Towards airborne seismic data acquisition

Thomas Rapstine & Paul Sava

Center for Wave Phenomena and Dept. of Geophysics, Colorado School of Mines, Golden CO 80401
email trapstin@mines.edu

ABSTRACT

The ability to acquire ground strong motion data from a moving airborne platform without contacting the ground leads to exciting new exploration opportunities. Existing stationary laser and camera technology may be used to achieve this goal, if the platform motion can be reliably measured and compensated for. Sensors measure a combination of platform and ground motion, and signal separation is required to recover the ground motion. Machine learning techniques may be used for signal separation, but rely on a large amount of training data for successful and robust separation. Training data for platform motion, such as a hovering unmanned aerial vehicle (UAV), are not available. Therefore, we construct a strategy for generating synthetic platform motion training data using a path defined by cubic splines. The cubic spline path gives access to an analytic form of the position, velocity, and acceleration of the platform. Since the noise in simulated platform motion observations must reflect reality, we perform platform sensor noise analysis using experimental data. We find the expected platform sensor noise to be correlated, non-white, non-Gaussian and on the same level as the expected ground motion signals. The noise is sparsely represented in the wavelet domain, which may give a basis for signal separation.

Key words: land acquisition, UAV, signal separation, machine learning

1 INTRODUCTION

The ability to gather seismic data from an airborne platform would open new opportunities in seismology, and potentially in seismic exploration. Many current sensor technologies such as laser doppler vibrometer (LDV) systems, stereo camera systems, and radar systems allow motion to be measured without contact. These systems could potentially be mounted on an airborne platform to acquire seismic data. However, a limiting use-case of the named systems is the assumption of a stationary measurement platform. If these systems are mounted on an airborne platform, then the observed signal reflects the combined motion of the ground and of the platform itself. However, if the platform motion is known exactly, then the motion of the ground may be recovered using stereo cameras and a coordinate transformation (Rapstine et al., 2017). In practice, the the platform motion is measured within a level of uncertainty, leading to an uncertain estimate of ground motion. A key problem is, therefore, separating the platform signal from the ground signal using a noisy estimate of the platform motion. Before separating these signals, it is necessary to understand platform motion and the associated uncertainties.

The position and orientation, or pose, of an airborne platform, such as an unmanned aerial vehicle (UAV), is commonly found by fusing measurements from internal and external sensors. The fusion process and sensor noise measurements introduce uncertainty into the resulting pose. This process is routinely undertaken in airborne navigation, where the position and orientation of a platform is estimated using information from multiple sensors. A common external source of information comes from Global Positioning Systems (GPS), which provide the position of an object within a range of uncertainty. There are many complex factors which influence the GPS data noise including environment, sensor, and object dynamics (Groves, 2013; Misra and Enge, 2006). Internal navigation systems (INS) may also provide the pose of a platform at high sample rate but contain their own uncertainty, which accumulates over time (Titterton and Weston, 2004; Flenniken et al., 2005). INS sensors may include accelerometers and gyroscopes which give estimates of the acceleration and angular velocity of an object. The orientation information from a gyroscope plus the acceleration information may be integrated to determine the velocity and position of an object over time. The integration causes constant accelerometer bias error to manifest as linear error in velocity and quadratic error in time due to numerical integration (Titterton and Weston, 2004). In addi-

tion, many factors contribute to INS noise which have complex non-linear relationships (Flenniken et al., 2005). The propagating INS error is generally suppressed by incorporating information from external GPS. In short, the sources and character of the noise in INS and GPS, and therefore in the fused UAV position and orientation, are complex. However, by combining the information from GPS and INS sensors, the effects of their respective weaknesses are minimized and their strengths combined.

A strength of GPS is that it directly measures position, however at a low sample rate. Therefore the bias position error from GPS is a constant, as opposed to effects of accelerometer bias error from INS, which quadratically increases with time. A major strength of INS measurements is their high sample rate compared to GPS, and their availability in remote or isolated regions where GPS is unavailable. A common way to combine data from GPS and INS sensors is to use one of the many flavors of Kalman filters (Kalman, 1960). Using Kalman filters to combine INS and GPS data requires knowledge about each sensor's error distribution, which may vary greatly depending on the sensor and environment. The performance of a Kalman filter depends on choosing the appropriate error model, which in practice may be hard to determine. An alternative for combining INS and GPS data is to approach their combination as a machine learning problem.

A variety of neural networks strategies have been deployed to "learn" how to navigate a vehicle with INS data alone when GPS is unavailable (Chiang et al., 2003; Noureldin et al., 2003; Hiliuta et al., 2004; Sharaf et al., 2005; Noureldin et al., 2007, 2011; Xu et al., 2010). In these works, a neural network is trained to predict the error in the INS position estimate. The estimated positions from INS are used as input to a neural network; the output is a prediction of INS-derived position error. GPS data are used as the true position for on-line training of the neural network. The on-line training is occurrs while a vehicle is in the field acquiring both GPS and INS data. When the GPS data are not available, training stops, the network predicts the INS position error, and the INS derived position is corrected using the predicted error. A variety of network inputs have been explored by others, for example inputting INS position estimates at varying times or including accelerometer, gyroscope, and attitude information (Noureldin et al., 2011). In addition, different types of machine learning algorithms were investigated, including artificial neural networks, neuron-wavelet networks (Noureldin et al., 2003), radial basis function neural networks (Semeniuk and Noureldin, 2006), adaptive fuzzy neuron-fuzzy inference systems Abdel-Hamid et al. (2007), and support vector machines Xu et al. (2010). The success of machine learning strategies for fusing motion measurements leads us to propose using similar strategies for separating motion measurements.

Motion signals represent a sequence of information; and thus they are naturally ordered. Common machine learning strategies for fusing motion signals do not explicitly leverage the order during the learning process. The order of information in a time series is therefore an important characteristic of the data. A variant of a neural network that learns based on the inputs and their ordering is the long short-term memory (LSTM) network (Hochreiter and Schmidhuber, 1997). LSTM networks are a practical extension of recurrent neural networks (RNNs), and have been applied successfully to character and word prediction (Sutskever et al., 2011; Mikolov et al., 2013). LSTM networks are specifically designed to avoid practical problems that occur when training RNNs. These problems are rooted in the backpropagation algorithm used to update neurons during training. The backpropagation algorithm is a minimization process, such as gradient descent. The minimization process relies on partial derivatives that represent the change in a neuron with respect to each neuron it is connected to. The partial derivatives are multiplied together according to the chain rule. In RNNs, there are many connected neurons, and therefore many partial derivatives combine into an exceedingly small or large number. This problem is referred to as the "vanishing" or "exploding gradient" problem, and is discussed in detail in Bengio et al. (1994). LSTM networks, as their name indicates, have the concept of memory embedded in the neurons which comprise the network. By adding the concept of memory to the network, the backpropagation process stabilizes. The concept of memory makes LSTM networks well suited for processing time series data, where the order of information in the series carries significant meaning.

In this report, we present steps towards constructing a supervised learning strategy for separating UAV and ground motion signals. In order to design inputs to a learning algorithm, a major step is characterizing the signals we wish to separate, along with the noise present in the signals. The separation algorithm must be robust with respect to the observed noise signals. We analyze the noise from a high-end INS sensor in detail to define an appropriate signal separation strategy. We show that the INS noise is correlated, non-Gaussian, and not white or band-limited in the frequency domain. However, the noise may be represented sparsely in the wavelet domain, which suggests that it may be separated from noisy observations of ground motion taken from a moving airborne platform.

