

AN ANALYSIS OF WIND RETRIEVAL ALGORITHMS FOR SMALL UNMANNED AERIAL SYSTEMS

Timothy Bonin^{*1,2}, Brett Zielke¹, Wascar Bocangel^{2,3}, Wayne Shalamunenc⁴, and Phillip Chilson^{1,2}

¹School of Meteorology, University of Oklahoma, Norman, OK

²Atmospheric Radar Research Center (ARRC), University of Oklahoma, Norman, OK

³School of Electrical and Computer Engineering, University of Oklahoma, Norman, OK

⁴Central Oklahoma Radio Control Society, Norman, OK

1. INTRODUCTION

Unmanned Aerial Systems (UASs) are quickly becoming important tools for monitoring and researching the lower atmosphere, primarily the planetary boundary layer (PBL) (Holland et al. 2001; Spiess et al. 2007; van den Kroonenberg et al. 2008; Reuder et al. 2009). UASs are relatively inexpensive and have the ability to collect samples with high spatial and temporal resolution. Flight plans can be customized to obtain particular data depending on the phenomenon under examination. For example, a flight plan with a quick ascent rate could be used to rapidly penetrate the PBL and study the convective boundary layer late in the day. A flight with a very slow ascent rate might be used to acquire high vertical resolution sampling over a shallow layer to study an early morning inversion. Such a flight plan could then be repeated every thirty minutes with each subsequent flight modified to achieve slightly higher altitudes in order to observe an early morning transition from a very stable boundary layer to a convective boundary layer.

Since UASs are increasingly being utilized for atmospheric observations and research, accurate instrumentation and observation methods must be developed for these platforms. Moreover, the instrumentation must be small and lightweight due to weight restrictions. Recently, the Atmospheric Radar Research Center (ARRC) at the University of Oklahoma (OU) has begun developing a small UAS, the SMARTSonde (Small Multi-function Autonomous Research and Teaching Sonde), for planetary boundary layer research.

The SMARTSonde is capable of measuring thermodynamic parameters such as pressure, temperature, relative humidity, and amounts of trace gases such as ozone. While these parameters can be directly measured during flights from instrumentation, more complex methods are necessary to determine the wind speed and direction from SMARTSonde data.

Currently, three different algorithms are used to calculate the wind profile from SMARTSonde flights. The algorithms are largely based on measurements from an onboard GPS unit. A pitot tube is also being investigated as a supplement to the GPS measurements as a means of improving performance of the algorithms. Derived wind speeds could be used to complement the thermodynamic variables to calculate boundary layer stability parameters, such as the gradient Richardson number. Wind profiles could also potentially be useful for severe weather events, for example, low-level wind shear data obtained during pre-storm environments as a means of determining the likelihood of tornadogenesis. The combination of thermodynamic variables with wind information could be used to study transition periods, inversions, and the low-level jet. Stability indices during these studies could be calculated with the supplementary wind data.

2. BACKGROUND

2.1. Instrumentation to Study the Planetary Boundary Layer

Due to the large impact of diurnal fluctuations on the structure of the PBL, it evolves faster than many other meteorological phenomena, which has important industrial impacts. Examples include air quality in urban areas and PBL winds for wind turbines (Kondragunta et al. 2008; Seaman

* Corresponding author address: Timothy Bonin, School of Meteorology, University of Oklahoma, 120 David L. Boren Blvd., Norman, OK 73072. E-mail: tim.bonin@ou.edu

and Michelson 2000; Endlich et al. 1982; Emeis et al. 2007). Comprehending the diurnal fluctuations in the wind is especially important for wind energy forecasting. Studying these phenomena can be difficult and might require a variety of in-situ measurements from instrumented towers, which can be relatively expensive to install and maintain. Although fixed location measurements from towers are well suited for studying some aspects of the PBL, their scope is somewhat limited. Rawinsondes can also be used to study the vertical structure of the PBL, but they too are relatively expensive and one has no control of the flight direction.

Radars, lidars, sodars, wind profilers, and other remote sensing tools are used to measure PBL variables as well. Most remote sensors can operate continuously without much human intervention, but one must rely on retrieval algorithms to obtain data and thermodynamic variables are difficult to measure with most remote sensing instruments. Thus, the SMARTSonde, along with other UASs in development, are significant and unique instruments in the boundary layer research community. These in-situ instruments are capable of measuring both thermodynamic variables and the wind, while keeping costs low and providing flexibility to the user.

2.2. UAS Systems and Wind Profiling

There are a number of ongoing projects around the world that are using small aircraft to study various weather phenomena. A wide variety of airframes and instrumentation are in use for a large number of applications. Researchers have also devised several methods for measuring the wind for the different platforms.

The Aerosonde was one of the first small robotic aircrafts that was used for atmospheric research. In the early 1990s, work began on the platform design for atmospheric research (Holland et al. 1992, 2001). The Aerosonde is able to measure temperature, pressure, humidity, trace gas concentrations, wind, and several other variables. However, the wind measurement method is proprietary for AAI Corporation, the company that builds and manages the Aerosonde flights. Although, flight time for the Aerosonde can be purchased through the AAI Corporation; the UAS itself cannot be purchased.

