



Filtering of pulsed lidar data using spatial information and a clustering algorithm

Alcayaga, Leonardo

Published in:
Atmospheric Measurement Techniques

Link to article, DOI:
[10.5194/amt-13-6237-2020](https://doi.org/10.5194/amt-13-6237-2020)

Publication date:
2020

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Alcayaga, L. (2020). Filtering of pulsed lidar data using spatial information and a clustering algorithm. *Atmospheric Measurement Techniques*, 13(11), 6237-6254. <https://doi.org/10.5194/amt-13-6237-2020>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



Filtering of pulsed lidar data using spatial information and a clustering algorithm

Leonardo Alcayaga

DTU Wind Energy, Frederiksborgvej 399, 4000 Roskilde, Denmark

Correspondence: Leonardo Alcayaga (lalc@dtu.dk)

Received: 24 November 2019 – Discussion started: 20 February 2020

Revised: 18 July 2020 – Accepted: 25 August 2020 – Published: 20 November 2020

Abstract. Wind lidars present advantages over meteorological masts, including simultaneous multipoint observations, flexibility in measuring geometry, and reduced installation cost. But wind lidars come with the “cost” of increased complexity in terms of data quality and analysis. Carrier-to-noise ratio (CNR) has been the metric most commonly used to recover reliable observations from lidar measurements but with severely reduced data recovery. In this work we apply a clustering technique to identify unreliable measurements from pulsed lidars scanning a horizontal plane, taking advantage of all data available from the lidars – not only CNR but also line-of-sight wind speed (V_{LOS}), spatial position, and V_{LOS} smoothness. The performance of this data filtering technique is evaluated in terms of data recovery and data quality against both a median-like filter and a pure CNR-threshold filter. The results show that the clustering filter is capable of recovering more reliable data in noisy regions of the scans, increasing the data recovery up to 38 % and reducing by at least two-thirds the acceptance of unreliable measurements relative to the commonly used CNR threshold. Along with this, the need for user intervention in the setup of data filtering is reduced considerably, which is a step towards a more automated and robust filter.

1 Introduction

Long-range scanning wind lidars are useful tools, and their adoption has grown rapidly in recent years in wind energy applications (Vasiljević et al., 2016). Scanning wind lidars can measure time evolution and spatial characteristics of wind fields over large domains at a lower cost of installation than meteorological masts. Nevertheless, atmospheric con-

ditions and instrument noise can have an important impact on the data quality. For long-range scanning lidars this becomes an important issue due to the lack of additional instruments placed over the measurement area that would be useful to compare data quality since noise can contaminate large portions of the scanning domain. The most commonly used criterion to retrieve reliable observations is a threshold on values of the carrier-to-noise ratio, CNR, threshold that will depend on site conditions, experimental setup, and the instrument manufacturer (Gryning et al., 2016; Gryning and Floors, 2019). Even though the CNR threshold retrieves quality observations, its application might result in large numbers of good data rejected in regions far from the instrument, where CNR has decreased rapidly with distance. To cope with this issue Meyer Forsting and Troldborg (2016) and Vasiljević et al. (2017) have proposed filters based on the smoothness and continuity of the wind field. Such filters work by detecting discrete or anomalous steps (above a certain threshold predefined by the user) in line-of-sight wind speed, V_{LOS} , compared to its local (moving) median. Beck and Kühn (2017) first and then Karagali et al. (2018), in an adapted version, follow a different approach (here called a KDE filter, from kernel density estimate) based on the statistical self-similarity of the data, which, in simple terms, means that reliable observations are alike and will be located close together in the observational space. The probability density distribution of observations (estimated via KDE) in a dynamically normalized V_{LOS} – CNR space shows that measurements likely to be valid are located in a high-data-density region. Observations sparsely distributed beyond a boundary defined by a threshold in the acceptance ratio or the ratio between the probability density of any observation and the maximum probability density over the whole set of measure-

ments are finally identified as noise. Both approaches need the definition of one or more thresholds and a window size, either in time for the KDE filter or in space for the wind field smoothness approach. These parameters are dependent on different characteristics of the data, like the lidar scanning pattern for instance.

Both approaches miss important and complementary information, either neglecting the strength of the signal backscattering (quantified by CNR) or the spatial distribution and smoothness of the wind field. Moreover, in both approaches the position of observations is not taken into account, information that can shed light on areas permanently showing anomalous values of V_{LOS} or CNR, like hard targets. Including all these features within the smoothness approach is difficult since CNR is not a smooth field like V_{LOS} . Moreover, considering smoothness and position in the KDE filter results in a computationally costly kernel density estimation if we look for an optimal bandwidth parameter in a higher-dimensional space with a fine resolution of the kernel density estimate.

Data self-similarity – over any scale in the case of fractals or a range of them in real situations (Mandelbrot, 1983) – is closely related to clustering techniques (Backer, 1995), which can classify large data sets with many different features at a relatively low computational cost. The KDE filter approach shares some characteristics with the popular k -means clustering algorithm MacQueen (1967) since they define one (or several for k -means) specific group of data belonging to a unique category (or cluster) whose size and location in the observational space will depend on data density or, more specifically, on a kernel density estimation. The main difference between these two algorithms is the way they treat sparse data points that fall in low-density regions. Unlike the KDE filter, which rejects noise via the acceptance ratio, the k -means clustering algorithm assigns sparse points to the cluster with the nearest center no matter if they are outliers or present unlikely values from a physical point of view.

The Density Based Spatial Clustering for Applications with Noise algorithm, or DBSCAN (Ester et al., 1996; Pedregosa et al., 2011), introduced in Sect. 4.3, presents several advantages over k -means in detecting clusters in a higher dimensional space: it introduces the notion of noise or sparsely distributed observations, it does not need prior knowledge of the number of clusters in the data and it is capable of identifying clusters of arbitrary shape. To the best of our knowledge, this is the first time that this type of clustering algorithm is applied to identify nonreliable observations from pulsed lidars. This approach, which can be understood as a natural extension of the KDE filter, is compared to the smoothness-based filter on two types of data: synthetic wind fields data as a controlled test case and real data.

This paper is organized as follows: Sect. 2 describes the real data used to test the different filtering approaches, and Sect. 3 presents the synthetic data used during a controlled

test as well as the methodology to obtain it. Section 4 then gives a description of the different filters applied in this study to both data sets to continue with the definition of the performance tests in Sect. 5. In Sect. 6 the performance tests are presented along with a discussion on their validity and quality. Section 7 discusses the quality of the methodology behind the tests and the advantages and disadvantages of the proposed approach. Section 8 presents the conclusions of this study.

2 Real data: Østerild balconies experiment

The filtering techniques presented here were tested on lidar measurements made at the Test Centre Østerild located in northern Jutland, Denmark (see Fig. 1). Known as the Østerild balconies experiment (Mann et al., 2017; Karagali et al., 2018; Simon and Vasiljevic, 2018), this measuring campaign aimed to characterize horizontal flow patterns above a flat, heterogeneous forested landscape at two heights relevant for wind energy applications, covering an area of around 50 km² and a wide range of scales in both time and space.

The experiment consist of two measuring phases (see Table 1) with two long-range WindScanners performing plan position indicator (PPI) scanning patterns, aligned in the north–south axis and installed at 50 m a.g.l. during phase 1 and 200 m a.g.l. in phase 2. WindScanners (Vasiljević et al., 2016) consist of two or more spatially separated lidars, which are synchronized to perform coherent scanning patterns, allowing the retrieval of two- or three-dimensional velocity vectors at different points in space. These experiments were conducted between April and August of 2016 (Simon and Vasiljevic, 2018). In each phase, the northern and southern lidars scanned in the west and east direction relative to the corresponding meteorological masts where they were installed. The data used in this study originated from both phases of the experiment, with PPIs pointing to the west. For more details about the experiment, lidars, and terrain characteristics see Karagali et al. (2018), Vasiljević et al. (2016), and Simon and Vasiljevic (2018).

This data set is well suited to test different data filtering techniques. A large measurement area will be affected by local terrain and atmospheric conditions, like clouds or large hard targets. Moreover, at this scale lidars reach their measuring limitations since the backscattering from aerosols decreases rapidly with distance (Cariou, 2015).

3 Synthetic data

Assessing and comparing the performance of filters is challenging with no reference available to verify that rejected or accepted observations are reliable or bad observations. This is especially difficult for long-range scanning lidars since their measurements cover large areas and, due to spatial variability, a valid reference would need several sec-

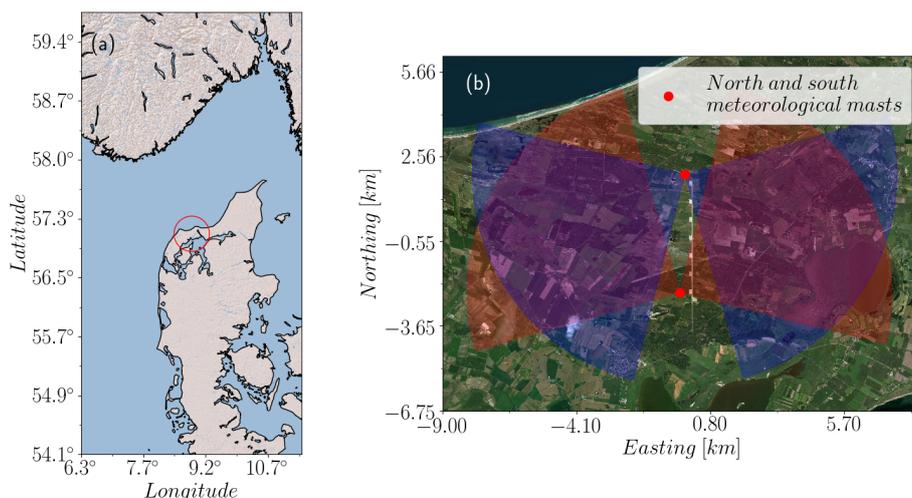


Figure 1. (a) Location of the Østerild turbine test center, place of the balconies experiments, northern Jutland, Denmark (© 2009 Esri). (b) Detail of the test center site, with the location of the meteorological masts where north (blue) and south (red) WindScanners were installed. During the measurement campaign the PPI scans both pointed west in some periods and east in others (© 2017 DigitalGlobe, Inc.).

Table 1. Characteristics of the balconies experiment, from Karagali et al. (2018). The scans are neither instantaneous nor totally synchronous, with a horizontal sweep speed of 2° s^{-1} in the azimuth direction in a range of 90° and a total time of 45 s per scan.