2 SEPARATION TARGETS

Our goal is to recover the ground motion from airborne observations, which are contaminated by the effects of a moving and rotating UAV. The UAV is assumed to be hovering above the ground as the observations are acquired. The model we choose for representing observations of the UAV and ground at a specific time is

$$\widehat{\mathbf{u}} = \mathbf{u} + \varepsilon_{\mathbf{u}} \tag{1}$$

$$\widehat{\mathbf{r}} = \mathbf{r} + \varepsilon_{\mathbf{r}} \tag{2}$$

$$\widehat{{}^{t}\mathbf{g}} = {}^{t}\mathbf{g} + \varepsilon_{\mathbf{g}}. \tag{3}$$

The notation $\widehat{(\cdot)}$ denotes a noisy measurement. We denote by $\mathbf{u}$ and $\mathbf{r}$ the true position and Euler angle orientation, respectively, of the UAV in a reference coordinate system. ${}^{t}\mathbf{g}$ denotes the position of a point on the ground in the UAV's current coordinate system at time t. Additive observation noise is denoted

by $\varepsilon_{\mathbf{u}}, \varepsilon_{\mathbf{r}}, \varepsilon_{\mathbf{g}}$ for the UAV position, orientation, and ground position, respectively. Fusing INS and GPS measurements give the observation $\widehat{\mathbf{u}}$ and $\widehat{\mathbf{r}}$ while a separate sensor, such as a stereo camera or LDV, give the observation ${}^{t}\widehat{\mathbf{g}}$. These observations are linked through a coordinate transformation between the reference coordinate frame and the UAV's current coordinate system.

We aim to remove, or separate, the effects of the coordinate system transformation in order to recover the position of the ground $\mathbf{g}$ in the reference coordinate system. If no noise is present, the position of the ground in the reference coordinate system is recovered by

$$\mathbf{g} = \mathbf{R}\left({}^{t}\mathbf{g}\right) + \mathbf{u}, \tag{4}$$

where $\mathbf{R}$ represents a rotation operator formed using the Euler angles contained in $\mathbf{r}$. We note that there are various ways to represent rotation using a matrix, whether it is by rotation about a certain axis, using Euler angles, or by using quaternions (see Diebel (2006) for various representations). For simplicity, we refer to $\mathbf{R}$ as a rotation matrix. Equation 4 exactly recovers the position of the ground, but represents an unrealistic scenario where no noise is present in our measurements. By solving for ${}^{t}\mathbf{g}$, we obtain an expression of measurement of the ground motion from the UAV as a function of the true position of the ground $\mathbf{g}$

$${}^{t}\mathbf{g} = \mathbf{R}^{T}\left(\mathbf{g} - \mathbf{u}\right). \tag{5}$$

Rotation may be represented by orthogonal matrices, so we use $\mathbf{R}^{-1} = \mathbf{R}^{T}$. Equation 5 shows that, if no noise is present ($\varepsilon_u = \varepsilon_r = \varepsilon_g = 0$), we observe a rotated version of the difference between the position of the ground and UAV. Adding the noise term $\varepsilon_{\mathbf{g}}$, equation 5 simply becomes

$${}^{t}\widehat{\mathbf{g}} = \mathbf{R}^{T}\left(\mathbf{g} - \mathbf{u}\right) + \varepsilon_{\mathbf{g}}. \tag{6}$$

Equation 5 along with equations 1 and 2 represent our observations in terms of the signals we wish to separate. Using a machine learning approach for separation requires many realistic training examples of these noisy observations.

3 GENERATING TRAINING DATA

In preparation for supervised learning, we anticipate the need for a large amount of training data in the form of UAV and ground motion described in the previous section. The training data must consist of the true motion and the perceived noisy observations of the UAV and ground motion.

A large amount of real ground motion data is available online through the Center for Engineering Strong Motion Data (Center for Engineering Strong Motion Data, 2016). Therefore, 3D displacement signals from this open database may be used as training data for the ground motion ($\mathbf{g}$). An example of such a signal is in Figure 1. Unfortunately, a similar data set does not exist for UAV motion, so we are required to simulate a hovering UAV.

We first describe a strategy for simulating hovering UAV motion. Then we analyze an IMU data time series to determine the character of noise to simulate the UAV position and orientation observations.

3.1 Simulating a hovering UAV

We simulate a hovering UAV that is translating and rotating as a function of time. The position of the hovering UAV can be represented by a path in space. Along the path of the UAV, ideally we simulate accelerometer and gyroscope data in the form of acceleration and angular velocity. We therefore simulate the first and second time derivatives of position to have access to velocity and acceleration. We achieve this requirement by constructing a path that is analytically defined by a cubic spline in time, which has two continuous derivatives by definition. Control points points for splining, are defined using 3D recursive walk scheme within a bounding ellipsoid.

The control points are generated by first specifying the number of desired control points n in the 3D path sequence. Two control points are randomly selected inside a specified ellipsoid. The midpoint between the two points is computed. A random perturbation is then added to the midpoint, which gives a third control point. We use random perturbations drawn from a Gaussian with zero mean and a scaled variance denoted $\mathcal{N}(0, ns)$, where n is the length of the sequence at a particular step and s is a scale factor. The process of computing midpoint and perturbing is recursively applied between the original two control points and the third perturbed midpoint. In practice, this recursive algorithm is applied in-place to a preallocated array. The first and last entries of the array represent the first two control points in the ellipsoid. The control points for the procedurally defined random path within an ellipsoid are generated using pseudocode in Algorithm 1.

Using Algorithm 1, we obtain the samples representing the coordinates of the UAV in the coordinate system of the UAV at time zero, denoted $\mathbf{u}$. We obtain samples representing the orientation of the UAV, denoted $\mathbf{r}$, using a similar algorithm except that we ensure that the angles lie between upper and lower bounds. In this way, we can procedurally define the position and orientation, along with their first two time derivatives, of the UAV and generate as many training sets as we wish. To simulate observations, realistic noise must be added to the position and orientation data. A statistical and time-frequency analysis of real position and orientation noise follows, which gives insight into the characteristics of the noise we must simulate.

3.2 Observation noise in UAV position and orientation

The statistical and time-frequency behavior of UAV position and orientation observation noise influences our choice of a signal separation algorithm. Here, we analyze real noise from a high-end Honeywell HG4930 IMU sensor.

The noise test is conducted in an isolated environment over about one hour at a sampling rate of 100 Hz while the sensor is stationary (Mazal, 2018). The experiment results in

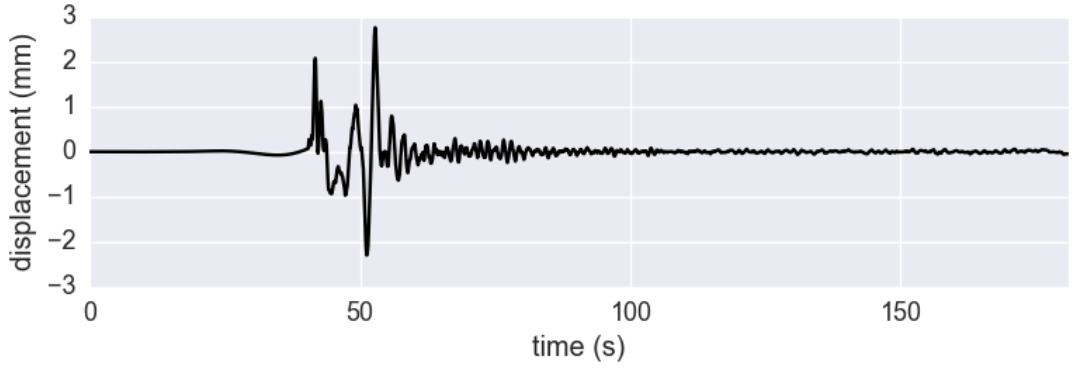


Figure 1. An example of a strong motion signal representing the position of the ground ($\mathbf{g}$) from USGS station OK034 outside of Cushing Oklahoma in fall 2016.

