

Sensors, Sensors, Everywhere ...

Chandrika Kamath

Lawrence Livermore National Laboratory, Livermore CA 94550

Sensors are increasingly being used to monitor complex systems, ranging from plasma physics experiments, such as tokamaks, to the power grid, airplanes, or a smart building. In these cases, the sensors allow us to control the normal operation of the grid or the building better, obtain forewarning of untoward incidents which may result in cascading power outages or damage to the experimental setup, and provide insights into the behavior of the plasma or the grid. Sensors in medical systems, such as X-rays, magnetic resonance imaging (MRI), and positron emission tomography (PET) scans, are being used to help diagnose ailments or to understand the effect of a particular treatment on a medical condition, such as the reduction in the size of a tumor when the patient is given a certain medication. Sensors are also used in non-destructive evaluation, where we want to observe the state of a system, such as the structural integrity of a bridge, without damaging it or taking it apart.

All these uses of sensors are the result of technological advancements that have resulted in the sensors becoming smaller, cheaper, and more versatile, allowing them to be deployed in unprecedented numbers and in places never even considered in the past. As a result, the data generated by these sensors have also increased in volume, being measured in terabytes or more. The data are also very complex and range from time series data to multi-spectral images taken over time, and streaming data that must be analyzed in real time. In this essay, we consider two illustrative applications where we need to analyze the data obtained from sensors. We describe the motivation for the analysis, the data, the challenges in the analysis, and potential avenues for solutions.

1 Integrating wind energy on the power grid

Renewable resources, such as wind and solar, are contributing an increasing percentage of our energy requirements. A challenge with these resources is that they are intermittent. There may be days where there is no wind and days where the wind speed suddenly increases by a large amount and remains at a high level for several hours. Or, the wind energy may be high during night time, when the demand for energy is relatively low. Until recently, when the percentage of wind energy on the power grid was small, this intermittency was not a problem. However, as renewable resources have started to contribute more than 20-30% of the energy on the grid in some regions, it is getting increasingly difficult for control room operators to ensure that the load and the generation are balanced at all times (see Figure 1(a)).

Wind energy is typically scheduled on the grid based on a forecast. Control room operators use a 0- to 6-hour ahead forecast to determine the amount of energy to schedule for the hours

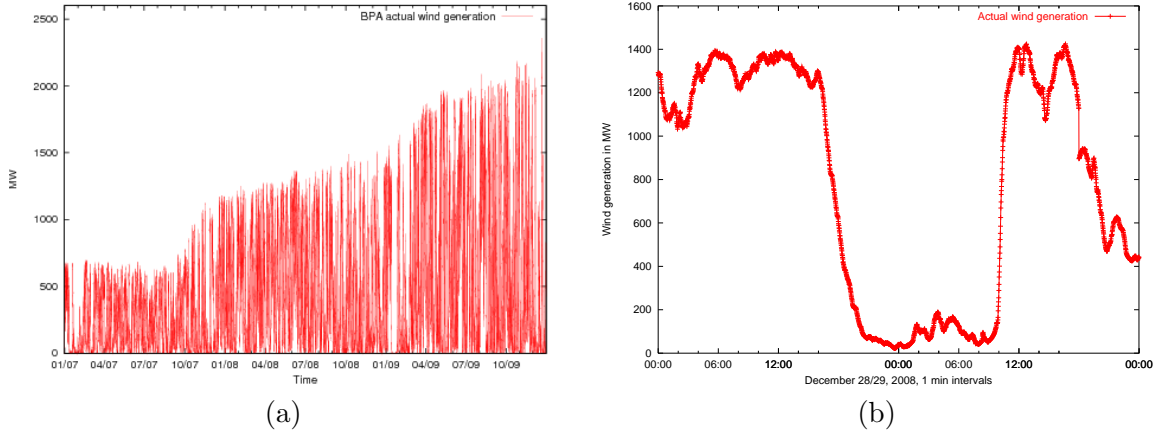


Figure 1: Wind generation from the wind farms in the mid-Columbia Basin providing energy to Bonneville Power Administration. (a) The increase in actual wind generation from 2007 to 2009. (b) Examples of positive and negative wind ramps during a two day period in December 2008.

ahead. The forecasts are updated hourly and, if required, appropriate changes made to the energy scheduled. Additional fine tuning is done in real time so that the load and the generation are balanced at all times. These wind power generation forecasts are obtained from numerical weather prediction simulations which predict the wind speeds for a time horizon of up to ten days. The wind speed is then converted into wind power generation for all the turbines in a wind farm.