In Norway, a group is working on using a small foam plane, the Small Unmanned Meteorologi-

cal Observer (SUMO), for atmospheric boundary layer research (Reuder et al. 2009). The wind speed has been estimated for the SUMO by essentially taking the difference between the maximum and minimum ground relative speed around a circle and dividing that difference by two. This requires constant throttle so that the airspeed is constant.

A Chinese group has also experimented with using a UAS to obtain soundings of the atmosphere. Mhuqing et al. (2004) built a small aircraft and created an autopilot system for the instrument. They also created a new algorithm to measure the wind. It was based on the fact that if the UAV flies in a circle by maintaining a constant roll angle, the circle pattern that the plane flies in shifts downwind overtime at the same rate as the wind speed.

More complex methods for measuring the wind have also been devised. While previous methods used only the GPS derived data to calculate the wind speed and direction, others have used additional instrumentation to find the wind profile. For example, a five-hole pitot tube has been used in conjunction with an inertia measurement unit (IMU) for calculating a three-dimensional wind value (van den Kroonenberg et al. 2008). Using the five-hole probe, the wind measurements could be made at up to 40 Hz, allowing for quantification of turbulence. The standard deviation of the wind speed estimates were usually $\sim 0.5 \text{ ms}^{-1}$ and for wind direction were typically $\sim 10^\circ$.

Many different wind algorithms have been developed for specific platforms around the world and most should work across platforms provided the proper instrumentation is available. For example, the methods developed by Mhuqing et al. (2004) or Reuder et al. (2009) should work with any UAS equipped with a GPS unit. Unfortunately different wind algorithms have not been compared against each other to assess their relative accuracies.

2.3. SMARTSonde Data Set

For all SMARTSonde flights used in this study, the Paparazzi autopilot system (http://paparazzi.enac.fr/wiki/Main_Page) was used to both control the plane in-flight and transmit data down to a ground control station (GCS). The autopilot system is open source, which allows users to create flight plans for any airframe. Fig. 1 depicts a screenshot of the GCS program used

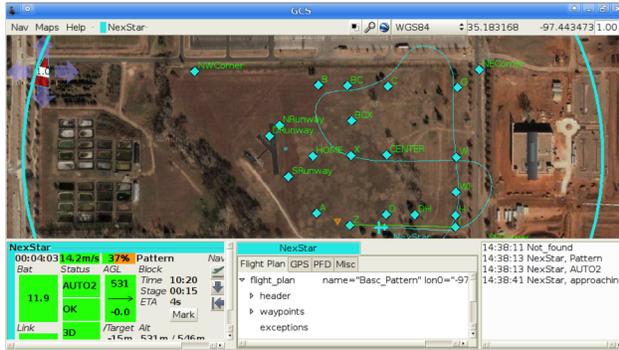


Figure 1: Screenshot of the GCS while the SMARTSonde is in flight. The plane's location is shown with its track by the thin light blue line. In the bottom left, the height, ground relative speed, climb rate, and other statuses are shown.

on a ground based computer to track the plane and monitor its status. In this case, the plane was flying back and forth over a set of waypoints; the track of the plane can be seen by the light blue line behind the plane.

Several different flight plans have been created for the SMARTSonde project, one of which is demonstrated in Fig. 1. That type of flight can be used to examine variations in thermodynamic properties of the atmosphere on a constant height surface, such as over different land surfaces. Another flight plan that is often used to retrieve vertical profiles of the boundary layer involves a helical ascent pattern. The throttle and pitch are held at a constant value, thereby maintaining a nearly constant airspeed. An example of this pattern, with corresponding thermodynamic variables, can be seen in Fig. 2. This helical ascent plan is used for all of the flights in this study. Profiles of the PBL obtained from these flights are compared to a local rawinsonde observation, the OUN site (WMO station number 72357). The SMARTSonde flights have been conducted at the Central Oklahoma Radio Control Society (CORCS) airfield, which is less than 1 km from the OUN site.

Thermodynamic data are shown in Fig. 2. These data are communicated back to the ground based computer from the SMARTSonde's instrument package through the autopilot interface. The SMARTSonde is always equipped with a SHT75 (Sensirion) for temperature and humidity measurements and an SCP1000 (VTI Technologies) to measure static atmospheric air pressure. Work is currently underway to install an Aeroqual SM50 (Aeroqual Ltd.) on the plane to measure atmo-

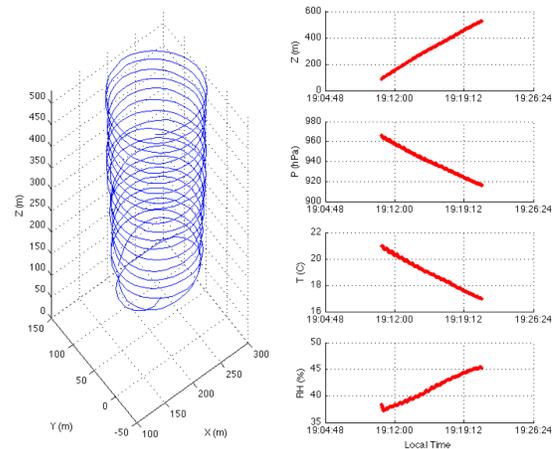


Figure 2: Example helical ascent flight with constant throttle and constant pitch. On the right, from top to bottom, the height AGL, pressure, temperature, and relative humidity are shown.

spheric gas concentrations, in particular to examine levels of ozone. A pitot tube (Eagle Tree) is also installed on the wing of the SMARTSonde in order to measure the plane's airspeed. However, this measurement cannot currently be communicated back to the GCS in real-time.

