Windlab Limited - LIDAR for Wind Forecasts Project
LESSONS LEARNT REPORT #1

Funding Agreement Details

Recipient Name	Windlab Limited
Project Commencement Date	03 April 2019
Project Completion Date	13 July 2020
Primary Contact Name	Keith Ayotte
Contact Telephone Number	0433 108 393 (Mobile) 6175 4607 (Office)
Contact Email	keith.ayotte@windlab.com
Reporting Period	April/October 2019

LESSONS LEARNT

Key Lesson/s
These points follow from the supporting information supplied below: 1) Theory and practice show that randomness essentially dilutes the performance of all-time series prediction algorithms. The objective of our early work is to understand the origins of the randomness and, where possible, define a strategy for minimising its effects. Doing this gives the prediction algorithms the best chance of being able to predict the time series into the future. This applies to both the wind turbine SCADA measurements and the upstream LIDAR measurements. 2) The placement of the LIDARs is key to optimising the correlation between upstream measurements and measured wind speed (or generation) at any turbine location. The placement of measurement locations gives a pattern of influence that is quite variable across a geographic area the size of even a small wind farm. As such, careful choice of measurement locations is key to coverage of all turbines. 3) The measured hub height wind speed and the resultant generation contain a high level of randomness. This randomness varies throughout the day. This turbulence/randomness is strongest during the daytime when the sun's heating provides buoyancy for the creation of turbulent eddies. This is seen very consistently in the diurnally changing variance in both the horizontal and vertical wind speeds and is expected to continue and become stronger as we move into summer when the heating is strongest. Any forecast algorithm needs to accommodate this diurnal change in variance. 4) 4.) The LIDARs used in the project measure horizontal and vertical wind speeds every minute at several levels up to 133m. This capability has been key to building a complete picture of the atmospheric processes that are in play across the wind farm as we are trying to make generation predictions.

Implications for Future Projects

Points 1 and 3 above are generally applicable to forecasting any atmospheric process on time scales of a few to several minutes. The naturally occurring random variation in atmospheric boundary layer flows caused by turbulent motions have time scales of order seconds to tens of minutes. In the context of 5min forecasts, only the largest of these variations will be predictable. Therefore, significant effort must be made in making and analysing high quality measurements up front in order to clearly understand the processes at work in determining the state of the atmosphere 5 minutes hence.

Knowledge Gaps Identified

To date the project has focused on data analysis aimed at understanding wind speed/generation at each turbine as it relates to that of each of the other turbines as well as the upstream LIDAR wind measurements. The turbulent state of the atmospheric boundary layer has a strong effect on the predictability of the wind and therefore wind farm generation. Understanding the time scales associated with turbulent motions under various atmospheric conditions is likely to be key to being able to appropriately filter out their effects and increasing the accuracy of the 5-minute forecast. Given this understanding, as the project moves forward the focus will turn to applying this understanding to a number of forecast algorithms to determine the most appropriate to the task.

See *Supporting Information* below

Supporting Information

Evidence of the underlying principle

This project makes use of an idea that underpins all fluid flow. That simple idea is that the properties of fluid flow, including velocity that we are seeking to predict, are transported along with that flow to locations downstream in the flow. This leaves open the opportunity to measure upstream in the flow to estimate future values of the property being predicted, in this case wind speed. For example, if the air is moving at 10m/s, in five minutes a parcel of air will move 3km (10m/s X 300s=3000m). Therefore, if we measure at a distance 3km upstream, the properties of the flow will reach us in five minutes time. To see the effect of this more clearly, Figure 1 shows the wind speed measured at a wind turbine at the Kiata Wind Farm (blue line) and a lidar located approximately 3km upstream from the turbine in the direction from which the wind is blowing at the time (orange line).

The data are at 1min intervals. The arrows show the peaks in the wind speed that are roughly 5min apart in time.