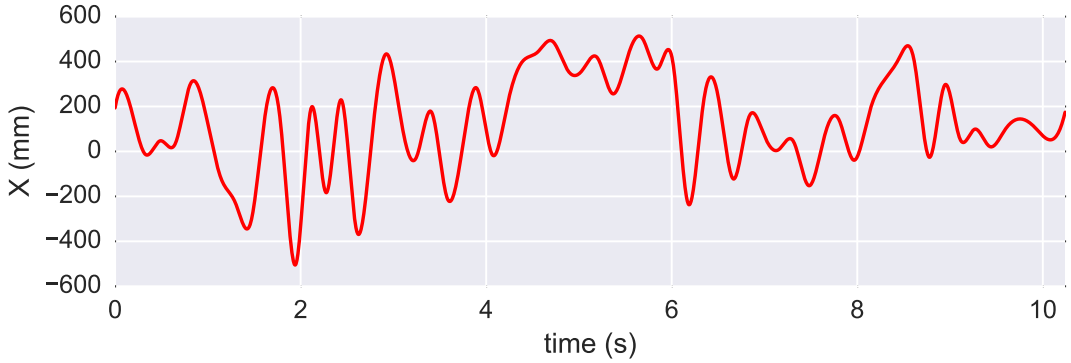
$n = length(\mathbf{x})$;
if $n \leq 2$ **then**
  return;
else
  $i = floor(n/2)$;
  $mx = (\mathbf{x}[0] + \mathbf{x}[-1])/2$;
  $my = (\mathbf{y}[0] + \mathbf{y}[-1])/2$;
  $mz = (\mathbf{z}[0] + \mathbf{z}[-1])/2$;
  $\mathbf{x}[i] = mx + \mathcal{N}(0, ns)$;
  $\mathbf{y}[i] = my + \mathcal{N}(0, ns)$;
  $\mathbf{z}[i] = mz + \mathcal{N}(0, ns)$;
  while *not in ellipsoid* **do**
    $\mathbf{x}[i] = mx + \mathcal{N}(0, ns)$;
    $\mathbf{y}[i] = my + \mathcal{N}(0, ns)$;
    $\mathbf{z}[i] = mz + \mathcal{N}(0, ns)$;
  end
  $get_control(\mathbf{x}[: i+1], \mathbf{y}[: i+1], \mathbf{z}[: i+1], s)$;
  $get_control(\mathbf{x}[i :], \mathbf{y}[i :], \mathbf{z}[i :], s)$;
end

Algorithm 1: The recursive function $get_control(\mathbf{x}, \mathbf{y}, \mathbf{z}, s)$ for generating random control points within an ellipsoid for splining.

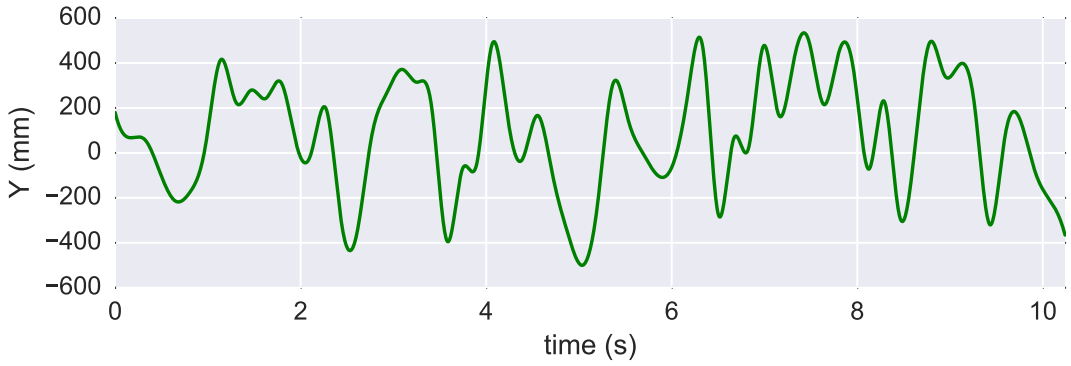
a large number of samples (around 360,000) for each component of position and orientation. Accelerometer, gyroscope, and GPS data are professionally processed by Honeywell to derive position and orientation (Mazal, 2018). The position is represented in earth-centered, earth-fixed (ECEF) coordinates. The orientation is represented in a local level frame, which may be used to construct the rotation matrix $\mathbf{R}$. These data represent the uncertainty present in post-processed position and orientation of a UAV. The values of the first samples in position and orientation are subtracted from the time series. This subtraction shifts all noise values to a local coordinate system for analyzing the noise character change over one hour. Figures 3(a)-3(c) show the position noise and Figures 4(a) -4(c) show the orientation noise. The noise traces show oscillatory features, for example between 2000 and 3000 seconds in y-position in Figure 3(b). In addition, the noise appears to drift in each time series. To quantify the character of these features, we analyze the statistics and time-frequency behavior of the noise time series.

Summary statistics and empirical PDFs give concise insight on the statistical character of the UAV position and orientation noise. A variety of common summary statistics for position are in Table 1, with corresponding empirical PDFs shown in Figures 5 and 6. The mean values show that the noise bias, after one hour of observation, can be between 4-8 mm in position. The standard deviation of position estimates is less than 6 mm. The empirical PDFs in Figure 5 show a Gaussian character for the y-position. The x and z-positions appear bimodal; a consequence of the large values in the x-position time series after 3000. The interaction between the x, y, and z-positions may be explored using the empirical joint PDFs in Figure 7. The trends in the off-diagonal cross-plots show correlation between the time series, which is verified by the correlation coefficients in Table 3. The correlation coefficients are visualized for easier interpretation in Figure 9. For example, the cross-plot between the y and z positions shows a modest negative linear trend in the data values implying a negative linear correlation. Correlation implies the data are not statistically independent, which is a common assumption made in signal denoising and separation. Based on these observations we see the position noise is correlated and can potentially be non-Gaussian.

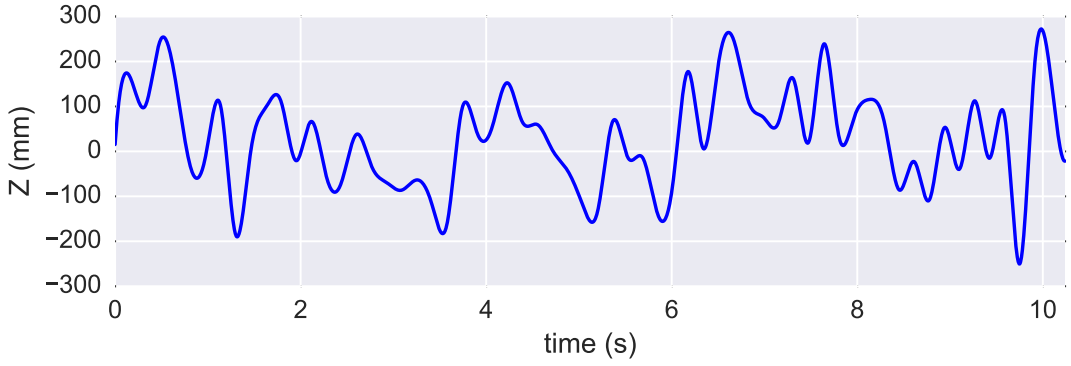
The coordinate transformations are defined by a translation in position and rotation in orientation, so a similar noise analysis we perform on the post-processed orientation data. Orientation is expressed as an angle of rotation about a particular axis. Rotation about the x-axis, y-axis, and z-axis are referred to as roll, pitch, and yaw, and denoted as ϕ, θ, and ψ, respectively. A summary of orientation statistics is in Table 2. The bias in orientation noise is below a millidegree, with the bias in ϕ being the lowest. The standard deviation of the orientation noise is less than 0.4 millidegrees. As with position noise, we observe modest correlation between the time series



(a)



(b)

