Homework 4 — Kriging and Wind Power

36-740, Fall 2026

Due at 6 pm on Thursday, 24 September 2026

The Problem

Wind power is a renewable, non-polluting energy source, but deciding where to put wind turbines is a complicated issue, with technical, political and social aspects[1]. On the technical side, an obvious issue is finding locations with high capacity to generate wind power. Here the basic rule of thumb is that the energy which can be extracted from a wind of speed v is proportional to the cube of the wind speed, v^3. (If you're curious, read the footnote[2].) Because wind speed fluctuates, power companies usually want locations with high average capacity over a typical year[3].

The data set on the course website contains measurements of the wind speed over a coastal region, taken at a grid of points, every 6 hours, over a period of several years. This homework consists of a series of questions which will build towards selecting a location for a new wind power plant.

Questions

1. (5) Where in the world is the data? Plot the locations of the grid points on a map. (The package `maps` may be helpful but isn't required.)
2. (5) Should you transform the data before before proceeding? Explain your reasoning. If you think a transformation is appropriate, do it now and do all further analyses on the transformed data.
3. (5) Create plots of the time series for each of the grid points. Comment on any common patterns you see in the data, and any anomalies or unusual features.
4. (5) Do your explorations suggest the need for any additional transformations or adjustments to the data? Explain, and, if you think they're called for, use the adjusted data going forward.
5. (5) Calculate a time average for each of the grid points, and plot the averages over space.
6. (10) Fit a stationary, isotropic covariance function for the averages. Plot and describe the shape of the covariance function.

[1] The readings on Canvas describe some of the controversies about particular wind power stations, or proposed stations, and some of the implications of wind power in general. One of them is directly relevant to our data.

[2] You can skip this note if you're willing to take the cube-of-wind-speed rule on faith. For those who are still here: The kinetic energy of a body of mass m and speed v is $\frac{1}{2}mv^2$. The kinetic energy per unit volume of a fluid, of density ρ, is thus $\frac{1}{2}\rho v^2$. The volume of fluid which flows past a turbine of cross-sectional area A in a time interval T is ATv. The total kinetic energy of the fluid flowing past the turbine is thus $\frac{1}{2}AT\rho v^3$. The energy per unit time ("power" in the jargon of physics) is thus $\frac{1}{2}A\rho v^3$. — This neglects considerations like whether the turbine is more or less efficient at different speeds, etc., so it's only start, but good enough for our purposes.

[3] Utility companies also need to worry about maximum wind speed (will the turbines break? can the electrical grid handle that much power coming in?), and, to a less extent, minimum speed (how much of the time will the turbine be idle?). We will ignore these complications, and just focus on average capacity.

7. (12) Using what you found in the previous two problems, and spatial linear prediction ("kriging"), calculate predicted averages for points on a rectangular grid, running from latitude 39.5 to 42.5, and longitude -69 to -72. (This extends the region the data was collected on.) Use as many grid points as you have time for; even very slow code should allow you to do a 30×30 grid. Plot these predictions on a map. Comment on any noticeable features of the map.

8. (12) Calculate the prediction variance for each of your grid points, and create a second map showing the standard deviation (not variance) of the predictions. Comment.

9. (10) Using your maps, where would you suggest locating a wind power station? What's the expected value of v^3 at this location? What is the standard error in this estimate?

10. (8) How precise is your suggestion? Refer to your previous work in explaining your answer.

11. (12) Your averages in problem 5 are based on a limited range of data, and therefore imprecise. For each location, calculate an annual average, the variance of those annual averages, and the implied variance for the over-all mean at that location, assuming annual averages are uncorrelated. (Why is that a not-crazy approximation?) Report the mean of the variances across locations. Re-do the kriging, incorporating the average variance as a "nugget effect" at distance 0. How much does the map of predicted values change? How much does the location of the predicted maximum change?

12. (1) How much time did you spend on this problem set?

Presentation Rubric (10): The text is laid out cleanly, with clear divisions between problems and sub-problems. The writing itself is well-organized, free of grammatical and other mechanical errors, and easy to follow. Plots are carefully labeled, with informative and legible titles, axis labels, and (if called for) sub-titles and legends; they are placed near the text of the corresponding problem. All quantitative and mathematical claims are supported by appropriate derivations, included in the text, or calculations in code. Numerical results are reported to appropriate precision. Any code displayed in the write-up is relevant, without dangling or unused fragments, and is indented, commented, and uses meaningful names. All parts of all problems are answered with coherent sentences, and raw computer code or output are only shown when explicitly asked for.