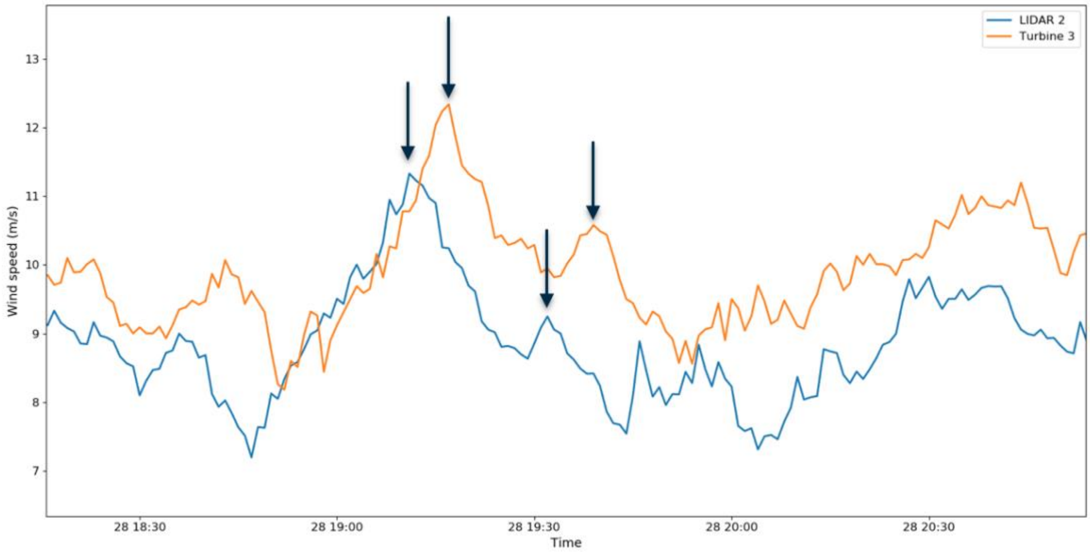

__Figure 1__ Wind speed time series measured a wind turbine and a LIDAR separated by roughly 3km along a line roughly parallels with the wind direction. Arrows show peaks in wind speed roughly 5min apart.

The LIDAR and turbine SCADA measurements represented in this figure point to significant potential for this principle to provide information useful in forecasting wind speeds and/or generation up to several minutes in advance and in so doing enhance the accuracy of wind forecasts. However, the above demonstrated principle comes with a number of challenges.

LIDAR placements

Figure 2 shows the turbine and LIDAR locations relative to a point in the geographic centre of the Kiata Wind Farm, with the turbine locations represented by stars and the LIDAR locations represented by crosses. The shading on the figure shows the relative probability that a parcel originating from either one of the LIDARs will be in that location five minutes after the measurement. The figure shows good coverage of the southern turbines by the southern LIDAR location. Though showing little turbine coverage, the northern LIDAR position provides a high-quality background measurement separated from the southern LIDAR by nearly 5km, with which cross-correlation of the two LIDAR measurements can be performed. This is key to understanding the distance downwind that measurements can be expected to be correlated and what that correlation looks like in a statistical sense. The LIDARs have been in these locations since their original installation, just over four months. It is expected that one or both of the LiDARs will be relocated in November of this year.

