Tutorial Workshop on Parameter Estimation for Biological Models

Computer Practical: Introduction to Bayesian Analysis

July, 2018

Brian Reich

The table below has the overall free throw proportion and results of free throws taken in pressure situations, defined as ``clutch'' (https://stats.nba.com/), for ten National Basketball Association players (those that received the most votes for the Most Valuable Player Award) for the 2016-2017 season.

Player	Overall	Clutch makes	Clutch attempts
Russell Westbrook	0.845	64	75
James Harden	0.847	72	95
Kawhi Leonard	0.880	55	63
LeBron James	0.674	27	39
Isaiah Thomas	0.909	75	83
Stephen Curry	0.898	24	26
Giannis Antetokounmpo	0.770	28	41
John Wall	0.801	66	82
Anthony Davis	0.802	40	54
Kevin Durant	0.875	13	16

There is a clear trend in the overall percentage and clutch percentage:

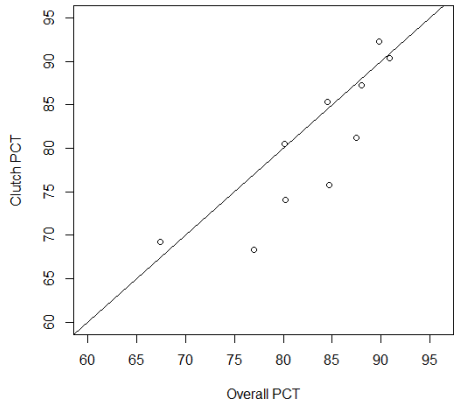

Define X and Y as the player's overall and clutch percentages, respectively, and fit the simple linear regression Y ~ Normal(a + bX,s) using the Gibbs sampler code: www4.stat.ncsu.edu/~reich/GIBBS_SLR.R

Fit the model and answer the following questions:

(1) Does the Gibbs sampler converge using default uninformative priors?
(2) Do you find a significant relationship between X and Y?
(3) What priors would you choose if you strongly believe a priori that there is a perfect relationship between X and Y? How do the results change when you use this prior?
(4) What priors would you choose if you strongly believe a priori that there is no correlation between X and Y? How do the results change when you use this prior?
(5) How do these results compare with least squares?
(6) Explain to your partner the difference between a sampling distribution and a posterior distribution.
(7) Explain to your partner how a Gibbs sampler works.