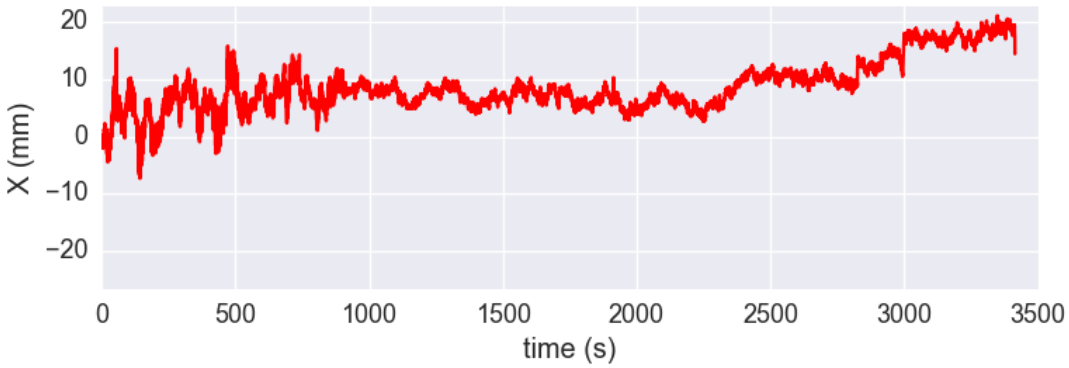


(c)

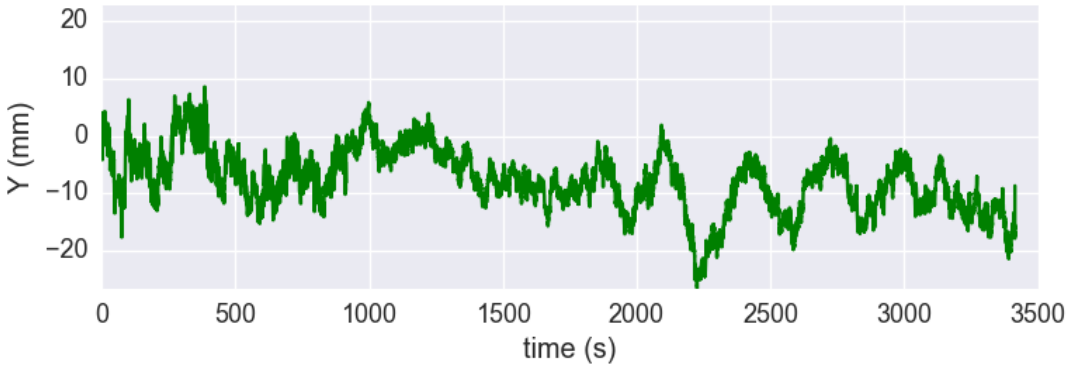
Figure 2. Simulated UAV position in an ellipsoid as a function on time using Algorithm 1 for (a) x, (b) y, and (c) z position. Representing the position using cubic splines gives an analytic form for the velocity and acceleration of the UAV during the simulation.

in Figure 8. The orientation PDFs do not appear to be Gaussian. For reference, the correlation is quantified by the orientation correlation coefficients in Table 3. We observe correlation between position and orientation.

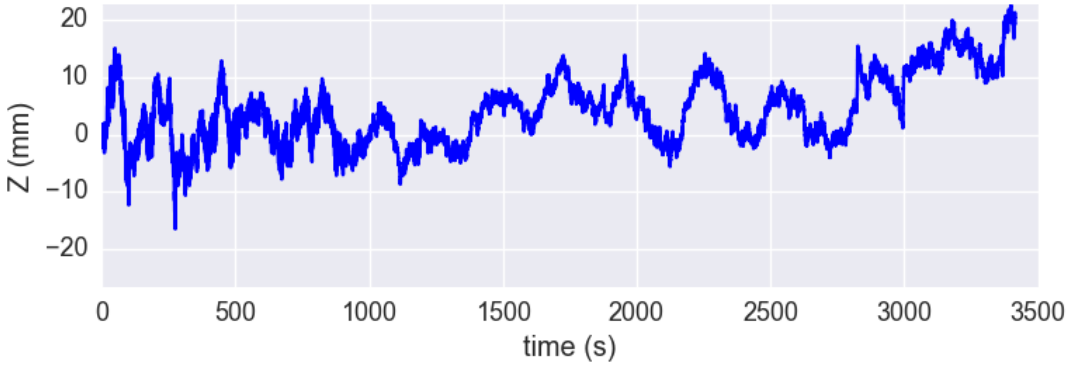
Aside from statistics, understanding the time-frequency behavior of the noise reveals periodic features in the noise as they change with time. The time-frequency behavior for position is shown in Figure 10 in the form of a short time Fourier Transform (STFT). The STFT allows frequency content to be represented as a function of time. Each plot in figure 10 shows



(a)



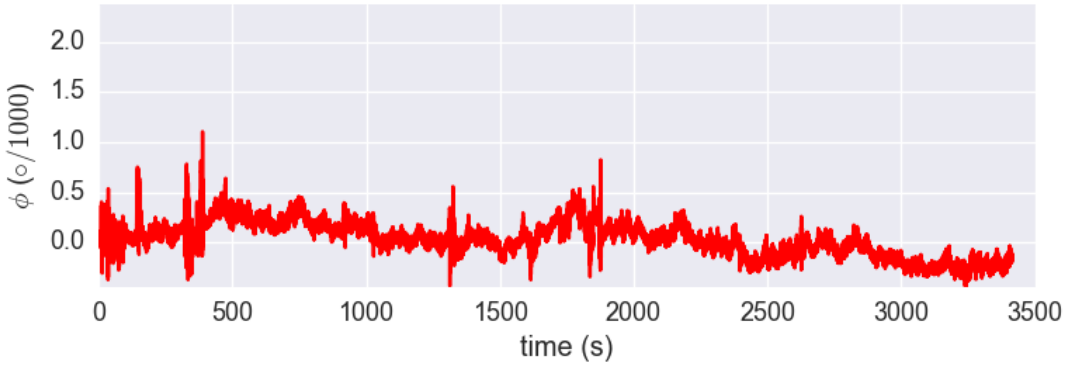
(b)



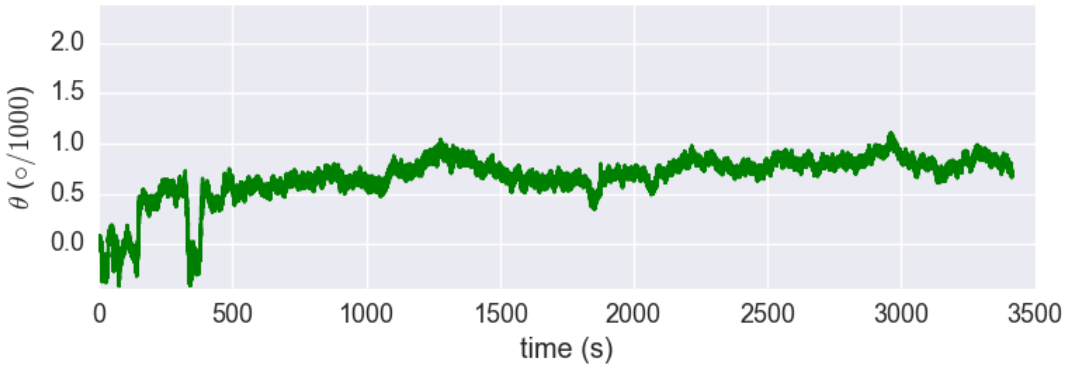
(c)

Figure 3. Position observation noise in millimeters for the (a) x, (b) y, and (c) z positions. The position noise amplitudes are greater than the ground motion signal amplitudes in figure 1.