Phase	Measurement start	Measurement end
50 m a.g.l. (1)	12 Apr 2016 12:45:41 UTC	17 Jun 2016 12:48:01 UTC
200 m a.g.l. (2)	29 Jun 2016 13:35:56 UTC	12 Aug 2016 09:09:55 UTC
Scanner	Location coordinates (m)	Scanning pattern, west
Southern lidar	492768.8 (E) 6322832.3 (N)	344–256°, 2° steps
Northern lidar	492768.7 (E) 6327082.4 (N)	196–284°, 2° steps

ondary anemometers scattered over the scanning area. Testing filters on a controlled and synthetic data set contaminated with well-defined noise presents an option to deal with this problem. In this study, the filters presented in Sect. 4 are tested on individual scans sampled from synthetic wind fields generated using the Mann turbulence spectral tensor model (Mann, 1994) and contaminated with procedural noise (Perlin, 2001).

3.1 Synthetic-wind-field generation

Synthetic PPI scans are sampled by a lidar simulator from synthetic wind fields generated via the Mann model (Mann, 1998) in a horizontal, two-dimensional square domain of 2048×2048 grid points, with dimensions of $9200 \text{ m} \times 7000 \text{ m}$. The generated turbulence fields are the result of input parameters of the turbulence spectral tensor model, namely length scale, L ; turbulence energy dissipation, $\alpha \epsilon^{2/3}$; and anisotropy, Γ . The fields generated correspond to wind speed fluctuations, to which the desired mean wind speed is subsequently added. Depending on the initial random seed used, different wind field realizations with the

Table 2. Synthetic wind field characteristics and parameters.

Parameter	Values
L , m	62, 125, 250, 500, 750, 1000
$\alpha \epsilon^{2/3}$, $\text{m}^{4/3} \text{ s}^{-2}$	0.025, 0.05, 0.075
Γ	0, 1, 2, 2.5, 3.5
Number of seeds used	10
Mean wind speed, U (m s^{-1})	15
Mean wind speed direction range (degrees)	90 to 270
Total number of scans generated	4305

exact same turbulence statistics can be generated. For details on wind field generation using the Mann model, refer to Mann (1998). Table 2 shows the range of values used for the generation of two-dimensional wind fields. Large values of $\alpha \epsilon^{2/3}$ or small-scale turbulence, for instance, mean that sudden spatial changes in wind speed are more likely, which increase the false identification of outliers. Mean wind direction, turbulence anisotropy, and length scale will also affect the sampling due to the lidars' measuring characteristics.

3.2 Lidar simulator

Lidar simulators have been presented previously by Stawarski et al. (2013) and Meyer Forsting and Troldborg (2016). They sample V_{LOS} values from wind fields generated via large-eddy simulation (LES), mimicking the operational principle of lidars by proper time and spatial (probe volume) averaging of the background wind field. The lidar simulator presented here follows the same principles, this time sampling from synthetic wind fields generated via the Mann model.

The simulator receives scanning pattern characteristics as input (beam range, range gate step, azimuth angle range, and azimuth angle steps) to generate a primary mesh with the sampling positions on top of background wind field. Following the measuring principle of the lidar, the V_{LOS} observed at each position in this mesh will represent averages of a continuum along each range gate step (due to probe volume averaging) and an average of many azimuth positions within the azimuth step due to the almost continuous sweep of the lidar's beam. The simulator mimics this by generating a secondary, refined mesh with N_r points in each range gate and N_ϕ beams within each azimuth step. The background wind field components, U and V , are then interpolated on this secondary mesh and projected on each refined beam to obtain V_{LOS} using Eq. (1), with θ being the corresponding beam azimuth angle.

$$V_{\text{LOS}} = \cos(\theta)U + \sin(\theta)V \quad (1)$$

The final step is the spatial (probe volume) averaging and the azimuth (sweeping) averaging around each position in the primary mesh. Spatial averaging is done by applying a weighting function to all V_{LOS} values along each refined beam. The weighting function used here is defined in Eq. (2), as in Banakh and Smalikho (1997) and Smalikho and Banakh (2013). This function will assign weights to each point in the refined beam according to its distance to the range gate position in the primary mesh, F , and the instrument probe volume parameters, namely, range gate length, Δp , and full width at half maximum, Δl (cf. Table 3). Here, $\text{Erf}(x)$ is the error function, and r_p is the beam width contribution to the volume averaging.

$$w = \frac{1}{2\Delta p} \left\{ \text{Erf} \left[\frac{(r - F) + \Delta p/2}{r_p} \right] - \text{Erf} \left[\frac{(r - F) - \Delta p/2}{r_p} \right] \right\}; \quad r_p = \frac{\Delta l}{2\sqrt{\ln(2)}} \quad (2)$$

The azimuth averaging is the arithmetic mean of the N_ϕ values of V_{LOS} at each range gate after spatial averaging. It represents the accumulation information of the backscattered signal spectra as they sweep an azimuth sector before estimation of the spectral peak and V_{LOS} .

3.3 Synthetic noise generation

The most simple noise that can be used to contaminate synthetic scans are sparse, uniformly distributed outliers. This noise, also known as salt-and-pepper noise, is easily detected and eliminated by median-like filters when extreme; discrete steps affect the smoothness of an image (Huang et al., 1979; Burger and Burge, 2008). Nevertheless, noise in real scans comes as regions of anomalously high and/or low V_{LOS} , and they can pass through the filter undetected. Procedural noise, introduced by Perlin (2001) to recreate synthetic textures on surfaces for computer graphics applications, creates regions of coherent noise that better resembles the spatial distribution of scanning lidars' measurements. For the two-dimensional case, the procedural noise function $N(x, y)$ maps two-dimensional coordinates, (x, y) , onto the range $[-1, 1]$ as follows,

- A two-dimensional grid of m by n elements is generated, and a pseudorandom, two-dimensional unit gradient, $\mathbf{g}_{ij} = (g_x, g_y)$, is assigned to each grid point (x_i, y_j) . The pseudorandomness rises from the fact that \mathbf{g}_{ij} are picked from a precomputed list of gradients with length $l \ll m \times n$. We select values from this list using the index permutation grid $p_{ij} \in \{0, \dots, l\}$, also with $m \times n$ elements. Then, \mathbf{g}_{ij} will correspond to the gradient in the position p_{ij} of the precomputed list. Elements p_{ij} are shuffled for each realization.
- For each grid point (x_i, y_j) enclosing (x, y) , a distance vector $\mathbf{d}_{i,j} = (x - x_i, y - y_j)$ is generated.
- Finally, the noise function is the sum of dot products, $N(x, y) = \sum_q w_q (\mathbf{g}_{ij}^q \cdot \mathbf{d}_{i,j}^q)$, for q grid points surrounding (x, y) . Weights w_q correspond to $w_q = C \frac{1}{\|\mathbf{d}_{i,j}^q\|}$, and C is a normalization constant to ensure that $N(x, y) \in [-1, 1]$.

The function $N(x, y)$ allows the generation of noisy regions, than can be distributed according to backscatter decay with distance. Three bands centered at 50 %, 70 %, and 90 % of the total beam length (and spanning the entire azimuth range) have an increasing fraction of noise, contaminating 30 %, 60 %, and 90 % of the observation, respectively. The noise amplitude is 35 (m s^{-1}), the limit of the observable range for the instruments described in Sect. 2. Figure 2c shows one contaminated scan and its increasing contaminated area as we move along the beams. The same figure shows the distribution of the noise generated by the algorithm after scaling and the probability distribution of contaminated synthetic V_{LOS} compared to real data with low values of CNR. The distribution of real data presents heavier tails than the ones generated, with higher probability of observing extreme values of V_{LOS} . Modeling real noise is difficult since the process that generates it depends on the measuring principle of the lidar and atmospheric conditions. The synthetic noise used

Table 3. The characteristics of the lidar simulator and real long-range lidar (Karagali et al., 2018; Vasiljević et al., 2016) used for the controlled test of the filters.

	Simulator	Real
Azimuth range	256–344°	256–344°
Azimuth step	2°	2°
Beam length	7000 (m)	7000 (m)
Range gate length, $\Delta\rho$	35 (m)	35 (m)
Full width at half maximum, Δl	75 (m)	75 (m)
Sweeping time per scan	Instantaneous	45 (s)
Primary mesh size (radial \times azimuth)	45 \times 198	–
Secondary mesh size at each range gate ($N_r \times N_\phi$)	21 \times 51	–
Total secondary mesh size ($N_r \times N_\phi$)	21 \times 51	–

here does not intend to be totally realistic but more subtle and smoother than the one observed in real measurements, making the identification of contaminated points more difficult.

4 Filtering techniques applied on real and synthetic data

4.1 CNR threshold

CNR thresholds are well known, and lidar manufacturers usually recommend values for rejection of signals with poor backscattering or hitting hard targets (Cariou, 2015). However, the selection of an appropriate threshold for CNR that assures data quality and good data recovery is not easy. Figures 3 and 4 show data from a scan with noisy observations from CNR values below -27 dB. Both extreme and limited values of V_{LOS} show low CNR values in the distant region of the scan and data loss results after the application of the CNR threshold (Fig. 4b). When a limit to V_{LOS} is applied instead, Fig. 4c shows that the smoothness in V_{LOS} is lost in the lower part of the scan. A conservative threshold of -24 dB is used here since the resulting V_{LOS} probability distribution shows very few outliers, and it can be used as a reference when the performance of the filters proposed is compared.

4.2 Median-like filter

The median filter arises as a viable option for detecting erroneous measurements since it is well known that this type of nonlinear filter is suited to detect and filter noise that presents distributions with large tails. Here we use an adaptation of the traditional median filter used in the image-processing community, closely related to the three-stage filtering technique described in Menke et al. (2019): observations are not replaced by the local moving median but excluded if the absolute difference between their value and the local moving median is above a certain threshold, $\Delta V_{LOS, threshold}$. Unlike Huang et al. (1979), the two-dimensional moving window is replaced by two one-dimensional window instances,

the first in the line of sight or the radial direction, r , and finally in the azimuth direction, θ , considering the polar coordinates of the scan. This simplification reduces the computation time significantly.

The input parameters of this filter will be the size (or number of elements) of the moving windows in the radial and azimuth directions, n_r and n_ϕ , respectively, and $\Delta V_{LOS, threshold}$. For fixed values of $\Delta V_{LOS, threshold}$, n_r , and n_ϕ , the spatial structure of wind speed fluctuations will have an important effect on the recovery rate and noise detection of this filter. A sensitivity analysis carried out using the metrics presented in Sect. 5.1 on the synthetic data set shows that the optimal set of parameters is $n_r = 5$, $n_\phi = 3$, and $\Delta V_{LOS, threshold} = 2.33 \text{ m s}^{-1}$ (See Appendix A). This set is used for both artificial and real data. The filter does not include a time window, and it is applied to individual scans.