3. WIND RETRIEVAL ALGORITHMS

There are three algorithms that are currently in use for the SMARTSonde to determine the wind characteristics. The primary use of these algorithms is to obtain a vertical profile of the mean wind. With this profile, wind shear in the PBL can be measured. Since all of the algorithms involve temporal averages, they are not useful for determining small fluctuations in the u and v components of the wind at one height at short timescales. The three algorithms currently under examination are the best curve fitting method, Nelder-Mead optimization method, and the Paparazzi autopilot output.

3.1. Best Curve Fitting

One method of retrieving the wind information from a SMARTSonde flight is by fitting a curve to the plane's ground relative speed, which is provided by the on-board GPS unit. This method is similar to the wind algorithm used by Reuder et al. (2009). Below, equation 1 describes the plane's

ground-relative speed, Y , based on other variables. In the equation, ψ is the airplane heading with north being 0° , θ is the wind direction using meteorological convention, a is the airspeed of the plane, and v is the wind speed. For this method, a can be treated as a constant since the plane is flying with a constant throttle value. The wind vector consists of a headwind and crosswind component, which are labeled. Overall, the plane's ground relative speed is based on the summation of its along stream component (s) and normal component (n).

$$Y^2 = \underbrace{(a + v\cos(\psi - \theta + 180))}_{\text{headwind}}^2 + \underbrace{(v\sin(\psi - \theta + 180))}_{\text{crosswind}}^2 \quad (1)$$

The values of ψ and Y are known in this equation since they are relayed to the ground control station from the GPS unit every second. With these data points, a polynomial curve fitting can be performed to determine the best values of a , v , and θ . This polynomial fitting can be done as frequently as desired. However, the fitting may not be representative of the actual wind speed if the algorithm is run on too few data points. To get the best results and the best fitting, it is best to have a large number of data points over a wide range of ψ . In order to get a representative fit of the data for the entire range of ψ from 0° to 360° , the polynomial fitting is performed for each circle that the SMARTSonde completes. For this method to work, since a is treated as a constant, the plane needs to fly at a constant throttle.

Fig. 3 shows how the ground relative speed generally changes with the direction that the plane is traveling. On this particular day, the prevailing winds were from the south. This is shown in the flight track by the faster ground-relative speed when the plane flew north. The SMARTSonde was moving slower when it was headed south. This would indicate that it had a tailwind flying north and a headwind going south, hence slowing it down. Utilizing this fact, a section of the flight can be examined using the best curve fitting method. Fig. 4 shows a section of the flight when the plane was ascending from around 810 m to 890 m above ground level (AGL), during which the plane made a complete circle. Equation 1 is fitted to the instantaneous ground-relative velocities, which is shown

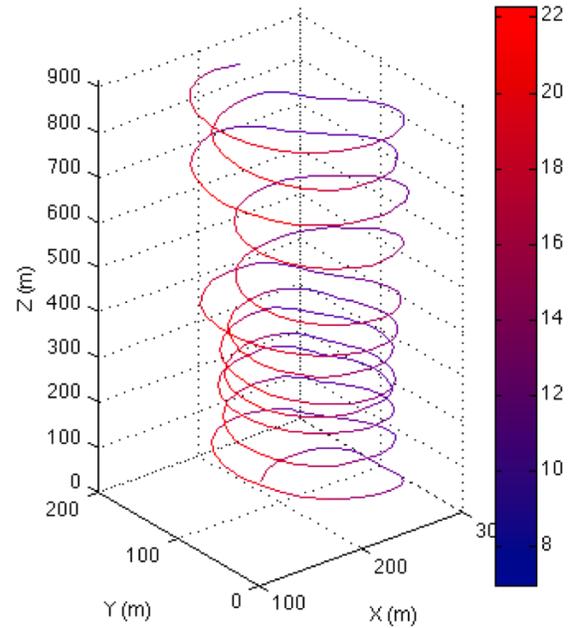


Figure 3: Example of how the ground relative speed, given by the color of the track, changes depending on the direction of travel. The speeds on the colorbar are given in ms^{-1} . The SMARTSonde was flying in a clockwise direction.

in figure 5. By applying this fitting, the wind speed v and the wind direction θ is found for the height of 850m AGL, since that is the average height during this flight segment. This type of fitting is then applied to every circle during the SMARTSonde's ascent. By doing this, a wind profile can be made for the PBL, as shown in Fig. 6 for this case.