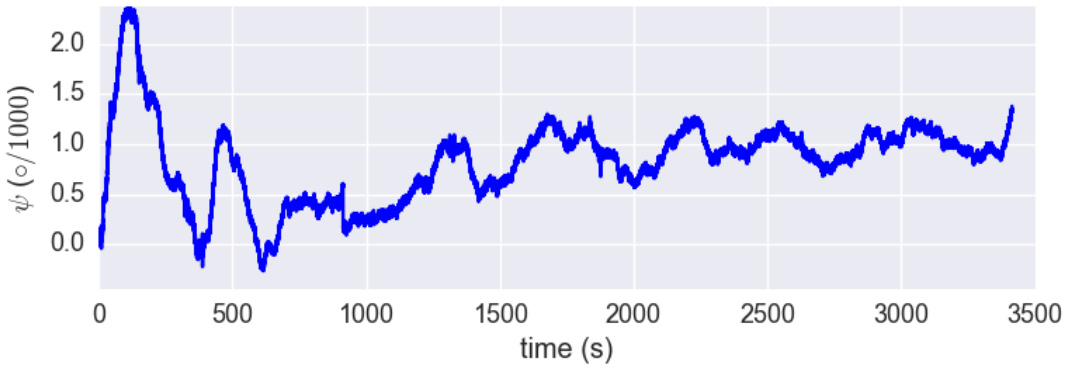
that there is a distinct change in the frequency content of the position noise at 900 seconds. The position noise is more white before 900 seconds and less white afterwards. This shows that the time-frequency behavior of the noise is not time-invariant, which is a common assumption made when denoising or separating signals. We note that the position is derived from acceleration, so the effects of a double integration in time (inverse frequency squared decay) are imprinted on the STFT; we do



(a)



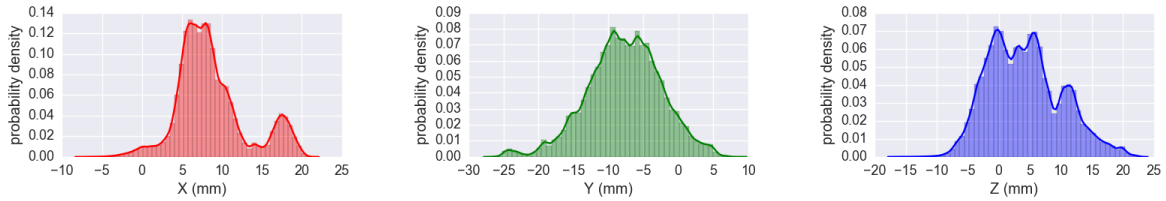
(b)



(c)

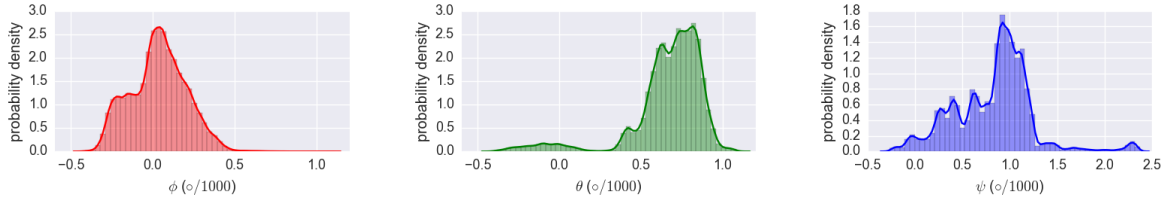
Figure 4. Orientation observation noise in thousandths of a degree for (a) ϕ , (b) θ, and (c) ψ. The noise in orientation observations has structure.

not expect a perfectly white STFT. The repeating bands after 900 seconds occur at 5 Hz intervals, and perhaps represent modal frequencies of the sensor. The change in character at 900 seconds and repeating 5 Hz bands are not viewed in the STFT of orientation in Figure 11. This implies that orientation noise is more time-invariant than the position noise. However, there appears to be a broad band spike in the STFT of ϕ and ψ in Figures 11(a) and 11(c) at 900 seconds. Perhaps the sensor was slightly rotated at 900 seconds, causing a change in posi-



(a) (b) (c)

Figure 5. Empirically derived PDFs for (a) x, (b) y, and (c) z positions. We observe non-Gaussian multi-modal behavior for position components.



(a) (b) (c)

Figure 6. Empirically derived PDFs for (a) ϕ, (b) θ, and (c) ψ. We observe non-Gaussian and multi-modal behavior for orientation components.

	min	max	median	std.	mean
X	-7.305	21.097	7.837	4.382	8.636
Y	-26.589	8.559	-7.647	5.312	-7.734
Z	-16.482	22.633	3.686	5.877	4.173

Table 1. Summary statistics for position in millimeters after shifting.

	min	max	median	std.	mean
ϕ	-0.444	1.106	0.038	0.170	0.034
θ	-0.420	1.111	0.712	0.232	0.663
ψ	-0.265	2.373	0.898	0.408	0.826

Table 2. Summary statistics for orientation in millidegrees

tion noise character. The STFT of the orientation noise is more persistent with time than the positional noise.

Based on observations we draw the following general guidelines regarding expected statistics in position and orientation noise over a one hour time period. The position noise could be Gaussian, but is not necessarily Gaussian, which is a common assumption made in noise models. In contrast, the orientation noise is not likely Gaussian over this time period. There are many samples in the time series, but perhaps an even longer experiment may cause the position and orientation noise to tend towards a Gaussian distribution due to the central limit theorem. There is evidence of correlation between the x, y, and z-position noise, and also correlation between the ϕ, θ, and ψ noise. The time-frequency behavior of the position noise is not time-invariant, and may contain higher energies at certain modal frequencies. In contrast, the orientation noise appears to be time-invariant and does not exhibit modal frequency beating.

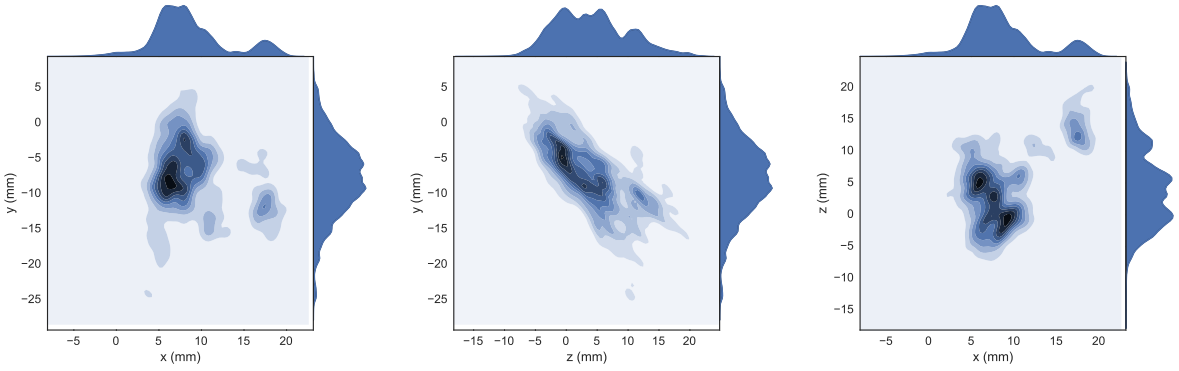
	X	Y	Z	ϕ	θ	ψ
X	1.000	-0.204	0.484	-0.640	0.520	0.110
Y	-0.203	1.000	-0.710	0.270	-0.343	-0.305
Z	0.484	-0.711	1.000	-0.397	0.242	0.314
ϕ	-0.640	0.270	-0.397	1.000	-0.373	-0.306
θ	0.520	-0.343	0.243	-0.372	1.000	-0.023
ψ	0.110	-0.305	0.314	-0.306	-0.023	1.000

Table 3. Correlation coefficients for position and orientation. Noise is correlated between position components and orientation components. The information in this table is easier to interpret visually in Figure 9.