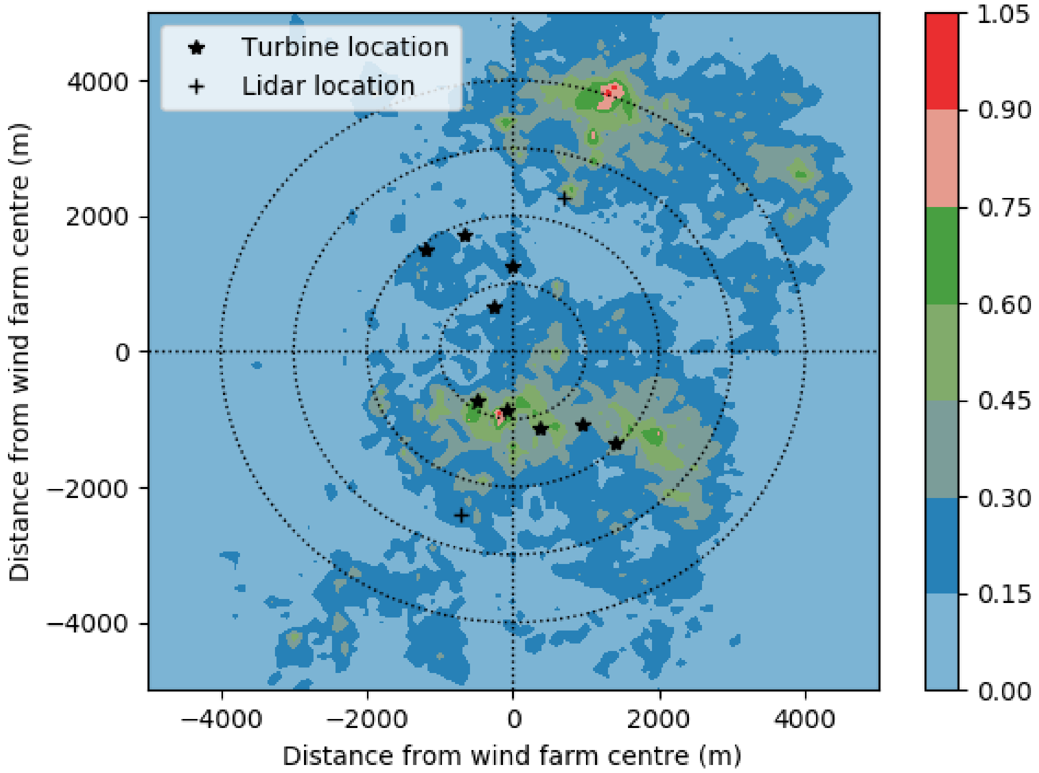

Figure 2 *LIDAR (crosses) and turbine (stars) placement relative to the geographic centre of the Kiata Wind Farm. Shading shows relative probability of air parcels in the flow being in a location five minutes after measurement.*

The relative probability distribution across the figure shows how the geographic coverage of any measurement position varies significantly spatially, making the choice of measurement locations critical for coverage of all of the wind turbines. The figure also suggests that a key location for measurement is roughly 2km to the NNW of the current southern placement. Having movable measurement platforms is clearly a significant advantage in the project.

The physical and statistical challenge of randomness in the atmosphere

The challenge in forecasting parameters such as wind speed or resultant generation is that the atmosphere contains a great deal of randomness. This randomness and its physical underpinnings and statistical nature occupy a significant portion of the branch of the atmospheric sciences called micrometeorology, which is largely the study of the ways in which turbulent (random) motions transport heat, moisture and momentum throughout the lowest few kilometres of the atmosphere. Wind turbines operate in this same area of the atmosphere both physically and phenomenologically.

Where this becomes important in the context of this project is that all time series prediction methods are sensitive to noise or randomness in either the time series being predicted or the series that are being used as inputs to the prediction. In the extreme, an attempt to predict an entirely random, and by that definition unpredictable, time series of values will result in what is essentially a persistence forecast – one in which the predicted value is on average the same as the current value. Further to this, adding sophistication to the prediction methodology yields very little in terms of increasing the skill of the forecast. As such, the task of forecasting 5-minutes ahead is therefore really one of identifying and separating the signal from the noise in the measurements and basing the forecast on that signal while discarding the effects of the natural randomness in the turbulent atmosphere.