However, forecasting wind speed accurately using numerical weather prediction models can be difficult, especially in regions where the terrain is complex and the meteorological processes controlling the wind speed are difficult to model. When the forecast is inaccurate, the operators schedule the wind energy for the upcoming hour based on the actual wind power generation for the previous hours or days, their past experience, and the current weather conditions. This is understandably difficult under normal operating conditions, but more so during ramp events, where the energy generated suddenly increases or decreases rapidly in response to changes in wind velocity (see Figure 1(b)).

In the case of positive ramps, the operators must either reduce other generation at short notice, sell the excess energy, or let the wind energy go to waste if the transmission lines cannot handle the sudden increase in energy. In case of a negative ramp event, the operators must have enough backup power to keep the load balanced. Having this additional back-up might not be cost-effective, especially if a negative ramp is predicted but does not occur.

In this problem, an obvious question to ask is the following - is there some way in which we can exploit the many sources of data better to enable the control room operators to make well-informed decisions? Afterall, there are frequently several weather stations in the vicinity of a wind farm or located upstream of the wind farm. Each of these weather stations measures several variables, such as the wind speed, wind gusts, temperature, pressure, relative humidity, and so on (see Figure 2).

One answer is to analyze historical data to gain further insights into the ramp events. We could consider just the wind power generation and extract statistics on the ramp events. How often do ramp events of a certain magnitude occur? Do they occur more often during certain times of the




Figure 2: Weather variables measured at 5 minute intervals for the first week of September 2010 at one of the weather stations in the mid-Columbia Basin region. (a) Relative humidity; (b) pressure; (c) wind direction; and (d) wind speed.

day or year? How does the frequency of these events change as the installed capacity of the wind farms increases? Do the positive ramps behave the same way as the negative ramps? While these are very simple calculations, they can none-the-less provide useful insights into wind generation and ramp events.

More sophisticated analyses are also possible. For example, we may ask which of the many weather variables available from different weather stations are the most informative? It is likely that many of the variables are correlated, for example, the temperature at two nearby weather stations at the same altitude. We can use Pearson's correlation coefficient to identify such variables and thus reduce the number of data streams that must be monitored. Another option is to use feature selection techniques to identify the variables associated with ramp events. For example, we can create a dataset that, for each day, has the daily average weather variables and a label which indicates if the day had ramp events or not. Then, for each weather variable, we could create two histograms - one for the values of the variable on the days with ramp events and the other for values

of the variable on days without ramp events. If the two histograms overlap, it is an indication that the variable is not very useful in separating days with and without ramp events. On the other hand, if the histograms are well separated, then the variable is likely to be an important indicator of ramp events.

We can also build predictive models to determine if the weather variables can be used to predict days when ramp events are likely. Simple models, such as decision trees, which are easy to interpret, can be quite helpful. More complex models, such as neural networks or support vector machines, could then be considered if they improve the accuracy.

There are several challenges in such analysis. Data from sensors are often of low quality, with several missing and incorrect values, as well as noise in the data. A few missing values can be interpolated, or predicted from other values, but a large number of missing values may make a sensor not very useful. Noise can be reduced by smoothing, for example, by convolution with a mean or a Gaussian filter. Sometimes, the data from different sensors are sampled at different rates. In such cases, we may need to determine what time interval is the most appropriate for analysis. For example, if some data are available at 10 minute intervals, while others are available at hourly intervals, we could perform the analysis by averaging the data so all variables are available for each hour. In some problems, we may also want predictions in real time. For example, instead of predicting days with ramp events, can we predict if there will be a ramp event in the next hour? As the time interval for the prediction reduces, the quality of the data becomes more important - missing values or noise in the data can make near-term prediction difficult. And, it goes without saying that the sensors must be in appropriate locations and of the appropriate type such that the variables they measure are indeed useful in the analysis. The identification of the locations and type of sensors (a problem referred to as sensor placement) offers other opportunities for the application of mathematical and statistical techniques.