It is important to note that the UAS does not need to fly in a circle to utilize this method. It could work with most patterns as long as the UAS turns at some point. However, the algorithm gets more frequent updates of the wind estimation when a circular flight path is used.

3.2. Nelder-Mead Optimization Method

The Nelder-Mead optimization method is another way to find the wind speed and direction. This method was originally used in the Paparazzi program; however, it is not well-documented and the software only produces output once every ten seconds. Originally written for sailplanes, this routine is written primarily to improve the performance of the autopilot in windy conditions. Again,

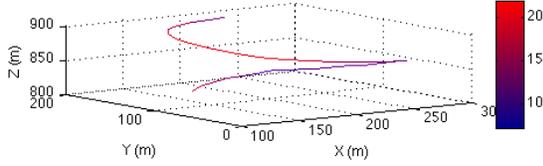


Figure 4: Same as figure 3 for a smaller flight segment.

this wind retrieval algorithm utilizes an optimization scheme using only the ground-relative velocity from the GPS unit. First, the airspeed a is defined as:

$$a = \frac{1}{n} \sum_{i=1}^n \|S(i) - W\| \quad (2)$$

In this equation, n is the number of the GPS measurements that are used in the optimization, which needs to be fairly small since the airspeed is assumed to be constant. S consists of the ground-relative plane speed measurements given by the GPS. W is the wind vector. Since $S - W - a = 0$, assuming perfect measurements and constant airspeed and wind speed, iterations for values of W can be made in order to find the correct value for the wind speed and direction, since its a three dimensional vector. To accomplish this, a quantity of the standard deviation, σ , must be minimized. This quantity is defined in equation 3.

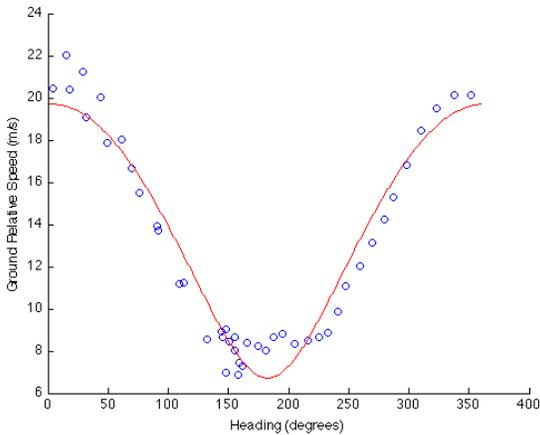


Figure 5: Example fitting of the equation in the best curve fitting method. This corresponds to the cutout in figure 4. The red curve is the best fit of equation 1 while the blue circles are the raw ground-relative velocities dependent on the heading direction.

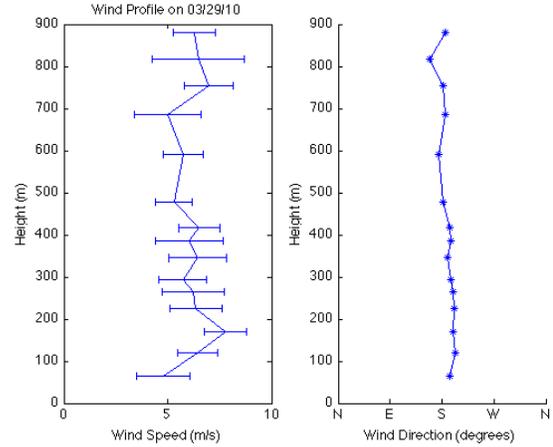


Figure 6: Wind profile by using the best curve fitting method on the flight shown in figure 4 on March 29, 2010. The horizontal lines in the wind speed are error bars. Their size depends on the average error of the best curve fitting compared to the true velocities.

$$\sigma = \frac{1}{n} \sum_{i=1}^n (\|S(i) - W\| - a)^2 \quad (3)$$

The Nelder-Mead optimization method will alter the value of the components of W based on the two previous values of σ and whether the last value was higher or lower than the previous iteration. If σ was lower, then the iteration is in the correct direction and the next estimate for W will be in the same direction as the last. Over a number of iterations, the optimization scheme converges onto the true values for the components of W until σ reaches a minimum. After this point, no more calculations are made and the wind vector is found.

For the current helical ascent flight plan, 151 GPS-derived values are used in this optimization scheme. Since the GPS measurements come in at 5Hz, this corresponds to around 30 seconds of flight time. This value is used based on experimentation with different values. Using a lower number of points, n , the wind data becomes much too noisy and invaluable. When using more instantaneous velocity values, small changes in the wind vector with height are not resolved.

Currently, the u and v components of the wind have been calculated accurately for SMARTSonde flights with this method. However, due to the fact that the w component is typically very small, it has not been able to be resolved using this method

since it is typically smaller than the noise in the GPS data.

3.3. *Paparazzi Wind Algorithm*

As mentioned previously in Section 3.2, the Paparazzi autopilot program yields the wind speed and direction at flight level. These values are given by the program every ten seconds. However, the algorithm that Paparazzi uses is not well-documented. Generally speaking, it uses some form of the Nelder-Mead optimization method. However, the number of points, n , that are used in the optimization is unknown.