4.3 Filtering using a clustering algorithm

If we represent lidar observations as m -dimensional vectors, with m the number of features or parameters of the data, measurements not affected by poor backscattering or noise will cluster together in regions of high data density, as shown in Figs. 3 and 4. The approach presented here identifies such clusters applying DBSCAN to data described by CNR, V_{LOS} , and, additionally, spatial location and smoothness features, which help to make clusters more distinguishable.

DBSCAN identifies clusters and noise based on two parameters: a neighborhood size, ε , and a minimum number of nearest neighbors, NN. The parameter ε is the Euclidean distance from one observation to the limits of a neighborhood that might contain NN (or more) nearest neighbors. Intuitively, these parameters will define the minimum density that a group of data points needs to have to be identified as a cluster. Observations within a cluster fall into the following categories.

- Core point: points q whose ε neighborhood contains NN or more points.

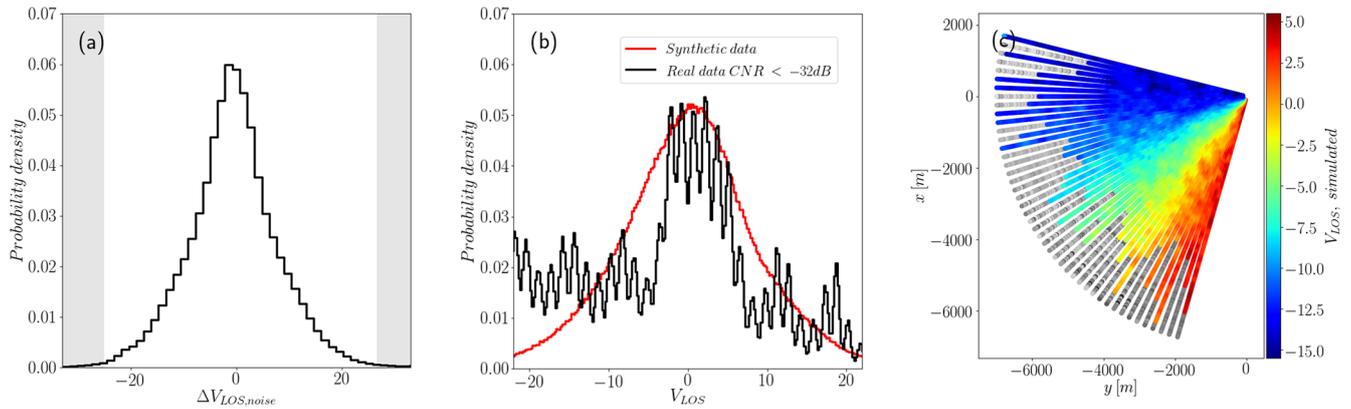


Figure 2. Procedural noise on synthetic scans. **(a)** Distribution of $\Delta V_{LOS,noise}$, the noise added. Maximum values are within the observable range $[-35, 35]$ ($m s^{-1}$). **(b)** Distribution of real V_{LOS} with low CNR values (black) and contaminated, synthetic $V_{LOS} + \Delta V_{LOS,noise}$ (red) for a mean wind direction facing the scan. **(c)** Individual scan showing the increasing fraction of added noise (gray) with distance.

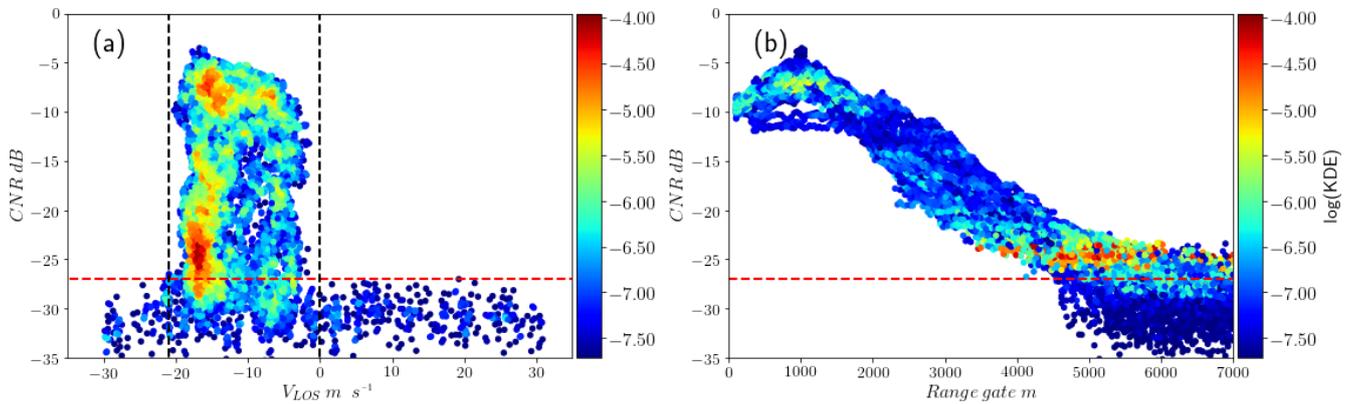


Figure 3. **(a)** CNR and V_{LOS} for one scan from the balconies experiment, including the probability density (KDE). Observations with $CNR > -27$ dB (dashed red line) show a limited range of V_{LOS} (dashed black line). A portion of observations with high probability density remains in the rejection area. **(b)** CNR v/s distance for the same data. Observations with low CNR values and high probability density can be found in the distant region of the scan.

- Direct density-reachable point: points p which are reachable by q by lying within its ϵ neighborhood.
- Density-reachable point: points p reachable by a point r through one or a set of directly connected core points q .

DBSCAN travels across data points identifying core points, border points (density-reachable points with at least one core point within the ϵ neighborhood) and noise, or points that do not belong to any of the categories described above. Finally, the algorithm separates clusters as individual groups of density-connected points. Figure 5 schematically shows these definitions and how the algorithm works.

The input parameters ϵ and NN have a significant influence on the number and characteristics of the clusters detected. For example, large ϵ together with a small NN will end up with sparse clusters that might include noise. In order to find the parameters separating the least dense cluster

from noise, we fix NN to a certain value k and determine ϵ from the data density distribution. The latter is well described by the k -distance function, $d_k(n)$, which represents the distances from all data points n to their respective k th nearest neighbor, sorted in ascending order. When k is 5, for instance, $d_5(n)$ in Fig. 6 shows sudden changes (or knees) that give some indications about the data density distribution. The knee highlighted represents a limit between a group of reliable observations and the one growing fast towards noisy data. The positions of these knees in the graph correspond to the peaks in the curvature of $d_k(n)$, $\kappa(n)$ in expression Eq. (3). In this expression primes correspond to the peaks in the curvature of $d_k(n)$ with respect to n . The continuous version of $d_k(n)$ is made by spline-fitting on a reduced set of uniformly distributed points over the original data set.

$$\kappa = (d_k(n))'' / (1 + (d_k(n))')^2)^{3/2} \tag{3}$$

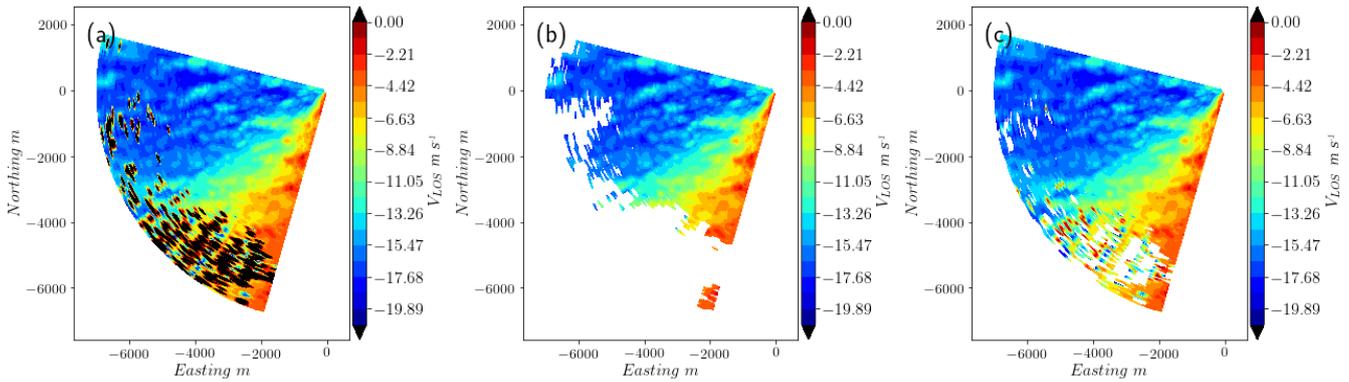


Figure 4. (a) Unfiltered scan with V_{LOS} values outside the range $[-21, 0]$ in Fig. 3a in black. (b) Filtered scan with $CNR > -27$ and the resulting data loss in the upper part of the scan. (c) V_{LOS} within the $[-21, 0]$ range, showing anomalous values in the lower part of the scan.

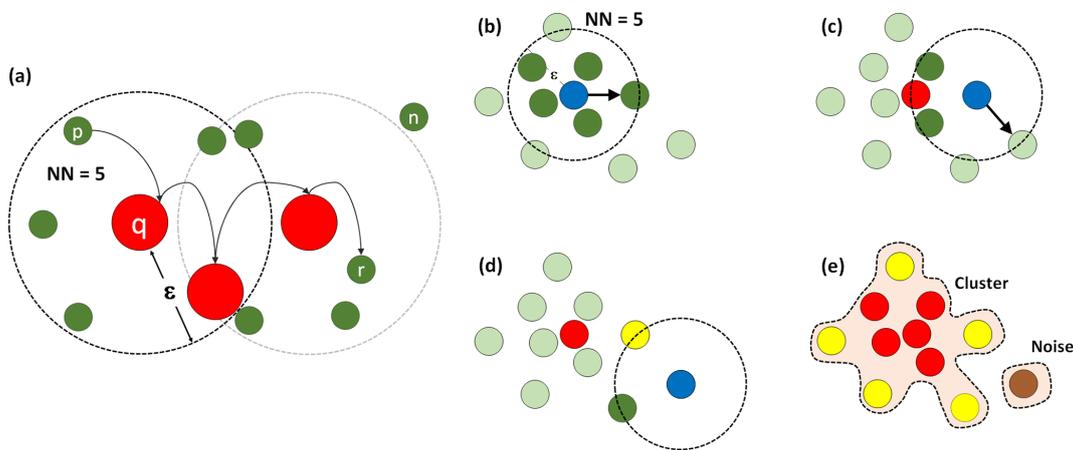


Figure 5. (a) DBSCAN algorithm definitions: direct density-reachable point p (reachable by the core point q) and density-reachable and density-connected points p and r . Here point n does not belong to any of these categories but noise. The DBSCAN algorithm working: (b) the current point being evaluated has the minimum number of nearest neighbors required, NN , within a neighborhood of size ϵ , classified as a *core point* (red). (c) The next point has fewer than NN neighbors, but one of them is a core point and becomes a *border point* (yellow). (d) A point with neither NN neighbors nor core points within ϵ , classified as *noise* (brown). (e) The final cluster and noise. The former is a collection of density-connected points.