2 Blob tracking in plasma physics experiments

As another example of sensor data, we consider sequences of images from the National Spherical Torus Experiment (NSTX: nstx-u.pppl.gov). NSTX is a magnetic fusion device in the form of a spherical torus (that is, a low aspect ratio tokamak) with several state-of-the-art diagnostics to help scientists understand magnetically-confined fusion reactors, such as ITER (iter.org). The success of such reactors is dependent on two key factors - first, the plasma has to be hot enough so that the particles can fuse, and second, it has to be confined long enough so that the particles do fuse. This can be challenging for many reasons, one of which is the presence of fine-scale turbulence which causes leakage of plasma and particles from the center of the reactor to the edge. This leakage could result in a significant heat loss from the plasma, loss of confinement of the particles, as well as erosion or vaporization of the containment wall of the reactor.

To understand the physical mechanisms behind this fine-scale turbulence, scientists have exploited 2-D optical imaging technology to obtain images of the edge turbulence in experimental reactors. These images have indicated the presence of “blobs” or “coherent” structures which “retain their geometry over many characteristic lengths of motion”. By characterizing these structures and tracking their motion over time, physicists can improve their understanding of edge turbulence and its role in the transport of heat and particles to the edge of the plasma.

Figure 3, shows how two different images from two sequences are processed for analysis. These sequences are obtained using an ultra-high speed, high resolution, PSI-5 camera, with each sequence consisting of 300 frames taken at 250,000 frames/second. Each 16-bit image is 64×64 pixels. The original raw images are often noisy, with isolated bright and dark pixels making it difficult to see the blobs in the data (as shown in the first column in Figure 3). The images are first processed using a 2-D 3×3 median filter to remove the noise spikes, followed by further smoothing using an 11×11 Gaussian filter with $\sigma = 0.4$. The resulting de-noised images are shown in the second column in Figure 3. Next, the background or quiescent intensity must be removed from the images. This is the intensity which would have been observed in the absence of the blobs. For short sequences, we can represent this quiescent emission profile by using either the mean or the median of the 300 frames that form the sequence. The third column in Figure 3 shows the quiescent image obtained by taking the median, and the fourth column shows the final result, after the background has been divided out from the de-noised image.




Figure 3: The processing of two images (top and bottom rows) from the National Spherical Torus Experiment. (a) Original image; (b) after filtering to remove noise; (c) the quiescent or background plasma represented using the median of the 300 frame sequence; (d) filtered image after dividing out the quiescent plasma.

A longer sequence of 15 consecutive images from a single sequence is shown in Figure 4 after each frame has been denoised and the quiescent image divided out. This clearly shows the movement of the blobs down and to the right over time.

The analysis task in this problem is to identify, extract, characterize, and track the coherent structures over time using a single algorithm with a single set of parameters (perhaps derived from the data) for all frames in the sequence. Our main motivation for the analysis is to compare the experimental data with theory so that we can validate or refine the theory. This is difficult as the coherent structures are poorly understood empirically and not understood theoretically. Consequently, the physics of these structures is under intensive investigation. From the analysis viewpoint, this unfortunately implies that we cannot compare the effectiveness of our blob-extraction methods by comparing our results with, say, a known ground-truth image, as there is no such image. What makes the problem even harder is that we cannot let our current understanding of the theory influence the results as the goal of the analysis is to validate the theory.




Figure 4: Fifteen consecutive frames from a longer sequence showing the movement of plasma “blobs” toward the bottom and right.

In the absence of ground-truth images, we need to consider alternative ways of evaluating different blob-extraction methods. For example, if an algorithm depends on many parameters making it difficult to select the best set of parameters, or the results of an algorithm are very sensitive to the values of the parameters, then it is unlikely to be a good candidate for extracting the blobs. This is because it would be difficult to ensure that the results are a true reflection of the data and not an artifact of the choice of algorithms or the parameters used. One idea we have explored is to select different methods which are not very sensitive to the parameter settings and see if we can obtain similar results when the methods are used to analyze an image sequence.

Designing robust algorithms for segmenting sequences of images and tracking objects over time is a challenging and interesting technical problem that occurs in many domains ranging from plasma physics to medicine and surveillance. As it has become increasingly inexpensive to obtain such images, mathematical and statistical techniques to address these problems will become necessary if we are to extract the useful information from these sequences.

3 Acknowledgments

The two problems described in this essay formed part of the Sapphire and StarSapphire projects at the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory. More information on the former is available at <https://computation.llnl.gov/casc/sapphire/> and on the latter at <https://computation.llnl.gov/casc/StarSapphire/>. I gratefully acknowledge the interactions with members of my project team and the interest of the application scientists in sharing their data and domain expertise with us.

LLNL-MI-542351 This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.