3.4. *Use of Airspeed Measurements in Algorithms*

Since the SMARTSonde is equipped with a pitot tube for measuring the plane's airspeed, it is possible to use the airspeed in order to improve algorithm performance. In addition, the plane's airspeed and throttle could be allowed to vary during flight if the airspeed measurements were an input into the algorithms. For the best curve fitting and Nelder-Mead algorithms, the airspeed measurements would be used for a at each individual point in the equations instead of using an average airspeed.

If the airspeed vector, instead of only the airspeed magnitude were known, there are additional methods that could be used to calculate the wind speed. However, in order to get the airspeed vector, the pitch, roll, and yaw angles all need to be known. The pitch and roll are currently measured from differential infrared sensors. An inertial measurement unit (IMU) would be needed to provide accurate values for yaw. Due to the cost, an IMU has not yet been used in the SMARTSonde.

Airspeed measurements are not used in the wind algorithms at this time, because the data from the pitot tube is not being relayed to the GCS in realtime. Instead, it is saved onto a data logger which has no time stamp of when the data was recorded. Because of this issue, it is extremely difficult to get the airspeed measurements to coincide with the GPS measurements coming into the GCS. Work is being done to have the airspeed measurements relayed to the GCS through the autopilot program. Once that is accomplished, the airspeed data can be used with the GPS measurements to improve the algorithm performance.

4. ALGORITHM PERFORMANCE AND COMPARISONS

The three wind algorithms have been used to make wind measurements for profiles of the PBL for comparison against a nearby rawinsonde and a mesonet station. The rawinsonde used for comparison wind profiles were from the Norman, OK site. As previously stated, this is less than a kilometer away from where the SMARTSonde flights for this study took place, making it a good site for comparison. In addition, the ground near the rawinsonde site is relatively flat, mitigating any terrain effects. The rawinsonde observations should closely match the SMARTSonde's observations.

The National Weather Center (NWC) mesonet observation was used for additional comparison. The mesonet station is maintained by the Oklahoma Climatological Survey (OCS) just to the east of the National Weather Center. The data are archived for public use with one minute resolution. The wind speed and direction is observed at 10 m above the ground. The main advantage of using the NWC mesonet observation is that the wind observation at the exact moment as the SMARTSonde's takeoff can be used for comparison. This will show if the wind changes significantly between the rawinsonde observation time and the time of the SMARTSonde flight.

The SMARTSonde has made a total of 34 helical ascent flights. These have taken place between February 2010 through early January 2011. Most of these flights occurred during periods when the synoptic weather conditions were fairly weak. Thus, conditions during the balloon launch should be similar to conditions during SMARTSonde flights, as long as the two times are within several hours of each other. Hence, it is reasonable to make direct comparisons.

Fig. 7 shows four typical wind profile plots calculated using the various methods. They are compared with the Norman sounding and mesonet station. Note that the rawinsonde observation is at 00Z for all of the comparison soundings; it is usually launched at 23:00Z. Going from figure 7a to 7d, the flights become progressively more recent. In addition, the flights started their helical ascent progressively closer to the ground with increasing confidence in the autopilot system. This allows for better observations near to the surface, which is critical during morning and evening transition periods.