When scans are very noisy, the selection of a proper value of ϵ is difficult since knees are located closer together, and a larger fraction of observations show a fast-growing $d_k(n)$, as expected. In this case, the fraction of points showing reliable CNR values is taken into account, and ϵ is estimated by expression Eq. (4). Here f_{CNR} corresponds to the fraction of observed CNR values within the range $[-24, -8]$, and the constants c_1 and c_2 are obtained from the upper and lower bounds of ϵ in the data, respectively.

$$\epsilon_{CNR} = c_1 f_{CNR} + c_2 \tag{4}$$

The set of features considered when filtering synthetic data does not include CNR because it is not available from the lidar simulator described in Sect. 3. For synthetic and real data sets we consider spatial location (azimuth and radial positions) and smoothness as additional features. The latter, ΔV_{LOS} , corresponds to the median difference in V_{LOS} be-

tween a specific position and its direct neighbors in one individual scan.

Since we consider features that vary significantly in magnitude (CNR and range gate distance for instance), we normalize the data before the application of DBSCAN. This step is necessary for the estimation of meaningful distances between observations, which is the basis of this approach. There are several ways to do this. Here, the data in each feature are centered by subtracting their median and scaled according to their interquartile range. This aims to minimize the influence from outliers in the normalization.

The clustering filter is implemented to be a non-supervised classifier and does not need more input parameters than the different features and the number of scans put together as a batch before filtering. The latter is set to three in this case to speed up calculations and avoid creating clusters from noisy regions. From this point of view, this filter is also as dy-

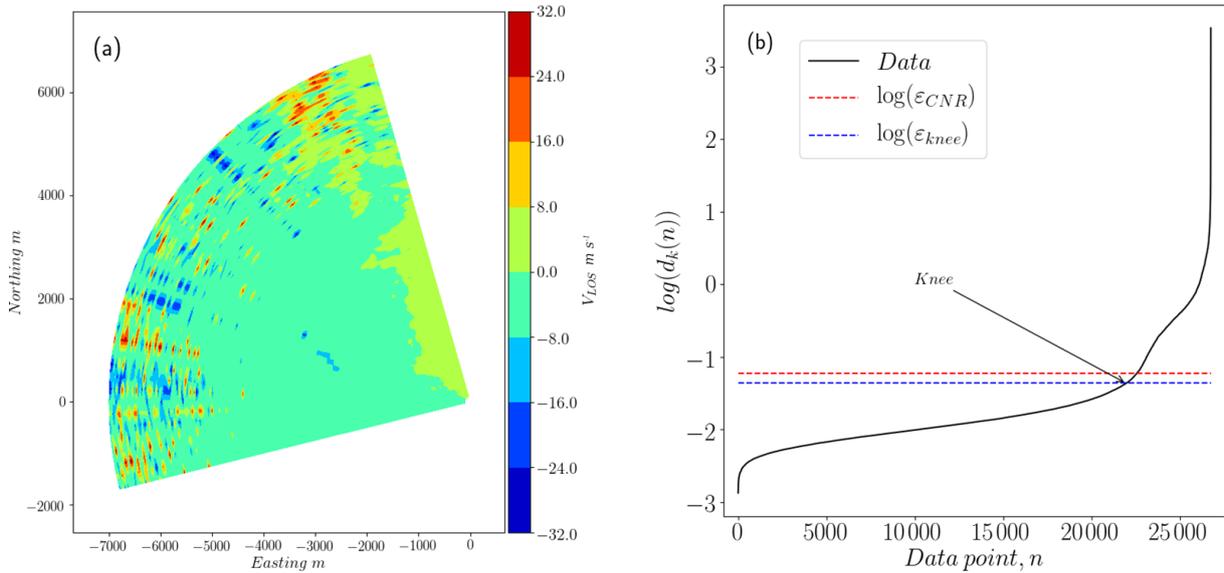


Figure 6. (a) Scan from the balconies experiment (phase 1) with 48 % of data points in the range of reliable observations with $CNR \in [-24, -8]$ dB. (b) Logarithm of sorted distances to the fifth-nearest neighbor for each point in a data set. The total number of observations corresponds to three consecutive scans, or 26 730 points. The sorted fifth distances show three knees separating three types of structures: reliable observations with distances below ϵ_{knee} , an overlapping region where the distance between points grows faster, and pure noise or nonstructured data.

dynamic – similar to Beck and Kühn (2017) – when applied to a real data set since it will consider the data structure within a period limited to 135 s (three scans of 45 s in our case), and characteristics of temporal evolution of the data are indirectly taken into account. For the synthetic data used in this test, more than one scan filtered per iteration gives enough data density in noisy and reliable areas of the observational space. We speculate that scans that are correlated in time will enhance the self-similarity of the data, thus improving the performance of the filter. Turbulence structures with length scales in a range between the range gate size and the scanning area size will evolve at a slower rate than the time elapsed between consecutive scans.

5 Performance metrics

5.1 Synthetic data

Equations (5) to (7) define three metrics to assess the performance of the filters given prior knowledge of the position and magnitude of noise in a controlled case with N observations. The fraction of noise detected, η_{noise} , quantifies the relative importance of true positives or the difference between observations identified as noise, N_{noise} , and false positives, N_{pos} , over the total number of contaminated observations. The fraction of good observations recovered, η_{recov} , gives an idea of the true negatives over the total number of non-contaminated observations, $N_{non-cont}$. True negatives are not equal to $N - N_{noise}$ since the latter might include false neg-

atives, N_{neg} . The relative importance of these two metrics for a given fraction of noise in a contaminated scan, f_{noise} , is quantified by η_{tot} , which takes into account cases with a large fraction of noise detected and low recovery rate and vice versa.

$$\eta_{noise} = \frac{N_{noise} - N_{pos}}{N_{cont}} \tag{5}$$

$$\eta_{recov} = \frac{N - (N_{noise} + N_{neg})}{N_{non-cont}} \tag{6}$$

$$\eta_{tot} = f_{noise}\eta_{noise} + (1 - f_{noise})\eta_{recov} \tag{7}$$

5.2 Real data

In the absence of reference measurements, the quality of the data retrieved after filtering is assessed by comparing the distribution of radial wind speeds for very reliable observations (with CNR values within the range of -24 to -8 dB) with the distribution of filtered observations that fall out of this range. Observations out of the reliable range population usually show a probability density function (or PDF) with heavier tails, like the PDFs in Fig. 7. Here we understand a heavy-tailed PDF as a distribution that slowly goes to 0 and shows higher probability density for values beyond the 3σ limit (or 3 SD limit) when compared to the normal distribution, evidence of a higher probability of occurrence of outliers or extreme values. The recovering rate of observations beyond the $[0.003, 0.997]$ quantile range of the reliable V_{LOS} (shaded area in Fig. 7) could shed light on the quality of the data retrieved by the filter.

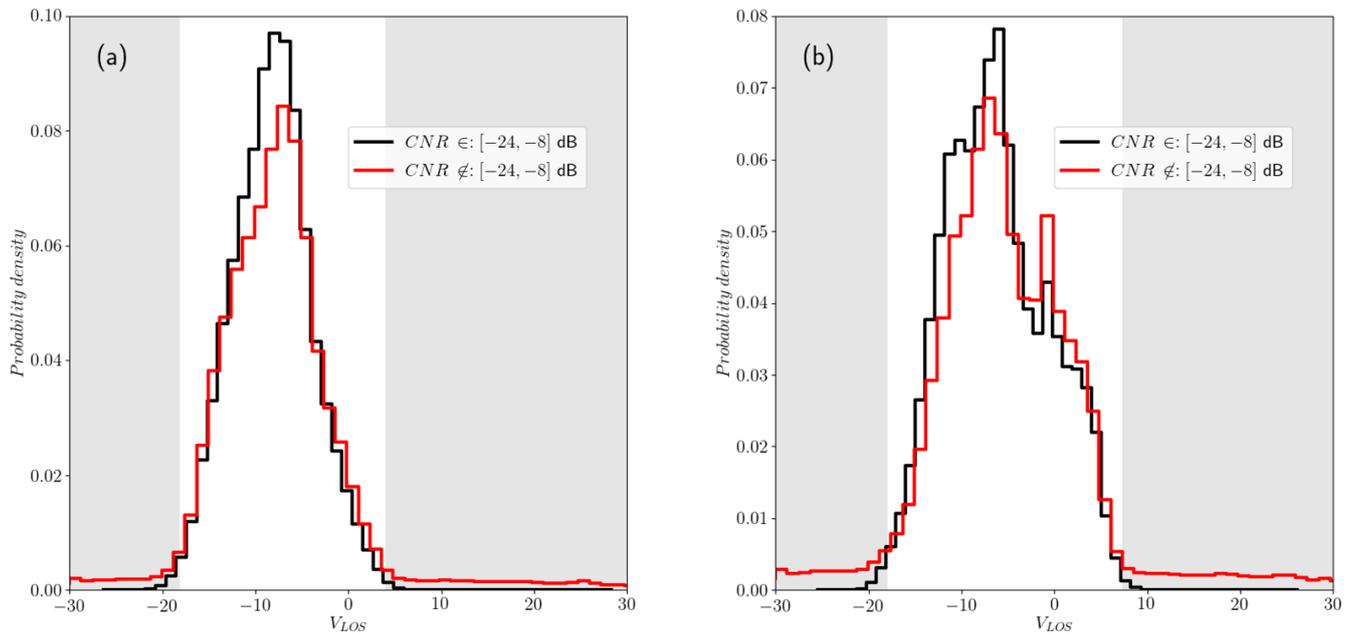


Figure 7. Probability density function of reliable observations of V_{LOS} (solid black line) and nonreliable observations (solid red line) for (a) phase 1 of the balconies experiment, with scans performed at 50 m a.g.l., and (b) phase 2 of the same campaign, with scans performed at 200 m a.g.l.

