PROBABILISTIC INFERENCE

Due date: 4/05/2018, 9am

Probabilistic inference is the process of using probability to assess a hypothesis using prior knowledge and existing evidence. In terms of minimization (fitting), probabilistic inference implies determining the posterior probability density functions (pdfs) of parameters of interest given the likelihood function (the model) and prior pdfs. Here are some good references to read up on probabilistic inference:

1. Sivia – Data Analysis: A Bayesian Tutorial,
   `http://aprsa.villanova.edu/files/sivia.pdf`
2. Hogg – Data Analysis Recipes,
   `https://arxiv.org/abs/1008.4686`,
   `https://arxiv.org/abs/1205.4446`,
   `https://arxiv.org/abs/1710.06068`
3. MacKay – Information Theory, Inference, and Learning Algorithms,
   `http://www.inference.org.uk/itprnn/book.html`

1. Build your own Metropolis MCMC algorithm for sampling the probability function. Make sure that the code can run arbitrary number of walkers simultaneously. Fit a linear function to the CO_2 data available on the course homepage. Plot the log-likelihood as a function of iteration, parameter posteriors, and parameter cross-section. Look up MCMC triangle plots for an idea of good visualization.
2. Using either autocorrelation or Rubin-Gelman test, justify the number of iterations necessary for the fit. Re-read Hogg's 2017 recipe for detailed instructions if you find yourself confused on how to do this.
3. Fit a silly linear function $f(x) = ax + b + c$ to the same data-set and replot all graphs. Do the results make sense?
4. Come up with a reasonable likelihood function to describe the CO_2 trend evident in the data. Comment extensively on the reliability of the fitted parameters.
5. Convert last week's assignment to probabilistic inference and determine parameter posteriors.