at the station locations, capturing both magnitude and directional mismatch. Timestamps exceeding the 75th percentile threshold are excluded (∼25% of observations).

RMSE on held-out snapshots

	Mean	Std
Total speed	9.38%	0.87%
Joint u/v	8.75%	2.46%

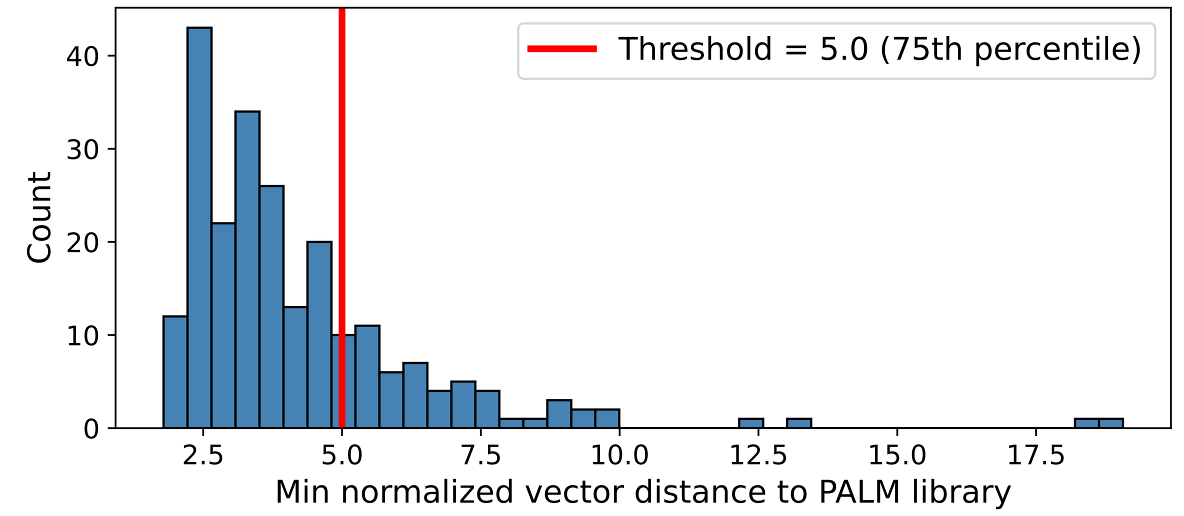

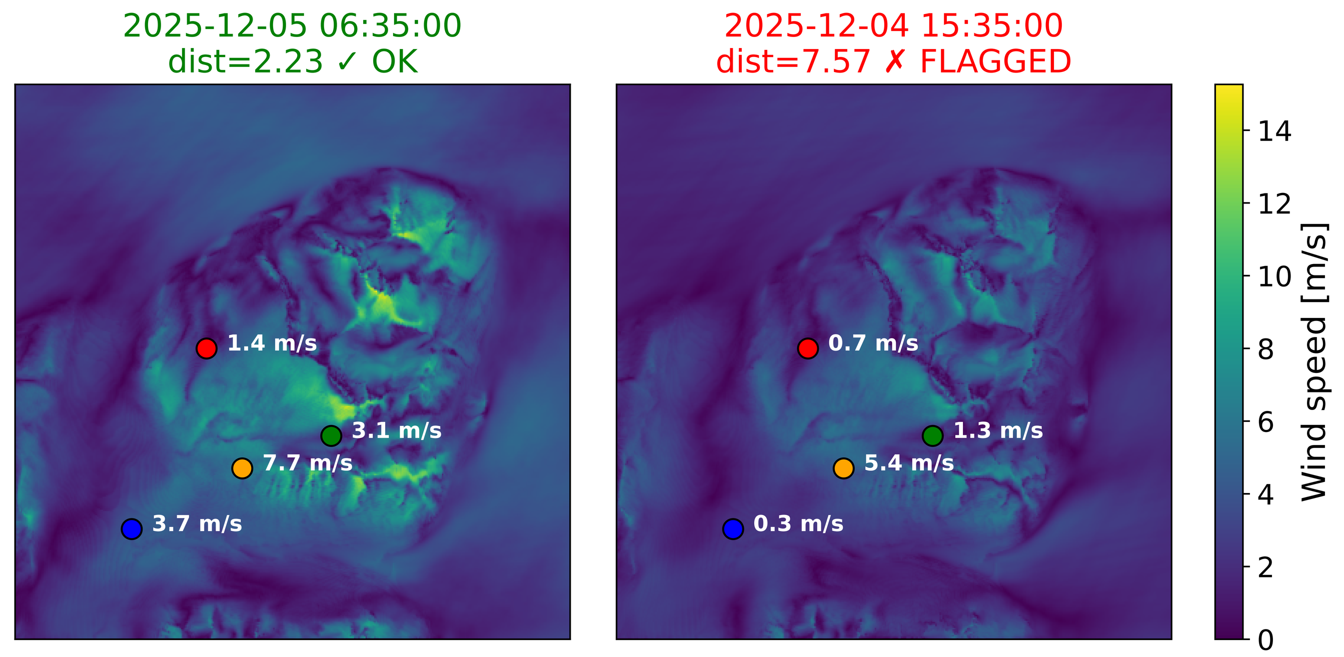

Reconstructed wind speed fields for two representative timestamps. Dots show the four station measurements used as input. Left: timestamp passing the quality gate. Right: flagged timestamp – no PALM snapshot closely matches the observed inter-station amplitude ratio.

Future work

- Expand the PALM library to multiple inflow speeds to improve amplitude coverage and reduce the fraction of flagged timestamps.
- Deploy directional weighting in future sensor campaigns to reduce reconstruction RMSE.
- Extend the framework to other meteorological variables and other complex terrain sites.

References:

anna.poltronieri@uit.no