Another metric is the similarity between the PDF of reliable and nonreliable data after filtering. The distance between both probability density functions can be compared with similarity metrics like the Kolmogorov–Smirnov test (Kolmogorov, 1933) or Kullback–Leibler (KL) divergence (Kullback and Leibler, 1951). The former test measures the statistical similarity between two random variables, X_1 and X_2 , by estimating the statistical distance, D (or K–S statistic), between their cumulative distribution functions, $F_1(x)$ and $F_2(x)$, as the supreme of their difference.

$$D_K = \sup_x \|F_1(x) - F_2(x)\| \quad (8)$$

The null hypothesis here is that two realizations are from the same distribution if the K–S statistic is such that its two-tailed p -value is above a certain level α . Because the number of data analyzed here is large – we analyzed over 20 000 scans for the two phases of the Østerild campaign, each with 8910 data points, over almost 10 d – this similarity test is very precise but also very strict, rejecting the null hypothesis for small deviations between $F_1(x)$ and $F_2(x)$. Nevertheless, the K–S statistics can be used to compare which probability distribution after filtering is closer to the one representing the reliable observations.

The KL divergence is a measure of similarity or overlapping of two distributions, P_1 and P_2 , with realizations X_1 and X_2 , respectively. It is used in different applications to shed light on the loss of information when X_1 is represented by P_2 or vice versa and is defined by the expression Eq. (9).

$$D_{KL} = \sum_x P_1(x) \log \left(\frac{P_2(x)}{P_1(x)} \right) \quad (9)$$

Both metrics will be used to estimate how the distribution of nonreliable observations of V_{LOS} is modified after filtering and if the new distribution is similar (or close in a statistical distance sense) to the probability density of reliable observations of the radial wind speed, shown in Fig. 7 for phases 1 and 2 of the measurement campaign, respectively.

Both performance metrics, the recovery rate of abnormal measurements in the tails of the PDF of reliable observations and its statistical distance to the PDF of filtered nonreliable observations, will be assessed for the median-like filter, the clustering filter, and also for data filtered with a CNR threshold of -29 dB following Gryning and Floors (2019).

6 Results

6.1 Synthetic data

In Fig. 8 we can see the result of the two filters applied to one synthetic scan contaminated with procedural noise. The contaminated observations are indicated by the gray area in this scan. Extreme values contaminating V_{LOS} are identified by both filters without problems, but subtle alterations of the original values of the scan are hard to detect for the median-like filter. The clustering filter performs very efficiently when detecting this type of contaminated observation and filters al-

most all the noise. Both filters repeat this behavior in all the synthetic scans used for this controlled test, as can be seen in Fig. 9, which shows the resulting metrics of the two filters applied to the whole synthetic data set. Looking at η_{tot} , both filters show similar mean values and spread, with the clustering filter performing slightly better. The difference becomes noticeable when we see η_{noise} , which for the clustering filter shows a mean value of 0.95, far larger than the 0.67 of the median-like filter. The latter result could be problematic if the median-like filter is used since noise contaminating the filtered scan will result in nonrealistic wind fields after reconstruction.

Both filters perform well when evaluated in terms of η_{rec} , with the median-like filter showing a higher mean fraction of good observations retrieved, 0.96, compared with the 0.89 of the clustering filter. This result is expected since the median-like filter is more permissive regarding fluctuations that can seem locally anomalous for the clustering filter.

6.2 Real data

The data set from the balconies experiment presents advantages for the clustering filter since the CNR value can be included as a feature in describing the data. Nevertheless, as mentioned already in Sect. 2, we do not count on any reference to assess the performance of the filter apart from the radial wind speed distribution of very reliable observations with CNR values within the range of -24 to -8 dB. As mentioned earlier, valid observations in this range might present a similar distribution. Figure 7 shows this distribution before filtering, shadowing the area of values of V_{LOS} that fall in the region beyond 99.7% of the total probability or 3σ limit, usually classified as outliers. Figures 10 and 11 show the recovery fraction for CNR and median-like and clustering filters when applied to data in the reliable and nonreliable CNR ranges for phases 1 and 2 of the Østerild experiment. Unlike the clustering filter, the CNR threshold and median-like filters show nonnegligible recovery rates beyond the 3σ limit, particularly significant in the former. This result is very much in line with the η_{noise} metric from the synthetic data. Within the 3σ range, the CNR and median-like filters perform slightly better than the clustering filter in terms of recovery fraction, in agreement with the results of η_{rec} in Sect. 6.1. Even though this might compensate the fact that CNR threshold and median-like filters fail to filter out the major part of outliers, increasing the availability of measurements, this difference does not make the PDF of the filtered data more similar to the PDF of reliable data, as Table 4 shows via the metrics D_K and D_{KL} . According to this metric, the PDF of the data after the application of the clustering approach looks more statistically similar to reliable observations. This table also shows D_K and D_{KL} of the nonreliable data before filtering, which in all cases is improved, except for D_K for median and CNR threshold filters during phase 2.

Table 4. Results of PDF similarity test of reliable and nonreliable data after filtering. The CNR = -29 dB threshold is also included (Gryning and Floors, 2019).

Phase 1	D_K	D_{KL}
Nonreliable data before filtering	0.097	0.134
CNR threshold > -29 dB	0.045	0.109
Median filter	0.047	0.126
Clustering filter	0.037	0.105
Phase 2		
Nonreliable data before filtering	0.110	0.126
CNR threshold > -29 dB	0.114	0.052
Median filter	0.117	0.057
Clustering filter	0.103	0.045

Figures 12–14 show the performance of the three different filters in different regions of the scan from, respectively, phase 1 and 2 of the experiment. When the spatial distribution of the recovery fraction is analyzed, we can see that the lowest values shown by the clustering filter are mostly located in the far region of the scan, which in general presents low CNR values. The spatial recovery rate during phase 1 also shows that the median-like and clustering filters are able to identify hard targets, which are also a source of bad observations. For scans recorded at 50 m a.g.l. in phase 1, backscatter is affected by a group of seven turbines located approximately in the middle of the scanning area, with one turbine touching the end of the southern beams of the scan and a meteorological mast located very close to the lidar. Figure 13 shows a detail of the recovery rate associated with the flow in the vicinity of the turbine group in which we can see that the clustering filter is able to better identify the turbine locations, recovering more data in the surroundings when compared to the median-like filter. The PDF of V_{LOS} in this area also shows more similarities between the data filtered with the clustering algorithm and observations with CNR values in the range $[-24, 8]$.

Table 5 shows a summary of the additional data available when the CNR = -29 dB threshold and the median-like and clustering filters are applied instead of the more conservative and restrictive CNR = -24 dB threshold filter. Additionally, this table shows the fraction of observations exceeding the 3σ limit that are recovered by the three filters. Even though the clustering filter shows a slightly lower fraction of additional data available when compared to the other filters, most of it comes from values within the 3σ region. Moreover the quality of the data recovered by the clustering approach seems to be higher when all these results are tested with the performance metrics defined in Sect. 5.

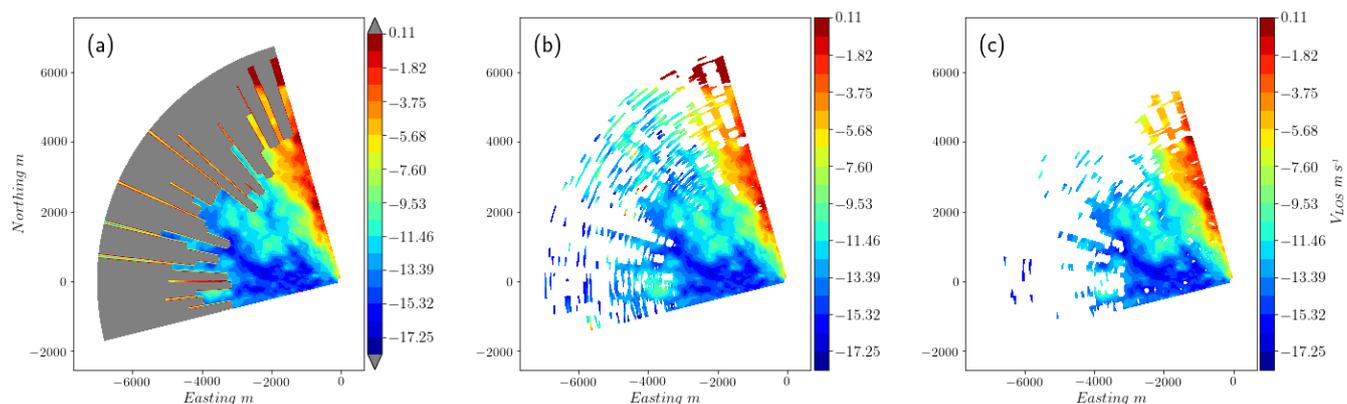


Figure 8. (a) Contaminated synthetic scan with noise indicated by gray area. (b) Scan filtered using the median-like approach. (c) Clustering filter.

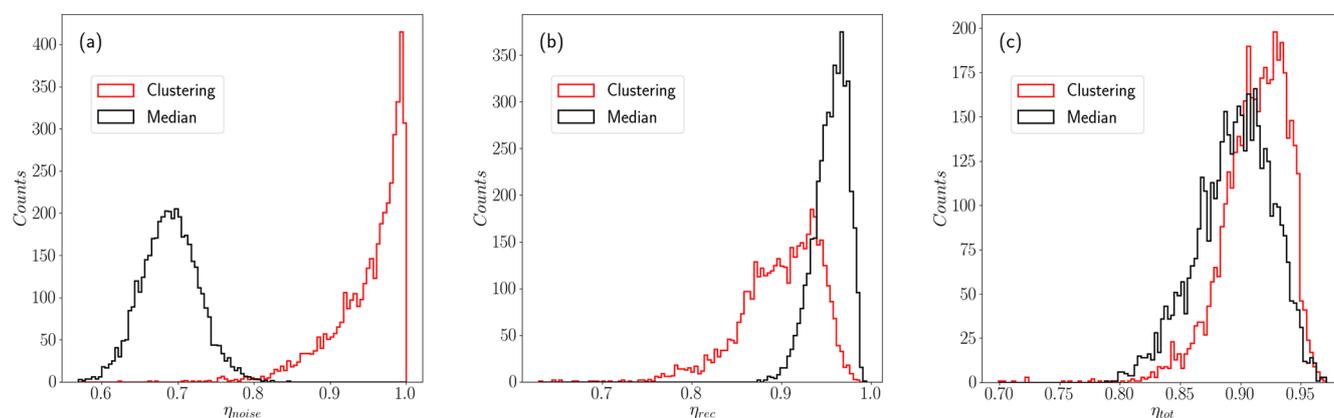


Figure 9. Histograms of the three performance indexes for the total number of synthetic scans. (a) Both filters show similar spread, but the clustering filter rejects a rather higher fraction of noise. (b) The higher recovery rate of the median-like filter and its narrower distribution are superior to the clustering algorithm; the cost is acceptance of more contaminated observations. (c) Both filters have similar mean values for η_{tot} around 0.9.