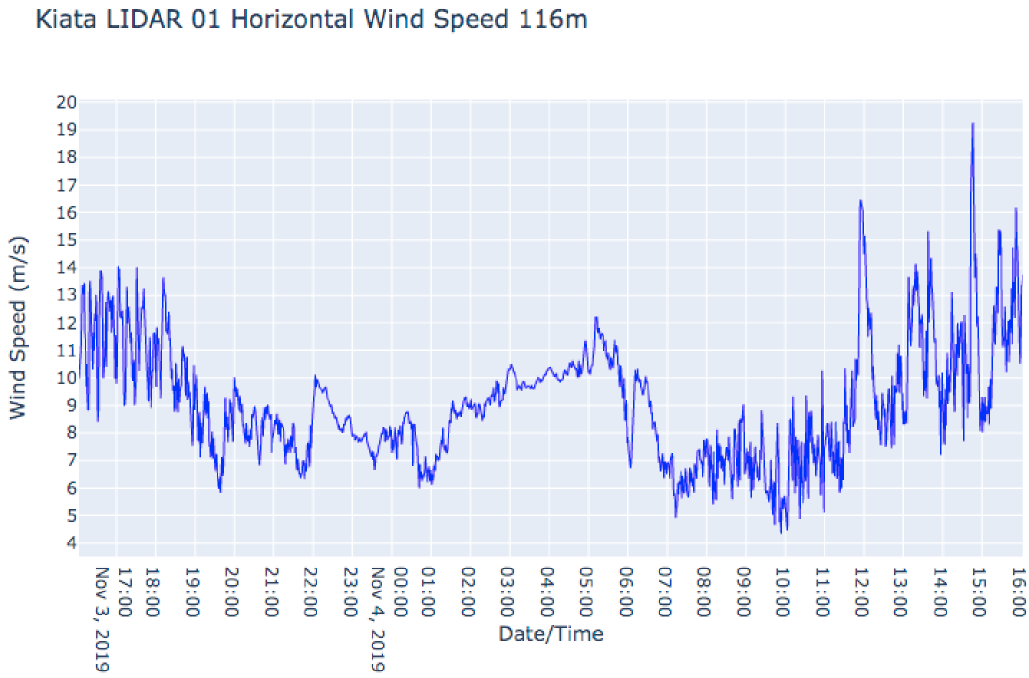

Figure 3 *Horizontal wind speed measured by a LIDAR at 116m above the surface for a 24-hr period. Strong turbulence/randomness is evident at the far left and right of the figure, where the sun's heating supplies buoyancy which in turn drives turbulent motions.*

Figure 3 shows the wind speed measured from LIDAR 1 at the Kiata Wind Farm at 116m above the ground. The turbulence during the daytime hours (far left and right in the figure) is easily seen in contrast to during the night, when the suns heating is not present to create buoyancy and the resulting turbulent motions. Consistent with our physical understanding of turbulent motions in the atmospheric boundary layer, we can also see the strong vertical component of the turbulent motions in Figure 4 during the same periods during which the horizontal wind speed is fluctuating the strongest. Fortunately, the sophisticated nature the LIDAR measurements we have allow us to understand more clearly the origins and the statistical nature of the randomness in the measured wind speed we are trying to predict.

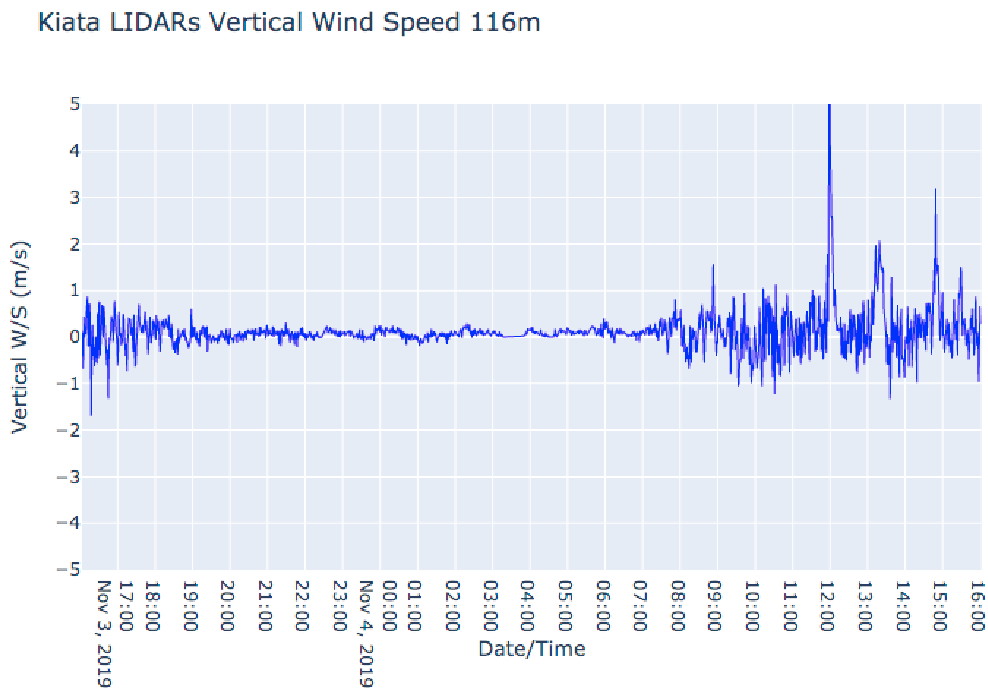

Figure 4 *Vertical wind speed measured by a LIDAR at 116m above the surface for a 24-hr. period. Strong turbulence/randomness is evident at the far left and right of the figure, where the sun's heating supplies buoyancy which in turn drives turbulent motions.*

Quantifying randomness

Figure 5 shows wind speed measured by one of the LIDARs at Kiata Wind Farm at 116m above the ground. The blue line is the unfiltered measurements while the red line is the result of passing a wavelet filter through the unfiltered signal, that is specifically chosen so that the statistical distribution of the residuals is close to Gaussian in form.