4 SEARCHING FOR SIGNAL SEPARATION CRITERIA

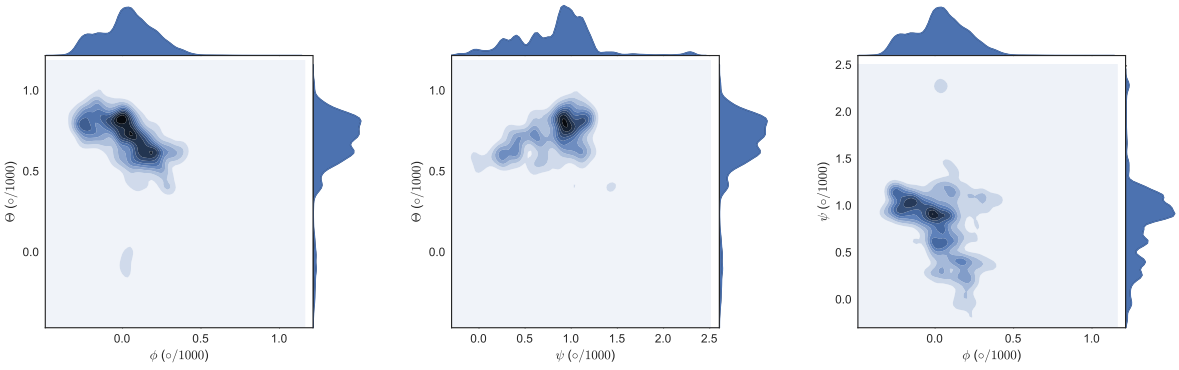
The noise analysis in Section 3.2 shows that the UAV position and orientation observation noise is at best on the same order of magnitude as our seismic signal in Figure 1. The noise is potentially correlated, non-white, and non-Gaussian. To remove this large noise signal from our measurements, we must find a clever way to quantify the differences between the noise and the signal. To do this, we express our noise and signal in the wavelet domain and search for separation criteria.

A zoomed in portion of 10 seconds of the position noise signal in Figure 14 shows the noise is piecewise smooth. Wavelets provide concise representations of piecewise smooth signals (Mallat, 2008), which are not concisely represented using more conventional Fourier bases. A concise representation of piecewise smooth signals is possible because wavelets have compact support in time and frequency, and may form an or-



(a) (b) (c)

Figure 7. Empirically derived joint and marginal PDFs for position. We observe non-Gaussian behavior for all components and modest correlation between components.



(a) (b) (c)

Figure 8. Empirically derived joint and marginal PDFs for orientation. We observe non-Gaussian behavior for all components and modest correlation between components.

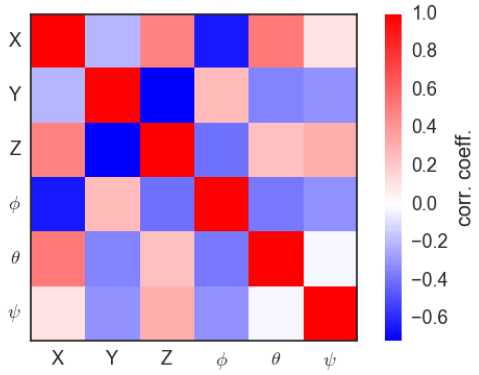
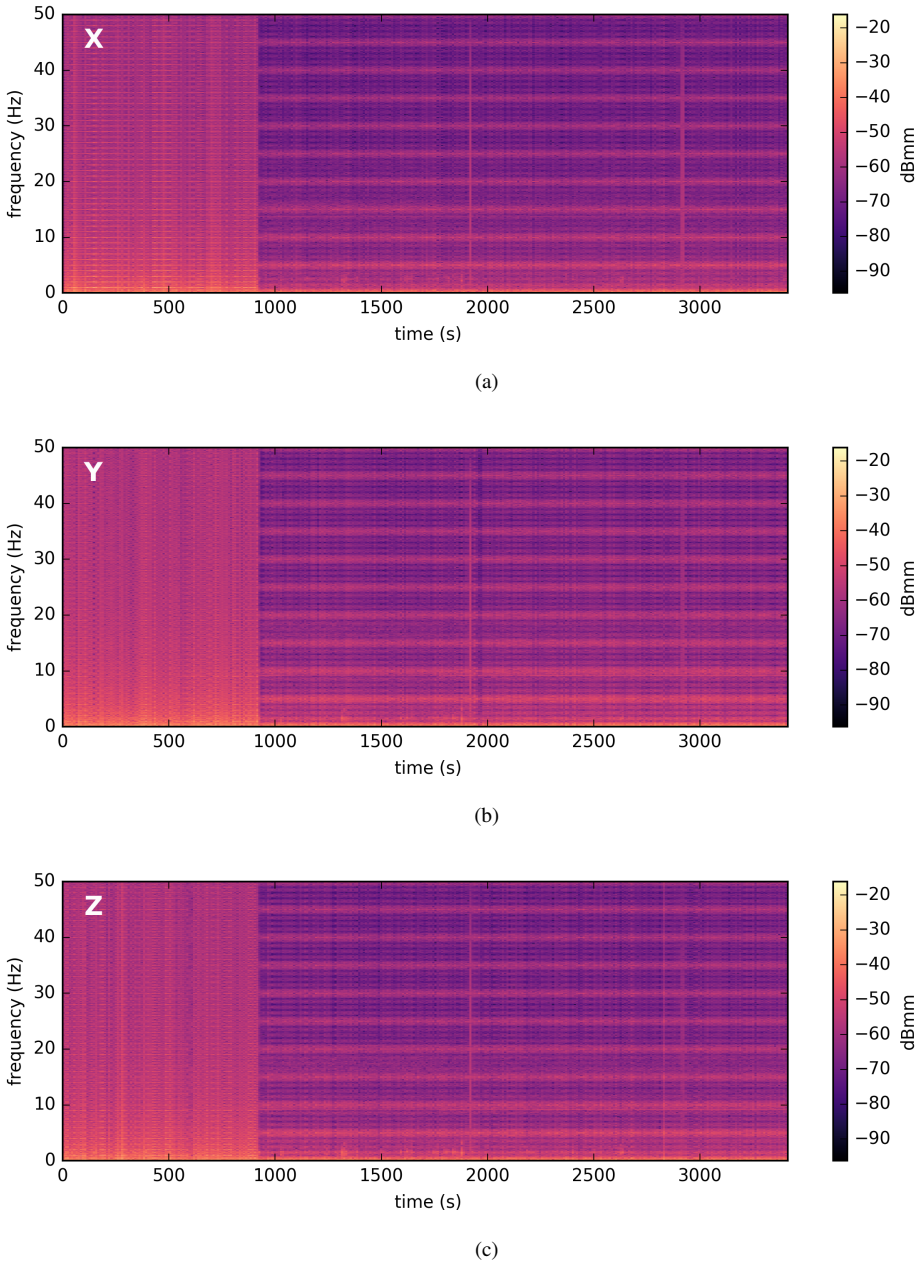


Figure 9. Visualizing correlation coefficients in Table 3 for position and orientation noise. Red shows positive correlation and blue show negative correlation.