7 Discussion

The metrics introduced in Sect. 5.1 attempt to evaluate two different capabilities of the filters: the quality and number of data recovered. In general these two metrics are in conflict every time a high rate of noise detection will decrease the data recovery. The metric η_{tot} attempts to quantify their relative importance regarding the noise fraction, which in this study is distributed in a relatively wide range but on average represents 20% of the total number of measurements per scan. The impact of the noise fraction distribution on the performance of the filters was not explored, and variations in its dispersion and mean value might be necessary. Regarding the synthetic scans, they do not allow the identification of outliers in the time domain because they are time-independent. Time-evolving synthetic turbulence fields would be necessary to generate scans correlated in time and enhance the self-similarity of the data. This might improve the performance of the clustering approach and allow the ad-

dition of a time dependence in the median-like filter, used already in Meyer Forsting and Troldborg (2016).

The synthetic wind fields used here do not consider the presence of hard targets. These anomalies in the wind field are observed by lidars as points with high CNR values and abnormal V_{LOS} . Assessing the performance of the filters in detecting such anomalies needs a more realistic model of the pulsed lidar. This lidar simulator would allow the generation of information normally available in real lidar measurements, like CNR, and the spread in the power spectra of the heterodyne signal, S_b . This additional information will benefit the performance assessment of the clustering filter and the simulation of hard targets. A more realistic lidar model was already implemented by Brousmiche et al. (2007), which can be used to further explore these aspects of the filtering process.

The data set analyzed from the balconies experiment corresponds to horizontal scans at 50 and 200 m a.g.l., limiting the analysis to one scanning pattern. Different scanning

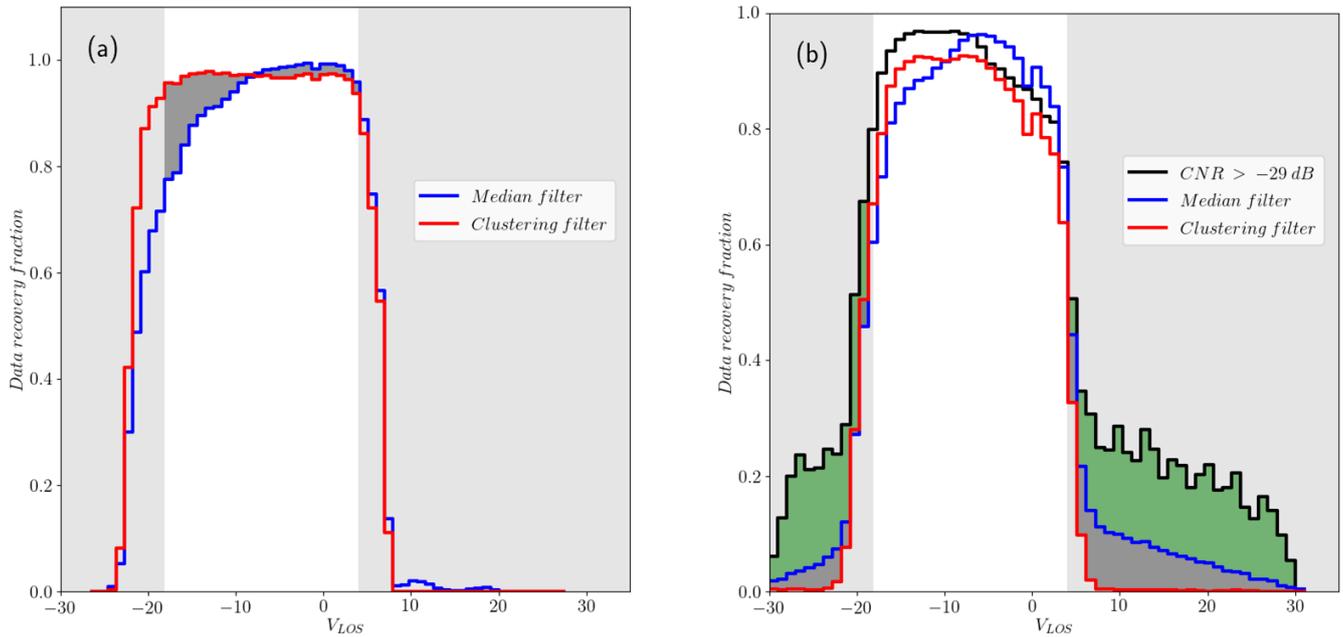


Figure 10. Distribution of recovery fraction per wind speed bin for phase 1 of the experiment of (a) reliable observations ($-24 < \text{CNR} < -8$) and (b) nonreliable data ($\text{CNR} < -24$ or $\text{CNR} > -8$) for the three types of filters. The shadowed area in both graphs corresponds to the region where observations exceed the 99.7% probability (or 3σ limit) in the PDF of reliable observations. The darker shadowed area highlights the additional fraction of extreme values nonfiltered by the median-like and CNR filters when the former uses the optimal input set $n_r = 5$, $n_\phi = 3$, and $\Delta V_{\text{LOS, threshold}} = 2.33 \text{ m s}^{-1}$.

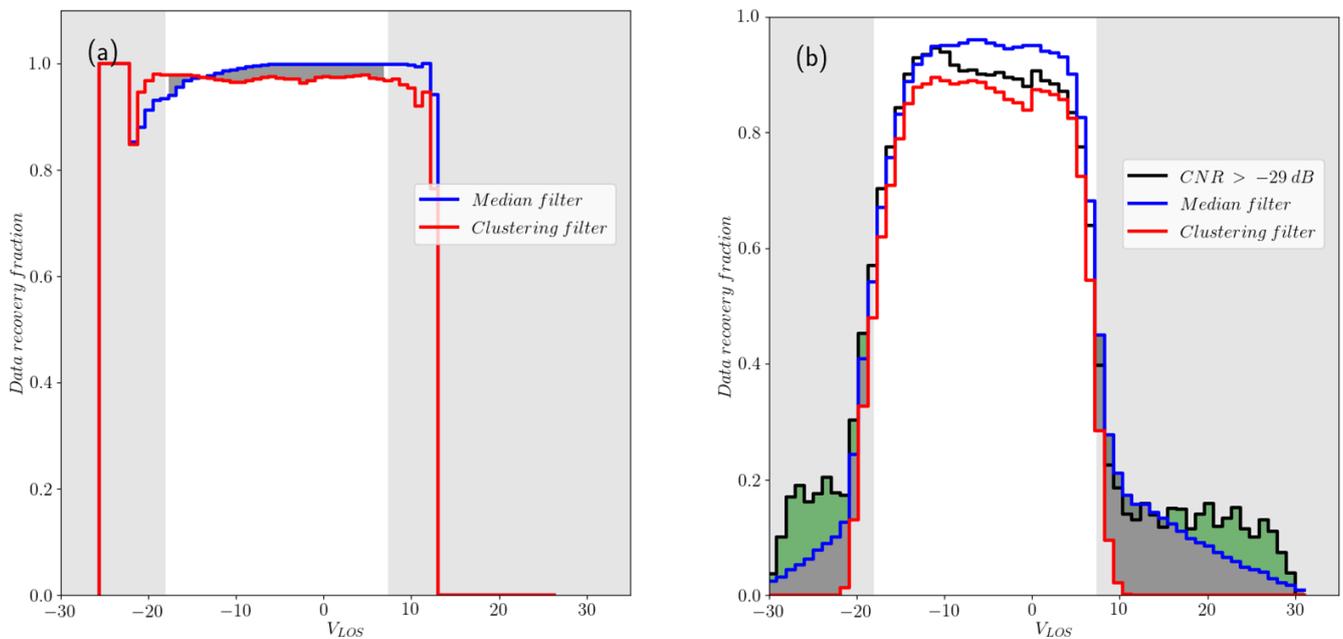


Figure 11. Distribution of recovery fraction per wind speed bin for phase 1 of the experiment of (a) reliable observations and (b) nonreliable data for CNR and median and clustering filters. The shadowed area in both graphs corresponds to the region where observations exceed the 3σ limit in the PDF of reliable observations. Again, the darker shadowed area highlights the additional fraction of extreme values nonfiltered by the median-like and CNR filters when the former uses the optimal input set $n_r = 5$, $n_\phi = 3$ and $\Delta V_{\text{LOS, threshold}} = 2.33 \text{ m s}^{-1}$.

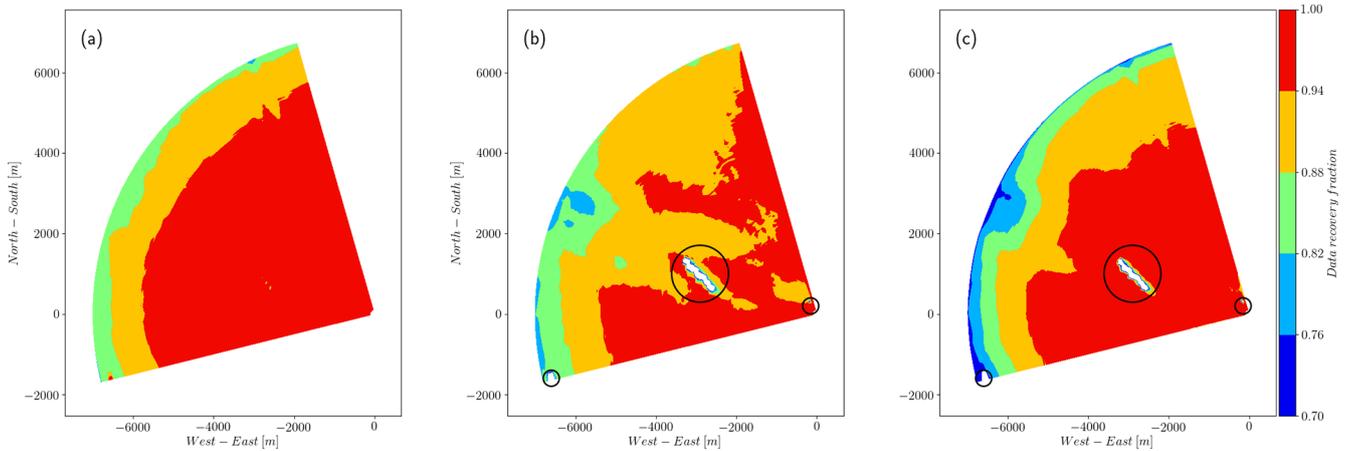


Figure 12. Total recovery fraction for phase 1 of the experiment. The noisy and far regions of the scans show a high recovery, above 80 %, for (a) the CNR > -29 dB threshold filter and (b) the median-like filter and below 75 % for (c) the clustering filter. The three groups of hard targets (turbines and one meteorological mast close to the lidar), which are identified by the median and clustering filter with recovery rates below 20 %, are highlighted.

