

## Probability Density Functions: Does the wind change from one year to the next?

Probability density functions (PDFs) are an essential part of statistical theory, and they have been used in turbulence studies for many years. However, the general applicability of PDFs is often hard to discern, and physical interpretations of PDFs are not easily found. The goal of this project is to help you gain some physical insight into PDFs and to let you test out formal ways to tell whether two PDFs differ.

Use NCEP 10-m wind data. The available data cover the globe at 6-hour intervals, with a year's worth of data in each file. The files are big, and may overload your computer. Thus, you may wish to choose a subset of this data for your analysis. (The NCO package can be used to subset NetCDF files before reading them into Matlab.)

1. Plot sample PDFs of surface winds. How would you characterize the wind PDFs? You will probably want to group data somehow to improve your statistics. To what extent are consecutive observations in time independent? Are geographically adjacent observations statistically independent?
2. Are the wind PDFs stationary in time? One formal mechanism for checking this is the Kolmogorov–Smirnov test (`kstest2` in matlab). How does this test work? What do you need to take into consideration in applying this test? How are your results influenced by the number of degrees of freedom in your data?
3. Another comparison test that you might consider is the chi-squared test. How do the results of this test differ from those of the Kolmogorov–Smirnov test?
4. **(bonus)** Investigations of the stochastic forcing of the ocean often look at the conditional probability of finding a given change in the wind, given a particular current value of the wind. What does this show?