Fig. 7a shows an event on February 12, 2010

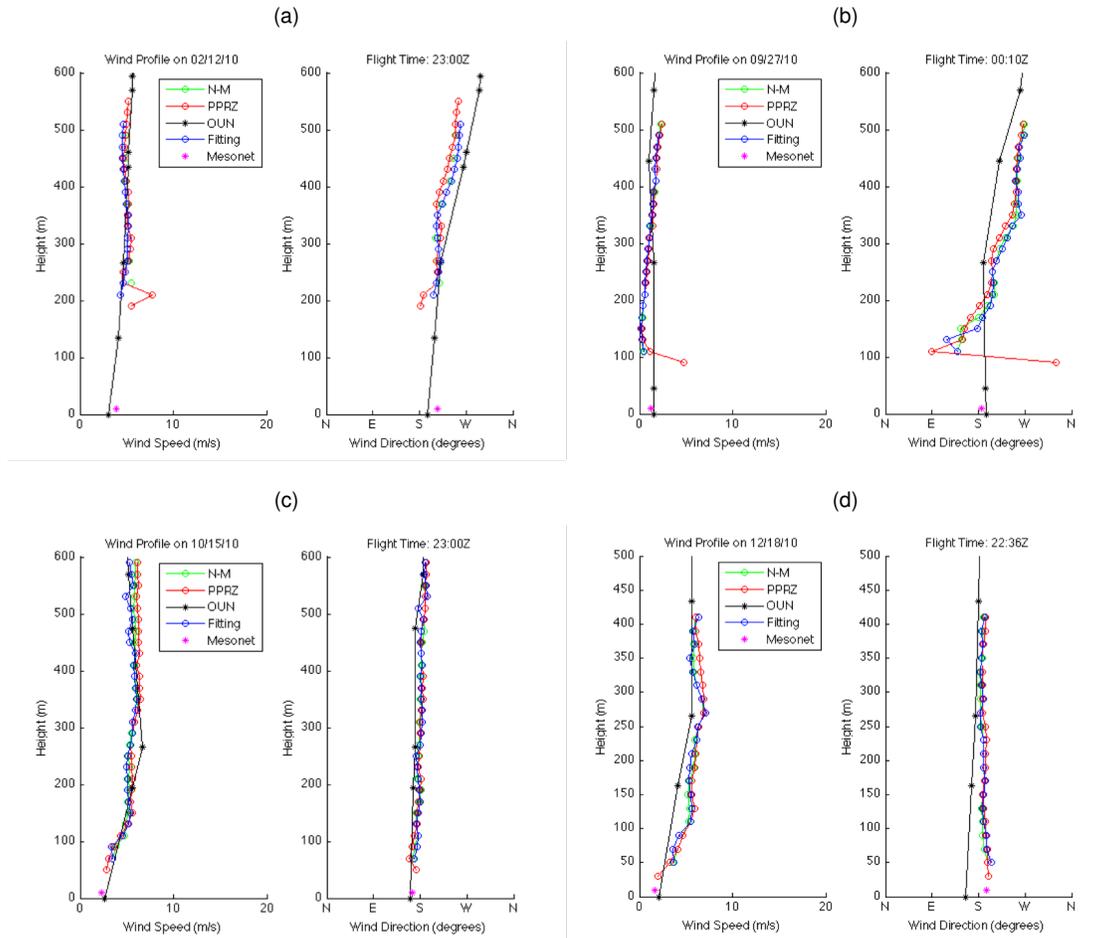


Figure 7: Four examples of typical wind profiles using the different algorithms compared against the mesonet station and the rawinsonde. N-M is the Nelder-Mead method, PPRZ is the paparazzi output, OUN is the Norman rawinsonde, Fitting is the best curve fitting method, Mesonet is the local mesonet observation at 10m AGL.

where there was a noticeable wind direction shift from southerly to westerly between 300-500 meters. This was associated with an elevated mixed layer moving over the cooler air near the surface. All of the SMARTSonde's algorithms agreed with the rawinsonde observation of this wind shift. Another wind shift also occurred in Fig. 7b. This day, the winds were much lighter than the previous day, but both the rawinsonde and the wind algorithms still showed the low-level wind shear. In Fig. 7c, a noticeable feature was the weaker winds that were observed below 100m AGL. This shows that as the plane has started its ascents from a lower altitude, the SMARTSonde's algorithms were able to resolve the weaker winds near the surface. In the final example in Fig. 7d, once again weaker winds are observed near the surface. In addition,

this flight took place ~30 minutes prior to the rawinsonde launch. There is disagreement in the wind direction between the rawinsonde observations and the SMARTSonde's observations below 250m AGL. However, the SMARTSonde's lowest observation, using any of the three algorithms, agrees quite well with the wind direction that the NWC mesonet reported at the takeoff time for the SMARTSonde flight.

The four example plots comparing algorithm output against the sounding, as well as the other 30 comparisons not shown here, generally illustrate good agreement between the rawinsonde observations and the different wind algorithms. In fact, the agreements between the three algorithms themselves is quite remarkable. In addition, no algorithm appeared to perform drastically better

than any other when compared against the rawinsonde. In order to better determine which algorithm worked best, mean errors between each algorithm and the rawinsonde were computed.

Tables 1–3 show the root mean squared errors of the difference between the algorithm outputs and the rawinsonde observations. Based on these numbers, both the Nelder-Mead optimization and the best fit curve methods provide wind measurements that are closer to the wind data provided by the rawinsonde in every category. The best fit curve method gives a more accurate value of the wind speed, while the Nelder-Mead optimization provides a better value for the wind direction. Both of those methods approximated the wind vector to within $\sim 1.7 \text{ ms}^{-1}$ of the rawinsonde's vector. Since the standard error for the rawinsonde system for u and v is 1 ms^{-1} , the errors for the wind vector are slightly larger than the error for the rawinsonde system. Considering that the algorithms' outputs have errors themselves, and that the wind changes between the time of the launch and the SMARTSonde flights, this is not surprising. It still shows that the Nelder-Mead optimization and best fit curve methods agree quite well with the rawinsonde observations.

The individual u and v components of the wind agree quite well for all the algorithms when evaluated against the rawinsonde, as shown in Fig. 8. Generally, the components lie around the one-to-one line. There are some offsets, especially at higher values. This is mainly due to one or two flights that have many data points that are slightly different than the sounding. Also, there are some indications that the algorithms have a slight bias in wind direction depending on whether the SMARTSonde is circling to the left or right. This could belie some of the offset from the one-to-one line. Overall, the points for each component for all the algorithms are in general agreement with the rawinsonde observations.

5. FUTURE WORK

Thus far, the wind algorithms have only been compared with local rawinsonde observations. In the near future, a sodar will be set up near the CORCS field which will provide an additional data source for comparison. This will provide a more direct comparison than to the rawinsonde, since it will be operating at the same time the SMARTSonde is flying. In addition, it will be located at the airfield used by the SMARTSonde. There may

also be some comparisons in the future with a Doppler lidar.