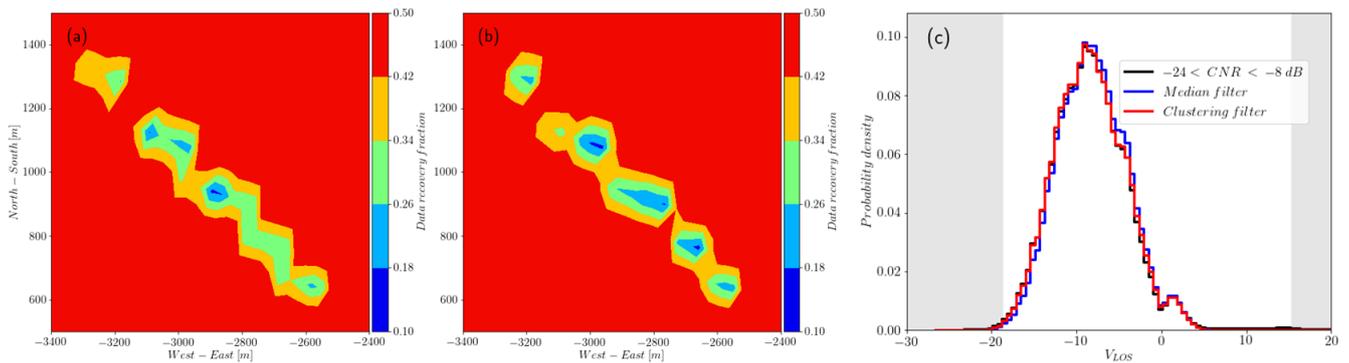


Figure 13. Detail of the recovery rate at the site of the turbines for the (a) median filter and (b) clustering filter. The recovery is lower in the flow regime of the turbine cluster (there are seven turbines in line) and higher in their surroundings for the clustering filter. Red denotes recovery rates of 0.5 or higher. (c) Probability density of V_{LOS} around the group of turbines.

patterns in vertical and horizontal planes as well as wind fields over different topography would make this analysis more general, thus shedding light on the capabilities of the filters presented here. This is especially critical regarding the median-like filter, which might require again a sensitivity analysis to select proper parameters that adapt to different scanning patterns and turbulence field characteristics. So far, $\Delta V_{LOS, threshold}$ showed a dependence on the L and $\alpha \varepsilon^{2/3}$ parameters during the sensitivity analysis presented in Sect. 6.1. Larger fluctuations in the V_{LOS} field, whether they come from larger turbulent structures, higher turbulence energy, or both, will need a larger value of $\Delta V_{LOS, threshold}$ to avoid the rejection of good measurements. Range height indicator (RHI) scanning patterns can pose the challenge of strong vertical shear and small turbulent structures that will need to reduce the window size n_r and n_ϕ for the median-like filter and the selection of a different set of features (or a new

definition for ΔV_{LOS}) for the clustering filter in order to keep reliable observations from being filtered out.

Regarding feature selection, the clustering filter could consider the spectral spreading of the heterodyne signal, S_b , and time variation of V_{LOS} in addition to features already used in this work to characterize and distinguish better clustering of good measurements. Nevertheless, due to the Euclidean distance definition, additional dimensions will make the data more sparse in higher dimensions, making it necessary to use more data points per filtering step (here we used only three scans at a time) to avoid the identification of good observations as spread-out, low-density noise. It is because of this that the application of a feature selection method might be necessary (Chandrashekar and Sahin, 2014).

Using the statistical distances D_K and D_{KL} as a metric for the filter performance might not be totally correct. At range gates far from the lidar, the distance between beams increases as well as the area covered by the accumulation of spectral

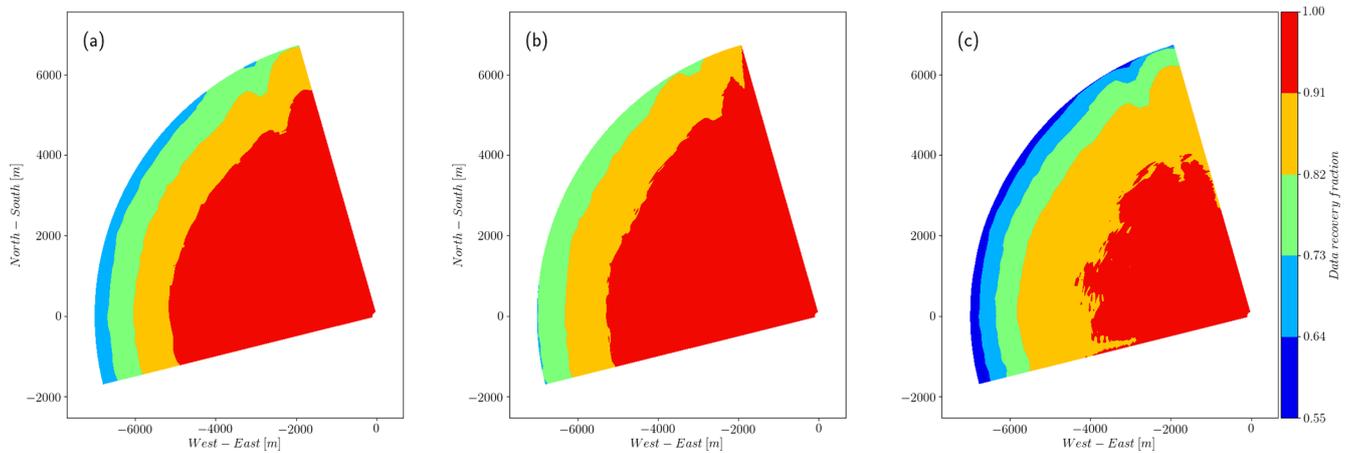


Figure 14. The total recovery fraction of observations for phase 2 of the experiment. The noisy and far regions of the scans show a high recovery, above 70 %, for (a and b) the CNR > −29 dB threshold and median-like filters, respectively. The recovery decreases to 55 % in the same region for the clustering filter, in line with the previous results, assuming that outliers (above the 3σ limit) and noise are more likely to be located here.

Table 5. Additional data recovered, relative to the number of observations in the reliable range of CNR, and fraction of data recovered with values beyond the 3σ range.

Phase 1, 3σ quantiles = [−18.16, 3.96] m s ^{−1}	Fraction of data recovered beyond 3σ	Additional data recovered
CNR threshold > −29 dB	27.1 %	23.4 %
Median filter	14.0 %	23.1 %
Clustering filter	8.6 %	22.1 %
Phase 2, 3σ quantiles = [−18.08, 7.35] m s ^{−1}		
CNR threshold > −29 dB	16.5 %	40.4 %
Median filter	12.6 %	42.4 %
Clustering filter	3.2 %	38.1 %

information in the azimuth direction. Averaging V_{LOS} over larger areas as we move forward through each beam might affect the statistics and the PDF of V_{LOS} (especially its spread) in the outer region of the scan. The fact that this region is where we usually find the nonreliable measurement group may make the PDFs of reliable and nonreliable observations somewhat different. These possible deviations need to be investigated further.

The selection of features and the number of scans put together per filtering step or iteration could also be automated using feature selection methods. Nevertheless, this would make the clustering filter more complex in its implementation and more computationally expensive, which is the main disadvantage of this approach compared to the median-like filter. Very efficient median filters can achieve a computational complexity up to $\mathcal{O}(n)$, with n being the number of observations in the data set. Depending on the data structure, DBSCAN shows a computational complexity from $\mathcal{O}(n \log(n))$ to $\mathcal{O}(n^2)$. If the distance between points is in general smaller than ε , the first limit can be achieved, but

clusters with different densities makes the algorithm less efficient. In the data analyzed here, having clusters with different densities is not an issue. Nevertheless, for nonhomogeneous flows, scans might persistently show regions with V_{LOS} , CNR, or other features with noticeably different values. It may then be necessary to revisit the clustering algorithm used and implement an ε -independent clustering approach, like OPTICS (Ankerst et al., 1999) for instance.

8 Conclusions

The CNR threshold filtering has been the common approach to retrieve reliable observations from lidar measurements. In this work we compared this approach against two alternative techniques: a median-like filter based on the assumption of smoothness of the wind field, hence in the smoothness of the radial wind speed observed by a wind lidar, and a clustering filter based on the assumption of self-similarity of the observations captured by the wind lidar and the possibility

of clustering them in groups of good data and noise. A controlled test was carried out on the last two approaches using a simple lidar simulator that sampled scans from synthetic wind fields later contaminated with procedural noise. The results indicate that the clustering filter is capable of detecting more added noise than the median-like filter at a good recovery rate of noncontaminated data. When the three filters are tested on real data, the clustering approach shows a better performance when identifying abnormal observations, increasing the data availability between 22 % and 38 % and reducing the recovery of abnormal measurements between 70 % and 80 % when compared to a CNR threshold. This is an important result because it increases the spatial coverage of the data which can be used later for wind field reconstruction and wind data analysis, especially in the far region of the scan, which covers the largest measured area.

Even though the median-like filter is computationally efficient, it needs an optimal definition of input parameters, which are dependent on the turbulence characteristics of the wind field. The clustering filter is more robust in this sense because it is capable of automatically adapting its parameters to the structure of the data. This is a step forward to a more robust and automated processing of data from lidars, which ideally should be independent of the turbulence characteristics of the measured wind field or the scanning pattern used.

Appendix A: Sensitivity analysis on median-like filter parameters

Figure A1 shows contours that present the optimal value for η_{tot} among all possible values of $\Delta V_{\text{LOS, threshold}}$ and n_{ϕ} for $n_r = 5$, the optimal window size in the radial direction. Large $\Delta V_{\text{LOS, threshold}}$ results in large η_{recov} but poor results for η_{noise} and the opposite for values of the threshold, as expected. The metric η_{tot} then becomes relevant to determine the optimal combination of parameters. From the contours it is possible to see that the performance in terms of the η_{tot} metric is less sensitive to n_{ϕ} than $\Delta V_{\text{LOS, threshold}}$. Even though the results here show average metrics for all the scans simulated, the optimal value of $\Delta V_{\text{LOS, threshold}}$ increases with the turbulence energy and length scale parameters, which is problematic because it requires previous knowledge of turbulence characteristics that usually are not available before reconstruction and, more importantly, data filtering.