(a)

(b)

(c)

Figure 10. Short-time Fourier transform of position. We see a distinct change in character around 900 seconds, which may be caused by the sensor rotating slightly.

thobasis, whereas general windowed sines and cosines do not form an orthobasis. In addition, wavelets may be customized to have certain desirable mathematical properties. One example is the Daubechies family of wavelets, which have the most compact support in time possible for a prescribed number of vanishing moments (Daubechies, 1988). The lowest number of vanishing moments is 1, which gives a Haar wavelet that ideally represents piecewise constant signals. p=2 vanishing moments ideally represents piecewise linear signals, p=3 best represents piecewise quadratic signals, and so forth. Our noise time series look mostly piecewise linear, so choosing a wavelet with a lower number of vanishing moments is justified. The discrete wavelet transform (DWT) projects a signal onto an orthobasis of wavelets; the wavelet coefficients represent the inner products of the signal on each orthobasis vector. The signal may be reconstructed by computing a linear combination of the orthobasis vectors using the wavelet coefficients as

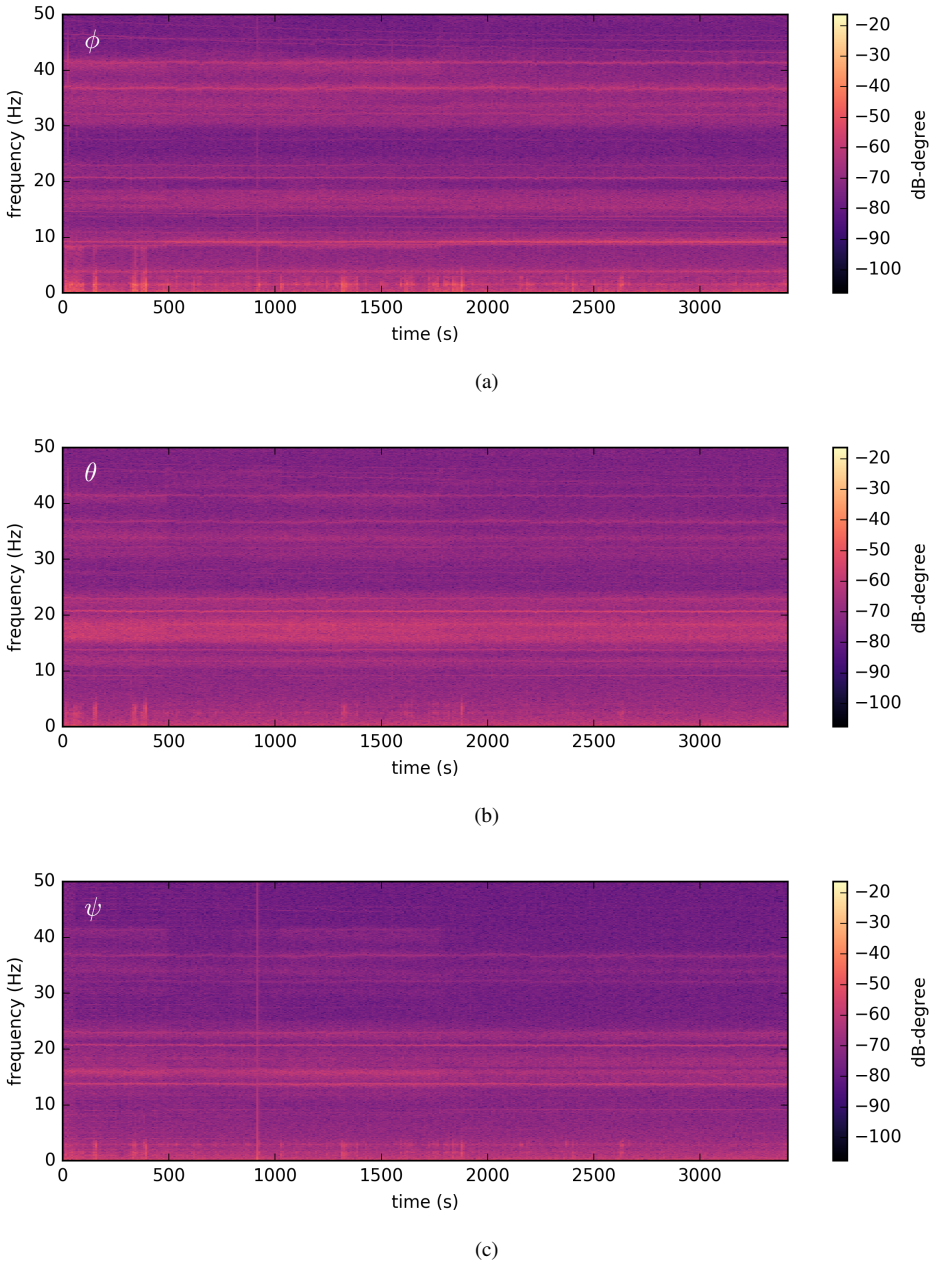


Figure 11. Short-time Fourier transform of orientation. We observe a broadband spike in the frequency content around 900 seconds, which may indicate the sensor was rotated and caused the change in position noise character after 900 seconds.

weights. This reconstruction is similar to reconstructing a signal from its Fourier transform using only a few frequencies.

Reconstructing the position noise time series using the largest 10% of the discrete wavelet transform (DWT) coefficients shows excellent recovery of the noise signals in Figure 14. A more comprehensive view of how well wavelets represent the noise is achieved by reconstructing the noise signals and measuring the signal-to-noise ratio (SNR) as a function of kept wavelet coefficients. Here, we use the definition for SNR between a clean signal $\mathbf{x}$ an noisy signal $\widetilde{\mathbf{x}}$ as

$$SNR(\mathbf{x}, \widetilde{\mathbf{x}}) = 20 \log_{10} \left(\frac{||\mathbf{x}||_2}{||\mathbf{x} - \widetilde{\mathbf{x}}||_2} \right). \quad (7)$$

Figures 12 shows how SNR increases with the number of kept coefficients for position and orientation. In general, high SNR is achieved by representing the noise signals using several

thousand wavelet coefficients. This result is excellent given our original noise time series is 360,000 samples long.

Observing the decay in sorted wavelet coefficients gives another way to analyze how the noise signal is expressed in the wavelet domain. In Figure 13(a), we see the largest 20 wavelet coefficients for the z-position noise time series. The decay is computed for a variety of Daubechies wavelets with a varying number of vanishing moments. The number of wavelet vanishing moments increases with the blue shade in the plot. We observe that a low number of vanishing moments gives a faster wavelet coefficient decay, and therefore represents our signal more concisely. This is expected, since our noise signal is mostly piecewise linear and therefore reconstructing using a wavelets with a low number of vanishing moments is appropriate. The wavelet coefficient decay for our seismic signal in Figure 1 is shown in 13(b). Unfortunately in general, the wavelet coefficients from out signal are lower than our noise.

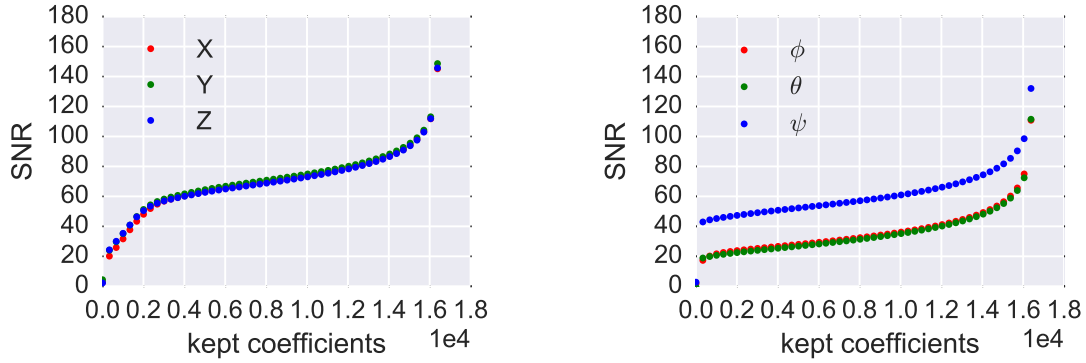
The wavelet coefficient decay plots and reconstruction errors show that the UAV position observation noise is more sparse than the ground motion signal in the wavelet domain. It would be beneficial to find a domain where the converse was true; a domain where the ground motion signal was sparser than the noise. If the noise and signal have sparse representations in sufficiently different domains, then we may potentially be able to use morphological component analysis (MCA) to separate the signals (Starck et al., 2010). An alternative is to leverage the fact that the noise and signal are statistically independent with respect to one another. Given the property of statistical independence, an appropriate signal separation technique may be Independent Component Analysis (ICA) (Comon, 1994). ICA separates sources by maximizing the independence of the estimated sources. There are a variety of ICA algorithms based on different measures of independence such as the InfoMax principle (Bell and Sejnowski, 1995) or maximum likelihood (Cardoso, 1997). In general, ICA works by solving a minimization problem to find both a linear mixing matrix of sources and the sources themselves. Prior information on source mixing may be used to help solve the minimization problem. However, a limitation of many ICA algorithms is the assumption that the sources are linearly mixed using a time-invariant transform. The rotation of the UAV changes with time, the mixing of the noise and signal vary with time in this scenario. Future work focuses on separating noise and signal, by expressing the signal in a domain where it is sparse.