The wind algorithms will also be used in conjunction with the thermodynamic data from the SMARTSonde in order to compute stability parameters, such as gradient Richardson number profiles. This combination of the thermodynamic and dynamic data will be very useful in future boundary layer studies. For example, the early morning transition and the low level jet dissipation could be more thoroughly studied with this instrumentation.

6. CONCLUSIONS

Overall, the three algorithms that were created for the SMARTSonde provided reasonably accurate results when compared against local rawinsonde observations. Some additional verification of the algorithms and more comparisons with other instrument systems are needed. It appears as though both the Nelder-Mead optimization and the best fit curve methods performed better than the paparazzi autopilot output. At this point, either could be used in order to accurately obtain a wind profile.

Since each algorithm performs somewhat differently, the choice of algorithm could be dictated in part by the flight plan selected for a particular experiment. For the helical ascent flight plan, the best fit curve method would work best since it obtains independent observations quicker than the Nelder-Mead method. The SMARTSonde can complete a circle in less than thirty seconds, which is what the Nelder-Mead method essentially needs for averaging. On the other hand, if the flight plan includes more straight paths, the Nelder-Mead method could give faster updates since it could take a long time for the UAS to turn a full 360 degrees. Either way, both the Nelder-Mead and best curve fitting algorithms give valid wind profiles.

7. REFERENCES

- Emeis, S., M. Harris, and R. M. Banta, 2007: Boundary-layer anemometry by optical remote sensing for wind energy applications. *Meteorol. Z.*, **16** (4), 337–347.
- Endlich, R., F. Ludwig, C. Bhumralkar, and M. Es-toque, 1982: A diagnostic model for estimating winds at potential sites for wind turbines. *J. Appl. Meteorol.*, **21** (10), 1441–1454.

Table 1: Mean Errors for Best Fit Curve vs. Rawinsonde

Time from launch of radiosonde (number of flights)	Wind Speed (m/s)	Wind Direction (degrees)	u-comp (m/s)	v-comp (m/s)	V-vector (m/s)
< 20 minutes (7)	1.13	13.71	1.24	1.12	1.67
< 1 hour (22)	1.14	16.08	1.18	1.29	1.75
All flights (34)	1.30	15.07	1.20	1.30	1.77

Table 2: Mean Errors for Nelder-Mead Optimization vs. Rawinsonde

Time from launch of radiosonde (number of flights)	Wind Speed (m/s)	Wind Direction (degrees)	u-comp (m/s)	v-comp (m/s)	V-vector (m/s)
< 20 minutes (7)	1.38	12.44	1.16	1.34	1.77
< 1 hour (22)	1.24	14.83	1.07	1.33	1.71
All flights (34)	1.37	13.92	1.10	1.33	1.72

Table 3: Mean Errors for Paparazzi Output vs. Rawinsonde

Time from launch of radiosonde (number of flights)	Wind Speed (m/s)	Wind Direction (degrees)	u-comp (m/s)	v-comp (m/s)	V-vector (m/s)
< 20 minutes (7)	1.46	15.08	1.45	1.52	2.10
< 1 hour (22)	1.35	15.9	1.31	1.43	1.94
All flights (34)	1.42	14.97	1.29	1.38	1.89