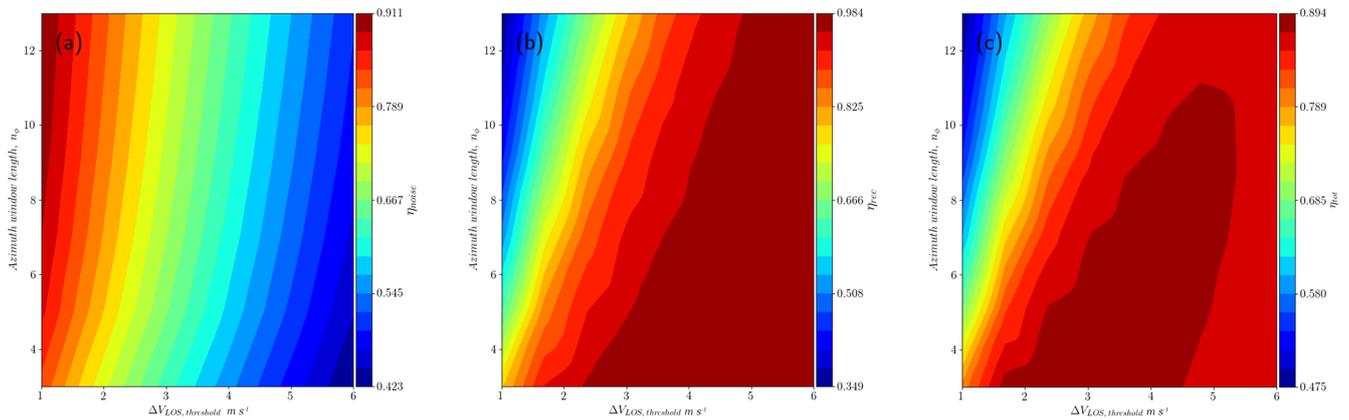


Figure A1. Contours of performance metrics for $n_r = 5$ over the $\Delta V_{\text{LOS, threshold}}-n_{\phi}$ space. Each point in the contour plot corresponds to the mean value of (a) η_{noise} , (b) η_{rec} and (c) η_{tot} among all the 4305 synthetic scans filtered. The optimal value corresponds to $n_r = 5$, $n_{\phi} = 3$, and $\Delta V_{\text{LOS, threshold}} = 2.33 \text{ m s}^{-1}$.

Code and data availability. The synthetic data and code are available at <https://doi.org/10.5281/zenodo.4014151> (Alcayaga, 2020). Real data can be found at <https://doi.org/10.11583/DTU.7306802.v1> (Simon and Vasiljevic, 2018).

Author contributions. LA implemented the analysis on synthetic and real data and wrote all sections of this paper.

Competing interests. The author declares that there is no conflict of interest.

Acknowledgements. I would like to thank Ioanna Karagali for comments and guidance through the data from the Østerild balconies experiments; Robert Menke for the first version of the median-like filter; and Ebba Dellwik, Nikola Vasiljevic, Gunner Larsen, Mark Kelly, and Jakob Mann for the very useful and clarifying feedback during the analysis and writing process.

Review statement. This paper was edited by Murray Hamilton and reviewed by two anonymous referees.

References

- Alcayaga, L.: Lidar data filtering algorithms, Zenodo, <https://doi.org/10.5281/zenodo.4014151>, 2020.
- Ankerst, M., Breunig, M. M., Kriegel, H.-P., and Sander, J.: OP-TICS: Ordering Points To Identify the Clustering Structure, in: Proc. ACM SIGMOD'99 Int. Conf. on Management of Data, 1–3 June 1999, Philadelphia, Pennsylvania, USA, pp. 49–60, ACM Press, 1999.
- Backer, E.: Computer-assisted Reasoning in Cluster Analysis, Prentice Hall International (UK) Ltd., Hertfordshire, UK, 1995.
- Banakh, V. A. and Smalikho, I. N.: Estimation of the Turbulence Energy Dissipation Rate from the Pulsed Doppler Lidar Data, *Atmos. Ocean. Opt.*, 10, 957–965, 1997.
- Beck, H. and Kühn, M.: Dynamic Data Filtering of Long-Range Doppler LiDAR Wind Speed Measurement, *Remote Sens.*, 9, 561, <https://doi.org/10.3390/rs9060561>, 2017.
- Brousmiche, S., Bricteux, L., Sobieski, P., Macq, B., and Winckelmans, G.: Numerical simulation of a heterodyne Doppler LiDAR for wind measurement in a turbulent atmospheric boundary layer, in: 2007 IEEE International Geoscience and Remote Sensing Symposium, 23–27 July 2007, Barcelona, Spain, <https://doi.org/10.1109/IGARSS.2007.4423420>, 2007.
- Burger, W. and Burge, M. J.: Digital Image Processing – An Algorithmic Introduction using Java, Texts in Computer Science, Springer, London, UK, 2008.
- Cariou, J.: Remote Sensing for Wind Energy, chap. Pulsed lidars, DTU Wind Energy, Denmark, 131–148, 2015.
- Chandrasekar, G. and Sahin, F.: A Survey on Feature Selection Methods, *Comput. Electr. Eng.*, 40, 16–28, <https://doi.org/10.1016/j.compeleceng.2013.11.024>, 2014.
- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X.: A Density-based Algorithm for Discovering Clusters a Density-based Algorithm for Discovering Clusters in Large Spatial Databases with Noise, in: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD'96, 2–4 August 1996, Portland, Oregon, 226–231, 1996.
- Gryning, S. and Floors, R.: Carrier-to-Noise-Threshold Filtering on Off-Shore Wind Lidar Measurements, *Sensors*, 19, 592, <https://doi.org/10.3390/s19030592>, 2019.
- Gryning, S., Floors, R., Peña, A., Batchvarova, E., and Brümmner, B.: Weibull Wind-Speed Distribution Parameters Derived from a Combination of Wind-Lidar and Tall-Mast Measurements Over Land, Coastal and Marine Sites, *Bound-Lay. Meteorol.*, 159, 329, <https://doi.org/10.1007/s10546-015-0113-x>, 2016.
- Huang, T., Yang, G., and Tang, G.: A fast two-dimensional median filtering algorithm, *IEEE T. Acoust. Speech*, 27, 13–18, <https://doi.org/10.1109/TASSP.1979.1163188>, 1979.
- Karagali, I., Mann, J., Dellwik, E., and Vasiljević, N.: New European Wind Atlas: The Østerild balconies experiment, *J. Phys. Conf. Ser.*, 1037, 052029, <https://doi.org/10.1088/1742-6596/1037/5/052029>, 2018.
- Kolmogorov, A.: Sulla determinazione empirica di una legge di distribuzione, *Inst. Ital. Attuari, Giorn.*, 4, 83–91, <https://ci.nii.ac.jp/naid/10010480527/en/> (last access: 4 September 2020), 1933.
- Kullback, S. and Leibler, R. A.: On Information and Sufficiency, *Ann. Math. Stat.*, 22, 79–86, 1951.
- MacQueen, J.: Some methods for classification and analysis of multivariate observations, in: Proceedings Fifth Berkeley Symp. on Math. Statist. and Prob., vol. 1: Statistics, 21 June–18 July 1965 and 27 December 1965–7 January 1966, Berkeley, California, USA, 281–297, 1967.
- Mandelbrot, B.: The fractal geometry of nature, W. H. Freeman and Comp., New York, USA, 1983.
- Mann, J.: The spatial structure of neutral atmospheric surface-layer turbulence, *J. Fluid Mech.*, 273, 141–168, <https://doi.org/10.1017/S0022112094001886>, 1994.
- Mann, J.: Wind field simulation, *Probabilist. Eng. Mech.*, 13, 269–282, [https://doi.org/10.1016/S0266-8920\(97\)00036-2](https://doi.org/10.1016/S0266-8920(97)00036-2), 1998.
- Mann, J., Angelou, N., Arnqvist, J., Callies, D., Cantero, E., Arroyo, R. C., Courtney, M., Cuxart, J., Dellwik, E., Gottschall, J., Ivanell, S., Kühn, P., Lea, G., Matos, J. C., Palma, J. M. L. M., Pauscher, L., Peña, A., Sanz, R. J., Söderberg, S., Vasiljevic, N., and Veiga, R. C.: Complex terrain experiments in the new european wind atlas, *Philos. T. R. Soc. A*, 375, 20160101, <https://doi.org/10.1098/rsta.2016.0101>, 2017.
- Menke, R., Vasiljević, N., Wagner, J., Oncley, S. P., and Mann, J.: Multi-lidar wind resource mapping in complex terrain, *Wind Energ. Sci.*, 5, 1059–1073, <https://doi.org/10.5194/wes-5-1059-2020>, 2020.
- Meyer Forsting, A. and Troldborg, N.: A finite difference approach to despiking in-stationary velocity data - tested on a triple-lidar, *J. Phys. Conf. Ser.*, 753, 072017, <https://doi.org/10.1088/1742-6596/753/7/072017>, 2016.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E.: Scikit-learn: Machine Learning in Python, *J. Mach. Learn. Res.*, 12, 2825–2830, 2011.

- Perlin, K.: Noise hardware, in: Real-Time Shading SIGGRAPH, Course Notes, Baltimore, Maryland, USA, 2001.
- Simon, E. and Vasiljevic, N.: Østerild Balconies Experiment (Phase 2), figshare, <https://doi.org/10.11583/DTU.7306802.v1>, 2018.
- Smalikho, I. N. and Banakh, V. A.: Accuracy of estimation of the turbulent energy dissipation rate from wind measurements with a conically scanning pulsed coherent Doppler lidar. Part I. Algorithm of data processing, *Atmos. Ocean. Opt.*, 26, 404–410, <https://doi.org/10.1134/S102485601305014X>, 2013.
- Stawiarski, C., Träumner, K., Knigge, C., and Calhoun, R.: Scopes and Challenges of Dual-Doppler Lidar Wind Measurements – An Error Analysis, *J. Atmos. Ocean. Tech.*, 30, 2044–2062, <https://doi.org/10.1175/JTECH-D-12-00244.1>, 2013.
- Vasiljević, N., Lea, G., Courtney, M., Cariou, J., Mann, J., and Mikkelsen, T.: Long-Range WindScanner System, *Remote Sens.*, 8, 896, <https://doi.org/10.3390/rs8110896>, 2016.
- Vasiljević, N., L. M. Palma, J. M., Angelou, N., Carlos Matos, J., Menke, R., Lea, G., Mann, J., Courtney, M., Frölen Ribeiro, L., and M. G. C. Gomes, V. M.: *Perdigão 2015: methodology for atmospheric multi-Doppler lidar experiments*, *Atmos. Meas. Tech.*, 10, 3463–3483, <https://doi.org/10.5194/amt-10-3463-2017>, 2017.