5 CONCLUSION

In anticipation of using a supervised machine learning algorithm, we describe a strategy to simulate UAV and ground motion training data along with their respective noisy observations. Noise analysis of INS sensor data shows that observation noise in UAV position has a similar magnitude as the strong motion signal, and is correlated, non-white, and non-Gaussian. The noise appears to be sparsely represented in the wavelet domain. The sparse representation may be leveraged to remove the noise from observations. The search for separating criteria continues with a focus on representing the noise and the signal in sparse domains. When a sufficient separation criteria is found, training data with realistic noise may be readily generated based on the findings in this report.

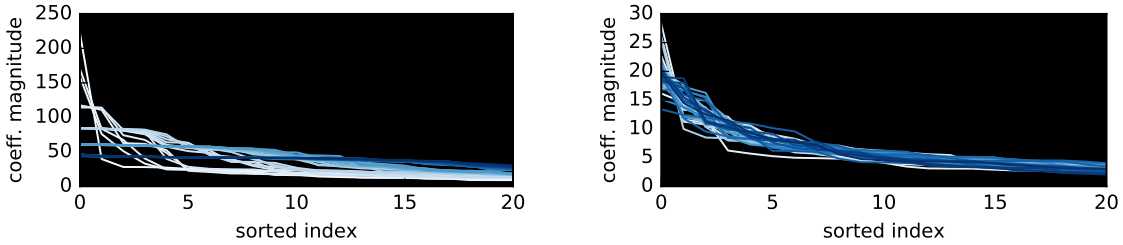
6 ACKNOWLEDGEMENTS

We thank the sponsor companies of the Center for Wave Phenomena, whose support made this research possible. We also thank Honeywell for gathering the HG4930 noise data and permitting results to be shared.



(a) (b)

Figure 12. SNR as a function of kept wavelet coefficients for (a) position noise, and (b) orientation noise. High SNR is achieved by reconstructing using a relatively small number of coefficients.



(a) (b)

Figure 13. Sorted Daubachies wavelet coefficient decay for vertical component of (a) noise and (b) an earthquake signal. Decay shown for multiple Daubachies wavelets. The number of wavelet vanishing moments increases with intensity of the color blue. Larger vanishing moment wavelets show slower coefficient decay.

REFERENCES

Abdel-Hamid, W., A. Noureldin, and N. El-Sheimy, 2007, Adaptive fuzzy prediction of low-cost inertial-based positioning errors: IEEE Transactions on Fuzzy Systems, **15**, 519–529.

Bell, A. J., and T. J. Sejnowski, 1995, An information-maximization approach to blind separation and blind deconvolution: Neural computation, **7**, 1129–1159.

Bengio, Y., P. Simard, and P. Frasconi, 1994, Learning long-term dependencies with gradient descent is difficult: IEEE transactions on neural networks, **5**, 157–166.

Cardoso, J.-F., 1997, Infomax and maximum likelihood for blind source separation: IEEE Signal processing letters, **4**, 112–114.

Center for Engineering Strong Motion Data, 2016, USGS Station OK034: www.strongmotioncenter.org.

Chiang, K.-W., A. Noureldin, and N. El-Sheimy, 2003, Multisensor integration using neuron computing for land-vehicle navigation: GPS solutions, **6**, 209–218.

Comon, P., 1994, Independent component analysis, a new concept?: Signal processing, **36**, 287–314.

Daubechies, I., 1988, Orthonormal bases of compactly supported wavelets: Communications on pure and applied mathematics, **41**, 909–996.

Diebel, J., 2006, Representing attitude: Euler angles, unit quaternions, and rotation vectors: Matrix, **58**, 1–35.

Flenniken, W., J. Wall, and D. Bevly, 2005, Characterization of various IMU error sources and the effect on navigation performance: Ion Gnss, 967–978.

Groves, P. D., 2013, Principles of GNSS, inertial, and multi-sensor integrated navigation systems: Artech house.

Hiliuta, A., R. Landry, and F. Gagnon, 2004, Fuzzy corrections in a GPS/INS hybrid navigation system: IEEE Transactions on Aerospace and Electronic Systems, **40**, 591–600.

Hochreiter, S., and J. Schmidhuber, 1997, Long short-term memory: Neural computation, **9**, 1735–1780.

Kalman, R., 1960, A new approach to linear filtering and prediction problems, tins: ASME, J. Basic, 34–45.

Mallat, S., 2008, A wavelet tour of signal processing: the sparse way: Academic press.

Mazal, C., 2018: personal communication.

Mikolov, T., I. Sutskever, K. Chen, G. S. Corrado, and

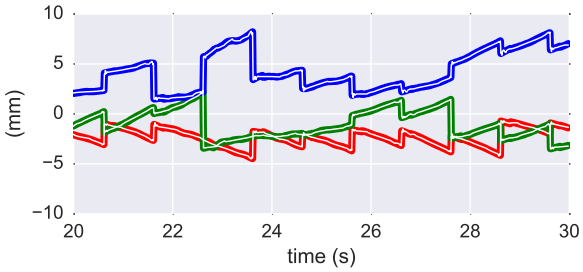


Figure 14. Zooming-in on the position noise time series shows the piecewise smooth character of position noise. The white curves show reconstructed noise signals using only the largest 10% of wavelet coefficients.

J. Dean, 2013, Distributed representations of words and phrases and their compositionality: Advances in neural information processing systems, 3111–3119.

Misra, P., and P. Enge, 2006, Global positioning system: Signals, measurements and performance second edition: Massachusetts: Ganga-Jamuna Press.

Noureldin, A., A. El-Shafie, and M. Bayoumi, 2011, GPS/INS integration utilizing dynamic neural networks for vehicular navigation: Information Fusion, **12**, 48–57.

Noureldin, A., A. El-Shafie, and M. R. Taha, 2007, Optimizing neuro-fuzzy modules for data fusion of vehicular navigation systems using temporal cross-validation: Engineering Applications of Artificial Intelligence, **20**, 49–61.

Noureldin, A., A. Osman, and N. El-Sheimy, 2003, A neuro-wavelet method for multi-sensor system integration for vehicular navigation: Measurement science and technology, **15**, 404.

Rapstine, T., P. Sava, and E. Arias, 2017, Airborne seismic data acquisition using stereo vision: Journal of Unmanned Vehicle Systems.

Semeniuk, L., and A. Noureldin, 2006, Bridging GPS outages using neural network estimates of ins position and velocity errors: Measurement science and technology, **17**, 2783.

Sharaf, R., A. Noureldin, A. Osman, and N. El-Sheimy, 2005, Online INS/GPS integration with a radial basis function neural network: IEEE Aerospace and Electronic Systems Magazine, **20**, 8–14.

Starck, J.-L., F. Murtagh, and J. M. Fadili, 2010, Sparse image and signal processing: wavelets, curvelets, morphological diversity: Cambridge university press.

Sutskever, I., J. Martens, and G. E. Hinton, 2011, Generating text with recurrent neural networks: Proceedings of the 28th International Conference on Machine Learning (ICML-11), 1017–1024.

Titterton, D., and J. L. Weston, 2004, Strapdown inertial navigation technology: IET, **17**.

Xu, Z., Y. Li, C. Rizos, and X. Xu, 2010, Novel hybrid of LS-SVM and Kalman filter for GPS/INS integration: The Journal of Navigation, **63**, 289–299.