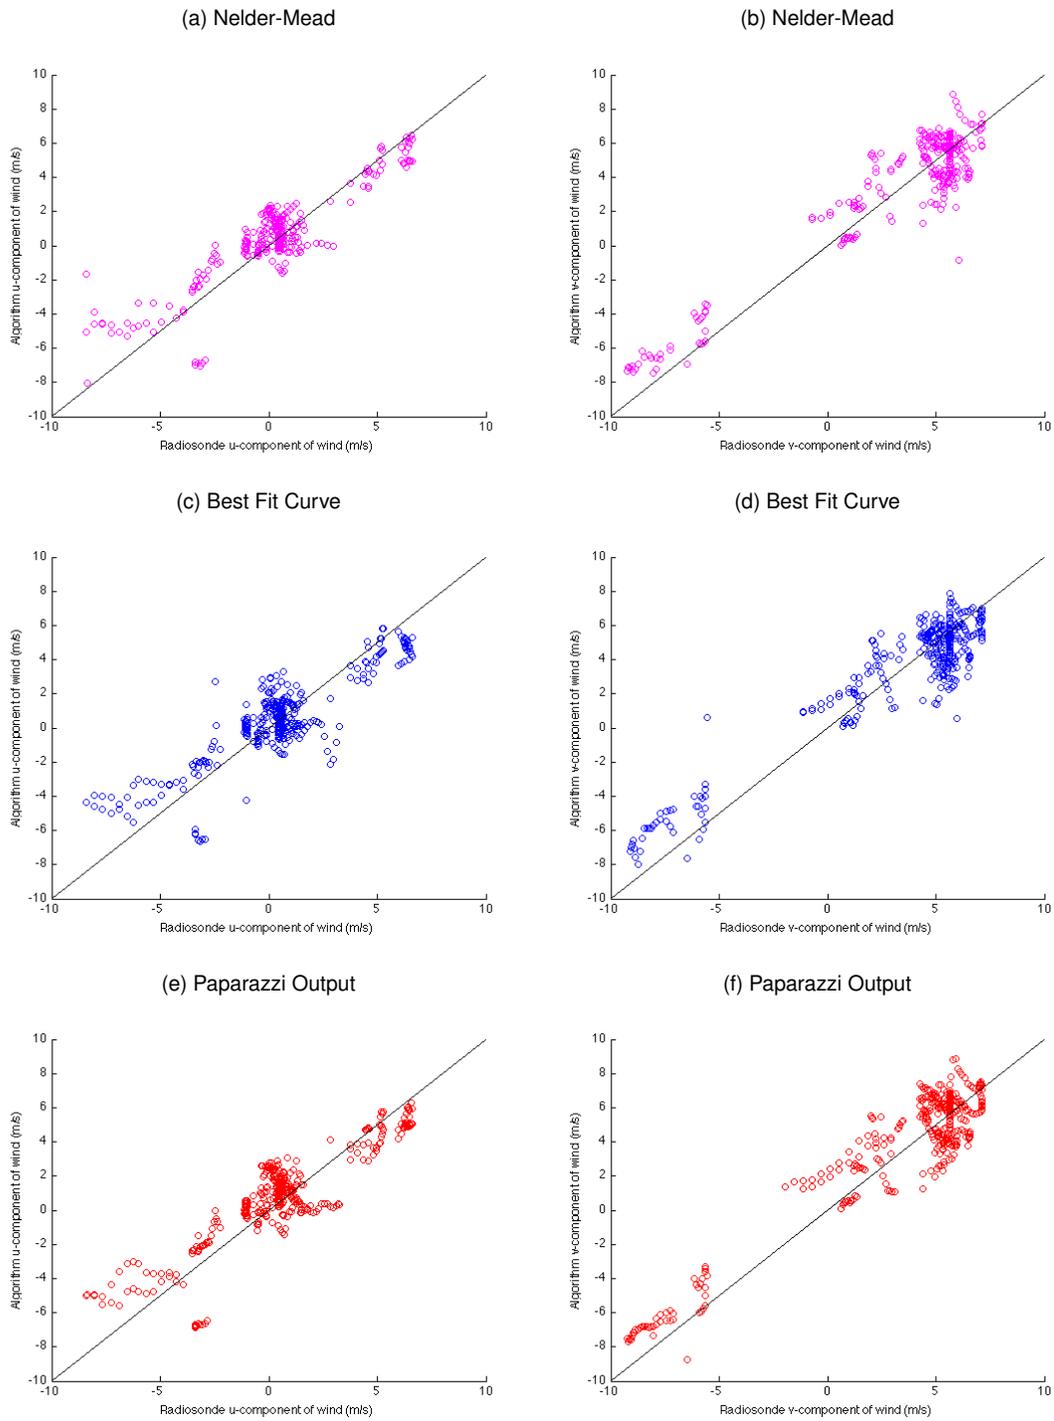


Figure 8: Scatter plots comparing the u and v components of the wind from the algorithms to the radiosonde.

- Holland, G., et al., 2001: The aerosonde robotic aircraft: A new paradigm for environmental observations. *Bull. Amer. Meteor. Soc.*, **82 (5)**, 889–901.
- Holland, G. J., T. McGeer, and H. Youngren, 1992: Autonomous aerosondes for economical atmospheric soundings anywhere on the globe. *Bull. Amer. Meteor. Soc.*, **73 (12)**, 1987–1998.
- Kondragunta, S., et al., 2008: Air quality forecast verification using satellite data. *J. Appl. Meteor. Climatol.*, **47 (2)**, 425–442.
- Mhuqing, M., C. Hongbin, W. Gai, P. Yi, and L. Qiang, 2004: A miniature robotic plane meteorological sounding system. *Advances in Atmospheric Sciences*, **21 (6)**, 890–896.
- Reuder, J., P. Brisset, M. Jonassen, M. Müller, and S. Mayer, 2009: The Small Unmanned Meteorological Observer SUMO: A new tool for atmospheric boundary layer research. *Meteorol. Z.*, **18 (2)**, 141–147.
- Seaman, N. L. and S. A. Michelson, 2000: Mesoscale meteorological structure of a high-ozone episode during the 1995 NARSTO-Northeast study. *J. Appl. Meteorol.*, **39 (3)**, 384–398.
- Spiess, T., J. Bange, M. Buschmann, and P. Vörsmann, 2007: First application of the meteorological Mini-UAV 'M²AV'. *Meteorol. Z.*, **16 (2)**, 159–169.
- van den Kroonenberg, A., T. Martin, M. Buschmann, J. Bange, and P. Vörsmann, 2008: Measuring the wind vector using the autonomous mini aerial vehicle M²AV. *J. Atmos. Ocean. Tech.*, **25 (11)**, 1969–1982.