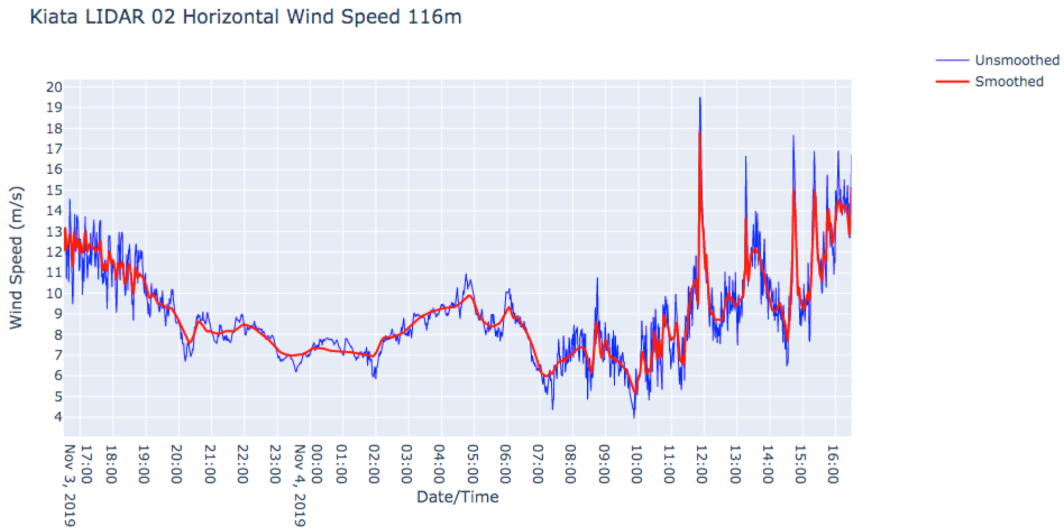

Figure 5 *Smoothed (red) and unsmoothed (blue) horizontal wind speed measured by a LIDAR at 116m above the surface for a 24-hr. period. Smoothing is done with wavelet filter.*

Figure 6 shows the probability distribution of the residuals from the smoothing process. The red line in the plot is an exact Gaussian distribution with the same standard deviation as the measured distribution. The measured distribution is clearly near-Gaussian, suggesting the residuals are largely random in nature, as we would expect turbulent motions to be. Choosing different type and levels of filtering give residual distributions that are more/less Gaussian. This points to the sensitivity of the filtering process to the type/weight of filtering used.

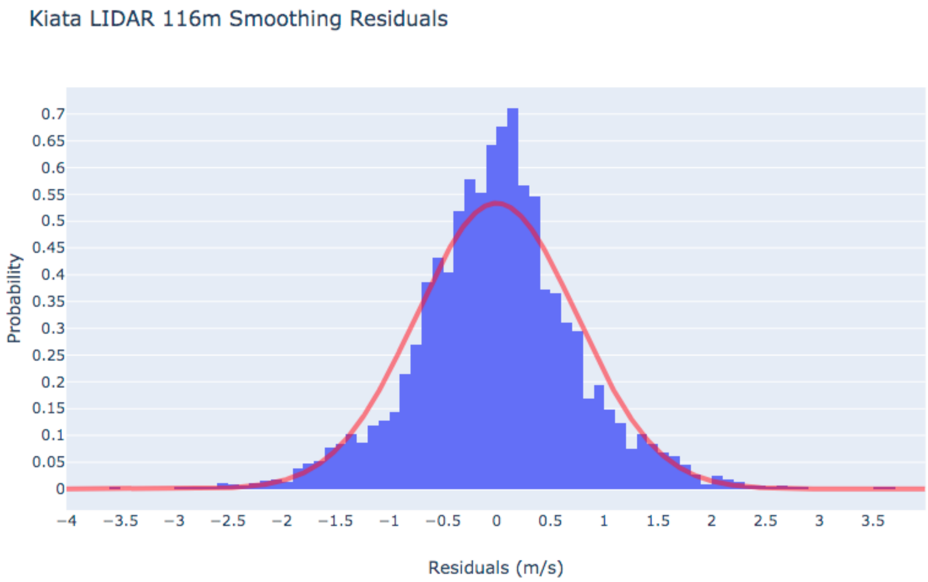

Figure 6 *Probability distribution of residuals from smoothing of horizontal wind speed shown in Figure 5 above.*